Statistical Inference:
777 Exercises, 77 Stories, and Solutions

Nils Lid Hjort
University of Oslo

Emil Aas Stoltenberg
University of Oslo and the BI Norwegian Business School

– This PartOne version, for STK 4090 / 9090,
with preliminary versions of Exercises from Chs 1-9, Appendix
last touched by Nils and by Emil, 12-March-2025 –

©Nils Lid Hjort and Emil Aas Stoltenberg, 2025

Some technical stuff

ISBN - Numbers numbers

The Kioskvelter Project

This is a draft of our book-to-be and it may not be reproduced
or transmitted, in any form or by any means, without permission.

1234-5678

To my somebody
– N.L.H.

To my somebody
– E.A.S.

Preface

This book builds on Hello, here is some text without a meaning. This text should show what a printed text will look like at this place. If you read this text, you will get no information. Really? Is there no information? Is there a difference between this text and some nonsense like "Huardest gefburn"? Kjift – not at all! A blind text like this gives you information about the selected font, how the letters are written and an impression of the look. This text should contain all letters of the alphabet and it should be written in of the original language. There is no need for special content, but the length of words should match the language.

(xx then three-four paragraphs here, on the carrying ideas behind and structure of the book: *exercises* and *stories*. a partly flipped classroom, with direct participation from the first pages of each chapter. there will be *solutions to all exercises*, not physically placed inside the book, but rather on the book's website-to-be, perhaps url'd `www.mn.uio.no/math/english/research/projects/HjortStoltenberg`. That website will also have all datasets and code is `R` and `python` to carry out all analyses, for the construction of each of the book's figures, etc.

the Hjort-Stoltenberg book website

(xx if we're clever with the 777 exercises, 77 stories, we should mention Stigler's 7 pillars. x)

(xx briefly on on prerequisties: linear algebra, with matrix theory, etc.; calcululs, with functions of one or more variables, partial derivatives, etc.; programming, in `R` or `Python` or other appropriate language, both for running common algorithms inside relevant pacakages, and for programming one's own functions, for simulation, etc.)

(xx crisp clear prose here, regarding segments of readers and how they can manouevre through the material. overall: from beginning master's level, in statistics, probability theory, data science, machine learning, and upwards, to PhD level and more. xx) (i) The Linear Readers, who will benefit from having the stamina to work through chapter by chapter (ideally also exercise for exercise), and appropriate subsets of our stories. These readers will be at a high master or PhD level. (ii) The Statistical Stories Readers, for those who already know the basics on statistical models, parameter estimation and testing, some Bayes, etc. (iii) Our book is also for the specialists inside certain themes, who wish to learn even more.

(xx crisp clear prose here, regarding courses and teaching. below we help readers and instructors by also providing short lists of relevant stories, for the different types of courses using our book. xx) Several types of courses can be taught from this book.

five courses which can be taught from our book

(i) Hard-core statistical inference, with parametric models, etc.: Chs. 1, half of 2, then most of 3, 4, half of 5, 6, 7, 8; a selection of Stories.

(ii) Large-sample theory, the careful probablity theory leading to CLT and more, with applications in statistics: Chs. 1, 2, 5; half of 9, a selection of Stories.

(iii) Empirical processes, convergence, approximations, applications in statistics: Chs. 1, 2, 5, 9; a selection of Stories.

(iv) Survival and event history analysis: Chs. 1, the essence of 2, 3, 4, 5, then the full 9; a selection of Stories.

(v) Model selection and model averaging: Chs. 1, the essence of 2, 3, 4, 5, then the full 11; a selection of Stories.

(vi) Bayesian statistics and confidence distributions: Chs. 1, the essence of 3, 5, then the full 6, 7, 15; a selection of Stories.

(vii) A special course on statistical applications can also be given, with emphasis on the stories, and less concern for all the technical details. This is also a way to apprciate and to learn about modern statistics, concepts, modelling, computation, scientific context.

The authors owe special thanks to Céline Cunen, Gudmund Hermansen, Tore Schweder, for having contributed significantly to several of our Statistical Stories, and also for always pleasant and inspiring long-term collaborations. Deep thanks are also due to a long list of colleagues and friends, who have taken part in discussions and rounds of clarification of relevance to various exercises and stories in our book: Marthe Aastveit, Patrick Ball, Bear Braumoeller, Gerda Claeskens, Aaron Clauset, Dennis Cristensen, Ingrid Dæhlen, Arnoldo Frigessi, Ingrid Glad, Håvard Hegre, Aliaksandr Hubin, Ingrid Hobæk Haff, Kristoffer Hellton, Anders Dahl Hjort, Bjørn Jamtveit, Martin Jullum, Vinnie Ko, Alexander Koning, George Lowther, Ian McKeague, Per Mykland, Per August Moen, Jonas Moss, Håvard Mokleiv Nygård, Lars Olsen, Sam Power, Oskar Høgberg Simensen, Catharina Stoltenberg, Gunnar Taraldsen, Ingrid Van Keilegom, Lars Walløe, Jonathan Williams, Lan Zhang.

We have also benefitted, directly and indirectly, through the collective efforts of several grander wide-horizoned funded projects: the *FocuStat: Focus Driven Statistical Inference with Complex Data* 2014-2019 project (led by Hjort) at the Department of Mathematics, University of Oslo, funded by the Norwegian Research Council; the *Stability and Change* 2022-2023 project (led by Hjort and Hegre), funded by and hosted at the Centre for Advanced Study (CAS), Academy of Science and Letters, Oslo; and *Integreat: The Norwegian Centre for Knowledge-Driven Machine Learning* 2023-2033 Centre of Excellence (led by Frigessi and Glad), Oslo, funded by the Norwegian Research Council. We finally acknowledge with gratitude a partial support stipend from the Norwegian Non-Fiction Writers and Translators Association (Norsk faglitterær forfatter- og oversetterforening).

Nils Lid Hjort and Emil Aas Stoltenberg
Blindern, some day in 2025

Contents

Preface iii

Contents v

I Short & crisp 1

1 Statistical models 3

2 Large-sample theory 51

3 Parameters, estimators, precision, confidence 113

4 Testing, sufficiency, power 147

5 Minimum divergence and maximum likelihood 181

6 Bayesian inference and computation 241

7 CDs, confidence curves, combining information 265

8 Loss, risk, performance, optimality 295

9 Brownian motion and empirical processes 319

10 Survival and event history analysis 337

11 Model selection 355

12 Markov chains, Markov processes, and time series 381

13 Estimating densities, hazard rates, regression curves 399

14 Bootstrapping 415

15 Bayesian nonparametrics 417

16 Statistical learning 419

II Stories **423**

i Demography, Epidemiology, Medicine **425**

ii Art, History, Literature, Music **465**

iii Economics, Political Science, Sociology **505**

iv Biology, Climate, Ecology **543**

v Sports **559**

vi Simulated stories **589**

vii Miscellaneous stories **611**

III Appendix **639**

A Mini-primer on measure and integration theory **641**

B Overview of stories and data **685**

References **705**

Name index **717**

Subject index **719**

Part I

Short & crisp

I.1

Statistical models

In this chapter we study families of distributions and densities that we are to meet time and again in this book. A partial list includes the uniform, normal and multinormal, chi-squared, the t and the F, Gamma, exponential, Weibull, Beta, Dirichlet, Poisson, compound Poisson, binomial, multinomial, geometric, Pareto, Gumbel, logistic. These families have parameters, with values to be set for certain studies or illustrations, or for purposes of confidence setting and tests; more generally these parameters are estimated from data, as we return to in several later chapters. We also learn fruitful ways of extending and mixing given families of distributions; in such fashions the classical models can be building blocks for forming new ones. Mathematical techniques for deriving crucial properties include those of moment-generating functions, convolutions, and double expectation.

Key words: distributions, models, moments, moment-generating functions, parameters, quantiles, sums, transformations

The aim of this chapter is to go through a generous list of parametric statistical models, from the well-known distributions connected with the normal model, to the Beta and the Gamma, to the binomial, Poisson, and negative binomial for discrete data, etc., along with deriving their basic properties. These models turn up repeatedly in later chapters and in our Statistical Stories, with variations, as direct models for data, or as building blocks for more complicated constructions. The normal and multinormal distributions play important roles, also because these become fruitful simple-to-use approximations to sometimes much more complicated exact distributions.

These models, for probability theory and statistics, rely on deeper mathematical constructions and considerations, with random variables being measurable functions on probability spaces, measure and integration theory, etc. For this book it has been practical to organise that body of mathematical theory in Appendix A. For the present chapter on models we take certain notions and basic definitions for granted, with background and more detail in that appendix. Thus we deal here with classes of distributions, parameters, probability densities, cumulative distribution functions, conditional and marginal distributions, means and variances, quantiles, correlations, and so on. In particular, any nonnegative function $f(y)$ integrating to 1 over some interval is a probability density over that interval; if Y is drawn from this f it has a cumulative distribution function (c.d.f.) $F(y) = \Pr(Y \le y)$, which in the continuous case is the same as $\int_{-\infty}^{y} f(y')\,\mathrm{d}y'$; the mean of a random variable Y drawn from this distribution is $\mathrm{E}\,Y = \int y f(y)\,\mathrm{d}y$; its median is $F^{-1}(\frac{1}{2})$; its variance $\mathrm{Var}\,Y = \mathrm{E}\,(Y - \mathrm{E}\,Y)^2$; the covariance between two random

variables X and Y with means a and b is $\operatorname{cov}(X,Y) = \operatorname{E}(X-a)(Y-b) = \operatorname{E} XY - ab$; etc. We also deal with sums of random variables, drawn from the same or different distributions. In one of its classical forms, the density $h(z)$ of $Z = X + Y$, where X and Y are independent with densities $f(x)$ and $g(y)$, is the convolution formula

$$h(z) = \int f(z-y)g(y)\,\mathrm{d}y = \int f(x)g(z-x)\,\mathrm{d}x, \tag{1.1}$$

see Ex. A.23 for more. There are also occasions in this chapter where the rules for double expectation and variance, i.e.

$$\operatorname{E} X = \operatorname{E}\operatorname{E}(X \mid Y) \quad \text{and} \quad \operatorname{Var} X = \operatorname{E}\operatorname{Var}(X \mid Y) + \operatorname{Var}\operatorname{E}(X \mid Y), \tag{1.2}$$

see Ex. A.27, come in handy.

In addition to defining and presenting a list of useful models, and diving into their properties and inter-connections, we develop certain tools, useful also in later chapters. These include transformations (see Ex. 1.12), moment-generating functions (see Ex. 1.31, with more in Ex. A.36), characteristic functions (see Ex. 1.35), conditional distributions, mixtures, and simulation. (xx make sure we have a little bit on simulation. xx) Also included is material on the general exponential family class, which has several of the classic models as special cases (see Ex. 1.55, with follow-up material in Ch. 4).

Importantly, several of the central models worked with in this introduction chapter find uses inside wider contexts, e.g. for regression situations, as we shall return to in later chapters. To indicate that direction of model building, below we learn about a random variable Y having a gamma distribution with parameters (a,b), which we write as $Y \sim \operatorname{Gam}(a,b)$; see Ex. 1.9. Now suppose there is a dataset consisting of measurements $(x_{i,1}, x_{i,2}, y_i)$ for individuals $i = 1,\ldots,n$, where the main outcome Y_i is influenced by the covariates $x_{i,1}, x_{i,2}$. Then a gamma regression model could take the form $Y_i \mid (x_{i,1}, x_{i,2}) \sim \operatorname{Gam}(a_i, b)$, with $a_i = \exp(\beta_0 + \beta_1 x_{i,1} + \beta_2 x_{i,2})$. The traditional multiple linear gression model is of a similar type, with Y_i given the covariates having normal distributions, with mean function linear in the covariates. Tools developed in later chapters may then be applied to estimate parameters, with confidence intervals, testing, comparisons, prediction, etc.

(xx just a few pointers to later chapters. CLT. normal approximations. estimation, testing. calibrate with what's in the abstract. we may point to more complex models, making clear that these classic families of distributions are often used as stepping stones. could point to Markov chains etc., but not really touching these in this chapter. also: take care with mentions of limit distributions and CLT, which we may choose to touch here and there, but details come in Ch. 2. xx)

(xx as of 12-March-2025, we have a little fortellerproblem: we do mgf and characteristic functions in Appendix A, good, but then need just a bit of basic models there, complete with $M(t)$ formulae for the normal, the binomial, just a few more. but those models are more formally introduced here in Ch1. so how to deal with this. xx)

(xx one para to point to Main Stuff, also signalled with the fermat. xx)

Normal, bi- and multinomial, exponential, gamma, mixing

Ex. 1.1. *The normal distribution.* The perhaps most famous and broadly useful distribution in probability theory and statistics is the normal distribution, also called the Gaußian distribution. It is also a building block for various inferred and related models

and distributions, as we learn later in the chapter. In its standard form, before we add on two more parameters, the normal density is

$$\phi(x) = \frac{1}{\sqrt{2\pi}} \exp(-\tfrac{1}{2}x^2) \quad \text{on the real line.}$$

We call this the standard normal distribution, and write $X \sim \mathrm{N}(0,1)$ to indicate this. It is standard in statistics and probability theory to use $\phi(x)$ for its density and $\Phi(x)$ for its cumulative distribution function (c.d.f.).

(a). There are myriad ways of demonstrating that $1/(2\pi)^{1/2}$ is the correct constant here, i.e. that $I = \int \exp(-\frac{1}{2}x^2)\,\mathrm{d}x = (2\pi)^{1/2}$. You are allowed to take this for granted, but attempt to show it via expressing I^2 as a double integral, featuring $\exp\{-\frac{1}{2}(x^2+y^2)\}$, and then substituting $x = r\cos\theta$ and $y = r\sin\theta$, followed by the use of double integration tools from calculus.

(b). Show that for X a standard normal, its mean is zero and its variance is one.

(c). With X a standard normal, consider $Y = \xi + \sigma X$, with ξ any number and σ positive. Show that its mean and standard deviation are μ and σ, and that its density
Gauß can be written

$$f(y) = \phi\Big(\frac{y-\xi}{\sigma}\Big)\frac{1}{\sigma} = \frac{1}{(2\pi)^{1/2}\sigma} \exp\Big\{-\tfrac{1}{2}\Big(\frac{y-\xi}{\sigma}\Big)^2\Big\}.$$

We write $Y \sim \mathrm{N}(\xi, \sigma^2)$ to indicate this distribution. Show that $\Pr(\xi - 1.96\,\sigma \le Y \le \xi + 1.96\,\sigma) = 0.95$. Find the c such that $\Pr(|Y-\xi| \le c\sigma) = 0.50$.

(d). With X a standard normal, consider $Z = X^2$. Find its c.d.f., and show that its density becomes $g(z) = (2\pi)^{-1/2}\exp(-\frac{1}{2}z)/\sqrt{z}$. We learn about the chi-squared distribution in Ex. 1.47; this X^2 has such a chi-squared distribution, with degrees of freedom equal to 1, which we write as $X^2 \sim \chi^2_1$.

(e). Consider X_1, X_2, X_3 being independent and standard normal. Work out the means and variances of X_1^2, $X_1^2 + X_2^2$, $X_1^2 + X_2^2 + X_3^2$. Simulate say 10^4 realisations of these distributions, check their histograms, and describe their different behaviour close to zero.

(f). Consider the enigmatic density $f(x) = e^{-\pi x^2}$ on the real line, featuring and combining the eternal mathematical constants e and π, integrating to 1. What is its standard deviation, and what is the probability that an X with this distribution is inside $[-1,1]$?

(g). A statistician is helping clients to predict the future outcome Y of some variable. The predictions $\widehat{Y}$ are such that $\widehat{Y} - Y$ is approximately a zero-mean normal with standard deviation 5.0. What is the probability that the prediction is inside $(-1,1)$ of the future outcome?

(h). (xx some pointers, placed here or elsewhere. point to mgf already, for the linear combination property. then a simple question to illustrate this. xx)

Ex. 1.2. *Normal sums.* Sums of independent normals have themselves normal distributions. This is clearly easiest to demonstrate via moment-generating functions (m.g.f.s), see Ex. 1.32 below, but it is worth doing this via convolution formulae too.

(a). Let X and Y be independent standard normals. Show that $X+Y \sim \mathrm{N}(0,2)$, via the convolution formula (1.1). With a bit more algebraic work, show that if $X_1 \sim \mathrm{N}(\xi_1, \sigma_1^2)$ and $X_2 \sim \mathrm{N}(\xi_2, \sigma_2^2)$ are independent, then $X_1 + X_2 \sim \mathrm{N}(\xi_1 + \xi_2, \sigma_1^2 + \sigma_2^2)$.

(b). Generalise this: show that if $X_i \sim \mathrm{N}(\xi_i, \sigma_i^2)$ for $i = 1, \ldots, m$, and these are independent, then $Z = \sum_{i=1}^m a_i X_i$ is also normal, with mean $\sum_{i=1}^m a_i \xi_i$ and variance $\sum_{i=1}^m a_i^2 \sigma_i^2$.

(c). Sometimes 2 + 2 might be 5, for extremely high values of 2. If independent X and Y are 2 and 2, but observed with with some Gaussian noise on top, of standard deviation level σ, find a formula $p(\sigma)$ for the probability that $X + Y$ is outside $[3.5, 4.5]$, i.e. 3 or 5 or even farther away from 4 when rounded off to the nearest integer. Plot this function.

Ex. 1.3. *Binomial distribution.* One of the more old, classic, and deservedly famous distributions in probability and statistics is the binomial. If there is a fixed probability $p = P(A)$ of a certain event A taking place, in a certain type of experiment, then the number Y of times A is seen, in n independent experiments, is the binomial, which we write as $Y \sim \mathrm{binom}(n, p)$.

the binomial distribution

(a). Show that

$$\Pr(Y = y) = \binom{n}{y} p^y (1-p)^{n-y} \quad \text{for } y = 0, 1, \ldots, n.$$

This involves the essential combinatorial fact that the number of ways one may place precisely y 1s in a total of n possible position is $\binom{n}{y} = n!/\{y!\,(n-y)!\}$. Explain that Y can be expressed as $X_1 + \cdots + X_n$, where X_i is a simple 0-1 variable, with $\Pr(X_i = 1) = p$, and where these are independent. Such X_i are called *Bernoulli variables*. Use this to prove the classic formulae np and $np(1-p)$ for mean and variance. Also, deduce the $\Pr(Y = y)$ formula from the $Y = \sum_{i=1}^n X_i$ description.

Bernoulli variables

(b). If the first question to ask concerning a distribution is about its centre (its mean, or perhaps its median), and the second is about its spread (its standard deviation, or perhaps a different measure, like its interquartile range), then the third question would be about its skewness, the degree of asymmetry. The classical skewness definition of a distribution, or equivalently of a random variable Y having that distribution, is $\mathrm{skew} = \mathrm{E}\, W^3$, where $W = (Y - \mathrm{E}\, Y)/(\mathrm{Var}\, Y)^{1/2}$ is the normalised version of Y, i.e. linearly transformed to have mean zero and standard deviation one. Show for the binomial (n, p) case that its skewness is

the skewness

$$\mathrm{skew} = \mathrm{E}\Big[\frac{Y - np}{\{np(1-p)\}^{1/2}}\Big]^3 = \frac{1-2p}{\{np(1-p)\}^{1/2}},$$

which for fixed p goes to zero with rate $1/\sqrt{n}$ (i.e. $\sqrt{n}\,\mathrm{skew}$ tends to a positive constant). Briefly discuss what this entails regarding the degree of asymmetry for the binomial distribution.

(c). After the skewness comes the so-called kurtosis, defined as $\mathrm{kurt} = \mathrm{E}\, W^4 - 3$, with W as in the previous point. The minus 3 is there in order for the kurtosis to be zero for the normal distribution; show that this is the case. Then show that

the kurtosis

$$\mathrm{kurt} = (1/n)[1/\{p(1-p)\} - 6],$$

for the binomial, and comment.

Ex. 1.4. *Trinomial probabilities.* (xx emil looks it over and checks if this is suitable here in Ch1, perhaps before Ex. 1.5. if not in App A. xx) Consider the so-called trinomial distribution for a random pair (X, Y), with probability mass function

$$f(x,y) = \frac{n!}{x!\,y!\,(n-x-y)!} p^x q^y (1-p-q)^{n-x-y} \quad \text{for } x \geq 0,\ y \geq 0,\ x+y \leq n.$$

Here n is the total count number, p, q the probabilities of events of type One and Two in repeated experiments, with $p+q<1$. With $Z = n - X - Y$ representing the number of events of type Three (not One, not Two), this is a model for the numbes of events One, Two, Three in n independent experiments; hence the trinomial name. See also Ex. 1.5.

(a). Verify that what is here called the probability mass function is the same as the density of the distribution with respect to counting measure on the set of (x, y) with $x \geq 0$, $y \geq 0$, $x + y \leq n$ – or, for that matter, with respect to counting measure on the set of all pairs (x, y) with $x \geq 0$, $y \geq 0$.

(b). Show by summing over the y that the distribution of X becomes a binom(n, p) from Ex. 1.3.

(c). Show that $Y \,|\, (X = x) \sim \text{binom}(n - x, q/(1-p))$. Give a formula for $\mathrm{E}\,(Y \,|\, X = x)$, and deduce the formula for $\mathrm{E}\,X$ from this. Find also the covariance between X and Y, using this scheme of conditioning with respect to $X = x$ first. Deduce that the correlation between them is $-\{p/(1-p)\}^{1/2}\{q/(1-q)\}^{1/2}$.

(d). Find a formula for $\Pr(X \leq x_0, Y \leq y_0)$, expressed as a sum over $x \in \{0, 1 \ldots, x_0\}$ (as opposed to a double sum over lots of (x, y) pairs). For a setup with $n = 50$, $(p, q) = (0.22, 0.33)$, compute the probability $\Pr(X \leq 15, Y \leq 20)$.

Ex. 1.5. *The multinomial model.* The binomial model, with basic properties treated in Ex. 1.3, is about sorting and counting events in two categories; if $Y \sim \text{binom}(n, p)$, then also $n - Y \sim \text{binom}(n, 1-p)$. The *multinomial model* is the natural extension of also Ex. 1.4 to more than two and three categories. Suppose there are n independent experiments, where each time one (and only one) of the events $A_1, \ldots, A_k$ takes place, with the same probabilities $p_1, \ldots, p_k$ for each experiment. Let then $Y = (Y_1, \ldots, Y_k)$, with Y_j counting the number of times A_j occurred, for $j = 1, \ldots, k$. Of course $Y_1 + \cdots + Y_k = n$, and $p_1 + \cdots + p_k = 1$, so there are $k - 1$ free parameters in the model.

(a). Show that $Y_j \sim \text{binom}(n, p_j)$, and deduce that we already know $\mathrm{E}\,Y_j = np_j$ and $\mathrm{Var}\,Y_j = np_j(1-p_j)$, even before we start working on the joint distribution of $(Y_1, \ldots, Y_k)$.

the multinomial model

(b). Show that the joint probability distribution becomes

$$f(y_1, \ldots, y_k) = \Pr(Y_1 = y_1, \ldots, Y_k = y_k) = \frac{n!}{y_1! \cdots y_k!} p_1^{y_1} \cdots p_k^{y_k}$$

for nonnegative $(y_1, \ldots, y_k)$ with sum n. The first factor $n!/(y_1! \cdots y_k!)$ is a combinatorial one, the number of different ways one may place '1' in y_1 positions, '2' in y_2 positions, etc., up to 'k' in y_k positions. Note that this generalises the classic $n!/(y_1!\,y_2!) = \binom{n}{y_1}$ for the binomial case, the number of ways one may place '1' in y_1 ways (and hence '2' in $n - y_1$ ways) in a list $1, \ldots, n$.

(c). Show that each pair has a trinomial distribution, e.g.

$$\Pr(Y_1 = y_1, Y_2 = y_2) = \frac{n!}{y_1!\, y_2!\, (n - y_1 - y_2)!} p_1^{y_1} p_2^{y_2} (1 - p_1 - p_2)^{n - y_1 - y_2}$$

for $y_1 \geq 0, y_2 \geq 0, y_1 + y_2 \leq n$. Note that formulae from Ex. 1.4 therefore apply to pairs (Y_i, Y_j) here. For $i \neq j$, show that $\mathrm{cov}(Y_i, Y_j) = -np_ip_j$, and find the correlation between Y_i and Y_j.

(d). Among the most used acronyms of statistical parlance is *i.i.d.*, for idependent and identically distributed. Explain in the present setup that $Y = Z_1 + \cdots + Z_n$, where $Z_1, \ldots, Z_n$ are i.i.d., with Z_i taking values $(1, 0, \ldots, 0), \ldots, (0, \ldots, 0, 1)$ with probabilities $p_1, \ldots, p_k$. Derive again the formulae for means, variances, covariances, starting with this representation. i.i.d.

Ex. 1.6. *Histograms.* Suppose data $Y_1, \ldots, Y_n$ are i.i.d. from some continuous density f over some interval $[a, b]$. Create disjoint cells $C_1, \ldots, C_k$, with $C_j = (a_{j-1}, a_j]$, for $a = a_0 < \cdots < a_k = b$. Let N_j count the number of data points in cell j. The *histogram*, associated with the chosen cells, is then histogram

$$\widehat{f}(x) = \widehat{p}_j / h_j \quad \text{for } x \in C_j, \text{ with lengths } h_j = a_j - a_{j-1},$$

where $\widehat{p}_j = N_j/n$ estimates $p_j = \Pr(Y_i \in C_j)$. In most automatic histogram algorithms, the width is taken constant across cells. The notation $\widehat{f}(x)$ indicates that beyond being an effective way of showing the essential spread and shape of the data, it is *an estimate* of the underlying $f(x)$.

(a). Carry out some simulations, samling n datapoints from the standard normal, creating histograms with k cells. For small and big n, play with k small, big, and about right. In `R`, you may use `hist(y,breaks=20,prob=T)`, etc.

(b). Show that $(N_1, \ldots, N_k)$ is multinomial. Taking a constant cell width h, for simplicity, explain that we may write $p_j = \int_{C_j} f \, \mathrm{d}y = f(\zeta_j)h$, for a suitable ζ_j inside cell C_j. From this derive that for $x \in C_j$, we have $\mathrm{E}\, \widehat{f}(x) = f(\zeta_j)$ and $\mathrm{Var}\, \widehat{f}(x) = f(\zeta_j)\{1 - f(\zeta_j)h\}/(nh)$. Argue that for the histogram to achieve $\mathrm{E}\, \widehat{f}(x) \to f(x)$ and $\mathrm{Var}\, \widehat{f}(x) \to 0$, as n increases, we need $h \to 0$ and $nh \to \infty$.

Ex. 1.7. *Hazard rates and survival functions.* Here and below we shall partly follow the implied tradition of using say T and $f(t)$ and $F(t)$, for random variables with their densities and c.d.f.s, rather than say Y and $f(y)$ and $F(y)$, when these relate to *time*. – Consider a random variable T on the halfline $[0, \infty)$, with density f and c.d.f. F. Classes of such distributions are sometimes most conveniently or fruitfully defined and discussed in terms of their hazard or cumulative hazard functions, as opposed to their densities and c.d.f.s, as we outline here; see also Ch. 10.

(a). For any nonnegative random variable T, the survival curve, or function, is $S(t) = \Pr(T \geq t)$. For T continuous, $S(t) = 1 - F(t)$, but explain that in the general case, we need $S(t) = 1 - F(t-)$, with $F(t-)$ the left hand limit of F at position t. Show that integrated survival is equal to mean survival:

$$\mathrm{E}\, T = \int_0^\infty S(t) \, \mathrm{d}t = \int_0^\infty t \, \mathrm{d}F(t). \tag{1.3}$$

(You may write the latter as $\int_0^\infty \int_0^\infty I(s < t) \, \mathrm{d}s \, \mathrm{d}t$ and then use the Fubini theorem.) Show from this that $\mathrm{E}\, T^p = \int_0^\infty \{1 - F(t^{1/p})\} \, \mathrm{d}t$, for any positive p.

hazard rate function

(b). Show that

$$\Pr(T \in [t, t+\varepsilon] \,|\, T \geq t) = h(t)\varepsilon + O(\varepsilon^2) \quad \text{for small } \varepsilon, \tag{1.4}$$

in terms of the so-called *hazard rate* function $h(t) = f(t)/\{1-F(t)\}$. With T interpreted as the time to a certain event, the function $h(t)$ describes the chance of this event taking place in the next instance, among those having survived up to t.

(c). So we may deduce the hazard rate from the density. Starting instead with $h(t)$, define first *the cumulative hazard* $H(t) = \int_0^t h(s)\,\mathrm{d}s$, and show that $F(t) = 1 - \exp\{-H(t)\}$. The function $S(t) = \Pr(T \geq t) = \exp\{-H(t)\}$ is important in its own right, and is called *the survival function.*

(d). Suppose an individual has survived up to time t_0. Show that

$$\Pr(T \geq t \,|\, T \geq t_0) = \frac{S(t)}{S(t_0)} = \exp[-\{H(t) - H(t_0)\}] \quad \text{for } t \geq t_0.$$

Show that the median lifetime, for such an invidual having lived up to t_0, is $t^* = H^{-1}(H(t_0) + \log 2)$.

Ex. 1.8. *The exponential distribution.* The *exponential distribution* is a simple but important one, in probability theory and statistics, which with positive parameter θ has the density $f(t,\theta) = \theta\exp(-\theta t)$ for $t > 0$. We write $T \sim \mathrm{Expo}(\theta)$ to indicate this.

the exponential distribution

(a). Show that the c.d.f. becomes $F(t,\theta) = 1 - \exp(-\theta t)$, and find the median. Show also that we may write $T = T_0/\theta$, where T_0 has the unit exponential distribution with density $\exp(-t_0)$. Show that T has mean and variance $1/\theta$ and $1/\theta^2$.

(b). Using Ex. 1.7, show that the hazard rate is constant, $h(t) = \theta$, and that the cumulative hazard rate is $H(t) = \theta t$. Show also that the exponential distribution is the only one where the hazard rate is constant.

(c). Show that the median survival time is $(\log 2)/\theta$. If an individual has survived up to time t_0, what is the median survival time?

(d). Assume certain light bulbs have a longevity distribution with the property that $\Pr(T \geq t_0 + t \,|\, T \geq t_0)$ does not depend on t_0. Argue that such light bulbs may be sold as if they were brand new, as long as they are still alive. Show that their distribution must be exponential.

the memoryless property

Ex. 1.9. *The Gamma distribution.* The gamma function is important in various branches in mathematics, probability theory, and statistics, and is defined as $\Gamma(a) = \int_0^\infty x^{a-1}\exp(-x)\,\mathrm{d}x$ for a positive. We may hence define a family of probability densities via $g_0(t,a) = \Gamma(a)^{-1}t^{a-1}\exp(-t)$ for $t > 0$. This is called the *Gamma distribution* with shape parameter a.

the gamma function

(a). With T_0 having this density, and b a positive scale parameter, show that $T = T_0/b$ has density

$$g(t,a,b) = \{b^a/\Gamma(a)\}t^{a-1}\exp(-bt) \quad \text{for } t > 0.$$

This is the two-parameter $\mathrm{Gam}(a,b)$ distribution. Verify that $\Gamma(1) = 1$, that $\Gamma(a+1) = a\Gamma(a)$ for all $a > 0$, and that $\Gamma(m) = (m-1)!$ for $m = 1, 2, \ldots$.

the Gamma distribution

(b). When T has the $\mathrm{Gam}(a,b)$ distribution, show that the mean and variance are a/b and a/b^2. Find also that $\mathrm{E}\,T^p = \{\Gamma(a+p)/\Gamma(a)\}/b^p$, valid for any p, as long as $p > -a$. Use this to show that the skewness and kurtosis become equal to skew $= 2/a^{1/2}$ and kurt $= 6/a$. Finally, regarding moments, find that the inverse gamma distributed variable $1/T$ has mean $b/(a-1)$ and finite variance $b^2/\{(a-1)^2(a-2)\}$, as long as $a > 2$.

(c). Verify that for $a = 1$ we have the exponential distribution, with density $b\exp(-bt)$ and cumulative $1-\exp(-bt)$. Show that $a = 2$ gives density $b^2t\exp(-bt)$ and cumulative $1-\exp(-ty)(1+bt)$. More generally, show that the cumulative is

$$\int_0^t g(s,a,b)\,\mathrm{d}s = 1-\exp(-bt)\Big\{1+bt+\frac{(bt)^2}{2!}+\cdots+\frac{(bt)^{a-1}}{(a-1)!}\Big\}$$

for the case of a being an integer.

(d). For a an integer, give an explicit expression for the hazard function $h(t,a,b)$, as per (1.4), and show that it converges to b as time increases. Show that this is the case also for any a, i.e. not only for integers; it increases from zero to b, if $a > 1$, and decreases from infinity to b, if $a < 1$.

(e). Let T_1, T_2 be independent and exponential with the same θ. Show that $T_1 + T_2 \sim \mathrm{Gam}(2,\theta)$. With $T_1,\ldots,T_k$ seen as the independent waiting times between events, show that the time to event k is a $\mathrm{Gam}(k,\theta)$.

(f). With $T_1 \sim \mathrm{Gam}(a_1,b)$ and $T_2 \sim \mathrm{Gam}(a_2,b)$ independent, show that $T_1+T_2 \sim \mathrm{Gam}(a_1+a_2,b)$. Generalise. This may indeed be accomplished via the convolution formulae from Ex. A.23, but as for other instances it becomes easier to prove such statements via m.g.f.s; see Ex. 1.31–1.32.

Ex. 1.10. *Mixing the exponential.* Sometimes waiting time type data do not follow an exact exponential distribution, but rather one characterised as a mixture of such; T given θ has the $\mathrm{Expo}(\theta)$ distribution, but the values of θ vary from occasion to occasion, or from individual to individual.

(a). Suppose indeed that $T\,|\,\theta \sim \mathrm{Expo}(\theta)$ but that θ has some density $g(\theta)$. Show that the density of T then becomes $f(t) = \int_0^\infty \theta\exp(-\theta t)g(\theta)\,\mathrm{d}\theta$.

(b). Suppose the distribution of θ is such that $1/\theta$ has mean value $1/\theta_0$ and a positive standard deviation τ. Show, starting with $\mathrm{E}\,(T\,|\,\theta) = 1/\theta$ and $\mathrm{Var}\,(T\,|\,\theta) = 1/\theta^2$, that $\mathrm{E}\,T = 1/\theta_0$ and $\mathrm{Var}\,T = 1/\theta_0^2 + 2\tau^2$; see (1.2). The case of a very tight distribution for the θ corresponds to τ small, which again means the case of a constant rate θ_0 for all.

(c). A convenient class of distributions for θ is the Gamma, with parameters (a,b), from Ex. 1.9. Its mean and variance are a/b and a/b^2; now find also the mean and variance of $1/\theta$. Show that the density of T can be written

$$f(t,a,b) = \int_0^\infty \theta\exp(-\theta t)\frac{b^a}{\Gamma(a)}\theta^{a-1}\exp(-b\theta)\,\mathrm{d}\theta = \frac{ab^a}{(b+t)^{a+1}},$$

and also that its c.d.f. is

$$F(t,a,b) = 1-\Big(\frac{b}{b+t}\Big)^a = 1-\frac{1}{(1+t/b)^a}.$$

(d). (xx a bit more. the hazard rate function $h(t) = f(t)/\{1 - F(t)\}$ is decreasing. expressions for quantiles, $t_0(p, a, b) = b\{1/(1-p)^{1/a} - 1\}$, solution to $F(t) = p$, to be used for Story iii.4 for fitting the distribution to the 95 between-war-times. xx)

(e). Find an expression for the hazard rate function $h(t, a, b) = f(t, a, b)/\{1 - F(t, a, b)\}$, and comment on its form, compared to the exponential case.

(f). Find the mean and variance of T, for the $g(t, a, b)$ distribution. This might be used to estimate (a, b) from data. (xx could point to Story iii.4, perhaps with better calibration. xx)

Ex. 1.11. *Gamma-mixing the gamma.* A given parametric distribution may sometimes be fruitfully extended by placing a separate distribution on one of its parameters. The following is an illustration.

(a). Consider a distribution which for given individuals is a gamma, but where the scale parameter varies between individuals. Specifically, suppose $Y \,|\, b \sim \text{Gam}(a_0, b)$ and that b has a distribution with $\text{E}\, 1/b = 1/b_0$ and $\text{Var}\, 1/b = \tau^2$. Show that Y has mean a_0/b_0 and variance $a_0/b_0^2 + (a_0 + a_0^2)\tau^2$.

(b). For the special case of $b \sim \text{Gam}(c, d)$, thus leading to a 3-parameter model, find the density $f(y, a, c, d)$ for Y. (xx work a bit with parametrisation here; big (c, d) correspond to old gamma. the following to be cleaned and sent to solutions. xx)

$$\bar f(y) = \int_0^\infty \frac{b^a}{\Gamma(a)} y^{a-1} \exp(-by) \frac{d^c}{\Gamma(c)} b^{c-1} \exp(-db)\, \text{d}b = \frac{d^c}{\Gamma(c)} \frac{\Gamma(a+c)}{\Gamma(a)} \frac{y^{a-1}}{(d+y)^{a+c}}.$$

Transformations, uniform, Pareto, Cauchy, Beta, Dirichlet

Ex. 1.12. *Transformation from X to Y.* We often encounter transformations, from one variable X to another Y, also in the vector case. We need formulae for how the density $g(y)$ of the Y can be found in terms of the density $f(x)$ for X.

(a). In the one-dimensional case, suppose $X = h(Y)$, equivalently $Y = h^{-1}(X)$, where h is smooth and increasing. Show that $\Pr(Y \le y) = \Pr(X \le h(Y))$, with density formula

$$g(y) = f(h(y))h'(y).$$

Show also that if $x = h(y)$ is continuous and decreasing, the formula becomes $g(y) = f(h(y))|h'(y)|$. Write down density formulae for the variables $Y_1 = \exp(X)$, $Y_2 = 3.33 - 2.22\, X$, $Y_3 = \log X$ (assuming for that case that X is positive).

(b). Show that if X is normal, then a linearly transformed $Y = a + bX$ is also normal. Show that if $X \sim \text{Gam}(a, b)$, with density proportional to $x^{a-1} \exp(-bx)$, then $Y = bX \sim \text{Gam}(a, 1)$.

(c). Suppose then that $X = (X_1, \ldots, X_p)^{\text{t}}$ and $Y = (Y_1, \ldots, Y_p)^{\text{t}}$ are vectors, with transformations binding them together,

$$\begin{pmatrix} X_1 \\ \vdots \\ X_p \end{pmatrix} = \begin{pmatrix} h_1(Y_1, \ldots, Y_p) \\ \vdots \\ h_p(Y_1, \ldots, Y_p) \end{pmatrix}, \quad \begin{pmatrix} Y_1 \\ \vdots \\ Y_p \end{pmatrix} = \begin{pmatrix} h_1^{-1}(X_1, \ldots, X_p) \\ \vdots \\ h_p^{-1}(X_1, \ldots, X_p) \end{pmatrix}.$$

We write this as $X = h(Y)$ and $Y = h^{-1}(X)$, for short. It is assumed that these systems of equations have unique solutions, and that the transformations are smooth, with continuous partial derivatives. In particular, the so-called Jacobi matrix

$$J(y) = \frac{\partial h(y)}{\partial y} = \frac{\partial h(y_1, \ldots, y_p)}{\partial y_1 \cdots \partial y_p},$$

having $\partial h_i(y)/\partial y_j$ as its (i, j) component, exists, and is continuous, with a non-zero determinant $\det(J(y))$. – Now, if X has density $f(x)$, show that

$$\Pr(Y \in B) = \int_{h(B)} f(x)\,\mathrm{d}x = \int_B f(h(y))|\det(J(y))|\,\mathrm{d}y.$$

This shows that Y has density $g(y) = f(h(y))|\det(J(y)))|$. This is essentially the multi-dimensional 'integration by substitution' formula of calculus.

(d). Let X, Y be independent standard exponentials, and consider $(U, V) = (X + Y, X/Y)$. Show that these become independent, having densities $u\exp(-u)$ and $1/(v + 1)^2$. More generally, suppose in this construction of (U, V) that X, Y are independent from the $\mathrm{Gam}(a, 1)$. Show that $U \sim \mathrm{Gam}(2a, 1)$, that V has density $\{\Gamma(2a)/\Gamma(a)^2\}v^{a-1}/(v+1)^{2a}$, and that they are independent.

Ex. 1.13. *Transformations to the uniform.* We say that a variable U is uniform on the interval $[a, b]$ if its density is constant over that interval, i.e. $1/(b - a)$, and zero outside. In particular, we write $U \sim \mathrm{unif}(0, 1)$ to indicate a variable with the uniform distribution on the unit interval.

(a). For such a $U \sim \mathrm{unif}(0, 1)$, find its mean and variance. Find the probabilities that U lands in $[0.03, 0.04]$, or in $[0.77, 0.78]$.

(b). Let F be a continuous and increasing c.d.f. for a variable X. Show that $U = F(X)$ is uniform on the unit interval. Conversely, we may start with $U \sim \mathrm{unif}(0, 1)$ and map to $X' = F^{-1}(U)$. Show that X' has distribution F. This is hence a simulation recipe, generating realisations from a given and perhaps not standard distribution.

(c). Consider the c.d.f. $F(x) = (x/10)^{3.33}$ on $[0, 10]$. Simulate 10^4 independent X_i from this distribution, by transforming uniforms. Make a fine histogram, with the density $f(x)$ plotted alongside.

Ex. 1.14. *Polar and spherical coordinates.* Points in the plane may be transformed from Cartesian to polar coordinates, and vice versa; similarly, but with more delicate details, points in the three-dimensional space can be represented in spherical coordinates.

(a). Suppose X and Y are independent and standard normal, and transform to polar coordinates, $X = R\cos A$ and $Y = R\sin A$. Find the density $g(r, a)$ for (R, A), with R positive and $A \in [0, 2\pi]$. Show in particular that length R and angle A become independent, with A having a uniform distribution on $[0, 2\pi]$ (i.e. the flat density $1/(2\pi)$ over that interval). Find also the distribution of $Z = Y/X = \tan A$; see also Ex. 1.18.

(b). Explain that a somewhat idiosyncratic but nevertheless feasible way of simulating pairs of standard normals is to draw R from the density $h_0(r) = r\exp(-\frac{1}{2}r^2)$ and A from the uniform over $[0, 2\pi]$, followed by the polar transformation. Do this in your computer,

generating say 10^4 pairs (X_i, Y_i). Then tweak your code, drawing R from some density h different from h_0. Show that this amounts to the joint density

$$f(x,y) = h((x^2+y^2)^{1/2})(x^2+y^2)^{-1/2}/(2\pi).$$

Take e.g. $h(r) = \frac{1}{2}\gamma r^{\gamma-1}\exp(-\frac{1}{2}r^\gamma)$, with γ close to but not equal to 2. Generate and plot 10^4 points (x,y) via this recipe, perhaps via the inverse c.d.f. trick of Ex. 1.13; how can you spot that these data are *not* independent normal pairs?

(c). Then go from plane to space. Start with (X,Y,Z) a triple of independent standard normals, and transform to (R,A,B) via

$$X = R\sin A\cos B, \quad Y = R\sin A\sin B, \quad Z = R\cos A.$$

Set up the Jacobi matrix $\partial(x,y,z)/\partial(r,a,b)$ of partial derivatves, and show that its determinant becomes $r^2\sin a$. Use this to show that the joint density of (R,A,B) can be written in product form $g_1(r)g_2(a)g_3(b)$, with g_1 the density proportional to $r^2\exp(-\frac{1}{2}r^2)$, $g_2(a) = \frac{1}{2}\sin(a)$ over $[0,\pi]$, and $g_3(b) = 1/(2\pi)$ over $[-\pi,\pi]$.

(d). Just to see operationally that transformations go two ways, simulate say 10^5 standard normal triples (X,Y,Z) via simulating triples (R,A,B) from the densities found away. Check that the resulting (X,Y,Z) look adequate. As for the plane question, tweak your code to draw R from a somewhat different distribution than proportional to $r^2\exp(-\frac{1}{2}r^2)$, and explain that you have invented a generalisation of the triple independent normal.

Ex. 1.15. *Ordering exponentials.* Let Y_1, Y_2, Y_3 be independent unit exponentials, with density $f(y)\exp(-y)$ for y positive, and order them, to $Y_{(1)} < Y_{(2)} < Y_{(3)}$. Then define the so-called spacings between them, $Z_1 = Y_{(1)}$, $Z_2 = Y_{(2)} - Y_{(1)}$, $Z_3 = Y_{(3)} - Y_{(2)}$.

(a). Show first that the joint density of $(Y_{(1)}, Y_{(2)}, Y_{(3)})$ is $3!\,f(y_{(1)})f(y_{(2)})f(y_{(3)})$ on the set $y_{(1)} < y_{(2)} < y_{(3)}$. Find then the joint density for (Z_1,Z_2,Z_3), and show that they are independent.

(b). Then generalise, considering i.i.d. unit exponentials $Y_1,\ldots,Y_n$, ordered into $Y_{(1)} < \cdots < Y_{(n)}$. Work with the scaled spacings $V_1 = nY_{(1)}$, $V_2 = (n-1)(Y_{(2)} - Y_{(1)})$, up to $V_{n-1} = 2(Y_{(n-1)} - Y_{(n-2)})$, $V_n = Y_{(n)} - Y_{(n-1)}$. Show that

$$Y_{(1)} = \frac{V_1}{n}, Y_{(2)} = \frac{V_1}{n} + \frac{V_2}{n-1}, \ldots, Y_{(n)} = \frac{V_1}{n} + \frac{V_2}{n-1} + \cdots + \frac{V_{n-1}}{2} + \frac{V_n}{1},$$

and then show that in fact $V_1,\ldots,V_n$ are i.i.d. unit exponentials.

(c). The Euler constant $\gamma_e = 0.5772...$ is defined as the limit of $1+1/2+\cdots+1/n-\log n$. Use the above to show that $Y_{(n)} = \max_{i\le n} Y_i$ has mean close to $\log n + \gamma_e$, and variance converging to $\pi^2/6$.

Ex. 1.16. *The Pareto distribution.* Consider a variable T defined on the range $T \ge t_0$, for some positive t_0, with c.d.f. $F(t) = 1-(t_0/t)^\theta$ for $t \ge t_0$, for some positive parameter θ. This is the Pareto distribution, and we may write $T \sim \text{pareto}(t_0,\theta)$ to indicate this.

the Pareto distribution

(a). Use the mean formula of (1.3) to show that the mean of a Pareto is $t_0\theta/(\theta-1)$, for $\theta > 1$. Show furthermore that the variance is $t_0^2\theta/\{(\theta-1)^2(\theta-2)\}$, for $\theta > 2$, and that the median is $\text{med}(T) = t_0 2^{1/\theta}$.

(b). Show that the density is $f(t,\theta) = \theta t_0^\theta / t^{\theta+1}$ for $t \geq t_0$; you may use this to find the mean formula again. Furthermore, explain that the hazard rate is $h(t) = \theta/t$.

(c). Show that $\Pr(T \geq t \,|\, T \geq t_1) = (t_1/t)^\theta$, for $t \geq t_1$. So $T \,|\, (T \geq t_1)$ is Pareto (t_1, θ). Deduce from this that

$$\text{med}(T - t \,|\, T \geq t) = ct = (2^{1/\theta} - 1)t \quad \text{for all } t.$$

This is linked to the so-called Lindy Effect: the longer the life is observed to be (of a company, an idea, a party), the longer is the remaining lifetime expected to be (note that $c > 1$). Show in fact that the Pareto distribution is the only one with the property that $\text{med}(T - t \,|\, T \geq t) = ct$ for all t. the Lindy Effect

(d). Show that $Y = \log(T/t_0)$ is exponential with parameter θ. So a representation of the Pareto is $T = t_0 \exp(Y)$, with $Y \sim \text{Expo}(\theta)$. You may use this to find the mean formula once more.

Ex. 1.17. *Maxima of i.i.d. samples.* To illustrate one of the many ways in which suitable start models may be generalised and extended, consider $Y_1, \ldots, Y_n$ i.i.d. from some distribution with c.d.f. F and density f, and from these define $M_n = \max_{i \leq n} Y_i$. (xx brief pointer to order statistics and sample quantiles in Chs 2, 3. xx)

(a). Show that M_n has c.d.f. $G_n(z) = F(z)^n$ with density $g_n(z) = nF(z)^{n-1} f(z)$. Show that its median is $\mu_m = F^{-1}((\frac{1}{2})^{1/m})$. Compute and draw these, for say $n = 1, \ldots, 10$, for the case of the Y_i being standard normal, tracking also the medians μ_m.

(b). Before pursuing some maxima, consider a situation with X_1 and X_2 being independent, with c.d.f.s H_1, H_2 and densities h_1, h_2. Show that

$$\Pr(X_1 \leq X_2) = \int H_1(x_2) h_2(x_2)\, \mathrm{d}x_2 = \int \{1 - H_2(x_1)\} h_1(x_1) \mathrm{d}x_1.$$

(c). Long jumpers A and B have jumps being $\text{N}(7.90, 0.10^2)$ and $\text{N}(7.80, 0.15^2)$, respectively (on the metres scale). What is the probability that A jumps longer than B, with one jump each? Consider also p_m, the probability that B after m jumps is better than A after $m = 1, 2, 3, 4, 5, 6$ jumps (in the sense of 'best jump so far'), assuming that all jumps are independent. Show that $p_m = \int F_1(z)^m m F_2(z)^{m-1} f_2(z)\, \mathrm{d}z$, in terms of the two distributions. Compute p_m both via numerical integration and via simulations.

Ex. 1.18. *Ratios and the Cauchy.* If (X, Y) has a certain distribution, what happens to the ratio $V = Y/X$?

(a). Suppose that X and Y are independent with the same density f on $(0, \infty)$. Show that $V = Y/X$ has density $g(v) = \int_0^\infty x f(x) f(vx)\, \mathrm{d}x$. With X and Y independent from the same exponential distribution, show that $g(v) = 1/(1+v)^2$.

(b). With X and Y independent from the same Gamma (a, b), show that $V = Y/X$ has density $\{\Gamma(2a)/\Gamma(a)^2\}\, v^{a-1}/(1+v)^{2a}$.

(c). Suppose now that X and Y are independent from the same density f, symmetric around zero. Show that V has density $g(v) = 2\int_0^\infty x f(x) f(vx)\, \mathrm{d}x$. For the special case of a ratio of two independent standard normals, show that the Cauchy

$$g(v) = (1/\pi)/(1+v^2) \quad \text{with c.d.f.} \quad G(v) = \tfrac{1}{2} + (1/\pi) \arctan v.$$

This is the Cauchy distribution (in its standard form). Show that it has no mean. Find its interquartile range. Deduce the Y/X distribution also via Ex. 1.14(a).

(d). There is an infinitely long straight line ahead of you, with distance r to the nearest point. You kick a ball towards the line, with some angle $A \in (-\pi/2, \pi/2)$. Show that it crosses the line at position $X = r\tan(A)$. Show that with A having density g and c.d.f. G, over $(-\frac{1}{2}\pi, \frac{1}{2}\pi)$, then X has c.d.f. $F(x) = G(\arctan(x/r))$. When A is uniform over that interval, show that X/r is standard Cauchy. Investigate what happens if A is uniform over a tighter interval, like $(-\pi/4, \pi/4)$.

(e). What distribution does the random angle A need to have, in order for X to be normal? For $r = 1$ unit away from the line, show that the peculiar density $g(a) = \phi(\tan(a))/\cos^2(a)$ leads to a standard normal X. Graph g to see that it is symmetric and bimodal, with modes at $\pm$ 45 degrees.

Ex. 1.19. *The Beta distribution.* An important class of distributions, over the unit interval $(0, 1)$, is the *Beta distribution*, with two positive parameters. We write $p \sim \text{Beta}(a, b)$ if its density is

the Beta distribution

$$\text{be}(p, a, b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} p^{a-1}(1-p)^{b-1} \quad \text{for } p \in (0, 1).$$

(a). As an introductory exercise, of separate interest, suppose X and Y are independent Gamma variables with parameters $(a, 1)$ and $(b, 1)$. Construct from these the sum $Z = X + Y$ and ratio $P = X/(X + Y)$. Finding the joint density for (P, Z), demonstrate that Z and P are independent, with $Z \sim \text{Gam}(a + b, 1)$ and P having precisely the $\text{be}(p, a, b)$ density. This in particular shows that the integration constant is the correct one, i.e. that $\int_0^1 p^{a-1}(1-p)^{b-1}\,\mathrm{d}p = \Gamma(a)\Gamma(b)/\Gamma(a+b)$.

(b). Compute and display a few of these densities, for (a, b) of your choice. Note that the uniform is the special case of $(a, b) = (1, 1)$.

(c). Show that $\text{E}\, p = p_0 = a/(a + b)$ and that $\text{Var}\, p = p_0(1 - p_0)/(a + b + 1)$.

(d). Find a formula for $\text{E}\, p^m$, for $m = 1, 2, \ldots$, in terms of (a, b), and in terms of the reparametrisation $(cp_0, c(1 - p_0))$. Use this to find

$$\text{E}\,(p - p_0)^3 = \frac{2p_0(1-p_0)(1-2p_0)}{(c+1)(c+2)},$$

with a consequent formula for the skewness. For fixed mean p_0, show that the skewness tends to zero with increasing c.

(e). Examine the particular Beta$(\frac{1}{2}, \frac{1}{2})$ distribution. Show that its density and c.d.f. become $f(p) = (1/\pi)/\{p(1-p)\}^{1/2}$ and $F(p) = (2/\pi)\arcsin(\sqrt{p})$. Find is quantile $F^{-1}(q)$. Plot these functions.

Ex. 1.20. *The Dirichlet distribution.* Let $G_1, \ldots, G_k$ be independent and Gamma distributed, with parameters $(a_1, 1), \ldots, (a_k, 1)$. With $G = G_1 + \cdots + G_k$ their sum, consider the random ratios

$$(X_1, \ldots, X_{k-1}) = (G_1/G, \ldots, G_{k-1}/G).$$

It inherits a distribution, with density $h(x_1, \ldots, x_{k-1})$, worked with below, in the simplex where each $x_i \geq 0$ and $x_1 + \cdots + x_{k-1} < 1$. Taking also $X_k = G_k/G = 1 - (X_1 + \cdots + X_{k-1})$ on board, we have a vector $(X_1, \ldots, X_k)$ of random probabilities summing to 1 over its k categories. Its distribution has a name: it's the Dirichlet distribution, with k categories, and parameters $(a_1, \ldots, a_k)$, which we write as $X \sim \text{Dir}(a_1, \ldots, a_k)$.

the Dirichlet distribution

(a). Suppose $(X_1, X_2, X_3, X_4, X_5, X_6) \sim \mathrm{Dir}(a_1, a_2, a_3, a_4, a_5, a_6)$. Show that $(X_1 + X_4 + X_6, X_2, X_3 + X_5) \sim \mathrm{Dir}(a_1 + a_4 + a_6, a_2, a_3 + a_5)$. Generalise and formalise this summing-over-cells property of the Dirichlet distribution.

(b). With $X \sim \mathrm{Dir}(a_1, \ldots, a_k)$, show that each $X_i \sim \mathrm{Dir}(a_i, a - a_i)$, with $a = a_1 + \cdots + a_k$, and that this is the same as a $\mathrm{Beta}(a_i, a - a_i)$. Show from this that

$$\mathrm{E}\, D_i = \xi_i = a_i/a, \quad \mathrm{Var}\, D_i = \xi_i(1 - \xi_i)/(a + 1),$$

Show also the $\mathrm{cov}(D_i, D_j) = -\xi_i \xi_j/(a + 1)$ for $i \neq j$.

(c). We have been able to derive certain basic properties above, without really needing an expression for the density of a Dirichlet vector. We tend to this now, using the transformation machinery of Ex. 1.12. Show in fact, starting with $(G_1, \ldots, G_k)$ and transforming to $(X_1, \ldots, X_{k-1}, G)$, (i) that $(X_1, \ldots, X_{k-1})$ has the density

$$h(x_1, \ldots, x_{k-1}) = \frac{\Gamma(a)}{\Gamma(a_1) \cdots \Gamma(a_k)} x_1^{a_1 - 1} \cdots x_{k-1}^{a_{k-1} - 1} (1 - x_1 - \cdots - x_{k-1})^{a_k - 1}$$

over the simplex; (ii) that $G \sim \mathrm{Gam}(a, 1)$; (iii) that these are independent; and (iv) that this also verifies the implied formula for integrating $x_1^{a_1 - 1} \cdots x_{k-1}^{a_{k-1} - 1} (1 - x_1 - \cdots - x_{k-1})^{a_k - 1}$ over the simplex. Note that these efforts and results generalise the findings of Ex. 1.19(a).

Ex. 1.21. *Dirichlets inside Dirichlets.* The summing-over-cells property of the Dirichlet, see Ex. 1.20, has other consequences and angles, and we shall learn here that long Dirichlet vectors might be split into Dirichlet parts via Dirichlet cuts.

(a). Start with A, B, C independent gammas with parameters $(a, 1)$, $(b, 1)$, $(c, 1)$. Form from these $X = A/(A + B)$, $Y = (A + B)/(A + B + C)$, $Z = A + B + C$. Explain that we already know that $X \sim \mathrm{Beta}(a, b)$, $Y \sim \mathrm{Beta}(a + b, c)$, $Z \sim \mathrm{Gam}(a + b + c, 1)$. Show in fact that X, Y, Z also are independent. This means working out their joint distribution; establish first that the inverse transform is $A = XYZ$, $B = (1 - X)YZ$, $C = (1 - Y)Z$, and that the associated Jacobi determinant becomes yz^2. Explain that all of this leads to the product Beta representation

$$\frac{A}{A + B + C} = \frac{A}{A + B} \frac{A + B}{A + B + C} = \mathrm{Beta}(a, b)\, \mathrm{Beta}(a + b, c) = \mathrm{Beta}(a, b + c).$$

(b). (xx nils rant, to be edited. link to pinned-down Dirichlet processes. xx) Suppose $(X_1, \ldots, X_k, Y_1, \ldots, Y_\ell)$ is $\mathrm{Dir}(a_1, \ldots, a_k, b_1, \ldots, b_\ell)$, which we may represent as $X_i = G_i/(S + T)$, $Y_j = H_j/(S + T)$, with the G_i and H_j being independent gamma variables with the appropriate parameters, and sums $S = \sum_{i=1}^k G_i$ and $T = \sum_{j=1}^\ell H_j$. Explain that $W = \sum_{i=1}^k X_i = S/(S + T)$ is a $\mathrm{Beta}(a, b)$, with $a = \sum_{i=1}^k a_i$ and $b = \sum_{j=1}^\ell b_j$, and show that

$$X_i = \frac{G_i}{S + T} = W \frac{G_i}{S} = W X_i' \quad \text{and} \quad Y_j = \frac{H_j}{S + T} = (1 - W) \frac{H_j}{T} = (1 - W) Y_j',$$

where $X' = (X_1', \ldots, X_k') \sim \mathrm{Dir}(a_1, \ldots, a_k)$ and $Y' = (Y_1', \ldots, Y_\ell') \sim \mathrm{Dir}(b_1, \ldots, b_\ell)$, independent of W. We learn that the long Dirichlet vector (X, Y) may be split into two separate Dirichlet vectors $X' = X/w$ and $Y' = Y/(1 - w)$, by conditioning on $W = w$.

(c). Now generalise to longer vectors. Let $X = (X_1, \ldots, X_k)$ be a $\mathrm{Dir}(a_1, \ldots, a_k)$, where each X_i is a vector $(X_{i,1}, \ldots, X_{i,m_i})$, with corresponding $a_i = (a_{i,1}, \ldots, a_{i,m_i})$. Write $W_i = \sum_{j=1}^{m_i} X_{i,j}$ for the sum over X_i. Explain first that $(W_1, \ldots, W_k) \sim \mathrm{Dir}(b_1, \ldots, b_k)$, where $b_i = \sum_{j=1}^{m_i} a_{i,j}$ is the sum over a_i. Show next that with $X_i' = X_i/W_i$, then $X_i' \sim \mathrm{Dir}(a_i)$, and that $X_1', \ldots, X_k'$ are independent of $(W_1, \ldots, W_k)$.

(d). The previous point implies that a long Dirichlet vector may be split into independent Dirichlet components, in several ways. In the setting above, starting with the long $(X_1, \ldots, X_k)$, suppose we condition on $(W_1, \ldots, W_k) = (w_1, \ldots, w_k)$, a given probability vector. With $X_i' = X_i/w_i$, show that $X_1', \ldots, X_k'$ are independent, with $X_i' \sim \mathrm{Dir}(a_i)$. This is the sorcerer's apprentice property of the Dirichlet distribution; pin it down, into subsets with given sums, and see that each part is a scaled Dirichlet again.

Ex. 1.22. *The Beta-binomial distribution.* Sometimes comparable binomial experiments may be modelled and analysed jointly, but where the success probability is not the same across studies. The $p = \Pr(\text{girl})$ may e.g. vary from family to family, see Story i.2.

(a). Suppose in general terms that $Y \,|\, p \sim \mathrm{binom}(n, p)$, and that p has a distribution with mean p_0 and standard deviation τ_0. Using the double expectation rule (1.2), show that

$$\mathrm{E}\, Y = np_0 \quad \text{and} \quad \mathrm{Var}\, Y = np_0(1-p_0) + n(n-1)\tau_0^2.$$

Hence the extra-binomial component of the variance, the $n(n-1)\tau_0^2$, becomes more noticeable with increasing n. The case of $\tau_0 = 0$ corresponds to the usual binomial.

(b). Suppose $Y \,|\, p \sim \mathrm{binom}(n, p)$ and that $p \sim \mathrm{Beta}(a, b)$. Show that this leads to the distribution

$$\begin{aligned} \Pr(Y = y) &= \int_0^1 \binom{n}{y} p^y (1-p)^{n-y} g(p, a, b)\, \mathrm{d}p \\ &= \binom{n}{y} \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \frac{\Gamma(a+y)\Gamma(b+n-y)}{\Gamma(n+a+b)} \quad \text{for } y = 0, 1, \ldots, n. \end{aligned}$$

Give formulae for the mean and variance of Y. For the special case of the uniform for p, show that all outcomes for Y are equally likely.

Ex. 1.23. *The Dirichlet-multinomial distribution.* Here we deal with the natural extension of the Beta-binomial setup of Ex. 1.22, from the case of two categories to more than two.

(a). Let $Y = (Y_1, \ldots, Y_k)$, for given probability vector $p = (p_1, \ldots, p_k)$, have a multinomial $(n, p_1, \ldots, p_k)$ model, as per Ex. 1.5. Assume then that the p is not fixed, but with p_i variances $\tau_{0,i}^2$ around mean $p_{0,i}$. Show that Y_i, marginally, has mean $np_{0,i}$ and variance $np_{0,i}(1-p_{0,i}) + n(n-1)\tau_{0,i}^2$.

(b). Let in particular $p \sim \mathrm{Dir}(cp_0)$, with parameters $cp_0 = (cp_{0,i}, \ldots, cp_{0,k})$. Show that

$$\mathrm{Var}\, Y_i = \{n + n(n-1)/(c+1)\} p_{0,i}(1-p_{0,i}) = \frac{c+n}{c+1} np_{0,i}(1-p_{0,i}),$$

with a clear overdispersion factor with respect to multinomial variation.

(c). Show that the marginal distribution of $(Y_1, \ldots, Y_k)$, now overdispersed compared to the multinomial, becomes

$$\begin{aligned}\bar{f}(y_1, \ldots, y_k) &= \int \frac{n!}{y_1! \cdots y_k!} p_1^{y_1} \cdots p_{k-1}^{y_{k-1}} (1 - p_1 - \cdots - p_{k-1})^{y_k} \\ &\qquad g(p_1, \ldots, p_{k-1}) \, \mathrm{d}p_1 \cdots \mathrm{d}p_{k-1} \\ &= \frac{n!}{y_1! \cdots y_k!} \frac{\Gamma(c)}{\Gamma(cp_{0,1}) \cdots \Gamma(cp_{0,k})} \frac{\Gamma(cp_{0,1} + y_1) \cdots \Gamma(cp_{0,k} + y_k)}{\Gamma(c + n)}.\end{aligned}$$

(d). For the case of Dirichlet parameters $cp_0 = (1, \ldots, 1)$, show that all outcomes $(y_1, \ldots, y_k)$ have the same probability, and find a formula for how many different outcomes there can be.

Laws of small numbers, Poisson, geometric, negative binomial

Ex. 1.24. *The Poisson distribution.* Counting a high number of events with small probabilities leads to the Poisson distribution, which we define here. A count variable Y is said to have the Poisson distribution with parameter $\theta > 0$ if

$$\Pr(Y = y) = \exp(-\theta)\theta^y / y! \quad \text{for } y = 0, 1, 2, \ldots.$$

We write $Y \sim \text{Pois}(\theta)$ to indicate this.

(a). Show that the probabilities indeed sum to 1. Verify next that $\mathrm{E}\, Y = \theta$, $\mathrm{E}\, Y(Y - 1) = \theta^2$ and show from this that the variance is equal to the mean. Show further that $\mathrm{E}\, Y(Y-1)(Y-2) = \theta^3$, $\mathrm{E}\, Y(Y-1)(Y-2)(Y-3) = \theta^4$, and with further algebra that for $W = (Y - \theta)/\sqrt{\theta}$, we have skew $= \mathrm{E}\, W^3 = 1/\sqrt{\theta}$, kurt $= \mathrm{E}\, W^4 - 3 = 1/\theta$. Show also that $\mathrm{Var}\,(Y - \theta)^2 = 2\theta^2 + \theta$.

(b). With $Y \sim \text{Pois}(\theta)$, what is the most probable outcome? What is the probability that Y is odd?

(c). Show that the sum of two independent Poisson variables is Poisson, with parameter equal to the sum of the two parameters. Generalise.

(d). Consider $Y \sim \text{binom}(n, p)$, and assume that n grows, while p becomes small, in the fashion of $np \to \theta$. Show that Y then tends to the $\text{Pois}(\theta)$ distribution, in the sense that the point probabilities converge. See also Ex. 2.8 for a fuller picture.

(e). In some event counting applications there are more zeros than predicted by the Poisson, leading naturally to a more general model with

$$\Pr(Y = 0) = p_0, \quad \Pr(Y = y) = (1 - p_0) \exp(-\theta)\theta^y / y! \,/\{1 - \exp(-\theta)\} \quad \text{for } y \geq 1.$$

Verify that these probabilities sum to 1, and find expressions for the mean and variance. This model is sometimes called the zero-inflated Poisson, since situations with $p_0 > \exp(-\theta)$ are prevalent, but also cases with $p_0 < \exp(-\theta)$ are allowed. Simulate say 1000 datapoints from the model with $\theta = 3.00$ and $p_0 = 0.25$, and check the histogram.

Ex. 1.25. *The geometric distribution.* Suppose Y has the distribution with point probabilities $f(y) = (1 - p)^{y-1} p$ for $y = 1, 2, \ldots$. This is the geometric distribution, and we write $Y \sim \text{geom}(p)$ to indicate this.

(a). Show that the probabilities $f(y)$ indeed sum to 1. Suppose independent experiments are carried out, each time with probability p that a certain event A takes place. With Y the first time A happens, show that $Y \sim \text{geom}(p)$.

(b). Show that Y has mean $1/p$ and variance $(1-p)/p^2$, via direct summation of $\sum_{y=1}^{\infty} yf(y)$ etc. If Y is the number of times you need to roll a six-sided die until it shows a '6', find the mean and the standard deviation.

(c). Another way of finding the mean and variance is as follows. With probability p, $Y = 1$; with complementary probability $1-p$, $Y = 1 + Y'$, with Y' having the same distribution as Y. Show that this leads to $\mathrm{E}\, Y = p + (1-p)(1+\mathrm{E}\, Y)$ and solve. Use this representation to also find the variance. Show also that $\mathrm{E}\,(Y - 1/p)^3 = (1-p)(2-p)/p^3$.

(d). Show that $\Pr(Y \geq y) = (1-p)^{y-1}$, derive a formula for $\Pr(Y \geq y_0 + y \,|\, Y \geq y_0)$, and comment.

(e). A simple related distribution is when one starts counting at 0, not at 1, so to speak. Show that with $Y \sim \text{geom}(p)$, as defined above, the variable $Y_0 = Y - 1$ has point probabilities $\Pr(Y_0 = y) = q^y p$ for $y = 0, 1, \ldots$, writing $q = 1 - p$. Show that Y_0 has mean $(1-p)/p$ and variance $(1-p)/p^2$.

(f). Suppose Y given p has geometric probabilities $(1-p)^{y-1}p$ for $y = 1, 2, \ldots$, but that p itself stems from a uniform distribution. Find the distribution for Y; note that this implies $1/2 + 1/6 + 1/12 + 1/20 + \cdots = 1$.

Ex. 1.26. *Time to last event.* Suppose independent geometric experiments are carried out by m individuals, say players throwing dice until they get the '6'. How long time does it take until the event in question has taken place, for all individuals?

(a). Let $Y_1, \ldots, Y_m$ be the time needed for the m individuals to see the event in question. Show that $F(y) = \Pr(Y_i \leq y) = 1 - (1-p)^y$ for $y = 1, 2, \ldots$. For $Z_m = \max(Y_1, \ldots, Y_m)$, show that its c.d.f. is $H_m(z) = F(z)^m$, and use this to find a formula for the median time to final event: $z^*_m = \log(1 - (\frac{1}{2})^{1/m})/\log(1-p)$ (which may be rounded to the nearest integer).

(b). For $p = 1/6$, compute also the mean $\mathrm{E}\, Z_m$, for say $m = 1, \ldots, 100$. Plot this, along with the median times z^*_m.

Ex. 1.27. *Mixing the Poisson.* Suppose observations come from Poisson mechanisms, but with different parameters, forming their own distribution. There are several versions and uses of such Poisson overdispersion models. (xx pointer to Poisson regression with overdispersion, perhaps in Ch5. xx)

(a). Suppose $Y \mid \theta \sim \text{Pois}(\theta)$ but that θ has a distribution with mean θ_0 and variance τ_0^2. Show that Y has mean θ_0 and variance $\theta_0 + \tau_0^2$.

(b). Specialise to the case of $\theta \sim \text{Gam}(a, b)$, see Ex. 1.9. Show that $\mathrm{E}\, Y = \theta_0 = a/b$ and that $\mathrm{Var}\, Y = \theta_0(1 + 1/b)$. Argue that with a large b we come back to pure Poisson. Show also that the marginal distribution of Y becomes

$$f(y, a, b) = \frac{\Gamma(a+y)}{\Gamma(a)\, y!} \frac{b^a}{(b+1)^{a+y}} = \frac{\Gamma(a+y)}{\Gamma(a)\, y!} \Big(\frac{b}{b+1}\Big)^a \Big(\frac{1}{b+1}\Big)^y \quad \text{for } y = 0, 1, 2, \ldots.$$

negative binomial

We are discovering the general *negative binomial* distribution in the process, of the form

$$g(y,a,p) = \frac{\Gamma(a+y)}{\Gamma(a)\,y!}(1-p)^y p^a \quad \text{for } y = 0,1,\ldots, \tag{1.5}$$

for parameters $a > 0$, $p \in (0,1)$; see Ex. 1.28 for more details.

Ex. 1.28. *The negative binomial.* We met the negative binomial distribution in Ex. 1.27 and now point to other features and constructions.

(a). Let X_1, X_2 be independent from the geometric distribution $q^x p$ for $x = 0,1,\ldots$, with $q = 1-p$. Show that $Y = X_1 + X_2$ has distribution $\Pr(Y = y) = (y+1)q^y p^2$, for $y = 0,1,\ldots$. For $Y = X_1 + X_2 + X_3$ a sum of three such independent geometric variables, show that $\Pr(Y = y) = \binom{y+2}{2} q^y p^3$ for $y = 0,1,\ldots$.

(b). Generalise to the case of $Y = X_1 + \cdots + X_a$, the sum of a independent geometric variables, each with $q^x p$ for $x = 0,1,\ldots$. Show that

$$\Pr(Y = y) = \binom{y+a-1}{a-1} q^y p^a = \frac{\Gamma(y+a)}{\Gamma(a)\,y!}(1-p)^y p^a \quad \text{for } y = 0,1,\ldots,$$

i.e. the negative binomial with parameters (a,p). Deduce that the number of ways in which one may find nonnegative numbers $x_1,\ldots,x_a$ with a given sum y is $\binom{y+a-1}{a-1} = (y+a-1)!/\{(a-1)!\,y!\}$. In how many ways may one find 5 nonnegative numbers with sum 100? And with 10 nonnegative numbers with sum 100?

(c). How do we know that the negative binomial probabilities (1.5) sum to one, also when the a is a non-integer? Deduce from this that

$$\sum_{y=0}^{\infty} \frac{\Gamma(y+a)}{\Gamma(a)} \frac{u^y}{a!} = \frac{1}{(1-u)^a} \quad \text{for } u \in (0,1).$$

Show that $\mathrm{E}\,Y = aq/p$, $\mathrm{Var}\,Y = aq/p^2$.

(d). (xx the step from $Y = X_1 + \cdots + X_a$, counting from zero, to $Y' = X_1' + \cdots + X_a'$, counting each from one, so that $Y' \geq a$. and just a bit more. reason for the negative binomial term. xx)

(e). In one of the episodes of the television series *Siffer* (NRK, 2011), programme leader Jo Røislien announced he would flip his coin and land 'krone' ten times in a row – which he then proceeded to do. He looked a bit tired, though; he had just kept on doing this, complete with his opening statement, until he had achieved the ten krone in a row event, and then showed only this crowning minute on tv. About how many times did he need to flip his coin, in total, before he (and his camera man) could show that final string of crowns? Simulate the process, and give a histogram of say 1,000 realisations.

Ex. 1.29. *Conditioning on Poisson sums and a generalised binomial.* We start with Poisson sums and see a connection to the binomial. Using generalised Poissons then lead to generalisations of the binomial.

(a). Let X and Y be independent Poissons with parameters θ_1, θ_2. Show that X given $X+Y = n$ is a binomal (n,p), with $p = \theta_1/(\theta_1+\theta_2)$. Generalise to the case of $X_1,\ldots,X_m$ being independent Poissons, with parameters $\theta_1,\ldots,\theta_m$. Show that their conditional distribution given $X_1 + \cdots + X_m = n$ is a multinomial with count n and probabilities $(p_1,\ldots,p_m)$, where $p_j = \theta_j/(\theta_1 + \cdots + \theta_m)$.

(b). To generalise the above, and in its turn also the binomial distribution, consider the gamma-mixed Poissons of 1.27. Specifically, let $X \,|\, \theta_1 \sim \text{Pois}(\theta_1)$ and $Y \,|\, \theta_2 \sim \text{Pois}(\theta_2)$, with gamma distributions $(c\theta_{0,1}, c)$ and $(c\theta_{0,2}, c)$ for θ_1 and θ_2; when c is large, this means tight concentration around $\theta_{1,0}, \theta_{0,2}$, and we're back to Poisson. Show that the conditional distribution of X given $X + Y = n$ may be written

$$f(x \,|\, n) \propto \binom{n}{x} \Gamma(c\theta_{0,1} + x)\Gamma(c\theta_{0,2} + n - x) \quad \text{for } x = 0, 1, \ldots, n,$$

where '$\propto$' means 'proportional to'. For a particular case, explain that with $c\theta_{0,1} = c\theta_{0,2} = 1$, X has the uniform distribution on $0, 1, \ldots, n$.

(c). It may not be easy to sum the terms above directly, to find the normalisation constant, but show via expressions in Ex. 1.22 that we in fact must have

$$f(x \,|\, n) = \binom{n}{x} \frac{\Gamma(c\theta_{0,1} + c\theta_{0,2})}{\Gamma(c\theta_{0,1})\Gamma(c\theta_{0,2})} \frac{\Gamma(c\theta_{0,1} + x)\Gamma(c\theta_{0,2} + n - x)}{\Gamma(c\theta_{0,1} + c\theta_{0,2} + n)}$$

for $x = 0, 1, \ldots, n$. In yet other words, we have reinvented the Beta-binomial distribution with gamma-mixing of Poisson parameters. Also, for large c, we are back to plain binomial (n, p), with $p = \theta_{0,1}/(\theta_{0,1} + \theta_{0,2})$.

(d). To invent other generalisations of the binomial, therefore, we might attempt other extensions of the Poisson. One such is to let X and Y have point probabilities proportional to $\theta_1^x/(x!)^\gamma$ and $\theta_2^x/(y!)^\gamma$, for some $\gamma \neq 1$; we work with such models in Ex. 4.34 and Story iv.7. Show here that X given $X + Y = n$ has the distribution

$$f(x \,|\, n) \propto \binom{n}{x}^\gamma p^x (1-p)^{n-x} \quad \text{for } x = 0, 1, \ldots, n,$$

where $p = \theta_1/(\theta_1 + \theta_2)$. Compute these probabilities, for say $n = 50$, $p = 0.33$, and some values of γ around 1. Check then numerically $np_0 = \text{E}\, X$ and the dispersion ratio $\rho = \text{Var}\, X/\{np_0(1-p_0)\}$, to learn that there is overdispersion and underdispersion, compared to the binomial case, for $\gamma < 1$ and $\gamma > 1$, respectively.

Moments, moment-generating functions, characteristic functions

Ex. 1.30. *Moments.* Consider a random variable X with c.d.f. F. Its mean is $\text{E}\, X = \int x \,\text{d}F(x)$, and we may of course define higher moments.

(a). Use results of Ex. A.16 to show that $\text{E}\, X^k$, first seen as $\int y \,\text{d}G_k(y)$, the mean of the variable $Y = X^k$ with distribution G_k inherited from F, is also the same as $\int x^k \,\text{d}F(x)$; thus there is no ambiguity there.

(b). Suppose $h(x)$ is convex, on a domain in which X falls with probability 1. This means $h(\lambda x + (1-\lambda)y) \leq \lambda h(x) + (1-\lambda)h(y)$ for all $\lambda \in [0, 1]$. Show that $h(\text{E}\, X) \leq \text{E}\, h(X)$. This is Jensen's inequality. Show also that equality holds only in the linear case, with $h(x) = a + bx$.

(c). For $r < s$, show that $(\text{E}\,|X|^r)^{1/r} \leq (\text{E}\,|X|^s)^{1/s}$, i.e. $h(r) = (\text{E}\,|X|^r)^{1/r}$ is a non-decreasing function in r. In particular, note that if $|X|$ has a finite s-moment, then all moments of smaller order are also finite. Illustrate by computing and graphing $h(r)$ and $\log h(r)$ for the case of $X \sim \text{Expo}(1)$.

(d). For X a standard normal, show that $\mathrm{E}\,|X|^p = 2^{p/2}\Gamma(\frac{1}{2}(p+1))/\sqrt{\pi}$ for $p \geq 0$, a formula that also can be written $(\frac{1}{2})^{p/2}\Gamma(p+1)/\Gamma(\frac{1}{2}p+1)$. Compute and graph the function $h(r)$ for this case. For p an even integer, the formula simplifies to $(\frac{1}{2})^{p/2}p!/(\frac{1}{2}p)!$.

(e). For a random variable X with finite fourth moment, we have defined its skewness skew and kurtosis kurt in Ex. 1.3. Give expressions skew $= h_3(\mu_1,\mu_2,\mu_3)$ and kurt $= h_4(\mu_1,\mu_2,\mu_3,\mu_4)$ in terms of the moments $\mu_j = \mathrm{E}\,X^j$, and also expressions skew $= h_3^*(\mu_2^*,\mu_3^*)$ and kurt $= h_4^*(\mu_2^*,\mu_3^*,\mu_4^*)$ in terms of the centralised moments $\mu_j^* = \mathrm{E}\,(X-\mu_1)^j$.

Ex. 1.31. *Moment-generating functions.* For a random variable Y, with distribution P, its m.g.f. is

$$M(t) = \mathrm{E}\,\exp(tY) = \int \exp(ty)\,\mathrm{d}P(y),$$

defined for each t at which the expectation exists. The moment-generating function is useful for finding and characterising distributions, for finding their moments, for handling the distributions of sums of variables, and in connection with distributional limits. When Y has a density $f(y)$ (with respect to Lebesgue measure), we have $M(t) = \int \exp(ty)f(y)\,\mathrm{d}y$, and if it is discrete with pointmasses $f(y)$ for sample space S, say, then $M(t) = \sum_{y\in S}\exp(ty)f(y)$. The expectation operator is more general, however, and $M(t)$ is perfectly defined also for intermediate cases where Y can have both discrete and continuous parts; see Ex. A.16.

(a). For a standard normal $Y \sim \mathrm{N}(0,1)$, show that $M(t) = \exp(\frac{1}{2}t^2)$. When $Y \sim \mathrm{N}(\xi,\sigma^2)$, derive $M(t) = \exp(\xi t + \frac{1}{2}\sigma^2 t^2)$.

(b). A moment-generating function has that name since it generates moments; we indeed have $M'(0) = \mathrm{E}\,Y$, $M''(0) = \mathrm{E}\,Y^2$, etc., under mild conditions, see Ex. A.36 for details. Use this to find the first four moments of the standard normal.

(c). For $Y \sim \mathrm{Expo}(\theta)$, show that $M(t) = 1/(1-t/\theta)$, for $t < \theta$.

(d). For $Y \sim \mathrm{Gam}(a,b)$, with density $\{b^a/\Gamma(a)\}y^{a-1}\exp(-by)$, show that $M(t) = \{b/(b-t)\}^a$, for $t < b$. In particular, $M(t) = 1/(1-t)^a$ for $\mathrm{Gam}(a,1)$.

(e). Suppose Y is equal to zero with probability 0.90, but a standard normal with probability 0.10. Find the $M(t)$, and generalise.

(f). For the binomial (n,p), show that $M(t) = \{1-p+p\exp(t)\}^n$.

(g). For $Y \sim \mathrm{Pois}(\theta)$, find $M(t) = \exp\{\theta(e^t-1)\}$. Use this, with Ex. 1.27, to find $M(t)$ also for the negative binomial (a,p). (xx hm, should give the formula here. xx)

(h). Let $Y = \pm 1$ with probabilities $\frac{1}{2},\frac{1}{2}$. Show that

$$M(t) = \cosh(t) = \tfrac{1}{2}(e^t + e^{-t}) = 1 + (1/2)t^2 + (1/4!)t^4 + (1/6!)t^6 + \cdots.$$

(i). For the uniform distribution on the unit interval, show that $M(t) = \{\exp(t)-1\}/t$, for $t \neq 0$, and with $M(0) = 1$. For Y having the uniform distribution on the $[-1,1]$ interval, show that

$$M(t) = \frac{\exp(t)-\exp(-t)}{2t} = \frac{\sinh t}{t},$$

and that this function may be written as the infinite sum $1 + (1/3!)t^2 + (1/5!)t^4 + \cdots$.

Ex. 1.32. *Distribution of sums via moment-generating functions.* Importantly, the m.g.f. $M(t) = \mathrm{E}\,\exp(tY)$, if it exists in a neighbourhood around zero, characterises the distribution; variables whose m.g.f.s are identical in such a neighbourhood have identical distributions. See details in Ex. A.36.

(a). Suppose $M(t) = (\frac12)^{25}\{1+\exp(t)\}^{25}$. What is the underlying distribution? Similarly, if $M(t) = 0.99 + 0.01\,\exp(\frac12 t^2)$, what is the distribution?

(b). With X and Y being independent, with m.g.f.s M_1 and M_2, show that the sum $Z = X + Y$ has m.g.f. $M_X(t)M_Y(t)$.

(c). With X and Y independent and standard normal, show that $X+Y$ is $\mathrm{N}(0,2)$. Redo the questions of Ex. 1.2.

(d). Redo a part of Ex. 1.24, showing that sums of independent Poissons are Poisson. Similarly, if X, Y, Z are independent binomials $(n_1, p), (n_2, p), (n_3, p)$, identify the distribution of $W = X + Y + Z$.

(e). If $X_1, \ldots, X_n$ are independent gamma variables, with parameters $(a_1, b), \ldots, (a_n, b)$, find the distribution of $X_1 + \cdots + X_n$.

the Laplace distribution

Ex. 1.33. *The Laplace distribution.* The Laplace or double exponential distribution, in its simplest form, has density $f_0(y) = \frac12 \exp(-|y|)$, on the real line; note the cusp at its centre point zero.

(a). Let V_1 and V_2 be independent standard exponentials. Show that $Y = V_1 - V_2$ has this density $f_0(y)$. Deduce from this that its m.g.f. is $M_0(t) = 1/(1-t^2)$, for $|t| < 1$.

(b). More generally, consider $Y = V_1 - V_2$ where these two are independent and Expo(θ). Show that Y has density $f(y) = \frac12\theta\exp(-\theta\,|y|)$, with zero mean and variance $2/\theta^2$. Also, show that its m.g.f. is $M(t) = 1/\{1-(t/\theta)^2\}$ for $|t| < \theta$. The Laplace with variance 1 is hence that with $\theta = \sqrt{2}$.

(c). Suppose X for given σ is a $\mathrm{N}(0,\sigma^2)$, but that the variance $V = \sigma^2$ has some distribution. Show that the m.g.f. for such a normal scale mixture becomes $M(t) = \mathrm{E}\,\exp(tX) = M_V(\frac12 t^2)$, where $M_V(s)$ is the m.g.f. for V. In particular, show that if $X \,|\, \sigma \sim \mathrm{N}(0,\sigma^2)$ and $\sigma^2 \sim \mathrm{Expo}(1)$, then X has the Laplace distribution with variance 1.

(d). If $X \,|\, V \sim \mathrm{N}(0,V)$, and V has density $g(v)$, show that X has density $f(x) = \int_0^\infty \phi(x/v^{1/2})(1/v^{1/2})g(v)\,\mathrm{d}v$. Translate the result above to the interesting formula

$$\begin{aligned}\int_0^\infty \frac{1}{(2\pi)^{1/2}}\exp\Big(-\tfrac12\frac{x^2}{v}\Big)\frac{\exp(-v)}{v^{1/2}}\,\mathrm{d}v &= 2\int_0^\infty \frac{1}{(2\pi)^{1/2}}\exp\{-(\tfrac12 x^2/w^2 + w^2)\}\,\mathrm{d}w\\ &= \tfrac12\sqrt{2}\exp(-\sqrt{2}|x|).\end{aligned}$$

Use this to find a formula for the integral $\int_0^\infty \exp\{-(av^2 + b/v^2)\}\,\mathrm{d}v$.

Ex. 1.34. *Summing uniforms.* [xx to be polished. xx] can we simply invert $\{\exp(t) - 1\}^n/t^n$, for $n = 2, 3, \ldots$? $\mathrm{E}\,N = e$ for N, the number of uniforms to sum to reach 1. xx]

(a). Let $U_1, U_2, \ldots$ be i.i.d. uniforms on the unit interval. Show that the densities for $U_1 + U_2$ and $U_1 + U_2 + U_3$ may be written

$$f_2(x) = \begin{cases} x & \text{when } 0 \le x \le 1, \\ 2 - x & \text{when } 1 \le x \le 2, \end{cases}$$

and

$$f_3(x) = \begin{cases} \frac{1}{2}x^2 & \text{when } 0 \le x \le 1, \\ \frac{1}{2}(-2x^2 + 6x - 3) & \text{when } 1 \le x \le 2, \\ \frac{1}{2}(3 - x)^2 & \text{when } 2 \le x \le 3. \end{cases}$$

Show that f_2 is continuous, but that its derivative has a jump at position $x = 1$. Show however that f_3 is smoother, with a continuous derivative.

(b). Generalise the above to the case of $X = U_1 + \cdots + U_n$; show that its density may be written

$$f_n(x) = \frac{1}{(n-1)!} \sum_{j=0}^{\lfloor x \rfloor} (-1)^j \binom{n}{j} (x - j)^{n-1}$$

for $0 \le x \le n$, where $\lfloor x \rfloor$ is the so-called floor function, the largest integer equal to or to the left of x (so $\lfloor 2.99 \rfloor = 2$, $\lfloor 3.00 \rfloor = 3$, $\lfloor 3.01 \rfloor = 3$). In particular, show that $f_n(x) = x^{n-1}/(n-1)!$ for $0 \le x \le 1$. This is sometimes called the Irwin–Hall distribution, from the related and concurrently published papers Irwin (1927), Hall (1927).

(c). Draw the densities f_1, f_2, f_3, f_4, f_5 in a diagram, and comment on their forms.

(d). Let now N be the number of uniforms U_j needed in order for their sum to exceed 1. Show that

$$\Pr(N > n) = \Pr(U_1 + \cdots + U_n < 1) = F_n(1),$$

with F_n the cumulative for f_n. Also show that in fact $F_n(1) = 1/n!$. Use this to deduce that $\Pr(N = n) = 1/\{n\,(n-2)!\}$ and that $\mathrm{E}\,N = e$. As a follow-up computational exercise, simulate a high number of such N to calculate the mathematical constant e to three decimal places. Check if you get Ibsen's and Tolstoy's birth years right.

Ex. 1.35. *Characteristic functions.* Above we have found multiple uses for moment-generating functions, as in Ex. 1.31 and 1.32. Another very useful transform is the characteristic function, worked with in App. A, see Ex. A.39, which is also a crucial technical tool in the development of the large-sample theory of Ch. 2. For a variable with distribution P, its definition is $\varphi(t) = \mathrm{E}\,\exp(itX) = \int \{\cos(tx) + i\sin(tx)\}\,\mathrm{d}P(x)$, with $\mathrm{d}P(x)$ to be read as $f(x)\,\mathrm{d}x$ when P has a density f. Notably, this $\varphi(t)$ always exists, also for distributions without means, etc. We note here the inversion theorem from Ex. A.46, that if $\int |\varphi(t)|\,\mathrm{d}t$ is finite, then there is a density $f(x) = (2\pi)^{-1} \int \exp(-itx)\varphi(t)\,\mathrm{d}t$, see Ex. A.46, which is also seen to be continuous.

(a). For the standard normal, show that $\varphi(t) = \exp(-\frac{1}{2}t^2)$, and that $\varphi(t) = \exp(-\frac{1}{2}\sigma^2 t^2)$ for the $\mathrm{N}(0, \sigma^2)$.

(b). Show that $\varphi(t)$ is real, and then equal to $\mathrm{E}\cos(tX)$, if and only if the distribution is symmetric around zero.

(c). For U a unit exponential, show that $\varphi(t) = 1/(1-it)$. Deduce that $X = U - V$, with its Laplace distribution, has $\varphi(t) = 1/(1+t^2)$, and that

$$\tfrac{1}{2}\exp(-|x|) = (2\pi)^{-1}\int \cos(tx)\frac{1}{1+t^2}\,\mathrm{d}t.$$

(d). Show from this, essentially by changing from (x,t) to (t,x), that the standard Cauchy has $\varphi(t) = \exp(-|t|)$.

(e). In Ex. 1.18(d) we saw that if a ball is kicked from distance r to an infinitely long straight line, with random angle A being uniform on $(-\frac{1}{2}\pi, \frac{1}{2}\pi)$, then the position where the ball crosses the line is $X = rX_0$, where X_0 is Cauchy. Now assume r is not fixed, but comes from a unit exponential distribution. Show that X, a scale mixture of Cauchys, has characteristic function $\phi(t) = 1/(1+|t|)$.

(f). For V a uniform on $[-1,1]$, show $\varphi(t) = (\sin t)/t$. Prove from this that even though $\int (\sin t)/t\,\mathrm{d}t = \pi$, the integral $\int |(\sin t)/t|\,\mathrm{d}t$ must be infinite. Show that

$$f(x) = (2\pi)^{-1}\int \cos(tx)\Big(\frac{\sin t}{t}\Big)^2\,\mathrm{d}t$$

is equal to the triangular density on $[-2,2]$.

(g). Just as for m.g.f.s, sums of independent components are associated with products of characteristic functions. Show that if X, Y, Z are independent, with $\varphi_1(t), \varphi_2(t), \varphi_3(t)$, then $X+Y+Z$ has $\varphi_1(t)\varphi_2(t)\varphi_3(t)$ as its characteristic function. When $X_1, \ldots, X_n$ are i.i.d. with characteristic function $\varphi(t)$, show that $(X_1 + \cdots + X_n)/\sqrt{n}$ has $\varphi_n(t) = \varphi(t/\sqrt{n})^n$. We learn in Ch. 2 that if the X_i have mean zero and finite variance σ^2, then $\varphi_n(t) \to \exp(-\frac{1}{2}\sigma^2 t^2)$. In this connection, you may numerically compute

$$f_n(x) = (2\pi)^{-1}\int \cos(tx)\Big(\frac{\sin t/\sqrt{n}}{t/\sqrt{n}}\Big)^n\,\mathrm{d}t$$

and verify that it tends to the $\mathrm{N}(0, 1/12)$ density; cf. Ex. 1.34.

(h). Show that there cannot be two i.i.d. variables with sum having the uniform distribution. (xx but pointer to story with $U = V_1 + V_2$, two independent but with different distributions. xx)

Ex. 1.36. *Characterisation theorems.* (xx to be repaired carefully. start with symmetry, then remove that assumption. may push these two things to Ch2. xx)

(a). If X, Y are independent with the same distribution, and $(X+Y)/\sqrt{2} \sim X$, show that X must be a zero-mean normal. (xx pointer to CLT. xx)

(b). For another characterisation lemma, suppose that $(X+Y)/2 \sim X$, and show by induction that this implies the curious property that with $X_1, \ldots, X_n$ i.i.d. from such a model, then the sample mean $\bar{X}_n$ has the same distribution for every n. Show that X must be Cauchy. Why is this not contradicting the LLN?

Ex. 1.37. *Cumulants and the cumulant-generating function.* For a variable X with m.g.f. $M(t)$, assumed finite in an interval around zero, it is sometimes fruitful to work with the *cumulant-generating function* $K(t) = \log M(t)$. When expanded in a power series around zero, with

cumulants

$$K(t) = K'(0)t + \tfrac{1}{2}K''(0)t^2 + \cdots = \sum_{j=1}^{\infty} \frac{\kappa_j}{j!} t^j,$$

we call the coefficients $\kappa_j = K^{(j)}(0)$ the *cumulants* of the distribution.

(a). For the $\mathrm{N}(\xi, \sigma^2)$ distribution, show that $K(t) = \xi t + \frac{1}{2}\sigma^2 t^2$. For the unit exponential distribution, show $K(t) = t + t^2/2 + t^3/3 + \cdots$, with $\kappa_j = (j-1)!$ for $j \geq 1$.

(b). If X has mean ξ, write $X = \xi + X_0$, Show that $K(t) = \xi t + K_0(t)$, where $K_0(t) = \log M_0(t)$, with M_0 the m.g.f. of the zero-mean variable X_0. Hence the cumulants κ_j for X are the same as the cumulants $\kappa_{j,0}$ for X_0, for $j \geq 2$.

(c). Via successive derivatives of $M(t) = \exp\{K(t)\}$, show that

$$\begin{aligned} M' &= MK', \\ M'' &= M\{K'' + (K')^2\}, \\ M''' &= M\{K''' + 3K''K' + (K')^3\}, \\ M'''' &= M\{K'''' + 4K'''K' + 6K''(K')^2 + 3(K'')^2 + (K')^4\}. \end{aligned}$$

Write ξ and σ^2 for the mean and variance of X. From the equations, show first that $K'(0) = \kappa_1 = \xi$ and that $K''(0) = \kappa_2 = \sigma^2$, and next that

$$\mathrm{E}\, X^3 = \kappa_3 + 3\xi\kappa_2 + \xi^3, \quad \mathrm{E}\, X^4 = \kappa_4 + 4\xi\kappa_3 + 6\xi^2\kappa_2 + 3\kappa_2^2 + \xi^4.$$

(d). Skewness and kurtosis were defined in Ex. 1.3. We may now find useful expressions for these in terms of the cumulants. Show first via formulae above that

$$\mathrm{E}\,(X-\xi)^3 = \kappa_3, \quad \mathrm{E}\,(X-\xi)^4 = \kappa_4 + 3\sigma^4.$$

Explain that this leads to $\mathrm{skew}(X) = \kappa_3/\sigma^3$, $\mathrm{kurt}(X) = \kappa_4/\sigma^4$.

Ex. 1.38. *Generating functions.* Moment-generating functions, studying distributions via the transformation $M(t) = \mathrm{E}\exp(tX)$, have several close relatives, which might be more convenient for certain classes of distributions. It is e.g. common to use *Laplace transformations* $L(s) = \mathrm{E}\exp(-sX)$ for distributions on $[0, \infty)$, then studied for $s \geq 0$. Here we work through the basic properties of *generating functions*, primarily used for distributions on the nonnegative integers. If $\Pr(Y = j) = p_j$, for $j = 0, 1, 2, \ldots$, define $G(s) = \mathrm{E}\, s^Y = \sum_{j=0}^{\infty} p_j s^j = p_0 + p_1 s + p_2 s^2 + \cdots$, called the generating function for that distribution, or for variables having that distribution.

Laplace transforms

generating functions

(a). Show that $G(s) = M(\log s)$, for s such that the latter exists. Demonstrate that $G(s)$ is finite, for $|s| < 1$, and also for $s = 1$. Find the generating functions for (i) the binomial (n, p); (ii) the Poisson with parameter θ; (iii) the geometric with $\Pr(Y = j) = q^{j-1}p$ for $j \geq 1$, with $q = 1 - p$; answers are

$$G_1(s) = (1 - p + ps)^n, \quad G_2(s) = \exp\{-\theta(1-s)\}, \quad G_3(s) = \frac{ps}{1-qs},$$

the latter valid for $|s| < 1/q$.

(b). Returning to the general case, give an expression for $G'(s)$, show that $G'(1) = \mathrm{E}\, Y$, and that $G''(1) = \mathrm{E}\, Y(Y-1)$. Find the mean and variance for the Poisson using generating functions.

(c). Suppose X and Y are random variables taking on values in $\{0, 1, 2, \ldots\}$, and that their generating functions are equal, on an interval around zero. Show that X and Y must have identical distributions.

(d). Show that if X, Y, Z are independent, with generating functions G_1, G_2, G_3, then the generating function for $X + Y + Z$ is $G_1(s)G_2(s)G_3(s)$. Show from this, and the previous point, yet again, that a sum of independent Poissons is a Poisson.

Ex. 1.39. *Getting Y from two copies of X.* Let X_1, X_2 be independently drawn from some distribution on the nonnegative integers, and from these draw Y from the binomial $(X_1 + X_2, \frac{1}{2})$.

(a). Writing ξ and σ^2 for the mean and variance for the X distribution, set up formulae for the conditional mean and variance of Y, given X_1, X_2. From these show that $\mathrm{E}\, Y = \xi$, $\mathrm{Var}\, Y = \frac{1}{2}\xi + \frac{1}{2}\sigma^2$.

(b). Write $G_0(s) = \mathrm{E}\, s^X$ for the generating function of the X_i. Show that the generating function for Y may be written $G(s) = G_0(\frac{1}{2} + \frac{1}{2}s)^2$. When $X_i \sim \mathrm{binom}(m, p)$, show from this that $Y \sim \mathrm{binom}(2m, \frac{1}{2}p)$. Then show that if the X_i are Poisson, then Y reproduces the same distribution.

Ex. 1.40. *Sums of random lengths and the compound Poisson.* Let $X_1, X_2, \ldots$ be i.i.d., from a distribution with mean ξ, variance σ^2, and m.g.f. $M_0(t)$. Consider then a random sum of these random elements; $Z = \sum_{i=1}^{N} X_i$, where N has some distribution with mean λ, variance τ^2, and generating function $G(s) = \mathrm{E}\, s^N$. We define Z as zero if $N = 0$.

(a). Show that Z has m.g.f. $M(t) = G(M_0(t))$. Show that Z has mean $\lambda\xi$ and variance $\lambda\sigma^2 + \xi^2\tau^2$.

compound Pòisson

(b). Consider the so-called compound Poisson variable $Z = \sum_{i=1}^{N} X_i$, where the X_i are i.i.d. with m.g.f. $M_0(s)$ and $N \sim \mathrm{Pois}(\lambda)$. Show that the m.g.f. of Z may be written

$$\mathrm{E}\, \exp(tZ) = \mathrm{E}\, M_0(s)^N = \exp[\lambda\{M_0(t) - 1\}],$$

and that this leads to mean $\lambda\xi$ and variance $\lambda(\xi^2 + \sigma^2)$. Find an expression for $\mathrm{E}\, \exp(tZ)$ for the particular case of $X_i \sim \mathrm{Gam}(a, b)$.

(c). Consider now $Z = \sum_{i=1}^{N} X_i$, where the X_i are i.i.d. unit exponentials. For given N, the Z is a $\mathrm{Gam}(N, 1)$. With $N \sim \mathrm{geom}(p)$, show that $Z \sim \mathrm{Expo}(p)$. Generalise.

(d). Using terminology and results from Ex. 1.37, show that the cumulant-generating function for the compound Poisson is $K(t) = \lambda\{M_0(t) - 1\}$, with cumulants $\kappa_j = \lambda\, \mathrm{E}\, X^j$. Show from this that $\mathrm{skew}(Z) = (1/\lambda)^{1/2}\gamma_3$ and $\mathrm{kurt}(Z) = (1/\lambda)\gamma_4$, with $\gamma_3 = \mathrm{E}\, X^3/(\mathrm{E}\, X^2)^{3/2}$ and $\gamma_4 = \mathrm{E}\, X^4/(\mathrm{E}\, X^2)^2$. For a Poisson sum $Z = \sum_{i=1}^{N} X_i$ of i.i.d. standard normals, so that $Z \,|\, (N = n) \sim \mathrm{N}(0, n)$, find the variance and kurtosis.

Ex. 1.41. *The logarithmic distribution.* Consider a variable X with point probabilities $\Pr(X = x) = c(p)^{-1}p^x/x$ for $x = 1, 2, \ldots$, with p a parameter in $(0, 1)$. This distribution is sometimes called the logarithmic distribution.

(a). Show that we must have $c(p) = -\log(1-p)$. Find expressions for the m.g.f. $M(t)$, its mean, and its variance. Comment on the cases where p is close to zero, or close to one. Show also for its generating function that $G(s) = \mathrm{E}\, s^X = c(ps)/c(p)$.

(b). Consider $Z = \sum_{i=1}^N X_i$, with the X_i being i.i.d. with this logarithmic distribution, and N is Poisson, with parameter expressed as $\lambda c(p)$. Find the mean and variance of Z, and show that its distribution is a negative binomial. We learn that the negative binomial is inside the class of compound Poissons.

The multinormal, the t, the chi-squared, the F

Ex. 1.42. *Mean and variance matrix for a random vector.* (xx calibrate with Story vii.1. xx) Consider a random vector of length p, say $Y = (Y_1, \ldots, Y_p)^{\mathrm{t}}$. Its mean vector is defined as the vector of means, i.e. $\xi = \mathrm{E}\, Y = (\mathrm{E}\, Y_1, \ldots, \mathrm{E}\, Y_p)^{\mathrm{t}}$, and its *variance matrix* Σ (also often called the covariance matrix), of dimension $p \times p$, has the variances on the diagonal and the covariances outside.

(a). Show that the elements of Σ are $\mathrm{cov}(Y_i, Y_j)$, for $i, j = 1, \ldots, p$, and that we may write $\Sigma = \mathrm{E}\,(Y-\xi)(Y-\xi)^{\mathrm{t}}$. If we transform Y to $Z = AY + b$, with A a $m \times p$ matrix and b a vector of length m, show that $\mathrm{E}\, Z = A\xi + b$ and that $\mathrm{Var}\, Z = A\Sigma A^{\mathrm{t}}$. Explain that this generalises the usual rule $\mathrm{Var}\,(aY) = a^2 \mathrm{Var}\, Y$ for one-dimensional variables.

(b). Let X_1, X_2, X_3, X_4 be i.i.d. standard normals. Set up the variance matrix for $(X_1, X_1 + X_2, X_1 + X_2 + X_3, X_1 + X_2 + X_3 + X_4)^{\mathrm{t}}$.

(c). For the multinomial model studied in Ex. 1.5, with $Y = (Y_1, \ldots, Y_k)^{\mathrm{t}}$ counting the number of n events that fall in categories $1, \ldots, k$, with probabilities $p_1, \ldots, p_k$, show that $\mathrm{E}\, Y = np$, where p is the vector of $p_1, \ldots, p_k$. Show also that its variance matrix can be written $\Sigma = D - pp^{\mathrm{t}}$, where D is diagonal with elements p.

(d). For the multinomial model, with $\mathbf{1}$ the vector $(1, \ldots, 1)^{\mathrm{t}}$, show that $\mathbf{1}^{\mathrm{t}}\Sigma\mathbf{1} = 0$. Explain that this is related to the linear relationships $\sum_{j=1}^k p_j = 1$ and $\sum_{j=1}^k Y_j = n$. The matrix is hence not of full rank, and not invertible. Work however with the shorter vector $Y_0 = (Y_1, \ldots, Y_{k-1})^{\mathrm{t}}$, and show that its variance matrix may be written $\Sigma_0 = D_0 - p_0 p_0^{\mathrm{t}}$, where $p_0 = (p_1, \ldots, p_{k-1})^{\mathrm{t}}$ and D_0 is diagonal with these elements. Show that $\Sigma_0^{-1} = D_0^{-1} + (1/p_k)\mathbf{1}_0\mathbf{1}_0^{\mathrm{t}}$, where $\mathbf{1}_0$ is the vector of $k-1$ 1s. For any vector $x = (x_1, \ldots, x_k)$ with sum zero, and with x_0 being the shortened version $(x_1, \ldots, x_{k-1})$, demonstrate that $x_0^{\mathrm{t}}\Sigma_0^{-1}x_0 = \sum_{j=1}^k x_j^2/p_j$.

Ex. 1.43. *The multinormal distribution.* Let $Y = (Y_1, \ldots, Y_p)^{\mathrm{t}}$ be a random vector of length p. We say that it is multinormally distributed, with mean vector ξ and variance matrix Σ, which needs to be positive definite, provided its joint density can be written

$$f(y) = (2\pi)^{-p/2}|\Sigma|^{-1/2} \exp\{-\tfrac{1}{2}(y-\xi)^{\mathrm{t}}\Sigma^{-1}(y-\xi)\},$$

where the domain for y is all of $\mathbb{R}^p$. We write $Y \sim \mathrm{N}_p(\xi, \Sigma)$ to indicate this distributon.

(a). Show that $\int y f(y)\, \mathrm{d}y$ indeed is equal to ξ, so calling it the mean vector is appropriate. Show also that $\mathrm{E}\,(Y-\xi)(Y-\xi)^{\mathrm{t}}$, calculated from the density, is equal to Σ.

(b). Show that if $Y \sim \mathrm{N}_p(\xi, \Sigma)$, then $Y - \xi \sim \mathrm{N}_p(0, \Sigma)$.

(c). Assume now that A is an invertible $p \times p$ matrix, and consider the transformation $Z = AY$. Show that if $Y \sim \mathrm{N}_p(\xi, \Sigma)$, then $Z = AY \sim \mathrm{N}_p(A\xi, A\Sigma A^{\mathrm{t}})$.

(d). By the spectral decomposition theorem of linear algebra, there is an orthonormal matrix P, with $PP^{\mathrm{t}} = I = P^{\mathrm{t}}P$, such that $P\Sigma P^{\mathrm{t}} = D = \mathrm{diag}(\lambda_1, \ldots, \lambda_p)$, with these values being the eigenvalues of Σ. Show that $Z = P(Y - \xi)$ has components $Z_1, \ldots, Z_p$ which are independent, with $Z_j \sim \mathrm{N}(0, \lambda_j)$.

(e). Show that a vector Y is multinormal if and only if all linear combinations are normal. In particular, if $Y \sim \mathrm{N}_p(\xi, \Sigma)$, then $V = c^{\mathrm{t}}Y = c_1Y_1 + \cdots + c_pY_p$ is normal $\mathrm{N}(c^{\mathrm{t}}\xi, c^{\mathrm{t}}\Sigma c)$. [xx need to say something careful about allowing constants to be seen as normal, with zero variance. xx]

(f). Generalise point (c) to state that with any matrix A, of size say $q \times p$, the transformed $Z = AY$ is a multinormal $\mathrm{N}_q(A\xi, A\Sigma A^{\mathrm{t}})$.

(g). For the binormal case, with means ξ_1, ξ_2, standard deviations σ_1, σ_2, and correlation ρ, show that the density may be written

$$f(x,y) = \frac{1}{2\pi}\frac{1}{\sigma_1\sigma_2(1-\rho^2)^{1/2}}\exp\Big[-\tfrac{1}{2}\frac{1}{1-\rho^2}\Big\{\Big(\frac{x-\xi_1}{\sigma_1}\Big)^2 + \Big(\frac{y-\xi_2}{\sigma_2}\Big)^2 - 2\rho\Big(\frac{x-\xi_1}{\sigma_1}\Big)\Big(\frac{x-\xi_1}{\sigma_1}\Big)\Big\}\Big].$$

Show that X and Y are independent if and only if the correlation is zero.

(h). We learn that the situation is easy and clean for the multinormal case, where independence is equivalent to zero correlation. This is in general more complicated. (i) Let $X \sim \mathrm{N}(0,1)$ and $Y = DX$, with D a random sign with equal probabilities for $-1, 1$. Show that Y also is standard normal, that the correlation is zero, but that they are dependent. (ii) For another example, again with $X \sim \mathrm{N}(0,1)$, let $Y = X$ if $|X| \le a$ but $Y = -X$ if $|X| > a$. Show that Y is standard normal. Show that the correlation $\rho(a)$ goes from -1 to 1 as a goes from zero to infinity; thus there is an a_0 for which the correlation is zero. Also, find this a_0 numerically (answer: 1.537). In these zero-correlation constructions, with normal marginals, the point is that there is not joint binormality.

Ex. 1.44. *The multinormal and conditional distributions.* Consider a multinormally distributed vector, of length $p+q$, blocked into subvectors of sizes p and q. Let us write this as

$$\begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix} \sim \mathrm{N}_{p+q}\Big(\begin{pmatrix} \xi_1 \\ \xi_2 \end{pmatrix}, \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}\Big).$$

(a). By carrying out a linear transformation, and using results from Ex. 1.43, show that

$$Z = Y_1 - \Sigma_{12}\Sigma_{22}^{-1}Y_2 \sim \mathrm{N}_p(\xi_1 - \Sigma_{12}\Sigma_{22}^{-1}\xi_2, \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}),$$

and that this Z is independent of Y_2.

(b). Show that the distribution of Y_1 given $Y_2 = y_2$ must be multinormal. Derive the important formulae for the conditional mean and variance,

$$\mathrm{E}\,(Y_1 \,|\, y_2) = \xi_1 + \Sigma_{12}\Sigma_{22}^{-1}(y_2 - \xi_2), \qquad \mathrm{Var}\,(Y_1 \,|\, y_2) = \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}.$$

Note that the conditional mean is a linear function in y_2 and that the conditional variance matrix is constant, not depending on y_2.

(c)**.** Now study the simplest two-dimensional prototype case, with

$$\begin{pmatrix} X \\ Y \end{pmatrix} \sim \mathrm{N}_2(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}).$$

Show that $Y \,|\, x \sim \mathrm{N}(\rho x, 1-\rho^2)$ and that $X \,|\, y \sim \mathrm{N}(\rho y, 1-\rho^2)$. Discuss implications for situations where an easy to measure X might be a proxy for a harder to come by Y. How can you estimate Y from X, and with what precision?

(d)**.** In generalisation of this special binormal case, consider a binormal

$$\begin{pmatrix} Y \\ X \end{pmatrix} \sim \mathrm{N}_2(\begin{pmatrix} \xi \\ \xi_0 \end{pmatrix}, \begin{pmatrix} \sigma^2 & \rho\sigma\sigma_0 \\ \rho\sigma\sigma_0 & \sigma_0^2 \end{pmatrix}).$$

In particular, the X is seen as stemming from a $\mathrm{N}(\xi_0, \sigma_0^2)$ distribution. Show that $Y \,|\, (X = x)$ is normal, with constant variance $\sigma^2(1-\rho^2)$ and linear mean function $\mathrm{E}\,(Y \,|\, x) = \xi + \rho(\sigma/\sigma_0)x$. This is essentially the linear regression model, for $Y \,|\, x$, here derived as a consequence of binormality. Without pretensions of having precise numbers, guess the binormal parameters in your population for (X, Y) being height and weight, and deduce from this a model predicting a person's weight from his or her height.

(e)**.** Generalise this to situations with a $(p+1)$-dimensional normal distribution for $(Y, X_1, \ldots, X_p)$, with Y seen as the main outcome, influenced by covariates $X_1, \ldots, X_p$. Show that Y given these covariates is normal, with a constant variance, and find the linear conditional mean function $\mathrm{E}\,(Y \,|\, X_1, \ldots, X_p)$. This is essentially the linear multiple regression model; see Ex. 3.34.

(f) (xx a bit here on regression towards the mean. xx)**.**

Ex. 1.45. *How tall is Nils?* Assume that the heights of Norwegian men above the age of twenty follow the normal distribution $\mathrm{N}(\xi, \sigma^2)$ with $\xi = 180$ cm and $\sigma = 9$ cm.

(a)**.** Given this information only, what is your point estimate of his height, and what is your 95 percent prediction interval?

(b)**.** Assume now that you learn that his four brothers are actually 195 cm, 207 cm, 196 cm, 200 cm tall, and furthermore that correlations between brothers' heights in the population of Norwegian men is equal to $\rho = 0.80$. Use this information about his four brothers to revise your initial point estimate of his height, and provide the updated 95 percent prediction interval. Is Nils a statistical outlier in his family?

(c)**.** Suppose that Nils has n brothers and that you learn their heights. Give formulae for the updated normal parameters ξ_n and σ_n, in the conditional distribution of his height given these extra pieces of information. Use this to clarify the following statistical point: Even if you get to know all facts concerning 99 brothers, there should be a limit to your confidence in what you may infer about Nils.

Ex. 1.46. *Simulating from the multinormal.* For given ξ and Σ, how can we simulate realisations y from the $\mathrm{N}_p(\xi, \Sigma)$? We take it for granted here that yes, we can simulate from the one-dimensional $\mathrm{N}(\xi, \sigma^2)$, e.g. via `xi + sigma*rnorm(100)` in `R`. For simplicity we take $\xi = 0$ here, without loss of generality.

there are simulation algorithms in packages, like `mvrnorm` om `R`, but the recipe here works better for large dimension

(a). From details of Ex. 1.43, there is a spectral decomposition, with $P\Sigma P^{\mathrm{t}} = D$, with the diagonal D having the eigenvalues of Σ and P orthonormal, i.e. $PP^{\mathrm{t}} = I_n = P^{\mathrm{t}}P$, with I_n the identity matrix. Show that with $D^{1/2}$ the diagonal with the square-roots of the eigenvalues, and with $\Sigma^{1/2} = P^{\mathrm{t}}D^{1/2}P$, this is a matrix which when squared becomes Σ.

(b). Explain from this that with $N = (N_1, \ldots, N_p)^{\mathrm{t}}$ a string of i.i.d. standard normals, $Y = \Sigma^{1/2}N \sim \mathrm{N}_p(0, \Sigma)$.

(c). Explain from Ex. 1.44 that a perhaps long multinormal vector (Y_1, Y_2), with lengths p and q, can be simulated in blocks; first Y_1, then $Y_2 \,|\, Y_1$. This is sometimes practical in high dimensions.

(d). To attempt this in practice, simulate say $n = 50$ copies $Y_i = (Y_{i,1}, \ldots, Y_{i,p})^{\mathrm{t}}$, in dimension say $p = 10$, with all variables being i.i.d. standard normal. From these compute the empirical variance matrix Σ_n, with empirical variances on the diagonal and covariances outside, i.e. $\sigma_{j,k} = (n-1)^{-1}\sum_{i=1}^n (Y_{i,j} - \bar{Y}_j)(Y_{i,k} - \bar{Y}_k)$ for $j, k = 1, \ldots, p$. Then simulate from the multinormal $\mathrm{N}_p(0, \Sigma_n)$. Check to see if their empirical variance matrix is close to Σ_n again.

Ex. 1.47. *The chi-squared.* This exercise goes through some basic properties of the chi-squared; see also Ex. 1.52 for its eccentric cousin, the noncentral or eccentric chi-squared.

(a). We say that a nonnegative variable X has the chi-squared distribution, with degrees of freedom m, and write $X \sim \chi^2_m$ for this, when its density takes the form

$$g_m(x) = \frac{1}{2^{m/2}\Gamma(m/2)} x^{m/2-1} \exp(-\tfrac{1}{2}x) \quad \text{for } x > 0.$$

Show that its m.g.f. becomes $M(t) = (1-2t)^{-m/2}$, for $t < \frac{1}{2}$. Show further that $\mathrm{E}\, X = m$, $\mathrm{Var}\, X = 2m$, and that the skewness, i.e. $\mathrm{E}\, W^3$ with $W = (X - \mathrm{E}\, X)/(\mathrm{Var}\, X)^{1/2} = (X-m)/(2m)^{1/2}$, is $(8/m)^{1/2}$. From the density, show $\mathrm{E}\,(\chi^2_m)^p = 2^p\Gamma(m/2+p)/\Gamma(m/2)$.

(b). Via the m.g.f., show the simple and basic convolution property for the chi-squared, that if $X_1, \ldots, X_n$ are independent and chi-squared distributed with degrees of freedom $m_1, \ldots, m_n$, then the sum $Z = \sum_{i=1}^n X_i$ is chi-squared too, with degrees of freedom $\sum_{i=1}^n m_i$. A generalisation is given in Ex. 1.52.

(c). If N is standard normal, show that $X = N^2 \sim \chi^2_1$. Establish that if $X \sim \chi^2_m$, with m a natural number, then it may be represented as $X = N_1^2 + \cdots + N_m^2$, in terms of independent standard normals $N_1, \ldots, N_m$. Note, however, that the χ^2_m with the density $g_m(x)$ above may be used also when m is not a natural number.

(d). There are connections between the chi-squared and the Gamma distribution (see Ex. 1.9). Show that the $\mathrm{Gam}(\frac{1}{2}m, \frac{1}{2})$ is the χ^2_m; that if $Y \sim \mathrm{Gam}(a, b)$, then $X = 2bY \sim \chi^2_{2a}$; and that if $Z \sim \chi^2_m$, then $\frac{1}{2}Z \sim \mathrm{Gam}(\frac{1}{2}m, 1)$.

(e). Consider independent $X \sim \chi^2_a$ and $Y \sim \chi^2_b$. Show, perhaps via Gamma distribution ratios in Ex. 1.12, that $R = X/(X+Y) \sim \mathrm{Beta}(\frac{1}{2}a, \frac{1}{2}b)$.

(f). When $X \sim \chi^2_m$, show that $\mathrm{E}\log X = \log 2 + \psi(\frac{1}{2}m)$, where $\psi(x) = \Gamma'(x)/\Gamma(x)$ is the digamma function.

(g). Consider $W = AB$, a product of independent standard normals. Use (1.2) to show that its m.g.f. is $M(t) = 1/(1-t^2)^{1/2}$, for $|t| < 1$. Deduce from this that W has the same distribution as $\frac{1}{2}(K - L)$, where K and L are independent with χ^2_1 distributions.

(h). Then consider $W = AB + CD$, where A, B, C, D are independent standard normals. Show that W has density $\frac{1}{2}\exp(-|w|)$, i.e. the Laplace distribution worked with in Ex. 1.33.

(i). Consider the second degree equation $x^2 + Bx + C = 0$ from school. If there are many such equations, with B and C being independent standard normal, how many of these equations will have both roots real?

Ex. 1.48. *Tell me about X and $X + Y$.* In exercises above we have seen that sums of independent normals, Poissons, chi-squares are respectively normal, Poisson, chi-square. Sometimes we meet questions going the other way.

(a). Suppose X and Y are independent, and that you learn the distributions of X and $Z = X + Y$. Show that Y must have m.g.f. $M_Y(t) = M_Z(t)/M_X(t)$. As an illustration, suppose $X \sim \mathrm{N}(0, 1)$ and $X + Y \sim \mathrm{N}(0, 2)$. Show that $Y \sim \mathrm{N}(0, 1)$.

(b). Explain similarly that if X and Y are independent, with X and $X+Y$ being Poisson, then Y is also Poisson.

(c). Suppose X and Y are independent, that $X \sim \chi^2_a$ and that $Z = X + Y \sim \chi^2_{a+b}$. Show that the only possibility is then the expected one, that $Y \sim \chi^2_b$.

(d). Consider $X_1, \ldots, X_n$ i.i.d. standard normal. Writing as usual $\bar{X}$ for the sample mean, we know that $\sqrt{n}\bar{X}$ is also standard normal. Show that the vector of $X_i - \bar{X}$ is independent of $\bar{X}$. Writing $Z_0 = \sum_{i=1}^n (X_i - \bar{X})^2$, show from $\sum_{i=1}^n X_i^2 = Z_0 + n\bar{X}^2$ that $Z_0 \sim \chi^2_{n-1}$. (xx calibrate with other things; this is giving a simpler proof of the lemma below, with no orthogonal transformations. xx)

(e). (xx for later, perhaps in Ch3: also nice for p-dimensional things with $Q(\theta) = \sum_{j=1}^k (y_j - \theta)^{\mathrm{t}} \Sigma_j^{-1} (y_j - \theta)$, where $Q_{\min} = Q(\widehat{\theta}) \sim \chi^2_{(k-1)p}$, independent of $\widehat{\theta} = A^{-1} \sum_{j=1}^k \Sigma_j^{-1} y_j$. xx)

Ex. 1.49. *Orthonormal transformations.* [xx check and calibrate with chi-squared things, to get the order right. restructure text. xx] We have seen in Ex. 1.43 that a multinormal vector can be sent via a linear transformation to independent one-dimensional normal components, and vice versa. This also leads to useful characterisation and representation theorems involving independence. In the present exercise we shall e.g. find a proof that the sample mean $\bar{Y}$ and the sample variance statistic $Z = \sum_{i=1}^n (Y_i - \bar{Y})^2$ are independent; this fact, which does not hold outside the normal family, is actively used in Ex. 3.4, and will also be utilised elsewhere, like in Ex. 3.6 xx).

(a). Suppose $X = (X_1, \ldots, X_n)$ is a vector with i.i.d. and standard normal components, and let A be an orthonormal matrix, which means $AA^{\mathrm{t}} = I = A^{\mathrm{t}}A$. In yet other words, each row of A and each column of A has length 1, rows are ortogonal, as well as columns. Show that $Y = AX$ must have components $Y_1, \ldots, Y_n$ which are also i.i.d. and standard normal. – Here you may also use the general transformation formula of Ex. 1.12.

(b). To exemplify the above, show that if X_1, X_2 are independent and standard normal, then also Y_1, Y_2, where

$$\begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix} = \begin{pmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & -1/\sqrt{2} \end{pmatrix} \begin{pmatrix} X_1 \\ X_2 \end{pmatrix} = \begin{pmatrix} (X_1 + X_2)/\sqrt{2} \\ (X_1 - X_2)/\sqrt{2} \end{pmatrix},$$

must be independent and standard normal.

(c). When A is orthonormal, show that it preserves length, so $\|Au\| = \|u\|$, for any vector u; here $\|u\|$ is Euclidean length, so $\|u\|^2 = u_1^2 + \cdots + u_n^2$.

(d). Let again $X_1, \ldots, X_n$ be i.i.d. standard normals. Construct an orthogonal matrix A by letting its first row be $(1/\sqrt{n}, \ldots, 1/\sqrt{n})$, and define $Y = AX$. Then show that $Y_1 = \sqrt{n}\bar{X} = \sum_{i=1}^n X_i/\sqrt{n}$, and that $Z_0 = \sum_{i=1}^n (X_i - \bar{X})^2 = \sum_{i=2}^n Y_i^2$.

(e). Conclude from this that (i) $\sqrt{n}\bar{X} \sim \mathrm{N}(0,1)$, (ii) $Z_0 \sim \chi^2_{n-1}$, and (iii) $\bar{X}$ and Z_0 are idependent. This was proven more directly in Ex. 1.48 above, via decompositions of chi-squares.

(f). Show that this implies the following classical and important properties, starting with an independent sample $Y_1, \ldots, Y_n$ from the $\mathrm{N}(\xi, \sigma^2)$: The statistics $\bar{Y}$ and $Z = \sum_{i=1}^n (Y_i - \bar{Y})^2$ are independent, with $\bar{Y} \sim \mathrm{N}(\xi, \sigma^2/n)$ and $Z \sim \sigma^2 \chi^2_{n-1}$. Show also from this that the classical empirical variance

$$\widehat{\sigma}^2 = \frac{1}{n-1} \sum_{i=1}^n (Y_i - \bar{Y})^2 \tag{1.6}$$

is unbiased for the population variance, i.e. $\mathrm{E}\,\widehat{\sigma}^2 = \sigma^2$. Construct an unbiased estimator for σ, of the type $c_n \widehat{\sigma}$. [xx point to generalisations for the linear regression model. xx]

(g). Consider the general multinormal distribution $Y \sim \mathrm{N}_p(\xi, \Sigma)$, with invertible Σ. Show that $K = (Y - \xi)^{\mathrm{t}} \Sigma^{-1} (Y - \xi) \sim \chi^2_p$. Suppose $Y = y_{\mathrm{obs}}$ is observed, that Σ is known, but ξ unknown. Give a confidence region R such that $\xi \in R$ with probability 90 percent. How does this region shrink, if you observe 100 vectors from the multinormal, rather than merely 1?

Ex. 1.50. *The t distribution.* Consider independent variables $X \sim \mathrm{N}(0,1)$ and $K \sim \chi^2_m$. The ratio $t = X/(K/m)^{1/2}$ is then said to have the t distribution, with m degrees of freedom. We write $t \sim t_m$ to indicate this.

(a). Looking at the denominator first, show that K/m has mean 1 and variance $2/m$, indicating that it will be close to 1 with high probability when m is moderate or big. Explain that this implies that t is not far from having a standard normal distribution for growing m. Finer details regarding this are in Ch. 2.

(b). Find the mean and variance of t.

(c). Show that its density can be written

$$g_m(x) = \frac{\Gamma((m+1)/2)}{\Gamma(m/2)} \frac{1}{\sqrt{m\pi}} \frac{1}{(1 + x^2/m)^{(m+1)/2}}.$$

Show that $g_m(x)$ tends to the standard normal density $\phi(x)$ as m increases, matching the intuition above. Show also that the Cauchy distribution, see Ex. 1.18, is the t_1 distribution.

(d). Find also the skewness and kurtosis for the t_m distribution. In particular, show that the latter is kurt $= 6/(m-4)$ for $m > 4$.

(e). Assume $Y_1, \ldots, Y_n$ are i.i.d. $\mathrm{N}(\xi, \sigma^2)$. With $\widehat{\sigma}$ the empirical standard deviation, from (1.6), show that

$$t = (\bar{Y} - \xi)/(\widehat{\sigma}/\sqrt{n})$$

has the t-distribution with $n-1$ degrees of freedom. This is the classic t-statistic dating all the way back to Student (1908),

Ex. 1.51. *The noncentral t distribution.* Consider $X \sim \mathrm{N}(\xi, 1)$, independent of $K \sim \chi^2_m$. The ratio $t = X/(K/m)^{1/2} = (\xi + \mathrm{N}(0,1))/(\chi^2_m/m)^{1/2}$ is then said to follow that noncentral t distribution, written $t \sim t_m(\xi)$, with degrees of freedom m and noncentrality parameter ξ. Its distribution is a complex one, but both the density $g_m(t, \xi)$ and c.d.f. $G_m(t, \xi)$ are implemented in various software packages, as e.g. `dt(t,m,ksi)` and `pt(t,m,ksi)` in `R`.

(a). Explain that the classical t distribution of Ex. 1.50 corresponds to the special case of $\xi = 0$. Via arguments in that exercise, explain also that the $t_m(\xi)$ becomes close to a $\mathrm{N}(\xi, 1)$ for growing m.

(b). Find expressions for the mean and variance of the $t_m(\xi)$.

(c). Consider again an i.i.d. sample $Y_1, \ldots, Y_n$ from $\mathrm{N}(\xi, \sigma^2)$, with sample mean $\bar{Y}$ empirical standard deviation $\widehat{\sigma}$, as in Ex. 1.50. Show that $t = \sqrt{n}\bar{Y}/\widehat{\sigma}$ has the $t_{n-1}(\xi/\sigma)$ distribution. We shall need this for studying the power of the t test in Ex. 4.12.

Ex. 1.52. *The noncentral chi-squared.* Consider also the so-called noncentral chi-squared distribution, say $K \sim \chi^2_m(\lambda)$, with λ the excentre or eccentricity parameter; the case of $\lambda = 0$ corresponds to the ordinary $K \sim \chi^2_m$. It is the distribution of $Y_1^2 + \cdots + Y_m^2$, where the Y_i are independent normals, with $Y_i \sim \mathrm{N}(\xi_i, 1)$, and $\lambda = \sum_{i=1}^m \xi_i^2$.

(a). Show that the m.g.f. of $K \sim \chi^2_m(\lambda)$ may be written

$$M(t) = \mathrm{E}\exp(tK) = \frac{\exp\{\lambda t/(1-2t)\}}{(1-2t)^{m/2}} \quad \text{for } t < \tfrac{1}{2}.$$

(b). Also the noncentral chi-squared distributions have convolution properties, generalising those of Ex. 1.47. If $K_i \sim \chi^2_{m_i}(\lambda_i)$, and these are independent, for $i = 1, \ldots, n$, show that $\sum_{i=1}^n K_i$ is another noncentral chi-squared, with degrees of freedom $\sum_{i=1}^n m_i$ and excentre parameter $\sum_{i=1}^n \lambda_i$.

(c). Another property which can be established, remarkbly without yet having seen the density of a $\chi^2_m(\lambda)$, is the following, using Ex. 1.48: if X and Y and independent, with $X \sim \chi^2_{m_1}(\lambda_1)$ and $X + Y \sim \chi^2_{m_1+m_2}(\lambda_1 + \lambda_2)$, then by necessity $Y \sim \chi^2_{m_2}(\lambda_2)$.

(d). Its density can be expressed in several ways; show that this is one such valid formula:

$$f(k, m, \lambda) = \sum_{j=0}^{\infty} \{\exp(-\tfrac{1}{2}\lambda)(\tfrac{1}{2}\lambda)^j/j!\}\, g_{m+2j}(k),$$

where $g_{m+2j}(k)$ is the χ^2_{m+2j} density. In other words, the noncentral chi-squared is a Poisson mixture of central chi-squared distributions. Show that this entails the representation $K \,|\, (J = j) \sim \chi^2_{m+2j}$, where $J \sim \mathrm{Pois}(\frac{1}{2}\lambda)$. Also non-integer values of m are allowed here.

(e). Establish that for $K \sim \chi^2_m(\lambda)$, we have $\mathrm{E}\, K = m + \lambda$ and $\mathrm{Var}\, K = 2m + 4\lambda$. Show also that the skewness of K becomes $2^{3/2}(m + 3\lambda)/(m + 2\lambda)^{3/2}$. What is required in order for this skewness to tend to zero?

(f). With N a standard normal, let $K = (\lambda^{1/2} + N)^2$, which has the $\chi^2_1(\lambda)$ distribution. Consider the normalised variable

$$\frac{K - (1 + \lambda)}{(2 + 4\lambda)^{1/2}} = \frac{N^2 + 2\lambda^{1/2} N - 1}{(2 + 4\lambda)^{1/2}}.$$

Work out its m.g.f. and show that it tends to $\exp(\frac{1}{2}t^2)$ for growing λ.

(g). More generally, with $K \sim \chi^2_m(\lambda)$, work out a formula for the m.g.f. $M(t)$ for $Z = \{K - (m + \lambda)\}/(2m + 4\lambda)^{1/2}$. For any fixed m, show that $M(t) \to \exp(\frac{1}{2}t^2)$ as λ grows, and comment on this finding.

Ex. 1.53. *Noncentral chi-squared for empirical variances.* We saw in Ex. 1.49 that if $X_1, \ldots, X_n$ are i.i.d. $\mathrm{N}(a, 1)$, with common a, then $Z = \sum_{i=1}^n (X_i - \bar{X})^2 \sim \chi^2_{n-1}$, with consequences for the empirical variance estimator. Here are some fruitful generalisations.

(a). Let the X_i have non-identical means, $X_i \sim \mathrm{N}(a_i, 1)$. Show that $Z \sim \chi^2_{n-1}(\lambda)$, with noncentrality parameter $\lambda = \sum_{i=1}^n (a_i - \bar{a})^2$.

(b). Assume now that $X_i \sim \mathrm{N}(a_i, 1/m_i)$ for $i = 1, \ldots, n$, perhaps reflecing sample sizes m_i for different groups, and with $M = \sum_{i=1}^n m_i$. Consider $Z = \sum_{i=1}^n m_i (X_i - \widetilde{X})^2$, with $\widetilde{X} = \sum_{i=1}^n (m_i/M) X_i$. Show that $Z = \sum_{i=1}^n m_i X_i^2 - M\widetilde{X}^2$, and that its distribution is a $\chi^2_{n-1}(\lambda)$, with $\lambda = \sum_{i=1}^n m_i (a_i - \widetilde{a})^2$, where $\widetilde{a} = \sum_{i=1}^n (m_i/M) a_i$.

(c). Let $X \sim \mathrm{N}_p(\xi, \Sigma)$. Show that $Z = X^{\mathrm{t}} \Sigma^{-1} X \sim \chi^2_p(\xi^{\mathrm{t}} \Sigma^{-1} \xi)$.

Ex. 1.54. *The F distribution.* As we have seen Ex. 1.49, with a normal sample $X_1, \ldots, X_n$, the distribution of the classical empirical variance estimator $\widehat{\sigma}^2 = (n - 1)^{-1} \sum_{i=1}^n (X_i - \bar{X})^2$ is governed by $\widehat{\sigma}^2 \sim \sigma^2 \chi^2_m / m$, where $m = n - 1$ is the degrees of freedom. Suppose there are two independent samples, from normal distributions $\mathrm{N}(\xi_1, \sigma_1^2)$ and $\mathrm{N}(\xi_2, \sigma_2^2)$, of sample sizes n_1 and n_2, with estimators $\widehat{\sigma}_1^2$ and $\widehat{\sigma}_2^2$.

(a). Let $\rho = \sigma_1/\sigma_2$, the ratio of standard deviations. For the ratio of the two empirical variances, show that

$$R^2 = \frac{\widehat{\sigma}_1^2}{\widehat{\sigma}_2^2} = \rho^2 F, \quad \text{where } F \sim \frac{\chi^2_{m_1}/m_1}{\chi^2_{m_2}/m_2},$$

with degrees of freedom $m_1 = n_1 - 1$ and $m_2 = n_2 - 1$, and with the two chi-squareds being independent. We say that F has the F distribution, or Fisher distribution, with degrees of freedom (m_1, m_2), and write $F \sim F(m_1, m_2)$.

(b). Show that when $F \sim F(m_1, m_2)$, then $1/F \sim F(m_2, m_1)$. Show furthermore that

$$\mathrm{E}\, F = \frac{m_2}{m_2 - 2}, \quad \mathrm{E}\, F^2 = \frac{m_1 + 2}{m_1} \frac{m_2^2}{(m_2 - 2)(m_2 - 4)},$$

these expressions being finite when $m_2 > 2$ and $m_2 > 4$, respectively. Find also an expression for the variance. Verify that both $\mathrm{E}\, F$ and $\mathrm{E}\, F^2$ tend to 1 as the degrees of freedom increase.

(c). The main aspects of the F distribution have been worked out, above, from the constructive definition given in (a), without actually needing any formula for its density; also, probabilities are found using software packages, like `pf(x,m1,m2)` in `R`. Once in a while one needs the density function, however. Show first that the cumulative function can be written

$$\Pr(F \le x) = \Pr(\chi^2_{m_1}/m_1 \le x\chi^2_{m_2}/m_2) = \int_0^\infty G(xy(m_1/m_2), m_1)g(y, m_2)\,\mathrm{d}y,$$

in terms of the cumulative $G(\cdot, m)$ and densiy $g(\cdot, m)$ of the χ^2_m. Then take the derivative to get

$$h(x, m_1, m_2) = \int_0^\infty g(xy(m_1/m_2), m_1)y(m_1/m_2)g(y, m_2)\,\mathrm{d}y.$$

Complete the math to land at

$$h(x, m_1, m_2) = \frac{\Gamma(\frac12(m_1+m_2))}{\Gamma(\frac12 m_1)\Gamma(\frac12 m_2)}(m_1/m_2)^{m_1/2}\frac{x^{m_1/2-1}}{\{1+(m_1/m_2)x\}^{(m_1+m_2)/2}}.$$

The exponential family class

Ex. 1.55. *The exponential family class.* Many parametric models fall under the wide umbrella of the exponential family class, which we treat in this and the folllowing exercises. This will be properly generalised and extended down the road, but we start with this definition: Suppose Y has model density of the form

exponential family

$$\begin{aligned} f(y,\theta) &= \exp\{\theta_1 T_1(y) + \cdots \theta_p T_p(y) - k(\theta_1,\ldots,\theta_p)\}h(y) \\ &= \exp\{\theta^{\mathrm{t}} T(y) - k(\theta)\}h(y), \end{aligned} \tag{1.7}$$

for appropriate functions $T_1(y),\ldots,T_p(y)$ and $h(y)$; the $k(\theta)$ function is there to secure integration to one. We require that there is no fixed linear relationship among $T_1(y),\ldots,T_p(y)$, and also that the support of the distribution, the smallest closed set having probability 1, is the same, for all parameter values. We then say Y is of the exponential family class, with *data functions* $T(y) = (T_1(y),\ldots,T_p(y))^{\mathrm{t}}$ and *natural parameters* $\theta = (\theta_1,\ldots,\theta_p)^{\mathrm{t}}$.

(a). Before we start developing the general theory for the full class, we verify that a few classic models are under its umbrella. For the following models, write the model density in a form matching (1.7). (i) $Y \sim \mathrm{binom}(n,p)$. (ii) $Y \sim \mathrm{Pois}(\theta)$. (iii) $Y \sim \mathrm{Beta}(a,b)$. (iv) $Y \sim \mathrm{Gam}(a,b)$. (v) $Y \sim \mathrm{N}(\xi,\sigma^2)$, with known σ (and see Ex. 1.56). (vi) Let (X,Y,Z) be trinomial (n,p,q,r), with $r = 1-p-q$. First work with $f(x,y,p,q)$, and then with the submodel where $p = a^2$, $q = 2a(1-a)$, $r = (1-a)^2$.

(b). Show that we must have

$$k(\theta) = \log\Big(\int \exp\{\theta^{\mathrm{t}} T(y)\}h(y)\,\mathrm{d}\nu(y)\Big),$$

assumed to be finite for at least some θ, and with $\mathrm{d}\nu(y)$ the dominating measure (often enough simply $\mathrm{d}y$, for models with continuous densities). Let in fact Θ be the set of θ such that $k(\theta)$ is finite, called the natural parameter region. Show that Θ is a convex set.

natural parameter region

(c). The score function of a model is $u(y,\theta) = \partial \log f(y,\theta)/\partial\theta$; see Ex. 5.14 for more on this. For the class of models studied here, show that $u(y,\theta) = T(y) - \xi(\theta)$, where

$$\xi(\theta) = \frac{\partial k(\theta)}{\partial\theta} = \frac{\int T(y)\exp\{\theta^{\rm t}T(y)\}h(y)\,{\rm d}y}{\int \exp\{\theta^{\rm t}T(y)\}h(y)\,{\rm d}y}.$$

Show that the score function must have mean zero, which here means $\mathrm{E}_\theta\, T(Y) = \xi(\theta)$. Show also that $\mathrm{Var}_\theta\, T(Y) = \partial^2 k(\theta)/\partial\theta\partial\theta^{\rm t}$, giving variances and covariances of the $T_j(Y)$ in one matrix formula.

(d). We come back to general likelihood theory in Ch. 5, but for now show that if $Y_1,\ldots,Y_n$ are i.i.d. from the exponential family density, then the logarithm of the joint density can be written

$$\ell_n(\theta) = \sum_{i=1}^n \log f(Y_i,\theta) = n\{\theta^{\rm t}\bar T - k(\theta)\},$$

with $\bar T = (1/n)\sum_{i=1}^n T(Y_i)$ the vector of averages $\bar T_j = (1/n)\sum_{i=1}^n T_j(y_i)$. If this is a family with say $p = 3$ parameters, and $n = 10000$, then the full relevant information required for computing the ℓ_n function is captured in the 3 averages $\bar T_1, \bar T_2, \bar T_3$. This is related to the sufficiency concept, returned to in Ch. 5. Show also that this $\ell_n(\theta)$ is a concave function.

Ex. 1.56. *The normal and binormal exponential family members.* The normal and binormal (and multinormal) classes of distributions belong under the exponential family umbrella.

(a). Consider first the normal (ξ,σ^2) model. Explain that its density may be written

$$f(y) = \frac{1}{(2\pi)^{1/2}}\exp\Big(-\tfrac12\frac{y^2}{\sigma^2} + \frac{\xi y}{\sigma^2} - \tfrac12\frac{\xi^2}{\sigma^2}\Big),$$

and argue from this that the normal is of exponential family type, with natural parameters ξ/σ^2 and $1/\sigma^2$, and data functions equivalent to (y, y^2).

(b). Consider then (X,Y) binormal, with its five parameters $\xi_1,\xi_2,\sigma_1,\sigma_2,\rho$, see Ex. 1.43. Show that the distribution is of the exponential family type, with natural parameters

$$\frac{\xi_1}{(1-\rho^2)\sigma_1^2}, \frac{1}{(1-\rho^2)\sigma_1^2}, \frac{\xi_2}{(1-\rho^2)\sigma_2^2}, \frac{1}{(1-\rho^2)\sigma_2^2}, \frac{\rho}{(1-\rho^2)\sigma_1\sigma_2}.$$

In the terminology of exponential families, identify also the associated $T_1(x,y),\ldots,T_5(x,y)$.

(c). Show that $(\xi_1,\xi_2,\sigma_1,\sigma_2,\rho)$ is in one-to-one correspondence with the five exponential class natural parameters above.

(d). Suppose now that $\xi_1 = \xi_2$ and $\sigma_1 = \sigma_2$, so this distribution for (X,Y) has now three parameters. Show that it is inside the exponential family, and identify its natural parameters.

Ex. 1.57. *A log-linear density model on the unit interval.* The machinery of the exponential family class makes it easy to construct new models, starting with relevant data functions T_j, in the language of Ex. 1.55.

(a). For densities on the unit interval, start with data function $T(y) = y - \frac{1}{2}$, to define $f(y,\theta) = \exp\{\theta T(y) - k(\theta)\}$ for $y \in [0,1]$. Find the $k(\theta)$, a formula for the c.d.f. $F(y,\theta)$, and the mean and variance for Y having this density.

(b). Then go on to

$$f(y,\theta_1,\theta_2) = \exp\{\theta_1(y-\tfrac{1}{2}) + \theta_2(y-\tfrac{1}{2})^2 - k(\theta_1,\theta_2)\},$$

i.e. using $T_j(y) = (y-\frac{1}{2})^j$ as data functions, for $j = 1,2$. Set up an integral function for $k(\theta_1,\theta_2)$ and some algorithms for determining the means, variances, and covariance, for $T_1(Y), T_2(Y)$, for any given (θ_1,θ_2).

(c). Spell out the necessary details for the third order log-linear density model, which uses $T_j(y) = (y-\frac{1}{2})^j$ for $j = 1,2,3$. Again, set up algorithms so that you may compute the means, variances, covariances for T_1, T_2, T_3, for any $(\theta_1,\theta_2,\theta_3)$.

Yet other models: log-normal, Weibull, Gompertz, Gumbel, et al.

Ex. 1.58. *The log-normal distribution.* Starting with $X \sim \mathrm{N}(\xi,\sigma^2)$, the variable $Y = \exp(X)$ is said to be a *log-normal*, and we write $Y \sim \mathrm{logN}(\xi,\sigma^2)$ to indicate this.

the log-normal distribution

(a). Consider the view that the distribution should or could have been named the exponormal instead – would you agree? Show that with $Y \sim \mathrm{logN}(\xi,\sigma^2)$, its mean and variance are $\exp(\xi + \frac{1}{2}\sigma^2)$ and $\{\exp(2\sigma^2) - \exp(\sigma^2)\}\exp(2\xi)$.

(b). Show that its density may be written $\phi_\sigma(\log y - \xi)/y = \sigma^{-1}\phi(\sigma^{-1}(\log y - \xi))/y$, for $y > 0$. Find its mode.

(c). Assume that $Y \,|\, \xi \sim \mathrm{logN}(\xi,\sigma^2)$, and that $\xi \sim \mathrm{N}(\xi_0,\tau^2)$. Show that marginally, $Y \sim \mathrm{logN}(\xi_0, \sigma^2+\tau^2)$. Make explicit the connection to Ex. 1.64.

(d). Show that a product of independent log-normals is log-normal. Suppose $Y_1,\ldots,Y_n$ are i.i.d. from the $\mathrm{logN}(\xi,\sigma^2)$ distribution. Explain what happens to their harmonic mean, $Z_n = (Y_1\cdots Y_n)^{1/n}$.

(e). Assume a random time variable T has the $\mathrm{logN}(0,1)$ distribution. Find a formula for its hazard rate $h(t)$, and show that $h(t) \doteq (\log t)/t$ for growing t. Plot the exact hazard rate, along with its approximation, and comment.

Ex. 1.59. *The Weibull distribution.* The Weibull distribution, with positive parameters (a,b), has c.d.f. $F(t) = 1 - \exp\{-(t/a)^b\}$ for $t \geq 0$. The b is called the shape parameter, with a a scale parameter. The Weibull generalises the exponential distribution, which is the special case of $b = 1$. Other parametrisations are sometimes convenient, as with $1 - \exp(-ct^b)$.

Ernst Hjalmar Waloddi Weibull

(a). Find a formula for the median, and more generally for the q-quantile $F^{-1}(q)$. Show that the density can be written $f(t) = \exp\{-(t/a)^b\}bt^{b-1}/a^b$ for $t > 0$. Find also a formula for the hazard rate, and draw this in a diagram, for $b = 0.9, 1.0, 1.1$, say for $a = 1$.

(b). Use the general mean formula of (1.3) to work through the details of

$$\mathrm{E}\,T^p = \int_0^\infty \Pr(T^p \geq u)\,\mathrm{d}u = \int_0^\infty \exp\{-(u^{1/p}/a)^b\}\,\mathrm{d}u = a^p\,\Gamma(1+p/b).$$

Show that this leads to mean $a\,\Gamma(1+1/b)$ and variance $a^2\,\{\Gamma(1+2/b)-\Gamma(1+1/b)\}^2$. With T from the Weibull (a,b), plot the function $\mathrm{sd}(T)/\mathrm{E}\,T$ as a function of b. (xx write out. also reparametrisation, with $1-\exp(-ct^b)$. point to story. xx)

(c). Show that $V=(T/a)^b\sim\mathrm{Expo}(1)$, and use this to give a recipe for simulating outcomes from any Weibull.

Ex. 1.60. *The Gompertz distribution.* The Gompertz distribution, with positiver parameters (a,b), has hazard rate $h(t)=a\exp(bt)$.

(a). Find the cumulative hazard rate, the c.d.f., and the density. Find also a formula for the median, and more generally for the q quantile, expressed via (a,b). (xx pointer to Story ii.1. more; round off. xx)

(b). Suppose an individual has survived up to time t_0. Show that her cumulative hazard rate, for the remaining lifetime, is $H(t)-H(t_0)=(a/b)\{\exp(bt)-\exp(bt_0)\}$. Give a formula for $t^*(t_0)$, her median survival time. (xx then brief application of this, for Norwegian women, using perhaps rough estimates of (a,b), via data from Human Mortality Index. give $(t_0,t^*(t_0))$ as a graph, for women born in perhaps 1900, 1960, 2020. can also be a Story. xx)

the Gumbel distribution

Ex. 1.61. *The Gumbel distribution.* Here we work with the Gumbel distribution, useful e.g. in models for extreme values.

(a). Let $X_1,\ldots,X_n$ be i.i.d. from the standard exponential distribution. Show that their maximum value M_n has c.d.f. $\{1-\exp(-m)\}^n$. Deduce that $M_n-\log n$ has c.d.f. $G_n(u)=\{1-(1/n)\exp(-u)\}^n$, for all $u\geq-\log n$.

(b). Show that the limit c.d.f. for $M_n-\log n$ becomes $G(u)=\exp\{-\exp(-u)\}$, and that this defines a c.d.f. on the full line. This is called the *Gumbel distribution*. Find its median and interquartile range.

(c). Find its density $g(u)=\exp\{-u-\exp(-u)\}$, and draw it in a diagram, along with the densities g_n for say $n=10,20,30$, for $M_n-\log n$.

(d). With U having the Gumbel distribution, show that its m.g.f. becomes $M(t)=\Gamma(1-t)$, for $t<1$. To find the moments, show by taking derivatives of $\Gamma(a)=\int_0^\infty v^{a-1}\exp(-v)\,\mathrm{d}v$ at position $a=1$ that

$$\mathrm{E}\,U^p=(-1)^p\Gamma^{(p)}(1)=(-1)^p\int_0^\infty(\log v)^p\exp(-v)\,\mathrm{d}v.$$

These may be found numerically. You may also work with the cumulant-generating function, where a certain connection can be made to the Riemann zeta function, namely

$$K(t)=\log\Gamma(1-t)=\gamma_e t+\sum_{j=2}^\infty\frac{\zeta(j)}{j}t^j$$

for $|t|<1$. Here $\gamma_e=0.5772...$ is the Euler constant, see Ex. 1.15, and $\zeta(j)=\sum_{n=1}^\infty 1/n^j$. In the terminology of Ex. 1.37, show that $\kappa_1=\gamma_e$, $\kappa_2=\zeta(2)$, $\kappa_3=2\,\zeta(3)$, $\kappa_4=6\,\zeta(4)$. With $\zeta(2)=\pi^2/6$, $\zeta(3)=\pi^3/25.7943$, $\zeta(4)=\pi^4/90$, establish that $\mathrm{E}\,U=\gamma_e$, $\mathrm{Var}\,U=\sigma^2=\pi^2/6$, $\mathrm{skew}(U)=2\,\zeta(3)/\sigma^3\doteq1.1395$, $\mathrm{kurt}(U)=12/5=2.4$.

(e). To appreciate the perhaps strange-looking connection from Gumbel moments to the zeta function, show from efforts of Ex. 1.15 that

$$M_n - \log n = \sum_{i=1}^{n} W_i/i + a_n$$

where $W_i = V_i - 1$, for i.i.d. unit exponentials V_i, and $a_n \to \gamma_e$. Find formulae for the mean, variance, skewness, kurtosis of $M_n - \log n$ based on this, and take their limits; these will agree with formulae found above.

Ex. 1.62. *The logistic distribution.* Consider the logistic distribution (in its standard form), with c.d.f. $H(x) = \exp(x)/\{1 + \exp(x)\}$, over the real line.

the logistic distribution

(a). Show that H indeed is a proper c.d.f., and that its density is $h(x) = \exp(x)/\{1 + \exp(x)\}^2 = H(x)\{1 - H(x)\}$, symmetric around zero. Find its interquartile range.

(b). With X having the logistic distribution, show that its m.g.f. $M(t)$ becomes

$$\int \frac{\exp\{(t+1)x\}}{\{1+\exp(x)\}^2}\,\mathrm{d}x = \int_0^\infty \frac{u^t}{(1+u)^2}\,\mathrm{d}u = \int_0^1 v^t(1-v)^{1-t}\,\mathrm{d}t = \Gamma(1+t)\Gamma(1-t),$$

for $|t| < 1$. Show from this that actually $X = U - V$, in terms of independent Gumbel distributed U and V, see Ex. 1.61. Demonstrate this fact also directly, via convolution. Use this representation to show that $\mathrm{Var}\, X = \pi^2/3$ and that its kurtosis is $\mathrm{kurt} = 6/5 = 1.2$.

(c). Here we managed to find moments of the logistic via the Gumbel difference representation. We may also use the intriguing formula $\Gamma(1+t)\Gamma(1-t) = \pi t/\sin(\pi t)$, a consequence of Euler's reflection formula. Take derivatives of this function to find the variance and kurtosis formulae.

(d). A scaled version of the logistic has c.d.f. $H(x, \tau) = H(x/\tau) = \exp(x/\tau)/\{1 + \exp(x/\tau)\}^2$, for a suitable positive τ. Show that the variance is $\tau^2\pi^2/3$, which is 1 for $\tau = \sqrt{3}/\pi$. Draw a figure with the density, alongside the standard normal, and comment.

Ex. 1.63. *The logistic-normal distribution on the unit interval.* The most prominent model for distributions on the unit interval is the Beta, also useful in various guises and roles outside its direct use for fitting datasets; see Ex. 1.19. Here we work with a different class, useful for versions of logistic regression, which we meet in Ch. 5. Let $H(u) = \exp(u)/\{1 + \exp(u)\}$ the logistic transform, worked with in Ex. 1.62, taking any real u to the unit interval.

(a). Solve $H(u) = v$ to find the inverse transform $H^{-1}(v) = \log\{v/(1-v)\}$. Start with X a standard normal and define $V = H(X)$. Show that its c.d.f. is $G(v) = \Pr(X \le H^{-1}(v)) = \Phi((H^{-1}(v)))$, with density

$$g(v) = \phi(H^{-1}(v))\frac{1}{v(1-v)} = \frac{1}{(2\pi)^{1/2}}\exp[-\tfrac{1}{2}\{\log v - \log(1-v)\}^2]\frac{1}{v(1-v)}.$$

(b). Generalise to the case where $X \sim \mathrm{N}(\xi, \sigma^2)$, where the density becomes

$$g(v, \xi, \sigma) = \phi((H^{-1}(v) - \xi)/\sigma)\frac{1}{\sigma}\frac{1}{v(1-v)}.$$

We call this the logistic-normal with parameters (ξ, σ). Draw c.d.f.s and densities of this type, for some combinations of parameters. Explain that if $X_1, \ldots, X_p$ has a joint multinormal distribution, then the random probability

$$p(X_1, \ldots, X_p) = H(\beta_0 + \beta_1 X_1 + \cdots + \beta_p X_p)$$

has the logistic-normal distribution, and identify its parameters. In this sense the logistic-normal class is closed under linear combinations of the underlying $H^{-1}(V)$, whereas the Beta class does not have similar properties.

(c). Suppose $Y \,|\, p$ is binom(n, p), and that p is logistic-normal. There are no closed-form expressions for the resulting probabilities

$$\bar{f}(y) = \binom{n}{y} \int_0^1 p^y (1-p)^{n-y} g(p, \xi, \sigma)\, \mathrm{d}p \quad \text{for } y = 0, 1, \ldots, n,$$

but it may be worked with numerically. Display a few of these $\bar{f}(y)$, for say $n = 100$ and a few values of (ξ, σ).

(d). We may also work the other way here: suppose $V = H(X)$ is Beta$(cp_0, c(1-p_0))$; what is then the distribution for X?

Ex. 1.64. *A normal with a normal mean is normal.* (xx preliminary version; need just a few editorial decisions regarding where to place it, and how. xx) The normal distribution has a convenient coherence type property: if X given its mean parameter is normal, and this mean parameter itself is normal, then X, marginally, is again normal. This is also related to what is found in Ex. 1.44.

(a). Consider first independent X_1 and X_2 with densities f_1 and f_2. Show that $\int f_1 f_2\, \mathrm{d}x$ is the density of $X_1 - X_2$, evaluated at zero. Write then $\phi_\sigma(x - \xi) = \sigma^{-1}\phi(\sigma^{-1}(x-a))$ for the N(ξ, σ^2) density. Show that

$$\int \phi_{\sigma_1}(x - \xi_1)\phi_{\sigma_2}(x - \xi_2)\, \mathrm{d}x = \phi_{(\sigma_1^2 + \sigma_2^2)^{1/2}}(\xi_1 - \xi_2).$$

(b). Assume that X given ξ has mean ξ and variance σ^2, and further that ξ stems from its own distribution, with mean ξ_0 and variance τ^2. Show that X, marginally, has mean ξ_0 and variance $\sigma^2 + \tau^2$. Then specialise to the normal case, with $X \,|\, \xi$ and ξ having normal distributions. Show that X indeed also is normal, (i) by integrating out the ξ, with respect to its distribution, and also (ii) by arguing via $X = \xi + \varepsilon = \xi_0 + \delta + \varepsilon$, where ε and δ are zero-mean normals with variances σ^2 and ε^2. (xx point to connected thing for logN. xx)

(c). Suppose $Y \,|\, (x_1, x_2)$ is normal N$(a + b_1 x_1 + b_2 x_2, \sigma^2)$, as in linear regression models we will study in later chapters; see e.g. Ex. 3.34. Assume then that (x_1, x_2) themselves have a distribution, in its space of covariate pairs, and that this distribution is binormal. Show that Y, marginally, is normal, and give formulae for its mean and variance.

Ex. 1.65. *Mixing the normal scale.* (xx at the moment nils thinks this exercise will go away, partly with material in Story v.4 and partly elsewhere. point back to and calibrate with Ex. 1.33. xx) Suppose $X \sim \mathrm{N}(\xi, \sigma^2)$ for given parameters (ξ, σ), but that there are background mechanisms producing these (ξ, σ). In various settings this leads to good 'mixtures of normals' models for actually observed data.

(a). Suppose a given individual has his ξ and that his associated X is a $\mathrm{N}(\xi, \sigma^2)$. Assume next that in a population of such X, there is a distribution $\xi \sim \mathrm{N}(\xi_0, \sigma^2_{\text{extra}})$ of their means. Show that an X sampled from that population is a $\mathrm{N}(\xi_0, \sigma^2 + \sigma^2_{\text{extra}})$. From a statistical modelling viewpoint we have simply 'put in more in the σ', perhaps stretched its interpretation a little, without inventing or having to invent a new model for the observed X, per se. Also, without knowing more, or perhaps having a separate experiment, we cannot identify the components of the observed variance.

(b). Now turn attention to the scale. With the mean ξ kept fixed, but σ having some density $\pi(\sigma)$, show that X has density $\bar f(x) = \int (1/\sigma)\phi((x-\xi)/\sigma)\pi(\sigma)\,\mathrm{d}\sigma$, and the variance of X is the mean of the distribution of σ^2. Assume for simplicity of presentation that $\xi = 0$, and work with the case where the distribution of σ is such that $1/\sigma^2 \sim \mathrm{Gam}(a, b)$ (it is common to express this by saying that σ^2 has an inverse gamma distribution). Work out that

$$\bar f(x) = \frac{1}{(2\pi)^{1/2}} \frac{b^a}{\Gamma(a)} \frac{\Gamma(a+\frac12)}{(b+\frac12 x^2)^{a+1/2}}.$$

It is useful to transform this to a member of the well-known distributions, to facilitate computations of probabilities etc. Show therefore that

$$V = (\tfrac12 X^2/b)/(1+\tfrac12 X^2/b) \sim \mathrm{Beta}(\tfrac12, a),$$

and express the c.d.f. of X in terms of the c.d.f. of this Beta distribution: $\Pr(|X| \le x) = \mathrm{Be}((\frac12 x^2/b)/(1+\frac12 x^2/b), \frac12, a)$. Check these details via simulations.

Ex. 1.66. *Normal mixtures.* [xx to come. mean, variance, skewness. xx] Suppose Y is such that with probability p_j, it is a normal (ξ_j, σ_j^2), with probabilities $p_1, \ldots, p_k$ summing to 1. Its density may be written $f(y) = \sum_{j=1}^k p_j \phi_{\sigma_j}(y - \xi_j)$, where $\phi_\sigma(u) = \sigma^{-1}\phi(\sigma^{-1}u)$ is the density of a $\mathrm{N}(0, \sigma^2)$. Such distributions are called normal mixtures.

(a). With J taking values $1, \ldots, k$, with probabilities $p_1, \ldots, p_k$, let $Y \,|\, (J = j) \sim \mathrm{N}(\xi_j, \sigma_j^2)$. Show that this Y has the density above; this amounts to a way of representing and interpreting a normal mixture.

(b). From $\mathrm{E}\,(Y \,|\, J) = \xi_J$ and $\mathrm{Var}\,(Y \,|\, J) = \sigma_J^2$, show that

$$\mathrm{E}\,Y = \bar\xi = \sum_{j=1}^k p_j \xi_j, \quad \mathrm{Var}\,Y = \mathrm{E}\,(Y-\bar\mu)^2 = \sum_{j=1}^k p_j \sigma_j^2 + \sum_{j=1}^k p_j (\xi_j - \bar\xi)^2.$$

(c). (xx someting; display a few. xx)

Ex. 1.67. *The hypergeometric distribution.* You draw a sample of n items from a bag of N, which has A of Type One and $B = N - A$ of type Two. Consider X, the number among the sampled n which are of Type One.

(a). Show that X has distribution

$$f(x) = \Pr(X = x) = \binom{A}{x}\binom{B}{n-x}\Big/\binom{N}{n}.$$

For which x is this positive? Explain the identity $\sum_{x=0}^A \binom{A}{x}\binom{B}{n-x} = \binom{A+B}{n}$.

(b). Show that $\mathrm{E}\, X = nA/N = np$, with $p = A/N$ the proportion of Type One in the bag. This may be done work with $\sum_{x=0}^{n} x f(x)$, or by writing $X = J_1 + \cdots + J_n$, with J_i and indicator for selected item i being a Type One or not.

(c). Explain that if one samples one item at the time, followed by replacing the item, then the $J_1, \ldots, J_n$ above are independent Bernoulli variables with probability $p = A/N$, leading in that case to binomial variance $np(1-p)$. For the present hypergeometric setting, where n items sampled in one go without replacement, the J_i are dependent; show that $\mathrm{cov}(J_1, J_2) = p(A-1)/(N-1) - p^2 = -p(1-p)/(N-1)$. Deduce that the variance formula becomes $\mathrm{Var}\, X = c_n np(1-p)$, with c_n being the shrinking factor $(N-n)/(N-1)$. This may be accomplished working algebraically with $\mathrm{E}\, X(X-1)$, or via the representation above.

(d). (xx just a bit more. comparing with binomial. xx)

a copula, coupling the binormal with other distributions

Ex. 1.68. *Building bivariate dependence models.* (xx something here. need to have something beside the multinormal for dependence. xx) Let X and Y have densities $f_1(x)$ and $f_2(y)$, with c.d.f.s $F_1(x)$ and $F_2(y)$. To model dependence between them, the idea pursued here is mapping them to the normal scale, then using the binormal model, and mapping back.

(a). Write $X = F_1^{-1}(\Phi(U))$ and $Y = F_2^{-1}(\Phi(V))$, where (U, V) is binormal with zero means, unit variances, and correlation ρ. Show that X and Y indeed have densities f_1 and f_2. Writing $g_\rho(u, v)$ for the binormal density, show that the joint density can be written

$$f(x,y) = f_1(x) f_2(y) R_\rho(\Phi^{-1}(F_1(x)), \Phi^{-1}(\Phi(y))), \quad \text{with } R_\rho(u,v) = \frac{g_\rho(u,v)}{\phi(u)\phi(v)}.$$

with the dependence factor R_ρ of course being 1 when $\rho = 0$. Show in fact that

$$R_\rho(u,v) = \frac{1}{(1-\rho^2)^{1/2}} \exp\Big\{-\tfrac{1}{2}\frac{1}{1-\rho^2}(\rho^2 u^2 + \rho^2 v^2 - 2\rho uv)\Big\}.$$

(b). Now consider $f_1(x) = \exp(-x)$ and $f_2(y) = \exp(-y)$, i.e. two unit exponentials. Construct a bivariare pair (X, Y) via the recipe above, with a density $f_\rho(x, y)$ having unit exponential marginals. Compute the correlations $\mathrm{corr}(X, Y)$ as a function of ρ, which might be easiest via simulations. (xx simulate (X, Y) for some values of ρ. note very different cases, $\rho = -0.99$ and $\rho = 0.99$. the range of correlations is $(-0.64, 1)$. the density formula is this. xx)

$$f(x,y) = \exp(-x-y) R_\rho\big(\Phi^{-1}(1-\exp(-x)), \Phi^{-1}(1-\exp(-y))\big).$$

(c). (xx one more example. for the uniform case, we have $\mathrm{corr}(X, Y)$ quite close to ρ, but they remain different. xx)

Ex. 1.69. *Alice and Bob correlate their binomials.* Here we show how correlated binomials may be constructed, leading also to correlated random walks.

(a). Alice flips her fair coin n times, with i.i.d. 0-1 outcomes $A_1, \ldots, A_n$. Bob has two coins, and mixes between them depending on Alice's outcomes: if $A_i = 1$, he uses the plus-coin with probability $\frac{1}{2} + a$ for heads; and if $A_i = 0$ he uses his minus-coin with $\frac{1}{2} - a$ for heads. With B_i his outcome, show that $\Pr(B_i = 1) = \frac{1}{2}$, and argue therefore that

both X_n and Y_n are binomial $(n, \frac{1}{2})$ variables, where $X_n = \sum_{i=1}^n A_i$ and $Y_n = \sum_{i=1}^n B_i$ are the number of heads for Alice and for Bob. Show that the correlation between these two binomials is $2a$.

(b). In the little story-telling above, Bob observes Alice's outcomes, one by one, which then influence his choice between two coins; Alice doesn't even need to be aware of Bob's existence. Explain however that we from observed pairs of coin flips (A_i, B_i) never can see the difference between that scenario and the alternative one, that Bob is the one flipping his fair coin, without caring for Alice, before she chooses between two biased ones. This is arguably an instance of what Breiman (2001) alludes to as the Rashomon Effect (from a Japanese movie in which different persons report very differently about something they have all observed): data alone cannot help us uncover which of the chains of action have been at work. See Story i.9 for further discussion. Show indeed that as long as Alice and Bob have a joint scheme of producing outcomes $(0,0), (0,1), (1,0), (1,1)$, with probabilities respectively $\frac{1}{4}(1+a), \frac{1}{4}(1-a), \frac{1}{4}(1-a), \frac{1}{4}(1+a)$, then (X_n, Y_n) have the correlated binomial distribution.

the Rashomon Effect: different models may offer equally good explanations

(c). Find a way to compute $f(x,y) = \Pr(X_n = x, Y_n = y)$, for $x, y = 0, 1, \ldots, n$.

(d). Leaving the Rashomon aspects to the side, generalise the first setup to the case of two correlated binomials (n, p), where p is not necessarily $\frac{1}{2}$. Take indeed $\Pr(A_i = 1) = p$ and then $\Pr(B_i = 1 \,|\, A_i = 1) = p + a$, $\Pr(B_i = 0 \,|\, A_i = 0) = 1 - p + ap/(1-p)$, for $a < \min(p, 1-p)$, and show that this works properly. What is the correlation between X_n and Y_n?

(e). Show that $\sqrt{n}(X_n/n - p, Y_n/n - p)$ tends in distribution to a binormal zero-mean (X, Y), with variances $p(1-p)$ and covariance ap.

(f). (xx brief pointer to two correlated random walks, Ch9, with two correlated Brownian motions. also good ML exercise, finding $\widehat{a}$ based on having observed Alice and Bob random walks, easy $\sqrt{n}(\widehat{a} - a)$, test for $a = 0$, etc.)

More material, extras, leftovers (as of 12-March-2025)

Ex. 1.70. *The purely noncentral excentric chi-squared.* The noncentral chi-squared, with its somewhat complicated density expressed as an infinite sum, etc., admits a simple representation, as a pure chi-squared plus a pure noncentral part, as shown in Hjort (1988a).

(a). Consider a variable K from the noncentral chi-squared distribution, say $K \sim \chi^2_m(\lambda)$. We found its m.g.f. in Ex. 1.52. Deduce that *if* there is a representation, of the form $K = K_0 + U$, with $K_0 \sim \chi^2_m$ and a certain purely eccentric variable U, then we must have $M(t) = \mathrm{E}\exp(tU) = \exp\{\lambda t/(1-2t)\}$. In particular, such a U will then have a distribution depending only on λ, the same for each m.

(b). Let us use the occasion to demonstrate the moment-generating aspect of moment-generating functions. Still under the assumption that there actually is a purely eccentric U, as above, find the four first moments of U, using $\mathrm{E}\,U^j = M^{(j)}(0)$, the derivatives of M at zero. Show indeed that $\mathrm{E}\,U = \lambda$, $\mathrm{Var}\,U = 4\lambda$, and that its skewness is $3/\lambda^{1/2}$.

(c). Indeed, there is such a U. You may check Hjort (1988a) to find constructive arguments to *deduce* the distribution of U, as opposed to taking the answer and then verifying

that it is correct. Here we are content with the verification: show that with U having the distribution

$$\Pr(U \le u) = \sum_{j=1}^{\infty} p(j)\Gamma_{2j}(u), \quad \text{with} \quad p(j) = \exp(-\tfrac{1}{2}\lambda)(\tfrac{1}{2}\lambda)^j/j!,$$

where $\Gamma_{2j}(u) = \Pr(\chi^2_{2j} \le u)$ is the cumulative for the χ^2_{2j}, its m.g.f. indeed becomes the desired $\exp\{\lambda t/(1-2t)\}$. We allow also Γ_0 here, for the chi-squared with zero degrees of freedom, for the variable which is equal to zero with probability one. This is a Poisson mixture of χ^2_{2j} distributions. In particular, U has a positive pointmass at zero, $\Pr(U=0) = \exp(-\frac{1}{2}\lambda)$.

(d). Use the representation $U \,|\, (J=j) \sim \chi^2_{2j}$, with $J \sim \text{Pois}(\frac{1}{2}\lambda)$, to find the mean, the variance, the skewness, the kurtosis of U. Use these formulae again to find the mean, variance, skewness, kurtosis for $K = K_0 + U$, the noncentral $\chi^2_m(\lambda)$.

(e). What is required, of (m,λ), for $K \sim \chi^2_m(\lambda)$ to have a normal limit distribution (when properly normalised)?

(f) (xx just a bit more. for small λ, an approximation is $(1-\frac{1}{2}\lambda)\chi^2_m + \frac{1}{2}\lambda\chi^2_{2m+2}$. note its use for power studies of certain tests. there is a full purely noncentral process, say $\{U(\lambda)\colon \lambda \ge 0\}$, perhaps in chapter 8. xx).

Ex. 1.71. *Uniforms and exponentials.* There are various interconnections between uniforms and exponentials.

(a). With $U \sim \text{unif}(0,1)$, show that $V = -\log U \sim \text{Expo}(1)$. Deduce that with $U_1, \ldots, U_n$ i.i.d. from the uniform, we have $W_n = -\log(U_1 \cdots U_n) \sim \text{Gam}(n,1)$.

(b). Start with $G_2 \sim \text{Gam}(2,1)$, the sum of two unit exponentials. Show that with $B \sim \text{unif}(0,1)$, the product BG_2 is a G_1, a unit exponential. Explain that this also shows that with three uniforms, we have the curious result that $(U_1U_2)^{U_3}$ is also a uniform.

(c). Write G_n for a $\text{Gam}(n,1)$. Explain that its distribution is determined from the moment function $A_n(s) = \text{E}\, G_n^s = \Gamma(n+s)/\Gamma(n)$ for $s \ge 0$. To determine the independent random factor B needed to make BG_3 into a G_1, explain that we require

$$\text{E}\,(BG_3)^s = \text{E}\, B^s\, \Gamma(3+s)/\Gamma(3) = \Gamma(1+s).$$

Then show that this holds when we choose $B \sim \text{Beta}(1,2)$. Explain that with independent uniforms U_1, U_2, U_3 and this B, we have $(U_1U_2U_3)^B \sim \text{unif}$.

(d). Then generalise, showing first that with $B \sim \text{Beta}(1, n-1)$ independent of a $G_n \sim \text{Gam}(n,1)$, we have $BG_n \sim G_1$, and then that with i.i.d. uniforms, we have $(U_1 \cdots U_n)^B \sim \text{unif}$.

Ex. 1.72. *The noncentral F distribution.* A ratio of independent chi-squared distributions, modulo a multiplication factor, has the F distribution, as seen in Ex. 1.54. There are certain statistical uses of the extended construction (xx give pointers; power; cc(λ); more) when the nominator is a noncentral chi-squared. Formally, if $X \sim \chi^2_{m_1}(\lambda)$ and $Y \sim \chi^2_{m_2}$ are independent, then

$$F = \frac{X/m_1}{Y/m_2} = \frac{\chi^2_{m_1}(\lambda)/m_1}{\chi^2_{m_2}/m_2} \sim F(m_1, m_2, \lambda),$$

termed a noncentral F distribution with degrees of freedom (m_1, m_2) and excentre parameter λ.

(a). Find the mean and variance of such an F.

(b). Suppose $Y_1, \ldots, Y_n$ are i.i.d. $\mathrm{N}(\xi, \sigma^2)$, and consider $F = n\bar{Y}^2/\widehat{\sigma}^2$. Show that $F \sim F(1, n-1, n\xi^2/\sigma^2)$.

Ex. 1.73. *The Tweedie distribution.* As a special case of the construction of Ex. 1.40, study $Z = \sum_{i=1}^{N} X_i$, where $N \sim \mathrm{Pois}(\lambda)$ and the X_i are i.i.d. from a $\mathrm{Gam}(a, b)$.

(a). Show that the c.d.f. can be expressed as

$$H(z) = \sum_{n=0}^{\infty} p(n, \lambda) G(z, na, b) = p(0, \lambda) + \sum_{n=1}^{\infty} p(n, \lambda) G(z, na, b),$$

with $p(n, \lambda)$ the Poisson and $G(z, na, b)$ the c.d.f. of the $\mathrm{Gam}(na, b)$. In particular, show that $\Pr(Z = 0) = \exp(-\lambda)$, and that $\mathrm{E}\, Z = \lambda a/b$, $\mathrm{Var}\, Z = \lambda\{(a/b)^2 + a/b^2\}$.

(b). Simulate say $n = 500$ outcomes for this distribution, for parameter values (λ, a, b) you decide on. Check if you can estimate these values based on the simulated sample.

(c).

Ex. 1.74. *The Beta Prime distribution.* (xx keep it brief. essentially just $Y = X/(1-X)$ with $X \sim \mathrm{Beta}(a, b)$. xx)

Ex. 1.75. *The Dagum distribution.* [xx to be done. three para. formulae for quantiles etc. xx]

Ex. 1.76. *Yet another two-parameter survival data model.* (xx a short rant. xx) The gamma function $\Gamma(t)$ has its minimum at position $t_0 = 1.4616$, and then increases, rather rapidly, from that point onwards. Define a distribution via its survival function $S(t) = \{\Gamma(at_0)/\Gamma(a(t_0 + t))\}^b$ for $t \geq 0$. Show that its cumulative hazard rate becomes $H(t) = b\{\log\Gamma(a(t_0+t)) - \log\Gamma(at_0)\}$, with hazard rate $h(t) = ba\psi(a(t_0+t))$. plot some curves.

Ex. 1.77. *Leftovers.* [xx we push smaller things here, when they belong better in later chapters. in particular, the present chapter should be free of estimators and confidence and inference, and also of lilmits. xx]

(a). If Y_n and Y have m.g.f. s M_n and M, then $M_n(t) \to M(t)$ for all t in a neighbourhood around zero is sufficient for $Y_n \to_d Y$. [xx repair this, point to Ch. 2. since we need to point to what convergence in distribution is. xx]

(b). In particular, if the m.g.f. $M_n(t)$ of some Y_n tends to $\exp(\frac{1}{2}t^2)$ for all t close to zero, then $Y_n \to_d \mathrm{N}(0, 1)$. Illustrate this principle for the following situation: Let $Y_1, Y_2, \ldots$ be i.i.d. from the simple symmetric two-point distribution, where Y_i takes on values $-1, 1$ with probabilities $\frac{1}{2}, \frac{1}{2}$. Show that the normalised mean $Z_n = \sqrt{n}\bar{Y}_n = n^{-1/2}\sum_{i=1}^{n} Y_i$ has m.g.f. $\cosh(t/\sqrt{n})^n$, and show that it converges to $\exp(\frac{1}{2}t^2)$. You have now proved the Central Limit Theorem for this special case.

Ex. 1.78. *Moment-generating functions for two or more variables.* (xx needs rounding off and pointers to other matters. xx) The one-dimensional $M(t)$ for a variable Y dealt with in Ex. 1.31, A.36 generalises neatly to the case of several variables. For a vector $Y = (Y_1, \ldots, Y_p)^{\mathrm{t}}$ the joint m.g.f. is

$$M(t) = M(t_1, \ldots, t_p) = \mathrm{E}\, \exp(t^{\mathrm{t}} Y) = \mathrm{E}\, \exp(t_1 Y_1 + \cdots + t_p Y_p),$$

involving of course the joint distribution of the random vector.

(a). Assume $Y_1, \ldots, Y_p$ are actually independent, with m.g.f.s $M_1, \ldots, M_p$. Show that $M(t_1, \ldots, t_p) = M_1(t_1) \cdots M_p(t_p)$.

(b). Assume $Y \sim \mathrm{N}_p(\xi, \Sigma)$. Show that $M(t) = \exp(\xi^{\mathrm{t}} t + \frac{1}{2} t^{\mathrm{t}} \Sigma t)$.

(c). Let Y be standard normal. Show that

$$M(s,t) = \mathrm{E}\, \exp(sY + tY^2) = \exp\{\tfrac{1}{2} s^2/(1-2t)\}/(1-2t)^{1/2}$$

for $t < \frac{1}{2}$. This characterises the joint distribution of (Y, Y^2). Comment on the cases $s = 0$ and $t = 0$.

(d). (xx do the trinomial too. xx)

Ex. 1.79. *Product of two normals.* (xx we will see later how this pans out; connect to a cc(ϕ) in Ch. 7. xx) For independent $X \sim \mathrm{N}(a, 1)$ and $Y \sim \mathrm{N}(b, 1)$, consider their product XY. Can its distribution be close to a normal?

(a). We start with the cleanest case, letting $Z = UV$ with U and V independent standard normals. Show that its m.g.f. is

$$M(t) = \mathrm{E}\, \exp(tUV) = \mathrm{E}\, \exp(\tfrac{1}{2} t^2 Y^2) = \frac{1}{(1-t^2)^{1/2}} = \frac{1}{(1-t)^{1/2}(1+t)^{1/2}},$$

for $|t| < 1$. Deduce from this that Z has the same distribution as $\frac{1}{2}(K - L)$, where K and L are independent with χ^2_1 distributions.

(b). For the non-zero-mean case, write $X = a + U$ and $Y = b + V$, show that $Z = XY - ab = aV + bU + UV$, and that XY has variance $\sigma^2 = 1 + a^2 + b^2$. Work also out its skewness and kurtosis; in particular, show that $\mathrm{E}\, Z^3 = 6ab$. (xx might also get insight from $(X+Y)^2 - (X-Y)^2 = 4XY$, leading to

$$Z = (1/4)\{\mathrm{N}(a+b, 2)^2 - \mathrm{N}(a-b, 2)^2\} = \tfrac{1}{2}\{\chi^2_1(\tfrac{1}{2}(a+b)^2) - \chi^2_1(\tfrac{1}{2}(a-b)^2)\}.$$

need to write up these things with Dennis-Nils for $\phi = ab$, needed for comparing two Fillers. xx)

(c). Establish first that $\mathrm{E}\,\{\exp(tZ) \,|\, (U = u)\} = \exp(tbu) \exp\{\frac{1}{2} t^2 (a+u)^2\}$, and use results from Ex. 1.78 to show that Z has m.g.f.

$$M(t) = \frac{1}{(1-t^2)^{1/2}} \exp\{\tfrac{1}{2} a^2 t^2 + \tfrac{1}{2} t^2 (b+at)^2/(1-t^2)^{1/2}\} \quad \text{for } |t| < 1.$$

(d). In order to check if Z/σ might have a distribution not far from the standard normal, plot the function $\log M(t)$, for a few choices of (a, b), and inspects its closeness to $\frac{1}{2} t^2$. (xx round this off. xx)

(e). (xx might drop this, or give a pointer to something in Ch7, re difference of Fieller ratios. xx) Consider the parameters $x_{0,1} = -a_1/b_1$ and $x_{0,2} = -a_2/b_2$, the positions at which two lines $a_1 + b_1 x$ and $a_2 + b_2 x = 0$. It's tricky to test $x_{0,1} = x_{0,2}$, but that hypothesis is equivalent to $a_1 b_2 = a_2 b_1$. May hence construct $Z = \widehat{a}_1 \widehat{b}_2 - \widehat{a}_2 \widehat{b}_1 = Z_1 - Z_2$, say. I find something like

$$\begin{aligned} \mathrm{E}\,\exp(tZ) \;=\; & \frac{1}{(1-t^2)^{1/2}} \exp\{\tfrac{1}{2}a_1^2 t^2 + \tfrac{1}{2}t^2(b_2 + a_1 t)^2/(1-t^2)^{1/2}\} \\ & \frac{1}{(1-t^2)^{1/2}} \exp\{\tfrac{1}{2}a_2^2 t^2 + \tfrac{1}{2}t^2(b_1 - a_2 t)^2/(1-t^2)^{1/2}\}. \end{aligned}$$

the point is that the skewness has disappeared, and that $(Z_1 - Z_2)/\tau$ becomes closer to normal, with $\tau^2 = 2 + a_1^2 + b_1^2 + a_2^2 + b_2^2$. still a hard thing to do very well.

Ex. 1.80. *Nonparametric models.* Above we have seen a wide range of parametric models, each model indexed via a finite and perhaps small number of parameters. Models can also be nonparametric, however, with fewer modelling assumptions placed on their outcomes. (xx just a bit. data with unknown mean and unknown variance, but saying nothing more. symmetry. unimodal. shape constraints. xx)

Ex. 1.81. *Extra: m.g.f. only for $t \le 0$.* (xx to be made into an exercise inside terrain of Laplace transforms. as mgf, we have $M(t)$ finite only for $t \le 0$, but moments exist. xx) Consider the variable $X = Y^2$, where Y is a unit exponential.

(a). Show that X has density $\exp(-x^{1/2})/(2x^{1/2})$, and that that its moment sequence is $\mathrm{E}\,X^k = (2k)!$ for $k = 1, 2, \ldots$. Show however that its m.g.f. $M(t)$ is finite only for $t \le 0$. Hence its derivatives at zero do not exist properly either, only derivatives from the left.

(b). To find the $M(t)$ it's more convenient to work with its Laplace transform $L(s) = \mathrm{E}\,\exp(-sX)$, for $s \ge 0$; see Ex. 1.38. Show that

$$L(s) = \int_0^\infty \exp(-sy^2 - y)\,\mathrm{d}y = \exp(1/(4s))(2\pi)^{1/2}\{1 - \Phi(1/(2s)^{1/2})\}/(2s)^{1/2}$$

for $s > 0$, with $L(0) = 1$. (xx So the $1 + \sum_{k=1}^\infty (-1)^k s^k\,\mathrm{E}\,X^k/k!$ formula leads to $1 + \sum_{k=1}^\infty (-1)^k\{(2k)!/k!\}\,s^k$?, but this looks strange, with too big terms. perhaps this is merely an instance of Taylor power series not working, now that we only have right derivatives, not full derivatives. xx)

(c). (xx could bring in the Feller example, along the lines of $f_a(x) = c\exp(-x^{1/4})(1 + a\psi(x))$, with the $\psi(x)$ constructed so cleverly that $\int x^k \psi(x)\,\mathrm{d}x = 0$ for all k. xx)

Ex. 1.82. *Extra: Cumulants and the CLT.* (xx nils rant; might land in Ch2. xx) Consider $X_1, X_2, \ldots$ being i.i.d. with mean zero, variance one, and m.g.f. $M(t) = \exp\{K(t)\}$ finite in a neighbourhood around zero, with

$$K(t) = \tfrac{1}{2}t^2 + \sum_{j \ge 3} (\kappa_j/j!)\,t^j \quad \text{for } |t| \le \delta.$$

We call K the cumulant-generating function for the X_i distribution. With $Z_m = (X_1 + \cdots + X_m)/m^{1/2}$, show that its m.g.f. becomes

$$M_m(t) = M(t/m^{1/2})^m = \exp\{mK(t/m^{1/2})\}.$$

Its cumulant-generating function is hence

$$K_m(t) = \tfrac{1}{2}t^2 + \sum_{j\geq 3} \frac{\gamma_j}{j!\, m^{j/2}} t^j.$$

Hence all coefficients are smaller than for K itself, and $K_m(t) \to \frac{1}{2}t^2$. (xx this is more or less the CLT. xx)

Notes and pointers

(xx to come. a bit old literature, but crisply, and not systematic. brief genesis of the normal, a few sentences on the chi-squared, Pearson (1900), the t, Student (1908), the F, the Dirichlet, more. we also point to essential things in later chapters. point out that the normal is famous and useful also because of a host of approximation methods. xx)

(xx mention things for Irwin (1927) and Hall (1927) and before. xx)

(xx where do we have a precise theorem on $M_X = M_Y$ implying $F = G$? inversion formula? Jensen is Johan Jensen. Vector form too?)

(xx find the little fun fact from Hjort (1994), regarding t and midway between cauchy and normal. xx)

The chi-squared is on the list over deservedly famous distributions in probability theory and statistics, and stems from Karl Pearson's famous 1900 paper, 'On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling'. [xx a bit more: he establishes the chi-square distribution, the test carrying the chi-squared name, and sets up a rigorous concepttual framework for hypothesis testing. xx]

(xx point to uses of expofamily in both Chs. 4 and 5. xx)

I.2

Large-sample theory

The broad themes of this chapter are the concepts, details, methods, results, applications pertaining to three modes of convergence for random variables: convergence in probability, convergence almost surely, convergence in distribution. The first two have to do with random variables X_n coming close to some limit X, with increasing n, typically indexed by sample size; often the limit is merely a constant. The chief result here, with various extensions and uses, is the Law of Large Numbers, that the empirical average of a sequence of observations tends to the expected value of the underlying distribution. The third mode of convergence rather involves the distribution of X_n coming close to the distribution of some limit X, with the Central Limit Theorem being a prime statement. These machineries also lead to practically useful approximations; the idea is that a complicated distribution may be approximated by something much simpler. This is in particular helped by a collection of approximation methods called the delta method. The theory is developed first for functions of i.i.d. sequences, involving tools of moment-generating functions and characteristic functions, along with various probability inequalities. It is then extended to cover cases of independent variables from non-equal distributions, culminating in the Lindeberg theorem, giving precise conditions under which a sum of independent components approaches normality. Methods and results from this chapter are crucial for developing the likelihood theory of Ch. 5, and also for several later chapters.

Key words: central limit theorems, characteristic functions, delta method, large-sample approximations, laws of large numbers, Lindeberg conditions, modes of convergence

In this chapter we study convergence of sequences $(X_n)_{n\geq 1}$ of random elements. In statistics, the index n typically refers to the sample size, and X_n is some function of the n data points available, for example an estimator. The modes of convergence we study take place as n grows without bounds; hence the name large-sample theory.

A random element X is a function defined on a probability space $(\Omega, \mathcal{F}, \Pr)$, and taking values in some space $\mathcal{X}$. Throughout this chapter, and indeed throughout this book, $\mathcal{X}$ is a metric space equipped with the Borel-σ-algebra, i.e., the smallest σ-algebra containing all the open sets. When $\mathcal{X}$ is a subset of the real line, then X is a random variable; when $\mathcal{X}$ is a subset of $\mathbb{R}^k$ for some $k \geq 1$, X is a random vector, so a random vector of dimension one is a random variable; and when $\mathcal{X}$ is a set if functions, X is called a random or stochastic process. In this chapter we concentrate mainly on convergence of random variables and vectors, with the more involved themes of convergence of stochastic

processes studied in Ch. 9. Several of the results of the present chapter are valid for stochastic processes, however; which ones will be clear from the context.

Applications of large-sample theory are plentiful in probability and statistics, partly to understand crucial phenomena better, and partly to provide fruitful and practical approximations; an estimator or a statistic might have a very complicated exact distribution, but have a simple to use and sometimes accurate large-sample approximation. The key convergence concepts, with ensuing approximations and applications, are as follows.

First, if $X_n = (X_{n,1}, \ldots, X_{n,k})$ is a sequence of random vectors, and $a = (a_1, \ldots, a_k)$ is some point, we say that X_n *converges to* a *in probability*, written $X_n \to_{\rm pr} a$, if for each positive ε,

$$\Pr(\|X_n - a\| \geq \varepsilon) \to 0, \quad \text{as } n \text{ tends to infinity.} \tag{2.1}$$

Here $\|x - a\| = (\sum_{j=1}^k (x_j - a_j)^2)^{1/2}$ is the Euclidean distance, and hence ordinary distance in the one-dimensional case. If $\widehat{\theta}_n$ is an estimator of the value θ_0, we say that the estimator is *consistent* if $\widehat{\theta}_n \to_{\rm pr} \theta_0$. Various results from the present chapter will indeed have immediate applications in Ch. 3. consistency of an estimator

Second, a stronger version of convergence is X_n converges to a *almost surely*, that is, the convergence occurs with probability one. This means that

$$\Pr(\lim_n \|X_n - a\| = 0) = 1, \tag{2.2}$$

and we write $X_n \to_{\rm a.s.} a$. We say that an estimator $\widehat{\theta}_n$ is strongly consistent for θ_0 if $\widehat{\theta}_n \to_{\rm a.s.} \theta_0$. A strong achievement indeed is the *strong Law of Large Numbers* (LLN), which says that the LLN

$$\text{if } X_1, X_2, \ldots \text{ are i.i.d. with finite mean } \xi, \text{ then } \bar{X}_n = n^{-1} \sum_{i=1}^n X_i \to_{\rm a.s.} \xi, \tag{2.3}$$

with no further assumptions required. One may readily prove weaker versions of the LLN, as in Ex. 2.10(a) and Ex. 2.15, but with the development of sharper tools, and separate valuable results along the way, we reach the strong LLN in Ex. 2.58–2.59. It immediately has many applications and uses, as we shall see.

We note that limits in probability and almost surely can easily be defined also when the limit is a random vector and not merely a constant, just by replacing a with a random vector X in (2.1)–(2.2). Most often, though, these two convergence concepts are used for cases when the limit is a constant.

The third and statistically speaking most important concept is that of *convergence in distribution*. If the random vectors X_n and X of dimension $k \geq 1$ have c.d.f.s F_n and F, we say that X_n converges in distribution to X if $X_n \to_d X$ $F_n \Rightarrow F$

$$F_n(x) \to F(x), \quad \text{for all continuity points } x = (x_1, \ldots, x_k) \text{ of } F. \tag{2.4}$$

We write $X_n \to_d X$ or $F_n \Rightarrow F$ to indicate convergence in distribution, allowing for simplicity also statements like $X_n \to_d {\rm N}(0,1)$. This form of convergence is also called *weak convergence* and *convergence in law*. The bigger sibling of the LLN is the Central Limit Theorem (CLT): if $X_1, X_2, \ldots$ are i.i.d. random variables with variance σ^2, then the CLT

$$\sqrt{n}(\bar{X}_n - {\rm E}\, X_1) \to_d {\rm N}(0, \sigma^2), \tag{2.5}$$

again with no further assumptions needed beyond a finite variance. Below we go considerably further, however, in detail, in extensions, in applications. In particular, with additional tools and efforts we reach the Lindeberg theorem, with precise necessary conditions for a sum of independent variables from different distributions to approach normality. Such results are, for example, used to establish approximate normality for estimators in regression models.

Reaching the CLTs, with variations, requires hard mathematical work, with various technical details to sort out. When the main theorems have been established, however, along with further tools, their actual use for statistical applications might be relatively straightforward. In particular, functions of approximately normal variables are also approximately normal, as we learn in the subsection on delta methods.

Modes of convergence

Ex. 2.1. *Implications.* Of the three modes of convergence, convergence almost surely is the strongest, convergence in distribution the weakest, with convergence in probability lying somewhere in the middle. In the following, X_n and X are random vectors. In statistical applications, $(X_n)_{n\geq 1}$ will often be a sequence of estimators, and the limit a constant, e.g., the parameter of a model.

(a). Show that $X_n \to_{\text{a.s.}} X$ if and only if the event $\cap_{n\geq 1} \cup_{k\geq n} \{\|X_k - X\| \geq \varepsilon\}$ has probability zero for all $\varepsilon > 0$.

(b). Show that if $X_n \to_{\text{a.s.}} X$, then $X_n \to_{\text{pr}} X$.

(c). Show that if $X_n \to_{\text{pr}} X$ then $X_n \to_d X$.

(d). Let Y be a nonnegative random variable. Show that $\Pr(Y \geq a) \leq \mathrm{E}\,(Y)/a$ for any $a > 0$, which is Markov's inequality. Deduce that if $\mathrm{E}\,\|X_n - X\| \to 0$, then $X_n \to_{\text{pr}} X$.

Markov's inequality

(e). Show that $\|X_n - X\| \to_{\text{a.s.}} 0$ is equivalent to $|X_{n,j} - X_j| \to_{\text{a.s.}} 0$ for each j, and that the same is true for convergence in probability.

Ex. 2.2. *The Borel–Cantelli lemma and convergence almost surely.* Let $A_1, A_2, \ldots$ be a sequence of events. Consider $A_{\text{i.o.}} = \cap_{n\geq 1} \cup_{m\geq n} A_m$, the full-sequence event corresponding to the A_n occurring infinitely often.

(a). Let $A_1, A_2, \ldots$ be a sequence of events, and let N be the total number of occurrences of the A_i. Show that $\mathrm{E}\,N = \sum_{i=1}^{\infty} \Pr(A_i)$.

(b). Show that if $\sum_{n=1}^{\infty} \Pr(A_n)$ is convergent, then $\Pr(A_{\text{i.o.}}) = 0$. So sooner or later there will be a finite (but random) n, such that none of the A_m will ever occur, for $m > n$.

The Borel–Cantelli lemma

(c). Assume in addition that the $A_1, A_2, \ldots$ are independent events. Show that if $\sum_{n=1}^{\infty} \Pr(A_n)$ is divergent, then $\Pr(A_{\text{i.o.}}) = 1$. In particular, for the case of independent events, there can't be say a 50 percent chance that there will be infinitely many occurrences. To show this, you may prove and use the inequality $1 + x \leq \exp(x)$, valid for all $x \in \mathbb{R}$.

(d). Let $X_1, X_2, \ldots$ be i.i.d. from the unit exponential distribution. Will there be infinitely many cases with $X_i \geq 0.99 \log i$, with $X_i \geq \log i$, with $X_i \geq 1.01 \log i$?

(e). Let $X_1, X_2, \ldots$ be i.i.d. standard normal. Show that

$$\Pr(X_i \geq a) = \frac{\phi(a)}{a}\{1 + o(1)\}, \quad \text{as } a \to \infty,$$

and use this to show that there will be infinitely many cases with $|X_i| \geq (2 \log i)^{1/2}$, but only finitely many cases with $|X_i| \geq (2.001 \log i)^{1/2}$. The quantity $\{1 - \Phi(a)\}/\{\phi(a)/a\}$ is sometimes called the Mills ratio.

Ex. 2.3. *The converses do not hold.* The implications arrows in Ex. 2.1(b)–(c) only point in one direction, as we will now see.

(a). Let $X_1, X_2, \ldots$ be Bernoulli random variables with success probabilities $p_1, p_2, \ldots$. Show that $X_n \to_{\text{pr}} 0$ as long as $p_n \to 0$, but that $X_n \to_{\text{a.s.}} 0$ takes more. What is the probability for having infinitely many $X_n = 1$, for $p_n = 1/n^{0.99}$, for $p_n = 1/n$, for $p_n = 1/n^{1.01}$? Conclude that we can have convergence in probability, but not almost surely.

(b). Let U be unif $(0, 1)$, and divide the unit interval in $2, 3, \ldots$ pieces: $A_1 = [0, 1/2]$, $A_2 = (1/2, 1]$, $A_3 = [0, 1/4]$, $A_4 = (1/4, 1/2]$, $A_5 = (1/2, 3/4]$, and so on. Define the random variables $X_n = I(U \in A_n)$, and show that $X_n \to_{\text{pr}} 0$, but not almost surely.

(c). Find an example where $X_n \to_d X$, but there is not convergence in probability.

(d). Let X_n be a sequence of binary random variable with distributions $\Pr(X_n = a_n) = p_n$ and $\Pr(X_n = 0) = 1 - p_n$. Construct the a_n and p_n sequences such that $X_n \to_{\text{pr}} 0$, but $\mathrm{E}\,|X_n|$ does not converge to zero.

(e). Similarly, let $X_n = \theta + Z_n$, for $\theta \in \mathbb{R}$, where $Z_1, Z_2, \ldots$ are independent random variable with distribution $\Pr(Z_n = 2^n) = \Pr(Z_n = -2^n) = 1/(3n)$ and $\Pr(Z_n = 0) = 1 - 2/(3n)$. Show that $X_n \to_{\text{pr}} \theta$, but that $\mathrm{E}\,|X_n - \theta|$ does not converge to zero.

Ex. 2.4. *Partial converses.* As we have just seen, the implication arrows of Ex. 2.1 only point in one direction. We do, however, have certain partial converses.

(a). As as partial converse to Ex. 2.1(b), show that if $X_n \to_{\text{pr}} X$, then there is a subsequence $(n_k)_{k \geq 1}$ such that $X_{n_k} \to_{\text{a.s.}} X$. To construct such a subsequence, you may use the Borel–Cantelli lemma.

(b). Prove the converse of (a), to obtain the lemma saying that $X_n \to_{\text{pr}} X$ if and only if any subsequence (X_{n_k}) has a further subsequence $(X_{n_{k_j}})$ such that $X_{n_{k_j}} \to_{\text{a.s.}} X$.

(c). As a partial converse to Ex. 2.1(c), show that if $X_n \to_d a$ for some constant $a \in \mathbb{R}$, then $X_n \to_{\text{pr}} a$. Generalise to higher dimensions.

(d). As a partial converse to Ex. 2.1(d), we have that if $X_n \to_{\text{pr}} X$ and $\|X_n\|$ is bounded by a random variable Y such that $|Y|^p$ is integrable for some $p \geq 1$, then $\mathrm{E}\,\|X_n - X\|^p \to 0$. This is a version of Lebesgue's dominated convergence theorem, see Ex. A.11(d), and you may consult that exercise before solving this one.

(e). Another partial converse to Ex. 2.1(d) is provided by the following lemma. If X_n is a sequence of nonnegative random vectors such that $X_n \to_{\text{pr}} X$ and $\mathrm{E}\,X_n \to \mathrm{E}\,X$, with X integrable, then $\mathrm{E}\,\|X_n - X\| \to 0$. Prove it first in one dimension, where you may use that $|y| = y + 2\max(-y, 0)$, then generalise.

Ex. 2.5. *Uniform integrability.* Yet another partial converse to Ex. 2.1(d) is obtained by introducing the concept of *uniform integrability*. A sequence $X_1, X_2, \ldots$ of random variables is uniformly integrable if

$$\lim_{M\to\infty} \sup_{n\geq 1} \mathrm{E}\,|X_n| I(|X_n| > M) = 0. \tag{2.6}$$

As we show in this exercise, if $(X_n)_{n\geq 1}$ is uniformly integrable and $X_n \to_{\mathrm{pr}} X$, then $\mathrm{E}\,|X_n - X| \to 0$. Thus, compared to Ex. 2.4(d) (which is essentially Lebesgue's dominated convergence theorem), the uniform integrability assumption takes the place of the domination assumption imposed in the dominated convergence theorem. Moreover, as we show in (c) and (i), uniformly integrability is a bit weaker than boundedness, and turns out to be necessary for convergence in mean.

(a). Show that $\mathrm{E}\,|X| < \infty$ is equivalent to $\lim_{M\to\infty} \mathrm{E}\,|X| I(|X| > M) = 0$. This explains why the property in (2.6) is called *uniform* integrability.

(b). Show that if $(X_n)_{n\geq 1}$ is uniformly integrable, then $\sup_n \mathrm{E}\,|X_n| < \infty$. Exhibit a sequence of random variables for which the converse does not hold.

(c). Show that if $(X_n)_{n\geq 1}$ is dominated by an integrable Y, then it is uniformly integrable. Construct a sequence of random variables that is uniformly integrable, but that is not dominated by an integrable random variable.

(d). Let $a, b \in \mathbb{R}$ and suppose that $(X_n)_{n\geq 1}$ and $(Y_n)_{n\geq 1}$ are uniformly integrable. Show that $aX_n + bY_n$ is uniformly integrable.

(e). Let $(X_n)_{n\geq 1}$ be a uniformly integrable sequence of random variables such that $X_n \to_{\mathrm{a.s.}} X$. Show that the limit X must be integrable, and that $\mathrm{E}\,|X_n - X| \to 0$. To show the integrability part, consult Fatou's lemma Ex. A.11(c). For the convergence in mean part, the trick is to truncate $Y_n = |X_n - X|$ and take it from there.

(f). Show that in (e), it is sufficient that X_n converges in probability to X.

(g). Show that if $(X_n)_{n\geq 1}$ is such that $\sup_n \mathrm{E}\,|X_n|^{1+\delta} < \infty$ for some $\delta > 0$, then it is uniformly integrable.

(h). Show that $(X_n)_{n\geq 1}$ is uniformly integrable if and only if

$$\lim_{M\to\infty} \limsup_{n\to\infty} \mathrm{E}\,|X_n| I(|X_n| > M) = 0,$$

and X_n is integrable for each n.

(i). The claim in (e) has a converse. Show that if $(X_n)_{n\geq 1}$ and X are nonnegative and integrable, $X_n \to_{\mathrm{a.s.}} X$ and $\mathrm{E}\,X_n \to \mathrm{E}\,X$, then $(X_n)_{n\geq 1}$ is uniformly integrable.

(j). Deduce from (i) that if $(X_n)_{n\geq 1}$ and X are integrable (but not necessarily nonnegative), $X_n \to_{\mathrm{a.s.}} X$ and $\mathrm{E}\,|X_n - X| \to 0$, then $(X_n)_{n\geq 1}$ is uniformly integrable.

(k). Show that if X_n^2 and Y_n^2 are uniformly integrable, then $X_n Y_n$ is uniformly integrable. Show also that if X_n^p is uniformly integrable, then so is X_n^r for all $1 \leq r \leq q$.

Ex. 2.6. *Scheffé's lemma.* If we replace the random vectors in Ex. 2.4(e) with the densities $(f_n)_{n\geq 1}$ and f, then the limit is automatically integrable, and we obtain the lemma typically known by the name of Scheffé.

Scheffé's lemma

(a). Let Y_n and Y be random vectors with densities f_n and f with respect to a dominating measure μ. Show that if $f_n(y) \to f(y)$ for μ-almost all y, then $\int |f_n - f| \, d\mu \to 0$, and consequently $Y_n \to_d Y$.

(b). Suppose X_n has the Gamma distribution with parameters (a_n, b), where $a_n \to 1$. Show that X_n tends in distribution to an exponential with mean $1/b$.

(c). Show that under the conditions of (a) we have $\sup_B |\Pr(Y_n \in B) - \Pr(Y \in B)| \to 0$, the supremum being over all Borel sets B. Since $Y_n \to_d Y$ only requires $\Pr(Y_n \in B) \to \Pr(Y \in B)$ for sets of the form $B = (-\infty, y]$ (or, more generally for all continuity sets B, see Ex. 2.18), convergence $\Pr(Y_n \in B) \to \Pr(Y \in B)$ uniformly in B is stronger than what is required for convergence in distribution: For a classical example thereof, consider the sequence $Y_n = 1/n$. Clearly, $Y_n \to_d 0$, but show that $\sup_B |\Pr(Y_n \in B) - \Pr(Y \in B)| = 1$.

(d). For the case of Y_n and Y on the integers $0, 1, 2, \ldots$, show that $p_n(j) = \Pr(Y_n = j)$ converges to $p(j) = \Pr(Y = j)$ for each j, if and only if $Y_n \to_d Y$.

(e). If Y_n and Y have densities f_n and f, and $Y_n \to_d Y$, we should expect $f_n \to f$. This is not alway happening, however. Consider the case of $F_n(y) = y + (n\pi)^{-1} \sin(n\pi y)$. Plot the F_n and its density f_n, for some n. Show that $Y_n \to_d \text{unif}$, but that $f_n(y)$ does not converge to 1 for all y. It is also instructive to transform to the approximate normal scale, via $X = \Phi^{-1}(Y)$. Show that X then has density $g_n(x) = \phi(x)\{1 + \cos(n\pi\Phi(x))\}$, with very notable oscillations, but where the c.d.f. $G_n(x)$ nevertheless tends to the standard normal.

Ex. 2.7. *From discrete to continuous.* Often enough discrete distributions have continuous limits.

(a). Let X_n have distribution $\Pr(X_n = j/n) = 1/(n+1)$ for $j = 0, 1, \ldots, n$. Show that $X_n \to_d X$, where X has the uniform distribution on the unit interval. Show also that $\sup_B |\Pr(X_n \in B) - \Pr(X \in B)|$ does *not* converge to zero (cf. Ex. 2.6(c)).

(b). With perhaps similar techniques as in (a), consider X_n with distribution $\Pr(X_n = j/n) = j/\{n(n+1)/2\}$ for $j = 1, \ldots, n$. Find its limit distribution.

(c). Let $X_1, X_2, \ldots$ be independent Bernoulli random variables with success probability p. Only using (2.4), show that $\sqrt{n}(\bar{X}_n - p) \to_d \text{N}(0, p(1-p))$. You ought to use Stirling's formula $n! \sim \sqrt{2\pi n}(n/e)^n$ here; we prove that formula via probabilistic tricks in Ex. 2.35.

Ex. 2.8. *Many small probabilities give a Poisson.* The Law of Small Numbers says that if we sum a high number of 0-1 variables, with each having a small probability of 1, then we're close to a Poisson. This was touched upon in Ex. 1.24, but here we give a clearer picture, with more to come in Ex. 2.29, with natural setups yielding Poisson limits.

der Gesetz der kleinen Zahlen

(a). Suppose $Y_n \sim \text{binom}(n, p_n)$, with p_n becoming small with growing n in a way which has $np_n \to \lambda$. Show that $Y_n \to_d \text{Pois}(\lambda)$. As a small demonstration, take $n = 100$ and $p = 0.05$. Compare the exact and the limiting probabilities.

(b). Suppose (X_n, Y_n) has the trinomial distribution, with parameters (n, p_n, q_n), see Ex. 1.4, where we also found the correlation between them. Assume now that p_n, q_n become small with n, such that $np_n \to \lambda_1$, $nq_n \to \lambda_2$. Show that the correlation betweem X_n and Y_n tends to zero, and that $(X_n, Y_n) \to_d (X, Y)$, where X and Y are independent and Poisson with parameters λ_1, λ_2. Generalise to the multinomial situation.

Ex. 2.9. *Maximum of uniforms.* Let $X_1, \ldots, X_n$ be i.i.d. random variables with the uniform distribution on $(0, \theta)$, and set $M_n = \max_{i \le n} X_i$.

(a). Show that $M_n \to_{\rm pr} \theta$, that is, the maximum of the observations is consistent for the unknown endpoint.

(b). Find the limit distribution of $V_n = n(\theta - M_n)$, and use this result to find an approximate 90 percent confidence interval for θ, i.e., a random interval $[L_n, U_n]$ such that $\Pr(L_n \le \theta \le U_n)$ is approximately 0.90.

Convergence in probability and tail bounds

Ex. 2.10. *Markov, Chebyshov, and the Law of Large Numbers.* In view of the definition of convergence in probability in (2.1), it is certainly useful to find mathematically manageable bounds for so-called tail probabilities, i.e., $\Pr(|X| \ge a)$ for a random variable X. There are several such, as we learn here, with more to come in Ex. 2.16. Their uses include assessing how likely it might be that an estimator is some distance off its target.

(a). Markov's inequality was proven in Ex. 2.1(d): If X is nonnegative, then $\Pr(X \ge a) \le \mathrm{E}\,(X)/a$ for each $a > 0$. More generally, if $h(x)$ nonnegative and nondecreasing, with $h(a)$ positive, show that $\Pr(X \ge a) \le \mathrm{E}\,h(X)/h(a)$. From this, deduce the Chebyshov inequality: if X has finite variance (but is not necessarily nonnegative), then

the Chebyshov inequality

$$\Pr(|X - \mathrm{E}\,X| \ge \varepsilon) \le \mathrm{Var}\,(X)/\varepsilon^2, \quad \text{for } \varepsilon > 0.$$

(b). With $X_1, \ldots, X_n$ being i.i.d. from the same distribution as X, with mean ξ and standard deviation σ, show that for the empirical mean $\bar{X}_n$ that $\Pr(|\bar{X}_n - \xi| \ge \varepsilon) \le \sigma^2/(n\varepsilon^2)$. Since the right hand side tends to zero as n tends to infinity, you have now proven the version of the *Law of Large Numbers* (LLN) that says that the empirical mean of n i.i.d. random variables with finite second moments converges in probability to its expectation as n tends to infinity.

the LLN via Chebyshov

(c). Let $Y_1, Y_2, \ldots$ be i.i.d. random variables with expectation θ and finite variance σ^2. Consider the weighted average $\widehat{\theta}_n = \sum_{i=1}^n w_i Y_i / \sum_{i=1}^n w_i$, for nonnegative and fixed weights w_i. Give a condition for consistency ensuring $\widehat{\theta}_n \to_{\rm pr} \theta$, in terms of these weights. What happens for $w_i = 1/i$, and for $w_i = 1/i^{1.5}$?

(d). If X has mean ξ, and a finite fourth moment, show that $\Pr(|X - \xi| \ge \varepsilon) \le \mathrm{E}\,|X - \xi|^4/\varepsilon^4$. For $X \sim \mathrm{N}(\xi, \sigma^2)$, deduce that $\Pr(|X - \xi| \ge \varepsilon) \le 3\sigma^4/\varepsilon^4$.

(e). With $X_1, \ldots, X_n$ i.i.d. from a distribution with finite fourth moment, write $\gamma_4 = \mathrm{E}\,\{(X_i - \xi)/\sigma\}^4 - 3$ for its kurtosis. Show that

$$\mathrm{E}\,|\bar{X}_n - \xi|^4 = \frac{\sigma^4}{n^2}\{3 + (1/n)(\gamma_4 - 3)\}.$$

Show hence that $\Pr(|\bar{X}_n - \xi| \ge \varepsilon) \le 3.01\,\sigma^4/(n^2\varepsilon^2)$, for all large enough n. When is this a sharper result than that of the Chebyshov inequality?

Ex. 2.11. *The binomial and the empirical distribution function.* For i.i.d. observations $Y_1, \ldots, Y_n$, we form *the empirical cumulative distribution function*, the e.c.d.f., with $F_n(t) = n^{-1} \sum_{i=1}^n I(Y_i \le t)$. Plotting F_n for a given dataset is informative, also for comparing with given distributions.

(a). Before returning to F_n, consider $X_n \sim \text{binom}(n,p)$ and the familiar ratio $\widehat{p}_n = X_n/n$. Show that indeed $\widehat{p}_n \to_{\text{pr}} p$. Historically speaking, this is the earliest clear LLN statement.

(b). Coming back to the empirical c.d.f., explain that $F_n(t)$ for given t is just a binomial proportion, and that $F_n(t) \to_{\text{pr}} F(t)$ for each t. In particular, the events $F_n(t) \le F(t)-\varepsilon$ and $F_n(t) \ge F(t)+\varepsilon$ both have probabilities going to zero, for positive ε.

(c). We come back to finer analysis of F_n in Ex. 2.61 and 9.25, but find the variance of $F_n(t)$ and the covariance betweem $F_n(t_1)$ and $F_n(t_2)$.

Ex. 2.12. *Continuous mapping for convergence in probability.* The theorem known as the continuous mapping theorem (which will be proved in its entirety during the course of this chapter) says that the three modes of convergence introduced above are preserved under a continuous mapping. In this exercise we look at the convergence in probability part of this theorem.

(a). Let X_n be a sequence of random variables such that $X_n \to_{\text{pr}} a$, with a being a constant. Show that if g is a function continuous at point $x = a$, then $g(X_n) \to_{\text{pr}} g(a)$.

(b). Suppose more generally that $X_n \to_{\text{pr}} X$, with the limit being a random variable such that $\Pr(X \in C) = 1$, for $C \subset \mathbb{R}$. Show that if g is uniformly continuous on C, then $g(X_n) \to_{\text{pr}} g(X)$.

(c). Show that in (b) it is sufficient that g is continuous on C. There are several ways of showing this, one involves the subsequences result from Ex. 2.4(a).

(d). Generalise the above to sequences $X_n = (X_{n,1}, \ldots, X_{n,k})$ of random vectors and continuous functions $g\colon \mathbb{R}^k \to \mathbb{R}^m$.

Ex. 2.13. *Toeplitz lemma and convergence in probability.* Toeplitz lemma says that if $a_1, a_2, \ldots$ is a nonnegative sequence such that $\sum_{j=1}^{\infty} a_j = \infty$ and $x_1, x_2, \ldots$ is a sequence such that $x_n \to x$, then $\sum_{j=1}^{n} a_j x_j / \sum_{j=1}^{n} a_j \to x$ as $n \to \infty$. If we set $a_j = 1$ for all j, the lemma says that $\bar{x}_n = n^{-1} \sum_{i=1}^{n} x_i \to x$ as long as $x_n \to x$. In this exercise we look into the conditions under which this carries over to the situation where x_n is replaced by a sequence X_n of random variables converging in probability to some random or constant limit. The example in (b) and the lemma in (c) are from Linero and Rosalsky (2013).

(a). Prove the Toeplitz lemma.

(b). Let X_n be a sequence of independent random variables with distributions $\Pr(X_n = 0) = 1 - n^{-1}$ and $\Pr(X_n = n) = n^{-1}$. Show that $X_n \to_p 0$, and that the empirical means $\bar{X}_n = n^{-1} \sum_{i=1}^{n} X_i$ is such that its characteristic function

$$\text{E}\, \exp(it\bar{X}_n) \to \exp\left(\int_0^1 \frac{e^{itx} - 1}{x} \, \text{d}x \right),$$

which you may peek ahead to Ex. 2.36 to show. As a consequence of Ex. 2.28, the limit here is the charactertistic function of a nondegenerate random variable, say Z. We are thus in the situation that $X_n \to_p 0$, while $\bar{X}_n \to_d Z$, which shows that Toeplitz lemma does not translate directly to convergence in probability.

(c). Suppose that X_n is a uniformly integrable sequence of random variables such that $X_n \to_p 0$, and let $\bar{X}_n = n^{-1}\sum_{i=1}^n X_i$. Show that $\bar{X}_n \to_p X$. You may look back at Ex. 2.5(f) before attempting to prove this.

Ex. 2.14. *Quantiles.* Here we learn that the empirical quantiles are consistent for their population counterparts. Further informative analyses are in Ex. 3.21, including convergence in distribution, but here we limit attention to direct convergence in probability.

(a). Suppose first that $U_1, \ldots, U_n$ are i.i.d. from the uniform distribution on the unit interval, and let M_n be the empirical median. With the empirical c.d.f. F_n from Ex. 2.11, use the fact that $F_n(t) \to_{\rm pr} t$, for each t, to show that $M_n \to_{\rm pr} \frac{1}{2}$. Then generalise to the case of observations $Y_1, \ldots, Y_n$ from a distribution with continuous and strictly increasing c.d.f. F, showing for the empirical median that $M_n \to_{\rm pr} F^{-1}(\frac{1}{2})$.

floor function

(b). More generally, for a $q \in (0,1)$, let $Q_n = Q_n(q)$ be the empirical q-quantile, defined here to be $Y_{(\lfloor nq \rfloor)}$, the $\lfloor nq \rfloor$ order statistic, where $\lfloor x \rfloor = \max\{m \in \mathbb{Z}\colon m \le x\}$. Show that Q_n is consistent for the population quantile $F^{-1}(q)$.

(c). Show that the interquartile range, the 0.75 quantile minus the 0.25 quantile, is consistent for the population interquartile range $F^{-1}(0.75) - F^{-1}(0.25)$. Show more generally that if $\theta = g(F^{-1}(q_1), \ldots, F^{-1}(q_r))$ is any continuous function of a finite number of quantiles, then $\widehat{\theta} = g(Q_n(q_1), \ldots, Q_n(q_r))$ is consistent for θ.

Ex. 2.15. *Improving on the weak LLN.* We have seen in Ex. 2.10(a) that the LLN holds if the distribution has a finite variance. Here we get rid of the finite variance condition. For the almost sure version of the LLN see Ex. 2.58.

(a). Let $X_1, \ldots, X_n$ be i.i.d. with finite mean ξ. Truncate the random variables involved, i.e., write $X_i = X_i I_{|X_i| \le M} + X_i I_{|X_i| > M}$ for some $M > 0$, and show that for any $\varepsilon > 0$,

$$\Pr(|\bar{X}_n - \xi| \ge \varepsilon) \le \frac{8M}{n\varepsilon^2}\,\mathrm{E}\,|X| + \frac{4}{\varepsilon}\,\mathrm{E}\,|X-\xi| I_{|X-\xi|>M},$$

law of large numbers

(b). Let $X_1, X_2, \ldots$ be i.i.d. random variables with finite mean ξ, i.e., no assumption is made about its second moment. Show that $\bar{X}_n = n^{-1}\sum_{i=1}^n X_i \to_p \xi$.

(c). Let $X_1, X_2, \ldots$ be i.i.d. random variables with finite variance σ^2. Show that the empirical variance $n^{-1}\sum_{i=1}^n (X_i - \bar{X}_n)^2$ is consistent for σ^2.

Ex. 2.16. *Further tail bound inequalities.* In Ex. 2.10 we learned about the Markov and Chebyshov inequalities; here we work out further tail bounds for our toolboxes.

(a). Suppose X has a finite m.g.f. $M(t) = \mathrm{E}\,\exp(tX)$, as per Ex. 1.31. Show that

$$\Pr(X \ge a) \le q(a) = \min\{t \ge 0\colon \exp(-ta)M(t)\}.$$

Writing $M(t) = \exp\{K(t)\}$, show that this leads to $q(a) = \exp\{K(t_a) - at_a\}$, where t_a is the solution to $K'(t_a) = a$, if positive. Apply this to $X \sim \mathrm{N}(0,1)$ and a positive, and show that $\Pr(X \ge a) \le \exp(-\frac{1}{2}a^2)$. Show that this is indeed sharper than the tail bound $1/a^2$, from the simpler Chebyshov inequality, for all $a > 0$.

(b). For $X \sim \mathrm{N}(0,1)$ and a positive, show that

$$\Pr(|X| \ge a) \text{ is smaller than each of } \frac{1}{a^2}, \frac{3}{a^4}, \frac{15}{a^6}, 2\exp(-\tfrac{1}{2}a^2).$$

(xx a bit more here, rounding it off. more inequalities in Ch. 2. xx)

(c). Let $X \sim \chi^2_m$, which has mean and variance m and $2m$. Consider $p_m(a) = \Pr(X \geq m + am^{1/2})$. Show that $p_m(a) \leq q_m(a)$, where

$$\begin{aligned} q_m(a) &= \min\{t \geq 0 \colon \exp(-ta - t/\sqrt{m})(1 - 2t/\sqrt{m})^{-m/2} \\ &= \exp(-\tfrac{1}{2}a\sqrt{m})(1 + a/\sqrt{m})^{m/2}. \end{aligned}$$

Compare this bound both with bounds from the Markov inequality, and with the exact limit of $p_m(a)$, as m grows.

(d). Let $X_1, X_2, \ldots$ be i.i.d. with mean zero and variance one, so that $\sqrt{n}\bar{X}_n \to_d \mathrm{N}(0,1)$, but you don't need to use that here. Assume the m.g.f. $M(t) = \mathrm{E}\exp(tX_1) = \exp\{K(t)\}$ is finite. Show that

$$\Pr(\sqrt{n}\bar{X}_n \geq a) \leq \frac{M(t/\sqrt{n})^n}{\exp(ta)} = \exp\{nK(t/\sqrt{n}) - ta\},$$

for each t. (xx then a bit more. tail inequality. not too far from good bound $\exp(-\frac{1}{2}a^2)$. briefly mention and point to large deviations theory. xx)

Ex. 2.17. *Bernshteĭn and Weierstraß.* In c. 1885, Karl Weierstraß proved one of the fundamental and insightful results of approximation theory, that any given continuous function can be approximated uniformly well, on any finite interval, by polynomials. Here we give a constructive and relatively straightforward proof of the Weierstraß theorem, involving so-called Bernshteĭn polynomials. Let $g \colon [0,1] \to \mathbb{R}$ be continuous, and construct

$$B_n(p) = \mathrm{E}_p\, g(X_n/n) = \sum_{j=0}^{n} g(j/n)\binom{n}{j} p^j (1-p)^{n-j} \quad \text{for } p \in [0,1],$$

where $X_n \sim \mathrm{binom}(n,p)$. Note that $B_n(p)$ is a polynomial of degree n.

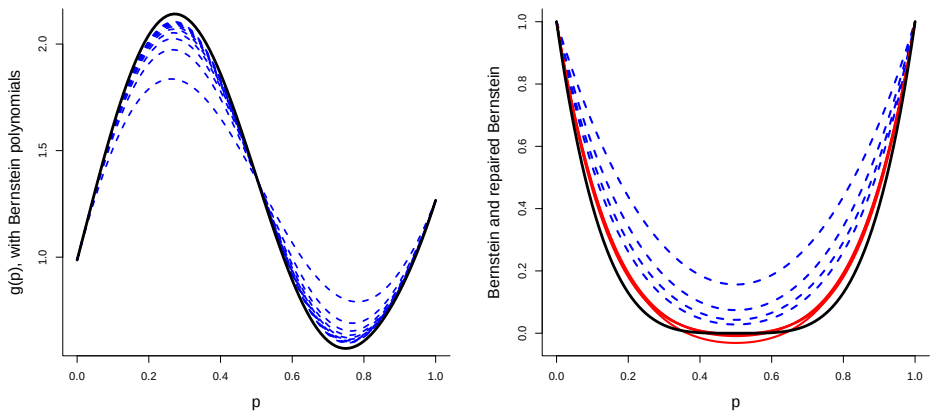


Figure 2.1: Left panel: the given non-polynomial function $g(p)$ (full black curve), along with approximating Bernshteĭn polynomials of order $10, 20, \ldots, 100$. Right panel: for the function $g_2(p) = (2p-1)^4$, the Bernshteĭn polynomials B_n (dashed curves), along with the bias-modified ones B_n^* (full curves), of order $4, 6, 8, 10$.

(a). Show that $B_n(p) \to_{\rm pr} g(p)$, for each p. Then show that the convergence is actually uniform. In some detail, for $\varepsilon > 0$, find $\delta > 0$ such that $|x-y| < \delta$ implies $|g(x)-g(y)| < \varepsilon$ (which is possible, as a continuous function on a compact interval is always uniformly continuous). Then show

$$|B_n(p) - g(p)| \le \varepsilon + 2M \Pr(|X_n/n - p| \ge \delta),$$

with M a bound on $|g(p)|$. Show from this that indeed $\max_p |B_n(p) - g(p)| \to 0$.

(b). Consider the marvellous function

$$g_1(p) = \sin(2\pi p) + \exp(1.234 \sin^3 \sqrt{p}) - \exp(-4.321 \cos^5 p^2)$$

on the unit interval. Compute the Bernshteĭn polynomials of various orders, and display these in a diagram, alongside the curve of g. Construct a version of Figure 2.1, left panel, which does this for $n = 10, 20, \ldots, 90, 100$. How high n is needed for the maximum absolute difference to creep below 0.01?

(c). Assuming smoothness of the g function, and writing $\widehat{p} = X_n/n$ for the binomial fraction, use Taylor expansion $g(\widehat{p}) = g(p) + g'(p)(\widehat{p}-p) + \frac12 g''(p)(\widehat{p}-p)^2 + O_{\rm pr}(|\widehat{p}-p|^3)$ to show that $B_n(p) = g(p) + \frac12 g''(p)p(1-p)/n + o(1/n)$. This invites a repaired Bernshteĭn polynomial, of the form $B_n^*(p) = B_n(p) - \frac12 C_n(p)p(1-p)/n$, with $C_n(p) = {\rm E}_p g''(X_n/n)$. Explain that this is still a polynomial, now of order $n+2$, and show in examples that it often succeeds better than B_n in coming close to the target $g(p)$. Make a version of Figure 2.1, right panel, which shows the B_n and the B_n^*, of order $4, 6, 8, 10$, for the case of $g_2(p) = (2p-1)^4$. Check here that $\max_p |B_n^*(p) - g(p)|$ becomes rather smaller than $\max_p |B_n(p) - g(p)|$.

(d). Assuming g has a derivative, construct a confidence band of the type $B_n(p) \pm 1.96\,\widehat{\sigma}(p)/\sqrt{n}$, for a certain $\widehat{\sigma}(p)$, with the property that it for each given p covers the underlying $g(p)$ with probability tending to 0.95.

(e). Let now $g(x,y)$ be an arbitrary function on the unit simplex, $\{(x,y)\colon x \ge 0, y \ge 0, x+y \le 1\}$. Construct a mixed polynomial $B_n(x,y)$ of degree n such that it converges uniformly to g on the simplex.

Convergence in distribution

Ex. 2.18. *The Portmanteau theorem.* So far we have taken as our definition of $X_n \to_d X$ that we have pointwise convergence $F_n(x) = \Pr(X_n \le x) \to \Pr(X_n \le x) = F(x)$ at all points x where F is continuous. A type of convergence we denote $F_n \Rightarrow F$. A limitation of this definition is that it only works for random vectors, and since we soon enough want to study convergence in distribution in spaces other than $\mathbb{R}^k$, we need a more general definition. Let $(\mathcal{X}, d)$ be a metric space equipped with its Borel-σ-algebra $\mathcal{B} = \mathcal{B}(\mathcal{X})$ (i.e., the smallest σ-algebra containing all the open sets), and let P_n and P be probability measures on the measurable space $(\mathcal{X}, \mathcal{B})$, typically the distributions of some random elements X_n and X taking values in $\mathcal{X}$. We say that P_n *converges weakly* to P, denoted $P_n \Rightarrow P$, if

weak convergence $P_n \Rightarrow P$

$$\int g \,{\rm d}P_n \to \int g \,{\rm d}P, \tag{2.7}$$

for all bounded and continuous functions $g\colon \mathcal{X} \to \mathbb{R}$, a collection of functions we denote $C_b(\mathcal{X})$. When P_n and P are the distributions of some random elements X_n and X,

$C_b(\mathcal{X})$

we might just write $\mathrm{E}\, g(X_n) \to \mathrm{E}\, g(X)$ for the convergence in (2.7). If X_n and X are random vectors, so that $\mathcal{X} = \mathbb{R}^k$ for some $k \geq 1$, then (2.7) is equivalent to $X_n \to_d X$ as defined in (2.4), a fact we prove in (e) of this exercise. As will become clear as we proceed, however, the definition in (2.7) is vastly more general than (2.4), in particular, it works when $\mathcal{X}$ is a function space, see Ch. 9. It is also often much easier to work with. In this exercise we first prove the 'bare bones', or classical, Portmanteau theorem, valid for all types of metric spaces, then proceed to proving its equivalence with $X_n \to_d X$. Ex. 2.19 presents additional equivalent statements. The classical Portmanteau theorem states that the following four statements are equivalent:

(i) $P_n \Rightarrow P$;

(ii) $\limsup_n P_n(F) \leq P(F)$ for every closed set F;

(iii) $\liminf_n P_n(B) \geq P(B)$ for every open set B;

(iv) $P_n(C) \to P(C)$ for every set C that is P-continuous, in the sense that $P(\partial C) = 0$, where $\partial C = \bar{C} \setminus C^\circ$ is the boundary of C (the closure minus the interior);

We prove the Portmanteau theorem through a string of subexercises.

(a). For any set A define $d(x, A) = \inf\{d(x, y)\colon y \in A\}$, i.e., the distance from x to A, and define the function

$$\eta(t) = \begin{cases} 1, & \text{if } t \leq 0, \\ 1 - t, & \text{if } 0 \leq t \leq 1, \\ 0, & \text{if } t \geq 1. \end{cases}$$

For some $\varepsilon > 0$, the function $\eta(d(x, A)/\varepsilon)$ will be used to approximate the indicator function I_A. Let $F = [a, b]$ be a closed interval on the real line, and make a sketch $\eta(d(x, F)/\varepsilon)$. Show that for any closed set F, that $x \mapsto \eta(d(x, A)/\varepsilon)$ is bounded and continuous, and that for $F^\varepsilon = \{x\colon d(x, F) < \varepsilon\}$

$$I_F(x) \leq \eta(d(x, A)/\varepsilon) \leq I_{F^\varepsilon}(x).$$

(b). Use the inequalities from (a) to show that (i) implies (ii) and (iii).

(c). Show that (ii) and (iii) imply (iv).

(d). To show that (iv) implies (i), take f to have its values inside $[0, 1]$, without loss of generality, write

$$\mathrm{E}\, f(X_n) = \int \int_0^1 I\{y \leq f(x)\}\, \mathrm{d}y\, \mathrm{d}P_n(x) = \int_0^1 \Pr\{f(X_n) \geq y\}\, \mathrm{d}y,$$

and finish up the argument. You may prove and use that for a continuous function f, we have the inclusions $\overline{f^{-1}(B)} \subset f^{-1}(\bar{B})$ and $f^{-1}(B^\circ) \subset (f^{-1}(B))^\circ$, which entails the inclusion $\partial f^{-1}(B) \subset f^{-1}(\partial B)$.

(e). Finally, consider the setting where have c.d.f.'s $F_n(x) = P_n(-\infty, x]$ and $F(x) = P(-\infty, x]$. Show that for $P_n \Rightarrow P$ it is necessary and sufficient that $F_n \Rightarrow F$. To show the sufficiency part, express the open set $B \subset \mathbb{R}$ as a union of the disjoint open intervals $(a_1, b_1), (a_2, b_2), \ldots$, and deduce that $F_n \Rightarrow F$ implies (iii) of the Portmanteau theorem.

Ex. 2.19. *More in the Portmanteau.* It is often convenient to have a smaller class of functions to work with than the bounded and continuous one. In this exercise we show that $P_n \Rightarrow P$ is equivalent to $\int f \,\mathrm{d}P_n \to \int f \,\mathrm{d}P$ for all bounded and continuous functions with bounded and continuous derivatives of all orders. This equivalence can be shown in several ways, the trick we use here is from Pollard (1984).

(a). As a preliminary that will be useful below, for $k = 1, 2, \ldots$ consider the functions $g_k(x, u) = |x - u|^k \exp(-\frac{1}{2}(x - u)^2)$. For a given $\varepsilon > 0$ define

$$K_k(u) = \begin{cases} |u + \varepsilon|^k \exp(-\frac{1}{2}(u + \varepsilon)^2), & u < -k^{1/2} - \varepsilon, \\ k^{k/2} \exp(-\frac{1}{2}k), & -k^{1/2} - \varepsilon \le u < k^{1/2} + \varepsilon, \\ |u - \varepsilon|^k \exp(-\frac{1}{2}(u - \varepsilon)^2), & k^{1/2} + \varepsilon \le u, \end{cases}$$

and show that $g_k(x, u) \le K_k(u)$ for all $x \in (-\varepsilon, \varepsilon)$ and that $\int K_k(u) \,\mathrm{d}u < \infty$.

(b). Let $f \colon \mathbb{R} \to \mathbb{R}$ be a bounded and continuous function. With X having any distribution, $Z \sim \mathrm{N}(0, 1)$ independent of X, and $\sigma > 0$, show that $\mathrm{E}\, f(X + \sigma Z) = \mathrm{E}\, f_\sigma(X)$ where (see Ex. A.40)

$$f_\sigma(x) = \frac{1}{\sqrt{2\pi}\sigma} \int f(u) \exp(-\tfrac{1}{2}(x - u)^2/\sigma^2) \,\mathrm{d}u.$$

Show that $f_\sigma(x)$ is bounded and continuous, and has bounded and continuous derivatives of all orders.

(c). Let $Z_1, Z_2, \ldots$ be a standard normal random variables with Z_i independent of X_i for each $i \ge 1$. Show that $X_n \to_d X$ if and only if $X_n + \sigma Z \to_d X + \sigma Z$ for each $\sigma > 0$.

(d). From (b) and (c), conclude that $X_n \to_d X$ if and only if $\mathrm{E}\, f(X_n) \to \mathrm{E}\, f(X)$ for all bounded and continuous functions, having bounded and continuous derivatives of all orders. This version of the Portmanteau theorem will be of use in Ex. 2.31.

continuous mapping

Ex. 2.20. *Continuous mapping for convergence in distribution.* Let $X_1, X_2, \ldots$ and X be k-dimensional random vectors, and $g \colon \mathbb{R}^k \to \mathbb{R}^m$ a function that is continuous on a set $C \subset \mathbb{R}^k$ such that $\Pr(X \in C) = 1$. Suppose that $X_n \to_d X$.

(a). Assume that g is continuous on all of $\mathbb{R}^k$, and show that $g(X_n) \to_d g(X)$.

(b). Now suppose that C is a proper subset of $\mathbb{R}^k$, and let F be a closed set. The challenge here is that while the sets $\{g(X_n) \in F\}$ and $\{X_n \in g^{-1}(F)\}$ are identical, the set $g^{-1}(F)$ need not be closed, as g is only assumed continuous on C. Use the fact that a set is included in its closure followed by the Portmanteau theorem to establish that

$$\limsup_n \Pr(X_n \in g^{-1}(F)) \le \Pr(X \in \overline{g^{-1}(F)}).$$

It then remains to prove that $\Pr(X \in \overline{g^{-1}(F)}) \le \Pr(X \in g^{-1}(F))$ as the reverse inequality is immediate from the inclusion of a set in its closure. To show this inequality, start by combining the continuity of g on C and the fact that a closed set contains all its limits points to show that $\overline{g^{-1}(F)} \cap C \subset g^{-1}(F)$, and take it from there.

Ex. 2.21. *Skorokhod's theorem.* If $X_n \to_d X$, we can construct random variables Y_n and Y on a common probability space such that $Y_n \sim X_n$ and $Y \sim X$ and $Y_n \to_{\text{a.s.}} Y$. This result simplifies many proofs.

probability transform

(a). Let $U \sim \text{unif}(0,1)$. For any c.d.f. F, define $X(u) = \inf\{x\colon F(x) \geq u\}$ for $0 < u < 1$, so $X(u) = F^{-1}(u)$ whenever F^{-1} exists. Show that $X(U) \sim F$.

(b). Suppose that $F_n(x) \to F(x)$ for all continuity points x of F. Consider the unit interval equipped with its Borel σ-algebra and Lebesgue measure, say μ. On this space, construct the random variables $Y_n(\omega) = \inf\{x\colon F_n(x) \geq \omega\}$ and $Y(\omega) = \inf\{x\colon F(x) \geq \omega\}$, and show that Y_n converges μ-almost surely to Y.

(c). Show that if $X_n \to_d X$ and $(X_n)_{n\geq 1}$ is uniformly integrable, then X is integrable and $\mathrm{E}\, X_n \to \mathrm{E}\, X$.

(d). Prove the continuous mapping theorem, Ex. 2.20(b), once over, this time using Skorokhod's theorem.

(e). Use Skorokhod's theorem to prove that (2.4) implies (i) and (iv) of the Portmanteau theorem.

Ex. 2.22. *The Cramér–Slutsky rules.* The usefulness of the three results in (c) below, together known as the Cramér–Slutksy rules, will become abundantly clear as we progress.

(a). Show that if $X_n \to_d X$, and $Y_n - X_n \to_{\text{pr}} 0$, then also $Y_n \to_d X$. This says that variables which are essentially close, for growing n, have identical limit distributions.

(b). Show that if X_n and Y_n are sequences of random vectors such that $X_n \to_d X$ and $Y_n \to_{\text{pr}} a$, for a random variable X and a constant a, then $(X_n, Y_n) \to_d (X, a)$.

(c). Show that if $X_n \to_d X$ and $Y_n \to_{\text{pr}} a$, as above, then Cramér–Slutsky

(i) $X_n + Y_n \to_d X + a$;
(ii) $X_n Y_n \to_d Xa$;
(iii) $X_n / Y_n \to_d X/a$, provided $a \neq 0$.

Explain why rules (i) and (ii) also hold when Y_n and a are matrices.

(d). Let $X_1, \ldots, X_n$ be i.i.d. Bernoulli(p). Looking back at Ex. 2.7(c) and Ex. 2.15(c), show that $\sqrt{n}(\bar{X}_n - p)/\widehat{\sigma}^2 \to_d \mathrm{N}(0,1)$, where $\widehat{\sigma}^2 = \widehat{p}(1-\widehat{p})$.

(e). Suppose $X_n \to_d X$, in dimension p, and that $A_n \to_{\text{pr}} A$, a symmetric matrix, in the sense that $A_{n,i,j} \to_{\text{pr}} A_{i,j}$ for each component (i,j). Show that $X_n^{\text{t}} A_n X_n \to_d X^{\text{t}} A X$.

(f). Assume $X_n \to_d \mathrm{N}_p(0, \Sigma)$, with a positive definite Σ, and that there is a consistent estimator $\widehat{\Sigma}$ for this matrix. Show that $X_n^{\text{t}} \widehat{\Sigma}^{-1} X_n \to_d \chi^2_p$.

(g). The claim in (b) is in general not true if the constant limit in probability a is replaced with a (nondegenerate) random variable. Construct an example where $X_n \to_d X$ and $Y_n \to_{\text{pr}} Y$, for a random variable Y, but (X_n, Y_n) does not converge in distribution.

Ex. 2.23. *Stochastic O_{pr} and o_{pr} symbols.* The symbols $O_{\text{pr}}(\cdot)$ and $o_{\text{pr}}(\cdot)$ are the stochastic analogues of $O(\cdot)$ and $o(\cdot)$. Let X_n be a sequence of random variables. Then $X_n = o_{\text{pr}}(1)$ if $X_n \to_{\text{pr}} 0$; and $X_n = O_{\text{pr}}(1)$ if X_n is bounded in probability, i.e., for any $\varepsilon > 0$ there is an M such that $\Pr(|X_n| > M) < \varepsilon$ for all n. If r_n is another sequence of random variables, then $X_n = o_p(r_n)$ if $X_n = r_n o_{\text{pr}}(1)$, and $X_n = O_{\text{pr}}(r_n)$ if $X_n = r_n O_{\text{pr}}(1)$. Typically, r_n will be a sequence of constants, $r_n = n^{-1}$ and $r_n = n^{-1/2}$ for example, appear frequently in parametric statistics.

(a). Show that $X_n = o_{\mathrm{pr}}(r_n)$ if and only if $\Pr(|X_n| \geq \varepsilon r_n) \to 0$ for each $\varepsilon > 0$.

(b). Show that $X_n = O_{\mathrm{pr}}(r_n)$ if and only if for any $\varepsilon > 0$ there are M and n_0 such that $\Pr(|X_n| \geq M r_n) < \varepsilon$ for all $n \geq n_0$.

(c). Show that if $\sqrt{n}X_n \to_d \mathrm{N}(0, \sigma^2)$, then $X_n = O_{\mathrm{pr}}(1/\sqrt{n})$. More generally, show that if $n^\alpha X_n \to_d Z$, then $X_n = O_{\mathrm{pr}}(n^{-\alpha})$.

(d). Show that $o_{\mathrm{pr}}(1) + o_{\mathrm{pr}}(1) = o_{\mathrm{pr}}(1)$; $o_{\mathrm{pr}}(1) + O_{\mathrm{pr}}(1) = O_{\mathrm{pr}}(1)$; $o_{\mathrm{pr}}(1)O_{\mathrm{pr}}(1) = o_{\mathrm{pr}}(1)$; $1/(1 + o_{\mathrm{pr}}(1)) = O_{\mathrm{pr}}(1)$; $o_{\mathrm{pr}}(O_{\mathrm{pr}}(1)) = o_{\mathrm{pr}}(1)$; and $o_{\mathrm{pr}}(1)/(1 + o_{\mathrm{pr}}(1)) = o_{\mathrm{pr}}(1)$.

(e). Let $r\colon \mathbb{R}^k \to \mathbb{R}$ be a function such that $g(0) = 0$, and X_n a sequence of random vectors. Show that if $r(x) = o(\|x\|^p)$ as $x \to 0$, then $R(X_n) = o_p(\|X_n\|^p)$; and that if $r(x) = O(\|x\|^p)$ as $x \to 0$, then $r(X_n) = O_p(\|X_n\|^p)$.

(f). Let $g\colon \mathbb{R} \to \mathbb{R}$ be a function such that $g(\theta) - g(\theta_0) = g'(\theta_0)(\theta - \theta_0) + o(|\theta - \theta_0|)$ as $\theta \to \theta_0$, and let $\widehat{\theta}_n$ be an estimator sequence such that $\sqrt{n}(\widehat{\theta}_n - \theta_0)$ converges in distribution. Show that $\sqrt{n}(g(\widehat{\theta}_n) - g(\theta_0)) = g'(\theta_0)\sqrt{n}(\widehat{\theta}_n - \theta_0) + o_p(1)$ using the rules about $o_{\mathrm{pr}}(\cdot)$ and $O_{\mathrm{pr}}(\cdot)$ that you have derived above in the above exercises.

(g). Suppose that $X_n = o_{\mathrm{pr}}(n^{-\alpha})$ for some $\alpha > 0$, and that $n^\alpha X_n$ is dominated by an integrable random variable Y, or is uniformly integrable (which is weaker, as we recall from Ex. 2.5). Show that $\mathrm{E}\, X_n = o(n^{-\alpha})$.

Ex. 2.24. *Tightness.* Consider the sequence $(X_n)_{n\geq 1}$ of random variables with distribution $\Pr(X_n = x) = 1/3$ for $x = 0, 1/2, n$. You may verify that the sequence $F_n(x) = \Pr(X_n \leq x)$ of distribution functions converges pointwise the limiting function,

$$G(x) = \begin{cases} 0, & x < 0, \\ \frac{1}{3}, & 0 \leq x < \frac{1}{2}, \\ \frac{2}{3}, & x \geq \frac{1}{2}. \end{cases}$$

This function, albeit a limit of distributions functions, is *not* itself a distribution function. The problem is that $G(\infty) = 2/3$, meaning that one third of the probability mass of $F_n(x)$ has escaped to inifinity. Tightness is a condition ensuring that a limit of distribution functions is itself a distribution function. A sequence $X_1, X_2, \ldots$ of random variables is tightness tight if for any $\varepsilon > 0$ there exists a constant K so that

$$\Pr(|X_n| > K) < \varepsilon, \quad \text{for all } n.$$

A tight sequence of random variables is also said to be bounded in probability.

(a). Show that $(X_n)_{n\geq 1}$ is tight if and only if for any $\varepsilon > 0$ there are a and b such that

$$\Pr\{X_n \in (a, b]\} > 1 - \varepsilon, , \quad \text{for all } n,$$

and that this is the same as requiring that for any $\varepsilon > 0$ we can find a and b such that $F_n(a) < \varepsilon$ and $F_n(b) > 1 - \varepsilon$ for all n, in terms of the c.d.f.s. F_n of X_n.

(b). Show that $(X_n)_{n\geq 1}$ is tight if and only if for any $\varepsilon > 0$ there is a $K > 0$ so that $\limsup_n \Pr(|X_n| > K) < \varepsilon$.

(c). Show that $(X_n)_{n\geq 1}$ is tight if and only if for any sequence of constants $a_1, a_2, \ldots \geq 0$ tending to zero, $a_n Y_n \to_p 0$.

(d). Suppose that for some $\delta > 0$ there is an $M > 0$ and an $n_0 \geq 1$ such that $\mathrm{E}\,|X_n|^\delta \leq M$ for all $n \geq n_0$. Show that $(X_n)_{n\geq 1}$ is tight.

(e). Show that if $(X_n)_{n\geq 1}$ is uniformly integrable, then it is tight.

(f). Let $(X_n)_{n\geq 1}$ be a tight sequence or random variables with c.d.f.s $(F_n)_{n\geq 1}$. Suppose that there is a non-decreasing and right-continuous function F such that $F_n(x) \to F(x)$ at each continuity point of F. Show that F must be a c.d.f., and hence $X_n \to_d X$, where the random variable X has distribution F.

Ex. 2.25. *Helly's theorem* ... says that for any sequence $(F_n)_{n\geq 1}$ of c.d.f.'s there is a non-decreasing and right-continuous function F and a subsequence $(F_{n_k})_{k\geq 1}$ such that $F_{n_k}(x) \to F(x)$ at each x where F is continuous. The limit function F need not be a c.d.f., as it may lack one of the three defining properties of a c.d.f. We know from Ex. 2.24(f), however, that if $(F_n)_{n\geq 1}$ is a tight sequence, then F is a bona fide c.d.f. In this exercise we prove Helly's theorem for c.d.f.'s, first on the real line, then on $\mathbb{R}^k$.

(a). Let $\mathbb{Q} = \{q_1, q_2, \ldots\}$ be the rational numbers, and consider the infinite array

$$\begin{array}{ccccc} F_1(q_1) & F_2(q_1) & F_3(q_1) & \ldots \\ F_1(q_2) & F_2(q_2) & F_3(q_2) & \ldots \\ F_1(q_3) & F_2(q_3) & F_3(q_3) & \ldots \\ \vdots & \vdots & \vdots & \end{array}$$

Since $F_n(q_1)$ lies between zero and one for all n, the the Bolzano–Weierstrass theorem says that there is a convergent subsequence $(F_{n_{1,k}}(q_1))_{k\geq 1}$, say $F_{n_{1,k}}(q_1) \to G(q_1)$ as $k \to \infty$. The sequence $(F_{n_{1,k}}(q_2))_{k\geq 1}$ is also bounded between zero and one, so we can extract a further subsequence $(n_{2,k})_{k\geq 1} \subset (n_{1,k})_{k\geq 1}$ such that $F_{n_{2,k}}(q_2)$ converges to a limit $G(q_2)$ as $k \to \infty$. Continuing in this manner we get the array

the diagonal method

$$\begin{array}{ccccc} n_{1,1} & n_{1,2} & n_{1,3} & \ldots \\ n_{2,1} & n_{2,2} & n_{2,3} & \ldots \\ n_{3,1} & n_{3,2} & n_{3,3} & \ldots \\ \vdots & \vdots & \vdots & \end{array}$$

where each row is included in the row above, i.e., $(n_{j,k})_{k\geq 1} \subset (n_{j-1,k})_{k\geq 1}$ for $j = 2, 3, \ldots$. Argue that the diagonal sequence $n_k = n_{k,k}$ is such that $F_{n_k}(q) \to G(q)$ for each $q \in \mathbb{Q}$.

(b). Define the function

$$F(x) = \inf\{G(q) \colon q > x, q \in \mathbb{Q}\}.$$

Since $F(x)$ is an infinum over an set that shrinks with growing x, $F(x)$ is non-decreasing. Also, as $F(x)$ is an infinum we can for any x and any $\varepsilon > 0$ find an $s \in \mathbb{Q}$ with $s > x$ such that $G(s) < F(x) + \varepsilon$. Use this to prove that $F(x)$ must be right-continuous.

(c). Let x be a continuity point of F (see Ex. A.15(c) for the why such points exist). Argue that for any $\varepsilon > 0$ we can find $y \in \mathbb{R}$ and $s, r \in \mathbb{Q}$ such that $y < r < x < s$ and

$$F(x) - \varepsilon \leq G(r) \leq G(s) \leq F(x) + \varepsilon.$$

Use these inequalities to prove that $\lim_k F_{n_k}(x) = F(x)$, proving Helly's theorem.

(d). Let $\mathbb{Q}^k = \{q_1, q_2, \ldots\}$ be an enumeration of the k-dimensional vectors with rational components. Mimic the one-dimensional proof and generalise Helly's theorem to sequences of c.d.f.'s on $\mathbb{R}^k$. The changes you need to make are essentially notational.

Ex. 2.26. *Prokhorov's theorem...* comprises (a) and (b) of this exercise. With Helly's theorem at hand, part (a) was essentially proved in Ex. 2.24(f), but we include it here for completeness.

(a). If $(X_n)_{n\geq 1}$ is tight, then there is a subsequence $(n_k)_{k\geq 1}$ such that $X_{n_k} \to_d X$ for some random variable X.

(b). If $X_n \to_d X$ for some random variable X, then $(X_n)_{n\geq 1}$ is tight.

(c). Suppose that $(X_n)_{n\geq 1}$ is tight, and that every subsequence of X_{n_k} that converges weakly at all, converges to the same random variable X. Show that then $X_n \to_d X$.

Ex. 2.27. *Characteristic functions converging to a characteristic function.* Let $(X_n)_{n\geq 1}$ be a sequence of random variables with characteristic functions $\varphi_1(t), \varphi_2(t), \ldots$. Suppose that $\varphi_n(t) \to \varphi(t)$, and that we recognise $\varphi(t)$ as the characteristic function of some random variable X. In such cases we can conclude that $X_n \to_d X$ without further ado.

(a). Suppose that X_n and X have characteristic functions $\varphi_n(t)$ and $\varphi(t)$, respectively, and that $\varphi_n(t) \to \varphi(t)$, where $\varphi(t)$ is a characteristic function. As an auxiliary assumption, suppose that $\varphi_n(t)$ is dominated by a function $g(t)$ such that $\int g(t)\,\mathrm{d}t < \infty$. By Ex. A.40(h) the X_n then have densities

$$f_n(x) = \frac{1}{2\pi}\int \exp(-itx)\varphi_n(t)\,\mathrm{d}t.$$

Show that under these assumptions $X_n \to_d X$.

(b). Next, we reduce the general case to the special case studied in (a), by using the result from Ex. 2.19(c). Suppose that X_n and X have characteristic functions $\varphi_n(t)$ and $\varphi(t)$, respectively, and that $\varphi_n(t) \to \varphi(t)$. Let Z be a standard normal random variable, and σ some positive constant. Show that $X_n + \sigma Z \to_d X + \sigma Z$, and conclude that $\varphi_n(t) \to \varphi(t)$ implies $X_n \to_d X$.

Ex. 2.28. *Lévy's continuity theorem and some more.* In Ex. 2.27 we made two crucial assumptions: that $\varphi_n(t) \to \varphi(t)$, *and* that the limiting function φ is indeed a characteristic function. What if we drop the second assumption and only assume that the characteristic functions φ_n converge to *some* function φ? Lévy's continuity theorem says that if this limiting function is continuous at zero, then it must be a characteristic function. In more terse prose, if $\varphi_n(t) = \mathrm{E}\exp(itX_n)$ and $\varphi_n(t) \to \varphi(t)$ with φ some function continous at zero, then φ is the characteristic function of some random variable X, and $X_n \to_d X$. The key to this theorem is that the continuity of φ at zero entails that $(X_n)_{n\geq 1}$ is tight. In fact, given that $\varphi_n(t) \to \varphi(t)$ for some function $\varphi(t)$, there is equivalence between the following three statements:

(i) $\varphi(t)$ is a characteristic function,

(ii) the X_n sequence is tight;

(iii) $\varphi(t)$ is continuous at zero;

and, as we know from Ex. 2.27, the convergence $\varphi_n(t) \to \varphi(t)$ combined with (iii) implies that convergence in distribution takes place.

The implication (i)⇒(iii) is immediate from the uniform continuity of characteristic functions, see Ex. A.39(a); and (i)⇒(ii) follows from combining Ex. 2.26(b) and Ex. 2.27. In (a) of this exercise we prove (ii)⇒(i); while (iii)⇒(i), which is Lévy's continuity theorem, is proven in (b)–(c).

(a). Suppose that X_n has characteristic function φ_n for $n = 1, 2, \ldots$, and that $(X_n)_{n\geq 1}$ is tight. Show that if $\varphi_n(t) \to \varphi(t)$ for *some* function φ, then $\varphi(t)$ is a characteristic function.

(b). Start by applying Fubini's theorem (see Ex. A.13) to show that if X has characteristic function φ and cumulative distribution function F, then

$$\int_{-\varepsilon}^{\varepsilon} \{1 - \varphi(t)\}\,\mathrm{d}t = 2\varepsilon \int \big(1 - \frac{\sin(x\varepsilon)}{x\varepsilon}\big)\,\mathrm{d}F(x).$$

In particular, the integral of $\varphi(t)$ on a symmetric interval around zero is really a real number, that is, the complex component disappears. Deduce that

$$\frac{1}{\varepsilon}\int_{-\varepsilon}^{\varepsilon} \{1 - \varphi(t)\}\,\mathrm{d}t \geq 2\int_{|x\varepsilon|\geq c} \big(1 - \frac{\sin(x\varepsilon)}{x\varepsilon}\big)\,\mathrm{d}F(x) \geq 2(1 - 1/c)\Pr\{|X| \geq c/\varepsilon\},$$

with the value $c = 2$ yielding the tail inequality

$$\Pr\{|X| \geq 2/\varepsilon\} \leq \frac{1}{\varepsilon}\int_{-\varepsilon}^{\varepsilon} \{1 - \varphi(t)\}\,\mathrm{d}t. \tag{2.8}$$

(c). If we now have a collection of random variables, where their characteristic functions have approximately the same level of smoothness around zero, then we should get tightness from (2.8). Assume that $X_1, X_2, \ldots$ have characteristic functions $\varphi_1, \varphi_2, \ldots$, and that $\varphi_n(t)$ converges to *some* function $\varphi(t)$ that is continuous at zero. For a given $\varepsilon' > 0$, find $\varepsilon > 0$ such that $|1 - \varphi(t)| \leq \varepsilon'$ for $|t| \leq \varepsilon$. Show that

$$\limsup_{n\to\infty} \Pr\{|X_n| \geq 2/\varepsilon\} \leq \frac{1}{\varepsilon}\int_{-\varepsilon}^{\varepsilon} \{1 - \varphi(t)\}\,\mathrm{d}t \leq 2\varepsilon'.$$

We've hence found a broad interval, namely $[-2/\varepsilon, 2/\varepsilon]$, inside which each X_n lies, with high enough probability, which means that $(X_n)_{n\geq 1}$ is tight, by Ex. 2.24(b).

(d). As a bonus, show that if $X_1, X_2, \ldots$ is a tight sequence of random variables with characteristic functions $\varphi_1, \varphi_2, \ldots$, then $(\varphi_n)_{n\geq 1}$ is uniformly equicontinuous. That is, for each $\varepsilon > 0$ there is a $\delta > 0$ such that $|t - s| < \delta$ implies $|\varphi_n(t) - \varphi_n(s)| < \varepsilon$ for all n.

Ex. 2.29. *Showing convergence in distribution via moment generating functions.* Here we spell out an m.g.f. analogue of the Lévy continuity theorem, see Ex. 2.28. Let $(X_n)_{n\geq 1}$ have m.g.f.s $\{M_n(t)\}_{n\geq 1}$ and suppose that all these exist on the same interval $[-a, a]$.

(a). Expand on Ex. 2.16(a) to show that

$$\Pr(|X_n| \geq K) \leq \exp(-Ka)\{M_n(a) + M_n(-a)\}, \quad \text{for all } n.$$

Show that if $M_n(t)$ converges to *some* function $M(t)$ for all t in some interval around zero that contains $[-a, a]$, then $(X_n)_{n\geq 1}$ is tight. Deduce that $X_n \to_d X$, where X is a random variable with m.g.f. $M(t)$. Note that since $x \mapsto \exp(tx)$ is not bounded, this last part is more involved than in the characteristic function case, cf. Ex. 2.28(a).

(b). For illustration, consider $X_1, \dots, X_n$ being independent 0-1 Bernoulli variables with $p_i = \Pr(X_i = 1)$. Show that if $\max_{i \le n} p_i \to 0$ and $\sum_{i=1}^n p_i \to \lambda$, then $\sum_{i=1}^n X_i \to_d \text{Pois}(\lambda)$. Note that this generalises the classical binomial to Poisson theorem, see Ex. 2.8, and also that proving this result directly from the point probabilities would be difficult.

(c). Suppose $X_1, X_2, \dots$ are independent Bernoulli with $p_i = i/n$, and consider $Y_n(t) = \sum_{i \le t\sqrt{n}} X_i$. Show that $Y_n(t) \to_d \text{Pois}(\frac{1}{2}t^2)$. The limit is actually a full Poisson process in t, with independent increments (see Ch. 9).

Ex. 2.30. *Showing convergence in two steps.* Suppose one wishes to prove that $X_n \to_d X$, but that technical issues make it easier to first prove that an approximation to X_n converges to an approximation to X. With a suitable extra condition this might suffice. Here we suppose that the random elements involved take values in a common (separable) metric space $(\mathcal{X}, d)$.

(a). For the approximations $A_{n,k}$ to X_n and A_k to X, suppose (i) that $A_{n,k} \to_d A_k$, for each k; (ii) that $A_k \to_d X$ as $k \to \infty$; and also (iii) that

$$\lim_{k\to\infty} \limsup_{n\to\infty} \Pr(d(X_n, A_{n,k}) \ge \varepsilon) = 0 \quad \text{for each } \varepsilon > 0.$$

Consider the closed sets F and $F^\varepsilon = \{x\colon d(x, F) \le \varepsilon\}$, and use the Portmanteau theorem to show that $X_n \to_d X$.

(b). Suppose somebody clever has managed to prove the CLT for bounded i.i.d. variables (for example using m.g.f.s): If $A_1, A_2, \dots$ are i.i.d. with mean zero and variance one, and $|A_i| \le k$, then $Z_n = \sqrt{n}\bar{A}_n \to_d \text{N}(0,1)$. How can you then prove that the CLT is valid also for unbounded i.i.d. random variables, as long at only their variance is finite?

Central limit theorems

Ex. 2.31. *The central limit theorem, Lindeberg's proof.* Let $X_1, X_2, \dots$ be i.i.d. random variables with mean zero and unit variance, define $X_{n,i} = X_i/\sqrt{n}$, and let $Z \sim \text{N}(0,1)$. We will show that $\text{E}\, f(\sum_{i=1}^n X_{n,i}) \to \text{E}\, f(Z)$ for all bounded and continuous functions with bounded and continuous derivatives of all orders, which, by Ex. 2.19 is equivalent to $\sum_{i=1}^n X_{n,i} \to_d \text{N}(0,1)$. Introduce $Z_{n,i} = Z_i/\sqrt{n}$ for $i = 1, 2, \dots, n$ and $n \ge 1$, where the $Z_1, Z_2, \dots$ are i.i.d. standard normal random variables independent of the $X_1, X_2, \dots$. In particular, $\sum_{i=1}^n Z_{n,i} \sim \text{N}(0,1)$ by Ex. A.40, for example.

(a). What is in essence Lindeberg's idea, back in 1922, was to show that $\text{E}\, f(\sum_{i=1}^n X_{n,i}) - \text{E}\, f(\sum_{i=1}^n Z_{n,i})$ tends to zero by replacing the summands $X_{n,i}$ by the Gaussiand summands $Z_{n,i}$, one by one. Show that by so doing the difference $f(\sum_{i=1}^n X_{n,i}) - f(\sum_{i=1}^n Z_{n,i})$ can be expressed as the (telescoping) sum

$$\begin{aligned} f(\sum_{i=1}^n X_{n,i}) - f(\sum_{i=1}^n Z_{n,i}) &= \sum_{k=1}^n \{f(A_k) - f(A_{k-1})\} \\ &= \sum_{k=1}^n \{f(R_k + X_{n,k}) - f(R_k + Z_{n,k})\}, \end{aligned}$$

where $A_k = \sum_{i=1}^k X_{n,i} + \sum_{i=k+1}^n Z_{n,i}$ and $R_k = \sum_{i=1}^{k-1} X_{n,i} + \sum_{i=k+1}^n Z_{n,i}$, so that $A_n = \sum_{i=1}^n X_{n,i}$ and $A_0 = \sum_{i=1}^n Z_{n,i}$.

(b). Let f be a bounded and continuous functions with bounded and continuous derivatives up to and including the third order. Show that for any $\varepsilon > 0$ there are $\delta > 0$ and K such that

$$|f(x+y) - f(x) - f'(x)y - \tfrac{1}{2}f''(x)y^2| \le \varepsilon y^2 I(|y| < \delta) + Ky^2 I(|y| \ge \delta),$$

uniformly in x and y.

(c). Verify that for each $k = 1, \ldots, n$ and $n \ge 1$, $X_{n,k}$ is independent of R_k and that $Z_{n,k}$ is independent of R_k. Use this and the findings from (a) and (b) to show that given $\varepsilon > 0$ we can choose $\delta > 0$ and find K such that

$$|\mathrm{E}\, f(\sum_{i=1}^n X_{n,i}) - \mathrm{E}\, f(\sum_{i=1}^n Z_{n,i})| \le 2\varepsilon + K\sum_{k=1}^n \mathrm{E}\, X_{n,k}^2 I_{|X_{n,i}|\ge\delta} + K\sum_{k=1}^n \mathrm{E}\, Z_{n,k}^2 I_{|Z_{n,i}|\ge\delta}.$$

Please conclude, and you will have shown the CLT for i.i.d. random variables.

(d). Suppose $Y_1, Y_2, \ldots$ are i.i.d. with finite mean ξ and standard deviation σ. Show that the random sum, normalised to have mean zero and variance one, i.e.,

$$\Big(\sum_{i=1}^n Y_i - n\xi\Big)/\sqrt{n\sigma^2} = \sqrt{n}(\bar{Y}_n - \xi)/\sigma,$$

tends to the $\mathrm{N}(0,1)$ in distribution.

Ex. 2.32. *The Lindeberg theorem.* We now extend the central limit theorem to independent random variables that do not necessarily have the same distribution. Lindeberg CLT

(a). Let $X_1, X_2, \ldots$ be independent mean zero random variables with variances $\sigma_1^2, \sigma_2^2 \ldots$, and form $X_{n,j} = X_j/B_n$, where $B_n^2 = \sigma_1^2 + \cdots + \sigma_n^2$. In view of the remainder term in Ex. 2.31(c), the Lindeberg condition is natural: Assume that

$$\sum_{j=1}^n \mathrm{E}\, X_{n,j}^2 I(|X_{n,j}| \ge \varepsilon) \to 0, \quad \text{as } n \to \infty, \text{ for each } \varepsilon > 0. \tag{2.9}$$

Show that the Lindeberg condition implies that $\max_{i\le n} \sigma_i^2/B_n^2 \to 0$ as $n \to \infty$; and proceed to show that $\sum_{j=1}^n X_{n,j} \to_d \mathrm{N}(0,1)$ by making the appropriate modifications to the proof for the i.i.d. case studied in Ex. 2.31.

(b). Suppose $X_1, X_2, \ldots$ are i.i.d. with mean zero and variance σ^2, then form from these $Z_n = \sum_{i=1}^n iX_i/(\sum_{i=1}^n i^2)^{1/2}$. Show that $Z_n \to_d \mathrm{N}(0, \sigma^2)$. Generalise.

Ex. 2.33. *Convergence in distribution for bounded variables.* Suppose variables X_n all live on a bounded interval, which we for simplicity and without loss of generality below take to be $[0,1]$.

(a). With X a candidate limit variable, also on $[0,1]$, we know from the Portmanteau theorem of Ex. 2.18 that $X_n \to_d X$ if and only if $\mathrm{E}\, g(X_n) \to \mathrm{E}\, g(X)$ for each bounded continuous g. Show via the Weierstraß polynomial approximation theorem of Ex. 2.17 that this is also equivalent to $\mathrm{E}\, X_n^k \to \mathrm{E}\, X^k$ for each k.

(b). Then show that this again is equivalent to convergence of the m.g.f.s, i.e. $M_n(t) \to M(t)$ for each t. For an illustration with an infinite sum of Bernoullis, check Story vi.2.

(c). We saw in Ex. 2.7 that if X_n is uniform on $1/n, 2/n, \ldots, n/n$, then $X_n \to_d U$, the uniform on the unit interval. Show from this that $s_n(k) = (1/n^{k+1})(1^k + 2^k + \ldots + n^k) \to 1/(k+1)$ for each k.

Ex. 2.34. *Proving the CLT with moment generating functions.* Here we prove a central limit theorem for i.i.d. random variables $X_1, X_2, \ldots$ whose m.g.f. $M(t) = \mathrm{E}\exp(tX_1)$ exist on an interval around zero. See Ex. A.36 for what this assumption entails about the random variables. Throughout, we assume that $\mathrm{E}\,X_1 = 0$ and $\mathrm{Var}\,X_1 = \sigma^2$. The CLT proof via characteristic functions is clearly more general, but in various cases it is simply more convenient to give proofs via m.g.f.s.

(a). Show that $M(t) = 1 + \frac{1}{2}\sigma^2 t^2 + o(t^2)$ as $t \to 0$.

(b). Show that $\sqrt{n}\bar{X}_n = n^{-1/2}\sum_{i=1}^n X_i$ has moment generating function of the form

$$M_n(t) = M(t/\sqrt{n})^n = \{1 + \tfrac{1}{2}\sigma^2 t^2/n + o(t^2/n)\}^n,$$

and conclude that $M_n(t) \to \exp(\frac{1}{2}t^2\sigma^2)$ as $n \to \infty$. Explain, in view of Ex. 2.29, that this proves the CLT, under these m.g.f. conditions.

the Stirling approximation

Ex. 2.35. *Proving the Stirling formula via the CLT.* The approximation formula

$$n! \doteq n^n e^{-n}\sqrt{2\pi n}, \quad \text{in the sense of} \quad \lim_{n\to\infty} \frac{n!}{n^n \exp(-n)(2\pi n)^{1/2}} = 1,$$

named after J. Stirling (1692–1770), though A. de Moivre stated the formula earlier

is a famous one. Here we shall prove this formula via the CLT, specifically for Poisson variables and gamma variables.

(a). If $X_n \sim \mathrm{Pois}(n)$, show that $Z_n = (X_n - n)/\sqrt{n} \to_d Z$, a standard normal. Show that $\exp(-n)(1 + n/1 + n^2/2! + \cdots + n^n/n!) \to \frac{1}{2}$.

(b). Write $p_n(j) = \Pr(X_n = j)$ for the Poisson probabilities. Work out that $\mathrm{E}\max(X_n - n, 0) = \sum_{j\geq n}(j-n)p_n(j) = np_n(n)$, perhaps by discovering a telescoping series. Show from this that

$$\mathrm{E}\max(0, Z_n) = \sum_{j\geq n} \frac{j-n}{\sqrt{n}} \exp(-n)\frac{n^j}{j!} = \sqrt{n}\exp(-n)n^n/n! \to \mathrm{E}\max(0, Z),$$

and that the right hand side is $1/(2\pi)^{1/2}$. Deduce Stirling from this.

(c). Whereas the arguments above lead to a clean proof of Stirling, attempt also the following. With ε small, argue that

$$\sum_{n\leq j\leq n+\varepsilon\sqrt{n}} \exp(-n)\frac{n^j}{j!} \doteq \frac{1}{(2\pi)^{1/2}}\varepsilon.$$

Prove Stirling from this.

(d). To learn more, from the arguments above, show that

$$a_n(c) = \mathrm{E}\,Z_n\, I(0 \leq Z_n \leq c) = \sum_{n\leq j\leq n+c\sqrt{n}} \frac{j-n}{\sqrt{n}} \exp(-n)\frac{n^j}{j!} \to a(c),$$

where $a(c) = \mathrm{E}\, Z\, I(0 \le Z \le c)$; the previous case corresponds to $c = \infty$. For the limit, show that $a(c) = \phi(0) - \phi(c) = (2\pi)^{-1/2}\{1 - \exp(-\frac{1}{2}c^2)\}$. For the left hand side, write $[x]$ for the integer value of x, and use the telescoping series trick to show that

$$a_n(c) = \frac{1}{\sqrt{n}}\{np_n(n) - np_n([n + c\sqrt{n}])\} = \sqrt{n}\exp(-n)\frac{n^n}{n!}\Big\{1 - \frac{n^{[c\sqrt{n}]}}{(n+1)\cdots([n+c\sqrt{n}])}\Big\}.$$

Deduce from this that

$$\frac{n^{[c\sqrt{n}]}}{(n+1)\cdots([n+c\sqrt{n}])} \to \exp(-\tfrac{1}{2}c^2),$$

for any $c > 0$. You may also try to show this directly, but here we get it from the Poisson and probability theory.

(e). Consider now gamma distributed variables $X_n \sim \mathrm{Gam}(n, 1)$, with densities $g_n(x) = \Gamma(n)^{-1}x^{n-1}\exp(-x)$ for x positive. Explain via the CLT that $Z_n = (X_n - n)/\sqrt{n} \to_d Z$, a standard normal. Show from this that we for each (a, b) have

$$L_n(a, b) = \mathrm{E}\, Z_n\, I(Z_n \in [a, b]) = \int_{n+a\sqrt{n}}^{n+b\sqrt{n}} \frac{x - n}{\sqrt{n}} g_n(x)\,\mathrm{d}x \to L(a, b),$$

where $L(a, b) = \int_a^b z\phi(z)\,\mathrm{d}z = \phi(a) - \phi(b)$.

(f). Show that $xg_n(x) = ng_{n+1}(x)$, and that this leads to

$$L_n(a, b) = \sqrt{n}\int_{n+a\sqrt{n}}^{n+b\sqrt{n}} (g_{n+1} - g_n)\,\mathrm{d}x.$$

With G_n the c.d.f. for g_n, show via expressions found in Ex. 1.9 that $G_n(t) - G_{n+1}(t) = \exp(-t)t^n/n!$. From this deduce that

$$L_n(0, \infty) = \sqrt{n}\int_n^\infty (g_{n+1} - g_n)\,\mathrm{d}x = \sqrt{n}\exp(-n)n^n/n! \to 1/(2\pi)^{1/2},$$

proving the Stirling again. Working with $L_n(0, b)$, deduce also that

$$\exp(-b\sqrt{n})(1 + b/\sqrt{n})^n \to \exp(-\tfrac{1}{2}b^2).$$

This may be verified independently; the present point is that these formulae, including the Stirling, follow from the CLT and probability theory for sums of gammas.

Ex. 2.36. *Two useful lemmas.* In proving the central limit theorems below, the following two lemmas will be useful.

(a). Demonstrate that with $z_{n,1}, z_{n,2}, \ldots$ a sequence of real numbers coming closer to zero, we have $\prod_{i=1}^n (1 + z_{n,i}) \to \exp(z)$, provided (i) $\sum_{i=1}^n z_{n,i} \to z$; (ii) $\max_{i \le n} |z_{n,i}| \to 0$; and (iii) $\sum_{i=1}^n |z_{n,i}|$ stays bounded. It may be helpful to show first that

$$\log(1 + x) = x - \tfrac{1}{2}x^2 + \tfrac{1}{3}x^3 - \cdots = x + x^2K(x),$$

with $K(x)$ a continuous function such that $|K(x)| \le 1$ for $|x| \le \frac{1}{2}$, and $K(x) \to -\frac{1}{2}$ when $x \to 0$. These statements are valid also when the $z_{n,i}$ and the x are complex numbers inside the unit ball, in which case the logarithm is the natural complex extension of the real logarithm.

(b). Let $x_1, \ldots, x_n$ and $y_1, \ldots, y_n$ be complex numbers such that $|x_j|, |y_j| \leq 1$ for all j. Show that $|x_1x_2 - y_1y_2| \leq |x_1 - y_1| + |x_2 - y_2|$, and proceed by induction to show that

$$|\prod_{j=1}^{n} x_j - \prod_{j=1}^{n} y_j| \leq \sum_{j=1}^{n} |x_j - y_j|, \quad \text{for all } n.$$

Ex. 2.37. *The Lindeberg CLT via truncations.* Here we present a proof of the CLT for independent random variables using m.g.f.s., but without assuming that the m.g.f.s. of the random variables involved necessarily exist. Since the m.g.f. of a bounded random variable always exists (prove it), the trick is to truncate the random variables involved: Let $X_1, X_2, \ldots$ be independent random variables with mean zero and variances $\sigma_1^2, \sigma_2^2, \ldots$. Write $B_n^2 = \sigma_1^2 + \cdots + \sigma_n^2$ and $X_{n,j} = X_j/B_n$ for $j = 1, \ldots, n$. We are to show that $\sum_{j=1}^n X_{n,j}$ converges in distribution to a standard normal random variable, provided the Lindeberg condition (2.9) is satisfied, which we assume throughout.

(a). For some $\varepsilon > 0$, define $\mu_{n,j} = \mathrm{E}\, X_{n,j} I(|X_{n,j}| < \varepsilon)$, as well as the random variables $Y_{n,j} = X_{n,j} I(|X_{n,j}| < \varepsilon) - \mu_{n,j}$ and $v_{n,j} = X_{n,j} I(|X_{n,j}| \geq \varepsilon) + \mu_{n,j}$. Verify that

$$X_{n,j} = Y_{n,j} + v_{n,j},$$

and also $|Y_{n,j}| \leq 2\varepsilon$, the latter implying that the m.g.f. of $Y_{n,j}$ exist for all j and n. Show that

$$\mathrm{E}\, \exp(tY_{n,j}) = 1 + \tfrac{1}{2}t^2 \mathrm{E}\, Y_{n,j}^2 + r_{n,j}(t),$$

where the remainders are such that $|\sum_{j=1}^n r_{n,j}(t)| \leq \varepsilon \exp(2\varepsilon)$ for $|t| \leq 1$.

(b). By Ex. 2.36(a), showing that $\prod_{j=1}^n \mathrm{E}\, \exp(tY_{n,j}) \to \exp(\frac{1}{2}t^2)$ now comes down to verifying that (i) $\sum_{j=1}^n \mathrm{E}\, Y_{n,j}^2 \to 1$; (ii) $\max_{j \leq n} \mathrm{E}\, Y_{n,j}^2 \to 0$ and $\max_{j \leq n} |r_{n,j}| \to 0$; and (iii) $\limsup_{n \to \infty} |\mathrm{E}\, \exp(tY_{n,j}) - 1|$ is finite. Please verify these conditions, remembering that it is sufficient to consider $|t| \leq 1$.

(c). Finally, show that $\sum_{j=1}^n v_{n,j} \to_{\mathrm{pr}} 0$ and conclude that $\sum_{j=1}^n X_{n,j} \to_d \mathrm{N}(0,1)$.

Ex. 2.38. *Proving the CLT with characteristic functions.* It can be argued that the most elegant, unified, and general proof of the CLT is obtained by the use of characteristic functions. This is because, contrary to the m.g.f., the characteristic function of a random variable always exists.

(a). Show that if X has finite mean ξ, then its characteristic function satisfies $\varphi(t) = 1 + i\xi t + o(t)$ as $t \to 0$. Also, its derivative exists, with $\varphi'(t) = \mathrm{E}\, iX \exp(itX)$, and in particular $\varphi'(0) = i\xi$ (see Ex. A.39 for more details).

(b). Show similarly that if X has finite variance σ^2, then

$$\varphi(t) = 1 + i\xi t - \tfrac{1}{2}(\xi^2 + \sigma^2)t^2 + o(t^2) \quad \text{as } t \to 0.$$

(c). If $X_1, X_2, \ldots$ are i.i.d. with mean zero and finite variance σ^2, then show that $Z_n = \sqrt{n}\bar{X}_n = n^{-1/2} \sum_{i=1}^n X_i$ has characteristic function of the form

$$\varphi_n(t) = \{1 - \tfrac{1}{2}\sigma^2 t^2/n + o(1/n)\}^n.$$

Prove the CLT from this.

Ex. 2.39. *The Lindeberg CLT, yet again.* Let $(X_{n,i})_{i\le n,n\ge 1}$ be a triangular array of mean zero random variables with variances $(\sigma^2_{n,i})_{i\le n,n\ge 1}$ such that $X_{n,1},\dots,X_{n,n}$ are independent for each n. We say that the triangular array is row wise independent.

(a). Let $(Y_{n,i})_{i\le n,n\ge 1}$ be a triangular array of row wise independent mean zero normally distributed random variables with variances $\mathrm{Var}(Y_{n,i}) = \sigma^2_{n,i}$ for all i and n. Using the second lemma of Ex. 2.36, show that for any $\varepsilon > 0$

$$|\mathrm{E}\,\exp(it\sum_{i=1}^n X_{n,i}) - \mathrm{E}\,\exp(it\sum_{i=1}^n Y_{n,i})| \le \tfrac{1}{3}|t|^3\varepsilon\sum_{i=1}^n \sigma^2_{n,i} + L_n(\varepsilon) + L'_n(\varepsilon),$$

where $L_n(\varepsilon) = \sum_{i=1}^n \mathrm{E}\,X^2_{n,i}I(|X_{n,i}|\ge\varepsilon)$ and $L'_n(\varepsilon) = \sum_{i=1}^n \mathrm{E}\,Y^2_{n,i}I(|Y_{n,i}|\ge\varepsilon)$.

(b). Deduce that $\sum_{i=1}^n X_{n,i} \to_d \mathrm{N}(0,1)$ provided $\sum_{i=1}^n \sigma^2_{n,i} \to 1$ and the Lindeberg condition holds.

Ex. 2.40. *Lindeberg and Liapunov and some more.* In this exercise we summarise various implications pertaining to the Lindeberg condition. Let $X_1, X_2, \dots$ be independent mean zero random variables with variances $\sigma^2_1, \sigma^2_2, \dots$, set $B^2_n = \sigma^2_1 + \cdots + \sigma^2_n$, and define $X_{n,i} = X_i/B_n$. Recall that the Lindeberg condition is satisfied if for each $\varepsilon > 0$, we have $\sum_{i=1}^n \mathrm{E}\,X^2_{n,i}I(|X_{n,i}|\ge\varepsilon) \to 0$ as n tends to infinity.

(a). Show that if the Lindeberg conditon holds, then $\max_{i\le n}\sigma^2_i/B^2_n \to 0$.

(b). Show that if B^2_n is convergent, then the Lindeberg condition fails.

(c). In view of (b), we assume that B_n diverges. Show that the Lindeberg condition holds for uniformly bounded random variables, i.e., $|X_i| \le M$ for all i.

(d). Show that if $\mathrm{E}\,|X_i|^{2+\delta} < M$ for all i and some $\delta > 0$, still assuming that B_n diverges, of course.then the Lindeberg condition holds.

(e). Suppose that there is a sequence of constants $K_1, K_2, \dots$ such that $|X_i| < K_i$ for each i, and that $K_n/B_n \to 0$. Show that the Lindeberg conditon holds.

(f). Show that the Liapunov condition, that is $\sum_{i=1}^n \mathrm{E}\,|X_{n,i}|^{2+\delta} \to 0$ for some $\delta > 0$, implies the Lindeberg condition.

(g). Suppose that $n^{-1}B^2_n \to \sigma^2 > 0$, and show that then

$$\frac{1}{n}\sum_{i=1}^n \mathrm{E}\,X^2_i I(|X_i| \ge \sqrt{n}\varepsilon) \to 0,$$

for each $\varepsilon > 0$ if and only if the Lindeberg condition holds; and that under these conditions $\sqrt{n}\bar{X}_n \to_d \mathrm{N}(0,\sigma^2)$, where $\bar{X}_n = n^{-1}\sum_{i=1}^n X_i$.

(h). Let $Y_1, Y_2, \dots$ be i.i.d. random variables with finite second moments. Show that $\max(Y_1,\dots,Y_n)/\sqrt{n}$ tends to zero in probability.

(i). Show that if $Y_1, Y_2, \dots$ are i.i.d. with mean μ and variance σ^2, then the Lindeberg condition holds. Thus, the Lindeberg CLT contains the CLT for i.i.d. random variables.

Ex. 2.41. *Limiting normality of linear combinations of i.i.d. variables.* Let $\varepsilon_1, \varepsilon_2, \ldots$ be i.i.d. from some distribution with mean zero and finite variance σ^2. For a sequence of multiplicative constants $a_1, a_2, \ldots$, consider

$$Z_n = \frac{\sum_{i=1}^n a_i \varepsilon_i}{B_n} = \sum_{i=1}^n (a_i/B_n)\varepsilon_i, \quad \text{with } B_n^2 = \sigma^2 \sum_{i=1}^n a_i^2,$$

which has mean zero and variance 1. The question is what should be demanded of the a_i sequence, to ensure that $Z_n \to_d \mathrm{N}(0,1)$ (even if the ε_i distribution might be looking say skewed and multimodal and strange).

(a). Let $D_n = \max_{i \le n} |a_i|/B_n$. Writing G for the distribution of ε_i, show that

$$\sum_{i=1}^n \mathrm{E}\Big|\frac{a_i\varepsilon_i}{B_n}\Big|^2 I\Big(\Big|\frac{a_i\varepsilon_i}{B_n}\Big| \ge \delta\Big) \le \sum_{i=1}^n \frac{a_i^2}{B_n^2}\mathrm{E}\,\varepsilon_i^2\, I(D_n|\varepsilon_i| \ge \delta) \le \int_{|u|\ge \delta/D_n} u^2\,\mathrm{d}G(u),$$

and conclude that $Z_n \to_d \mathrm{N}(0,1)$ provided $D_n \to 0$.

(b). Under a variety of setups, one actually has $D_n \to 0$, which is hence not at all a strict condition. Verify that the condition holds, and hence limiting normality, in the following cases: (i) $a_i = 1$ (which corresponds to the plain CLT); (ii) all $|a_i|$ inside some positive $[b,c]$ interval; (iii) $a_i = i$; (iv) $a_i = i^2$ (and generalise); (v) $a_i = 1/\sqrt{i}$. Show however that the condition does not hold for $a_i = 1/i$.

(c). A grasshopper sits at zero and then starts jumping, to the right and left with equal probability, and with jump sizes $1, 2, 3, \ldots$. With S_n her position after n jumps, show that $S_n/n^{3/2}$ has a normal limit. Though she keeps on passing zero, she will not be in that vicinity; show that $\Pr(|S_n| \le cn^{1.49}) \to 0$, for each c.

(d). [xx fikse xx]Another important case to understand well is when the a_i can be considered an i.i.d. sequence, drawn from their own distribution. Show that $D_n \to_{\mathrm{pr}} 0$ if the a_i distribution has finite variance. (xx nils thinks this is if and only if, actually. what happens with Z_n if the a_i are drawn from say the $1/|x|^2$ distribution, for $|x| \ge 1$? xx)

Ex. 2.42. *Higher-order expansions of m.g.f.s.* (xx to be polished. the aim is to give more info for CLT, of the type $|F_n(t) - \Phi(t)| \le c/\sqrt{n}$ for some c. xx) [xx intro text here, perhaps by Nils. (xx nils pushes an earlier thing from Ch2 to this place; then we edit and prune and clean. xx)]

(a). Consider a variable Y, with m.g.f. $M(t) = \mathrm{E}\exp(tY)$, assumed to be finite in at least a neighbourhood around zero. We have seen in Ex. A.36 that $\mathrm{E}\,Y^r = M^{(r)}(0)$. Write ξ and σ^2 for the mean and variance of Y. Show that $M(t) = 1 + \xi t + o(h)$, for $|t|$ small. Taking a Taylor expansion to the next step, show that $M(t) = 1 + \xi t + \frac{1}{2}(\xi^2 + \sigma^2)t^2 + o(t^2)$. Deduce also that $\log M(t) = \xi t + \frac{1}{2}\sigma^2 t^2 + o(t^2)$.

(b). We may also take the expansion to the third order, but it is simpler and more insightful to proceed from $Y = \xi + Y_0$, with Y_0 having mean zero. Show that

$$M(t) = \exp(t\xi)\,\mathrm{E}\exp(tY_0) = \exp(t\xi)\{1 + \tfrac{1}{2}\sigma^2 t^2 + \tfrac{1}{6}\gamma_3 t^3 + o(|t|^3)\},$$

where $\gamma_3 = \mathrm{E}\,(Y - \xi)^3$.

(c). Consider $Y_1, \ldots, Y_n$ i.i.d. from a distribution with mean zero and m.g.f. $M(t)$ being finite around zero. Show that $Z_n = \sqrt{n}\bar{Y}$ has

$$\begin{aligned} M_n(t) &= \mathrm{E}\,\exp(tZ_n) = M(t/\sqrt{n})^n \\ &= \{1 + \tfrac{1}{2}\sigma^2 t^2/n + \tfrac{1}{6}\gamma_3 t^3/n^{3/2} + o(|t|^3/n^{3/2})\}^n. \end{aligned}$$

Show from this that under the assumptions given, $\log M_n(t) = \frac{1}{2}\sigma^2 t^2 + \frac{1}{6}\gamma_3 t^3/\sqrt{n} + o(1/\sqrt{n})$. Explain why this is a proof of the CLT (via criteria given in Ex. **??**, with attention to certain further details in Ch 3 xx).

(d). (xx round off, point to CLT, identify remainder term with skewness. xx)

Joint convergence in distribution

Ex. 2.43. *Characterisitc functions of random vectors and Cramér–Wold.* With $X = (X_1, \ldots, X_k)^{\mathrm{t}}$ a random vector, in dimension p, we define its characteristic functions as

$$\varphi(t_1, \ldots, t_k) = \mathrm{E}\,\exp(it^{\mathrm{t}}X) = \mathrm{E}\,\exp\{i(t_1X_1 + \cdots + t_kX_k)\}$$

for $t = (t_1, \ldots, t_k)^{\mathrm{t}}$. Let $F(x) = F(x_1, \ldots, x_k)$ be the c.d.f. of X, and $\varphi(t) = \varphi(t_1, \ldots, t_k)$ its characteristic function. From Ex. A.41 we have the inversion formula

$$F(b) - F(a) = \lim_{\sigma \to 0} \frac{1}{(2\pi)^k} \int_{\mathbb{R}^k} \left(\prod_{i=1}^{k} \frac{\exp(-it_jb_j) - \exp(-it_ja_j)}{-it_j} \right) \varphi(t) \exp(-\tfrac{1}{2}t^{\mathrm{t}}t\sigma^2)\,\mathrm{d}t,$$

valid for all continuity points $a = (a_1, \ldots, a_k)$ and $b = (b_1, \ldots, b_k)$ of F, using the shorthand notation $\mathrm{d}t = \mathrm{d}t_1 \cdots \mathrm{d}t_k$. Ex. A.41 also contains other results pertaining to the c.f. $\varphi(t_1, \ldots, t_p)$ that you may need in this exercise.

(a). Let $X_n = (X_{n,1}, \ldots, X_{n,p})^{\mathrm{t}}$ be a sequence of random vectors with characteristic functions $\varphi_n(t_1, \ldots, t_p)$. Suppose that $\varphi_n(t_1, \ldots, t_p) \to \varphi(t_1, \ldots, t_p)$, where φ is the characteristic function of some random vector $X = (X_1, \ldots, X_p)^{\mathrm{t}}$. Mimic the steps taken in Ex. 2.27 to show that then $X_n \to_d X$.

(b). Show that $X_n \to_d X$ if and only if $a^{\mathrm{t}}X_n \to_d a^{\mathrm{t}}X$ for each $a \in \mathbb{R}^k$, i.e., if and only if all linear combinations converge. This is the Cramér–Wold theorem, also known as the Cramér–Wold device. Cramér–Wold device

(c). Show that a random pair (X_n, Y_n) converges in distribution to the binormal $\mathrm{N}_2(0, \Sigma)$, with Σ having diagonal elements 1, 1, and correlation ρ, if and only if $aX_n + bY_n \to_d \mathrm{N}(0, a^2 + b^2 + 2\rho ab)$ for each (a, b).

Ex. 2.44. *Tightness, Prokhorov, and the continuity theorem in* $\mathbb{R}^k$. A sequence of random vectors $X_n = (X_{n,1}, \ldots, X_{n,k})^{\mathrm{t}}$ is tight if for any $\varepsilon > 0$ there is a K such that

$$\Pr(\|X_n\| > K) < \varepsilon \quad \text{for all } n.$$

(a). Let F_n is the c.d.f. of X_n, and show that X_n being tight is equivalent to that for every $\varepsilon > 0$ there are $a = (a_1, \ldots, a_k)$ and $b = (b_1, \ldots, b_k)$ such that $F_n(b) - F_n(a) > 1 - \varepsilon$ for all n, where .

(b). Show that all the $X_{n,1}, \ldots, X_{n,k}$ are tight if and only if X_n is tight. Deduce that X_n is tight if and only if any linear combination $a^{\mathrm{t}}X_n = a_1X_{n,1} + \cdots + a_pX_{n,p}$ is tight.

(c). Show that if the sequence $(X_n)_{n\geq 1}$ of random vectors is tight, then the corresponding sequence of characteristic functions $\{\varphi_n(t_1,\ldots,t_p)\}_{n\geq 1}$ is uniformly equicontinuous.

Prokohorov's theorem in $\mathbb{R}^k$

(d). Prokohorov's theorem, proven for the one-dimensional case in Ex. 2.26, extends to the multi-dimensional case (and further): Demonstrate that (i) if $X_n \to_d X$, then X_n is tight; and (ii) that if $(X_n)_{n\geq 1}$ is tight, then there is a subsequence $(n_k)_{k\geq 1}$ such that $(X_{n_k})_{k\geq 1}$ converges in distribution. Recall that Helly's theorem for random vectors was proven in Ex. 2.25(d)

(e). Let φ_n and F_n be, respectively, the characteristic function and the c.d.f. of the random vector $X_n = (X_{n,1},\ldots,X_{n,k})$ for $n = 1, 2, \ldots$. Amend the derivations of Ex. 2.28 to show that for any $\varepsilon > 0$, and with $A_\varepsilon = (-\varepsilon,\varepsilon)\times\cdots\times(-\varepsilon,\varepsilon)$,

$$\begin{aligned}\int_{A_\varepsilon}\{1-\varphi_n(t)\}\,\mathrm{d}t &\geq 2^k\varepsilon^k\int_{\cap_{j=1}^k\{|\varepsilon x_j|\geq c_k\}}\Big(1-\prod_{j=1}^k\frac{\sin(x_j\varepsilon)}{x_j\varepsilon}\Big)\,\mathrm{d}F_n(x_1,\ldots,x_k)\\ &\geq 2^k\varepsilon^k(1-1/c_k^k)\Pr\big(\cap_{j=1}^k\{|X_{n,j}|\geq c_k/\varepsilon\}\big),\end{aligned}$$

Lévy's continuity theorem in dim. k

so that with $c_k = 2/(2^k-1)^{1/k}$ we get the inequality

$$\Pr\big(\cap_{j=1}^k\{|X_{n,j}|\geq c_k/\varepsilon\}\big)\leq\frac{1}{\varepsilon^k}\int_{-\varepsilon}^{\varepsilon}\cdots\int_{-\varepsilon}^{\varepsilon}\{1-\varphi_n(t_1,\ldots,t_k)\}\,\mathrm{d}t_1\cdots\mathrm{d}t_k.$$

Combine this inequality with (b) to show that if $\phi_n(t_1,\ldots,t_k)$ converges to some limit function $\varphi(t_1,\ldots,t_p)$ that is continuous at $(t_1,\ldots,t_p) = (0,\ldots,0)$, then $(X_n)_{n\geq 1}$ is tight, and that consequently the limit function φ must be a c.f., which by Ex. 2.43(a) entails that $X_n \to_d X$, where X is a random variable with c.f. φ.

Ex. 2.45. *The multi-dimensional CLT.* In this exercise we prove two canonical central limit theorems for random vectors taking values in $\mathbb{R}^k$.

multivariate CLT

(a). Suppose $X_1, X_2, \ldots$ are i.i.d. random vectors with mean ξ and variance matrix Σ. Invoke the Cramér–Wold device to show that $\sqrt{n}(\bar{X}_n - \xi) \to_d \mathrm{N}_k(0,\Sigma)$.

multivariate Lindeberg theorem

(b). Let $X_1, X_2, \ldots$ be independent random vectors of dimension k, with means $\xi_1, \xi_2, \ldots$, and with finite positive definite variance matrices $\Sigma_1, \Sigma_2, \ldots$. Let $B_n = (\Sigma_1+\cdots+\Sigma_n)^{1/2}$, and write $Z_{n,i} = B_n^{-1}(X_i - \xi_i)$. Show that

$$Z_n = (\Sigma_1+\cdots+\Sigma_n)^{-1/2}\sum_{i=1}^n(X_i-\xi_i) = \sum_{i=1}^n Z_{n,i} \to_d \mathrm{N}_k(0, I_k),$$

provided the multivariate Lindeberg condition holds, i.e.,

$$\sum_{i=1}^n \mathrm{E}\,\|Z_{n,i}\|^2 I(\|Z_{n,i}\|\geq\varepsilon)\to 0,\quad \text{for each } \varepsilon>0. \tag{2.10}$$

(c). Show that the multi-dimensional Liapunov condition $\sum_{i=1}^n \mathrm{E}\,\|Z_{n,i}\|^{2+\delta}\to 0$ for some $\delta > 0$, implies the Lindeberg condition in (2.10).

(d). Suppose that B_n^2/n converges to some positive definite Σ, and that the Lindeberg condition $n^{-1}\sum_{i=1}^n \mathrm{E}\,\|X_i-\xi_i\|^2 I(\|X_i-\xi_i\|\geq\sqrt{n}\varepsilon)\to 0$ for each $\varepsilon > 0$, holds. Show that $n^{-1/2}\sum_{i=1}^n(X_i-\xi_i)\to_d \mathrm{N}_p(0,\Sigma)$.

(e). Suppose that B_n^2/n converges to some positive definite Σ, and that there is a K and $\delta > 0$ such that $\mathrm{E}\,|X_{i,j} - \xi_{i,j}|^{2+\delta} \le K$ for $j = 1, \ldots, k$ and for all $i \ge 1$. Show that $n^{-1/2} \sum_{i=1}^n (X_i - \xi_i) \to_d \mathrm{N}_p(0, \Sigma)$.

Ex. 2.46. *Joint limiting normality for empirical mean and variance.* Suppose $X_1, X_2, \ldots$ are i.i.d. with mean ξ and standard deviation σ. We assume that also the skewness and kurtosis are finite, $\gamma_3 = \mathrm{E}\,(X_i - \xi)^3/\sigma^3$ amd $\gamma_4 = \mathrm{E}\,(X_i - \xi)^4/\sigma^4 - 3$.

(a). Show from the two-dimensional CLT that

$$\begin{pmatrix} \sqrt{n}(\bar{X}_n - \xi) \\ \sqrt{n}(\widehat{\sigma}_0^2 - \sigma^2) \end{pmatrix} \to_d \mathrm{N}_2(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2 & \gamma_3\sigma^3 \\ \gamma_3\sigma^3 & \sigma^4(2+\gamma_4) \end{pmatrix}),$$

where $\widehat{\sigma}_0^2 = n^{-1} \sum_{i=1}^n (X_i - \xi)^2$.

(b). Then show that with $\widehat{\sigma}^2 = n^{-1} \sum_{i=1}^n (X_i - \bar{X})^2$, which is a 'real estimator', as opposed to $\widehat{\sigma}_0^2$, which uses ξ, then we have $\sqrt{n}(\widehat{\sigma}^2 - \widehat{\sigma}_0^2) \to_{\mathrm{pr}} 0$; see Ex. 2.22. Conclude that the two-dimensional limit distribution result above continues to hold with $\sqrt{n}(\widehat{\sigma}^2 - \sigma^2)$ replacing $\sqrt{n}(\widehat{\sigma}_0^2 - \sigma^2)$.

(c). Let in particular the distribution of the X_i be normal, so $X_i \sim \mathrm{N}(\xi, \sigma^2)$. Show that γ_3 and γ_4 are equal to zero, so the general result above simplifies to

$$\begin{pmatrix} \sqrt{n}(\bar{X}_n - \xi) \\ \sqrt{n}(\widehat{\sigma}^2 - \sigma^2) \end{pmatrix} \to_d \mathrm{N}_2(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2 & 0 \\ 0 & 2\sigma^4 \end{pmatrix}).$$

For another instructive application, involving the Gamma distribution, see Ex. 3.27.

(d). Let $X_1, \ldots, X_n$ be independent Bernoulli variables with $\Pr(X_i = 1) = p$. With $\widehat{p} = \bar{X} = Y/n$ the familiar binomial ratio, explain first that the empirical variance here is $\widehat{\sigma}^2 = \widehat{p}(1 - \widehat{p})$. Then use the results above to find the joint limiting distribution of $\sqrt{n}(\widehat{p} - p, \widehat{\sigma}^2 - p(1-p))$.

The delta method

Ex. 2.47. *The delta method: the basics.* There are important and often reasonably simple to use approximation methods, in probability theory and statistics, going by the name of *the delta method*. It is related to functions of variables often being approximately linear, if the variables in question are not too spread out, with consequences for approximate normality.

(a). We start out in dimension one. Suppose in generic terms that $\sqrt{n}(A_n - a) \to_d Z$, with random variables A_n and a constant a. The limits are taken with respect to the index n tending to infinity, where n typically but not always is the size of an underlying sample. Show first that $A_n \to_{\mathrm{pr}} a$. Now consider a function g, defined in at least a neighbourhood around a, and assumed to have a continuous derivative there. Show via the mean value theorem that $g(A_n) - g(a) = g'(B_n)(A_n - a)$, for some B_n between A_n and a, and use this to show that $\sqrt{n}\{g(A_n) - g(a)\} \to_d g'(a)Z$.

(b). For the typical case where the limit is a zero-mean normal, say $Z \sim \mathrm{N}(0, \tau^2)$, explain that $\sqrt{n}\{g(A_n) - g(a)\} \to_d \mathrm{N}(0, g'(a)^2\tau^2)$.

(c). For the vector case, assume $\sqrt{n}(A_n - a) \to_d Z$ in dimension p, and consider a smooth function $g(a) = g(a_1, \ldots, a_p)$, assumed to have continuous first order derivatives in a neighbourhood around a. Explain that $A_n \to_{\rm pr} a$, so A_n will be inside that neighbourhood with probability tending to one. With $c = g^*(a)$ the vector of derivatives $\partial g(a)/\partial a_1, \ldots, \partial g(a)/\partial a_p$, evaluated at a, show that $\sqrt{n}\{g(A_n) - g(a)\} \to_d c^{\rm t} Z = c_1 Z_1 + \cdots + c_p Z_p$. In particular, if $Z \sim {\rm N}_p(0, \Sigma)$, explain that the limit $c^{\rm t} Z$ is a ${\rm N}(0, c^{\rm t}\Sigma c)$. The main point is that the convergence in distribution, from the previous points, holds also jointly.

(d). Sometimes one also needs the vector-to-vector extension of results above, leading to the joint limit distribution of several such g functions. Let $g = (g_1, \ldots, g_q)^{\rm t} \colon \mathbb{R}^p \to \mathbb{R}^q$ be such a function, with component functions having first order derivatives at position a, yielding a Jacobi matrix $J = \partial g(a)/\partial a$ of dimension $q \times p$. Show that $\sqrt{n}\{g(A_n) - g(a)\} \to_d JZ$. In particular, if $Z \sim {\rm N}_p(0, \Sigma)$, then the limit is ${\rm N}_q(0, J\Sigma J^{\rm t})$.

(e). The delta method has been formulated here in terms of limit distributions, implying approximate normality for function of approximately normal variables. Explain that in the framework above, with $\sqrt{n}(A_n - a) \to_d {\rm N}_p(0, \Sigma)$, then $g(A_n)$ is approximately normal, with mean $g(a)$ and variance $c^{\rm t}\Sigma c/n$. The delta method hence gives approximations also to means, variances, and covariances of smooth functions of other variables.

(f). The $\sqrt{n}$ factor above comes from the most typical uses of these methods, where variances of estimators go to zero with speed $1/n$, in terms of sample size. Show however that the mathematics goes through for any increasing sequence; if $d_n(A_n - a) \to_d Z$, with d_n tending to infinity, then $d_n\{g(A_n) - g(a)\} \to_d g'(a)Z$. There are indeed situations where the rate might be $n^{2/3}$ or $n^{1/3}$.

Ex. 2.48. *Applying the delta method.* Here we exercise our delta method muscles, to see how the general recipes of Ex. 2.47 may be applied in a few situations. As is clear, once we have established limiting normality for certain quantities, e.g. via the CLT, there is a host of easy-to-harvest applications for functions of these start quantities.

(a). If $\sqrt{n}(A_n - a) \to_d {\rm N}(0, 1)$, find out what happens to $\sqrt{n}(A_n^3 - a^3)$, and to $\sqrt{n}\{\exp(kA_n) - \exp(ka)\}$, where k is a constant.

(b). For Y a binomial (n, p), we have of course ${\rm Var}\,\widehat{p} = p(1-p)/n$ for the classic estimator $\widehat{p} = Y/n$. Use the delta method to find approximations to the means, variances, and distributions of (i) the estimated odds ratio $\widehat{p}/(1 - \widehat{p})$; (ii) the estimated log-odds-ratio $\log \widehat{p} - \log(1 - \widehat{p})$; (iii) the transformed estimator $\widehat{\gamma} = 2\arcsin(\widehat{p}^{1/2})$. In particular, for the latter, show $\sqrt{n}(\widehat{\gamma} - \gamma) \to_d {\rm N}(0, 1)$, with $\gamma = 2\arcsin(p^{1/2})$.

(c). Suppose $\widehat{p}_1 = Y_1/n$ and $\widehat{p}_2 = Y_2/n$ are two binomial estimates, with the same sample size n. Then $\sqrt{n}(\widehat{p}_1 - p_1) \to_d Z_1$ and $\sqrt{n}(\widehat{p}_2 - p_2) \to_d Z_2$, where $Z_j \sim {\rm N}(0, p_j(1 - p_j))$. Find the approximate normal distribution of $\widehat{p}_1/\widehat{p}_2$, viewed as an estimator of p_1/p_2. Modify arguments appropriately to find a good approximation to the variance of $\widehat{p}_1/\widehat{p}_2$, and its approximate normal distribution, also in the case of unequal sample sizes, say n_1 and n_2. [xx pointer to Story i.1. xx]

(d). Suppose $Y_1, \ldots, Y_n$ are independent from the geometric distribution with $\Pr(Y_i = y) = (1-p)^{y-1}p$ for $y = 1, 2, \ldots$. We learned in Ex. 1.25 that the mean and variance are $1/p$ and $(1-p)/p^2$. Find first the limiting distribution of $\sqrt{n}(\bar{Y} - 1/p)$ and then that of $\sqrt{n}(\widehat{p} - p)$, where $\widehat{p} = 1/\bar{Y}$.

(e). Suppose (a,b) is a certain position on the map, where one only has estimates, say A_n and B_n, for its x- and y-coordinates. Assume these are independent, approximately unbiased, and approximate normal, after n measurements. We formalise a version of this as $\sqrt{n}(A_n - a) \to_d N_1$ and $\sqrt{n}(B_n - b) \to_d N_2$, the limit variables N_1, N_2 being independent and standard normal. Having observed A_n and B_n, explain how you can put up 90 percent confidence intervals for a and b separately. Construct also a 90 percent confidence circle for (a,b).

(f). Let us pass from Cartesian to polar coordinates, letting

$$R_n = \|(A_n, B_n)\| = (A_n^2 + B_n^2)^{1/2} \quad \text{and} \quad \widehat{\alpha}_n = \arctan(B_n/A_n),$$

seen as estimators of the length $r = \|(a,b)\|$ and angle $\alpha = \arctan(b/a)$. Find the limit distributions for $\sqrt{n}(R_n - r)$ and $\sqrt{n}(\widehat{\alpha}_n - \alpha)$, and show that these are independent in the limit.

(g). Suppose one observes $(A_n, B_n) = (4.44, 2.22)$, with $n = 100$. Construct and display an approximate 90 percent confidence circle for (a,b), and then approximate 90 percent confidence intervals for the length r and angle α. How can you construct confidence intervals for r and α jointly, say $I_{r,n}$ and $I_{\alpha,n}$, such that the probability that $(r \in I_{r,n}) \cap (\alpha \in I_{\alpha,n})$ converges to 0.90?

Ex. 2.49. *Limiting normality for multinomials.* Consider the multinomial setup of Ex. 1.5, with $(Y_1, \ldots, Y_k)$ counting the number of events of type $1, \ldots, k$ in n independent experiments, each time with probabilities $p = (p_1, \ldots, p_k)^{\mathrm{t}}$. Here we sort out the basic large-sample behaviour of the relative frequencies $\widehat{p}_j = Y_j/n$. This is used e.g. in the Karl Pearson 1900 Story vii.1.

(a). Show that these $\widehat{p}_j = Y_j/n$ are consistent, with $\sqrt{n}(\widehat{p}_j - p_j) \to_d \mathrm{N}(0, p_j(1-p_j))$. Show more generally that there is full joint convergence in distribution here; $X_n = \sqrt{n}(\widehat{p} - p) \to_d Z \sim \mathrm{N}_k(0, \Sigma)$, where Σ is the matrix with elements $\sigma_{j,\ell} = p_j \delta_{j,\ell} - p_j p_\ell$. It may be written $\Sigma = D - pp^{\mathrm{t}}$ with D diagonal with elements p_j. Verify that this is consistent with $\sum_{j=1}^k Z_j = 0$.

(b). For $\gamma = g(p_1, \ldots, p_k)$ any smooth function of the relative frequencies, with natural estimator $\widehat{\gamma} = g(\widehat{p}_1, \ldots, \widehat{p}_k)$, show that $\sqrt{n}(\widehat{\gamma} - \gamma) \to_d \mathrm{N}(0, \tau^2)$, with $\tau^2 = c^{\mathrm{t}} \Sigma c = c^{\mathrm{t}} D c - (c^{\mathrm{t}} p)^2$, where $c = \partial g(p)/\partial p$. Check what this says, for the case of $\gamma = p_1 + \cdots + p_k$.

(c). Consider (X, Y, Z) being trinomial (n, p, q, r). With $\widehat{p} = X/n$, $\widehat{q} = Y/n$, $\widehat{r} = Z/n$, find the limit distribution for $\widehat{\gamma} = \widehat{p}/(\widehat{q}\widehat{r})^{1/2}$, as well as for $\widehat{\delta} = 2 \arcsin(\widehat{p}^{1/2}) - 2 \arcsin(\widehat{q}^{1/2})$.

Ex. 2.50. *Smooth functions of means.* Suppose $Y_1, \ldots, Y_n$ are i.i.d. from some distribution. (xx point suitably to Ex. 3.21, with joint normality for empirical quantiles, after which methods of this exercise apply also to these; smooth functions of means and quantiles are approximately normal. xx)

(a). Consider sample averages $\bar{g}_j = n^{-1} \sum_{i=1}^n g_j(Y_i)$, for functions $g_1, \ldots, g_p$, with finite variances. With $\xi_j = \mathrm{E}\, g_j(Y_i)$ the means, show from the CLT that $\sqrt{n}(\bar{g}_1 - \xi_1, \ldots, \bar{g}_p - \xi_p)^{\mathrm{t}}$ tends to a zero-mean multinormal $(X_1, \ldots, X_p)$, with covariance matrix Σ, having

elements $\sigma_{j,k} = \text{cov}\{g_j(Y_i), g_k(Y_i)\}$. Let then $A(x_1, \ldots, x_p)$ be any smooth function of these means. With $\bar{A}_n = A(\bar{g}_1, \ldots, \bar{g}_p)$, show that

$$\sqrt{n}\{\bar{A}_n - A(\xi_1, \ldots, \xi_p)\} \to_d a_1 X_1 + \cdots + a_p X_p \sim \text{N}(0, \tau^2),$$

with $\tau^2 = a^{\text{t}} \Sigma a$, in terms of $a_j = \partial A(\xi_1, \ldots, \xi_p)/\partial \xi_j$.

(b). For an illustration, let $Y_1, \ldots, Y_n$ be i.i.d. from the unit exponential distribution. Show that the skewness and kurtosis are skew $= 2$ and kurt $= 6$. Identify the multinormal limit of $\sqrt{n}(\bar{g}_1 - 1, \bar{g}_2 - 2, \bar{g}_3 - 6, \bar{g}_4 - 24)$, for $\bar{g}_j = n^{-1} \sum_{i=1}^n Y_i^j$. Find the limit distribution for $\widehat{\sigma} = (\bar{g}_2 - \bar{g}_1^2)^{1/2}$. Then find the limit distribution of $\sqrt{n}(\text{skew}^* - 2)$ and $\sqrt{n}(\text{kurt}^* - 6)$, with the empirical skew$^* = n^{-1} \sum_{i=1}^n W_i^3$ and kurt$^* = n^{-1} \sum_{i=1}^n W_i^4 - 3$, with $W_i = (Y_i - \bar{g}_1)/\widehat{\sigma}$. Find also the joint binormal limit for skew*, kurt*.

Ex. 2.51. *Delta method calculus for the normal case.* Let $Y_1, \ldots, Y_n$ be i.i.d. from the normal $\text{N}(\xi, \sigma^2)$, with standard estimators $\widehat{\xi} = \bar{Y}$ and $\widehat{\sigma}^2 = (n-1)^{-1} \sum_{i=1}^n (Y_i - \bar{Y})^2$. In various exercises in Ch. 1 we have worked with *exact finite-sample* calculus, for certain basic parameters, like the mean, variance, the quantile $\gamma_q = \xi + z_p \sigma$. Here we show how the delta method, starting with the basic limit distributions for the two parameters, can be used to put up large-sample normal approximations for *any* functions of the parameters, in cases where it might be too hard to carry out exact finite-sample calculus.

(a). Since the skewness and the kurtosis for the normal are zero, show that the general result from Ex. 2.70 implies that $(\sqrt{n}(\bar{Y} - \xi), \sqrt{n}(\widehat{\sigma} - \sigma))^{\text{t}}$ tends to say $(A, B)^{\text{t}}$, with these being independent and zero-mean normals with variances σ^2 and $\frac{1}{2}\sigma^2$. Show this directly, from the normality assumptions, as opposed to deriving it as a special case of the general statement. Note also that the $\sqrt{n}(\bar{Y} - \xi) \sim \text{N}(0, \sigma^2)$ holds exactly, for each finite n.

(b). With $\alpha = g(\xi, \sigma)$, for any smooth function of the two parameters, the natural estimator is $\widehat{\alpha} = g(\bar{Y}, \widehat{\sigma})$. Show that

$$\sqrt{n}(\widehat{\alpha} - \alpha) \to_d cA + dB \sim \text{N}(0, (c^2 + \tfrac{1}{2} d^2)\sigma^2),$$

where c and d are the partial derivatives of g, evaluated at the position (ξ, σ). Show how this leads to construction of confidence intervals for the α parameter.

(c). Consider the probability $p = \Pr(Y \geq y_0) = 1 - \Phi((y_0 - \mu)/\sigma)$, for some given threshold y_0, and the associated estimator $\widehat{p} = 1 - \Phi((y_0 - \bar{Y})/\widehat{\sigma})$. Finr the limit distribution of $\sqrt{n}(\widehat{p} - p)$, and use this to put up a confidence interval for p, with coverage level converging to 0.90. Compare with the simpler estimator $p^* = n^{-1} \sum_{i=1}^n I(Y_i \geq y_0)$, the binomial proportion, which bypasses the normal assumption.

(d). Then consider the parameter $\kappa = \xi/\sigma$, the normalised mean (so its value is unchanged when one passes from say millimetres to metres). For $\widehat{\kappa} = \bar{Y}/\widehat{\sigma}$, show that $\sqrt{n}(\widehat{\kappa} - \kappa) \to_d \text{N}(0, 1 + \frac{1}{2}\kappa^2)$, and construct an approximate 90 percent confidence interval. For this particular parameter, check also Ex. 3.39 for an exact method. (xx check if $g(\kappa) = \sqrt{2} \log(\kappa + (2 + \kappa^2)^{1/2})$ is an ok variance stabilising transformation; $\sqrt{n}\{g(\widehat{\kappa}) - g(\kappa)\}$ should tend to $\text{N}(0, 1)$. xx)

(e). Assume ξ is positive, so that the so-called coefficient of variation $\gamma = \sigma/\xi$ has a natural interpretation. With $\widehat{\gamma} = \widehat{\sigma}/\widehat{\xi}$ the plug-in estimator, show that $\sqrt{n}(\widehat{\gamma} - \gamma)$

tends to a normal with variance $\tau^2 = \frac{1}{2}\gamma^2(1+2\gamma^2)$. Find also a variance stabilising transformation, from γ to $\gamma^* = h(\gamma)$, with the property that $\sqrt{n}(\widehat{\gamma}^* - \gamma^*) \to_d \mathrm{N}(0,1)$: one such is $h(\gamma) = \log\gamma - \log(1+(1+2\gamma^2)^{1/2})$. Explain how confidence intervals can be set for γ via this.

Ex. 2.52. *Estimating mean and standard deviation outside normality.* Let $Y_1, \ldots, Y_n$ be i.i.d. from a distribution with finite fourth moment, and consider the usual mean $\bar{Y}_n$ and empirical standard deviation S_n. Under normality we have precise finite-sample results regarding their distributions, see Ex. 1.49, but here we investigate behaviour outside normality.

(a). Let as on previous occasions γ_3 and γ_4 be the skewness and kurtotis of the distribution. Use the delta method, with previous results from Ex. 2.70, to show that

$$\begin{pmatrix} \sqrt{n}(\bar{Y}_n - \xi) \\ \sqrt{n}(\widehat{\sigma}_n - \sigma) \end{pmatrix} \to_d \mathrm{N}_2\Big(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \sigma^2 \begin{pmatrix} 1 & \frac{1}{2}\gamma_3 \\ \frac{1}{2}\gamma_3 & \frac{1}{2} + \frac{1}{4}\gamma_4 \end{pmatrix}\Big).$$

Explain that this implies $\Pr(|\bar{Y}_n - \xi| \le 1.96\,\widehat{\sigma}) \to 0.95$, even outside normality; the normality-based standard recipes regarding inference for the mean parameter is hence not muich hampered by that modelling assumption.

(b). The situation is different when it comes to inference for the standard deviation, however. Show that $\sqrt{n}(\widehat{\sigma} - \sigma)/\widehat{\kappa} \to_d \mathrm{N}(0,1)$, if $\widehat{\kappa}$ is a consistent estimator for $\kappa = (\frac{1}{2} + \frac{1}{4}\gamma_4)^{1/2}$. Construct indeed such an estimator; see also Ex. 3.12. For an application of this, with sample size up to a million, see Story vii.2.

Ex. 2.53. *Approximate variances and the delta method.* We have built a well-working apparatus around the delta method and seen various applications. Statements reached are in terms of precise limit distributions, though without going into the quality of the resulting approximations. The present exercise goes into the details of variances and covariances for functions of sample averages.

(a). Suppose $X_1, \ldots, X_n$ are i.i.d. with mean zero, variance σ^2, and finite skewness $\gamma_3 = \mathrm{E}\,(X_i/\sigma)^3$ and kurtosis $\gamma_4 = \mathrm{E}\,(X_i/\sigma)^4 - 3$. With $\bar{X}_n$ as usual being the average, show then that

$$\begin{aligned} \mathrm{E}\,\bar{X}_n^2 &= \sigma^2/n, \\ \mathrm{E}\,\bar{X}_n^3 &= (\sigma^3/n^2)\gamma_3, \\ \mathrm{E}\,\bar{X}_n^4 &= (\sigma^4/n^2)(3+\gamma_4/n), \\ \mathrm{Var}\,\bar{X}_n^2 &= (\sigma^4/n^2)(2+\gamma_4/n). \end{aligned}$$

(b). Let then $Y_1, \ldots, Y_n$ be i.i.d., with finite mean ξ, standard deviation σ, skewness γ_3, kurtosis γ_4. With $\bar{Y}_n$ the sample average, consider then the variable

$$Z_n = a_0 + a_1(\bar{Y}_n - \xi) + \tfrac{1}{2}a_2(\bar{Y}_n - \xi)^2.$$

Show that $\mathrm{E}\,Z_n = a_0 + \frac{1}{2}a_2\sigma^2/n$, and that

$$\mathrm{Var}\,Z_n = a_1^2\sigma^2/n + (1/n^2)\{\tfrac{1}{4}a_2^2\sigma^4(2+\gamma_4/n) + a_1a_2\sigma^3\gamma_3\}.$$

(c). Consider any smooth function $Z_n = g(\bar{Y}_n)$. Since $\bar{Y}_n$ is close to ξ with high probability, it makes sense to carry out a Taylor expansion,

$$Z_n = g(\xi) + g'(\xi)(\bar{Y} - \xi) + \tfrac{1}{2}g''(\xi)(\bar{Y} - \xi)^2 + \delta_n,$$

where δ_n is a smaller-sized remainder term; you may prove that $n^{3/2}\delta_n$ is bounded in probability, provided g has three derivatives in a neighbourhood around ξ. Show from the above that

$$\operatorname{Var} Z_n = g'(\xi)^2\sigma^2/n + (1/n^2)\{\tfrac{1}{4}g''(\xi)^2\sigma^4(2+\gamma_4/n) + g'(\xi)g''(\xi)\sigma^3\gamma_3\} + o(1/n^2).$$

There is hence a clear leading $O(1/n)$ term for the variance, with other terms being of size $O(1/n^2)$. Explain how this relates to the basics of the delta method from Ex. 2.47.

(d). Suppose A is a random variable with mean a and finite variance, and that $g(y)$ is smooth in a neighbourhood around a. Use the Taylor approximation

$$g(y) = g(a) + g'(a)(y-a) + \tfrac{1}{2}g''(a)(y-a)^2 + O(|y-a|^3),$$

valid for y close to a, to show that

$$\mathrm{E}\, g(A) \doteq g(a) + \tfrac{1}{2}g''(a)\operatorname{Var} Y, \quad \operatorname{Var} g(Y) \doteq g'(a)^2 \operatorname{Var} Y,$$

and indicate the sizes of the error terms involved.

Ex. 2.54. *The empirical correlation coefficient under binormality.* For i.i.d. data pairs (X_i, Y_i), the classical empirical correlation coefficient is

$$R_n = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\{\sum_{i=1}^n (X_i - \bar{X})^2\}^{1/2}\{\sum_{i=1}^n (Y_i - \bar{Y})^2\}^{1/2}} = n^{-1}\sum_{i=1}^n \frac{X_i - \bar{X}}{\widehat{\sigma}_1}\frac{Y_i - \bar{Y}}{\widehat{\sigma}_2}, \quad (2.11)$$

with $\widehat{\sigma}_1^2$ and $\widehat{\sigma}_2^2$ the empirical variances $n^{-1}\sum_{i=1}^n (X_i - \bar{X})^2$ and $n^{-1}\sum_{i=1}^n (Y_i - \bar{Y})^2$. Here we find the the limit distribution of R_n under binormality.

(a). Assume first that the (X_i, Y_i) pairs are from a zero-mean binormal with variances 1 and correlation $\rho \in (-1, 1)$; see Ex. 1.43. Use results from Ex. 1.44, including $Y_i \,|\, X_i \sim \mathrm{N}(\rho X_i, 1-\rho^2)$, to derive expressions for $\mathrm{E}\, X_i^2 Y_i^2$, $\mathrm{E}\, X_i^3 Y_i$, $\mathrm{E}\, X_i Y_i^3$. Use these to show that

$$\Sigma = \operatorname{Var}\begin{pmatrix} X_i^2 \\ Y_i^2 \\ X_iY_i \end{pmatrix} = \begin{pmatrix} 2 & 2\rho^2 & 2\rho \\ 2\rho^2 & 2 & 2\rho \\ 2\rho & 2\rho & 1+\rho^2 \end{pmatrix}.$$

Use the CLT to argue that

$$\begin{pmatrix} A_n \\ B_n \\ C_n \end{pmatrix} = \sqrt{n}\begin{pmatrix} n^{-1}\sum_{i=1}^n X_i^2 - 1 \\ n^{-1}\sum_{i=1}^n Y_i^2 - 1 \\ n^{-1}\sum_{i=1}^n X_iY_i - \rho \end{pmatrix} \to_d \begin{pmatrix} A \\ B \\ C \end{pmatrix} \sim \mathrm{N}_3(0, \Sigma).$$

With $R_{n,0} = C_n/(A_nB_n)^{1/2}$, use the delta method to show that $\sqrt{n}(R_{n,0} - \rho) \to_d Z = -\frac{1}{2}\rho A - \frac{1}{2}\rho B + C$, and that in fact $Z \sim \mathrm{N}(0, (1-\rho^2)^2)$.

(b). Then generalise to the situation where the (X_i, Y_i) pairs are i.i.d. from a zero-mean binormal, with standard deviations σ_1, σ_2 and correlation ρ. Show that we still have $\sqrt{n}(R_{n,0} - \rho) \to_d \mathrm{N}(0, (1-\rho^2)^2)$.

(c). Then go one step further, to the full five-parameter binormal situation, with unknown means ξ_1, ξ_2, standard deviations σ_1, σ_2, and correlation ρ. Argue first that we must have $\sqrt{n}(R_{n,1} - \rho) \to_d \mathrm{N}(0, (1-\rho^2)^2)$, where $R_{n,1}$ is as in (2.11) but using ξ_1, ξ_2 instead of $(\bar{X}, \bar{Y})$. Then, finally, show what we really are after, that $\sqrt{n}(R_n - \rho)$ must have the same limit distribution.

(d). Explain that if $h(\rho)$ is a smooth function, then $\sqrt{n}\{h(R_n) - h(\rho)\} \to_d h'(\rho)(1-\rho^2)\,\mathrm{N}(0,1)$. Show that with the particular choice $\zeta = \frac{1}{2}\log\{(1+\rho)/(1-\rho)\}$ the variance is being stabilised, and $\sqrt{n}(\widehat{\zeta} - \zeta) \to_d \mathrm{N}(0,1)$, where $\widehat{\zeta} = \frac{1}{2}\log\{(1+R_n)/(1-R_n)\}$. This is called Fisher's zeta. Show that $\widehat{\zeta} \pm 1.645/\sqrt{n}$ becomes an approximate 90 percent confidence interval for ζ, and transform this to an approximate 90 percent confidence interval for ρ.

variance stabilising transformation

Fisher's zeta

Ex. 2.55. *The empirical correlation coefficient, general case.* Here we use some of the arguments of Ex. 2.54 to find the limit distribution of the empirical correlation R_n of (2.11) also outside binormality. Assume $(X_1, Y_1), \ldots, (X_n, Y_n)$ are i.i.d. pairs from a distribution with means ξ_1, ξ_2, standard deviations σ_1, σ_2, and correlation ρ. Write $a_{j,k} = \mathrm{E}\, U_i^j V_j^k$ for cross moments of the standardised $U_i = (X_i - \xi_1)/\sigma_1$ and $V_i = (Y_i - \xi_2)/\sigma_2$, where it is assumed that fourth order moments $a_{4,0}$ and $a_{0,4}$ are finite.

(a). Show that $\sqrt{n}(R_n - R_{n,0}) \to_{\mathrm{pr}} 0$, where $R_{n,0}$ is as R_n, but using the real ξ_1, ξ_2 instead of their estimators $\bar{X}, \bar{Y}$. Show also that the distribution of R_n and $R_{n,0}$ must depend on ρ but not on $\xi_1, \xi_2, \sigma_1, \sigma_2$. We may hence carry out our large-sample investigation with the standardised (U_i, V_i) rather than the (X_i, Y_i). Work with

$$\begin{pmatrix} A_n \\ B_n \\ C_n \end{pmatrix} = \begin{pmatrix} \sqrt{n}(n^{-1}\sum_{i=1}^n U_i^2 - 1) \\ \sqrt{n}(n^{-1}\sum_{i=1}^n V_i^2 - 1) \\ \sqrt{n}(n^{-1}\sum_{i=1}^n U_i V_i - \rho) \end{pmatrix},$$

and show that $(A_n, B_n, C_n)^{\mathrm{t}} \to_d (A, B, C)^{\mathrm{t}} \sim \mathrm{N}_3(0, \Sigma)$, for the variance matrix Σ of $(U_i^2, V_i^2, U_i V_i)^{\mathrm{t}}$. Spell out the elements of this matrix, using the $a_{j,k}$. Check that this agrees with the Σ of Ex. 2.54 under binormality.

(b). Then show that $\sqrt{n}(R_n - \rho) \to_d Z = -\frac{1}{2}\rho A - \frac{1}{2}\rho B + C$, and give an expression for the limit distribution variance τ^2. Explain how τ may be estimated from the data, and how this leads to confidence intervals of the type $R_n \pm 1.96\,\widehat{\tau}/\sqrt{n}$ for ρ.

(c). For a concrete illustration, consider the joint density $f(x,y) = 1 + a(x - \frac{1}{2})(y - \frac{1}{2})$ for (x,y) in the unit square. Find the allowed parameter range for a, and a formula for the correlation coefficient ρ in terms of a. Then apply the above to find the limit distribution of $\sqrt{n}(R_n - \rho)$.

Ex. 2.56. *The delta method outside root-n terrain.* (xx to come. not always $\sqrt{n}(X_n - a) \to_d \mathrm{N}(0, \tau^2)$ terrain. different limits, different speeds. xx)

Ex. 2.57. *Stretching the delta method.* (xx to be filled in. with $\sqrt{n}(X_n - a) \to_d V$, we have $Z_n = \sqrt{n}\{g(X_n) - g(a)\} \to_d g'(a)V$ for a fixed $g(x)$. here we consider $Z_n = \sqrt{n}\{g_n(X_n) - g_n(a)\}$. with $g_n''(a) = o(\sqrt{n})$ we may still have the right approximation. example: $Z_n = \sqrt{n}\{\exp(c_n X_n) - 1\}$. xx)

Ex. 2.58. *The Strong Law of Large Numbers: the basics.* (xx to be cleaned. xx) Suppose $X_1, X_2, \ldots$ are i.i.d. from a distribution with finite $\mathrm{E}\,|X_i|$. Then the mean $\xi = \mathrm{E}\,X_i$ exists, and we are aiming to prove the strong LLN of (2.3), that the event

$$A = \{\bar{X}_n \to \xi\} = \cap_{\varepsilon>0} \cup_{n_0\geq 1} \cap_{n\geq n_0} \{|\bar{X}_n| \leq \varepsilon\}$$

has probability equal to one hundred percent. We may for simplicity and without loss of generality take $\xi = 0$ below.

(a). Show that A is the same as $\cap_{N\geq 1}\cup_{n_0\geq 1}\cap_{n\geq n_0}\{|\bar{X}_n| \leq 1/N\}$, and deduce in particular from this that A is actually measurable – so it does make well-defined sense to work with its probability.

(b). Show that if $\Pr(A_N) = 1$ for all N, then $\Pr(\cap_{N\geq 1} A_N) = 1$ – if you're fully certain about a countable number of events, then you're also fully certain about all of them, jointly. This is actually not true with a bigger index set: if $X \sim \mathrm{N}(0,1)$, then you're 100 percent sure that $B_x = \{X \text{ is not } x\}$ takes place, for each single x, but from this does it *not* follow that you should be sure about $\cap_{\text{all } x} B_x$. Explain why.

(c). Show that $\Pr(A) = 1$ if and only if $\Pr(B_{n_0}) \to 0$, for each $\varepsilon > 0$, where $B_{n_0} = \cup_{n\geq n_0}\{|\bar{X}_n| \geq \varepsilon\}$. In words: for a given ε, the probability should be very low that there is *any* $n \geq n_0$ with $|\bar{X}_n| \geq \varepsilon$.

(d). A simple bound is of course $\Pr(B_{n_0}) \leq \sum_{n\geq n_0} \Pr\{|\bar{X}_n| \geq \varepsilon\}$, so it suffices to show, if possible, under appropriate conditions, that $\sum_{n\geq 1} \Pr\{|\bar{X}_n| \geq \varepsilon\}$ is a convergent series. With finite variance σ^2, explain that the classic simple Chebyshov bound, see Ex. 2.10, does *not* solve any problem here.

(e). (xx calibrate better with Ex. 2.10. xx) Show, however, that if the fourth moment is finite, then

$$\Pr\{|\bar{X}_n| \geq \varepsilon\} \leq \frac{1}{\varepsilon^4}\mathrm{E}\,|\bar{X}_n|^4 \leq \frac{c}{\varepsilon^4}\frac{1}{n^2},$$

for a suitable c. So under this condition, which is moderately hard, we've proven the strong LLN.

(f). One may squeeze more out of the chain of arguments here, which we now indicate, without full details. Assume $\mathrm{E}\,|X_i|^r$ is finite, for some $r > 2$, like $r = 2.02$. Then one may show, via arguments in von Bahr (1965), that the sequence $\mathrm{E}\,|\sqrt{n}\bar{X}_n|^r$ is bounded. This leads to the bound

$$\Pr\{|\bar{X}_n| \geq \varepsilon\} \leq \frac{1}{(\sqrt{n}\varepsilon)^r}\mathrm{E}\,|\sqrt{n}\bar{X}_n|^r,$$

and these form a convergent series. We have hence proven (modulo the von Bahr thing) that the strong LLN holds for finite $\mathrm{E}\,|X_i|^{2+\varepsilon}$, an improvement over the finite $\mathrm{E}\,|X_i|^4$ condition. – To get further, trimming away on the conditions until we are at the Kolmogorovian position of only requiring finite mean, we need more technicalities; see the following Ex. 2.59.

Ex. 2.59. *The Strong Law of Large Numbers: nitty-gritty details.* This exercise goes through the required extra technical details, along with a few intermediate lemmas, to secure a full proof of the full LLN theorem: as long as $\mathrm{E}\,|X_i|$ is finite, the infinite sequence of sample means $\bar{X}_n$ will with probability equal to one hundred percent converge to $\xi = \mathrm{E}\,X_i$.

(a). We start with Kolmogorov's inequality: Consider independent zero-mean variables $X_1, \ldots, X_n$ with variances $\sigma_1^2, \ldots, \sigma_n^2$, and with partial sums $S_i = X_1 + \cdots + X_i$. Then

$$\Pr\{\max_{i\le n} |S_i| \ge \varepsilon\} \le \frac{\mathrm{Var}\,S_n}{\varepsilon^2} = \frac{1}{\varepsilon^2}\sum_{i=1}^n \sigma_i^2.$$

Note that this is a much stronger result than the special case of caring only about $|S_n|$, with $\Pr\{|S_n| \ge \varepsilon\} \le \mathrm{Var}\,S_n/\varepsilon^2$, which is the Chebyshov inequality. To prove it, work with the disjoint decomposition

$$A_i = \{|S_1| < \varepsilon, \ldots, |S_{i-1}| < \varepsilon, |S_i| \ge \varepsilon\} \quad \text{and} \quad A = \cup_{i=1}^n A_i = \{\max_{i\le n} |S_i| \ge \varepsilon\}.$$

Show that $\mathrm{E}\,S_n^2 \ge \mathrm{E}\,S_n^2 I(A) = \sum_{i=1}^n \mathrm{E}\,S_n^2 I(A_i)$, that

$$\mathrm{E}\,S_n^2 I(A_i) = \mathrm{E}\,(S_i + S_n - S_i)^2 I(A_i) \ge \varepsilon^2 \Pr(A_i),$$

and that this leads to the inequality asked for.

(b). Consider a sequence of independent $X_1, X_2, \ldots$ with means zero and variances $\sigma_1^2, \sigma_2^2, \ldots$. Show that if $\sum_{i=1}^\infty \sigma_i^2$ is convergent, then $\sum_{i=1}^\infty X_i$ is convergent with probability 1. – It suffices to show that the sequence of partial sums $S_n = X_1 + \cdots + X_n$ is Cauchy with probability 1. Show that this is the same as

$$\lim_{n\to\infty} \Pr\big[\cup_{i,j\ge n}\{|S_i - S_j| \ge \varepsilon\}\big] = 0 \quad \text{for each } \varepsilon > 0.$$

Use the Kolmogorov inequality to show this.

(c). A quick example to illustrate this result is as follows. Consider $X = X_1/10 + X_2/100 + X_3/1000 + \cdots$, a random number in the unit interval, with the X_i independent, and with no further assumptions. Show that X exists with probability 1.

(d). Prove that if $\sum_{i=1}^\infty a_i/i$ converges, then $\bar{a}_n = (1/n)\sum_{i=1}^n a_i \to 0$. To show this, consider $b_n = \sum_{i=1}^n a_i/i$, so that $b_n \to b$ for some b. Show $a_n = n(b_n - n_{n-1})$, valid also for $n = 1$ if we set $b_0 = 0$, and which leads to $\sum_{i=1}^n a_i = nb_n - b_0 - b_1 - \cdots - b_{n-1}$.

(e). From the above, deduce that if $X_1, X_2, \ldots$ are independent with means $\xi_1, \xi_2, \ldots$ and variances $\sigma_1^2, \sigma_2^2, \ldots$, and $\sum_{i=1}^\infty \sigma_i^2/i^2$ converges, then $\bar{X}_n - \bar{\xi}_n \to_{\text{a.s.}} 0$. Here $\bar{\xi}_n = (1/n)\sum_{i=1}^n \xi_i$.

(f). Use the above to show that if $X_1, X_2, \ldots$ are independent with zero means, and all variances are bounded, then indeed $\bar{X}_n \to_{\text{a.s.}} 0$. Note that this is a solid generalisation of what we managed to show in (xx calibrate xx) – first, the distributions are allowed to be different (not identical); second, we have landed at a.s. convergence with the mild assumption of finite and bounded variances, whereas we there needed the harsher conditions of finite fourth moments.

(g). We're close to the Pole. For i.i.d. zero mean variables $X_1, X_2, \ldots$, split them up with the little trick

$$X_i = Y_i + Z_i, \quad \text{with} \quad Y_i = X_i I(|X_i| < i), \quad Z_i = X_i I(|X_i| \geq i).$$

We have $\bar{X}_n = \bar{Y}_n + \bar{Z}_n$, so it suffices to demonstrate that $\bar{Y}_n \to_{\text{a.s.}} 0$ and $\bar{Z}_n \to_{\text{a.s.}} 0$ (since an intersection of two sure events is sure). Use the Borel–Cantelli lemma, in concert with $\mathrm{E}\,|X_i| = \int_0^\infty \Pr(|X_i| \geq x)\,\mathrm{d}x$, to show that only finitely many Z_i are non-zero. Then use previous results to demonstrate $\bar{Y}_n - \bar{\xi}_n \to_{\text{a.s.}} \to 0$ and $\bar{\xi}_n \to 0$, where $\bar{\xi}_n$ is the average of $\xi_i = \mathrm{E}\,Y_i$.

(h). So we've managed to prove the Strong LLN; good. Attempt also to prove the interesting converse that if $\mathrm{E}\,|X_i| = \infty$, then the sequence of sample means is pretty erratic indeed:

$$\Pr\{\limsup_{n\to\infty} \bar{X}_n = \infty\} = 1, \quad \Pr\{\liminf_{n\to\infty} \bar{X}_n = -\infty\} = 1.$$

Simulate a million realisations from the density $f(x) = 1/x^2$, for $x \geq 1$, in your nearest computer, display the sequence of $\bar{X}_n$ on your screen, and comment.

Ex. 2.60. *Yes, we converge with probability 1.* We've proven that the sequence of empirical means converges almost surely to the population mean, under the sole condition that this mean is finite. This half-automatically secures almost sure convergence of various other natural quantities, almost without further efforts.

(a). Suppose $X_1, X_2, \ldots$ are i.i.d. with finite variance σ^2. Show that the classical empirical standard deviation $\widehat{\sigma} = \{\sum_{i=1}^n (X_i - \bar{X}_n)^2/(n-1)\}^{1/2}$ converges a.s. to σ. Note again that nothing more is required than a finite second moment.

(b). Suppose the third moment is finite, such that the skewness $\gamma_3 = \mathrm{E}\,\{(X-\xi)/\sigma\}^3$ is finite. Show that $\widehat{\gamma}_{3,n} = (1/n)\sum_{i=1}^n (X_i - \bar{X}_n)^3/\widehat{\sigma}^3$ is strongly consistent for γ_3.

(c). Then suppose the fourth moment is finite, such that the kurtosis $\gamma_4 = \mathrm{E}\,\{(X - \xi)/\sigma\}^4 - 3$ is finite. Construct a strongly consistent estimator for this kurtosis.

(d). Assume that $(X_1, Y_1), (X_2, Y_2), \ldots$ is an i.i.d. sequence of random pairs, with finite variances, and define the population correlation coefficient in the usual fashion, as $\rho = \mathrm{cov}(X, Y)/(\sigma_1\sigma_2)$. Show that the usual empirical correlation coefficient R_n of (2.11) converges with probability one to ρ.

(e). Formulate and prove a suitable statement regarding almost sure convergence of smooth functions of means.

(f). Let $X_1, X_2, \ldots$ be an i.i.d. sequence of nonnegative variables such that $\mathrm{E}\,\log X_i = \xi$ is finite. Show that the harmonic means $H_n = (X_1 \cdots X_n)^{1/n}$ converge with probability 1 to $\exp(\xi)$.

Ex. 2.61. *Glivenko–Cantelli theorem.* For i.i.d. observations $Y_1, \ldots, Y_n$, we form the empirical c.d.f., as in Ex. 2.11, with $F_n(t) = n^{-1}\sum_{i=1}^n I(Y_i \leq t)$. We have seen that since $F_n(t)$ is just a binomial ratio, $F_n(t) \to_{\text{a.s.}} F(t)$, for each t. It is a remarkable fact that this convergence also takes place uniformly, with probability 1. This is the Glivenko–Cantelli theorem: with $D_n = \max_t |F_n(t) - F(t)|$, the max taken over all t in

the domain in question, we have $\Pr(D_n \to 0) = 1$. This means that regardless of any strange or complicated aspects of the distribution F, with enough data one will be able to learn these. See also Ex. 3.9 and 9.25 for more information, regarding the speed with which $D_n \to 0$.

Glivenko–Cantelli

(a). Choose $t_1 < \cdots < t_m$, creating a finite number of cells $[t_j, t_{j+1})$, where we take $t_0 = -\infty$ and $t_{m+1} = \infty$. With $A_{m,j}$ the event that $F_n(t_j) \to F(t_j)$, argue that $\Pr(\cap_{j=1}^m A_{m,j}) = 1$.

(b). Consider any t in the cell $[t_j, t_{j+1})$. Writing $D_n(t) = F_n(t) - F(t)$, use monotinicity of F_n and F to show that

$$D_n(t_j) - \{F(t_{j+1}) - F(t_j)\} \le F_n(t) - F(t) \le D_n(t_{j+1}) + F(t_{j+1} - F(t_j).$$

Deduce that

$$\max_{t_j \le t < t_{j+1}} |D_n(t)| \le B_m + C_m,$$

where $B_m = \max_{1 \le j \le m} |D_n(t_j)|$ and $C_m = \max_{1 \le j \le m} \{F(t_{j+1}) - F(t_j)\}$.

(c). Show that $\Pr(\limsup D_n \le C_m) = 1$.

(d). For each $\varepsilon > 0$, show that a partition into cells can be arranged, with high m if required, so that $C_m \le \varepsilon$. Conclude that $\Pr(D_n \to 0) = 1$.

(e). Choose some moderately complicated normal mixture, of the type $f = \sum_{j=1}^k p_j \mathrm{N}(\mu_j, \sigma_j^2)$; see Ex. 1.66. Then simulate a high number n of data from this distribution, and read off $D_n = \max_t |F_n(t) - F(t)|$. Check out how high n needs to be to have $D_n \le 0.01$, say, with high probability, in a few situations.

Contiguity

[xx couple of exercises xx]

Stable and conditional convergence, and nonnormal limits

Ex. 2.62. *Stable convergence.* This exercise introduces the notion of stable convergence of random variables, which is a form of convergence lying between convergence in probability and convergence in distribution (to paraphrase Jacod and Mémin (1981), an early article on the topic). On a probability space $(\Omega, \mathcal{F}, \Pr)$, let $(X_n)_{n \ge 1}$ be a sequence of random variables, and let $\mathcal{G} \subset \mathcal{F}$. The sequence X_n converges $\mathcal{G}$-stably to the random variable X, where X is defined on an extension $(\Omega \times \mathbb{R}, \mathcal{G} \otimes \mathcal{B}(\mathbb{R}), \Pr')$, if

$$\mathrm{E}\, \xi f(X_n) \to \mathrm{E}'\, \xi f(X), \tag{2.12}$$

for every bounded $\mathcal{G}$-measurable random variable ξ, and every bounded and continuous function f. See Ex. A.32 for extensions of probability spaces. Here $\mathrm{E}'(\cdot)$ denotes the expectation with respect to $\Pr'$. We write $X_n \to_{\mathcal{G}\text{-st.}} X$ to indicate this form of convergence.

(a). Show that if $X_n \to_{\mathcal{G}\text{-st.}} X$ then $X_n \to_d X$.

(b). To see that the converse of (a) is not true, consider the sequence $X_n = I(n \text{ odd})Y + I(n \text{ even})Y'$, where Y and Y' are i.i.d. random variables with common distribution F. Show that $X_n \to_d F$, but that X_n fails to converge $\mathcal{G} = \sigma(Y)$ stably, for example. This example is from Aldous and Eagleson (1978).

(c). Suppose that $(X_n, Y) \to_d (X, Y)$ for every $\mathcal{G}$-measurable Y. We'll soon see that this is equivalent to $X_n \to_{\mathcal{G}\text{-st.}} X$. For now, assume that the limit X is also $\mathcal{G}$-measurable, and show that then $X_n \to_p X$. – This illustrates that to have $\mathcal{G}$-stable convergence of X_n to X without also having convergence in probability, X must be realised in a fashion that does not render it $\mathcal{G}$-measurable, hence the extension of the original probability space.

(d). Show that the following are equivalent:

(i) $X_n \to_{\mathcal{G}\text{-st.}} X$;
(ii) $(X_n, Y) \to_d (X, Y)$ for every $\mathcal{G}$-measurable random vector Y;
(iii) $(X_n, Y) \to_{\mathcal{G}\text{-st.}} (X, Y)$ for every $\mathcal{G}$-measurable random vector Y;
(iv) $(X_n, Y_n) \to_d (X, Y)$ for every sequence $(Y_n)_{n\geq 1}$ of random variables and every $\mathcal{G}$-measurable Y such that $Y_n \to_p Y$;
(v) $\mathrm{E}\, I_G f(X_n) \to \mathrm{E}'\, I_G f(X)$ for every $G \in \mathcal{G}$ and every bounded and continuous f;
(vi) $\mathrm{E}\,\{f(X_n) \,|\, G\} \to \mathrm{E}'\,\{f(X) \,|\, G\}$, for every $G \in \mathcal{G}$ with $\Pr(G) > 0$, and every bounded and continuous f;
(vii) $\mathrm{E}\, I_G \exp(itX_n) \to \mathrm{E}'\, I_G \exp(itX)$ for every $G \in \mathcal{G}$.

To prove this, Ex. 2.19**??** and Ex. A.33(c) might be of help.

(e). [xx rewrite xx] Let Q_n and Q be versions of the conditional distributions of X_n and X given $\mathcal{G}$, respectively, and let $Q_n f = \int_{\mathbb{R}} f(x) Q_n(\cdot, \mathrm{d}x)$, with Qf similarly defined. Show that $X_n \to_{\mathcal{G}\text{-st.}} X$ is equivalent to $\mathrm{E}\,\xi Q_n f \to \mathrm{E}\,\xi Q f$, for every ξ and f as above.

(f). [xx cramer slutsky for stable convergence here xx]

Ex. 2.63. *Conditional convergence.*[xx introtext here xx]

(a). Suppose that $(X_n, Y_n) \to_d (X, Y)$. Show that we then also have marginal convergence, that is $X_n \to_d X$ and $Y_n \to_d Y$.

(b). Show that $(X_n, Y_n) \to_d (X, Y)$ is equivalent to

$$\mathrm{E}\,\{f(X_n) \,|\, Y_n \in B\} \to \mathrm{E}\,\{f(X_n) \,|\, Y_n \in B\},$$

for all bounded and continuous f and all continuity sets B of the distribution of Y such that $\Pr(Y \in B) > 0$ and $\Pr(Y_n \in B) > 0$ for all n.

(c). Suppose $(X_n, Y_n) \to_d (X, Y)$, a binormal zero-mean limit, say $\mathrm{N}_2(0, \Sigma)$, with

$$\Sigma = \begin{pmatrix} \sigma_X^2 & \rho\sigma_X\sigma_Y \\ \rho\sigma_X\sigma_Y & \sigma_Y^2 \end{pmatrix}.$$

Deduce from (b) that $X_n \,|\, (|Y_n| \leq \varepsilon) \to_d X \,|\, (|Y| \leq \varepsilon)$ for each $\varepsilon > 0$, i.e., that the conditional distribution of X_n given $|Y_n| \leq \varepsilon$ converges in distribution to that of X given $|Y| \leq \varepsilon$. Show that $X \,|\, (|Y| \leq \varepsilon) \to_d X \,|\, (Y = 0) \sim \mathrm{N}(0, (1-\rho^2)\sigma_X^2)$ as $\varepsilon \to 0$.

(d). Consider a distribution for X_i with mean zero and variance σ^2, where we indeed know that $\sqrt{n}\bar{X}_n \to_d \mathrm{N}(0,\sigma^2)$. Suppose X_i also has an integrable characteristic function $\varphi(t)$, implying by Ex. A.40 the existence of a smooth density f for X_i and also a density f_n for $\sqrt{n}\bar{X}_n$. Show that

$$f_n(z) = 1/(2\pi)\int \exp(-itz)\varphi(t/\sqrt{n})^n\,\mathrm{d}t \to \sigma^{-1}\phi(\sigma^{-1}z),$$

i.e., that there is convergence not merely for the cumulatives, but for the densities too. To do so, you may split the domain of integration in two parts $|t| > \varepsilon\sqrt{n}$ and $|t| \le \varepsilon\sqrt{n}$, for some $\varepsilon > 0$, and use Ex. A.42.

(e). Suppose next that X_i is discrete, without a continuous density; for a concrete example, consider $X_i = \pm 1$ with equal probabilities, and for which $\varphi(t) = \cos t$. Then $Z_n = \sqrt{n}\bar{X}_n$ does not have a density, but we may add a little Gaussian noise, to form $Z_n^* = \sqrt{n}\bar{X}_n + \xi_n$, with $\xi_n \sim \mathrm{N}(0,\varepsilon_n^2)$. Show that Z_n^* has density

$$f_n^*(z^*) = 1/(2\pi)\int \exp(-itz^*)\varphi(t/\sqrt{n})^n \exp(-\tfrac{1}{2}\varepsilon_n^2 t^2)\,\mathrm{d}t,$$

and that this again converges to the normal density $\sigma^{-1}\phi(\sigma^{-1}z^*)$ provided merely that $\varepsilon_n \to 0$.

(f). We may use the same trick also in the vector case. Specifically, with a p-dimensional X_i having zero mean and covariance matrix Σ, show (i) that if X_i has an integrable characteristic function, then the density for $Z_n = \sqrt{n}\bar{X}_n$ tends to the $\mathrm{N}_p(0,\Sigma)$ density; and (ii) that if X_i is discrete, without a density, then $Z_n^* = \sqrt{n}\bar{X}_n + \xi_n$, with some small Gaussian added noise $\xi_n \sim \mathrm{N}_p(0,\varepsilon_n^2\Sigma)$ with $\varepsilon_n \to 0$, has a density $f_n^*(z^*)$ which tends to the same $\mathrm{N}_p(0,\Sigma)$ density.

(g). Suppose $(X_n, Y_n) \to_d (X, Y)$, a zero-mean binormal. If there is also density convergence, with (X_n, Y_n) having a density $f_n(x,y)$ tending to the appropriate binormal density $f(x,y)$; and $f_n(x,y)$ as well as the marginals $f_{Y_n}(y)$ are continuous, show that $X_n \,|\, (|Y_n| \le \delta_n)$ tends to $X \,|\, (Y = 0)$, for an appropriate sequence $\delta_n \to 0$.

(h). Show that the same limiting distribution statement as in (g) holds also when (X_n, Y_n) has a discrete distribution, or a noncontinuous density, using the 'adding small Gaussian noise to get densities' trick.

Ex. 2.64. *Limiting normality of rank sums statistics.* (xx to be edited and polished; nils rant so far. we put it in if it looks smooth enough, and with a brief pointer to Story v.7. point to Swensen (1983). xx) In a population of n individuals, followed on some continuous scale, a subgroup of interest, of size m, has ranks $X_1, \ldots, X_m$. These form a randomly selected subset of size m from $\{1, \ldots, n\}$, with all such $\binom{n}{m}$ subsets equally likely. The rank sum $Z_n = X_1 + \cdots + X_m$ is the Wilcoxon statistic.

(a). Explain that one may write $Z_n = \sum_{i=1}^n iJ_i$, where the 0-1 variables J_i are such that precisely m of them are 1, and with all $\binom{n}{m}$ subsets of such 1s being equally likely. Find $\mathrm{E}\,J_i$, $\mathrm{Var}\,J_i$, $\mathrm{cov}(J_i, J_j)$ for $j \ne i$. Writing $p = m/n$ for the sample ratio, show using either of the representations $\sum_{i=1}^m X_i$ or $\sum_{i=1}^n iJ_i$ that

$$\mathrm{E}\,Z_n = \tfrac{1}{2}m(n+1) \doteq \tfrac{1}{2}n^2 p, \quad \mathrm{Var}\,Z_n = (1/12)(n+1)m(n-m) \doteq (1/12)n^3 p(1-p).$$

(b). We aim indeed at showing limiting normality of Z_n here, with both n and m becoming larger, with $m/n \to p$. Explain that this must mean

$$(Z_n - \tfrac{1}{2}n^2p)/n^{3/2} \to_d \mathrm{N}(0, (1/12)p(1-p)).$$

We cannot use CLT or Lindeberg for studying Z_n, since the X_i are dependent, as are the J_i. Consider however a different parallel setup, involving independent Bernoulli variables $J_1^*, \ldots, J_n^*$ with $\Pr(J_i^* = 1) = p = m/n$. Explain that the distribution of Z_n is the same as the distribution of $Z_n^* = \sum_{i=1}^n iJ_i^*$ given $\sum_{i=1}^n J_i^* = m$. Show now that

$$\begin{pmatrix} A_n \\ B_n \end{pmatrix} = \begin{pmatrix} (1/\sqrt{n})\sum_{i=1}^n (i/n)(J_i^* - p) \\ (1/\sqrt{n})\sum_{i=1}^n (J_i^* - p) \end{pmatrix} \to_d \begin{pmatrix} A \\ B \end{pmatrix} \sim \mathrm{N}_2(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1/3 & 1/2 \\ 1/2 & 1 \end{pmatrix}).$$

(c). Find the distribution of $A \,|\, (B = b)$, and show in particular that $A \,|\, (B = 0) \sim \mathrm{N}(0, 1/12)$. This gives a clear limiting normality statement for Z_n^*, conditional on $\sum_{i=1}^n J_i^* = m$, and nicely solves our Wilcoxon problem; show that it corresponds precisely to the $(Z_n - \frac{1}{2}n^2p)/n^{3/2}$ limiting statement above. (xx some extra care needed. but an easy and instructive way to show normality for Wilcoxon, and also for other related variables. to illustrate, find limiting normality for $X_1^{1/2} + \cdots + X_m^{1/2} = \sum_{i=1}^n i^{1/2}J_i$. xx)

(d). (xx if we manage: also a link via the uniform order statistic process, and to integrals of salt-and-pepper processes, with $W_n = \int_0^1 ([ns]/n)\,\mathrm{d}C_n(s)$, with the $\mathrm{d}C_n(s)$ being $\mathrm{d}s$ or 0 with probabilities p and $1-p$, but conditional on the random region $\int_0^1 \mathrm{d}C_n(s)$ being $p = m/n$. if we're lucky there is a limit expressible as integral or Brownian bridge, for Ch9. nils will attempt to fix this and at least make the idea more precise. xx)

Ex. 2.65. *Nonnormal limits.* (xx polish this. point to process version, with more results for hitting times, etc., in Ch. 9. xx) Normally limits are normal, but not always. Here we shall indeed work with variables with mean zero and variance one, where the sample averages have nonnormal limits. The basic construction is as follows. Let $U_1, U_2, \ldots$ be i.i.d., with mean zero and variance one, and with m.g.f. $M_0(s) = \mathrm{E}\exp(sU_i)$ finite in a neighbourhood around zero; in particular, all moments for the U_i are finite. Let independently of these $J_1, J_2, \ldots$ be independent Bernoulli variables with $\Pr(J_i = 1) = 1/i$, $\Pr(J_i = 0) = 1 - 1/i$. Then form

$$Z_n = \frac{1}{\sqrt{n}} \sum_{i=1}^n J_i \sqrt{i}\, U_i = \sum_{i=1}^n J_i \sqrt{i/n}\, U_i.$$

A picture to have in mind is that most of the terms will be zero, with non-zero contributions becoming both more rare and more big as time proceeds.

(a). Show that there will with probability one be infinitely many $J_i = 1$, i.e. non-zero terms in the Z_n sum as n grows.

(b). Show that the terms $J_i\sqrt{i}U_i$ have mean zero and variance one; hence also the normalised sample average Z_n has mean zero and variance one. Find also an expression for the kurtosis $\kappa_n = \mathrm{E}\,Z_n^4 - 3$ of Z_n, and show that $\kappa_n \to \frac{1}{2}a_4$, where $a_4 = \mathrm{E}\,U_i^4$. Compare this to what we are 'used to' from the Lindeberg theorem.

(c). We already know that if Z_n has a limit distribution, it can't be normal. Working with the m.g.f., show that

$$M_n(t) = \mathrm{E}\exp(tZ_n) = \prod_{i=1}^n \Big[1 + \frac{1}{i}\{M_0(t\sqrt{i/n}) - 1\}\Big],$$

for all t around zero for which $M_0(t)$ is finite.

(d). Show that

$$\prod_{i=1}^{n}\Big[1+\frac{1}{i}\{M_0(t\sqrt{i/n})-1\}\Big]\to\exp\Big\{\int_0^1\frac{M_0(t\sqrt{x})-1}{x}\,\mathrm{d}x\Big\}.$$

Work first with Special Case One, where we let U_i have the simple symmetric two-point distribution $\Pr(U_i=1)=\Pr(U_i=-1)=\frac12$. Find the limiting kurtosis for Z_n in this case. Show that $M_0(s)=\frac12e^s+\frac12e^{-s}=1+(1/2!)s^2+(1/4!)s^4+\cdots$, and use this to find an infinite-sum expression for the limit of $M_n(t)$. Have you now proved that Z_n has a limit distribution?

(e). Then work with Special Case Two, where the U_i have a double exponential distribution, of the form $f(u)=\frac12\sqrt{2}\exp(-\sqrt{2}|u|)$ on the real line (the $\sqrt{2}$ factor is there to ensure variance one). Find the m.g.f. $M_0(s)$ for the U_i, and then the m.g.f. $M(t)$ for the limit distribution of Z_n.

(f). For most cases, regarding the distribution for the U_i, it is hard to learn the explicit distribution for Z_n (even in cases where there might be a clear distribution for its limit). For Special Case Two, however, find the explicit distribution for Z_n, for any given n.

CLTs for dependent random variables

Ex. 2.66. *A CLT for 1-dependent variables.* (xx decide later if these few should be pushed to Ch. 12. xx) Consider a stationary sequence $Y_1,Y_2,\ldots$, with mean zero and variance one, being 1-dependent. Stationarity means $(Y_1,\ldots,Y_r)$ having the same distribution as $(Y_{i+1},\ldots,Y_{i+r})$, for any i and block lengths r, and 1-dependence means that Y_i,Y_{i+1} may be dependent, but $Y_1,\ldots,Y_i$ is independent of $Y_{i+2},Y_{i+3},\ldots$. This exercise reaches a CLT for $\sum_{i=1}^n Y_i$, representing a genuine extension of the usual CLT and Lindeberg theorems from independence.

(a). Writing $\rho=\mathrm{corr}(Y_i,Y_{i+1})$, show that $(1/k)\mathrm{Var}(Y_1+\cdots+Y_k)=1+2(1-1/k)\rho$ which then goes to $1+2\rho$ for increasing k.

(b). For a given block length k, split $Y_1+\cdots+Y_n$ into $[n/k]$ blocks, and write block j of these as U_j+V_j, with U_j as sum of $k-1$ consecutive observations and V_j the last one of that block. Write then

$$Z_n=(1/\sqrt{n})\sum_{i=1}^n Y_i=(1/\sqrt{n})\Big(\sum_{j=1}^{[n/k]}U_j+\sum_{j=1}^{[n/k]}V_j+E_n\Big),$$

with E_n any extra left after the $k[n/k]$ variables captured in these first $[n/k]$ blocks.

(c). Explain why $U_1,U_2,\ldots$ are independent, so that the usual CLT applies to these. Show that $(1/\sqrt{n})\sum_{j=1}^{[n/k]}U_j\to_d\mathrm{N}(0,\tau_k^2)$, with $\tau_k^2=(1/k)\mathrm{Var}\,(Y_1+\cdots+Y_{k-1})$.

(d). Then use Ex. 2.30 to prove that $(1/\sqrt{n})\sum_{i=1}^n Y_i\to_d\mathrm{N}(0,1+2\rho)$, i.e. a CLT for 1-dependent variables.

(e). Assume $X_1, X_2, \ldots$ are i.i.d. with mean zero and variance one. Consider $Z_n = (1/\sqrt{n})(X_1X_2 + X_2X_3 + \cdots + X_{n-1}X_n)$. Show that $Z_n \to_d \mathrm{N}(0,1)$. Show also that

$$Z_n' = (1/\sqrt{n}) \sum_{i=1}^{n-1} (X_i - \bar{X}_n)(X_{i+1} - \bar{X}_n)$$

has the same limit distribution, where $\bar{X}_n$ as usual is the sample mean.

Ex. 2.67. *A CLT for m-dependent variables.* In natural generalisation of Ex. 2.66, consider a stationary m-dependent sequence $Y_1, Y_2, \ldots$, with mean zero and variance σ^2. There is accordingly potential dependence among $Y_1, \ldots, Y_m$, but for any i, $(Y_1, \ldots, Y_i)$ is independent of $(Y_{i+m+1}, \ldots, Y_n)$.

(a). Writing $\mathrm{cov}(Y_i, Y_j) = \sigma^2 \rho(|j-i|)$, with the autocorrelation function $\rho(\cdot)$, show first that in general terms,

$$(1/n) \operatorname{Var}\Big(\sum_{i=1}^{n} Y_i\Big) = \sigma^2 \Big\{1 + 2\sum_{j=1}^{n} (1 - j/n)\rho(j)\Big\}.$$

Then show that for the case of m-dependence, for any $k \geq m$, we have $(1/k)\operatorname{Var}(Y_1 + \cdots + Y_k) \to \sigma^2\{1 + 2\sum_{j=1}^{m} \rho(j)\}$,

(b). Extend arguments and techniques from Ex. 2.66 to show that $(1/\sqrt{n})\sum_{i=1}^{n} Y_i$ tends to a zero-mean normal with variance $\sigma^2\{1 + 2\rho(1) + \cdots + 2\rho(m)\}$.

(c). (xx a bit on how the acf works for an i.i.d. sequence: $\sqrt{n}\widehat{\rho}(j) \to_d \mathrm{N}(0,1)$, for each j. xx)

Further material and leftovers, as of 12-March-2025

Ex. 2.68. *Approximate multinormality.* (xx an old exercise we have possibly forgotten to take in again; used in a couple of later exercises; need checking. xx) In Ex. 2.31, 2.34, and 2.38 we have proven the CLT for independent random variables. We also need to extend this machinery to the multi-dimensional case, also since there are many situations where a one-dimensional estimator or test statistic is a function of several components.

(a). To approach the notion of 'approximate multinormality' we need an associated notion of convergence in distribution to a multinormal distribution. A random vector $Y_n = (Y_{n,1}, \ldots, Y_{n,p})^{\mathrm{t}}$ converges to a limit distribution variable $Y = (Y_1, \ldots, Y_p)^{\mathrm{t}}$ in distribution if (and actually only if) all linear combinations converge accordingly:

$$Z_n = c^{\mathrm{t}} Y_n = c_1 Y_{n,1} + \cdots + c_p Y_{n,p} \to_d Z = c^{\mathrm{t}} Y = c_1 Y_1 + \cdots + c_p Y_p$$

for each vector $c = (c_1, \ldots, c_p)^{\mathrm{t}}$ [xx where do we prove this fact ? xx]. Show that $Y_n \to_d \mathrm{N}_p(0, \Sigma)$ if and only if $c^{\mathrm{t}} Y_n \to_d \mathrm{N}(0, c^{\mathrm{t}}\Sigma c)$ for all c.

(b). Prove that the multi-dimensional CLT then follows from the one-dimensional version: If $X_1, X_2, \ldots$ are i.i.d. from some p-dimensional distribution, with finite mean $\xi = \mathrm{E}\, X_i$ and variance matrix $\Sigma = \mathrm{Var}\, X_i$, then the normalised sum $Z_n = n^{-1}\sum_{i=1}^{n}(X_i - \xi) = \sqrt{n}(\bar{X}_n - \xi)$ tends to the multinormal $\mathrm{N}_p(0, \Sigma)$.

(c). Show that a random pair (X_n, Y_n) converges in distribution to the binormal $\mathrm{N}_2(0, \Sigma)$, with Σ having diagonal elements 1, 1, and correlation ρ, if and only if $aX_n + bY_n \to_d \mathrm{N}(0, a^2 + b^2 + 2\rho ab)$ for each (a, b).

(d). Suppose $X_1, X_2, \ldots$ are i.i.d. with mean ξ and standard deviation σ. We assume that also the skewness and kurtosis are finite, $\gamma_3 = \mathrm{E}\,(X_i-\xi)^3/\sigma^3$ amd $\gamma_4 = \mathrm{E}\,(X_i-\xi)^4/\sigma^4-3$. Show from the two-dimensional CLT that

$$\begin{pmatrix} \sqrt{n}(\bar{X}_n - \xi) \\ \sqrt{n}(\widehat{\sigma}_0^2 - \sigma^2) \end{pmatrix} \to_d \mathrm{N}_2\Big(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2, & \gamma_3\sigma^3 \\ \gamma_3\sigma^3, & \sigma^4(2+\gamma_4) \end{pmatrix}\Big),$$

where $\widehat{\sigma}_0^2 = n^{-1}\sum_{i=1}^n (X_i-\xi)^2$.

(e). Then show that with $\widehat{\sigma}^2 = n^{-1}\sum_{i=1}^n (X_i-\bar{X})^2$, which is a 'real estimator', as opposed to $\widehat{\sigma}_0^2$, which uses ξ, then we have $\sqrt{n}(\widehat{\sigma}^2-\widehat{\sigma}_0^2) \to_{\mathrm{pr}} 0$; see Ex. 2.22. Conclude that the two-dimensional limit distribution result above continues to hold with $\sqrt{n}(\widehat{\sigma}^2-\sigma^2)$ replacing $\sqrt{n}(\widehat{\sigma}_0^2-\sigma^2)$.

(f). Let in particular the distribution of the X_i be normal, so $X_i \sim \mathrm{N}(\xi, \sigma^2)$. Show that γ_3 and γ_4 are equal to zero, so the general result above simplifies to

$$\begin{pmatrix} \sqrt{n}(\bar{X}_n - \xi) \\ \sqrt{n}(\widehat{\sigma}^2 - \sigma^2) \end{pmatrix} \to_d \mathrm{N}_2\Big(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2, & 0 \\ 0, & 2\sigma^4 \end{pmatrix}\Big).$$

For another instructive application, involving the Gamma distribution, see Ex. 3.27.

Ex. 2.69. *Characteristic functions for vector variables.* With $X = (X_1, \ldots, X_k)^{\mathrm{t}}$ a random vector, in dimension k, we define its characteristic functions as

$$\varphi(t_1, \ldots, t_k) = \mathrm{E}\,\exp(it^{\mathrm{t}}X) = \mathrm{E}\,\exp\{i(t_1X_1 + \cdots + t_kX_k)\}$$

for $t = (t_1, \ldots, t_k)^{\mathrm{t}}$.

(a). Show that if the components are independent, then $\varphi(t_1, \ldots, t_k) = \varphi_1(t_1)\cdots\varphi_k(t_k)$, in terms of the individual characteristic functions $\varphi_1, \ldots, \varphi_k$.

(b). Show for the multinormal case, where $X \sim \mathrm{N}_k(\xi, \Sigma)$, that $\varphi(t) = \exp(it^{\mathrm{t}}\xi - \frac{1}{2}t^{\mathrm{t}}\Sigma t)$.

(c). (xx spell out basics for $\varphi_n(t) \to \varphi(t)$ as equivalent to $X_n \to_d X$. inversion formula. Perhaps leave details to the Appendix xx)

Ex. 2.70. *The Cramér–Wold device.* Suppose X_n and X are random variables in $\mathbb{R}^k$.

(a). Show that $X_n \to_d X$ if and only if $a^{\mathrm{t}}X_n \to_d a^{\mathrm{t}}X$ for each a, i.e. if and only if all linear combinations converge. This is the Cramér–Wold theorem. (xx proof via characteristic functions, so need to set up this properly first. xx)

(b). This very useful device, for proving convergence in distribution for random vectors, can now be used to prove the multivariate CLT: Suppose $X_1, X_2, \ldots$ are i.i.d. with mean ξ and variance matrix Σ. Show that $\sqrt{n}(\bar{X}_n - \xi) \to_d \mathrm{N}_k(0, \Sigma)$.

(c). Let $X_1, X_2, \ldots$ be independent random vectors in dimension k, with finite positive definite variance matrices $\Sigma_1, \Sigma_2, \ldots$. When will

$$Z_n = (\Sigma_1 + \cdots + \Sigma_n)^{-1/2}(X_1 + \cdots + X_n) \to_d \mathrm{N}_k(0, I_k)?$$

(xx give clear conditions. the multidimensional Lindeberg theorem. for each a, need $\sum_{i=1}^n a^{\mathrm{t}}X_i/(\sum_{i=1}^n a^{\mathrm{t}}\Sigma_i a)^{1/2} \to_d \mathrm{N}(0,1)$. to be used for regression models in Ch5 and later. give sets of easier conditions for the general Lindeberg to hold. xx)

Ex. 2.71. *Approximate variances and the delta method.* There are important and often reasonably simple to use approximation methods, in probability theory and statistics, going by the name of *the delta method*. It is related to functions of variables often being approximately linear, if the variables in question are not too spread out, with consequences for approximate normality. More details are discussed in Ex. **??**. Here we are content to work out the basics, and to see its uses in the following few exercises.

(a)**.** Suppose $X_1, \ldots, X_n$ are i.i.d. with mean zero, variance σ^2, and finite skewness $\gamma_3 = \mathrm{E}\,(X_i/\sigma)^3$ and kurtosis $\gamma_4 = \mathrm{E}\,(X_i/\sigma)^4 - 3$ (xx calibrate crosslink later xx). With $\bar X_n$ as usual being the average, show then that

$$\begin{aligned} \mathrm{E}\,\bar X_n^2 &= \sigma^2/n, \\ \mathrm{E}\,\bar X_n^3 &= (\sigma^3/n^2)\gamma_3, \\ \mathrm{E}\,\bar X_n^4 &= (\sigma^4/n^2)(3+\gamma_4/n), \\ \mathrm{Var}\,\bar X_n^2 &= (\sigma^4/n^2)(2+\gamma_4/n). \end{aligned}$$

(b)**.** Let then $Y_1, \ldots, Y_n$ be i.i.d., with finite mean ξ, standard deviation σ, skewness γ_3, kurtosis γ_4. With $\bar Y_n$ the sample average, consider then the variable

$$Z_n = a_0 + a_1(\bar Y_n - \xi) + \tfrac12 a_2(\bar Y_n - \xi)^2.$$

Show that $\mathrm{E}\,Z_n = a_0 + \frac12 a_2\sigma^2/n$, and that

$$\mathrm{Var}\,Z_n = a_1^2\sigma^2/n + (1/n^2)\{\tfrac14 a_2^2\sigma^4(2+\gamma_4/n) + a_1a_2\sigma^3\gamma_3\}.$$

(c)**.** Consider any smooth function $Z_n = g(\bar Y_n)$. Since $\bar Y_n$ is close to ξ with high probability (see Ex. 2.10), it makes sense to carry out a Taylor expansion,

$$Z_n = g(\xi) + g'(\xi)(\bar Y - \xi) + \tfrac12 g''(\xi)(\bar Y - \xi)^2 + \delta_n,$$

where δ_n is a smaller-sized remainder term – you may prove that $n^{3/2}\delta_n$ is bounded in probability, provided g has three derivatives in a neighbourhood around ξ. Show from the above that

$$\mathrm{Var}\,Z_n = g'(\xi)^2\sigma^2/n + (1/n^2)\{\tfrac14 g''(\xi)^2\sigma^4(2+\gamma_4/n) + g'(\xi)g''(\xi)\sigma^3\gamma_3\} + o(1/n^2).$$

(d)**.** There is hence a clear leading $O(1/n)$ term, for the variance, with other terms being of size $O(1/n^2)$. To the first order of approximation, show that

$$\mathrm{Var}\,Z_n = g'(\xi)^2\sigma^2/n + o(1), \quad \text{and} \quad \sqrt{n}(\bar Y_n - \xi) \xrightarrow{d} \mathrm{N}(0,\sigma^2).$$

(e)**.** This is a version of the so-called *delta method*. Try to prove, with direct methods, using the formal apparatus of Ex. 2.31 or 2.34

$$\sqrt{n}(A_n - a) \to_d Z \quad \text{implies} \quad \sqrt{n}\{g(A_n) - g(a)\} \to_d g'(a)Z.$$

If the limit Z is a $\mathrm{N}(0,\sigma^2)$, then $g'(a)Z$ is the $\mathrm{N}(0, g'(a)^2\sigma^2)$.

(f)**.** We come forcefully back to the delta method, with more details, generalisations, and uses, in Ex. **??**. Here we spell out the function of a vector version, since it is so immediately useful; see exercises below. – Suppose $\sqrt{n}(A_n - a) \to_d Z$, as random vectors of dimension k (xx calibrate with what is said previously in chapter on such

convergence; can go to linear combinations xx). With a smooth $g(A_n)$, defined at least in a neighbourhood around a, we have

$$\sqrt{n}\{g(A_n) - g(a)\} \to_d c^{\mathrm{t}} Z = c_1 Z_1 + \cdots + c_k Z_k,$$

where $c = \partial g(a)/\partial a$, the vector of partial derivatives $c_j = \partial g(a)/\partial a_j$, computed at position a. Show that if the limit Z of $\sqrt{n}(A_n - a)$ is multinormal, say $\mathrm{N}_k(\xi, \Sigma)$, then the limit $c^{\mathrm{t}} Z$ is normal $(c^{\mathrm{t}}\xi, c^{\mathrm{t}}\Sigma c)$. The vector version of the delta method is often given precisely in this form. (xx to form this as an exercise, go via linear combinations, which then needs us mentioning the cramer-wold, see Ex. 2.68, with details in Ex. 2.70. we make the readers understand that the vector case follows from the unidim case. xx)

(g). (xx The point is that once a start limit distribution result has been established, perhaps via the CLT, then a string of further limit distribution results come almost for free. simple illustration here; more in Ex. 2.72 and 2.51. xx)

Ex. 2.72. *Applying the delta method.* Here we exercise our delta method muscles, to see how the general recipes may be applied in a few situations.

(a). For Y a binomial (n, p), we have of course $\mathrm{Var}\,\widehat{p} = p(1-p)/n$ for the classic estimator $\widehat{p} = Y/n$. Use the delta method to find approximations to the means, variances, and distributions of (i) the estimated odds ratio $\widehat{p}/(1-\widehat{p})$; (ii) the estimated log-odds-ratio $\log\widehat{p} - \log(1-\widehat{p})$; (iii) the transformed estimator $\widehat{\gamma} = 2\arcsin(\widehat{p}^{1/2})$.

(b). Suppose $\widehat{p}_1 = Y_1/n$ and $\widehat{p}_2 = Y_2/n$ are two binomial estimates, with the same sample size n. Then $\sqrt{n}(\widehat{p}_1 - p_1) \to_d Z_1$ and $\sqrt{n}(\widehat{p}_2 - p_2) \to_d Z_2$, where $Z_j \sim \mathrm{N}(0, p_j(1-p_j))$. Find the approximate normal distribution of $\widehat{p}_1/\widehat{p}_2$, viewed as an estimator of p_1/p_2. Modify arguments appropriately to find a good approximation to the variance of $\widehat{p}_1/\widehat{p}_2$, and its approximate normal distribution, also in the case of unequal sample sizes, say n_1 and n_2. [xx pointer to Story i.1. xx]

(c). Suppose $Y_1, \ldots, Y_n$ are independent from the geometric distribution with $\Pr(Y_i = y) = (1-p)^{y-1}p$ for $y = 1, 2, \ldots$. We learned in Ex. 1.25 that the mean and variance are $1/p$ and $(1-p)/p^2$. Find first the limiting distribution of $\sqrt{n}(\bar{Y} - 1/p)$ and then that of $\sqrt{n}(\widehat{p} - p)$, where $\widehat{p} = 1/\bar{Y}$.

(d). Suppose (a, b) is a certain position on the map, where one only has estimates, say A_n and B_n, for its x- and y-coordinates. Assume these are independent, approximately unbiased, and approximate normal, after n measurements. We formalise a version of this as $\sqrt{n}(A_n - a) \to_d N_1$ and $\sqrt{n}(B_n - b) \to_d N_2$, the limit variables N_1, N_2 being independent and standard normal. Having observed A_n and B_n, explain how you can put up 90 percent confidence intervals for a and b separately. Construct also a 90 percent confidence circle for (a, b).

(e). Let us pass from Cartesian to polar coordinates, letting

$$R_n = \|(A_n, B_n)\| = (A_n^2 + B_n^2)^{1/2} \quad \text{and} \quad \widehat{\alpha}_n = \arctan(B_n/A_n),$$

seen as estimators of the length $r = \|(a, b)\|$ and angle $\alpha = \arctan(b/a)$. Find the limit distributions for $\sqrt{n}(R_n - r)$ and $\sqrt{n}(\widehat{\alpha}_n - \alpha)$, and show that these are independent in the limit. You may look ahead and use Ex. **??**(d).

(f). Suppose one observes $(A_n, B_n) = (4.44, 2.22)$, with $n = 100$. Construct and display an approximate 90 percent confidence circle for (a, b), and then approximate 90 percent confidence intervals for the length r and angle α. How can you construct confidence intervals for r and α jointly, say $I_{r,n}$ and $I_{\alpha,n}$, such that the probability that $(r \in I_{r,n}) \cap (\alpha \in I_{\alpha,n})$ converges to 0.90?

Ex. 2.73. *Estimating mean and standard deviation outside normality.* Let $Y_1, \ldots, Y_n$ be i.i.d. from a distribution with finite fourth moment, and consider the usual mean $\bar{Y}_n$ and empirical standard deviation S_n. Under normality we have precise finite-sample results regarding their distributions, see Ex. 1.49, but here we investigate behaviour outside normality.

(a). Let as on previous occasions γ_3 and γ_4 be the skewness and kurtotis of the distribution. Use the delta method, with previous results from Ex. 2.68, to show that

$$\begin{pmatrix} \sqrt{n}(\bar{Y}_n - \xi) \\ \sqrt{n}(\widehat{\sigma}_n - \sigma) \end{pmatrix} \to_d \mathrm{N}_2\Big(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \sigma^2 \begin{pmatrix} 1, & \frac{1}{2}\gamma_3 \\ \frac{1}{2}\gamma_3, & \frac{1}{2} + \frac{1}{4}\gamma_4 \end{pmatrix}\Big).$$

(b). Show that $\sqrt{n}(\widehat{\sigma} - \sigma)/\widehat{\kappa} \to_d \mathrm{N}(0, 1)$, if κ is a consistent estimator for $\kappa = (\frac{1}{2} + \frac{1}{4}\gamma_4)$. For an application of this, with sample size up to a million, see Story vii.2.

Ex. 2.74. *A curious dense random subset of the unit interval.* Here we construct a random subset A of the unit interval, where both A and its complement A^c have Lebesgue measure $\frac{1}{2}$, and where both are *dense*.

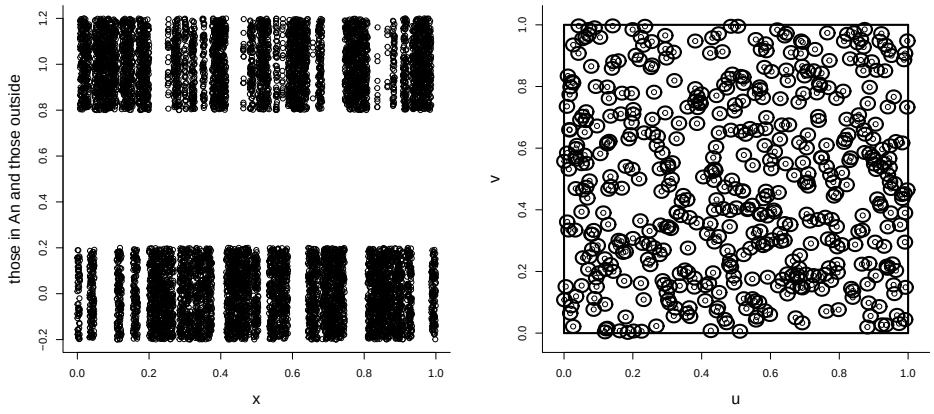


Figure 2.2: *The recipe of Ex. 2.74 is followed. Left panel: with $n = 50$, for the one-dimensional case, sets A_n and A_n^c are indicated. Right panel: for the two-dimensional case, the set A_n, the collection of points not covered by the union of the n tiny balls, has total area converging to $\frac{1}{2}$, here shown for $n = 500$.*

(a). Start with $u_1, \ldots, u_n$ i.i.d. uniform, then put tiny intervals around them, say $I_i = (u_i - \frac{1}{2}c/n, u_i + \frac{1}{2}c/n)$. Then define A_n as the set of points x outside all intervals $I_1, \ldots, I_n$; here we pedantically include also the small fraction of such points falling just outside the unit interval, i.e. in $(-\frac{1}{2}c/n, 0)$ or in $(1, 1 + \frac{1}{2}c/n)$. Show that $p_n(x) = \Pr(x \in A_n) = (1 - c/n)^n$, for each x, which tends to $\exp(-c)$. For the rest of the exercise we use

$c = \log 2$, for which $p_n(x) \to \frac12$. Construct a version of Figure 2.2 (left panel), there with $n = 50$. Increase the n to say 250, to see the visual difference.

(b). Consider the random area $M_n(t) = \int_0^t I(x \in A_n)\,\mathrm{d}x$ of non-covered points x over $[0,t]$, i.e. the fraction of points in that interval which lie in A_n. Show that its mean is $G_n(t) = \int_0^t p_n(x)\,\mathrm{d}x$, tending to $\frac12 t$, and that its variance goes to zero. Explain that this means $G_n(t) \to_{\rm pr} \frac12 t$, for each t.

(c). Explain that $G_n(t) \to_{\rm pr} \frac12 t$ implies the existence of a random set A which properly may be thought of as the limit of A_n, and that this set A has Lebesgue measure $\frac12$. (xx to be repaired. is this ok, or is it 'deeper than this'? xx)

(d). Let then $[b,c]$ be any interval. Show that the measure of $A_n \cap [b,c]$ tends to $\frac12(c-b)$. In this sense both the random A and its complement are everywhere dense in the unit interval.

(e). Carry out a similar programme for the two-dimensional case. For positions (u_i, v_i) uniformly distributed in the unit square, form tiny circles around them with radius $r_n = \{c/(\pi n)\}^{1/2}$. Let A_n be the set of points (x,y) avoiding all these balls. Show that $p_n(x,y) = \Pr((x,y) \in A_n) \to \frac12$, and that the size of this random subset tends to $\frac12$ in probability. Make a programme that plots this phenomenon, using e.g. $n = 500$, as with Figure 2.2.

(f). You're given the unit cube $[0,1]^3$. Create a system of tiny balls, centred at points $u_1, \ldots, u_n$ placed i.i.d. uniformly in the cube, with shrinking radii r_n. Decide on r_n such that the set A_n of points avoiding the n balls has Lebesgue measure tending to $\frac12$.

Ex. 2.75. *The Liapunov and Lindeberg theorems: main story.* Let $X_1, X_2, \ldots$ be independent zero-mean variables with at the outset different distributions $F_1, F_2, \ldots$ and hence different standard deviations $\sigma_1, \sigma_2, \ldots$. Below we also need their characteristic functions $\varphi_1, \varphi_2, \cdots$. The question is when we can rest assured that the normalised sum

Liapunov and Lindeberg theorems

$$Z_n = (X_1 + \cdots + X_n)/B_n = \sum_{i=1}^n X_i \Big/ \Big(\sum_{i=1}^n \sigma_i^2\Big)^{1/2}$$

really tends to the standard normal, as n increases.

(a). Show that Z_n has characteristic function

$$\kappa_n(t) = \mathrm{E}\,\exp(itZ_n) = \varphi_1(t/B_n)\cdots\varphi_n(t/B_n).$$

(b). From Ex. A.39 we have that $|\exp(ix) - 1 - ix - \frac12(ix)^2| \le \frac16|x|^3$ and that $\varphi_i(s) = 1 - \frac12\sigma_i^2 s^2 + o(s)$ for small s, so the essential idea is to write

$$\kappa_n(t) = \prod_{i=1}^n \{1 - \tfrac12\sigma_i^2 t^2/B_n^2 + \varepsilon_{n,i}(t)\},$$

and prevail until one has found conditions that secure convergence to the desired $\exp(-\frac12 t^2)$. In view of the first lemma in Ex. 2.36, this essentially takes (i) $\sum_{i=1}^n \varepsilon_{n,i}(t) \to 0$; (ii) $\max_{i\le n}\sigma_i^2/B_n^2 \to 0$ and $\max_{i\le n}|\varepsilon_{n,i}(t)| \to 0$; and (iii) $\sum_{i=1}^n |1 - \varphi_i(t/B_n)|$ staying bounded. Show that

$$\begin{aligned}|\varphi_i(s) - (1 - \tfrac12\sigma_i^2 s^2)| &= \Big|\int \{\exp(isx) - 1 - isx - \tfrac12(isx)^2\}\,\mathrm{d}F_i(x)\Big| \\ &\le \int |\exp(isx) - 1 - isx - \tfrac12(isx)^2|\,\mathrm{d}F_i(x) \le \tfrac16|s|^3\,\mathrm{E}\,|X_i|^3.\end{aligned}$$

(c). This leads to the условие Ляпунова version of the Lindeberg theorem: show that if the variables all have finite third order moments, with $B_n \to \infty$ and

$$\sum_{i=1}^n \mathrm{E}\Big|\frac{X_i}{B_n}\Big|^3 \to 0,$$

then $\kappa_n(t) \to \exp(-\frac{1}{2}t^2)$, which from Ex. 2.27 we know is equivalent to the desired $Z_n \to_d \mathrm{N}(0,1)$. This is (already) a highly significant extension of the CLT. If the X_i are uniformly bounded, for example, with B_n^2/n having a positive limit, which would rather often be the case, then the Lyapunov condition holds. It is also possible to refine arguments and methods to show that

$$\sum_{i=1}^n \mathrm{E}\Big|\frac{X_i}{B_n}\Big|^{2+\delta} \to 0, \quad \text{for some } \delta > 0,$$

is sufficient for limiting normality.

(d). The issue waits however for even milder and actually minimal conditions, and that is, precisely, the Lindeberg condition:

$$L_n(\varepsilon) := \sum_{i=1}^n \mathrm{E}\Big|\frac{X_i}{B_n}\Big|^2 I\Big\{\Big|\frac{X_i}{B_n}\Big| \geq \varepsilon\Big\} \to 0 \quad \text{for all } \varepsilon > 0. \tag{2.13}$$

Show that if the Lyapunov condition is in force, then the Lindeberg condition holds (so Lindeberg assumes less than Lyapunov).

Ex. 2.76. *The Lindeberg theorems: nitty-gritty details.* The essential story, regarding Lyapunov and Lindeberg, has been told in the previous exercise. Here we tend to the smaller-level but nevertheless crucial remaining details, in order for the ball to be shoven across the finishing line after all the preliminary work. Again, let $X_1, X_2, \ldots$ be independent, with distributions $F_1, F_2, \ldots$, zero means, standard deviations $\sigma_1, \sigma_2, \ldots$, and characteristic functions $\varphi_1, \varphi_2, \ldots$. The crucial random variable studied is

$$Z_n = \frac{X_1 + \cdots + X_n}{(\sigma_1^2 + \cdots + \sigma_n^2)^{1/2}} = \sum_{i=1}^n \frac{X_i}{B_n},$$

with $B_n^2 = \sum_{i=1}^n \sigma_i^2$. We assume that the Lindeberg condition holds, i.e. (2.13) is true.

(a). From Ex. 2.32 we have that $\max_{i\leq n}(\sigma_i^2/B_n^2) \to 0$, from the Lindeberg condition. Show that

$$\begin{aligned} |\varphi_i(t/B_n) - 1| &\leq \int |\exp(itx/B_n) - 1 - itx/B_n| \,\mathrm{d}F_i(x) \\ &\leq \tfrac{1}{2}t^2 \int (x/B_n)^2 \,\mathrm{d}F_i(x) \leq \tfrac{1}{2}t^2 \max_{i\leq n}(\sigma_i^2/B_n^2), \end{aligned}$$

so all $\varphi_i(t/B_n)$ are eventually inside radius say $\frac{1}{2}$ of 1. We are hence in a position to take the logarithm of $\kappa_n(t) = \mathrm{E}\exp(itZ_n)$, and work with $\log\kappa_n(t) = \sum_{i=1}^n \log\varphi_i(t/B_n)$, etc.; see the first lemma in Ex. 2.36.

(b). In continuation and refinement of arguments above, show that

$$r_n(t) = \varphi_i(t/B_n) - (1 - \tfrac{1}{2}\sigma_i^2 t^2/B_n^2),$$

can be bounded, as follows:

$$\begin{aligned}
|r_n(t)| &= \Big| \int \{\exp(itx/B_n) - 1 - itx/B_n - \tfrac12 (itx/B_n)^2\} \, \mathrm{d}F_i(x) \Big| \\
&\le \int |\exp(itx/B_n) - 1 - itx/B_n - \tfrac12 (itx/B_n)^2| \, \mathrm{d}F_i(x) \\
&\le \int_{|x|/B_n \le \varepsilon} \tfrac16 \frac{|t|^3 |x|^3}{B_n^3} \, \mathrm{d}F_i(x) + \int_{|x|/B_n > \varepsilon} \Big(\tfrac12 \frac{|t|^2 |x|^2}{B_n^2} + \tfrac12 \frac{|t|^2 |x|^2}{B_n^2} \Big) \, \mathrm{d}F_i(x) \\
&\le \tfrac16 |t|^3 \varepsilon \frac{\sigma_i^2}{B_n^2} + t^2 \, \mathrm{E} \Big| \frac{X_i}{B_n} \Big|^2 I\Big\{ \Big| \frac{X_i}{B_n} \Big| \ge \varepsilon \Big\}.
\end{aligned}$$

(c). Show that this leads to

$$\sum_{i=1}^n \big| \varphi_i(t/B_n) - (1 - \tfrac12 \sigma_i^2 t^2 / B_n^2) \big| \le \tfrac16 |t|^3 \varepsilon + t^2 \, L_n(\varepsilon),$$

with $L_n(\varepsilon)$ as defined in (2.13) (the Lindeberg condition), and by way of the first lemma in Ex. 2.36, that this secures what we were after, namely that

$$\prod_{i=1}^n \varphi_i(t/B_n) \to \exp(-\tfrac12 t^2),$$

and hence triumphantly $Z_n \to_d \mathrm{N}(0,1)$, under the Lindeberg condition only.

(d). Suppose that $\alpha > 0$ is so that $B_n / n^\alpha \to \sigma^2 > 0$, and that the Lindeberg condition

$$\frac{1}{n^\alpha} \sum_{i=1}^n \mathrm{E} \, X_i^2 I(|X_i| \ge n^{\alpha/2} \varepsilon) \to 0, \quad \text{for each } \varepsilon > 0,$$

holds. Show that then $n^{-\alpha/2} \sum_{i=1}^n X_i \to_d \mathrm{N}(0, \sigma^2)$.

Ex. 2.77. *Summing geometrically many terms.* Suppose $Y_1, Y_2, \ldots$ are i.i.d. with mean zero, variance σ^2, and m.g.f. $M_0(t)$. With $\Pr(N = n) = (1-p)^{n-1} p$ for $n \ge 1$, i.e. a geometric distribution, consider $Z_p = p^{1/2}(Y_1 + \cdots + Y_N)$, with the Y_i being independent of N.

(a). Show first that the generating function for N is $\mathrm{E} \, s^N = ps/\{1 - (1-p)s\}$ for $|s| < 1/(1-p)$; see Ex. 1.38. Show that Z_p has variance σ^2, and that its m.g.f. may be written

$$K_p(t) = \mathrm{E} \exp(tZ_p) = pM_0(p^{1/2}t)/\{1 - (1-p)M_0(p^{1/2}t)\}.$$

(b). Then use $M_0(t) = 1 + \frac12 \sigma^2 t^2 + o(|t|^2)$ to demonstrate that as $p \to 0$, with increasing number of terms $\mathrm{E} \, N = 1/p$, we have $K_p(t) \to 1/(1 - \frac12 \sigma^2 t^2)$ for $|t| < \sqrt{2}/\sigma$. This shows that $Z_p \to_d L_\sigma$, the Laplace distribution with standard deviation σ, see Ex. 1.33.

(c). For the particular case of a sum of randomly many normal terms, let the Y_i be i.i.d. standard normal. Show what $Z_p \,|\, N \sim \mathrm{N}(0, pN)$, and that $pN \to_d \mathrm{Expo}(1)$ as $p \to 0$. Explain how this matches Ex. 1.33.(c).

(d). (xx one or two further illustrations, with Laplace limit of $p^{1/2}(Y_1 + \cdots + Y_N)$. with $Y_i = Q_i - 1$, Poisson, we learn $p^{1/2}(V_N - N) \to_d L$, where $V_N \sim \mathrm{Pois}(N)$. similarly with $p^{1/2}(W_N - N)$, where $W_N \,|\, N \sim \chi^2_N$, randomly many degrees of freedom. xx)

(e). (xx something to simulate, to illustrate the cusp behaviour at the centre. with the Y_i having mean ξ and variance 1, we have

$$p^{1/2}(Y_1-\xi+\cdots+Y_N-\xi)=p^{1/2}N(\bar Y-\xi)=(pN)^{1/2}\,N^{1/2}(\bar Y-\xi)\to_d L_1.$$

this gives different inference for ξ, and different predictions for $\bar Y$, than what we're used to from normal terrain. we're also in scale mixtures of normal terrain, with random variance tending to a unit exponential. i'll look for variations. i do like the $pN\to_d \text{Expo}(1)$, since it gives the cool cusp in the limit for Z_p, but other variations for $pN\to_d V$ are ok too. xx)

(f). (xx just ranting away a bit until things settle. xx) More generally, with $Y_1, Y_2, \ldots$ being i.i.d. with zero mean and unit variance, consider

$$Z_p=p^{1/2}(Y_1+\cdots+Y_N)=(pN)^{1/2}N^{1/2}\bar Y_N,$$

with N having a distribution such that $pN\to_d V$, say, as $p\to 0$. From the CLT, $N^{1/2}\bar Y_N\to_d \text{N}(0,1)$, so this amounts to a situation with a normal limit, but a random variance, as in $X\,|\,V\sim\text{N}(0,V)$ but V random. Here $\text{E}\exp(tX)=\text{E}\exp(\frac12 t^2V)=M_0(\frac12 t^2)$, where $M_0(u)=\text{E}\exp(uV)$ is the mgf for V. Also, X has density

$$\bar f(x)=\int \phi(x/v^{1/2})(1/v^{1/2})\,\text{d}H(v),$$

with H the distribution of V. – The special case above amounts to $N\sim\text{geom}(1/p)$, where pN tends to the unit exponential, and where X gets the Laplace distribution. For another case, consider $N\,|\,\lambda\sim\text{Pois}(\lambda/p)$, and with λ having its own distribution with mean and variance λ_0 and τ_0^2, say. Show that $\text{E}\,pN=\lambda_0$ and that $\text{Var}\,pN\to\tau_0^2$. Also consider the special case for this setup where $\lambda\sim\text{Gam}(a,b)$. From

$$\text{E}\,\{\exp(-spN)\,|\,\lambda\}=\exp[-(\lambda/p)\{1-\exp(-sp)\}]$$

deduce

$$\text{E}\exp(-spN)=\frac{1}{[1+(1/b)(1/p)\{1-\exp(-sp)\}]^a}\to\frac{1}{(1+s/b)^a},$$

and that $pN\to_d V_{a,b}$, another $\text{Gam}(a,b)$. With the construction above, the normal scale mixture variable X has mgf $1/(1-\frac12 t^2/b)^a$, and density

$$\bar f(x)=\int \phi(x/v^{1/2})(1/v^{1/2})\frac{b^a}{\Gamma(a)}v^{a-1}\exp(-bv)\,\text{d}v.$$

This is a Laplace, for $a=1$. (xx but different, and interesting, for other a. round this off. xx)

(g). (xx i think we can use these tools to form a full Laplace process, for BNP use or otherwise. we should tune in to a $Z_p(t)=p^{1/2}(Y_1+\cdots+Y_{N(t)})$, with a clever $N(t)$. will look at $N(t)$ being a negative binomial process with mean t/p. nils thinks this works: (i) $\lambda\sim\text{Gam}(1,b)$, with mean $\lambda_0=1/n$. (ii) $N\,|\,\lambda\sim\text{Pois}(\lambda/p)$. write down $\text{E}\,(pN\,|\,\lambda)$ and $\text{Var}\,(pN\,|\,\lambda)$, then unconditional mean and variance for pN. (iii) find limit in distribution of pN. (iv) study $Z_p=p^{1/2}\sum_{i=1}^N Y_i$, close to $(pN)^{1/2}$ times a normal, etc. (v) make it into a full Laplace process, by having $\lambda_t\sim\text{Gam}(1,b_t)$. xx)

Ex. 2.78. *Local asymptotics.* The CLT and Lindeberg machineries yield normal limits and hence approximations in situations where independent observations come from given models. It is sometimes useful to extend such results to situations where observations stem from distributions close to, but not equal to, the postulated start models. The standard $\sqrt{n}$ speed of convergence for the CLT and relatives leads naturally to the notion of $O(1/\sqrt{n})$ neighbourhoods. If there is limiting zero-mean normality of variables like $\sqrt{n}(\widehat{\theta}-\theta_0)$, under a relevant null model at θ_0, then such variables typically have limiting non-zero-mean normal limits at such $O(1/\sqrt{n})$ alternatives.

(a). A simple setup illustrating such ideas is the following. Suppose $X_1,\ldots,X_n$ are i.i.d. from a distribution with mean $\xi+\delta/\sqrt{n}$ and variance $\sigma_n^2=\sigma^2+d/n$. Consider then $Z_n=\sqrt{n}(\bar{X}_n-\xi)$. Use the Lindeberg theorem, or a triangular version of the CLT, to demonstrate that $Z_n\to_d \mathrm{N}(\delta,\sigma^2)$.

(b). (xx for Ch. 5, an exercise with $\sqrt{n}(\widehat{\theta}-\theta_0)\to_d \mathrm{N}(b\delta, J^{-1})$, when data stem from $f(y,\theta_0)+\delta/\sqrt{n}h(y)$. natural special case: $f(y,\theta_0,\gamma_0\delta/\sqrt{n})$. xx)

(c). xx

Ex. 2.79. *Approximate normality when combining information sources.* (xx this is a nils rant, so far. it needs intro sentences. the point is partly that yes, Lindeberg gives us limiting normality of sums, but we also need consistent variance estimators. xx) To illustrate the general themes, in a situation exhibiting these general components, consider the following setup. There are Poisson parameters $\theta_1,\ldots,\theta_k$, with associated independent Poisson observations $y_j,\ldots,y_{j,m_j}$ for θ_j, leading to $\widehat{\theta}_j=\bar{y}_j=(1/m_j)\sum_{\ell=1}^{m_j}y_{i,\ell}$. The object is to make inference for the linear combination $\mu=a^{\mathrm{t}}\theta=\sum_{j=1}^k a_j\theta_j$, for which we use the estimator $\widehat{\mu}=a^{\mathrm{t}}\widehat{\theta}=\sum_{j=1}^k a_j\widehat{\theta}_j$, with variance

$$B_k^2=\mathrm{Var}\,\widehat{\mu}=\sum_{j=1}^k a_j^2\theta_j/m_j.$$

(a). For $Y\sim\mathrm{Pois}(\theta)$, show that $\mathrm{E}\,(Y-\theta)^3=\theta\ldots$, and that this implies that its skewness is $1/\theta^{1/2}$. Show also that with $Y_1,\ldots,Y_m$ i.i.d. from this distribution, we have $\mathrm{E}\,(\bar{Y}-\theta)^3=\theta/m$, with $\mathrm{skew}(\bar{Y})=1/(m\theta)^{1/2}$. Thus the skewness tends to zero, indicating limiting normality, as long as with θ, or m, or both, grow.

(b). Show furthermore that

$$\mathrm{skew}(\widehat{\mu})=\mathrm{E}\Big(\frac{\widehat{\mu}-\mu}{B_k}\Big)^3=\frac{\sum_{j=1}^k a_j^3\theta_j/m_j}{(\sum_{j=1}^k a_j^2\theta_j/m_j)^{3/2}}.$$

(xx then some Lindeberg things here, understanding when this tends to zero, leading to $Z_{k,0}=(\widehat{\mu}-\mu)/B_k\to_d \mathrm{N}(0,1)$. play a bit with a_j,m_j. xx)

(c). (xx then wish to find a case where variance is not well enough estimated. xx) We estimate the variance using $\widehat{B}_k^2=\sum_{j=1}^k a_j^2\widehat{\theta}_j/m_j$. To make inference for μ we need not merely the result of (b), but also relative consistency of the variance estimate. Show that $V_k=\widehat{B}_k^2/B_k^2$ has mean 1 and variance

$$\mathrm{Var}\,V_k=\frac{\sum_{j=1}^k a_j^4\theta_j/m_j^2}{(\sum_{j=1}^k a_j^2\theta_j/m_j)^2}.$$

(xx rigging the game so that $Z_{k,0} \to_d \mathrm{N}(0,1)$, but not Z_k. As a special case to consider, take a common $m_j = m_0$ for all sample sizes, $a_j = j$, and assume $\theta_j = 1/j$. What happens to B_k, $\widehat{B}_k$, V_k, and the natural ratio $Z_k = (\widehat{\mu} - \mu)/\widehat{B}_k$? xx)

(d). (xx then find the typical behaviour of V_k, to ensure also $Z_n = (\widehat{\mu} - \mu)/\widehat{B}_k \to_d \mathrm{N}(0,1)$. make connections to chapter 4 stuff on deviance and wilks. the Wilks thing is close to Z_n^2. xx)

Ex. 2.80. *Limiting normality of the sample variance matrix.* (xx can be better placed, inside Ch 3. results are used for Ex. **??**. xx) Consider i.i.d. vectors $Y_1, \ldots, Y_n$ from the multinormal $\mathrm{N}_p(\xi, \Sigma)$, first with known mean vector ξ, which we for convenience then set to zero. The estimated variance matrix is $\widehat{\Sigma} = (1/n) \sum_{i=1}^n Y_i Y_i^{\mathrm{t}}$.

(a). Write $\sigma_{j,k}$ for the elements of the $p \times p$ matrix Σ, and σ_j^2 for $\sigma_{j,j}$, the variance of component j of Y_i. Show that its estimator is $\widehat{\sigma}_j^2 = (1/n) \sum_{i=1}^n Y_{i,j}^2$, and that it has the distribution of $\sigma_j^2 \chi_n^2 / n$. Show also that $\sqrt{n}(\widehat{\sigma}_{j,j} - \sigma_{j,j}) \to M_{j,j}$, with this limit having the $\mathrm{N}(0, 2\sigma_j^4)$ distribution.

(b). Using first the one-dimensional CLT, show that $\sqrt{n}(\widehat{\sigma}_{j,k} - \sigma_{j,k})$ has a normal limit $M_{j,k}$, and find its variance.

(c). Then show that there is convergence in distribution of the full matrix, say $\sqrt{n}(\widehat{\Sigma} - \Sigma) \to_d M$, with $M = (M_{i,j})_{i,j=1,\ldots,p}$ multinormal with zero means, and that

$$\mathrm{cov}(M_{i,j}, M_{k,l}) = \sigma_{i,k}\sigma_{j,l} + \sigma_{i,l}\sigma_{k,j}.$$

(d). Assume Σ has full rank p. Show that limiting normality for $\widehat{\Sigma}$ implies limiting normality for $\widehat{\Sigma}^{-1}$, and that in fact $\sqrt{n}(\widehat{\Sigma}^{-1} - \Sigma^{-1}) \to_d M^* = -\Sigma^{-1} M \Sigma^{-1}$. Writing $\sigma^{j,k}$ for the elements of Σ^{-1}, show that $\mathrm{cov}(M_{i,j}^*, M_{k,l}^*) = \sigma^{i,k}\sigma^{j,l} + \sigma^{i,l}\sigma^{j,k}$.

(e). In case of an unknown mean vector, one uses the sample variance matrix $\widetilde{\Sigma} = (n-1)^{-1} \sum_{i=1}^n (Y_i - \bar{Y})(Y_i - \bar{Y})^{\mathrm{t}}$. Show that $\sqrt{n}(\widetilde{\Sigma} - \widehat{\Sigma}) \to_{\mathrm{pr}} 0$, where $\widehat{\Sigma} = n^{-1} \sum_{i=1}^n (Y_i - \xi)(Y_i - \xi)^{\mathrm{t}}$ uses the ξ. Deduce that $\sqrt{n}(\widetilde{\Sigma} - \Sigma) \to_d M$ and $\sqrt{n}(\widetilde{\Sigma}^{-1} - \Sigma^{-1}) \to_d M^* = -\Sigma^{-1} M \Sigma^{-1}$, i.e. with the same limits as above.

(f). (xx something to round if off. perhaps dimension 2. mention the Wishart distribution, but here we derive limits without knowing or using that. also, not lost at sea outside multinormality, but the covariance structure for the limit M becomes much more complicated. xx)

Ex. 2.81. *Maximum sample value of exponentials and the Gumbel distribution.* (xx nils reorganises some of these exercises which need tidying up. i now start with the gumbel and the maximum of exponentials, before taking up other gumbel related matters. xx)

the Gumbel distribution

We define the *Gumbel distribution* on the real line by its cumulative distribution function $G_0(u) = \exp\{-\exp(-u)\}$.

(a). Show that G_0 is indeed a cumulative distribution function, that its density is $g_0(u) = \exp\{-u - \exp(-u)\}$, and that its Laplace transform is $L_0(s) = \mathrm{E} \exp(-sU) = \Gamma(1+s)$, in terms of the gamma function.

the Euler–Mascheroni constant

(b). Use properties of the gamma function to show that the mean and variance of the Gumbel distribution is γ_e and $\pi^2/6$, where $\gamma_e = 0.5772...$ is the Euler constant. The latter has several equivalent definitions, among which is that $H_n - \log n \to \gamma_e$, where $H_n = 1 + 1/2 + \cdots + 1/n$ is the partial sum of the divergent harmonic series.

(c). With U having the Gumbel distribution, show also that its mode is 0 and that its median is $-\log(\log 2) = 0.367$. Find an expression for the q-quantile $G_0^{-1}(q)$. Show that $\Pr(-1.097 \le U \le 2.970) = 0.90$.

(d). The Gumbel distribution turns up in various contexts concerning extreme values. The simplest such case is as follows: let $X_1, \ldots, X_n$ be i.i.d. from the unit exponential model, with $M_n = \max_{i \le n} X_i$ their maximum. Show that $M_n - \log n \to_d U$, the Gumbel distribution.

(e). We may learn more about the distribution of M_n via first investigating the spacings. With $X_{(1)} < \cdots < X_{(n)}$ being the order statistics, let $D_1 = X_{(1)}$, $D_2 = X_{(2)} - X_{(1)}$, up to $D_n = X_{(n)} - X_{(n-1)}$. We have seen via Ex. 1.12 and 1.15 that the spacings are independent (for this special case of the exponential), with $D_i \sim \text{Expo}(n - i + 1)$. Show that this leads to the representation

$$X_{(n)} = D_1 + \cdots + D_n = V_1/1 + V_2/2 + \cdots + V_n/n,$$

with $V_1, \ldots, V_n$ being i.i.d. and unit exponential.

(f). Show from this that M_n has mean $H_n \doteq \log n + \gamma_e$ and variance $\sum_{i=1}^n 1/i^2$, tending to $\pi^2/6$. This is agreement with the Gumbel limit for $M_n - \log n$.

(g). Show that $M_n - H_n \to_d U - \gamma_e$, the zero-mean version of the Gumbel. Deduce from this that

$$\lim_{n\to\infty} \text{E} \exp\{-s(M_n - H_n)\} = \prod_{i=1}^{\infty} \frac{\exp(s/i)}{1 + s/i} = \Gamma(1+s)\exp(\gamma_e s).$$

This infinite-product form of the gamma function is actually equivalent to a famous formula by Weierstraß. Show also that

$$\sum_{i=1}^{\infty} \{s/i - \log(1 + s/i)\} = \gamma_e s + \log\Gamma(1+s) = \sum_{j=2}^{\infty} (-1)^j \frac{\zeta(j)}{j} s^j,$$

valid for $|s| < 1$, where $\zeta(j) = 1 + 1/2^j + 1/3^j + \cdots$ is Riemann's zeta function, at j. (xx two more sentences. here we derive these deep mathematical facts from a simple convergence in distribution result; could also go the other way, if we start with gamma function knowledge. xx)

(h). Yet another fruitful perspective on what we've learned above is in terms of an infinite sum of smaller and smaller exponentials. Consider independent exponentials $W_1, W_2, \ldots$, where $W_i \sim \text{Expo}(i)$, i.e. with mean $1/i$. Show that $W = \sum_{i=1}^{\infty}(W_i - 1/i)$ is finite, with probability one, and that its distribution is that of $U - \gamma_e$, the zero-mean Gumbel.

Ex. 2.82. *Weibull and Gamma maxima.* The basic result of Ex. 2.81, concerning the maximum of a sample of exponentials, leads to limit distribution results also for maxima from other distributions.

(a). Suppose that $X_1, \ldots, X_n$ are i.i.d. from the Weibull distribution, with cumulative function $F(x) = 1 - \exp\{-(x/a)^b\}$, for certain parameters (a, b). With M_m the sample maximum, show that $(M_n/a)^b - \log n$ tends to the Gumbel distribution.

(b). In two minutes, simulate $n = 1000$ values from the Weibull with $(a, b) = (1, \frac{1}{2})$. Guess in advance how large M_n will be, using the representation $M_n = a(\log n + U_n)^{1/b}$, where U_n tends to the Gumbel.

(c). Similarly consider the Gamma distribution with parameters $(a, b) = (2, 1)$, where the cumulative can be expressed as $F(x) = 1 - \exp(-x)(1+x)$, see Ex. 1.9. With M_n the maximum of a sample of size n from this distribution, show that $M_n - \log(1+M_n) - \log n$ tends to the Gumbel distribution. What is the approximate median for the M_n?

(d). More generally, suppose $F(x) = 1 - \exp\{-A(x)\}$, with $A(x)$ the cumulative hazard rate, and let again M_n be the maximum value from a sample of size n. Show that $A(M_n) - \log n$ tends to the Gumbel.

Ex. 2.83. *Maximum of independent geometric variables.* Let T have the geometric waiting time distribution with parameter p, i.e. with point probabilities $(1-p)^{t-1}p$ for $t = 1, 2, \ldots$. We write $T \sim \text{geom}(p)$ to indicate this distribution; see Ex. 1.25.

(a). Show that $V = pT$ has mean 1 and variance $1 - p$. Show also that if $p \to 0$, then $V = pT$ tends to the unit exponential in distribution. Give an approximate formula for the median of a geometric distribution with small p.

(b). Now suppose $V_1, \ldots, V_n$ are independent geometric waiting times with parameter $1/n$, hence with mean value n. With $Z_n = \max(V_1, \ldots, V_n)$ the time until all waiting times have been completed, we then have $Z_n/n = \max(V_1/n, \ldots, V_n/n)$, which is close in distribution to $\max(E_1, \ldots, E_n)$, with these E_i being independent unit exponentials. By results of Ex. 2.81 we should there expect $Z_n/n - \log n \to_d U$, the Gumbel. Show that indeed this holds. (xx needs some technicalities and a hint, see nils diehard. xx)

(c). (xx one more thing here. can we make something useful out of this, with mgf for Z_n/n. not easy. xx) we do have

$$\Pr(Z_n/n \le v) = \Pr(V_i \le nv)^n = [1 - (1 - 1/n)^{nv}]^n.$$

Ex. 2.84. *Collecting cards: how long time?* (xx nils will reorganise this a bit, after the abels taarn things. plan is basic things $T_1 + \cdots + T_n$ here, Gumbel limit, a bit more in next exercise, this being Ch4. then likelihood things in Story iv.6 about estimating n from $V_r = T_1 + \cdots + T_r$, time to having seen r different cards. then story about estimating n from observed V_r. xx) Consider a deck of n cards, with $X_1, X_2, \ldots$ independent draws from these, i.e. uniform on $\{1, \ldots, n\}$. How many such random draws are necessary, before you have seen all n cards? – There are several reformulations of this card collecting problem, and with other metaphors. You may think of a fair die, with n faces, and ask how many times you need to roll it until you've seen all faces.

(a). Show that the time needed, until we have seen all n cards, can be represented as $V_n = T_1 + \cdots + T_n$, where T_i is geometric with parameter $p_i = (n - i + 1)/n$. Hence $\mathrm{E}\, T_i = n/(n - i + 1)$, and the card finding process is easy in the beginning, then steadily harder. We may also re-order the T_i to $V_n = T'_1 + \cdots + T'_n$, where $T'_i \sim \text{geom}(i/n)$, which for some purposes is an easier representation.

(b). Let $(N_1, \ldots, N_n)$ be the number of times cards $1, \ldots, n$ have been seen, in the course of z independent random draws from the deck. Show that this is a multiomial with count z and probabilities $(1/n, \ldots, 1/n)$; in particular, $N_i \sim \text{binom}(z, 1/n)$. Show that the correlation between N_i and N_j is $-1/(n-1)$.

(c). Show also that another representation of V_n is as $\max(W_1, \ldots, W_n)$, where W_i is the first time $N_i \geq 1$. Show that $W_i \sim \text{geom}(1/n)$, with mean n. These are however dependent, so the Gumbel limit result of Ex. 2.83 does not immediately apply. Show that the correlation between W_i and W_j is small, however, namely $-\frac{1}{2}/(n-1)$, indicating that $(V_n - n\log n)/n$ should converge to the Gumbel, even with these waiting times being dependent. (xx polish wording here. xx)

(d). (xx nils will coordinate and calibrate this with what is placed in Ex. 1.25. xx) For T_i, with distribution $(1-p_i)^{t-1}p_i$ for $t = 1, 2, 3, \ldots$, show that

$$\mathrm{E}\, T_i = \frac{1}{p_i}, \quad \mathrm{Var}\, T_i = \frac{1-p_i}{p_i^2}, \quad \mathrm{E}\,(T_i - 1/p_i)^3 = \frac{(1-p_i)(2-p_i)}{p_i^3},$$

so the skewness of T_i is $\mathrm{E}\,(T_i - 1/p_i)^3/\sigma_i^3 = (2-p_i)/(1-p_i)^{1/2}$.

(e). Show that

$$\mathrm{E}\, V_n = n(1 + 1/2 + \cdots + 1/n) = nH_n \doteq n(\gamma + \log n),$$

using Ex. 2.81. Show also that

$$\mathrm{Var}\, V_n = \sum_{i=1}^n \Big(\frac{n^2}{i^2} - \frac{n}{i} \Big) \doteq n^2(\pi^2/6) - n(\gamma + \log n).$$

(f). (xx limit of skewness. not zero. xx) Now consider

$$U_n = \frac{V_n - \mathrm{E}\, V_n}{(\mathrm{Var}\, v_n)^{1/2}}, \quad U_{n,0} = \frac{V_n - n(\gamma + \log n)}{n\pi/\sqrt{6}},$$

and show that $U_n - U_{n,0} \to_{\mathrm{pr}} 0$. Show further that

$$\mathrm{E}\, U_n^3 = \frac{\sum_{i=1}^n \mathrm{E}\,(T_i - p_i)^3}{(\mathrm{Var}\, Z_n)^{3/2}} \doteq \frac{2n^3 \sum_{i=1}^n (1/i^3) + O(n^2)}{n^3\pi^3/6^{3/2}} \to \frac{2 \cdot 1.2021}{\pi^3/6^{3/2}} = 1.1396.$$

(g). So we're outside limiting normality; show indeed that the Lindeberg condition cannot hold here. (xx limit distribution. other things. xx)

(h). (xx check this. xx) With $U_n' = (V_n - n\log n)/n = \bar{T}_n - \log n$, show that

$$\begin{aligned} \mathrm{E}\, \exp(-sU_n') &= \exp(s\log n) \prod_{i=1}^n \mathrm{E}\, \exp\{-(s/n)T_i\} \\ &= n^s \prod_{i=1}^n \frac{(i/n)\exp(-s/n)}{1 - (1-i/n)\exp(-s/n)} \\ &= n^s \frac{(n!/n^n)\exp(-s)}{\prod_{i=1}^n \{1 - (i/n)\exp(-s/n)\}}. \end{aligned}$$

Then show that $U_n' \to_d U$. (xx hmm, have not landed this properly yet, but can be cool story. and if we prove $U_n' \to_d U$ in some other way, we are automatically deriving the side consequence

$$A_n(s) = \prod_{i=1}^n \{1 - (i/n)\exp(-s/n)\} \doteq \frac{n^s \exp(-n-s)(2\pi n)^{1/2}}{\Gamma(1+s)},$$

or

$$\prod_{i=1}^{n}\{\exp(s/n)-(i/n)\} \doteq \frac{n^s\exp(-n)(2\pi n)^{1/2}}{\Gamma(1+s)},$$

which for $s=0$ is Stirling. need a bit more work. xx)

(i). It is also useful to find the distribution of $G_n(v)=\Pr(V_n\le v)$ explicitly. Argue that $V_n\le v$ is equivalent to $A_1\cap\cdots\cap A_n$, where A_i is the event that i is seen in the course of the first z attempts. With $B_i=A_i^c$ its complement, that i has not been seen during these first v attempts. Use this to deduce that $1-G_n(v)$ can be written

$$\begin{aligned}\Pr(B_1\cup\cdots\cup B_n) &= \binom{n}{1}\Pr(B_1)-\binom{n}{2}\Pr(B_1\cap B_2)+\binom{n}{3}\Pr(B_1\cap B_2\cap B_3)-\ldots\\ &= \sum_{j=1}^{n}(-1)^{j-1}\binom{n}{j}(1-j/n)^v,\end{aligned}$$

for $v\ge n$. Use algebra to also derive

$$g_n(v)=\Pr(V_n=v)=\sum_{j=1}^{n}(-1)^{j-1}\binom{n-1}{j-1}(1-j/n)^{v-1}\quad\text{for } v\ge n.$$

Use $x^{n-1}+x^n+\cdots=x^{n-1}/(1-x)$ for $|x|<1$ to derive the identity

$$\sum_{j=1}^{n-1}(-1)^{j-1}\binom{n}{j}(1-j/n)^{n-1}=1.$$

(j). Use the case $(T_1=1,\ldots,T_n=1)$ to derive

$$\prod_{i=1}^{n}(i/n)=\frac{n!}{n^n}=\sum_{j=1}^{n-1}(-1)^{j-1}\binom{n-1}{j-1}(1-j/n)^{n-1},$$

and argue that these expressions are close to $\exp(-n)(2\pi n)^{1/2}$, by the Stirling approximation. Show via arrangements of this formula that $\sum_{j=0}^{n}(-1)^j\binom{n}{j}j^n=n!$. (xx $-4\cdot1+6\cdot16-4\cdot81+256=24$, etc. xx)

(k). (xx pointer here to a different story, where we estimate n based on how long time it took us to reach level r, i.e. $W_r=T_1+\cdots+T_r$. it might be a CD story with $C_r(n)=\Pr_n(W_r<W_{r,\mathrm{obs}})+\frac12\Pr_n(W_r=W_{r,\mathrm{obs}})$. how many Italians in my neighbourhood? xx)

Ex. 2.85. *The 2nd largest, 3rd largest, etc., for exponentials.* Let as in Ex. 2.81 $X_1,\ldots,X_n$ be i.i.d. from the unit exponential model. For the largest observation we saw there that $X_{(n)}-\log n\to U$, the Gumbel distribution with c.d.f. $\exp\{-\exp(-u)\}$. Here we shall work with the 2nd largest, the 3rd largest, etc.

(a). For a positive, consider W_a, with density

$$g_a(w)=\Gamma(a)^{-1}\exp\{-aw-\exp(-w)\} \tag{2.14}$$

on the real line. Show that $V_a=\exp(-W_a)$ has the gamma distribution with parameters $(a,1)$, and that the Laplace transform becomes $\mathrm{E}\exp(-tW_a)=\Gamma(a+t)/\Gamma(a)$. The Gumbel distribution is the case of $a=1$, so we may consider (2.14) a generalised Gumbel.

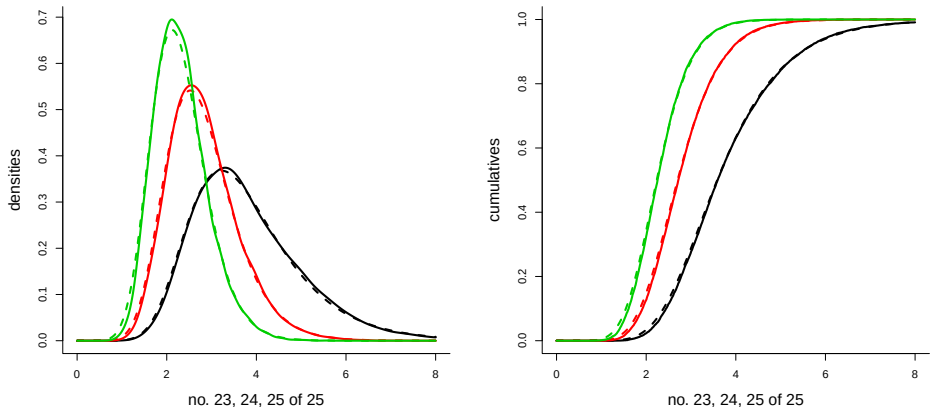


Figure 2.3: *For $n = 25$, the dashed curves are limit-approximation densities (left panel) and cumulatives (right panel) for order statistics 23, 24, 25 for i.i.d. exponentials, computed via the limits $\log n + W_i$ for $i = 1, 2, 3$. The full curves are the real densities and cumulatives, based on having simulated 10^4 outcomes for each.*

(b)**.** Deduce that $W_a = \log(1/V_a)$ has mean $-\psi(a)$ and variance $\psi'(a)$, with $\psi(a) = \partial \log \Gamma(a)/\partial a$ the digamma function. Show that the c.d.f.s for g_1, g_2, g_3 are

$$\begin{aligned} G_1(w) &= \exp\{-w - \exp(-w)\}, \\ G_2(w) &= \exp\{-w - \exp(-w)\}\{1 + \exp(-w)\}, \\ G_3(w) &= \exp\{-w - \exp(-w)\}\{1 + \exp(-w) + \tfrac{1}{2}\exp(-2w)\}. \end{aligned}$$

(c)**.** With order statistics $X_{(1)} < \cdots < X_{(n)}$, consider $W_{n,i} = X_{(n-i+1)} - \log n$, for given i; the case $i = 1$ is $W_{n,1} = X_{(n)} - \log n$ already considered in Ex. 2.81. Show that

$$\Pr(X_{(n-i+1)} - \log n \le w) = \Pr(U_{(n-i+1)} \le 1 - (1/n)\exp(-w)),$$

in terms of the order statistics for the uniform. Use the Beta connection of Ex. 3.17 to deduce that the density of $W_{n,i}$ may be written

$$g_{n,i}(w) = \text{be}(1 - (1/n)\exp(-w), n - i + 1, i)(1/n)\exp(-w)$$

in terms of the Beta density with parameters $(n - i + 1, i)$. Take the limit to prove that $W_{n,i} \to_d W_i$.

(d)**.** Construct a version of Figure 2.3, showing that the approximations based on the limit distributions work well already for $n = 25$. The dashed curves are using the limits, with $x_{(n-i+1)} = \log n + W_i$, wheras the dashed curves are the real densities and cumulatives, obtained by 10^4 simulations from $X_{(n-2)}, X_{(n-1)}, X_{(n)}$. (xx these are kernel estimates so need a pointer to Ch. 13. can we make a story out of this, insurance company cares for these most extreme outcomes. xx)

Ex. 2.86. *How many records are set?* Consider i.i.d. observations $X_1, X_2, \ldots$ from some continuous distribution on the line. Define $R_n = 1$ if X_n is bigger than all previous datapoints, i.e. a new record has been set.

(a). Explain first that the number of records set, in the course of the first n occasions, is $Z_n = \sum_{i=1}^n R_i$. Show that $\Pr(R_n = 1) = 1/n$, so that $\mathrm{E}\, Z_n = H_n = 1+1/2+\cdots+1/n$, the partial sum of the harmonic series. As we have noted earlier, $H_n - \log n \to \gamma_e = 0.5772...$, the Euler constant, so the number of records set is approximately $\log n + \gamma_e$.

(b). Show that $R_1, \ldots, R_n$ are independent, and deduce from this that $\mathrm{Var}\, Z_n = H_n - \sum_{i=1}^n 1/i^2 \doteq \log n + \gamma_e - \pi^2/6$. Use the Lindeberg theorem to show that $(Z_n - H_n)/H_n^{1/2} \to_d \mathrm{N}(0,1)$. Show also that $(Z_n - \log n)/(\log n)^{1/2} \to_d \mathrm{N}(0,1)$, which gives a simpler but somewhat more crude approximation to Z_n probabilities. Use these normal limits to give a prediction band, with $\Pr(a_n \le Z_n \le b_n) \to 0.95$.

(c). Now consider $A_n = Z_{2n} - Z_n$, the number of records set during $n, n+1, \ldots, 2n$. Show that A_n tends to the Poisson with mean $\log 2$. Generalise.

Ex. 2.87. *When the 00 box is hidden.* Consider a 2×2 table setup with counts $N_{0,0}, N_{0,1}, N_{1,0}, N_{1,1}$, corresponding to $N_{i,j}$ counting the cases of $(X = i, Y = j)$, for $i, j = 0, 1$, for two factors X and Y. We take the four counts to be a multinomial vector with probabilities $p_{0,0}, p_{0,1}, p_{1,0}, p_{1,1}$. Assume now that the 00 box is hidden, hence also the total number $N = N_{0,0} + N_{0,1} + N_{1,0} + N_{1,1}$; one has observed counts $N_{0,1}, N_{1,0}, N_{1,1}$, but not the $N_{0,0}$ in question. How can one estimate the hidden $N_{0,0}$, and then in its turn N? Such questions and estimation methods go back to Petersen (1896), who needed to estimate fish populations in the 1890ies. Versions of methods methods developed below is being used in Story iii.13 to estimate the number of killed persons in Srebrenica 1995. See also Ex. 5.62 for generalisations.

(a). Assume in this exercise that factors X and Y are independent, with $\Pr(X = 1) = p = p_{1,\cdot}$ and $\Pr(Y = 1) = q = p_{\cdot,1}$; we use '$\cdot$' notation to indicate that the index in question is being summed over. Show that

$$p_{0,0} = (1-p)(1-q), \quad p_{0,1} = (1-p)q, \quad p_{1,0} = p(1-q), \quad p_{1,1} = pq.$$

(b). Argue that $N_{1,\cdot}N_{\cdot,1}/N^2$ and $N_{1,1}/N$ are both valid estimates of $p_{1,1}$. Discuss conditions under which $N^* = N_{1,\cdot}N_{\cdot,1}/N_{1,1}$ is a reasonable estimator of N.

(c). The N is unknown, but we may still study the usual ratios $\widehat{p}_{i,j} = N_{i,j}/N$. Show that there is joint convergence in distribution, say $N^{1/2}(\widehat{p}_{i,j} - p_{i,j}) \to_d A_{i,j}$, as N increases, with the $A_{i,j}$ forming a four-dimensional mean zero normal. Give its variance matrix. Under independence, show that $(N^* - N)/N^{1/2} = N^{1/2}(N^*/N - 1)$ has limit distribution

$$\begin{aligned} U &= (1/p)(A_{1,0} + A_{1,1}) + (1/q)(A_{0,1} + A_{1,1}) - \{1/(pq)\}A_{1,1} \\ &= \{pA_{0,1} + qA_{1,0} + (p+q-1)A_{1,1}\}/(pq). \end{aligned}$$

This is a normal $(0, \tau^2)$; show that indeed $\tau^2 = (1-p)(1-q)/(pq)$. How can this be used to form a confidence interval for N? (More general schemes, for estimators and confidence intervals, are developed in See Ex. 5.62.)

(d). Show that the N^* leads to the natural estimator $\widehat{p} = N_{1,1}/N_{1,\cdot}$ for p. Find its approximate distribution, and assess how much is lost in precision by not knowing N. (xx check also with the implied $N^*_{0,0} = N^* - (N_{0,0} + N_{0,1} + N_{1,0})$. xx)

(e). The setup and methods above can be used in a variety of setups, for estimating the sizes of populations based on incomplete surveys; the N^* estimator above goes back

to Petersen (1896), estimating the number of fish based on capture-recapture surveys. Carry out a few simulation experiments, as follows. There are fish $\{1, \ldots, N\}$ in your pond. Your first catch, with fish being caught as in a binomial setup with probability p_1, gives the index set A_1; your captured fish are marked and released in the pond. Similarly your second catch, with catch probability p_2, gives index set A_2. By counting the numbers $N_{1,0}, N_{0,1}, N_{1,1}$ in the associated Venn diagram, estimate the total number of fish N (and your analysis should work without knowing p_1, p_2). Check if your 95 percent confidence interval captures the real N. To play with these methods, also to understand how the catch probabilities p_1, p_2 influence estimates and precision, there are helpful and easy-to-use algorithms in `R`, namely `intersect`, `union`, `setdiff`.

(f). (xx might bypass this point. but might include the case of three surveys, or leave it to $\ell_{\rm prof}(N)$ analysis in Ch 5. but we can ask for analysis of the estimator

$$n^*_{0,0,0} = \frac{n_{1,0,0}n_{0,1,0} + n_{1,0,0}n_{0,0,1} + n_{0,1,0}n_{0,0,1}}{n_{1,1,0} + n_{1,0,1} + n_{0,1,1}},$$

used in Lum et al. (2013). xx)

Notes and pointers

(xx to come. we point to certain famous things from the past: Kolmogorov (1933b), Lindeberg (1922), Borel and Cantelli. tail bounds. emil's extension of the Inlow (2010) paper, from CLT to Lindeberg. more on Lindeberg and the history of CLT developments in Cramér (1976), also see Schweder (1980, 1999). mention the classical text Billingsley (1968). xx)

(xx Weierstraß, Hveberg (2019), A generation or so later, such results had been generalised to so-called Stone–Weierstraß theorems, stating, in various forms, that certain classes of functions are rich enough to deliver uniform approximations to bigger classes of functions. This is useful also in branches of probability theory. xx)

(xx the first of the two lemmas: Nils beta 1990. Weierstrass: see also Hveberg (2019). our result xxxx if an extension of one by Inlow (2010). xx)

(xx include also something on Fra Preface i *Life and Times of CLT*: "For those who teach a course in probability whose objective is to prove the Central limit theorem of interest [...] is commentary on the characteristic function approach employed by Lyapunov versus the 'very simple' proof, as Le Cam describes it, given by Lindeberg". also, LeCam (1986). xx)

[xx for Scheffé: see Scheffé (1947), but also Kusolitsch (2010), who explains that the result is a special case of results published by F. Riesz in 1928. see also what Scheffe says in his paper about comments he got from Morse. xx]

(xx push this to Notes. xx) Inlow (2010) has shown how one can prove the usual CLT without the technical use of characteristic and hence complex functions. Essentially, he writes the X_i in question as $Y_i + Z_i$ with $Y_i = X_i\, I\{|X_i| \le \varepsilon\sqrt{n}\}$ and $Z_i = X_i\, \{|X_i| > \varepsilon\sqrt{n}\}$, after which 'ordinary' m.g.f.s may be used for the part involving the Y_i, yielding the normal limit, supplemented with analysis to show that the part involving the Z_i tends to zero in probability. – It is a non-trivial matter to extend Inlow's arguments, from the CLT to the Lindeberg theorem, but this is precisely what is done in Stoltenberg (2019). Check that note, on the book website, and make sure you understand its main tricks and steps.

(xx When Jarl Waldemar Lindeberg was reproached for not being sufficiently active in his scientific work, he said, 'Well, I am really a farmer'. And if somebody happened to say that his farm was not properly cultivated, his answer was, 'Of course my real job is to be a mathematics professor'. Hundred years ago!, i.e. in 1920, he published his first paper on the CLT, and in 1922 he generalised his findings to the classical Lindeberg Theorem, with the famous Lindeberg Condition, securing limiting normality of a sum of independent but not identically distributed random variables. He did not know about Ляпунов's earlier work, and therefore not about условие Ляпунова, the Lyapunov condition, which we treat below as a simpler-to-reach condition than the more general one of Lindeberg. Other luminaries whose work touch on these themes around the 1920ies and beyond include Paul Lévy, Harald Cramér, William Feller, and, intriguingly, Alan Turing who (allegedly) won the war and invented computers etc. xx)

(xx point to a couple of characterisation theorems books, kagan linnik rao, one more. xx)

(xx for Notes: The little $\log(1+x)$ lemma is stated, proven, and used in Hjort (1990b, Appendix). xx)

(xx the material is from Hjort and Pollard (1993) and Hjort (1986a). xx)

[xx For notes and pointers: Ex. 2.19(b) is from Jacod and Protter (2004, p. 166), they have it from Pollard (1984) xx].

ToDo notes, of 12-March-2025.

Clean and calibrate. Include a couple of classic nonparametric test procedures, like Wilcoxon, the sign test, more, to showcase the use of Lindeberg things to show limiting normality of such statistics too. point to Hjort and Pollard (1993) and Hjort (1986a).

Include some non-normal limits. Can take $\widehat{\mu}^* = \{1-c(D_n)\}\widehat{\mu}_{\rm narr} + c(D_n)\widehat{\mu}_{\rm wide}$ from model selection. And point to $n^{2/5}$ rates for f estimation.

I.3

Parameters, estimators, precision, confidence

With data observed from a statistical model, the theme of this chapter is that of constructing estimators for unknown statistical parameters, along with assessing their precision. This provides ways of comparing competing estimation methods. Basic concepts include the bias, the variance, the mean squared error of estimators. The development also naturally leads to the important notion and basic machinery of confidence intervals. General estimation methods covered here include the method of moments and the method of quantiles; these can also be combined. For regression setups, with response variables influenced by covariates, we go through the method of least squares. To understand and utilise the properties of classes of estimators in general models, we utilise the machinery of large-sample normal approximations, from Ch. 2. This also enables one to assess precision and to compare different competing estimators. Similar remarks apply also for the more versatile method of maximum likelihood, treated in Ch. 5.

Key words: approximations, confidence, estimators, linear regression, model parameters, moment matching, quantile matching, risk

Most statistical models have *parameters*, as we learn from the generous variety of models in Ch. 1. Parameters may then be fine-tuned, or estimated, from data, which is the grand theme of the present chapter. In generic terms, if a model has density $f(y,\theta)$, with $\theta = (\theta_1,\ldots,\theta_p)^{\rm t}$ its parameter vector, we use data $\mathcal{D}$ to construct *an estimator* $\widehat\theta = \widehat\theta(\mathcal{D})$. Thus $f(y,\widehat\theta)$ is the fitted model, which we use for interpretation and inference, themes we return to in more detail in later chapters. The data $\mathcal{D}$ can often be in the form of direct independent observations $y_1,\ldots,y_n$ from the model, but can also be different in character, involving censoring mechanisms, or measurement error.

A sensible minimum demand for an estimator is that it should tend to the right value, with increasing data volume. Formally, if $\widehat\theta = \widehat\theta_n$ is the estimator for some parameter whose true value is θ_0, based on a sample of size n, then we say that the estimator (or, more pedantically, the sequence of estimators) is *consistent* provided $\widehat\theta_n$ converges in probability to θ_0, i.e. $\widehat\theta_n \to_{\rm pr} \theta_0$ in the terminology of Ch. 2.

consistency of an estimator

One often needs estimates and inference methods for *focus parameters*, those of particular and context-driven interest, which are one-dimensional functions $\mu = \mu(\theta_1,\ldots,\theta_p)$ of the underlying model parameter vector. If $\widehat\mu$ is an estimator for this parameter, we often care about its mean, represented here as

focus parameter

$$\mathrm{E}_\theta\,\widehat\mu = \mu + b(\theta). \tag{3.1}$$

The footscript signals that the expectation operator is at work at the parameter position θ. The $b(\theta)$ is termed *the bias*; in various cases it is a function of μ only, but in general it depends on the full parameter vector θ. If $\mathrm{E}_\theta\,\widehat{\mu} = \mu$, at all positions θ, we say the estimator is *unbiased*. In addition to wishing for estimators with small bias, we care about its variability, and often about its *mean squared error*

unbiased estimator

mean squared error

$$\mathrm{mse}(\widehat{\mu},\theta) = \mathrm{E}_\theta\,(\widehat{\mu}-\mu)^2 = \mathrm{Var}_\theta\,\widehat{\mu} + b(\theta)^2, \tag{3.2}$$

the classic variance plus squared bias. This is a function of the unknown parameter, and gives a way of understanding and comparing performance for competing estimation schemes. When we can sort out the mathematics properly, depending on the situation at hand, we then choose estimators with smaller mse than those of competitors.

The $\mathrm{mse}(\widehat{\mu},\theta)$ of (3.2) is sometimes called the risk, or risk function, and relates to having chosen squared error $(\widehat{\mu}-\mu)^2$ as the underlying measure of quality. Other ways in which to compare and rank performance, involving also different quality functions and risk functions, will be dealt with in Ch. 8.

Fundamental and conspicuous instruments in the statistical toolkits, when summarising and reporting findings of investigations, are *confidence intervals* and *testing of null hypotheses*. The development in the present chapter deals with the former, regarding interpretation, construction, properties, performance, whereas the following chapter handles the latter, along with further connections.

Consider in general terms data y, perhaps a long vector or a data matrix, from a model with parameter $\theta = (\theta_1,\ldots,\theta_p)$, with $\mu = \mu(\theta_1,\ldots,\theta_p)$ a parameter of interest. Then $[L(y),U(y)]$ is a *confidence interval*, with confidence level α, like 0.90 or 0.95 or an even higher 0.99, provided

confidence interval

$$\Pr_\theta\{L(Y) \le \mu \le U(Y)\} = \alpha \quad \text{for all } \theta. \tag{3.3}$$

Thus $[L(Y),U(Y)]$ is a random interval, and with a high number of repeated situations it will capture the underlying μ a fraction α of the times. The reported confidence interval is $[L_{\mathrm{obs}},U_{\mathrm{obs}}] = [L(y_{\mathrm{obs}}),U(y_{\mathrm{obs}})]$, computed based on the actually observed data y_{obs}. Occasionally one also wants to construct *confidence regions* for a parameter vector, as opposed to confidence intervals for each of its components. Thus if (a,b) is a parameter pair in a model, a random region R, based on data, is a 95 percent confidence region provided $\Pr_\theta((a,b)\in R) = 0.95$ for all model parameters θ.

In various setups one can study distributions, biases, variances, confidence etc. quite accurately, as will be seen in many exercises below. Often enough this might be too complicated, however, and one relies instead on good approximations. There is indeed a host of normal approximations, sometimes with additional tools for finetuning these, as studied in Ch. 2, with yet more to come in Ch. 5. These methods may be understood, appreciated, seen in action, and used for new situations, without necessarily having been through each δ and ε of Ch. 2.

In this chapter we learn certain estimation principles, including those associated with the method of moments and the method of quantiles. There is also room for combining such methods, or for coming up with new estimators in unfamiliar waters. We go on to more advanced models and hence estimation methods in later chapters (and in several of our stories), but included below is the basics of linear regression and the least sum of squares methods. The more versatile and often well-performing method of *maximum likelihood* will be studied with care in Ch. 5.

[xx In this brief intro there should be a figure, conveying some basic ideas. We may snikinnføre confidence intervals, but that comes with more weight in Ch. 4, along with testing and power and p-values. we do mention a few key concepts here in intro, like unbiasedness, low variance, etc. xx]

Precise estimation in a few classical models

Ex. 3.1. *Mean squared error.* Suppose data lead to an estimator $\widehat{\mu}$ for a focus parameter $\mu = \mu(\theta)$, in a model with parameter θ.

(a). Verify the mse formula (3.2), and note its Pythagorean character. Make a little right-angle triangle figure with absolute bias and standard deviation for the two short sides and root-mse, rmse $= \text{mse}^{1/2}$, on the long side. These root operations bring the risk components down to the original scale of the measurements.

(b). For a simple situation, let $Y \sim \text{N}(\theta, 1)$, with θ to be estimated. Find formulae for the mean squared errors of the three estimators $0.9\,Y$, Y, $1.1\,Y$. Note the interplay between bias and variance.

(c). Generalise to the case of $Y_1, \ldots, Y_n$ being i.i.d. from $\text{N}(\theta, 1)$. Find $\text{mse}(\widehat{\theta}, \theta)$ for the three estimators $0.99\,\bar{Y}$, $\bar{Y}$, $1.01\,\bar{Y}$, with $\bar{Y}$ the sample average. Comment on what you find.

(d). In somewhat more general terms, consider an i.i.d. sample $Y_1, \ldots, Y_n$ from a distribution with unknown mean μ and variance σ^2. Show that $\bar{Y}$ is unbiased with variance σ^2/n. If your estimator for μ is $c_n\bar{Y}$, what is required of c_n, in order for the mean squared error to go to zero with growing n?

Ex. 3.2. *Binomial estimation.* Consider Y being binomial (n, p), as in Ex. 1.3.

(a). To estimate p the canonical choice is $\widehat{p} = Y/n$. Find its mean and variance, and a formula for $\text{mse}(\widehat{p}, p) = \text{E}_p\,(\widehat{p} - p)^2$.

(b). Then compare the simple binomial unbiased proportion with the estimator $\widehat{p}_B = (Y+1)/(n+2)$; this is arguably the very first Bayes estimator. Find its bias and variance, and a formula for $\text{mse}(\widehat{p}_B, p)$. Draw the two mse functions in a diagram, for say $n = 10$. When is the Bayes estimator better than the Y/n, according to this criterion?

Ex. 3.3. *Estimating the normal mean.* Suppose we have independent observations $Y_1, \ldots, Y_n$ from the normal distribution $\text{N}(\xi, \sigma^2)$.

(a). Prove that the sample average $\bar{Y} = n^{-1}\sum_{i=1}^n Y_i$ has the $\text{N}(\xi, \sigma^2/n)$ distribution. This $\bar{Y}$ is the canonical estimator for μ. Find also a clear formula for its risk, or mean squared error, namely $\text{mse}(\bar{Y}, \mu, \sigma) = \text{E}_{\mu,\sigma}(\bar{Y} - \mu)^2$. The subscript indicates that the mean operator is with respect to the probability mechanism dictated by (μ, σ).

(b). Then generalise the above somewhat, by finding the mean and variance also for the estimator $\widehat{\mu} = b\bar{Y}$, with b a constant (which might be close to 1). Use this to put up a clear expression for

$$\text{mse}(b\bar{Y}, \mu, \sigma) = \text{E}_{\mu,\sigma}(b\bar{Y} - \mu)^2.$$

Illustrate this, for values $b = 0.98, 1.00, 1.02$, and comment. For what values of the parameters (μ, σ) will the estimator $0.98\,\bar{Y}$ be better than the classic $\bar{Y}$? Are there values of the parameters where $1.02\,\bar{Y}$ is better than the plain $1.00\,\bar{Y}$?

(c). Suppose the starting assumptions about the data at hand are changed to merely saying that the Y_i are i.i.d. with mean μ and standard deviation σ, i.e. we avoid saying that the distribution of the error terms $\varepsilon_i = (Y_i - \mu)/\sigma$ needs to be exactly standard normal. How does this affect your findings and claims for the previous points?

Ex. 3.4. *Estimating the normal variance.* As in Ex. 3.3, suppose there are i.i.d. data $Y_1, \ldots, Y_n$ from the $\mathrm{N}(\xi, \sigma^2)$. Here we care about the standard deviation parameter σ. As we saw in Ex. 1.49, $Z = \sum_{i=1}^n (Y_i - \bar{Y})^2 \sim \sigma^2 \chi^2_m$, where $m = n - 1$. Also, the Z is stochastically independent of the sample mean $\bar{Y}$.

(a). Use the statement above to find the mean and variance of $\widehat{\sigma}^2 = cZ$ (where c ought to be about $1/n$). Find the mean squared error $\mathrm{mse}(cZ, \sigma) = \mathrm{E}_\sigma\,(cZ - \sigma^2)^2$. Check in particular the result for $c = 1/(n-1)$, the classical factor to make the estimator unbiased; for $c = 1/n$, which comes out of the maximum likelihood paradigm (see Ch. 5); and for $c = 1/(n+1)$. Find the best possible constant c for estimators of this type cZ, using the mean squared error on the σ^2 scale as criterion.

(b). Find also the mean and variance of $dZ^{1/2}$, seen as an estimator of σ, i.e. on the standard deviation scale, not that of the variance. Find an expression for $\mathrm{mse}(dZ^{1/2}, \sigma) = \mathrm{E}_\sigma\,(dZ^{1/2} - \sigma)^2$. Find the best d, according to this criterion.

(c). We may also finetune σ estimation on the log-scale. Examine the risk function $\mathrm{mse}(kZ^{1/2}, \sigma) = \mathrm{E}_\sigma\,\{\log(kZ^{1/2}) - \log\sigma\}^2$, and find the best value of k.

(d). A different solution to the issue of determining 'the best constant' when estimating σ, disregarding tradition and mathematical convenience, might be as follows. With $\widehat{\sigma}^2 = Z/(n-1)$ being the traditional sample variance, with $1/(n-1)$ selected to achieve unbiasedness on the σ^2 scale, consider $\sigma^* = c_n\widehat{\sigma}$, with c_n to be fine-tweaked perhaps a little bit away from 1. Find the c_n that makes $\mathrm{risk}(c_n\widehat{\sigma}, \sigma) = \mathrm{E}_\sigma\,|c_n\widehat{\sigma} - \sigma|$ smallest. This means relying on absolute error as loss function, and the solution needs numerical minimisation of a function which needs numerical integration. Give a table with these optimal c_n for say $n = 10, \ldots, 30$. Show that $c_n \to 1$ as n grows.

Ex. 3.5. *Confidence interval for an exponential rate.* Choose a sample size n, and simulate i.i.d. variables $Y_1, \ldots, Y_n$ from the exponential distribution, see Ex. 1.8, with parameter θ equal to say $\theta_0 = 3.33$.

(a). Construct a 90 percent confidence interval $[L, U]$ for θ. Check if θ_0 is contained in this interval, for the data you generated. Repeat the experiment say 100 times, and record how often the intervals contain θ_0. What is in fact the distribution of N, the number of the 100 intervals which cover the truth?

(b). In addition to checking whether the intervals cover the truth, compute the length $D = U - L$, and give a histogram of its distribution. Find $\mathrm{E}\,D$. Repeat also these experiments with a couple of other sample sizes, and comment.

Ex. 3.6. *Confidence interval for a normal variance.* Let $Y_1, \ldots, Y_n$ be i.i.d. from a normal distribution. How can we set up confidence intervals for the standard deviation σ (or, equivalently, for the variance σ^2)? Writing $m = n - 1$, the sample variance is $\widehat{\sigma}^2 = Z/m$, with $Z = \sum_{i=1}^n (Y_i - \bar{Y})^2 \sim \sigma^2 \chi^2_m$, from Ex. 1.49.

(a). Start with $[a, b] = [\Gamma_m^{-1}(0.05), \Gamma_m^{-1}(0.95)]$, an interval covering the χ^2_m with probability 0.90. Transform $a \le m\widehat{\sigma}^2/\sigma^2 \le b$ to the confidence interval $\text{ci} = [(m/b)^{1/2}\widehat{\sigma}, (m/a)^{1/2}\widehat{\sigma}]$. Show in detail that indeed $\Pr_\sigma(\sigma \in \text{ci}) = 0.90$.

(b). Most often one wishes to estimate and assess the σ parameter directly, being on the same scale as the measurements, but once in a while it would be more natural to communicate and interpret results on the variance scale. Show in suitable detail that $\Pr_\sigma((m/b)\widehat{\sigma}^2 \le \sigma^2 \le (m/a)\widehat{\sigma}^2) = 0.90$; confidence intervals can in this fashion be easily transformed, say from θ to $g(\theta)$, as here, from σ^2 to σ, or the other way around.

(c). The construction above is 'equitailed', starting with 0.05 probability to the left of a and 0.05 probability to the right of b. One might somewhat more generally use any $[a, b]$ with 0.90 probability for the χ^2_m, needing $\Gamma_m(b) - \Gamma_m(a) = 0.90$. The length of the 90 percent σ interval above is proportional to $1/a^{1/2} - 1/b^{1/2}$. Minimise this function, say for $m = 10, 20, 30, 40$. Compare these length-minimising 0.90 intervals with the simpler ones, and comment.

(d). Go through the analogue exercise, this time working with confidence intervals on the σ^2 scale, where the length is proportional to $1/a - 1/b$; minimise this, under the constraint $\Gamma_m(b) - \Gamma_m(a) = 0.90$, and compare with the above, for say $m = 10, 20, 30, 40$. The moral is that such delicate stretching doesn't matter so much, and in statistical practice we are typically content with the equitailed construction. See however the more principled construction of confidence distributions in 7.

Ex. 3.7. *Confidence interval for a normal mean.* Here we go through the basics for constructing confidence intervals for normal means. Since approximations to normality abound in statistical theory and practice, what we learn here quickly finds use also outside the strict normality assumptions.

(a). Start with the simplest prototype setup, a single Y from the $\text{N}(\xi, 1)$ model. Show that the random interval $[Y - 1.96, Y + 1.96]$ captures ξ with probability 0.95.

(b). Suppose $Y_1, \ldots, Y_n$ are i.i.d. from the $\text{N}(\xi, \sigma^2)$ model, at the moment with standard deviation parameter σ taken known. With $\bar{Y}$ as usual being the data average, show that $\sqrt{n}(\bar{Y} - \xi)/\sigma$ is standard normal, and deduce from this that $\bar{Y} \pm 1.96\,\sigma/\sqrt{n}$ is a 95 percent confidence interval for ξ.

(c). In most cases also the σ is unknown, however. Let $\widehat{\sigma}$ be the usual empirical standard deviation, see e.g. Ex. 3.4. Show that the natural construction

$$t = \frac{\bar{Y} - \xi}{\widehat{\sigma}/\sqrt{n}} = \frac{\sqrt{n}(\bar{Y} - \xi)}{\widehat{\sigma}} = \frac{\sqrt{n}(\bar{Y} - \xi)/\sigma}{\widehat{\sigma}/\sigma} \tag{3.4}$$

has a distribution not depending on the two parameters, and that this distribution, call it G_n, is symmetric around zero. Deduce also that with $t_{0,n} = G_n^{-1}(0.975)$, the random interval $\bar{Y} \pm t_{0,n}\widehat{\sigma}/\sqrt{n}$ covers ξ with probability 0.95.

(d). It is then a matter of finding and perhaps tabulating the distribution G_n of t. It is in fact the celebrated t distribution, with $\text{df} = n - 1$ degrees of freedom, see Ex. 1.50. But even without that specific knowledge detail, we could easily have simulated a high number for t, from (3.4), and read off the required quantile.

(e). In more generality, suppose β is a model parameter for which there is an estimator $\widehat{\beta}$ with distribution $\mathrm{N}(\beta, c_n^2\sigma^2)$, say, with a known factor c_n. Suppose also that there is a statistically independent estimator $\widehat{\sigma}$ for σ, with the property that $\widehat{\sigma}^2/\sigma^2 \sim \chi^2_m/m$, for a known m. Then show that $t = (\widehat{\beta} - \beta)/(c_n\widehat{\sigma}) \sim t_m$. Put up a 0.99 confidence interval for β based on this.

(f). Suppose $X_1, \ldots, X_{n_1}$ and $Y_1, \ldots, Y_{n_2}$ are random samples from two normal groups with equal variance, say $\mathrm{N}(\xi_1, \sigma^2)$ and $\mathrm{N}(\xi_2, \sigma^2)$. Find the distribution of the difference of sample means $D_n = \bar{X} - \bar{Y}$, and construct a confidence interval for $\delta = \xi_1 - \xi_2$. This is related to the classic two-sample t-test, see Ex. 4.4.

Ex. 3.8. *Normal quantiles: estimation and confidence.* Consider again the setup of Ex. 3.3, with a sample of Y_i from the normal $\mathrm{N}(\xi, \sigma^2)$ model. In the present exercise we care about quantiles, as opposed to 'only' the mean or the standard deviation.

(a). Show that the c.d.f. of Y_i may be written $F(y) = \Pr(Y_i \le y) = \Phi((y-\xi)/\sigma)$, with $\Phi(x)$ the c.d.f. for the standard normal (i.e. the `pnorm` function in `R`). Show that the q quantile $F^{-1}(q)$ is equal to $\gamma_q = \xi + z_q\sigma$, with $z_q = \Phi^{-1}(q)$. Thus the 0.95 quantile is $\gamma_{0.95} = \xi + 1.645\,\sigma$, etc.

(b). Find the mean and variance of the natural estimator $\widehat{\gamma}_q = \bar{Y} + z_q\,\widehat{\sigma}$, where $\widehat{\sigma} = (Z/m)^{1/2}$, with Z as in Ex. 3.4 and $m = n-1$.

(c). Show that we may write

$$\begin{aligned}\widehat{\gamma}_q - \gamma_q &\sim \xi + (\sigma/\sqrt{n})N + z_q\sigma(K_m/m)^{1/2} - \xi - z_q\sigma \\ &= \sigma[(1/\sqrt{n})N + z_q\{(K_m/m)^{1/2} - 1\}],\end{aligned}$$

in terms of $N \sim \mathrm{N}(0,1)$ and $K_m \sim \chi^2_m$, with these being independent. Verify that the pivot

$$W_{n,q} = \frac{\sqrt{n}(\widehat{\gamma}_q - \gamma_q)}{\widehat{\sigma}} \sim \frac{N + z_q\sqrt{n}\{(K_m/m)^{1/2} - 1\}}{(K_m/m)^{1/2}}$$

has a distribution not depending on ξ, σ. It can be simulated, for any given sample size n and quantile level q. With a_n such that $\Pr(-a_n \le W_{n,q} \le a_n) = 0.95$, convert this to a 95 percent confidence interval for γ_q. Note that for $q = 0.50$, the median case, the distribution of the $W_{n,q}$ is the t_m, the t with degrees of freedom $m = n-1$.

(d). Use the delta method of Ex. 2.47 to show that $\sqrt{n}\{(K_m/m)^{1/2} - 1\} \to_d \mathrm{N}(0, \frac{1}{2})$, and explain that $W_{n,q} \to_d \mathrm{N}(0, 1 + \frac{1}{2}z_q^2)$. Construct an approximate 95 percent confidence interval for γ_q based on this. Simulate a dataset of size e.g. $n = 50$, use both methods from this exercise to find 95 percent confidence intervals for the $\gamma_{0.75}$ quantile. Note that the delta method of Ex. 2.51 leads to approximate confidence intervals for *any* smooth $\gamma = g(\xi, \sigma)$; the point here is that a more exact analysis works for the quantile.

Ex. 3.9. *The empirical distribution function.* Assume there is an i.i.d. dataset $Y_1, \ldots, Y_n$ from an unknown distribution, with c.d.f. $F(t) = \Pr(Y_i \le t)$. We have learned about the empirical distribution function in Ex. 2.11 and 2.61, with $F_n(t) = n^{-1}\sum_{i=1}^n I(Y_i \le t)$ converging uniformly to the underlying F, with probability one; here we learn about how fast this happens, via convergence in distribution.

more on the empirical c.d.f.

(a). Explain that the empirical c.d.f. is the cumulative of the probability measure that puts probability mass $1/n$ at each data point. This is the natural nonparametric estimator of the unknown F.

(b). Construct a version of Figure 3.1, left panel, where $n = 100$ datapoints are simulated from the distribution $f = 0.50\,\mathrm{Expo}(r_1) + 0.50\,\mathrm{Expo}(r_2)$, with rates $r_1 = 2.00$ and $r_2 = 4.00$. Check how well the nonparametric estimator F_n succeeds in coming close to F.

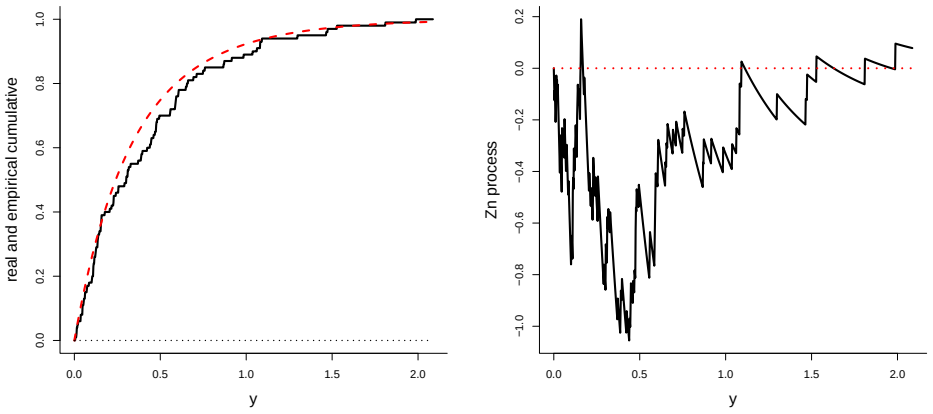


Figure 3.1: *Left panel: the real underlying data-generating $F(t)$ (dashed, red), with the empirical distribution function $F_n(t)$ (full line, black), computed from a sample of $n = 100$ data points from F. Right panel: the process $Z_n(t)$, computed for the data used for the same data. In 95 percent of such cases, the maximum absolute value of the Z_n process will be below 1.358.*

(c). Since we know so much about the binomial, we quickly learn a few basic properties of the F_n. Show that $F_n(t)$ is unbiased for $F(t)$, and that its variance is $F(t)\{1-F(t)\}/n$.

(d). Consider the process $Z_n(t) = \sqrt{n}\{F_n(t) - F(t)\}$. Show that it has mean zero, and that $Z_n(t) \to_d Z(t)$, say, where $Z(t)$ is a zero-mean normal with variance $F(t)\{1-F(t)\}$. Show also that

$$\mathrm{cov}\{Z_n(t), Z_n(t')\} = F(t)\{1 - F(t')\} \quad \text{for } t \le t'.$$

Compute and display the Z_n plot, using the same data values as for the previous figure; in other words, construct a version of Figure 3.1, right panel. Such plots may e.g. be used to check model adequacy – if the data come from a distribution not close to the F used to construct the plot, then the Z_n plot will deviate aignificantly from the zero line. To understand what might qualify as 'significantly different from the zero line' means we need theory for the behaviour of the full Z_n process, not merely the pointwise result that $Z_n(t) \to_d \mathrm{N}(0, F(t)\{1 - F(t)\})$. We return to such questions in Ch. 9.

Ex. 3.10. *Estimating the normal density.* Most often the statistical interest lies in estimating some parameter related to, or expressed through, the normal distribution, like the mean, spread, or quantile, as illustrated above. In some situations one wishes to estimate the density itself. Consider once again a sample $Y_1, \ldots, Y_n$ from the normal $\mathrm{N}(\xi, \sigma^2)$.

(a). For the parameter σ, we shall again use $Z = \sum_{i=1}^n (Y_i - \bar{Y})^2 \sim \sigma^2 \chi^2_m$, with $m = n-1$, as in several previous exercises. With the traditional default estimator $\widehat{\sigma}^2 = Z/(n-1)$ of (1.6), find formulae for the mean of $1/\widehat{\sigma}^2$ and $\log \widehat{\sigma}$.

(b). Construct unbiased estimators for $1/\sigma^2$ and for $\log \sigma$, and then for the log-density function $\log f(y, \xi, \sigma) = -\log \sigma - \frac{1}{2}(y-\xi)^2/\sigma^2 - \frac{1}{2}\log(2\pi)$. In (xx pointer to exercise Ch8) we also construct an unbiased estimator on the direct scale $f(y, \xi, \sigma)$. (xx may put in as hints, or simply have it our solutions; check details. xx) E $\log \widehat{\sigma} = \log \sigma + a_m$, with $a_m = \mathrm{E}\, \log(\chi^2_m/m) = \{\log 2 + \psi(\frac{1}{2}m)\}/(2m)$;

$$\mathrm{E}\, \frac{m}{m-2} \frac{(y-\widehat{\xi})^2}{\widehat{\sigma}^2} = \frac{(y-\xi)^2 + \sigma^2/n}{\sigma^2}.$$

(c). For independent samples $X_1, \ldots, X_{n_1}$ and $Y_1, \ldots, Y_{n_2}$ from normal populations $\mathrm{N}(\xi_1, \sigma_1^2)$ and $\mathrm{N}(\xi_2, \sigma_2^2)$, construct an unbiased estimator for $d(y) = \log\{f_2(y)/f_1(y)\}$.

Confidence via normal approximations

Ex. 3.11. *Confidence intervals via normal approximations.* (xx make connection to Wald tests to come in Ex. 4.5. xx) As we've already seen in various situations of Ch. 2, there are often estimators for interest parameters for which there is approximate normality. Then various recipes under strict normality can still be used, but as approximations.

(a). Suppose μ is such a parameter of interest, for which there is an estimator $\widehat{\mu}$, being approximately normal, in the mathematical sense of $\sqrt{n}(\widehat{\mu} - \mu) \to_d \mathrm{N}(0, \tau^2)$, for some appropriate limiting variance τ^2. Suppose also that there is a consistent estimator $\widehat{\tau}$ of τ, with $\widehat{\tau}/\tau \to_{\mathrm{pr}} 1$. Show that $Z_n = \sqrt{n}(\widehat{\mu} - \mu)/\widehat{\tau} \to_d \mathrm{N}(0,1)$; you may check with Ex. 2.22.

(b). Show under these mild and very frequently met assumptions that

$$\Pr(\mu \in \widehat{\mu} \pm 1.96\, \widehat{\tau}/\sqrt{n}) \to 0.95.$$

In other words, the $[\widehat{\mu} - 1.96\, \widehat{\tau}/\sqrt{n}, \widehat{\mu} + 1.96\, \widehat{\tau}/\sqrt{n}]$ is an *approximate* or *asymptotic* 95 percent interval for μ. Note the grand generality here; this simple construction works in a large variety of situations, also in nonparametric setups, cases with dependent data, etc.

(c). The simplest interesting application of this standard recipe is for the unknown mean ξ of a population. Verify via the CLT of Ch. 2 that $\sqrt{n}(\bar{Y} - \xi)/\widehat{\sigma} \to_d \mathrm{N}(0,1)$. Hence the t-based interval $\bar{Y} \pm 1.96\, \widehat{\sigma}/\sqrt{n}$, for which we have very precise probability computations under normality, is large-sample correct even if the data are not at all normal.

(d). Suppose (X, Y, Z) is trinomial (n, p, q, r), with $p + q + r = 1$. Construct an approximate 90 percent confidence interval for $d = q - p$. – Check that you see how similar and not complicated tasks can be tended to in the examples of Ex. 2.72.

Ex. 3.12. *Confidence intervals for the standard deviation, outside normality.* Consider i.i.d. data $Y_1, \ldots, Y_n$, from a distribution with mean ξ and standard deviation σ, from which we compute the classical sample mean $\widehat{\xi} = \bar{Y}$ and empirical variance $\widehat{\sigma}^2$ of (1.6). Here we illustrate the general large-sample methods by building confidence intervals for σ, with no assumptions on the distribution of the data, like normality. The only mild assumption we make is a finite fourth moment, in order for $\widehat{\sigma}$ to have a clear limit distribution. See Figure 3.2 for 100 simulated confidence intervals, all attempting to capture the true value, here $\sigma = 1$, for two different sample sizes.

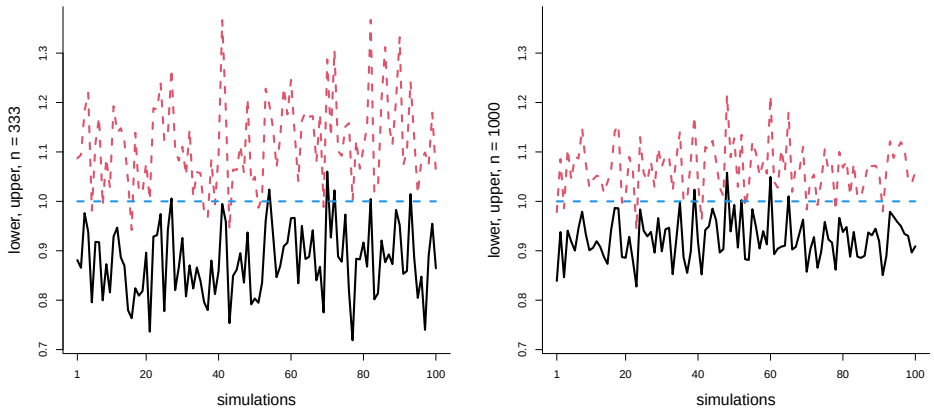


Figure 3.2: *Simulations, with datasets from the unit exponential, displaying lower and upper confidence points for 90 percent intervals; the intervals attempt to cover the true value $\sigma = 1$, and will succeed about 90 percent of the time. Left panel: with $n = 333$; right panel: with $n = 1000$.*

(a). Make sure you understand and can show that $\widehat{\xi}$ and $\widehat{\sigma}$ are consistent for ξ and σ, from the LLN theorems of Ch. 2.

(b). Via results from Ex. 2.46 and the delta method of Ch. 2, show that

$$\begin{pmatrix} \sqrt{n}(\widehat{\xi}-\xi) \\ \sqrt{n}(\widehat{\sigma}-\sigma) \end{pmatrix} \to_d \mathrm{N}_2(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \sigma^2 \begin{pmatrix} 1 & \frac{1}{2}\gamma_3 \\ \frac{1}{2}\gamma_3 & \frac{1}{2}+\frac{1}{4}\gamma_4 \end{pmatrix}),$$

in terms of skewness $\gamma_3 = \mathrm{E}\{(Y-\xi)/\sigma\}^3$ and kurtosis $\gamma_4 = \mathrm{E}\{(Y-\xi)/\sigma\}^3 - 3$.

(c). Show that $\widehat{\gamma}_4 = (1/n)\sum_{i=1}^n \{(Y_i - \bar{Y})/\widehat{\sigma}\}^4 - 3$ is consistent for γ_4, and use this to construct an approximate 90 percent confidence interval for σ. Note that this is a nonparametric procedure, totally free of other distributional assumptions, like normality; *if* one assumes normality, as an extra condition, one may do more, of course.

(d). For an illustration, consider the unit exponential distribution; show that the standard deviation is 1 and that the kurtosis is $\gamma_4 = 6$. Simulate a suitably high number of datasets of size $n = 333$ from this distribution. For each simulated dataset, compute $\widehat{\gamma}_4$, to check how close it is to γ_4, along with the approximate 90 percent confidence interval for σ. Construct a version of Figure 3.2 (left panel for $n = 333$ and right panel for $n = 1000$). Examine in particular the coverage of your intervals: how often do they contain the correct σ? Use simulations to check how big n must be, in order for $\widehat{\gamma}_4$ to be inside $[5.5, 6.5]$ with probability at least 95 percent.

(e). Generate a dataset of size $n = 333$ from the unit exponential, and construct an approximate 90 percent confidence ellipsoid on your screen for (ξ, σ). Check if it contains the true values.

Ex. 3.13. *Inference for the normal density, via large-sample methods.* Consider again i.i.d. data $Y_1, \ldots, Y_n$ from a normal density $f(y, \xi, \sigma)$. In Ex. 3.10 we constructed an

estimator $\widehat{f}(y)$ with the property that $\log \widehat{f}(y)$ is exactly unbiased for $\log f(y,\xi,\sigma)$, for any given n. Here we instead work out the large-sample approximations for the direct estimator $f^*(y) = f(y,\widehat{\xi},\widehat{\sigma})$, having plugged in the usual empirical mean and empirical standard deviation.

(a). We have seen in Ex. 2.51 that $(\sqrt{n}(\widehat{\xi}-\xi), \sqrt{n}(\widehat{\sigma}-\sigma))$ tends to (A,B), independent zero-mean normals with variances σ^2 and $\frac12\sigma^2$. For fixed y, wite $x = (y-\xi)/\sigma$. Use the delta method to explain that (i) $\sqrt{n}(\log\widehat{\sigma} - \log\sigma) \to_d (1/\sigma)B$; (ii) $\sqrt{n}((y-\widehat{\xi})/\widehat{\sigma} - (y-\xi)/\sigma) \to_d (1/\sigma)(-A - xB)$; (iii) $\sqrt{n}\{(y-\widehat{\xi})^2/\widehat{\sigma}^2 - x^2\} \to_d (1/\sigma)(-2xA - x^2B)$. Then combine these to show that

$$\sqrt{n}\{\log f^*(y) - f(y,\xi,\sigma)\} \to_d (1/\sigma)(-B + xA + \tfrac12 x^2 B),$$

a zero-mean normal with variance $\tau^2 = x^2 + (\frac12 x^2 - 1)^2 = \frac12(1+x^4)$.

(b). Show that the log-unbiased estimator $\log\widehat{f}(y)$ is close enough to this $\log f^*(y)$ to make the scaled difference $\sqrt{n}\{\log f^*(y) - \log\widehat{f}(y)\}$ go to zero in probability; explain that the two estimators therefore have the same limit distribution.

(c). Explain how a pointwise 90 percent confidence band can be constructed for the log-density, of the type $\log f^*(y) \pm 1.645\,\widehat{\tau}(y)$, which then may be transformed back to the density scale.

Ex. 3.14. *Variance of the variance estimator.* Let $Y_1,\ldots,Y_n$ be i.i.d., with mean ξ, variance σ^2, and finite kurtosis $\gamma_4 = \mathrm{E}\,(Y_i-\xi)^4/\sigma^4 - 3$.

(a). With $A = \sum_{i=1}^n (Y_i - \bar{Y})^2$, show that $\mathrm{E}\,A = (n-1)\sigma^2$. This says that $\widehat{\sigma}^2 = A/(n-1)$ is unbiased for the variance, regardless of the underlying distribution.

(b). If the Y_i are actually normal, then $A \sim \sigma^2\chi^2_m$, with $m = n-1$. Show that $\mathrm{Var}\,\widehat{\sigma}^2 = 2\sigma^4/(n-1)$.

(c). Outside normality, work out an expression for $\mathrm{Var}\,A$, and show that

$$\mathrm{Var}\,\widehat{\sigma}^2 = \Big(3 + \gamma_4 - \frac{n-3}{n-1}\Big)\frac{\sigma^4}{n} = \frac{2\sigma^4}{n-1} + \frac{\gamma_4\sigma^4}{n}.$$

simulate a high number of samples of size $n = 12$ from the t distribution t_m, with say $m = 6$, and 'verify' the formula.

(d). Check that the finite-sample formula above matches the large-sample limit for $\sqrt{n}(\widehat{\sigma}^2 - \sigma^2)$.

Ex. 3.15. *The binomial, the normal approximation, and confidence intervals.* Consider Y_n, a binomial (n,p). To invoke the CLT it is practical to use the $Y_n = X_1 + \cdots + X_n$ representation, in terms of X_i being i.i.d. Bernoullis.

(a). Use the CLT of Ch. 2 to deduce that the normalised variable

$$W_n = \frac{\sum_{i=1}^n (X_i - p)}{\{\mathrm{Var}\sum_{i=1}^n (X_i-p)\}^{1/2}} = \frac{Y_n - np}{\{np(1-p)\}^{1/2}} = \frac{\widehat{p}-p}{\{p(1-p)/n\}^{1/2}}$$

converges to the standard normal $\mathrm{N}(0,1)$ with increasing n. Discuss briefly the relevance of the skewness result from Ex. 1.3 in light of the limiting normality.

(b). With $\widehat{p}_B = (Y+1)/(n+2)$, show that the difference between W_n and $W_{n,B} = \sqrt{n}(\widehat{p}_B - p)$ is so small, for large n, that $W_{n,B}$ must have the same normal limit. The confidence intervals we construct below, based on $\widehat{p}$, can therefore alternatively be based on $\widehat{p}_B$.

binomial confidence

(c). Verify from the above that $\Pr(-1.96 \leq W_n \leq 1.96) \to 0.95$ as sample size increases, and use this to construct an interval, based on having observed $Y_n = y$ in a given experiment with known n, which covers the true p with probability approximately 95 percent.

asymptotic equivalence

(d). There are actually several constructions of such confidence intervals, with this property. Considering the basic estimator $\widehat{p} = Y/n$ again, write $\sigma_n^2 = p(1-p)/n$ for its variance, and $\widehat{\sigma}_n^2 = \widehat{p}(1-\widehat{p})/n$ for its estimated variance. Verify that both

$$W_n = (\widehat{p} - p)/\sigma_n \quad \text{and} \quad W_n' = (\widehat{p} - p)/\widehat{\sigma}_n$$

tend to the standard normal in distribution. Now show that the arguments above, used for W_n' in lieu of W_n, lead to the confidence interval $\widehat{p} \pm 1.96\,\widehat{\sigma}_n$ instead. Exemplify, with $n = 100$, for the three cases $y = 22$, $y = 55$, $y = 77$, where you compute both versions of the 95 percent confidence interval for p.

(e). Suppose certain details related to your applied research project require that you compute the probability p that $L \leq 1.33\,R$, where

$$L = \{(G_1/G)(G_2/G)(G_3/G)(G_4/G)\}^{1/4}, \quad R = \{(G_5/G)(G_6/G)(G_7/G)\}^{1/3},$$

in terms of $G_1, \ldots, G_8$ being i.i.d. from the χ^2_{12} distribution (the chi-squared with degrees of freedom equal to 12), and $G = \sum_{i=1}^8 G_i$. Since it's hard to find an exact formula, or an exact answer in other ways, you *simulate* a high number sim of such vectors $(G_1, \ldots, G_8)$, and check for each simulation whether the event just described takes place or not. How large should sim be, in order for your simulation based estimate of p to be correct to three decimal places? Carry out such simulations and thus find p. Display also a histogram of simulated L/R.

Ex. 3.16. *Confidence intervals via the delta method.* The delta method of Ex. 2.47 can be used to derive approximate normality of various estimators, after which the generic recipe of Ex. 3.11 may be used to form confidence intervals. Once limiting normality has been established, there are automatically limiting normality statements and uses for transformated parameters.

(a). Suppose $\sqrt{n}(\widehat{\mu} - \mu) \to_d \mathrm{N}(0, \kappa^2)$, for some parameter μ of interest. Confidence intervals then take the form $\widehat{\mu} \pm (c/\sqrt{n})\widehat{\kappa}$. With a smooth transformation to $\gamma = g(\mu)$ and $\widehat{\gamma} = g(\widehat{\mu})$, explain that we have $\sqrt{n}(\widehat{\gamma} - \gamma) \to_d g'(\mu)\kappa\,\mathrm{N}(0,1)$. If $\gamma = g(\mu)$ is monotone, with an inverse $\mu = g^{-1}(\gamma)$, explain that one via back-transformation can form the alternative interval

$$g^{-1}\big(\widehat{\gamma} \pm (c/\sqrt{n})g'(\widehat{\mu})\widehat{\kappa}\big).$$

Explain also how this simplifies for variance stabilising transformations, where $g'(\mu)\kappa = 1$.

(b). Illustrate the above for a binomial (n,p) situation, where we know from Ex. 2.72 that $\sqrt{n}(\widehat{\gamma} - \gamma) \to_d \mathrm{N}(0,1)$, with $\gamma = g(p) = 2\arcsin(p^{1/2})$: construct the interval on the scale of γ, and then back-transform to the p scale. Check numerically how close such intervals are to the standard ones, for a few combinations of n and binomial outcome y.

(c). Show via Taylor expansion that $g^{-1}(\widehat{\gamma}+\varepsilon) = \widehat{\mu} + (g^{-1})'(\widehat{\gamma})\varepsilon + O(\varepsilon^2)$ for small ε. Deduce that that the transformation intervals above become large-sample equivalent to the $\widehat{\mu} \pm (c/\sqrt{n})\widehat{\kappa}$ we started out with.

(d). (xx ranting. xx) So which g should one perhaps use? Suppose g_a is one such transformation, with a tweaking parameter a. Argue that the best choices of a are those for which $g_a(\widehat{\mu}) - g_a(\mu)$ are close to having zero-mean normal distributions. Play a bit with this idea, for the case of having normal approximation confidence intervals for σ in the $\mathrm{N}(\xi, \sigma^2)$ situation. The starting point here would be $\sqrt{n}(\widehat{\sigma} - \sigma) \to_d \mathrm{N}(0, \frac{1}{2}\sigma^2)$. Attempt the transformation $g_a(\sigma) = \sigma^a$; the task is to check accurate normality for the appopriate $\sqrt{n}(\widehat{\sigma}^a - \sigma^a)/\widehat{\kappa}_a$.

Quantiles and sample quantiles

Ex. 3.17. *Uniform ordering.* Consider $U_1, \ldots, U_n$ i.i.d. from the uniform distribution on the unit interval. Order these, to $U_{(1)} < \cdots < U_{(n)}$.

(a). Explain that in order to have $U_{(i)} \in [u, u+\mathrm{d}u]$, one needs $i-1$ points in $[0,u]$, 1 point in $[u, u+\mathrm{d}u]$, and $n-1$ points in $(u+\mathrm{d}u, 1]$. Show from this trinomial setup that $U_{(i)}$ has density

connection to Beta distributions

$$g_i(u) = \frac{n!}{(i-1)!\,1!\,(n-i)!} u^{i-1}(1-u)^{n-i} \quad \text{for } u \in (0,1).$$

Explain that $U_{(i)} \sim \mathrm{Beta}(i, n-i+1)$, see Ex. 1.22, and show that

$$\mathrm{E}\, U_{(i)} = p_i = i/(n+1) \text{ and that } \mathrm{Var}\, U_{(i)} = p_i(1-p_i)/(n+2).$$

(b). Explain further that $U_{(i)} \le p$ is equivalent to the binomial count Z of datapoints falling in $[0,p]$ being at least i, and that this gives rise to the following equation:

Beta and binomial connection

$$\frac{n!}{(i-1)!\,(n-i)!} \int_0^p u^{i-1}(1-u)^{n-i}\,\mathrm{d}u = \sum_{j=i}^{n} \binom{n}{j} p^j (1-p)^{n-j}.$$

(c). With $i < j$, show that $(U_{(i)}, U_{(j)})$ has density

$$g_{i,j}(u,v) = \frac{n!}{(i-1)!\,1!\,(j-i-1)!\,1!\,(n-j)!} u^{i-1}(v-u)^{j-i-1}(1-v)^{n-j}$$

for $u < v$. The idea behind the reasoning, and the ensuing notation, is that in order to see $U_{(i)} \in [u, u+\mathrm{d}u]$ and $U_{(j)} \in [v, v+\mathrm{d}v]$, there is a multinomial situation, with five boxes $[0,u], [u,u+\mathrm{d}u], [u+\mathrm{d}u, v], [v, v+\mathrm{d}v], [v+\mathrm{d}v, 1]$, inside which we need to find $i-1, 1, j-i-1, 1, n-j$ datapoints.

(d). For an i.i.d. uniform sample $U_1, \ldots, U_n$ on $[0,1]$, consider the uniform range $R_n = U_{(n)} - U_{(1)}$, where we know that $U_{(1)} \sim \mathrm{Be}(1,n)$. Show that given $U_{(1)} = u$, $U_{(n)}$ can be represented as $u + Z$, where Z is the maximum of another uniform sample, of size $n-1$, on $[0, 1-u]$. Use this to show that the c.d.f. of R_n can be expressed as $H_n(r) = nr^{n-1}(1-r) + r^n = nr^{n-1} - (n-1)r^n$, and show that this is the $\mathrm{Be}(n-1, 2)$ distribution.

(e). In general, if $Y_1, \ldots, Y_n$ are i.i.d. from a density f, show that the joint density for the full order statistic vector $(Y_{(1)}, \ldots, Y_{(n)})$ is $n! f(y_{(1)}) \cdots f(y_{(n)})$, on the set where $y_{(1)} < \cdots < y_{(n)}$. In particular, for order statistics from the uniform distribution, show that the joint density of $(U_{(1)}, \ldots, U_{(n)})$ is flat and equal to $n!$ on the set $u_{(1)} < \cdots < u_{(n)}$.

connection to exponential and Dirichlet

(f). Use this, in conjunction with Ex. 1.20, to demonstrate that

$$\begin{aligned}(U_{(1)}, U_{(2)}, \ldots, U_{(n)}) &=_d (D_1, D_1 + D_2, \ldots, D_1 + \cdots + D_n) \\ &=_d (V_1/S, (V_1 + V_2)/S, \ldots, (V_1 + \cdots + V_n)/S),\end{aligned}$$

with $V_1, \ldots, V_n, V_{n+1}$ being i.i.d. from the unit exponential, with sum $S = V_1 + \cdots + V_{n+1}$, and $(D_1, \ldots, D_n, D_{n+1})$ is a flat Dirichlet $(1, \ldots, 1, 1)$. The differences $D_i = U_{(i)} - U_{(i-1)}$ are called the *spacings*. We use '$=_d$' to signal equality in distributions. Show that this leads to the representation of the order statistics process as

$$U_{([nq])} = \sum_{i=1}^{[nq]} V_i / \sum_{i=1}^{n+1} V_i \quad \text{for } 0 \le q \le 1. \tag{3.5}$$

Here $[nq]$ is the largest integer less than or equal to nq. We have seen in Ex. 2.14 that $U_{([nq])} \to_{\mathrm{pr}} q$, via consistency properties of the empirical c.d.f. Now prove this result again, via the representation (3.5) and the law of large numbers.

Ex. 3.18. *Ratios of ordered uniforms.* (xx again, need checing and calibration, regarding what is told where. xx) Let $U_1, \ldots, U_n$ be an i.i.d. sample from the uniform distribution on the unit interval, and order these into $U_{(1)} < \cdots < U_{(n)}$. From these form the ratios

$$V_1 = U_{(1)}/U_{(2)},\ V_2 = U_{(2)}/U_{(3)},\ \ldots, V_{n-1} = U_{(n-1)}/U_{(n)},\ V_n = U_{(n)}/1.$$

(a). Show that the inverse transformation leads to the representation

$$U_{(n)} = V_n,\ U_{(n-1)} = V_n V_{n-1},\ \ldots, U_{(2)} = V_n V_{n-1} \cdots V_2,\ U_{(1)} = V_n V_{n-1} \cdots V_2 V_1.$$

(b). Find the joint probability density for $(V_1, \ldots, V_n)$, and show in fact that these are independent, with

$$V_1 \sim \mathrm{Beta}(1,1),\ V_2 \sim \mathrm{Beta}(2,1),\ \ldots, V_{n-1} \sim \mathrm{Beta}(n-1,1),\ V_n \sim \mathrm{Beta}(n,1).$$

(c). Independently of details above, we know from Ex. 3.17 that $U_{(i)}$ has a $\mathrm{Beta}(i, n-i+1)$ distribution. Explain that this $\mathrm{Beta}(i, n-i+1)$ can be represented as a product of different independent Beta variables.

(d). It is of course a somewhat cumbersome simulation recipe for generating a uniform sample, but it is a useful exercise, opening doors & minds to fruitful generalisations: For $n = 10$, say, generate ordered uniform samples of size n in your computer via the representation above, in terms of products of Beta variables. Carry out some checks to see that each single $U_{(i)}$ then has the right distribution, i.e. as described in (c).

(e). Work with the following generalisation of the construction above: Let $X_1, \ldots, X_n$ be an i.i.d. sample from the distribution with density $f(x) = ax^{a-1}$, i.e. a $\mathrm{Beta}(a,1)$. Again form the ratios $V_i = X_{(i)}/X_{(i+1)}$ as above, leading to $X_{(i)} = V_i V_{i+1} \cdots V_n$. Show that the V_i are again independent, now with $V_i \sim \mathrm{Beta}(ai, 1)$.

Ex. 3.19. *Min and max of two uniforms.* Suppose Y_1, Y_2 are i.i.d. from a density $f(y)$, and order them, to $V_1 < V_2$. (xx ask per august and martin why this particular probability calculation was of value. xx)

(a). Show that (V_1, V_2) has joint density $2f(v_1)f(v_2)$, on the set where $v_1 < v_2$.

(b). Then consider the special case of two datapoints from the uniform distribution on the unit interval, ordered to $V_1 < V_2$. Show that $R = V_1/V_2$ is another uniform on the unit interval, and that $W = V_2 - V_1$ is a Beta$(1, 2)$. Show that $\Pr(V_2 - V_1 \le c) = \Pr(V_2 - V_1 > c) = \frac{1}{2}$, for $c = 1 - 1/\sqrt{2} = 0.2929$.

(c). Find also the joint distribution for (R, W) here.

Ex. 3.20. *The sample median.* Let $Y_1, \ldots, Y_n$ be i.i.d. from a positive density f with true median $\theta = F^{-1}(\frac{1}{2})$.

(a). Suppose for simplicity that n is odd, say $n = 2m + 1$. Show that M_n has density of the form

$$g_n(y) = \frac{(2m+1)!}{m!\,m!} F(y)^m \{1 - F(y)\}^m f(y).$$

(b). Show then that the density of $Z_n = \sqrt{n}(M_n - \theta)$ can be written in the form $h_n(z) = g_n(\theta + z/\sqrt{n})/\sqrt{n}$. Prove that

$$h_n(z) \to (2\pi)^{-1/2} 2f(\theta) \exp\{-\tfrac{1}{2} 4 f(\theta)^2 z^2\},$$

where the Stirling approximation formula of Ex. 2.35 may be of use. The limit is the density $h(z)$ of the normal $\mathrm{N}(0, \tau^2)$, with $\tau = \frac{1}{2}/f(\theta)$. We have hence proved $Z_n \to_d \mathrm{N}(0, \tau^2)$, by Scheffé's lemma; see Ex. 2.6.

(c). So when is the sample mean best, and when might the sample median be the better estimator, when it comes to estimating the centre point θ of a symmetric density? Explain that this for large n is a matter of the ratio

$$\rho = \frac{\sigma}{\frac{1}{2}/f(\theta)} = 2\sigma f(\theta),$$

where σ is the standard deviation for f. Argue that if $\rho < 1$, then the sample mean is best, and that if $\rho > 1$, then the sample median is the best.

(d). Compare the limiting distributions for the sample mean and the sample median for the normal density, the double exponential density $\frac{1}{2}\exp(-|y|)$, and the Cauchy density $(1/\pi)/(1 + y^2)$.

(e). Consider the t distribution, with degrees of freedom ν, see Ex. 1.50. find an expression for the ratio $\rho = \rho(\nu)$, plot $(\nu, \rho(\nu))$ in a diagram, and comment. Show that $\rho(\nu)$ approaches $(2/\pi)^{1/2} = 0.7979$ for large ν. Show that for $\nu < 4.678$, there is roughness at the top, and the median is best; whereas for $\nu > 4.678$, there is a smoother density at the top, and the mean is best.

(f). Carry out a similar analysis for the binormal symmetric mixture model $f = \frac{1}{2}\,\mathrm{N}(-a, 1) + \frac{1}{2}\,\mathrm{N}(a, 1)$. For which values of a is the sample median a better estimator of the centre point then the sample mean? [xx later on, another chapter: the estimator which says $\widehat{\theta}$ is sample median if A_n and sample mean if A_n^c, where A_n is the event that $\frac{1}{2}/\widehat{f}(\widehat{\theta}_0) < \widehat{\sigma}$. xx]

Ex. 3.21. *Sample quantiles.* Suppose $Y_1, \ldots, Y_n$ are i.i.d. from some distribution with positive density f and c.d.f. F. The sample median estimates the population median $F^{-1}(0.50)$, and similarly the sample quantile $Q_n(q)$, at any prescribed level $q \in (0,1)$, estimates the population quantiles $F^{-1}(q)$. Built-in functions like `quantile(data,0.33)` in `R` find such sample quantiles directly, so users do not need the cumbersome linear interpolation fiddling between the two ordered observations coming closest to nq, or to care too much about ties in the data due to rounding-off errors. This exercise finds limit distributions for $\sqrt{n}\{Q_n(q) - F^{-1}(q)\}$. sample quantiles

(a). Suppose $U \sim$ unif and let $Y = F^{-1}(U)$. Show that Y has distribution F, and hence density f.

(b). Explain that the full order statistic vector $Y_{(1)} < \cdots < Y_{(n)}$ may be represented via a correspondingly ordered sample of the uniform, as $F^{-1}(U_{(1)}) < \cdots < F^{-1}(U_{(n)})$, with $U_{(i)}$ being the ith ordered observations in an i.i.d. sample $U_1, \ldots, U_n$ from the uniform, studied in Ex. 3.17. In particular, $Y_{(i)}$ has the same distribution as $F^{-1}(U_{(i)})$.

(c). This also means that if we work out basic approximation results for the order statistics from the uniform, we are a modest delta method step away from similar results for the general case of a density f. In particular, suppose we manage to show $\sqrt{n}(U_{([nq])} - q) \to Z_q$, for some Z_q. Show that this implies $\sqrt{n}\{Q_n(q) - F^{-1}(q)\} \to_d c(q)Z_q$, with $c(q) = (F^{-1})'(q) = 1/f(\mu_q)$, writing $\mu_q = F^{-1}(q)$.

(d). To illustrate this point in a simple case first, show from what we already know in Ex. 3.20 that $\sqrt{n}(U_{([0.50\,n])} - 0.50) \to_d \mathrm{N}(0, 0.50^2)$, for the uniform median. Then show for a general density f that $\sqrt{n}\{Q_n(0.50) - \mu\} \to_d \mathrm{N}(0, 0.50^2/f(\mu)^2)$, with $\mu = F^{-1}(0.50)$ the population median. Here we used an exact expression for the density of the median.

(e). There are actually several other ways of proving this median of uniforms result. Such an alternative approach is to use the representation (3.5). Explain that $U_{([0.50\,n])} = A_n/(A_n + B_n)$, with A_n and B_n the averages of the first and second half of i.i.d. variables $V_1, \ldots, V_{n+1}$ from the unit exponential. Use the CLT for the joint limit distributions of $\sqrt{n}(A_n - \frac{1}{2})$ and $\sqrt{n}(B_n - \frac{1}{2})$, and then use the delta method to land the $\mathrm{N}(0, 0.50^2)$ limit.

(f). Then generalise to the case of any given quantile level q. Show first that $\sqrt{n}(U_{[nq]} - q) \to_d \mathrm{N}(0, q(1-q))$, and then that the limit distribution is $\mathrm{N}(0, q(1-q)/f(\mu_q)^2)$ for $\sqrt{n}\{Q_n(q) - \mu_q\}$.

(g). For the case of two quantiles jointly, like the lower and upper sample quartiles, show for the uniform case that with $q_1 < q_2$,

$$\begin{pmatrix} \sqrt{n}\{Q_n(q_1) - q_1\} \\ \sqrt{n}\{Q_n(q_2) - q_2\} \end{pmatrix} \to_d \mathrm{N}_2\Big(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} q_1(1-q_1) & q_1(1-q_2) \\ q_1(1-q_2) & q_2(1-q_2) \end{pmatrix}\Big).$$

Prove this via the explicit density for $(U_{(i)}, U_{(j)})$, given in Ex. 3.17, using the Scheffé Lemma of Ex. 2.6.

(h). It is also instructive to use representation (3.5) via i.i.d. unit exponentials. Prove the joint binormal limit distribution statement from this. Then generalise to the case of r quantiles, for levels $q_1 < \cdots < q_r$. Show for the uniform case that there is a joint multivariate normal limit, with variances $q_j(1-q_j)$ and covariances $-q_j q_\ell$ for $j < l$.

Then carry out the transformation arguments needed to prove that for the case of an underlying positive density f, there is limiting joint normality for $\sqrt{n}(Q_{n,j} - \mu_j\}$, where the limit has variances $q_j(1-q_j)/f(\mu_j)^2$ and covariances $q_j(1-q_k)/\{f(\mu_j)f(\mu_k)\}$ for $j < k$. See Ex. 9.38 for convergence of a full quantile process.

(i)**.** So how variable are sample quantiles, how much bigger are variances out in the tails, compared to in the middle? Explain that the limiting standard deviation at quantile level q is $\mathrm{sd}(q) = c(q)\{q(1-q)\}^{1/2} = \{q(1-q)\}^{1/2}/f(\mu_q)$. For a few densities, including the uniform, the normal, the exponential, compute and plot this $\mathrm{sd}(q)$.

Ex. 3.22. *Good and bad estimators.* Suppose $X_1, \ldots, X_n$ are i.i.d. from the density $f(x,\theta) = \exp\{-(x-\theta)\}$ for $y \geq \theta$, i.e. a unit exponential starting at parameter θ.

(a)**.** Explain that we have $X_i = \theta + Y_i$, with the Y_i being i.i.d. from the unit exponential, and hene that the order statistics can be represented as $X_{(i)} = \theta + Y_{(i)}$, cf. Ex. 1.15.

(b)**.** For the smallest and largest observations, show that $\widehat{\theta}_A = X_{(1)} - 1/n$ and $\widehat{\theta}_B = X_{(n)} - H_n$ are unbiased estimators of θ, with $H_n = 1 + 1/2 + \cdots + 1/n$ the partial sum of the harmonic series. Find their variances.

(c)**.** (xx a bit more. spell out that $\widehat{\theta}_B$ is not consistent. a bit on $X_{(i)} - c_i$ too, where $c_i = 1/n + \cdots + 1/(n-i+1) = H_n - H_{n-i}$. median is ok. xx)

Ex. 3.23. *Smooth functions of quantiles.* (xx point to Ex. 2.50 for similar method and similar results. xx) Suppose $Y_1, \ldots, Y_n$ are i.i.d. from some distribution with continuous and positive density f. The inverse c.d.f. F^{-1} is then well-defined.

(a)**.** Consider some smooth function $B(\mu_1, \ldots, \mu_r)$ of the quantiles $\mu_j = F^{-1}(q_j)$. Via Ex. 3.21, identify first the multinormal limit $X \sim \mathrm{N}_r(0, \Sigma)$ of the vector with components $\sqrt{n}(Q_{n,j} - \mu_j)$, with $Q_{n,j} = F_n^{-1}(\mu_j)$ the sample quantiles. Then establish that with $\widehat{B} = B(Q_{n,1}, \ldots, Q_{n,r})$, we have $\sqrt{n}(\widehat{B} - B) \to_d c^{\mathrm{t}}X \sim \mathrm{N}(0, \tau^2)$, with $\tau^2 = c^{\mathrm{t}}\Sigma c$. Here c is the vector of $c_j = \partial B(\mu_1, \ldots, \mu_r)/\partial \mu_j$.

(b)**.** Suppose $Y_1, \ldots, Y_n$ are independent unit exponentials. With $p < q$, find the limit distribution of the ratio $\widehat{\gamma} = F_n^{-1}(q)/F_n^{-1}(p)$. (xx check: we do have an exercise with quartiles and $F_n^{-1}(1-\varepsilon) - F_n^{-1}(\varepsilon)$ for symmetric distributions? xx)

(c)**.** Suppose $Y_1, \ldots, Y_n$ are i.i.d. from the Weibull distribution with c.d.f. $F(y) = 1 - \exp(-y^b)$. Show that the quantile function is $\mu_p = F^{-1}(p) = d(p)^{1/b}$, with $d(p) = -\log(1-p)$. With $p < q$, find the limit distribution of

$$\sqrt{n}\Big[\frac{Q_{n,q}}{Q_{n,p}} - \Big\{\frac{d(q)}{d(p)}\Big\}^{1/b}\Big].$$

Argue that with $R_n = Q_{n,q}/Q_{n,p}$, the solution $\widehat{b}$ to $\log R_n = (1/b)\log\{d(q)/d(p)\}$ is a consistent estimator. Find the limit distribution of $\sqrt{n}(\widehat{b} - b)$.

(d)**.** (xx one more, with three or four quantiles at work. perhaps point to quantile pyramid things. xx)

(e)**.** (xx might wait until a Ch9 exercise with processes. xx) Consider quantiles levels $p < q < r$ for i.i.d. observations $U_1, \ldots, U_n$ from the uniform distribution on the unit

interval. For empirical quantiles $Q_{n,p}, Q_{n,q}, Q_{n,r}$, explain via Ex. 3.21 that

$$\begin{pmatrix} \sqrt{n}(Q_{n,p}-p) \\ \sqrt{n}(Q_{n,q}-q) \\ \sqrt{n}(Q_{n,r}-r) \end{pmatrix} \to_d \begin{pmatrix} A \\ B \\ C \end{pmatrix} \sim \mathrm{N}_3(0,\Sigma),$$

and identify the elements of Σ. From the empirical quantiles, form the two successive ratios $R_n = Q_{n,q}/Q_{n,p}$ and $S_n = Q_{n,r}/Q_{n,q}$. Show that

$$\begin{pmatrix} \sqrt{n}(R_n - q/p) \\ \sqrt{n}(S_n - r/q) \end{pmatrix} \to_d \begin{pmatrix} R \\ S \end{pmatrix} = \begin{pmatrix} (1/p)\{-(q/p)A + B\} \\ (1/q)\{-(r/q)B + C\} \end{pmatrix}.$$

Do the algebra to verify that R and S have variances $(1/p)^2\{p(1-q/p)^2 + q - p\}$ and $(1/q)^2\{q((1-r/q)^2 + r - q\}$, and that they are in fact independent.

Ex. 3.24. *Exercises with sample quantiles.* [xx various things, using the general results above. interquartile range $R_n = Q_n(0.75) - Q_n(0.25)$, e.g. for the normal. for the Cauchy, cool factoid: $\sqrt{n}(R_n - 2) \to_d \mathrm{N}(0, \pi^2)$. Limit distribution of sample median, given the two 0.25 and 0.75 quartiles. a little link to the nonparametric quantile processes of Hjort and Petrone (2007) and the more general quantile pyramids of Hjort and Walker (2009). also pointer to fuller process result in Ch. 9. the limit is $(F^{-1})'(q)W^0(q)$. xx]

(a). (xx perhaps pushed to BNP chapter. xx) Consider building a model for $V_1 < V_2 < V_3$ as follows: (i) $V_2 \sim \mathrm{unif}(0,1)$; (ii) given $V_2 = v_2$, let independently $V_1 \sim \mathrm{unif}(0, v_2)$ and $V_3 \sim \mathrm{unif}(v_2, 1)$. Show that (V_1, V_2, V_3) has mean $(1/4, 2/4, 3/4)$.

(b). Show that the densities of V_1 and V_3 become $g_1(v_1) = -\log v_1$ and $g_3(v_3) = -\log(1 - v_3)$ on the unit interval.

(c). Generalise to the case of $V_1 < \cdots < V_7$, thought of as random versions of the seven octiles: (i) $V_4 \sim \mathrm{unif}(0,1)$; (ii) given V_4, let $V_2 \sim \mathrm{unif}(0, V_4)$ and $V_6 \sim (V_4, 1)$, independently; (iii) given V_2, V_4, V_6, let V_1, V_3, V_5, V_7 be independent and uniform on the intervals $(0, V_2), (V_2, V_4), (V_4, V_6), (V_6, 1)$, respectively. Show that V_j has mean $j/8$, and find the densities for each individual V_j.

Ex. 3.25. *Which order statistics interval contains the true median?* Let $Y_1, \ldots, Y_n$ be i.i.d. from a positive and smooth density f, with cumulative F. With $Y_{(1)} < \cdots < Y_{(n)}$ the order statistics, which of the subintervals $(Y_{(i)}, Y_{(i+1)})$ will contain the true median, $\mu = F^{-1}(\frac{1}{2})$?

(a). Show that $p_i = \Pr\{\mu \in (Y_{(i)}, Y_{(i+1)})\} = \Pr\{\frac{1}{2} \in (U_{(i)}, U_{(i+1)})\}$, in terms of the order statistics from a uniform sample. We allow $i = 0, 1, \ldots, n$, here, for the $n+1$ possibilities for which interval shall contain μ, writing $u_{(0)} = 0$ and $u_{(n+1)} = 1$.

(b). Given $U_{(i)} = u$, show that the distribution of $U_{(i+1)}$ is the same as that of $u + (1-u)W$, where W is the smallest of $n - i$ observations from the uniforom in the unit interval. Use this to show that

$$\begin{aligned} p_i &= \Pr\{U_{(i)} < \tfrac{1}{2} < U_{(i+1)}\} = \int_0^{1/2} \Pr\{U_{(i+1)} > \tfrac{1}{2} \,|\, U_{(i)} = u\} g_i(u)\,\mathrm{d}u \\ &= \int_0^{1/2} \Big(\frac{\frac{1}{2}}{1-u}\Big)^{n-i} \mathrm{be}(u, i, n-i+1)\,\mathrm{d}u, \end{aligned}$$

involving a Beta density, as per Ex. 3.21. Show that this indeed leads to the explicit probability

$$p_i = (\tfrac12)^{n-i}\frac{n!}{(i-1)!\,(n-i)!}\int_0^{1/2} u^{i-1}\,\mathrm{d}u = \binom{n}{i}(\tfrac12)^n.$$

Hence we've reached the binomial probabilities, for a binom$(n,\frac12)$, via direct probability calculations. Try also to give a direct argument.

(c). (xx generalise to general quantile $\mu_p = F^{-1}(p)$. xx)

Moment matching methods

Ex. 3.26. *Moment matching estimators.* Suppose $Y_1,\ldots,Y_n$ are i.i.d. from some model $f(y,\theta)$, where $\theta = (\theta_1,\ldots,\theta_p)^{\mathrm{t}}$ is of dimension p. The *method of moments* consists in fitting the first p empirical moments to the theoretical ones. In detail, one computes

method of moments

$$M_1 = \bar Y = n^{-1}\sum_{i=1}^n Y_i,\ M_2 = n^{-1}\sum_{i=1}^n (Y_i-\bar Y)^2,\ \ldots,\ M_p = n^{-1}\sum_{i=1}^n (Y_i-\bar Y)^p,$$

and solves the p equations $M_1 = g_1(\theta),\ldots,M_p = g_p(\theta)$, where $g_1(\theta) = \mathrm{E}_\theta\, Y$, $g_2(\theta) = \mathrm{E}_\theta\,\{Y-g_1(\theta)\}^2$, up to $g_p(\theta) = \mathrm{E}_\theta\,\{Y-g_1(\theta)\}^p$.

(a). For one-parameter models, explain that this amounts to fitting the empirical and theoretical mean. If $Y_1,\ldots,Y_n$ are i.i.d. geom(p), see Ex. 1.25, use $\mathrm{E}\,Y_i = 1/p$ to find the method of moments estimator for p. For another application, assume $Y_1,\ldots,Y_n$ follow the distribution with c.d.f. y^θ on $[0,1]$. Find the method of moments estimator for θ.

(b). For two-parameter models, explain that the method of moments means fitting the empirical mean and variance to the theoretical ones. If $Y_1,\ldots,Y_n$ are i.i.d. Beta(a,b), see Ex. 1.19, find expressions for the method of moments estimators for a,b.

(c). Suppose $Y_1,\ldots,Y_n$ are i.i.d. from the uniform distribution on $[0,b]$. Find the moment estimator for b. Then assume that the data are from a uniform om $[a,b]$, with both parameters unknown. Find the moment estimators.

(d). Define now $h_j(\theta) = \mathrm{E}_\theta\, Y^j$ and $N_j = n^{-1}\sum_{i=1}^n Y_i^j$, for $j = 1,\ldots,p$. The 'method of direct moments' is to solve the equations $h_j(\theta) = N_j$ for $j = 1,\ldots,p$. For $p = 3$, set up the two systems of three equations with three unknowns, i.e. $g_j(\theta) = M_j$ and $h_j(\theta) = N_j$ for $j = 1,2,3$, and show that the solutions $\widehat\theta = (\widehat\theta_1,\widehat\theta_2,\widehat\theta_3)$ are the same. Show in fact that the two methods, fitting centralised or direct methods, are identical, in the general case, for $p \geq 2$.

(e). For a given application one may choose one's method based on convenience and practicality. Sometimes formulae for the direct moments are more easily found that those for the centralised moments. The case for using centralised moments is partly numerical safety; the M_j numbers may be much smaller than the N_j. Exemplify this by seeting up the two equations, with two unknowns, both for the centralised moments and for the direct moments, for the case of the Beta(a,b) distribution.

(f). Generate $n = 100$ data points via the equation $y_i = [\exp\{a(\xi+\sigma N_i)\}-1]/a$, for say $(a,\xi,\sigma) = (0.33, 0.55, 0.77)$, with the N_i being standard normal. This is a skewed

extension of the usual normal model, which corresponds to $a \to 0$ here. Find formulae for the first three moments for this distribution. From your data, use the method of moments to estimate the three parameters.

Ex. 3.27. *Moment fitting estimators for the Gamma distribution.* We now apply the moment matching principle of Ex. 3.26 to the Gamma model with parameters (a, b), with density proportional to $y^{a-1} \exp(-by)$ for y positive; see Ex. 1.9, where we also give the mean, variance, skewness, kurtosis.

(a). With $\bar{Y}$ and V_n the usual sample mean and sample variance, find explicit formulae for the moment estimators $\widehat{a}, \widehat{b}$; show that $\widehat{a} = \bar{Y}^2/V_n$ and $\widehat{b} = \widehat{Y}/V_n$.

(b). Use skewness and kurtosis formulae, in combination with Ex. 2.70, to show that

$$\begin{pmatrix} \sqrt{n}(\bar{Y} - a/b) \\ \sqrt{n}(V_n - a/b^2) \end{pmatrix} \to_d \begin{pmatrix} U \\ W \end{pmatrix} \sim \mathrm{N}_2(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} a/b^2 & 2a/b^3 \\ 2a/b^3 & (2+6/a)a^2/b^4 \end{pmatrix}).$$

Transform this, via the delta method, to find the limit distribution for the moment estimators.

Ex. 3.28. *Moment method estimation in a symmetric Beta.* Consider an i.i.d. dataset $Y_1, \ldots, Y_n$ from the symmetric Beta(a, a), with sample mean $\bar{Y}$ and sample variance $V_n = (n-1)^{-1} \sum_{i=1}^n (Y_i - \bar{Y})^2$.

(a). Show that Y_i has mean $\frac{1}{2}$ and variance $\sigma^2 = \frac{1}{4}/(2a+1)$. Explain that the moment method with the sample mean does not work in this situation, but that the moment method using the sample variance leads to $\widehat{a} = 1/(8V_n) - \frac{1}{2}$. Show that it is consistent.

(b). Via patient integration, show that the kurtosis is

$$\mathrm{kurt} = \int_0^1 \frac{(x-\frac{1}{2})^4}{\sigma^4} \frac{\Gamma(2a)}{\Gamma(a)^2} \{x(1-x)\}^{a-1} \, \mathrm{d}x - 3 = -\frac{6}{2a+3}.$$

Then use Ex. 2.46 to show that $\sqrt{n}(V_n - \sigma^2) \to_d \mathrm{N}(0, \tau^2)$, with

$$\tau^2 = \sigma^4(2 + \mathrm{kurt}) = \frac{1/16}{(2a+1)^2} \frac{4a}{2a+3}.$$

(c). Use the delta method to find the limiting distribution of $\sqrt{n}(\widehat{a} - a)$, and explain how this may be used to construct a 90 percent confidence interval for a. (xx later: compare variance to that of the ML method, with efficiency loss. check if the following is ok. xx)

$$\kappa^2 = (2a+1)^2 a/(2a+3).$$

Ex. 3.29. *Moment method estimators for the exponential family.* Consider an exponential family type model, studied in Ex. 1.55, with density of the form $f(y, \theta) = \exp\{\theta^{\mathrm{t}} T(y) - k(\theta)\} h(y)$. Here $T(y) = (T_1(y), \ldots, T_p(y))^{\mathrm{t}}$ is a collection of data functions. As we saw in the exercise pointed to, many classical models are special cases. Now suppose $Y_1, \ldots, Y_n$ are i.i.d. from such a model.

(a). Spell out the basic moment matching method, fitting empirical and model based moments for $Y_i, Y_i^2, \ldots, Y_i^p$.

(b). The moment matching idea is however flexible enough to allow us to choose other data functions than those associated with $Y_i, \dots, Y_i^p$. Show that matching the moments of $T_1(Y_i), \dots, T_p(Y_i)$ corresponds to solving the equations

$$\bar{T}_j = \xi_j(\theta) = \partial k(\theta)/\partial\theta_j \quad \text{for } j = 1, \dots, p,$$

with $\bar{T}_j = n^{-1}\sum_{i=1}^n T_j(Y_i)$.

(c). Consider the setup of Ex. 3.27, with the Gam(a, b) distribution. There we studied moment matching estimators based on sample mean $\bar{Y}$ and sample variance V_n. Show that the principle above, using the exponential family structure, leads to fitting $\bar{Y} = a/b$ and $n^{-1}\sum_{i=1}^n \log Y_i = \psi(a) - \log b$.

Ex. 3.30. *Mean and variance for background distribution for random sums.* Suppose there is a hidden background machine drawing first the number of dice N and then reporting the sum $Z = X_1 + \cdots + X_N$ of outcomes from having thrown those N dice (without reporting N).

(a). For the individual random X_i, show that the mean and variance are $\xi = 3.50$ and $\sigma^2 = 35/12 = 2.9167$. Use results from Ex. 1.40 to find formulae for the mean and variance of Z.

(b). If $Z_1, \dots, Z_m$ are observed, from such a two-layer-random machine, with mean 33.33 and standard deviation 12.12, estimate the mean and standard deviation for N.

Quantile matching methods

Ex. 3.31. *Quantile fitting estimators.* An alternative to the method of moments, described in Ex. 3.26, is to fit empirical and theoretical quantiles (actually in several ways). If $Y_1, \dots, Y_n$ are i.i.d. from a density $f(y, \theta)$, for a parameter vector of length p, with quantiles $Q(r, \theta) = F^{-1}(r, \theta)$, we choose quantile levels $r_1 < \cdots < r_p$, and solve the p equations $Q_n(r_j) = Q(r_j, \theta)$ with respect to the p unknown parameters, where $Q_n(r) = F_n^{-1}(r)$ is the empirical quantile.

method of quantiles

(a). Suppose the distribution to be fitted has c.d.f. $F(y) = y^\theta$ on the unit interval. Find the estimator corresponding to fitting the empirical to the theoretical median. Starting with the limit distribution for the median, see Ex. 3.20, find the limit distribution for $\sqrt{n}(\widehat{\theta} - \theta)$. More generally, find the estimator $\widehat{\theta}_r$ corresponding to fitting the r level quantile, and then the limit distribution for $\sqrt{n}(\widehat{\theta}_r - \theta)$. For large n, which value of r gives the most precise estimator?

(b). Consider $Y_1, \dots, Y_n$ from the location Cauchy density $f_0(y - \theta)$, with $f_0(x) = (1/\pi)/(1+x^2)$ the standard Cauchy. Its c.d.f. is $F_0(x) = \frac{1}{2} + (1/\pi)\arctan x$, see Ex. 1.18. Show that the r level quantile is $\mu + F_0^{-1}(r)$, and that this leads to the estimator $\widehat{\mu}_r = Q_n(r) - F_0^{-1}(r)$. Find the limit distribution for $\sqrt{n}(\widehat{\mu}_r - \mu)$. What quantile level r leads to the sharpest estimator?

(c). For data $Y_1, \dots, Y_n$ from a normal (ξ, σ^2), find estimators for the two parametrs using the median and interquartile range. Find the limit distributions. More generally, construct an estimator $\hat{\sigma}_r$ for σ based on the quantile difference $Q_n(1-r) - Q_n(r)$. Find the limit distribution and identify the best value of r.

Ex. 3.32. *Moment fitting and quantile fitting for the Weibull.* As a general illustration of moment and quantile fitting estimation methods, consider the Weibull distribution with c.d.f. $F(t) = 1 - \exp\{-(t/a)^b\}$ for $t \geq 0$, see Ex. 1.59.

(a). Take e.g. $(a, b) = (3.33, 1.44)$, and simulate $n = 100$ realisations. (i) Compute average and standard deviation for these, and compute estimates $(\widehat{a}_m, \widehat{b}_m)$. (ii) From 0.25 and 0.75 quantiles, fit the two relevant equations to compute $(\widehat{a}_q, \widehat{b}_q)$. Take the trouble to display three curves, the correct underlying cumulative hazard function $A(t) = (t/a)^b$ along with the two estimated versions.

(b). Repeat the experiment many times, to see how close the two $(\widehat{a}, \widehat{b})$ is to (a, b). Also, as an instance of a focused question, check how close the two median estimates $\widehat{m} = \widehat{a}(\log 2)^{1/\widehat{b}}$ are to the real median? Which of the two estimating schemes is best? We should point here to the likelihood methodologies of Ch. 5; the maximum likelihood method will be the winning strategy, beating both moment and quantile fitting, under model conditions.

Minimum sum of squares and linear regression

Ex. 3.33. *Linear regression and least squares estimation.* Consider observed pairs (x_i, Y_i) for $i = 1, \ldots, n$, where Y_i conditionally on covariate x_i is modelled to have mean $a_0 + bx_i$ and common variance σ^2. This is classical linear regression, widely used in theoretical and applied statistics, most often analysed and used with the additional assumption that the Y_i are normally distributed. Importantly, the model is extended to the case of multiple covariates below, see Ex. 3.34, 3.35, with more to come regarding statistical testing (xx in Ex. 4.36 and one more xx). To see some of these classical methods in action, check out Stories iii.1 and iv.1.

(a). It is helpful to reparametrise the regression line from $a_0 + bx_i$ to $a + b(x_i - \bar{x})$. Show that minimising the sum of squares $Q(a, b) = \sum_{i=1}^n \{y_i - a - b(x_i - \bar{x})\}^2$ leads to

$$\widehat{a} = (1/n)\sum_{i=1}^n y_i = \bar{y}, \quad \widehat{b} = \sum_{i=1}^n (x_i - \bar{x})y_i/M_n, \quad \text{with } M_n = \sum_{i=1}^n (x_i - \bar{x})^2.$$

(b). Show that $\widehat{a}$ and $\widehat{b}$ are unbiased, with zero covariance, and variances σ^2/n and σ^2/M_n.

(c). Let $Q_0 = \min_{a,b} Q(a, b) = \sum_{i=1}^n \{y_i - \widehat{a} - \widehat{b}(x_i - \bar{x})\}^2$ be the minimum sum of squares. Show that

$$Q(a, b) = \sum_{i=1}^n \{y_i - a - b(x_i - x)\}^2 = Q_0 + n(\widehat{a} - a)^2 + M_n(\widehat{b} - b)^2.$$

Use this to show that $\widehat{\sigma}^2 = Q_0/(n-2)$ is an unbiased estimator of σ^2.

(d). We have found natural and unbiased estimators for a, b, σ^2, without yet making assumptions of the underlying distributions for Y_i, beyond means and variance. Assume now, however, that the distributions are normal, so that $Y_i \sim \mathrm{N}(a + b(x_i - \bar{x}), \sigma^2)$. Show that (i) $\widehat{a} \sim \mathrm{N}(a, \sigma^2/n)$; (ii) $\widehat{b} \sim \mathrm{N}(b, \sigma^2/M_n)$; (iii) $\widehat{\sigma}^2 \sim \sigma^2\chi^2_m/m$, with $m = n - 2$; and (iv) that these three statistics are mutually independent. For this last point, it may be helpful to follow the line of proof for independence of sample mean and sample variance for normal data, used in Ex. 1.49, now with an orthonormal matrix A with first row $(1/\sqrt{n}, \ldots, 1/\sqrt{n})$ and second row $((x_1 - \bar{x})/M_n^{1/2}, \ldots, (x_n - \bar{x})/M_n^{1/2})$.

(e). Construct confidence intervals for b, for σ, and for $\mathrm{E}\,(Y \mid x_0) = a + b(x_0 - \bar{x})$, the mean value at a given position x_0.

Ex. 3.34. *Linear multiple regression and least squares.* The celebrated linear multiple regression model remains a cornerstone success story of theoretical and applied statistics. It is a bag of tools for investigating the extent to which covariates x influence the outcomes of certain interest variables Y. The standard formulation of the model is as follows. The data collected can be organised into (x_i, Y_i), for individuals or objects $i = 1, \ldots, n$, where $x_i = (x_{i,1}, \ldots, x_{i,p})^{\mathrm{t}}$ is of dimension p and y_i of dimension one. The model then postulates that

$$Y_i = x_i^{\mathrm{t}}\beta + \varepsilon_i = x_{i,1}\beta_1 + \cdots + x_{i,p}\beta_p + \varepsilon_i \quad \text{for } i = 1, \ldots, n,$$

where the ε_i are i.i.d. from the normal $\mathrm{N}(0, \sigma^2)$. Thus there are $p+1$ parameters at work here, the regression coefficients $\beta = (\beta_1, \ldots, \beta_p)^{\mathrm{t}}$ and the error distributiuon standard deviation σ. – Note that the very classical case of $y_i = a + bx_i + \varepsilon_i$, associated with a scatterplot of (x_i, Y_i), is a special case; see Ex. 3.33.

(a). With Y the vector of Y_i, ε the vector of ε_i, and X the $n \times p$ matrix having $(x_{i,1}, \ldots, x_{i,p})$ as its row i, show that

$$Y = X\beta + \varepsilon \sim \mathrm{N}_n(X\beta, \sigma^2 I),$$

with I the $n \times n$ identity matrix. This is a practical and compact linear algebra version of the model formulation. We do assume that X is of full rank p, so that the symmetric matrix $X^{\mathrm{t}}X$ is invertible. This amounts to there being at least p linearly independent covariate vectors in the X matrix; in particular, we must have $n \geq p$ to identify the β_j coefficients directly from data.

(b). The least squares estimator $\widehat{\beta}$ is the minimiser of $Q(\beta) = \|Y - X\beta\|^2 = \sum_{i=1}^n (Y_i - x_i^{\mathrm{t}}\beta)^2$. Show that $\sum_{i=1}^n (Y_i - x_i^{\mathrm{t}}\widehat{\beta})x_i = 0$. With Σ_n the $p \times p$ matrix $n^{-1}\sum_{i=1}^n x_i x_i^{\mathrm{t}} = n^{-1}X^{\mathrm{t}}X$, show that

$$\widehat{\beta} = (X^{\mathrm{t}}X)^{-1}X^{\mathrm{t}}Y = \Sigma_n^{-1} n^{-1} \sum_{i=1}^n x_i Y_i.$$

Prove also that $Q(\beta) = Q_0 + n(\widehat{\beta} - \beta)^{\mathrm{t}}\Sigma_n(\widehat{\beta} - \beta)$, with $Q_0 = \min_{\text{all } \beta} Q(\beta) = \sum_{i=1}^n \widehat{\varepsilon}_i^2$, writing $\widehat{\varepsilon}_i = Y_i - x_i^{\mathrm{t}}\widehat{\beta}$ for the estimated residuals.

(c). Show that $\widehat{\beta}$ is unbiased and that its variance matrix can be written $\sigma^2(X^{\mathrm{t}}X)^{-1} = (\sigma^2/n)\Sigma_n^{-1}$.

(d). Show also that $\widehat{\beta}$ has a multinormal distribution, so that in fact $\widehat{\beta} \sim \mathrm{N}_p(\beta, (\sigma^2/n)\Sigma_n^{-1})$. This is the key result about the least squares estimators. We also need precise information for estimating σ; see Ex. 3.35.

Ex. 3.35. *The residuals and their variances.* The setup is an in the previous Ex. 3.34, the $Y \sim \mathrm{N}_n(X\beta, \sigma^2 I)$ linear regression model. Above we focused on the least squared method and the ensuing properties for the estimators of the regression coefficients, and found $\widehat{\beta} \sim \mathrm{N}_p(\beta, (\sigma^2/n)\Sigma_n^{-1})$. We also need to deal carefully with estimators of σ, the residual standard deviation, also since we encounter statistics of the type $(\widehat{\beta}_j - \beta_j)/\widehat{\sigma}$.

(a). From the basic $Y = X\beta + \varepsilon$ we may define the estimated residuals as

$$\widehat{\varepsilon} = Y - X\widehat{\beta} = (I - H)Y, \quad \text{where} \quad H = X(X^{\mathrm{t}}X)^{-1}X^{\mathrm{t}},$$

the so-called hat matrix, of size $n \times n$. Show that H is symmetric and idempotent, which means that $H^{\mathrm{t}} = H$ and $H^2 = H$. This also implies $(I - H)H = 0$.

(b). Now consider the random minimum achieved by the $Q(\beta)$ which was used in the least squares operation,

$$Q_0 = \min\{Q(\beta)\colon \text{all } \beta\} = Q(\widehat{\beta}) = \|Y - X\widehat{\beta}\|^2 = \sum_{i=1}^{n}(Y_i - x_i^{\mathrm{t}}\widehat{\beta})^2.$$

The main result, arrived at below, is that $Q_0/\sigma^2 \sim \chi^2_m$, with degrees of freedom $m = n-p$, and that Q_0 is independent of $\widehat{\beta}$. Show first that

$$\begin{pmatrix} X\widehat{\beta} \\ \widehat{\varepsilon} \end{pmatrix} = \begin{pmatrix} HY \\ (I-H)Y \end{pmatrix} \sim \mathrm{N}_{2n}(\begin{pmatrix} X\beta \\ 0 \end{pmatrix}, \sigma^2 \begin{pmatrix} H & 0 \\ 0 & I-H \end{pmatrix}).$$

In particular, these two random vectors are independent; also, $Q_0 = \|\widehat{\varepsilon}\|^2 = Y^{\mathrm{t}}(I-H)Y$ is consequently independent of $X\widehat{\beta}$.

(c). Show that $(I-H)X = 0$, which implies $\widehat{\varepsilon} = (I-H)Y = (I-H)(Y-X\beta) = (I-H)\varepsilon$ and hence $Q_0 = \varepsilon^{\mathrm{t}}(I-H)\varepsilon$. We also reach the simple idenity

$$\|\varepsilon\|^2 = \varepsilon^{\mathrm{t}}H\varepsilon + \varepsilon^{\mathrm{t}}(I-H)\varepsilon,$$

where the left-hand side is a $\sigma^2\chi^2_n$ and the two terms on the right-hand side being independent. Show that the first term on the right-hand side is a $\sigma^2\chi^2_p$. Via independence and a moment-generating function argument show then that $Q_0 \sim \sigma^2\chi^2_{n-p}$ (this is essentially the basic trick of Ex. 1.48).

(d). With $Q_0 \sim \sigma^2\chi^2_m$, with degrees of freedom $m = n - p$, the natural estimator is $\widehat{\sigma}^2 = Q_0/m$. Show that it is unbiased and find its variance. Note that things we have learned in Ex. 1.49 and other exercises related to the empirical variance for a normal sample carry over to the the present linear regression model, simply with $n - p$ instead of $n - 1$.

Ex. 3.36. *Separate and joint confidence intervals.* Suppose a vector Y is observed from some $\mathrm{N}_p(\theta, \Sigma)$, with unknown θ and known positive definite Σ. How can we set confidence intervals for the individual θ_j, for given linear combinations, and then for all linear combinations jointly?

(a). Consider a given linear combination $\gamma = c^{\mathrm{t}}\theta = c_1\theta_1 + \cdots + c_p\theta_p$. Explain that $\widehat{\gamma} = c^{\mathrm{t}}Y \sim \mathrm{N}(\gamma, c^{\mathrm{t}}\Sigma c)$, and that the canonical 95 percent confidence interval for γ is $\widehat{\gamma} \pm 1.96\,\kappa(c)$, with $\kappa(c) = (c^{\mathrm{t}}\Sigma c)^{1/2}$. Explain what this entails for the components $\theta_1, \ldots, \theta_p$.

(b). Leaving the situation above aside for a minute, prove the following mathematical fact: for given x, and a given symmetric positive definite $p \times p$ matrix A, the maximum of $r(c) = c^{\mathrm{t}}x/(c^{\mathrm{t}}Ac)^{1/2}$, taken over all c, is equal to $(x^{\mathrm{t}}A^{-1}x)^{1/2}$. Do this by first seeing that the ratio $r(c)$ does not depend on the length of c, so the problem is equivalent to finding the maximum of $c^{\mathrm{t}}x$ over all c such that $c^{\mathrm{t}}Ac = 1$; you may proceed using Lagrange multipliers and a bit of linear algebra.

(c). Suppose now that one wishes to have intervals for *all* linear combinations, with confidence level 95 percent to hold jointly. To what extent does one need to stretch the $c^{\rm t}Y \pm 1.96\,\kappa(c)$ intervals, to secure such a joint-for-all 95 percent confidence? Show that

$$V = \max_{\text{all } c} \frac{|c^{\rm t}Y - c^{\rm t}\theta|}{(c^{\rm t}\Sigma c)^{1/2}} = \{(Y-\theta)^{\rm t}\Sigma^{-1}(Y-\theta)\}^{1/2},$$

and that $V^2 \sim \chi^2_p$. Explain then that with $z_p = \Gamma_p^{-1}(0.95)^{1/2}$, involving the upper 0.05 quantile of the χ^2_p, the intervals $I(c) = c^{\rm t}Y \pm z_p\,\kappa(c)$ have the desired property, that $\Pr\{c^{\rm t}\theta \in I(c) \text{ for all } c\} = 0.95$. Check how this z_p factor grows with p. For $p = 5$, for example, the separate confidence intervals need to be made 1.697 times wider, in order to secure that the confidence level is 95 percent for all linear combinations.

(d). (xx can point to Bonferroni method. in notes: Hjort (1988b). xx)

Ex. 3.37. *Confidence intervals for key parameters in linear regression models.* Consider the general linear regression setup of Ex. 3.34–3.35, where $Y_i = x_i^{\rm t}\beta + \varepsilon_i$ for $i = 1,\ldots,n$, and the ε_i being i.i.d. $\mathrm{N}(0,\sigma^2)$. The compact version of this is $Y \sim \mathrm{N}_n(X\beta, \sigma^2 I)$. We use notation from these exercises below, with least squares estimator $\widehat\beta$, $\widehat\sigma^2 = Q_0/(n-p)$, etc.

(a). Construct a 90 percent confidence interval for σ.

(b). From the general results obtained above, show that $\widehat\beta_j \sim \mathrm{N}(\beta_j, \sigma^2 r_j/n)$, with r_j the diagonal (j,j) element of Σ_n^{-1}. Show from this that $t_j = \sqrt{n}(\widehat\beta_j - \beta_j)/(r_j^{1/2}\widehat\sigma)$ has a t_m distribution, with degrees of freedom $m = n-p$, and construct a confidence interval for β_j from this. If a 99 percent confidence interval for β_j is outside zero, how can this be interpreted and used?

(c). More generally, consider $\gamma = c^{\rm t}\beta = c_1\beta_1 + \cdots + c_p\beta_p$, a linear combnination of the regression coefficients. Find the distribution of $\widehat\gamma = c^{\rm t}\widehat\beta$, and show that the interval $c^{\rm t}\widehat\beta \pm t_{m,0.975}\,\widehat\kappa(c)/\sqrt{n}$ covers the true $c^{\rm t}\beta$ with probability 95 percent, when $\widehat\kappa(c) = (c^{\rm t}\Sigma_n^{-1}c)^{1/2}\widehat\sigma$, with $t_{m,0.975}$ the upper 0.025 quantile of the t_m, with $m = n-p$ degrees of freedom. This may be used to assess the size of $\beta_1 - \beta_2$, of $\beta_1 - \frac12(\beta_2+\beta_3)$, and similar constrasts.

(d). The situation is different if one wishes to have confidence intervals for *all* linear combinations, jointly, with a guarantee that the overall coverage is e.g. 95 percent. Show using Ex. 3.36 that

$$V_n = \max_{\text{all } c} \frac{\sqrt{n}|c^{\rm t}\widehat\beta - c^{\rm t}\beta|}{(c^{\rm t}\Sigma_n^{-1}c)^{1/2}\widehat\sigma} = \{n(\widehat\beta-\beta)^{\rm t}\Sigma_n(\widehat\beta-\beta)\}^{1/2}/\widehat\sigma,$$

and that $V_n^2 \sim F_{p,m}$. Use this to show that the intervals $I_n(c) = c^{\rm t}\widehat\beta \pm z_{p,m}\,\widehat\kappa(c)/\sqrt{n}$, with $z_{p,m}$ the 0.975 quantile of the $F_{p,m}$ distribution, have the property that $\Pr\{c^{\rm t}\beta \in I_n(c) \text{ for all } c\} = 0.95$.

(e). Use the result about V_n to construct a confidence region, actually a confidence ellipsoid, for the full β vector.

Ex. 3.38. *Predicting the next y.* Suppose linear regression analysis has been carried out, for a given dataset $(x_1,y_1),\ldots,(x_n,y_n)$. How can we predict what happens with

Y_0, associated with another covariate vector x_0? This could be for an individual outside the dataset, or in a time context, speculating about the next datapoint in a sequence. An illustration of methods given below is in Story iv.1.

(a). Assume the regression data are of the form $y_i \sim \mathrm{N}(x_i^{\mathrm{t}}\beta, \sigma^2)$ for $i = 1, \ldots, n$, as with Ex. 3.34, and consider a new x_0, for which the not yet unobserved Y_0 is independent of the other data, and with distribution $\mathrm{N}(x_0^{\mathrm{t}}\beta, \sigma^2)$. Show that $\widehat{y}_0 = x_0^{\mathrm{t}}\widehat{\beta} \sim \mathrm{N}(x_0^{\mathrm{t}}\beta, x_0^{\mathrm{t}}\Sigma_n^{-1}x_0\,\sigma^2/n)$, and use this to form a confidence interval for the mean of Y_0, as opposed to for Y_0 itself.

(b). Then show that

$$Y_0 - \widehat{y}_0 \sim \mathrm{N}(0, \sigma^2(1 + n^{-1}x_0^{\mathrm{t}}\Sigma_n^{-1}x_0/n)).$$

Construct a prediction confidence interval for Y_0 based on this. Comment on the situation with a large n, and on the difference between confidence for Y_0 and its mean.

(c). Then consider the classic linear regression case with only one variable studied in Ex. 3.33, with $Y_i \sim \mathrm{N}(a + b(x_i - \bar{x}), \sigma^2)$. For predicting a not yet observed $Y_0 \sim \mathrm{N}(a + b(x_0 - \bar{x}), \sigma^2)$, demonstrate that

$$Y_0 - \{\widehat{a} + \widehat{b}(x_0 - \bar{x})\} \sim \mathrm{N}(0, \sigma^2\{1 + 1/n + (x_0 - \bar{x})^2/M_n\}),$$

with $M_n = \sum_{i=1}^n (x_i - \bar{x})^2$. Show that this leads to the 95 percent confidence band

$$Y(x_0) \in \widehat{a} + \widehat{b}(x_0 - \bar{x}) \pm c\,\widehat{\sigma}\{1 + 1/n + (x_0 - \bar{x})^2/M_n\}^{1/2},$$

for an appropriate range of x_0 values, with $c = t_{0.975,n-2}$ the t upper quantile. Illustrate this with a simulated dataset, to see both where the band is tight and where it becomes very broad. One learns that it is easier to predict outcomes related to x_0 closer to the middle of the covariate distribution.

Ex. 3.39. *Inference for mean by standard deviation.* (xx point to Ex. 7.4 with clean CD. xx) Consider $Y_1, \ldots, Y_n$ being i.i.d. from $\mathrm{N}(\xi, \sigma^2)$. Estimation and precise inference for ξ and for σ separately is classic terrain, via the t and the chi-squared distribution, and we have seen in Ex. 3.8 how we can also deal with quantiles $\xi + z_q\sigma$ in a precise fashion. For more complicated parameters $g(\xi, \sigma)$ we might turn to large-sample approximations, as per Ex. 2.51. This exercise concerns however precise inference for the parameter $\kappa = \xi/\sigma$.

(a). We start with the following cleanly stated problem: suppose $T \sim t_m(a)$, with the noncentral t, as per Ex. 1.51, with known m but unknown a. With t_{obs} the observed value of T, we form a confidence interval by collecting together all values a for which t_{obs} is not surprising, defined by $I = \{a\colon 0.05 \leq G_m(t_{\mathrm{obs}}, a) \leq 0.95\}$. Show that indeed $\Pr_a(a \in I) = 0.90$. Again, the $G_m(t, a)$ is a somewhat complicated function, but is numerically available in packages.

(b). Coming back to the normal sample, write as usual $\widehat{\xi} = \bar{Y}$ and $\widehat{\sigma}^2 = (n-1)^{-1}\sum_{i=1}^n (Y_i - \bar{Y})^2$, and form $\widehat{\kappa} = \widehat{\xi}/\widehat{\sigma}$. Show that $X = \sqrt{n}\widehat{\kappa} \sim t_m(\sqrt{n}\kappa)$, with $m = n - 1$ degrees of freedom. Show that

$$I = \{\kappa\colon 0.05 \leq G_m(\sqrt{n}\widehat{\kappa}, \sqrt{n}\kappa) \leq 0.95\}$$

is a 90 percent confidence interval for κ. In practice we find the interval by computing and monitoring the $G_m(\sqrt{n}\widehat{\kappa}, \sqrt{n}\kappa)$, as a function of κ.

(c). Choose your favourite (ξ, σ), generate a normal sample of size $n = 100$, and compute (i) the interval $\widehat{\kappa} \pm 1.645\,(1+\widehat{\kappa}^2)^{1/2}/\sqrt{n}$, from Ex. 2.51, and (ii) the exact interval from the recipe above. Compare and comment.

(d). Consider now a linear regression setup, as in Ex. 3.34–3.35, with independent $y_i = x_i^{\mathrm{t}}\beta + \varepsilon_i$ for $i = 1, \ldots, n$. With some given linear combination $c^{\mathrm{t}}\beta$ of the regression components, consider $\kappa = c^{\mathrm{t}}\beta/\sigma$. In the standard notation of the exercises pointed to, show that $\widehat{\kappa} = c^{\mathrm{t}}\widehat{\beta}/\widehat{\sigma}$ is such that $t = \sqrt{n}\widehat{\kappa} \sim t_m(\sqrt{n}\kappa)$. Explain how a precise 90 percent confidence interval can be set for κ.

(e). In such a linear regression setup, suppose x_a and x_b are certain covariate values, outside of $x_1, \ldots, x_n$, associated with unknown or not yet observed $Y_a = x_a^{\mathrm{t}}\beta + \varepsilon_a$ and $Y_b = x_b^{\mathrm{t}}\beta + \varepsilon_b$. Show that

$$p = \Pr(Y_b > Y_a) = \Phi(\kappa/\sqrt{2}), \quad \text{with } \kappa = (x_b - x_a)^{\mathrm{t}}\beta/\sigma,$$

Explain that you via a clear 90 percent confidence interval for κ also have a 90 percent confidence interval for p. (xx brief pointer to New Haven Story iv.1, where nils will put in estimate and interval for p, the probability that 1976 is warmer than 1907, based on data from 1912 to 1971. more on the details there, prediction vs. the mean change. xx)

Ex. 3.40. *Linear regression outside normality.* The aim here is to show and appreciate that the classical coefficient estimators in linear regression setups are still approximately normal, even when the error terms distribution is not normal. That this is so is essentially thanks to basic large-sample theory, as partly summarised in Ex. 3.11, and specifically to the Lindeberg type theorems of Ch. 2; see in particular Ex. 2.41. (xx pointer to other regression large-sample results in Ch. 5. xx)

(a). We first deal with a simple setup with a single regression coefficient. Suppose $y_i = x_i\beta + \varepsilon_i$ for $i = 1, \ldots, n$, with covariates x_i and error terms ε_i being i.i.d. from a zero-mean distribution with finite variance σ^2. Show that the estimator minimising $Q_n(\beta) = \sum_{i=1}^n (y_i - x_i\beta)^2$ is $\widehat{\beta} = \sum_{i=1}^n x_i y_i / M_n$, where $M_n = \sum_{i=1}^n x_i^2$. Show further that $\widehat{\beta}$ is unbiased with variance σ^2/M_n.

(b). Then consider $Z_n = M_n^{1/2}(\widehat{\beta} - \beta)$. Show that it has zero mean and variance σ^2, and that it can be written $\sum_{i=1}^n (x_i/M_n^{1/2})\varepsilon_i$. Deduce that $\widehat{\beta}$ is approximately normal, even if the ε_i are not normal, provided merely that $D_n = \max_{i \le n} |x_i|/M_n^{1/2} \to 0$. If in particular $(1/n)\sum_{i=1}^n x_i^2$ stays bounded, then the natural condition is $(1/\sqrt{n})\max_{i \le n} |x_i| \to 0$.

(c). Then consider the general linear regression model $Y_i = x_i^{\mathrm{t}}\beta + \varepsilon_i$ of Ex. 3.34, with the x_i being p-dimensional covariate vectors and β a p-dimensional vector of regression coefficients. We take the ε_i to be i.i.d. with mean zero and finite variance σ^2, but do not stipulate normality. The least squares estimator is $\widehat{\beta} = \Sigma_n^{-1}(1/n)\sum_{i=1}^n x_i Y_i$ with $\Sigma_n = n^{-1}\sum_{i=1}^n x_i x_i^{\mathrm{t}}$. It is unbiased with variance matrix $(\sigma^2/n)\Sigma_n^{-1}$, assumed to have full rank. Assume $\Sigma_n \to \Sigma$, a full rank matrix. Show that $Z_n = \sqrt{n}(\widehat{\beta} - \beta) \to_d \mathrm{N}_p(0, \Sigma)$, under the condition $R_n = (1/\sqrt{n})\max_{i \le n} \|x_i\| \to 0$.

(d). Assume now that the x_i are drawn i.i.d. from a distribution over the covariate space, with finite variance matrix. Show that $R_n \to_{\mathrm{pr}} 0$. (xx calibrate with what we have in 2.41. xx)

(e). Argue that the fine-tuned finite-sample confidence methods, developed in Ex. 3.34–3.35 under exact normality, continue to hold in the large-sample sense, even if the distributions are not normal. Give conditions for such results to hold.

(f). Suppose the same type of phenomenon is studied in both Denmark and Sweden, with the same meaning for Y and covariates x_1 and x_2. This leads to regression estimators $\widehat{\beta}_D$ and $\widehat{\beta}_S$ for the two analyses. Construct confidence intervals for $d_j = \beta_{D,j} - \beta_{S,j}$, for $j = 1, 2$.

Ex. 3.41. *Least squares estimation in other setups.* (xx to come. point is that minimising $Q(\theta) = \sum_{i=1}^n \{y_i - \xi_i(\theta)\}^2$ is a general principle, also outside linear regression. xx)

(a). (xx some easy cases. binomial. normal mean. poisson. xx)

(b). Assume $y_1, \ldots, y_n$ are i.i.d. and gamma distributed (a, b), with known a. Find the least squares estimator $\widehat{b}$. Then find the mean and variance of this estimator.

(c). (xx regression, with mean $a + bx_i$, then mean $a \exp(bx_i)$. xx)

Ex. 3.42. *Estimating with autocorrelated data.* Situations dealt with above have involved different models and questions, but have assumed independence between observations. In various setups, including those of time series, there is correlation, however, with consequences for estimators and their precision. The brief discussion below has implications for reading off estimates for key quantities in Markov Chain Monte Carlo settings in Bayesian analysis; see Ch. 6.

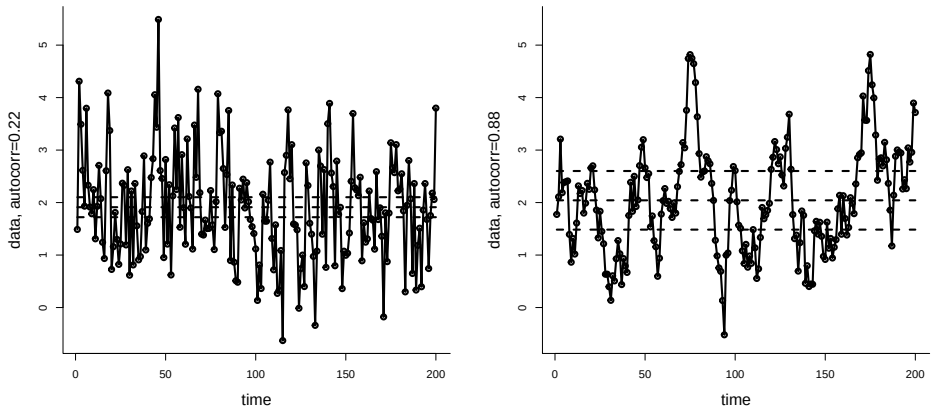


Figure 3.3: *Data with autocorrelation, with underlying* $\xi = 2.0$, $\sigma = 1.0$, *and then* $\rho = 0.22$ *(left panel) and* $\rho = 0.88$ *(right panel). The horizontal lines are for the estimate and the 95 percent confidence intervals for* ξ. *With higher autocorrelation, data are more connected, and estimates have larger variances.*

(a). Consider $V_1, \ldots, V_n$ i.i.d. with mean zero and variance one. From these we form a new series, with autocorrelation $\rho \in (-1, 1)$, as follows. Let $W_1 = V_1$ and then $W_i = \rho W_{i-1} + (1 - \rho^2)^{1/2} V_i$ for $i = 2, \ldots, n$. Show that these W_i have mean zero and variance one, but $\mathrm{cov}(W_i, W_j) = \rho^{j-i}$ for $i < j$. A time series, with three crucial parameters mean ξ, variance σ^2, and autocorrelation ρ, is then formed as $Y_i = \xi + \sigma W_i$.

(b). With such an autocorrelated time series $Y_1, \ldots, Y_n$, the natural estimator of ξ is still the sample average $\widehat{\xi} = \bar{Y}_n$, but its variance is affected. Show first in general terms that one for any $a \in (-1, 1)$ has

$$n^{-1} \sum_{1 \le i < j \le n} a^{j-i} = \frac{a}{1-a}\Big\{1 - \frac{1}{n} - \frac{1}{n}\frac{a}{1-a}(1-a^{n-1})\Big\} = \frac{a}{1-a}\{1 + O(1/n)\}.$$

Show that this leads to

$$\operatorname{Var} \bar{Y}_n = \frac{1}{n^2}\sum_{i,j} \operatorname{cov}(Y_i, Y_j) = \frac{\sigma^2}{n}\Big\{1 + 2\frac{\rho}{1-\rho} + O(1/n)\Big\} = \frac{\sigma^2}{n}\Big\{\frac{1+\rho}{1-\rho} + O(1/n)\Big\}.$$

Conclude that $\bar{Y}_n$ is unbiased and consistent for ξ, but with this adjusted variance.

(c). We're not aiming here for maximum generality (which would take us into the terrain of stationary time series); the main point to convey is that positively autocorrelated data still allow consistent estimation and inference for key quantities, but with bigger variances, wider confidence intervals, etc. Assume therefore that the V_i above are normal, implying that $(W_1, \ldots, W_n)$ and $(Y_1, \ldots, Y_n)$ have multinormal distributions. In addition to the basic moment structure, with $\mathrm{E}\, Y_i = \xi$, $\operatorname{Var} Y_i = \sigma^2$, $\operatorname{corr}(Y_i, Y_j) = \rho^{|j-i|}$, we can read off further variances and correlations, for Y_i^2, for $Y_i Y_{i+1}$, etc. Show that $\widehat{\sigma}^2 = n^{-1}\sum_{i=1}^n (Y_i - \bar{Y}_n)^2$ is consistent for σ^2, and that

$$\widehat{\rho} = \frac{1}{n}\sum_{i=1}^{n-1} \frac{Y_i - \widehat{\xi}}{\widehat{\sigma}} \frac{Y_{i+1} - \widehat{\xi}}{\widehat{\sigma}}$$

is consistent for ρ. Show that this implies $\sqrt{n}(\bar{Y}_n - \xi)/\widehat{g} \to_d \mathrm{N}(0, \sigma^2)$, with factor $\widehat{g}^2 = (1+\widehat{\rho})/(1-\widehat{\rho})$. Explain that this leads to these confidence intervals for ξ, adjusted for autocorrelation: $\widehat{\xi} \pm 1.96\, \widehat{g}\widehat{\sigma}/\sqrt{n}$. Construct a version of Figure 3.3, with $n = 200$ autocorrelated observations, with higher autocorrelation in the right panel, associated with wider confidence interval for the mean.

(d). Statisticians A and B need to estimate a mean parameter ξ with a certain precision, from observations Y_i with this mean and variance σ^2. A manages to get hold of i.i.d. data, and finds an estimate with a confidence interval after $n_A = 1000$ observations. B utilises a different setup for generating the data, and has to make do with a chain with autocorrelation $\rho = 0.90$. How much bigger does his sample size n_B need to be, in order to have the same precision as that obtained by his colleague?

More material, leftovers (as of 12-March-2025)

Ex. 3.43. *Jackknife estimation of the bias.* (xx nils jotting down; need polish, then connection to influence functions in Ch5. xx) Suppose $Y_1, \ldots, Y_n$ are i.i.d. from F, and that an estimator $\widehat{\theta}$, symmetric in its arguments, is used to estimate the parameter $\theta = \theta(F)$. Write $\mathrm{E}_F\, \widehat{\theta} = \theta(F) + b(F)$, with $b = b(F)$ the bias. If there is an easy explicit formula for $b(F)$, we may carry out bias-correction, using $\widehat{\theta}^* = \widehat{\theta} - \widehat{b}$, with $\widehat{b} = b(F_n)$. The following is a more generic and generally applicable strategy for bias estimation and hence bias correction. Define $\widehat{\theta}_{(i)} = \widehat{\theta}(Y_1, \ldots, Y_{i-1}, Y_{i+1}, \ldots, Y_n)$, i.e. with estimation carried out on the dataset of size $n-1$ where datum i has been taken out. From these so-called jackknife vaues, compute the average $\widehat{\theta}_{(\cdot)}$, and then $\widehat{b}_{\mathrm{jack}} = (n-1)(\widehat{\theta}_{(\cdot)} - \widehat{\theta})$.

(a). First study the case of the simple mean, with $\widehat{\xi} = \bar{Y}_n$. Do the algebra to show that $\widehat{\xi}_{(i)} = (n-1)^{-1}(n\bar{Y}_n - Y_i)$, and that this implies $\widehat{b}_{\rm jack} = 0$; the estimator is already unbiased. Then work with the variance estimator $V_n = n^{-1}\sum_{i=1}^n (Y_i - \bar{Y}_n)^2$, which we know has a negative bias. Again, do the algebra to find an expression for the jackknifed $V_{n,(i)}$, then the bias estimator, which becomes $\widehat{b} = -(n-1)^{-1}V_n$. Explain that this leads to the repaired $V_n^* = (n-1)^{-1}\sum_{i=1}^n (Y_i - \bar{Y}_n)^2$, the unbiased sample variance.

(b). Consider the skewness functional $\gamma(F) = \int (y-\xi(F))^3\,{\rm d}F(y)$, with $\xi(F) = \int y\,{\rm d}F(y)$ the mean. Explain that this leads to the estimator $\widehat{\gamma} = \gamma(F_n) = n^{-1}\sum_{i=1}^n (Y_i - \bar{Y})^3$. With some algebraic efforts, show that $\mathrm{E}_F\,\widehat{\gamma} = (1 - 3/n + 2/n^2)\gamma(F)$. Thanks to this explicit formula we know how to carry out easy debiasing, leading to the unbuased $\widehat{\gamma}_u$, say. Check what the jackknife scheme does, and verify, perhaps supplemented with simulations, that it essentially succeeds in knocking off the $-(3/n)\gamma(F)$ part.

(c). To understand the general principles involved, suppose $\mathrm{E}_F\,\widehat{\theta} = \theta(F) + c_1(F)/n + c_2(F)/n^2 + \cdots$. Work out that

$$\widehat{\theta}_{\rm jack} = \widehat{\theta} - \widehat{b}_{\rm jack} = n\widehat{\theta} - (n-1)\widehat{\theta}_{(\cdot)} \text{ has mean } \theta(F) - \frac{c_2(F)}{n(n-1)} + O\Big(\frac{1}{n^3}\Big),$$

i.e. the leading $1/n$ bias term of the first estimator has been knocked off.

(d). (xx perhaps one more example, with simulations. pointers to influence functione and leave-one-out things in Ch11; there we should point back to this exercise. xx)

(e). Suppose in general terms that $\widehat{\theta}$ is estimating θ, with the usual type of limiting normal behaviour, $\sqrt{n}(\widehat{\theta} - \theta) \to_d {\rm N}(0, \tau^2)$, say. Consider then the perhaps bias adjusted estimator $\widehat{\theta}^* = \widehat{\theta} - \widehat{c}/n$, where $\widehat{c}$ is consistent for some c. Show that also $\sqrt{n}(\widehat{\theta}^* - \theta)$ has the same limiting distribution. Explain that this means that bias corrections of the above type do not change the limit, but might repair the bias. In terms of mean squared errors, translate this to $n\,\mathrm{E}\,(\widehat{\theta} - \theta)^2 \doteq c^2/n + \tau^2$; it pays off to reduce the bias, but for large samples this does not matter so much.

Ex. 3.44. *Successive ratios for the uniform order statistics process.* Consider an i.i.d. sample $U_1, \ldots, U_n$ from the uniform distribution on the unit interval. For quantile levels $p < q < r$, let $Q_{n,p}, Q_{n,q}, Q_{n,r}$ be the empirical quantiles, and form from these the ratios $R_n = Q_{n,q}/Q_{n,p}$, $S_n = Q_{n,r}/Q_{n,q}$, $T_n = 1/Q_{n,r}$.

(a). Explain via Ex. 3.21 that

$$\begin{pmatrix} \sqrt{n}(Q_{n,p} - p) \\ \sqrt{n}(Q_{n,q} - q) \\ \sqrt{n}(Q_{n,r} - r) \end{pmatrix} \to_d \begin{pmatrix} A \\ B \\ C \end{pmatrix} \sim {\rm N}_3(0, \Sigma),$$

and identify the elements of Σ. Show that

$$\begin{pmatrix} \sqrt{n}(R_n - q/p) \\ \sqrt{n}(S_n - r/q) \\ \sqrt{n}(T_n - 1/r) \end{pmatrix} \to_d \begin{pmatrix} R \\ S \\ T \end{pmatrix} = \begin{pmatrix} (1/p)\{-(q/p)A + B\} \\ (1/q)\{-(r/q)B + C\} \\ (1/r)\{-(1/r)C\}. \end{pmatrix}$$

Do the algebra to verify that R, S, T have variances

$$(1/p)^2\{p(1-q/p)^2 + q - p\},\ (1/q)^2\{q(1-r/q)^2 + r - q\},\ (1/r)^2\{r(1-1/r)^2 + 1 - r\},$$

and that they are in fact independent.

(b)**.** It is useful to translate the analysis above to the terminology and results for the order statistics processes terminology of (xx Ch9 xx). We have indeed $(A,B,C) = (W^0(p), W^0(q), W^0(r))$, in terms of the Brownian bridge W^0, for which we also write $W^0(t) = W(t) - tW(1)$, with W the Brownian motion. With V_1, V_2, V_3 the independent increments $W(p), W(q) - W(p), W(r) - W(q)$, show that (xx care here xx)

$$\begin{aligned} R &= (1/p)\{-(q/p)W(p) + W(q)\} = (1/p)\{-(q/p-1)V_1 + V_2\}, \\ S &= (1/q)\{-(r/q)W(q) + W(r)\} = (1/q)\{-(r/q-1)(V_1+V_2) + V_3\}, \\ T &= (1/r)[-(1/r)\{W(r) - rW(1)] = ... \end{aligned}$$

(c)**.** An extended setup with quantiles levels $p_1 < \cdots < p_r$ is as follows. We have $\sqrt{n}(Q_{n,j} - p_j) \to_d W^0(p_j)$, and build successive ratios $R_{n,j} = Q_{n,j+1}/Q_{n,j}$, including also $R_{n,k+1} = 1/Q_{n,k}$. Show that $\sqrt{n}(R_{n,j} - p_{j+1}/p_j) \to_d R_j$, with these being independent, and with variances $\tau_j^2 = (1/p_j)^2\{p_j(p_{j+1}/p_j - 1)^2 + p_{j+1} - p_j\}$. Show also that these may be represented as

$$\begin{aligned} R_j &= (1/p_j)\{-(p_{j+1}/p_j)W(p_j) + W(p_{j+1})\} \\ &= (1/p_j)[-\{p_{j+1}/p_j - 1\}W(p_j) + W(p_{j+1}) - W(p_j)]. \end{aligned}$$

Write furthermore $W(p_j) = D_1 + \cdots + D_j$, in terms of independent increments $D_j \sim \mathrm{N}(0, p_{j+1} - p_j)$, and explain that

$$R_j = (1/p_j)[-(p_{j+1}/p_j)(D_1 + \cdots + D_j) + D_{j+1}\}.$$

(xx a bit more, for Ch9 process exercise. xx) With $p_j = j/k$, there is a nice independent increment process lying there, with $R_{n,j} = F_n^{-1}((j+1)/k) - F_n^{-1}(j/k)$ tending to $(j+1)/j = 1 + 1/j$, and $\sqrt{n}(R_{n,j} - (1+1/j)) \to_d R_j$,

$$R_j = (k/j)\{-(1/j)(D_1 + \cdots + D_j) + D_{j+1}\},$$

where $D_1, D_2, \ldots$ are independent $\mathrm{N}(0, 1/k)$. their variances are

$$\tau_j^2 = (k/j)^2[1/(jk) + 1/k] = \frac{k^2}{j^2}\frac{j+1}{jk} = \frac{k}{j^2}(1+1/j).$$

(d)**.** (xx tentative; might land in Ch9. xx) For fixed ε positive, consider the process

$$Z_n(p,\varepsilon) = \sqrt{n}\{Q_n(p+\varepsilon)/Q_n(p) - 1 - \varepsilon/p\}.$$

It converges to a normal process $Z(t,\varepsilon)$, where increments are independent over disjoint intervals. The variance level is

$$\tau^2(p,\varepsilon) = (1/p^2)\{(\varepsilon/p)^2 + \varepsilon\}.$$

Ex. 3.45. *Estimating a multinormal variance matrix.* (xx if included, then more polish, and calibration with other exercises. could have UMVU questions for $|\Sigma|$ and $\rho = |\Sigma|^{1/(2p)}$ in Ch8. xx)

(a)**.** We start out considering independent pairs $(X_1,Y_1), \ldots, (X_n,Y_n)$ of independent standard normals, and form from these the random 2×2 matrix

$$W = \sum_{i=1}^n \begin{pmatrix} X_i \\ Y_i \end{pmatrix} \begin{pmatrix} X_i \\ Y_i \end{pmatrix}^{\mathrm{t}} = \begin{pmatrix} \sum_{i=1}^n X_i^2 & \sum_{i=1}^n X_iY_i \\ \sum_{i=1}^n X_iY_i & \sum_{i=1}^n Y_i^2 \end{pmatrix} = \begin{pmatrix} W_{11} & W_{12} \\ W_{12} & W_{22} \end{pmatrix}.$$

Writing $M_n = n^{-1}\sum_{i=1}^n X_i^2$ and $\widehat{\gamma} = n^{-1}\sum_{i=1}^n X_iY_i/M_n$, explain that its determinant may be written $D_n = W_{11}W_{22,1}$, where $W_{11} \sim \chi^2_n$, and

$$W_{22,1} = W_{22} - W_{12}^2/W_{11} = \sum_{i=1}^n Y_i^2 - nM_n\widehat{\gamma}^2 = \sum_{i=1}^n (Y_i - \widehat{\gamma}X_i)^2.$$

Then use results and methods from Ex. 3.35 to demonstrate that $W_{22,1}$ is in fact independent of W_{11}, with a χ^2_{n-1} distribution, regardless of the values of the X_i. So D_n is a product of two independent χ^2_n and χ^2_{n-1}.

(b). Generalise the above to the following result: if $X_1, \ldots, X_n$ are i.i.d. $\mathrm{N}_p(0, I_p)$, then the random matrix $W_n^* = \sum_{i=1}^n X_iX_i^{\mathrm{t}}$ has a determinant D_n with distribution equal to that of $K_1 \cdots K_p$, where $K_1, \ldots, K_p$ are independent chi-squares, with degrees of freedom $n, n-1, \ldots, n-p+1$.

(c). Consider now $Y_1, \ldots, Y_n$ i.i.d. $\mathrm{N}_p(0, \Sigma)$, for some positive definite $p \times p$ variance matrix Σ. Writing $Y_i = \Sigma^{1/2}X_i$, with the X_i being $\mathrm{N}_p(0, I_p)$, show that $W_n = \sum_{i=1}^n Y_iY_i^{\mathrm{t}}$ has the same distribution as $\Sigma^{1/2}\, W_n^*\, \Sigma_n$, where W_n^* is as above. From this again, show that $|W_n| \sim |\Sigma|\, K_1 \cdots K_p$.

(d). Suppose $Y_1, \ldots, Y_n$ are i.i.d. from some p-dimensional $\mathrm{N}_p(\xi, \Sigma)$. For estimating Σ, use $\widehat{\Sigma} = A_n/(n-1)$, with $A_n = \sum_{i=1}^n (Y_i - \bar{Y})(Y_i - \bar{Y})^{\mathrm{t}}$, with as usual $\bar{Y} = n^{-1}\sum_{i=1}^n Y_i$. Show that $A_n = \sum_{i=1}^n Y_iY_i^{\mathrm{t}} - n\bar{Y}\bar{Y}^{\mathrm{t}}$ can be transformed to the representation $\sum_{i=2}^n Z_iZ_i^{\mathrm{t}}$, where the Z_i are i.i.d. with the $\mathrm{N}_p(0, \Sigma)$ distribution. From these efforts, show that the determinant of A_n has the distribution $|\Sigma|\, K_1 \cdots K_p$, with these being independent chi-squares, with degrees of freedom $n-1, \ldots, n-p$. In addition to clarifying that $\widehat{\Sigma}$ is unbiased for Σ, find unbiased estimators for $|\Sigma|$ and for $\rho = |\Sigma|^{1/(2p)} = (\sigma_1 \cdots \sigma_p)^{1/p}$, in terms of the standard deviations σ_j for the different componentss.

Ex. 3.46. *Comparing estimators for a skewed distribution on the unit interval.* Consider i.i.d. observations $Y_1, \ldots, Y_n$ from the model with c.d.f. $F(y, \theta) = y^\theta$ on the unit interval, with θ an unknown positive parameter. Various estimators may be constructed and then compared in terms of precision.

(a). We start with the moment estimator. Show that the density is $f(y, \theta) = \theta y^{\theta-1}$, and that its mean is $\theta/(\theta+1)$. Show that the moment method estimator becomes $\widehat{\theta}_m = \bar{Y}/(1-\bar{Y})$, in terms of the sample average $\bar{Y}$.

(b). Explain that the q level quantile is $y_q = q^{1/\theta}$, and that quantile matching leads to the estimator $\widehat{\theta}_q = \log q/\log Q_{n,q}$, with $Q_{n,q}$ the empirical quantile; a special case is the median matching estimator $\widehat{\theta}_{0.50} = -\log 2/M_n$, in terms of the median.

(c). Show that $Z_i = -\log Y_i \sim \mathrm{Expo}(\theta)$. Argue that moment matching on this scale leads to the estimator $\widehat{\theta}_e = 1/\bar{Z}$, with $\bar{Z} = n^{-1}\sum_{i=1}^n Z_i$.

(d). To compare the three estimators we find their limit distributions. Use the CLT to identify the limit of $\sqrt{n}\{\bar{Y} - \theta/(\theta+1)\}$, and then the delta method to establish that

$$\sqrt{n}(\widehat{\theta}_m - \theta) \to_d \mathrm{N}(0, \tau_m^2), \quad \text{where } \tau_m^2 = (1+\theta)^4[\theta/(\theta+2) - \{\theta/(\theta+1)\}^2].$$

(e). For the quantiles, show first that $\sqrt{n}(Q_{n,q} - q^{1/\theta}) \to_d V$, where V has variance $q(1-q)/\theta^2 q^{2-2/\theta}$. Again via the delta method, establish that

$$\sqrt{n}(\widehat{\theta}_q - \theta) \to_d -\frac{\theta^2}{\log q}\frac{1}{q^{1/\theta}}V \sim \mathrm{N}(0, \tau_q^2), \quad \text{where } \tau_q^2 = \frac{\theta^2}{(\log q)^2}\frac{1-q}{q}.$$

Examine this function numerically and find that the best value for q, leading to the smallest variance, is $q^* = 0.203$, with variance $1.544\,\theta^2$.

(f). For the transformed data estimator, show $\sqrt{n}(\widehat{\theta}_e - \theta) \to_d \mathrm{N}(0, \theta^2)$. Now compare the precision of the three estimators, and comment.

Ex. 3.47. *Exponential regression.* Something like $\lambda_i = \lambda_0 \exp(\beta x_i)$. Perhaps we need to write down a log-likelihood, but we don't push it hard, since that comes later. The point is to convey 'yes, we can put up regression models beyond the linear one'.

Ex. 3.48. *Nonparametric estimation.* [xx Something around estimation in iid setup of means, quantiles, functions thereof via delta method. A few points regarding $\widehat{\sigma}^2$, which remains a perfectly good estimator of σ^2 also outside normality, but then with a more complex variance formula, etc. xx]

Ex. 3.49. *Extra: empirical covariance matrix.* (xx we'll see later where and how to fit this in. could also be inside likelihood chapter.) Suppose $Y_1, \ldots, Y_n$ are i.i.d. vectors in dimension p, with mean ξ and covariance matrix Σ.

(a). With $A = \sum_{i=1}^n (Y_i - \bar{Y})(Y_i - \bar{Y})^{\mathrm{t}}$, show that Show that

$$\sum_{i=1}^n (Y_i - \xi)(Y_i - \xi)^{\mathrm{t}} = A + n(\bar{Y} - \xi)(\bar{Y} - \xi)^{\mathrm{t}}.$$

Deduce that $\widehat{\Sigma}_u = A/(n-1)$, the so-called empirical covariance matrix, is unbiased for Σ.

(b). (xx to be pushed to likelihood chapter. xx) Under multinormality, show that the log-likelihood function can be expressed as

$$\ell_n(\xi, \Sigma) = -\tfrac{1}{2} n \log|\Sigma| - \tfrac{1}{2} \sum_{i=1}^n (Y_i - \xi)^{\mathrm{t}} \Sigma^{-1} (Y_i - \xi).$$

Show that this, for any positive definite Σ, is maximised for $\widehat{\xi} = \bar{Y}$, leading to the profiled

$$\ell_{\mathrm{prof}}(\Sigma) = -\tfrac{1}{2} n \{\log|\Sigma| + \mathrm{Tr}(\Sigma^{-1}\Sigma^*)\},$$

in terms of $\Sigma^* = A/n$, i.e. with divisor n not $n-1$ as above. Show in general that if B is some given symmetric positive definite matrix, then the function $h(\Sigma) = \log|\Sigma| + \mathrm{Tr}(\Sigma^{-1}B)$ is minimised for $\widehat{\Sigma} = B$. Deduce that Σ^* is the ML estimator for Σ. (xx this might perhaps most easily be shown using the spectral decomposition $P\Sigma P^{\mathrm{t}} = D$, a diagonal. xx)

Ex. 3.50. *Squares and products of normals.* (xx leftover. this might not earn its place, depends on how dennis nils fieller difference work pans out. but it is a good illustration of good clean probability calculus, for a purpose.) For independent $\widehat{a} \sim \mathrm{N}(a, 1)$ and $\widehat{b} \sim \mathrm{N}(b, 1)$, consider

$$U = 2\,\widehat{a}\widehat{b} \quad \text{and} \quad V = \widehat{a}^2 + \widehat{b}^2.$$

Aspects of their joint distribution come up in the context of Fieller parameters, of the type a/b, etc.

(a). Show that $|U| \le V$. For

$$T_1 = (\widehat{a} + \widehat{b})/\sqrt{2} \sim \mathrm{N}((a+b)/\sqrt{2}, 1), \qquad T_2 = (\widehat{a} - \widehat{b})/\sqrt{2} \sim \mathrm{N}((a+b)/\sqrt{2}, 1),$$

show that these two are independent, and that

$$Z_1 = T_1^2 = \tfrac{1}{2}V + \tfrac{1}{2}U, \quad Z_2 = T_2^2 = \tfrac{1}{2}V - \tfrac{1}{2}U.$$

From $V = Z_1 + Z_2$ and $U = Z_1 - Z_2$, show that (U, V) has joint density

$$h(u,v) = g_1(\tfrac{1}{2}u + \tfrac{1}{2}v, \tfrac{1}{2}(a+b)^2)\, g_1(-\tfrac{1}{2}u + \tfrac{1}{2}v, \tfrac{1}{2}(a-b^2),$$

in terms of the density function $g_1(\cdot, \lambda)$ for the noncentral $\chi^2_1(\lambda)$ (xx mind the determinant; is it a factor 1/8? xx).

(b). Show that the conditional density for U given $V = v$ may be written

$$f_0(u \,|\, v) = \frac{g_1(\frac{1}{2}u + \frac{1}{2}v, \frac{1}{2}(a+b)^2)\, g_1(-\frac{1}{2}u + \frac{1}{2}v, \frac{1}{2}(a-b)^2)}{g_2(v, a^2+b^2)} \quad \text{for } u \in (-v, v).$$

The ratio $R = U/V$ has density $f(r \,|\, v) = v f_0(vr \,|\, v)$ for $r \in (-1, 1)$.

(c). (xx can we complete this? xx) For the difference between two Fieller parameters, i think we need something like $U - U'$ and $V + V'$. perhaps

$$\begin{aligned} U^* = U - U' &= Z_1 - Z_2 - (Z_3 - Z_4) = Z_1 - Z_2 - Z_3 + Z_4, \\ V^* = V + V' &= Z_1 + Z_2 + Z_3 + Z_4. \end{aligned}$$

which leads to

$$\tfrac{1}{2}(V^* + U^*) = Z_1 + Z_4, \quad \tfrac{1}{2}(V^* - U^*) = Z_2 + Z_3.$$

So we can write down the joint density for (U^*, V^*):

$$\begin{aligned} h^*(u^*, v^*) = \gamma_2(&\tfrac{1}{2}u^* + \tfrac{1}{2}v^*, \tfrac{1}{2}(a_1+b_1)^2 + \tfrac{1}{2}(a_2-b_2)^2) \\ &\gamma_2(-\tfrac{1}{2}u^* + \tfrac{1}{2}v^*, \tfrac{1}{2}(a_1-b_1)^2 + \tfrac{1}{2}(a_2+b_2)^2), \end{aligned}$$

modulo a determinant.

Ex. 3.51. *Proving LLN with only finite mean.* (xx if we like it, we include it inside the LLN subsection of Ch2. arguments below are correlated with those of Ex. 2.59, but not entirely the same, and here we don't need the Kolmogorov equality. xx) Let $X_1, X_2, \ldots$ be i.i.d. with finite $\mathrm{E}\,|X_i|$. Via a characterisation of the finite mean property we are able to demonstrate below that $\bar{X}_n = n^{-1}\sum_{i=1}^n X_i$ tends a.s. to $\mathrm{E}\,X_i$, i.e. the strong LLN, without assuming finite variance.

(a). We start with two lemmas of independent interest. First, a mathematical one: if $a_1, a_2, \ldots$ are numbers such that $\sum_{n=1}^\infty a_n/n$ is finite, show that $\bar{a}_n = n^{-1}\sum_{i=1}^n a_i \to 0$. Second, consider independent U_i, with means ξ_i and variances σ_i^2. Show that if $\sum_{n=1}^\infty \sigma_n^2/n^2$ is finite, then $\bar{U}_n - \bar{\xi}_n = n^{-1}\sum_{i=1}^n (U_i - \xi_i) \to_{a.s.} 0$.

(b). Write F for the c.d.f. of the X_i. Letting $C_i = [i-1, i)$ and then $a_i = \int_{C_i} |x|\,\mathrm{d}F(x)$, show that $\sum_{i=1}^\infty a_i = \mathrm{E}\,|X_1|$. Show next that $\int_{C_i} x^2\,\mathrm{d}F(x) \le i a_i$, and that $\sum_{n \ge i} 1/n^2 \le 1/(i-1)$ for $i = 2, 3, \ldots$. The latter may be worked with via $\int_i^\infty 1/u^2\,\mathrm{d}u = 1/i$.

(c). With these ingredients, show for the quantity $\lambda = \sum_{n=1}^{\infty} n^{-2} \int_{(-n,n]} x^2 \, \mathrm{d}F(x)$ that

$$\lambda \le \sum_{n=1}^{\infty} \sum_{i \le n} \frac{i a_i}{n^2} = \sum_{i=1}^{\infty} i a_i \sum_{n \ge i} \frac{1}{n^2} \le \pi^2/6 + \sum_{i=2}^{\infty} i a_i/(i-1).$$

The point is that finiteness of the mean implies finiteness of λ; show also the converse of that statement.

(d). Returning to the X_i, write these as $Y_i + Z_i$, with $Y_i = X_i I(X_i \in (-i, i])$ and $Z_i = X_i I(X_i \notin (-i, i])$, i.e. essentially sorting $|X_i|$ into smaller than or bigger than i. Taking averages leads to $\bar{X}_n = \bar{Y}_n + \bar{Z}_n$. Use Borel–Cantelli to show that there are only finitely many non-zero Z_i. Explain that the mean of $\bar{Y}_n$ tends to $\mathrm{E}\, X_1$, and not only with the variance tending to zero, but with $\sum_{n=1}^{\infty} (\mathrm{Var}\, Y_n)/n^2$ finite. Conclude that $\bar{X}_n \to_{a.s.} \xi = \mathrm{E}\, X_i$.

Notes and pointers

(xx to come. we point to various matters, genesis of crucial concepts, and also point to chapters ahead. explain that yes, we've touched and used CLT and delta method and a bit more here, but with details and more material to come in Ch 4. xx)

[xx CLT for binomials: associated with the famous names de Moivre (who showed a version of this in 1733) and Laplace (who had a clearer and more general proof in 1812). xx]

Briefly genesis of неравенцво Маркова, the Markov, and the неравенцво Чебышёва, the Chebyshov (often anglisised to Chebyshev, but his name was really Чебышёв). mention Kahneman et al. (2020).

I.4

Testing, sufficiency, power

In previous chapters we have learned about classes of distributions, their parameters, ways of estimating these from data, along with assessment of precision and confidence intervals. The present chapter goes on to the fundamental statistical reporting tool of hypothesis testing. Statistical testing of hypotheses are data-based rules for when to reject (and hence, when not to reject) a hypothesis about the parameters of a model. Theory is developed to construct such tests in quite general setups. These are constructed to have a certain significance level, like 0.05, the intended low probability of rejecting the hypothesis if it is in fact true. It also has a power function, the probability of rejection as a function of how far the model parameters might be from the hypothesis. We learn the basics of the Neyman–Pearson theory for optimal testing, and see it panned out for many situations, including the general setting of exponential families. The fruitful concept of sufficiency, the notion that a lower-dimensional vector of summaries contains all statistical information about a model, is also developed in this chapter, with implications for both estimation and testing. Parts of the theory of this chapter are generalised and extended in Chapters 5, 7, 8.

Key words: ancillarity, completeness, conditional tests, exponential family, factorisation theorem, Neyman–Pearson optimality, power, p-value, sufficiency, testing, Wald ratios

In the broad context of analysing data from models, the previous chapter dealt generally speaking with estimators of relevant parameters, their precision, and comparisons, leading in particular to confidence intervals. The present chapter develops the methods and applications further, in the direction of *statistical testing of null hypotheses*. Studying such tests involves their interpretation, construction, properties, performance, along with connections to confidence and yet other themes.

Consider in general terms data y stemming from a model with a parameter vector $\theta = (\theta_1, \ldots, \theta_p)$, and suppose one wishes to test the null hypothesis H_0 that θ is inside a well-defined subset Θ_0 of the full parameter region Θ. Precisely what constitutes a null hypothesis is a matter of scientific and statistical context, often reflecting the intentions and the overall aims of the data collection and its analysis. The null hypothesis is typically a statement, concerning the nature of the mechanisms studied, the incorrect rejection of which one attempts to avoid. A simple illustration is testing whether a particular regression coefficient is equal to zero (so H_0 could be '$\beta_3 = 0$'), or testing whether two parameters, perhaps for two groups, are equal (so H_0 would be '$\theta_A = \theta_B$').

A *test* for such a hypothesis is a rule saying 'H_0 is to be rejected if data y fall in the set R', along with the complementary statement 'H_0 is not rejected if data y fall in the set R^c'. We talk here of *the rejection region* R and *the acceptance region* R^c. A fundamental aspect of such a test to look for is its *significance level*, typically meant to be a relatively small probability, like 0.05 or 0.01 or even smaller. We say that the test has level α provided

the level of a test

$$\Pr_\theta(\text{reject } H_0) \le \alpha \quad \text{for all } \theta \in \Theta_0. \tag{4.1}$$

With level 0.01 one is guaranteed such a low chance of falsely claiming that H_0 is wrong, if it is indeed correct. So for the illustrations briefly alluded to above, when statisticians after careful analyses reject '$\beta_3 = 0$' or '$\theta_A = \theta_B$', therefore going on to claim that their alternatives are valid, with '$\beta_3 \neq 0$' and '$\theta_A \neq \theta_B$', these claims are seen as trustworthy (and may make into publications), since the probability of these claims being false is so low.

Often tests are carried out via appropriate *test statistics*, as when constructing a $T = T(y)$, a function of the data y, with the property that T ought to be inside some normal and well-understood range under H_0 conditions (we shall meet many cases of determining the *null distribution* of such test statistics), but bigger, if H_0 is wrong. In such cases, the rejection region takes the form $R = \{y\colon T(y) \ge t_0\}$, with t_0 the rejection threshold, chosen to have $\Pr_\theta(T(Y) \ge t_0) \le \alpha$ for all θ congruent with H_0.

test statistic, rejection threshold

Conceptually and operationally, there is a certain direction implied when setting up an hypothesis H_0 and its alternative. Rejecting H_0, with a test with low level testing level α, leads to a positive claim about the alternative being true; the observed data have landed in a region which if H_0 were true has low probability. On the other hand, *not rejecting* the null hypothesis should not be seen as 'verifying' H_0; one needs to be content with the not so bold statement 'the observed data do not provide sufficient evidence for claiming that H_0 does not hold'.

We are also keenly interested in *the power* of a test, which is the detection chance $\Pr_\theta(\text{reject } H_0)$ as a function of θ, in the alternative parameter domain $\Theta - \Theta_0$. Thus some tests are *stronger* than other tests with the same level, and we learn recipes for constructing such in the exercises below. The power of a test clearly depends on the quality of data, and typically of the sample size, as we shall see in exercises. Thus detecing that $\theta_A \neq \theta_B$, for parameters of two groups, in a setup where there really is a difference, becomes more likely with more data.

Below we also define, discuss, and use *p-values*, which are commonly quoted in most branches of applied statistics work, typically to indicate how clear a potential finding is. The idea is to quantify how unlikely it is, to observe what is actually observed, if some relevant null hypothesis H_0 is actually true. If the test is set up to reject the null if an appropriate test statistic T is sufficiently large, we're after $\Pr_\theta(T \ge t_{\rm obs})$ under H_0 conditions, with $t_{\rm obs}$ the observed T for the given dataset. Some care is needed since that probability might depend on parameters under H_0. The more careful version is

the p-value

$$p = \max\{\Pr_\theta(T \ge t_{\rm obs})\colon \theta \in \Theta_0\}. \tag{4.2}$$

A small p-value, like $p \le 0.01$, casts serious doubt on H_0, since the observed $t_{\rm obs}$ is so unlikely. A rephrasing of the testing scheme, with significance level say 0.01, is to reject H_0 if $p \le 0.01$.

A classic result for testing theory is the Neyman–Pearson Lemma, which in an idealised setup with just two possible densities identifies the most powerful method for testing one density against the other; see Ex. 4.7–4.9. This sharpens further questions for more general setups, and we manage to find optimal tests in a variety of setups, including for the broad exponential family class. This also necessitates exploring and developing certain themes of independent interest and use, those of sufficiency, ancillarity, completeness, conditional testing; see the string of exercises starting with Ex. **??**. A core idea in that terrain is that of data compression; for various models, a low-dimensional vector of data summaries contains all relevant statistical information.

(xx then one paragraph with pointers to other chapters and perhaps to a few of the stories. xx)

Testing, testing

Ex. 4.1. *Testing a null hypothesis.* Here we introduce the notion of null hypotheses and their testing in a few simple setups.

(a)**.** The probability $p = \Pr(A)$ of a certain event is meant to be $p_0 = 0.33$, if the machinery around it works as it should. To test this one carries out the relevant experiment $n = 100$ times, and the event takes place $y = 44$ times. Should you reject the 0.33 hypothesis? Show that with $Y \sim \text{binom}(n, p)$, the statistic $T = (Y - np_0)^2/\{np_0(1-p_0)\}$ is approximately a χ^2_1, under the null assumption that $p = p_0$. Show also that W tends to be bigger than a χ^2_1, if $p \neq p_0$. Show that when using T as test statistic, the p-value of (4.2) becomes $p = \Pr_{p_0}(T \geq t_{\text{obs}})$, with t_{obs} the value of T seen with $y_{\text{obs}} = 44$. Compute this p-value, and decide whether the $p = 0.33$ hypothesis should be rejected at the 0.05 level.

(b)**.** Suppose (X, Y, Z) is trinomial with sum n and probabilities (p, q, r); see Ex. 1.5. For concreteness, suppose a theory holds that $(p, q, r) = (0.20, 0.30, 0.50)$, and that $(X, Y, Z) = (42, 47, 111)$ is observed. Is this enough to reject the theory in question? Show that

$$W = (X - np)^2/(np) + (Y - nq)^2/(nq) + (Z - nr)^2/(nr)$$

has mean equal to 2, and that it is approximately a χ^2_2, using the multidimensional CLT for (X, Y, Z). Explain how this may be used to test the theory mentioned, and carry out the test. This is actually the classic Pearson chi-squared test for multinomials; see Story vii.1 for details, generalisations, and discussion.

(c)**.** You're rolling your die, but it takes you as many as $y_{\text{obs}} = 15$ rolls to get your first '6'. Does this make you suspect that the probability p of a '6' is not 1/6, but lower? With testing level $\alpha = 0.05$, what is the set of suspicious outcomes of Y, the number of throws to get the first '6'?

Ex. 4.2. *Connections from confidence intervals to testing.* Though confidence intervals and testing are two different reporting tools, when summarising inference, there are clear connections. Suppose μ is a parameter of inference, perhaps a function of model parameters, for which we can build both confidence intervals and tests.

(a)**.** Suppose one needs to test the one-point null hypothesis that $\mu = \mu_0$, a given value, and that $[L, U]$ is a 99 percent confidence interval. Show that the test consisting in rejecting, if μ_0 is outside this interval, has level 0.01.

(b). Suppose on the other hand that there is a well-defined 0.01 level test procedure for testing $\mu = \mu_0$, against $\mu \neq \mu_0$, for each candidate value μ_0. Gather together in a set A all the μ_0 values which are not rejected by the corresponding 0.01 level test. Show that $\Pr_\mu(\mu \in A) = 0.99$, making A a 99 percent confidence region.

(c). Go through the relevant details, from confidence interval to test and vice versa, for the simple prototype case of the observation being $Y \sim \mathrm{N}(\theta, 1)$.

Ex. 4.3. *Confidence intervals for quantiles.* Let $Y_1, \ldots, Y_n$ be i.i.d. from a continuous density, positive on its sample space. How can we construct confidence intervals for the median $\mu = F^{-1}(\frac{1}{2})$, and more generally for quantiles $\mu_q = F^{-1}(q)$? There are several approaches here, but here we give the basic method via the empirical c.d.f.

(a). Let F_n be the empirical c.d.f. for the data, see Ex. 3.9. If μ_0 is the true median, show that $F_n(\mu_0)$ is a simple B_n/n, with $B_n \sim \mathrm{binom}(n, \frac{1}{2})$, and that this implies $W_n(\mu_0) = \sqrt{n}\{F_n(\mu_0) - \frac{1}{2}\}/\frac{1}{2}$ being approximately a standard normal. Argue that a natural 0.05 level test for $\mu = \mu_0$ is to accept the hypothesis provided $|W_n(\mu_0)| \leq 1.96$.

(b). Following the general testing-to-confidence connection of Ex. 4.2, show that the associated 95 percent confidence interval becomes $\mathrm{ci}_n = \{\mu\colon |F_n(\mu) - \frac{1}{2}| \leq 1.96/(2\sqrt{n})\}$. In other words, we may read off the interval from a plot of F_n, without knowing or taking on board the details of the exact or approximate distribution of sample quantiles, as with Ex. 3.21.

(c). Generalise to the case of any quantile $\mu_q = F^{-1}(q)$. Show that the recipe above leads to the confidence interval $\mathrm{ci}_n = \{\mu_q\colon |F_n(\mu_q) - q| \leq z_\alpha\{q(1-q)\}^{1/2}/\sqrt{n}\}$, where $\Pr(|\mathrm{N}(0,1)| \leq z_\alpha) = \alpha$, the confidence level. For an illustration of these methods, check Story **??**, where we plot the empirical c.d.f. and read off confidence intervals for quantiles $F^{-1}(q)$ at levels 0.10, 0.50, 0.90, for the weight of mothers pre pregnancy.

(d). Discuss ways in which more accurate confidence intervals can be constructed, using the exact binomial distribution; for the median, for example, $nF_n(\mu) \sim \mathrm{binom}(n, \frac{1}{2})$. This leads to slight modification of the horizontal bands when reading off intervals from the empirical c.d.f.

Ex. 4.4. *t testing, one and two samples.* Testing the mean based on a sample of normal observations is a recurring problem, in several guises, and with the famous t test being the canonical procedure; details are given below. We also go through the basics for testing the difference of means for two normals samples. Due to the connections to confidence intervals discussed in Ex. 4.2 we also find accurate confidence intervals for the key parameters. Beyond their concrete relevance and repeated use in these standard setups, the t testing procedures are important since similar constructions can be worked with in large classes of more complicated setups, but then typically with approximations to key distributions, rather than the exact solutions found under these classic strict modelling assumptions.

(a). Suppose $X_1, \ldots, X_n$ are i.i.d. $\mathrm{N}(\xi, \sigma^2)$, so far assuming σ to be known, and that one wishes to test $H_0\colon \xi = \xi_0$ against the alternative that $\xi \neq \xi_0$, where ξ_0 is some appropriate given value, like zero. Using the exact distribution of $\bar{X}$, cf. Ex. 1.2, show that $Z = (\bar{X} - \xi_0)/(\sigma/\sqrt{n}) = \sqrt{n}(\bar{X} - \xi_0)/\sigma$ is standard normal under H_0, and that $|Z|$ will tend to be bigger than a normal if H_0 is not true. Explain that the test which rejects H_0 when $|Z| > 1.96$ has level 0.05.

(b). For the more realistic case of σ not being known, the natural construction is the t statistic $t = \sqrt{n}(\bar{X} - \xi_0)/\widehat{\sigma}$, with $\widehat{\sigma}$ the empirical standard deviation. Show from Ex. 1.50 that $t \sim t_{n-1}$ under H_0, and write down a precise 0.05 level test.

(c). Suppose now that the mean ξ for the population of X_i is to be compared with the mean η for another population, where we have i.i.d. data $Y_1, \ldots, Y_m \sim \mathrm{N}(\eta, \sigma^2)$. So the task is to test H_0 that $\xi = \eta$. Show first that $\bar{Y} - \bar{X} \sim \mathrm{N}(\eta - \xi, \sigma^2(1/m + 1/n)$. To build a t statistic we need an estimator for the variance. Writing $\widehat{\sigma}_1$ and $\widehat{\sigma}_2$ for the empirical deviances for samples 1 and 2, show that with

$$\widehat{\sigma}^2 = c_1\widehat{\sigma}_1^2 + c_2\widehat{\sigma}_2^2, \quad \text{using } c_1 = (n-1)/(n+m-2),\ c_2 = (m-1)/(n+m-2),$$

we have $\widehat{\sigma}^2 \sim \sigma^2\chi^2_{n+m-2}/(n+m-2)$, independent of $\bar{X} - \bar{Y}$. Conclude that $t = (\bar{X} - \bar{Y})/R \sim t_{n+m-2}$ under H_0, where $R = \widehat{\sigma}(1/m + 1/n)^{1/2}$.

(d). Use these building blocks to also construct a 90 percent confidence interval for $d = \xi - \eta$.

Ex. 4.5. *Wald tests.* Here we present the basics of so-called Wald tests, a general and practical way of forming tests via approximate normality. Such tests are used very routinely when looking for and reporting findings about regression coeffiecients in all standard regression models; see Story **??** for an illustration of this. The Wald tests can be constructed almost immediately, from the confidence interval construction of Ex. 3.11, via the interval-to-test connection of Ex. 4.2, but we tend to a few details to allow for potential simplifications of assumptions. – Assume $Y_1, \ldots, Y_n$ comprise the data (not necessarily i.i.d.), from a suitable model with vector parameter θ. Suppose further that $\mu = \mu(\theta)$ is a focus parameter, for which we need to test $\mu = \mu_0$, for a given null value μ_0.

(a). Suppose there is an estimator $\widehat{\mu}$ with the property that $\sqrt{n}(\widehat{\mu} - \mu) \to_d \mathrm{N}(0, \tau^2)$, and that there is a consistent estimator $\widehat{\tau}$ for this limit spread τ. Show as with Ex. 3.11 that $W_n = \sqrt{n}(\widehat{\mu} - \mu)/\widehat{\tau} \to_d \mathrm{N}(0,1)$. Show that the arguments go through, with a limiting standard normal, for $W_{n,0} = \sqrt{n}(\widehat{\mu} - \mu_0)/\widehat{\tau}_0$, at the null hypothesis, as long as $\widehat{\tau}_0$ is consistent for τ at this null position; also, technically speaking, we only need to establish $\sqrt{n}(\widehat{\mu} - \mu_0) \to_d \mathrm{N}(0, \tau^2)$ at the null hypothesis. Conclude that the test which rejects $\mu = \mu_0$ when $|W_{n,0}| \geq 1.96$ has level approaching 0.05 (and of course similarly with other chosen testing levels; level 0.01 corresponds to threshold 2.576, etc.).

(b). Explain that $p = \Pr(|N(0,1)| \geq W_{n,0,\mathrm{obs}})$ is an approximation to the exact p-value.

(c). Suppose $X \sim \mathrm{binom}(100, p)$, with a need for testing $p = 0.33$. Set up a Wald test, and compute the p-value, if $x_{\mathrm{obs}} = 44$. Suppose then that there is an additional $Y \sim \mathrm{binom}(100, q)$, and that one wishes to test $p = q$. Set up a Wald test, and compute the p-value, if (X, Y) are observed to be $(44, 55)$.

Ex. 4.6. *When confidence intervals for two parameters overlap.* Suppose confidence intervals for parameters a and b overlap. It might sound plausible that the null hypothesis $a = b$ will then not be rejected. To check aspects of this and related questions, consider a prototype setup where $\widehat{a} \sim \mathrm{N}(a, 1)$ and $\widehat{b} \sim \mathrm{N}(b, 1)$ are independent.

(a). With confidence level $1 - \alpha$, show that the canonical intervals are $\widehat{a} \pm z_0$ and $\widehat{b} \pm z_0$, with $z_0 = \Phi^{-1}(1 - \alpha/2)$, e.g. the standard 1.96 for level 0.95. Show that overlapping intervals corresponds to $|\widehat{d}| \leq 2z_0$, where $\widehat{d} = \widehat{b} - \widehat{a} \sim \mathrm{N}(b - a, 2)$

(b). Consider the test of $a = b$ which rejects when the intervals for a and b are disjoint. Show that this test has considerably lower significance level than α; for 95 percent intervals this level is 0.0056.

(c). Yes, it is possible to have overlapping 95 percent intervals and yet reject $a = b$ at level 5 percent. Identify the range of $\widehat{d}$ where this happens.

Neyman–Pearson optimal testing

Ex. 4.7. *The Neyman–Pearson Lemma: the basics.* Suppose data y come from a density f, where there are just two possibilities: either $f = f_0$, which is the null hypothesis to be tested, or $f = f_1$, the alternative. Here there's an optimal strategy, made clear by the so-called Neyman–Pearson Lemma, part of the 49-page landmark paper Neyman and Pearson (1933). For simplicity of presentation we consider the continuous case, where f_0 and f_1 are densities on the relevant sample space $\mathcal{Y}$ (which can be multidimensional). – A *test function* is a $T\colon \mathcal{Y} \to [0,1]$, with $T(y)$ the probability of rejecting f_0 if the the data take on value y. This setup even allows the possibility of an element of randomisation, as in 'if y turns out be 3.33 I throw some coins and I reject H_0 with probability 0.77'. Once in a blue while this might be of relevance, with discrete data, but in practice such a test function $T(y)$ takes on only values 1, for a rejection set R, and 0, for the complementary acceptance set R^c.

(a). Show that the probability of rejecting the null, if the null is true, can be written $\Pr_{f_0}(\text{reject}) = \int f_0 T \,\mathrm{d}y$.

(b). For a given testing level α, like 0.01, let T^* be the test which rejects when $f_1(y)/f_0(y) \geq c$, with c tuned such that

$$\Pr_{f_0}(T^* \text{ rejects}) = \int_{y\colon\, f_1(y)/f_0(y)\geq c} f_0(y)\,\mathrm{d}y = \alpha.$$

Let T be any other test function with the same level α. Show that the power difference at f_1 can be written

$$\pi_{T^*}(f_1) - \pi_T(f_1) = \int f_1(T^* - T)\,\mathrm{d}y = \int (f_1 - cf_0)(T^* - T)\,\mathrm{d}y.$$

(c). Show that among all possible tests, with level α, the T^* has the strongest detection power at position f_1. the Neyman–Pearson Lemma

(d). (xx a bit more, to cover discrete case; and we may also allow competitors with $\int f_0 T < \alpha$. xx)

Ex. 4.8. *The Neyman–Pearson Lemma: more details.* Here we tend to some further details, related to the Neyman–Pearson Lemma and its proof, in Ex. 4.7.

(a). For two positive densities f_0 and f_1 defined on the same sample space, as with the Neyman–Pearson Lemma, consider the event $A_c = \{f_1(Y)/f_0(Y) \geq c\}$. Show that the function $p(c) = \Pr_{f_0}(A_c) = \int_{f_1(y)\geq cf_0(y)} f_0(y)\,\mathrm{d}y$ is a continuous and monotone function, starting at 1 and ending at 0, when c travels through $[0,\infty)$. Hence deduce that there for given α really is a unique c in the Neyman–Pearson recipe.

(b). Illustrate the $p(c) = \Pr_{f_0}(A_c)$ in a few concrete situations, including (i) $f_0 \sim \mathrm{N}(0,1)$ and $f_1 \sim \mathrm{N}(1,1)$; (ii) $f_0 \sim \mathrm{Gam}(2.2, 3.3)$ and $f_1 \sim \mathrm{Gam}(3.3, 2.2)$. (iii) $f_0 \sim \mathrm{Gam}(3.3, 2.2)$ and $f_1 \sim \mathrm{Gam}(2.2, 3.3)$. For these situations, find the rejection threshold c such that the test rejecting when $f_1(y)/f_0(y) \geq c$ has significance level 0.05.

(c). For the three cases above, determine also the power at f_1. Note that the power at f_1 when testing f_0 is typically different from the power at f_0 when testing f_1.

Ex. 4.9. *The Neyman–Pearson Lemma: applications.* For the simple two-possibilities setup we learn from the Neyman–Pearson lemma that there is a clear recipe for setting up the optimal test for $f = f_0$ against $f = f_1$. Here are some examples.

(a). Suppose $Y \sim \mathrm{N}(\theta, 1)$. Show that the optimal test of level $\alpha = 0.01$, for testing $\theta = 0$ vs. $\theta = 1.234$, is to reject if $Y \geq z_{0.99} = 2.326$, the upper 0.01 point of the standard normal.

(b). In this situation, verify that one finds the very same optimal 0.01 level test, for any alternative point $\theta_1 > 0$. Hence the $Y \geq z_{0.99}$ test is *uniformly most powerful*, against all positive alternative.

UMP, uniformly most powerful

(c). Generalise this to the case of data $Y_1, \ldots, Y_n$ being i.i.d. from the normal $\mathrm{N}(\theta, \sigma^2)$, with known σ. Show that the test which rejects $\theta = 0$ against $\theta > 0$ when $Z_n = \sqrt{n}\bar{Y}/\sigma > z_{1-\alpha}$ is uniformly most powerful, among all tets with level α; here $z_{1-\alpha} = \Phi^{-1}(1-\alpha)$. Find its power function, and draw it in a diagram, for $\theta_0 = 1.234$, $\sigma = 1$, and for $n = 10, 20, 30$.

(d). Let $Y_1, \ldots, Y_n$ be i.i.d. from the $\mathrm{N}(\theta_0, \sigma^2)$ distribution, with θ_0 known. Consider the problem of testing $\sigma = \sigma_0$ against $\sigma > \sigma_0$, at level say 0.01, where σ_0 is a prescribed null value. Show that the test which rejects when $V_n = \sum_{i=1}^n (Y_i - \theta_0)^2/\sigma_0^2 \geq \gamma_{n,0.99}$, the 0.99 quantile of the χ^2_n, is uniformly optimal. Find its power function.

(e). Consider f_0, the standard normal, and $f_1(y) = \frac{1}{2}\sqrt{2}\exp(-\sqrt{2}|y|)$; they have both zero mean and unit variance. Find the optimal test for f_0 against f_1, with level 0.05, and find its detection power at f_1. Then do the opposite, constructing the best test at level 0.05 for f_1 against f_0, and find the power at f_0.

(f). Suppose $Y_1, \ldots, Y_n$ are i.i.d. from the exponential $\theta\exp(-\theta y)$. Find the strongest 0.10 level test for $\theta = \theta_0$ against an alternative $\theta_1 > \theta_0$. Your test will not depend on the θ_1, as long as it is to the right of θ_0; hence this test is uniformly most powerful against these alternatives.

Ex. 4.10. *Setting up the test via null simulations.* Consider the simple two-possibilities Neyman–Pearson setup, with data Y from a density f, where one tests $f = f_0$ vs. $f = f_1$, for given densities f_0 and f_1.

(a). Writing y_{obs} for the observed data, the strongest test consists in rejecting when $T(y_{\mathrm{obs}}) = f_1(y_{\mathrm{obs}})/f_0(y_{\mathrm{obs}}) \geq c$, with c finetuned to have level say $\alpha = 0.05$. Explain that to actually find the c, in a perhaps complicated situation, one may simulate a high number of $T(y^*) = f_1(y^*)/f_0(y^*)$, with y^* drawn from f_0, and then set c equal to the 0.95 quantile in this distribution.

(b). Explain that the Neyman–Pearson test may be defined and constructed using any monotone smooth transformation of the density ratio. In the simulation recipe above one may e.g. use simulations of $V(y) = \log f_1(y) - \log f_0(y)$, which might be numerically easier and safer, and then reject if the observed $V(y_{\rm obs})$ is larger than then 0.95 quantile point among the simulated $V(y^*)$.

(c). As an illustration, suppose the data are $Y_1, \ldots, Y_n$, with $n = 12$, and that one needs to test that their density f is standard normal, vs. the alternative that f is a t_ν, with $\nu = 10.0$ degrees of freedom. Show that the strongest test, with level 0.05, is to reject when $V = \sum_{i=1}^n \{\log f_\nu(Y_i) - \log \phi(Y_i)\} \geq c$, with $c = 0.636$.

(d). Explain that this null distribution simulation scheme also works in each case where one tests a specified $H_0\colon f = f_0$, also when the alternative hypothesis might be a bigger space than merely a given f_1, as long as the test statistic has been decided upon.

Ex. 4.11. *Density ratios and optimal testing: the normal and the Cauchy.* The Neyman–Pearson recipe is to reject when the density ratio $f_1(y)/f_0(y)$ is sufficiently big. This pans out differently in different situations, as illustrated here.

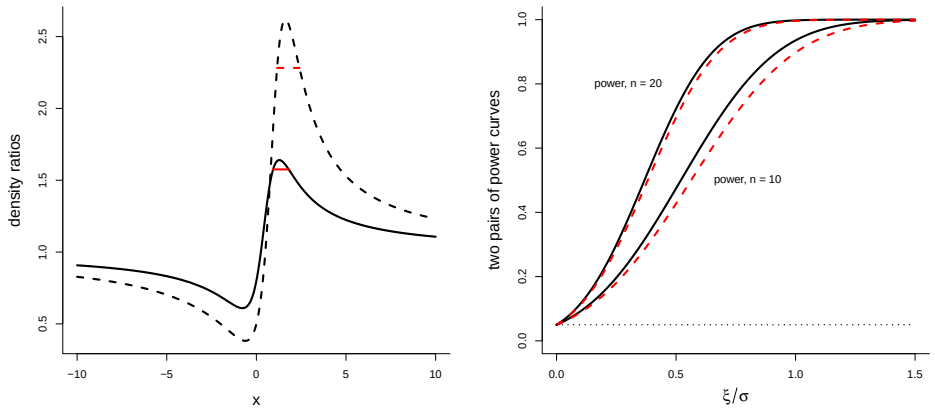


Figure 4.1: *Left panel: for the Cauchy density model, see Ex. 4.11, ratios $f(y,\theta_1)/f(y,0)$ (full line) and $f(y,\theta_2)/f(y,0)$ (slanted) are shown, for alternative values $\theta_1 = 0.50$ and $\theta_2 = 1.00$ to the null. Also indicated are the rejection intervals $[a_1, b_1]$ and $[a_2, b_2]$, for the optimal tests against θ_1 and θ_2. Right panel: two pairs of power. Testing $\xi = 0$ against $\xi > 0$, see Ex. 4.12, for a normal sample of size 10 (lower pair) and 20 (upper pair), at test level 0.05, as a function of ξ/σ. The lower power is for the t test (slanted, red); the upper power if for the normal based test (full, black), which requires that σ is known.*

(a). For a single observation Y, consider testing $f_0 = \mathrm{N}(0,1)$ against $f_1 = \mathrm{N}(\theta_1, 1)$, with θ_1 positive. Show that

$$\frac{f_1(y)}{f_0(y)} = \frac{\exp\{-\frac{1}{2}(y-\theta_1)^2\}}{\exp(-\frac{1}{2}y^2)} = \exp(\theta_1 y - \tfrac{1}{2}\theta_1^2).$$

Verify that this is a monotone function in y, regardless of the value of $\theta_1 > 0$. Argue that 'reject f_0 provided Y is big enough' becomes the uniformly optimal test. Exhibit the rejection threshold in $Y \geq c$, if the significance level is to be 0.10.

(b). The situation is rather different for the case of the Cauchy density $f(y,\theta) = (1/\pi)\{1+(y-\theta)^2\}^{-1}$. Suppose we wish to test $\theta = 0$ versus a positive θ_1. Show that

$$\frac{f_1(y)}{f_0(y)} = \frac{f(y,\theta_1)}{f(y,0)} = \frac{1+y^2}{1+(y-\theta_1)^2},$$

and draw this function in a diagram, for a few values of θ_1.

(c). For a concrete illustration, work through the alternative cases $\theta_1 = 0.50$ and $\theta_2 = 1.00$, for each case finding the rejection interval, say $[a_1, b_1]$ for the first and $[a_2, b_2]$ for the second, to give the optimal test, of level 0.10. (xx answers: $[0.933, 1.804]$ for θ_1; $[1.161, 2.388]$ for θ_2. construct a version of Figure 4.1, left panel. xx) – The point is that these rejection regions are different; the optimal test depends on the specific alternative, and there can be no uniformly optimal test.

(d). Again for the sake of concreteness, compute the optimal possible power, for any 0.10 level test, at $\theta_1 = 0.50$ and at $\theta_2 = 1.00$. Compare these powers with that of the simple test $Y \geq 3.078$.

(e). (xx there are two drastic differences here, between the simple normal and the non-simple Cauchy. the first is that the log-density-ratio $R(y,\theta_0,\theta_1) = \log f(y,\theta_1) - \log f(y,\theta_0)$ is monotone, for the normal, and not at all monotone for the Cauchy. the second is that of there being a simple one-dimensional sufficient statistic, in the case of $Y_1,\ldots,Y_n$ from the normal, whereas no such statistic exists for the Cauchy. where is sufficiency in kiosk? xx)

(f). (xx something about more regularity with n data points; above we just did $n = 1$. xx)

Ex. 4.12. *The t test and its power.* Suppose $Y_1,\ldots,Y_n$ are i.i.d. from the $\mathrm{N}(\xi,\sigma^2)$, with testing of $\xi = 0$ required against $\xi > 0$. This is simple and standard, in the case of σ being known, but needs the t test in the case of σ being unknown and estimated from the data, as we have seen in Ex. 4.4, Setting up the test uses the relevant t distribution, from $t = \sqrt{n}\bar{Y}/\widehat{\sigma} \sim t_{n-1}$, but for studying the power function also the noncentral t distribution is required.

(a). Consider first the case of σ being known. Show that $z = \sqrt{n}\bar{Y}/\sigma$ is standard normal under the null, and that the 0.05 level test becomes that of rejecting when $z \geq z_0 = \Phi^{-1}(0.95) = 1.645$.

(b). Show that at a given $\xi > 0$, we have $z \sim \mathrm{N}(\sqrt{n}\xi/\sigma, 1)$, and that this leads to the power function $\pi_{n,N}(\xi/\sigma) = \Phi(\sqrt{n}\xi/\sigma - z_0)$. Compute and display this power function for the case of $n = 10$ and $n = 20$, as with the black full curves of Figure 4.1 (right panel).

(c). Then consider the more complex situation where σ is not known, needing the empirical standard deviation $\widehat{\sigma}$ of (1.6). We have seen in Ex. 1.50 that

$$t = \frac{\sqrt{n}\bar{Y}}{\widehat{\sigma}} = \frac{\sqrt{n}\bar{Y}/\sigma}{\widehat{\sigma}/\sigma} \sim \frac{\mathrm{N}(0,1)}{(\chi^2_{\mathrm{df}}/\mathrm{df})^{1/2}},$$

with nominator and denominator being independent, and where the degrees of freedom is $\mathrm{df} = n-1$. The probability density $g_{\mathrm{df}}(x)$ and cumulative distribution $G_{\mathrm{df}}(x)$ of this t_{df}

distribution are moderately complicated, see the exercise mentioned, but that does not concern us much, as long as we can consult a table or run an algorithm to find associated quantiles and probabilities. Show hence that the t test, with level 0.05, must consist of rejecting when $t \geq t_0 = G_{\rm df}^{-1}(0.95)$. Using `qt(0.95,df)` in `R`, we find 1.833 and 1.729 for $n = 10$ and $n = 20$. Check that `qt(0.95,df)` becomes close to 1.645 as df increases, and explain why.

(d). For the power of the t test, show that $\pi_{n,t}(\xi/\sigma) = \Pr_{\xi/\sigma}(t \geq t_0)$, where

$$t = \frac{\sqrt{n}\bar{Y}}{\widehat{\sigma}} \sim \frac{\mathrm{N}(\sqrt{n}\xi/\sigma, 1)}{(\chi^2_{\rm df}/{\rm df})^{1/2}},$$

again with nominator and denominator independent. Explain that the power function therefore can be written

$$\pi_{n,t}(\xi/\sigma) = \Pr_{\xi/\sigma}(t \geq t_{0.m}) = 1 - G_m(t_{0,m}, \sqrt{n}\xi/\sigma),$$

with $G_m(x, \lambda)$ denoting the c.d.f. for this noncentral t with degrees of freedom m and noncentrality parameter λ; see Ex. 1.51. This function is complicated, but can easily be found numerically via e.g. `pt(x,m,lambda)` in `R`. Construct a version of Figure 4.1, right panel, perhaps with other sample sizes than 10, 20. Comment on your findings.

(e). Describe how the two-sided tests and power functions pan out here, whem $\xi = 0$ is to be tested against $\xi \neq 0$. Make a corresponding version of Figure 4.1, right panel, with the relevant two-sided power functions.

Ex. 4.13. *Establishing that a parameter is close to zero.* Suppose δ is some effect parameter, where the scientific context relates to establishing that it is close to zero. In such situations it might be natural to test $H_0\colon |\delta| \geq \varepsilon$, for some known small threshold ε, against the alternative that $\delta \in (-\varepsilon, \varepsilon)$. In some applications the δ might be the difference between two effect parameters, aiming to infer that these are close. This is turning the tables, somehow, compared to the more traditional setups, where the hypothesis to be tested is that a parameter is close to zero, against the alternative that it is some distance away.

(a). To study the main features of such a situation, consider a prototype setup with i.i.d. observations $Y_1, \ldots, Y_n$ being $\mathrm{N}(\delta, 1)$. Show that $\sqrt{n}\bar{Y} \sim \mathrm{N}(\sqrt{n}\delta, 1)$, which implies $Z_n = n\bar{Y}^2 \sim \chi^2_1(n\delta^2)$.

(b). Explain that it is natural to reject the null, and hence claim that $|\theta|$ is small, if $|\sqrt{n}\bar{Y}| \leq c$, or equivalently $Z_n \leq c^2$, for c calibrated to reach testing level e.g. $\alpha = 0.05$. This is equivalent to $Z_n \leq c^2$. Show that $\Pr_\delta(|\sqrt{n}\bar{Y}| \leq c) = \Gamma_1(c^2, n\delta^2)$, in terms of the c.d.f. for the noncentral χ^2_1.

(c). Explain that to have significance level 0.05, we need $\Gamma_1(c^2, n\varepsilon^2) = 0.05$. Find the full power curve $\pi_n(\delta) = \Pr_\delta(Z_n \leq c^2)$, and comment on the level of its maximum. Construct a version of Figure 4.2, left panel, for threshold $\varepsilon = 0.20$, for sample sizes $n = 200$ and $n = 400$.

Ex. 4.14. *Power and local power: a particular case.* This exercise studies a prototype situation in some detail; the type of calculations and results will be seen to be rather similar in a long range of different situations. Let $Y_1, \ldots, Y_n$ be i.i.d. data from $\mathrm{N}(\theta, \sigma^2)$.

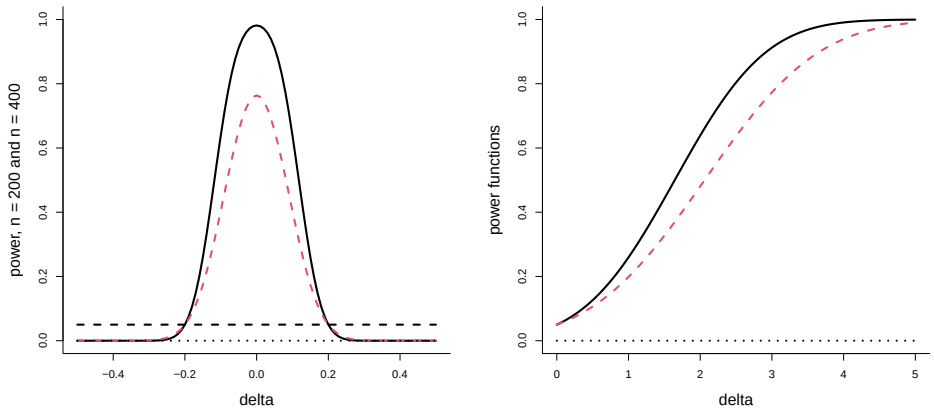


Figure 4.2: *Left panel: power functions for the bioequivalence tests of Ex. 4.13, with* $n = 200$ *(slanted curve) and* $n = 400$ *(full curve), and threshold* $\varepsilon = 0.20$*. Right panel: limiting local power functions for the two tests of Ex. 4.14, for* $\theta \leq \theta_0$ *against* $\theta > \theta_0$*, in terms of the local alternatives parameter* δ*, with* $\theta_n = \theta_0 + \delta/\sqrt{n}$*. The more powerful test is based on* $\bar{Y}$*, the other on the median* M_n*.*

One wishes to test $H_0 \colon \theta = \theta_0$ vs. the alternative that $\theta > \theta_0$, where θ_0 is a known value (e.g. 3.14). Two tests will be considered, based on respectively $\bar{Y}_n = n^{-1} \sum_{i=1}^n Y_i$ and the median M_n.

(a). For a given value of θ, prove that $\sqrt{n}(\bar{Y}_n - \theta) \to_d \mathrm{N}(0, \sigma^2)$ and $\sqrt{n}(M_n - \theta) \to_d \mathrm{N}(0, (\pi/2)\sigma^2)$. Note that the first result is immediate and actually holds with exactness for each n; the second result requires more care and follows from Ex. 3.20.

(b). Working under the null hypothesis $\theta = \theta_0$, show that

$$Z_n = \frac{\sqrt{n}(\bar{Y}_n - \theta_0)}{\widehat{\sigma}} \to_d \mathrm{N}(0,1), \quad Z_n^* = \frac{\sqrt{n}(M_n - \theta_0)}{(\pi/2)^{1/2}\widehat{\sigma}} \to_d \mathrm{N}(0,1),$$

where $\widehat{\sigma}$ is any consistent estimator of σ.

(c). With $z_{0.95} = \Phi^{-1}(0.95) = 1.645$, conclude from the above that the two tests that reject H_0 provided respectively

$$\bar{X}_n > \theta_0 + z_{0.95}\widehat{\sigma}/\sqrt{n} \quad \text{and} \quad M_n > \theta_0 + z_{0.95}(\pi/2)^{1/2}\widehat{\sigma}/\sqrt{n},$$

have the required asymptotic significance level 0.05; $\alpha_n = \Pr\{\text{reject } H_0 \,|\, \theta = \theta_0\} \to 0.05$. (There is one such α_n for the first test, and one for the other; both converge however to 0.05.)

(d). Our object is then to study the *local power*, the chance of rejecting the null hypothesis under alternatives of the type $\theta_n = \theta_0 + \delta/\sqrt{n}$, i.e. close to the null value. In generalisation of (b), show that

$$Z_n = \frac{\sqrt{n}(\bar{Y}_n - \theta_0)}{\widehat{\sigma}} \to_d \mathrm{N}(\delta/\sigma, 1), \quad Z_n^* = \frac{\sqrt{n}(M_n - \theta_0)}{(\pi/2)^{1/2}\widehat{\sigma}} \to_d \mathrm{N}((\pi/2)^{1/2}\delta/\sigma, 1),$$

where the convergence in question takes place under the indicated $\theta_0 + \delta/\sqrt{n}$ parameter values. Here we need to generalise the result of Ex. 3.20, to the case where θ_n moves with n. Writing F_n for the distribution of Y_i, you may do this by first showing that M_n has the same distribution as $F_n^{-1}(M_n^0) = \theta_0 + \delta/\sqrt{n} + \sigma\Phi^{-1}(M_n^0)$, where M_n^0 is the median for an i.i.d. sample from the uniform, and then using the delta method.

(e). Use these results to show that

$$\begin{aligned}\pi_n(\delta) &= \Pr\{\text{reject with } \bar{Y}_n \text{ test} \,|\, \theta_0 + \delta/\sqrt{n}\} \to \Phi(\delta/\sigma - z_{0.95}),\\ \pi_n^*(\delta) &= \Pr\{\text{reject with } M_n \text{ test} \,|\, \theta_0 + \delta/\sqrt{n}\} \to \Phi((2/\pi)^{1/2}\delta/\sigma - z_{0.95}),\end{aligned}$$

for the two power functions. Draw these in a diagram, and compare, as with Figure 4.2, right panel.

(f). Assume one wishes n to be large enough to secure that the power function is at least at level β for a certain alternative point θ_1. Using the local power approximation, show that the required sample sizes are respectively

$$n_A \doteq \frac{\sigma^2}{(\theta_1-\theta_0)^2}(z_{1-\alpha}+z_\beta)^2 \quad \text{and} \quad n_B \doteq \frac{\sigma^2/c^2}{(\theta_1-\theta_0)^2}(z_{1-\alpha}+z_\beta)^2$$

for tests A (based on the mean) and B (based on the median), with $c = \sqrt{2/\pi}$. With level $\alpha = 0.05$, compute these sample sizes for the case of $\beta = 0.95$ and $\theta_1 = \theta_0 + \frac{1}{2}\sigma$.

(g). One sometimes defines the ARE, the asymptotic relative efficiency of test B with respect to test A, as

$$\text{ARE} = \lim n_A(\theta_1,\beta)/n_B(\theta_1,\beta),$$

the limit in question in the sense of alternatives θ_1 coming closer to the null hypothesis at speed $1/\sqrt{n}$. Show that the ARE in this particular situation becomes $c^2 = 2/\pi = 0.6366$; test A needs only ca. 64 percent as many data points to reach the same detection power as B needs.

Ex. 4.15. *Power and local power: general results.* In Ex. 4.14 we examined two particular estimators and tests, for the mean parameter in the normal distribution, and found precise limit distributions and local power functions, for alternatives of order $1/\sqrt{n}$ from the null hypothesis. Here we go through the appropriate generalities. The setup is that of data, of sample size n, informing us about a key parameter θ, where the null hypothesis $H_0\colon \theta = \theta_0$ is to be tested against the alternative $\theta > \theta_0$, with θ_0 some specified value. We assume the data leads to one or more estimators $\widehat{\theta}$, with the basic premise that at the null parameter θ_0, we have $A_n = \sqrt{n}(\widehat{\theta} - \theta_0) \to_d \mathrm{N}(0,\kappa^2)$, for the appropriate κ.

(a). With $\widehat{\kappa}$ an estimator for κ, under H_0 conditions, show that $Z_n = \sqrt{n}(\widehat{\theta}-\theta_0)/\widehat{\kappa}_0 \to_d \mathrm{N}(0,1)$ under H_0. Explain that the test which rejects when $Z_n \geq z_{1-\alpha} = \Phi^{-1}(1-\alpha)$ has asymptotic testing level α, i.e. $\alpha_n = \Pr_0(Z_n \geq z_{1-\alpha}) \to \alpha$.

(b). We now make two more assumptions, concerning natural smoothness in a $O(1/\sqrt{n})$ size local neighbourhood around θ_0: (i) if the true parameter behind the data is $\theta_n = \theta_0 + \delta/\sqrt{n}$, then the same limit $\mathrm{N}(0,\kappa^2)$ obtains for $A_n' = \sqrt{n}\{\widehat{\theta} - (\theta_0 + \delta/\sqrt{n})\}$; (ii) the estimator $\widehat{\kappa}$ still converges to the same κ in probability. Show that $Z_n \to_d \mathrm{N}(\delta/\kappa, 1)$, and that this gives a limiting local power function,

$$\pi_n = \Pr(\text{reject} \,|\, \theta_0 + \delta/\sqrt{n}) \to \pi(\delta) = \Phi(\delta/\kappa - z_{1-\alpha}).$$

Show that this limit power function has derivative $\phi(z_{1-\alpha})/\kappa$ at zero.

(c). The above result is useful in its own right, providing an easy approximation to the power function for alternatives not far from the null hypothesis value. It may also be used for comparing different tests, built as above, for different estimators, say $\widehat{\theta}_j$. Assume the conditions above hold for these, with $\sqrt{n}(\widehat{\theta}_j - \theta_0) \to_d \mathrm{N}(0, \kappa_j^2)$, leading to estimators $\widehat{\kappa}_j$ for κ_j and to test statistics $Z_{n,j} = \sqrt{n}(\widehat{\theta}_j - \theta_0)/\widehat{\kappa}_j$. Explain that the best estimators lead to the best tests, with derivatives $\phi(z_{1-\alpha})/\kappa_j$ at zero.

(d). (xx briefly about ARE, as above. xx)

Ex. 4.16. *Optimal average power.* Supppose Y is observed, perhaps a full vector, from a density f, where one wishes to test the null hypothesis $f = f_0$, a given density. As we saw in Ex. 4.11, there are cases where there is no uniformly optimal test, against all or a subset of alternatives; the optimal test at f_1 might be different from the one at f_2. In one-parameter models, this is caused by the log-density-ratio not being monotone.

(a). Consider alternative densities $f_1, \ldots, f_m$, given nonnegative weights of importance $w_1, \ldots, w_m$. These may be taken to have sum 1. The *weighted average power* of a test T, at these points and with these weights, is

$$\bar{\pi}_T = \sum_{j=1}^m w_j \pi_T(f_j) = \sum_{j=1}^m w_j \int f_j(y) T(y)\,\mathrm{d}y,$$

with $\pi_T(f_j)$ the power at f_j. Let $\bar{f}(y) = \sum_{j=1}^m w_j f_j$. Show that this average power is maximised, among all test functions $T(y)$ with level α at the null, by the $T^*(y)$ which rejects H_0 when $\bar{f}(y)/f_0(y) \geq c$, with c tuned to give rejection probability α at the null.

(b). For a one-parameter model $f(y, \theta)$, consider testing of $\theta = \theta_0$ against $\theta > \theta_0$. For any test function $T(y)$ with rejection level α at the null, consider in general terms the weighted average power

$$\bar{\pi}_T = \int_{\theta > \theta_0} \pi_T(\theta)\,\mathrm{d}w(\theta),$$

with $\pi_T(\theta) = \int f(y, \theta) T(y)\,\mathrm{d}y$ the power of the test at position θ. Show that $\bar{\pi}_T = \int_{\theta > \theta_0} \bar{f}(y) T(y)\,\mathrm{d}y$, featuring the density $\bar{f}(y) = \int_{\theta > \theta_0} f(y, \theta)\,\mathrm{d}w(\theta)$. This is the model density averaged over all alternatives to the null, as weighted by the $\mathrm{d}w(\theta)$ measure.

(c). Show that the test maximising this weighted average power is rejecting the null if $\bar{f}(y)/f(y, \theta_0) \geq c$, with c tuned to have null level α.

(d). Suppose θ is one-dimensional and that the log-density $\log f(y, \theta)$ is monotone in θ. Explain that there then is a de facto uniformly optimal test for $\theta = \theta_0$ vs. $\theta > \theta_0$, and that the average power optimal test worked with here will be identical to this one, as long as the weight function is positive only for $\theta > 0$.

(e). (xx a little more. the marginal density, or predictive density, with a link to Bayes, but specifically with a 'prior' over the alternative space. can take Cauchy with $a \exp(-a\theta)$ over the halfline. xx)

(f). Consider the $\mathrm{N}(\xi, \sigma^2)$ model, and assume there is a need to test $f = \mathrm{N}(0, 1)$ against the alternative that $\xi > 0$, or $\sigma > 1$, or both. show first that the NP test against alternatives $(1.1, 2.2)$ and $(2.2, 3.3)$ are indeed different, so there is no uniformly most

powerful test. Then if a weight function $w(\xi,\sigma)$ is specified, in the domain $\xi>0, \sigma>1$, show that the average power optimal test consists in rejecting if

$$V(y_1,\ldots,y_n)=\frac{\bar f(y_1,\ldots,y_n)}{f_0(y_1)\cdots f_0(y_n)}=\int_{\xi>0,\sigma>1}\Big\{\prod_{i=1}^n\frac{f(y_i,\xi,\sigma)}{f(y_i,0,1)}\Big\}\,\mathrm{d}w(\xi,\sigma)$$

is larger than the treshold c, with this finetuned to be the upper quantile of its null distributon. As an illustration, suppose $n=10$ and that the weight function is flat on $[0,4]\times[1,6]$. Simulate from $W(y_1,\ldots,y_n)=\log V(y_1,\ldots,y_n)$ under the null hypothesis, as per Ex. 4.10, to find the rejection threshold, to have significance level 0.01.

(g). (xx connect this to Wilks things in Ch5, which will work as a good approximation when testing an inner point of the parameter region. xx)

Sufficiency, factorisation theorem, completeness

Ex. 4.17. *Sufficiency.* Flip a coin ten times and record the number of heads. It is intuitively clear that the number of heads in the ten tosses is as informative about the unknown $\theta=\Pr(\text{heads})$ of the coin, as also knowing the exact ordering in which the heads and tails occurred. In fact, knowing the ordering in which the heads and tails occurred appears redundant for making inference about θ. This leads us to the notion of a sufficient statistic. Let X be data with a distribution from a family $\mathcal{P}=\{P_\theta\colon\theta\in\Theta\}$ of distributions on the measurable space $(\mathcal{X},\mathcal{B})$. In colloquial terms, the statistic $T=T(X)$ is sufficient for θ if for every t and θ, the conditional distribution of X given $T=t$ is the same for all values of θ. More formally, let $T\colon\mathcal{X}\to\mathcal{Y}$ be a statistic, then T is sufficient for $\mathcal{P}$ (or for θ) if there is a function $Q_t(B)$ of $B\in\mathcal{B}$ and $t\in\mathcal{Y}$, such that for each B and $\theta\in\Theta$, we have $Q_t(B)=P_\theta(B\,|\,T=t)$ for $P_\theta T^{-1}$ almost all t, that is

sufficient statistic

$$\int_A Q_t(B)\,\mathrm{d}P_\theta T^{-1}(t)=\int_A P_\theta(B\,|\,T=t)\,\mathrm{d}P_\theta T^{-1}(t)=P_\theta(B\cap T^{-1}(A)),$$

for each θ and for all A in the Borel-σ-algebra $\mathcal{B}(\mathcal{Y})$ on $\mathcal{Y}$.

(a). Let $X_1,\ldots,X_n$ be i.i.d. Bernoulli(θ), and set $T=\sum_{i=1}^n X_i$. Show that for each t, $\Pr_\theta(X_1=x_1,\ldots,X_n=x_n\,|\,T=t)=1/\binom{n}{t}$ on $\{(x_1,\ldots,x_n)\colon x_i=0,1,\sum_{i=1}^n x_i=t\}$, and conclude that T is sufficient for θ.

(b). Let $Y_1,\ldots,Y_n$ be i.i.d. Pois(θ) and set $T=\sum_{i=1}^n Y_i$. Show that for each t

$$\Pr_\theta(Y_1=y_1,\ldots,Y_n=y_n\,|\,T=t)=\frac{t!}{n^t}\prod_{i=1}^n\frac{1}{y_i!},$$

on $\{(y_1,\ldots,y_n)\colon y_i=0,1,2,\ldots,\sum_{i=1}^n y_i=t\}$, and conclude that T is sufficient for θ.

(c). Let $X_1,\ldots,X_n$ be i.i.d. uniform on $(0,\theta)$ and set $T=\max_{i\le n}X_i$. Show that the conditional distribution of $X_1,\ldots,X_n$ given $T=t$ has density

$$f_\theta(x_1,\ldots,x_n\,|\,t)=\frac{1}{t^n},$$

on $\{(x_1,\ldots,x_n)\colon 0\le x_i\le t\text{ for all }i\}$, and conclude that T is sufficient for θ.

(d). Suppose that $Y \sim \mathrm{N}(0, \sigma^2)$ and let $Z = \pm 1$ with equal probability be independent of Y. Show that $Y' = Z|Y|$ has the same distribution as Y. Consider the function

$$Q_t(B) = \tfrac{1}{2} I_B(t) + \tfrac{1}{2} I_B(-t), \quad B \in \mathcal{B},\, t \geq 0,$$

and show, using that $Y' \sim Y$, that $|Y|$ is sufficent for σ^2.

(e). Let X and Y be independent $\mathrm{Expo}(\theta)$, and set $T = X + Y$. Let $U \sim \mathrm{unif}(0,1)$ be independent of X and Y, and define $X' = UT$ and $Y' = (1-U)T$. Show that $(X, Y) \sim (X', Y')$. From this, show that T is sufficient for θ.

(f). Let X be a random vector with a distribution from $\mathcal{P} = \{P_\theta \colon \theta \in \Theta\}$. Assume that T is sufficient for θ, and let $Q_t(B)$ be the function from the definition of sufficiency. Since $Q_t(B)$ is a version of $P_\theta(B \,|\, T = t)$, and X is a random vector, there exists a version of Q_t under $P_\theta T^{-1}$, $Q_t^\star$, such that $B \mapsto Q_t^\star(B)$ is a probability measure for all t [xx (see Notes and Pointers in the appendix) xx]. Given $T = t$, let X' be a random vector drawn from $Q_t^\star$. Show that $X' \sim X$.

Ex. 4.18. *The factorisation theorem.* In Ex. 4.17 we found sufficient statistics through qualified guesswork. The factorisation theorem provides a recipe for finding sufficient statistics through (minor) manipulations of the density of a sample. Suppose that the family $\mathcal{P} = \{P_\theta \colon \theta \in \Theta\}$ of distributions on $(\mathcal{X}, \mathcal{B})$ is dominated by a σ-finite measure μ. For each θ, let f_θ be the density of P_θ with respect to μ. The factorisation theorem says that the statistic $T \colon \mathcal{X} \to \mathcal{Y}$ is sufficient for $\mathcal{P}$ if and only if there exist nonnegative functions h and g_θ such that

$$f_\theta(x) = g_\theta(T(x))h(x),$$

for all $\theta \in \Theta$. We prove this theorem step by step.

(a). Assume that $\mathcal{P}$ is dominated by a probability measure P, and that for each $\theta \in \Theta$, the density $f_\theta = \mathrm{d}P_\theta/\mathrm{d}P$ is such that $f_\theta(x) = g_\theta(T(x))$, i.e., $h(x) = 1$ for all $x \in \mathcal{X}$. Show that T is sufficient for $\mathcal{P}$, and indeed also for the slightly larger family $\mathcal{P} \cup \{P\}$.

(b). For the converse of (a), assume that $\mathcal{P}$ is dominated by a probability measure P and that T is sufficient for $\mathcal{P} \cup \{P\}$. Let f_θ be a density of P_θ with respect to P, and show that $g_\theta(T(\cdot)) = \mathrm{E}_P(f_\theta \,|\, T)(\cdot)$ is also a density of P_θ with respect to P, i.e., that $\int_B g_\theta(T(x))\,\mathrm{d}P(x) = P_\theta(B)$ for all $B \in \mathcal{B}$. –Taken together, (a) and (b) prove the factorisation theorem in the case that $\mathcal{P}$ is dominated by (at least) one of its members, which is not wholly unusual, consider the family $\mathcal{P} = \{\mathrm{N}(\theta, 1) \colon \theta \in \mathbb{R}\}$, for example.

(c). To prove the factorisation theorem in full generality, we use the result from Ex. [xx abc xx] in the appendix, which says that if $\mathcal{P} = \{P_\theta \colon \theta \in \Theta\}$ is dominated by a σ-finite measure μ, then there is a probability measure of the form $P = \sum_{j=1}^\infty a_j P_{\theta_j}$ equivalent to $\mathcal{P}$, i.e., $P_\theta(A) = 0$ for all θ if and only if $P(A) = 0$. Now, assume that for each $\theta \in \Theta$, P_θ has density $g_\theta(T(x))h(x)$ with respect to μ, and show that T is sufficient.

(d). Conversely, assume that $\mathcal{P}$ is dominated by μ and that T is sufficient. Show that P_θ has density $g_\theta(T(x))h(x)$ with respect to μ.

(e). Try to make sense of the colloquial definition of sufficiency given in Ex. 4.17, and show that it says the same thing as the formal definition of sufficiency.

Ex. 4.19. *Some more sufficiency.* In this exercise we employ the factorisation theorem on some concrete examples, and touch on some important facts about sufficiency.

(a). Use the factorisation theorem to find sufficient statistics for i.i.d. samples from, respectively, the Bernoulli(θ) distribution; the $\mathrm{N}(\mu, 1)$ distribution; the $\mathrm{N}(\mu, \sigma^2)$ distribution with both parameters unknown; the $\mathrm{unif}(0, \theta)$ distribution; and, lastly, the $\mathrm{unif}(\theta, \theta + 1)$ distribution. What is particular about this latter?

(b). Let X, Y be i.i.d. from $\mathcal{P} = \{\mathrm{unif}(0, \theta) \colon \theta \geq 1\} \cup \{\mathrm{unif}\{1, 2, \ldots, \theta\} \colon \theta \in \mathbb{N}\}$. Show that $T = \max(X, Y)$ is sufficient for $\mathcal{P}$. Find also an expression for the function $Q_t(B)$ from the definition of sufficiency (see Ex. 4.17), and show that this function may be chosen so that $B \mapsto Q_t(B)$ is not a probability measure for all t.

(c). A sufficient statistic is not unique, and different sufficient statistics may provide varying degress of data compression. At one extreme are sufficient statistics not providing any compression of the data. Show that the full sample is sufficient.

(d). Let $X_1, \ldots, X_n$ be i.i.d. from an unknown distribution with c.d.f. belonging to a family of continuous c.d.f.s. Let $T = (X_{(1)}, \ldots, X_{(n)})$ be the order statistics, and show that it is sufficient.

(e). Show that any one-to-one transformation of a sufficient statistic is sufficient.

(f). Let $X_1, \ldots, X_n$ be i.i.d. from a member of $\mathcal{P} = \{\mathrm{N}(0, \sigma^2) \colon \sigma > 0\}$ and consider the statistics $T_1 = (X_1, \ldots, X_n)$, $T_2 = (X_1^2, \ldots, X_n^2)$, $T_3 = (X_1^2 + \cdots + X_k^2, X_{k+1}^2 + \cdots + X_n^2)$, and $T_4 = X_1^2 + \cdots + X_n^2$. Here, the data compression is increasing in the indices. Use the factorisation theorem to prove that all these statistics are all sufficient.

Ex. 4.20. *Sufficiency in exponential families* (xx need intro some technical but important things for the expo family. calibrate how and when we do sufficiency. xx)

(a). Suppose that Y stem from a member of the family $\mathcal{P} = \{P_\theta \colon \theta \in \Theta\}$, where each P_θ has density of the exponential family form $f_\theta(y) = \exp(\theta T(y) - k(\theta))h(y)$ with respect to some dominating measure μ. Show that $T = T(Y)$ is sufficient for the natural parameter θ; and that the distribution of T has a density of the exponential family form, say $g_\theta(t) = \exp(\theta t - k(\theta))$, with respect to νT^{-1}, where $\mathrm{d}\nu(y) = h(y)\,\mathrm{d}\mu(y)$.

(b). Let $Y \sim \mathrm{N}(0, \sigma^2)$. Write this normal density on exponential family form, and check the formula from (a) for the density $g_\theta(t)$ using the transformation methods of Ex. 1.12.

(c). Generalise (a) to the case where $Y_1, \ldots, Y_n$ are i.i.d. from a distribution in $\mathcal{P} = \{P_\theta \colon \theta \in \Theta\}$, where each P_θ has density $f_\theta(y) = \exp\{\theta^{\mathrm{t}} T(y) - k(\theta)\}h(y)$ with respect to some dominating measure μ. Show that $\bar{T} = (\bar{T}_1, \ldots, \bar{T}_p)^{\mathrm{t}}$, with $\bar{T}_j = (1/n)\sum_{i=1}^n T_j(Y_i)$ for $j = 1, \ldots, p$, is sufficient for the natural parameter $\theta = (\theta_1, \ldots, \theta_p)$; and that $\bar{T}$ has a distribution with density following the same exponential form with respect to an appropriate dominating measure.

(d). Here is a lemma that is useful for the following exercises. Let $\Pr_0$ and $\Pr_0$ be probability measures on the same measurable space $(\Omega, \mathcal{F})$, and let X and Y be random vectors with values in the measurable spaces $(\mathcal{X}, \mathcal{A})$ and $(\mathcal{Y}, \mathcal{B})$. Suppose that there are functions $a \colon \mathcal{X} \to (0, \infty)$ and $b \colon \mathcal{Y} \to [0, \infty)$ such that $\Pr_1$ has density $a(X)b(Y)$ with respect to $\Pr_0$. Show that the marginal distribution of Y under $\Pr_1$ is

$$\Pr_1 Y^{-1}(B) = \int_B b(y) \mathrm{E}_0 \left\{ a(X) \,|\, Y = y \right\} \mathrm{d}\Pr_0 Y^{-1}(y), \quad B \in \mathcal{B},$$

and that for each $A \in \mathcal{A}$,

$$\Pr_1(X \in A \,|\, Y = y) = \frac{\mathrm{E}_0\{I_A(X)a(X) \,|\, Y = y\}}{\mathrm{E}_0\{a(X) \,|\, Y = y\}}, \quad \Pr_1 Y^{-1} \text{ a.e.},$$

that is, under $\Pr_1 Y^{-1}$ the right hand side is a version of the conditional distribution of X given $Y = y$.

(e). Let $\{P_{a,b} \colon (a,b) \in \Theta = \Theta_a \times \Theta_b\}$ be a family of distributions having densities of the exponential family form $f_{a,b}(y) = \exp\{a^{\mathrm{t}}U(y) + b^{\mathrm{t}}V(y) - k(a,b)\}h(y)$ with respect to μ. Use (d) to show that the marginal distribution of $V = V(Y)$ has density

$$\exp(b^{\mathrm{t}}v - k(a,b)),$$

with respect to the measure

$$\mathrm{d}\lambda_a(v) = \exp(-b_0^{\mathrm{t}}v + k(a_0,b_0))\mathrm{E}_{a_0,b_0}\{\exp((a-a_0)^{\mathrm{t}}U) \,|\, V = v\}\,\mathrm{d}P_{a_0,b_0}V^{-1}(v),$$

where (a_0, b_0) is some fixed point in Θ. The important thing here is that V has a density belonging to an exponential family, not the exact form of the measure with respect to which this is a density.

(f). Use (d) to show that the conditional distribution of U given $V = v$ has density,

$$\exp(a^{\mathrm{t}}u - k_v(a)), \quad \text{with } k_v(a) = \log \mathrm{E}_{a_0,b_0}\{\exp((a-a_0)^{\mathrm{t}}U) \,|\, V = v\},$$

with respect to the measure $\mathrm{d}\nu_v(u) = \exp(-a_0^{\mathrm{t}}u)\,\mathrm{d}P_{a_0,b_0}(U^{-1}(u) \,|\, V = v)$. In particular, the conditional density of U given $V = v$ does not depend on b.

conditional sufficiency

(g). Show that, conditionally on $V = v$, U is sufficient for the family $\{P_a^{U \,|\, V=v} \,|\, a \in \Theta_a\}$. where $P_a^{U \,|\, V=v}$ is the conditonal distribution of U given $V = v$.

(h). Let $X \sim \mathrm{Expo}(b)$ and $U \sim \mathrm{Expo}(b+a)$ be independent, and define $V = X + Z$. By direct methods, for example the convolution formula in (1.1), find the density of the distribution of V, and show that the distribution of U given $V = v$ has density

$$f_a(u \,|\, v) = \frac{a\exp(-au)}{1 - \exp(-av)}, \quad 0 \le u \le v.$$

Verify that the expression from (e) and (f) gives these same densities (with respect to Lebesgue measure), as they should.

Ex. 4.21. *Minimal sufficiency.* In Ex. 4.19 we saw that for any model there are many different sufficient statistics, often with some providing more compression of the data than others. Since the purpose of a sufficient statistic is to compress the data, this naturally leads to a search for a sufficient statistic providing the maximum amount of data compression.If such a statistic exits, it is said to be minimal sufficient. Here is the definition: Let T be sufficient for $\mathcal{P} = \{P_\theta \colon \theta \in \Theta\}$. Then T is minimal sufficient if for any other sufficient statistic S, there is a measurable function g so that $P_\theta\{T = g(S)\} = 1$ for all $\theta \in \Theta$. Another way of saying this is that if T is such $S(x) = S(y)$ implies $T(x) = T(y)$ for any sufficient statistic S, then T is minimal sufficient.

minimal sufficient

(a). Show that any one-to-one function of a minimal sufficient statistic is a minimal sufficient statistic.

(b). Suppose that T is minimal sufficient, and let S be some sufficient statistic. Show that the σ-algebra generated by T must be contained in the σ-algebra generated by S.

(c). Show that if $T(x) = T(y)$ if and only if

$$\theta \mapsto \frac{f_\theta(x)}{f_\theta(y)} \text{ is constant in } \theta,$$

then T is minimal sufficient. Here we assume, for simplicity, that $f_\theta(y) > 0$ for all y. If this is not the case, then instead consider the case where $T(x) = T(y)$ if and only if $y \in \{y\colon f_\theta(y) = f_\theta(x)h(x,y)$ for all θ and some $h > 0\}$.

(d). With $X \sim \mathrm{N}(0, \sigma^2)$, show that the absolute value $|X|$ is minimal sufficient. Let $Y_1, \ldots, Y_n$ be i.i.d. $\mathrm{N}(\mu, \sigma^2)$, and show that $(\bar{Y}_n, S_n)$ is minimal sufficient, where $\bar{Y}_n = n^{-1}\sum_{i=1}^n X_i$ and $S_n = \sum_{i=1}^n (Y_i - \bar{Y})^2$. Let $Z_1, \ldots, Z_n$ be i.i.d. from a distribution with density $f_\theta(z) = \exp(-(z-\theta))$ for $z > \theta$ and $\theta \in \mathbb{R}$. Find a minimal sufficient statistic for θ.

(e). Let $g(x)$ be a positive and integrable function on $\mathbb{R}$, set $c(a,b)^{-1} = \int_a^b g(x)\,\mathrm{d}x$, and define $f_{a,b}(x) = c(a,b)g(x)I_{(a,b)}(x)$, and suppose that $X_1, \ldots, X_n$ are i.i.d. from the distribution with density $f_{a,b}(x)$. Find a minimal sufficient statistic for (a,b).

(f). Let $X_1, \ldots, X_n$ be i.i.d. from a distribution with density $f_\theta(x) = \frac{1}{2}\exp(-|x-\theta|)$ with $x, \theta \in \mathbb{R}$. Show that the order statistics are minimal sufficient.

Ex. 4.22. *Ancillary statistics.* Ancillarity is, in a sense, the opposite of sufficiency. If X stem from a member of $\mathcal{P} = \{P_\theta\colon \theta \in \Theta\}$, a statistic $V = V(X)$ is ancillary if its distribution is the same for all θ. In other words, V by itself does not provide any information about θ. This does not, however, mean that V should be disregarded when making inference on θ. It just means that if you only learn $V = v$, you have not learned anything about θ. ancillary statistic

(a). Let $Y \,|\, (V = v) \sim \mathrm{N}(\mu, \sigma_V^2)$ where V is Bernoulli $(\frac{1}{2})$. Why does it make the most sense to condition on $V = v$ when deriving a confidence interval for μ?

(b). Let X and Y be independent $\mathrm{unif}(\theta - \frac{1}{2}, \theta + \frac{1}{2})$, and set $T = \max(X,Y)$ and $V = |X - Y|$. Show that Show that (T,V) is minimal sufficient, and that V is ancillary.

(c). With T and V as in (b), compute the probability $\Pr(T - V \le \theta \le T)$ and the conditonal probability given $V = v$, $\Pr(T - \frac{1}{4}(1+v) \le \theta \le T + \frac{1}{4} - \frac{3}{4}v \,|\, V = v)$. Which of these two confidence intervals makes the most sense to you?

(d). Compute the probability $\Pr(T - V \le \theta \le T)$ and the conditonal probability given $V = v$, $\Pr(T - \frac{1}{4}(1+v) \le \theta \le T + \frac{1}{4} - \frac{3}{4}v \,|\, V = v)$. Which of these two confidence intervals makes the most sense to you?

(e). Let $(X_1, Y_1), \ldots, (X_n, Y_n)$ be i.i.d. with density $f_\theta(x,y) = \exp(-\theta x - y/\theta)$ for $x, y > 0$ and $\theta > 0$. Find the maximum likelihood estimator $\widehat{\theta}_n$ for θ. Show that $(\widehat{\theta}_n, V)$ is minimal sufficient, where $V = (\sum_{i=1}^n X_i)^{1/2}(\sum_{i=1}^n Y_i)^{1/2}$; and that V is ancillary.

(f). In many examples, the maximum likelihood estiamator is a minimal sufficient statistic. In (e) it is not. One way of measuring the information loss in only knowing $\widehat{\theta}_n$ compared to the full minimal sufficient statistic $(\widehat{\theta}_n, V)$, is by way of the Fisher information. Let $g_{\theta,n}$ be the density of $\widehat{\theta}_n$, and compute the difference $-\mathrm{E}\,\{\mathrm{d}^2/\mathrm{d}\theta^2 \sum_{i=1}^n \log f_\theta(X_i, Y_i)\} + \mathrm{E}\,\{\mathrm{d}^2/\mathrm{d}\theta^2 \sum_{i=1}^n \log g_{\theta,n}(\widehat{\theta}_n)\}$.

Ex. 4.23. *Completeness.* Often, models are so harmoniously constructed that there are clear one-to-one connections between estimators (perhaps based on a set of summary statistics) and estimands, in the sense that there for each estimand is only one unbiased estimator. Clarifying such regularity leads to the concept of *completeness*, which turns out to be useful also when coming to conditional testing and optimal power in exercises below. Technically, suppose some vector $T = (T_1, \ldots, T_p)^{\mathrm{t}}$ has a distribution $f(t, \theta)$, with the property that $\mathrm{E}_\theta\, h(T) = 0$ for all $\theta \in \Theta$ implies $\Pr_\theta\{h(T) = 0\} = 1$ for all θ, i.e. $h(t) = 0$ almost everywhere. We then say that T, or more formally its distribution, over the relevant parameter region, is complete.

complete

(a). Let $X \sim \mathrm{binom}(n, \theta)$, with $\theta \in (0, 1)$. Show that X is complete; zero is the only unbiased estimator of zero. (You may appeal to properties of power series.) Show that X is complete as long as the parameter region contains an open interval. Show similarly that if $X \sim \mathrm{geom}(p)$, see Ex. 1.25, then X is complete, again requiring only that the parameter range for p contains an open interval.

(b). With $Y_1, \ldots, Y_n$ i.i.d. from the uniform on $[0, \theta]$, consider $M = \max_{i \le n} Y_i$. Show that if the parameter region is the full halfline $\theta > 0$, then M is sufficient and complete, but that M is not complete if it is a priori known that $\theta \ge 1.234$, i.e. with a restricted parameter range.

(c). Suppose $Y_1, \ldots, Y_n$ are i.i.d. from the uniform distribution over $[\theta - 1, \theta + 1]$. Show that $(Y_{(1)}, Y_{(n)})$, the smallest and largest, is sufficient, but not complete.

(d). Suppose $Y_1, \ldots, Y_n$ are i.i.d. from the uniform distribution over $[a, b]$. Show that $(Y_{(1)}, Y_{(n)})$ is sufficient and complete. [xx check this xx]

(e). Consider Y being uniform on $\{1, \ldots, \theta\}$, where $\theta \in \{1, 2, \ldots\}$ is an unknown parameter. Show that Y is complete. When the parameter region is e.g. $\{4, 5, \ldots\}$, however, show that Y is not complete.

(f). If T is complete, show that any one-to-one transformation variable $T' = a(T)$ is also complete.

(g). Consider $Y_1, \ldots, Y_n$ an i.i.d. sample from the double exponential with density $f(y, \theta) = \frac{1}{2}\exp(-|y - \theta|)$. Show that the full set of order statistics $(Y_{(1)}, \ldots, Y_{(n)})$ is sufficient, but not complete; do this, by exhibiting two different unbiased estimators of θ.

Ex. 4.24. *Completeness for the exponential family.* For the large class of exponential family models, see Ex. 1.55 (xx and follow-up exercises above xx), there is a completeness lemma, as follows. Suppose $Y_1, \ldots, Y_n$ are i.i.d. from the model $f(y, \theta) = \exp\{\theta^{\mathrm{t}} T(y) - k(\theta)\}h(y)$, with $T = (T_1, \ldots, T_p)^{\mathrm{t}}$ and $\theta = (\theta_1, \ldots, \theta_p)^{\mathrm{t}}$ varying in an open set, then the vector of sample averages $(\bar{T}_1, \ldots, \bar{T}_p)^{\mathrm{t}}$ is not merely sufficient, as seen in Ex. 4.21, but also complete. We shall freely use this lemma. (xx but point to proof, check BickelDoksum or Johansen or Brown or Schervish. perhaps requiring analyting continuation arguments. point is that $a(\theta) = \mathrm{E}_\theta\, h(\bar{T})$ is a super smooth functions with all derivatives smooth. xx)

completeness lemma for exponential family

(a). Let $Y_1, \ldots, Y_n$ be i.i.d. from the $\mathrm{N}(\xi, 1)$. Show that the full set $(Y_1, \ldots, Y_n)$ is not complete, but that the sample mean $\bar{Y}$ is.

(b). Consider $Y_1, \ldots, Y_n$ i.i.d. from the $\mathrm{Gam}(a, b)$ model. Show that $(\sum_{i=1}^n Y_i, \sum_{i=1}^n \log Y_i)$ is sufficient and complete. Identify similarly a sufficient and complete pair of statistics for a sample from the $\mathrm{Beta}(a, b)$.

(c). Consider an i.i.d. sample $Y_1, \ldots, Y_n$ from the $\mathrm{N}(\xi, \sigma^2)$. Show that $(\sum_{i=1}^n Y_i, \sum_{i=1}^n Y_i^2)$ is sufficient and complete, and also that $(\bar{Y}, \widehat{\sigma})$ is sufficient and complete. Suppose then that the variance is postulated to be equal to the squared mean, so that the sample is from $\mathrm{N}(\theta, \theta^2)$. Construct two different unbiased estimators of θ, and show that this means that $(\bar{Y}, \widehat{\sigma})$ is not complete. You may similarly construct two different unbiased estimators of θ^2.

(d). Consider the linear regression model with $Y_i \sim \mathrm{N}(a + bx_i, \sigma^2)$, studied in Ex. 3.33. With $\widehat{a}$, $\widehat{b}$ the least squares estimators, and $Q_0 = \sum_{i=1}^n \{Y_i - \widehat{a} - \widehat{b}(x_i - \bar{x})\}^2$, show that the log-likelihood can be written $-n \log \sigma - \frac{1}{2}\{Q_0 + n(\widehat{a} - a)^2 + M_n(\widehat{b} - b)^2\}^2/\sigma^2$, with $M_n = \sum_{i=1}^n (x_i - \bar{x})^2$. Write this in the exponential family fashion, with natural parameters $1/\sigma^2, a/\sigma^2, b/\sigma^2$. Argue via the general exponential family results that $(Q_0, \widehat{a}, \widehat{b})$ is both sufficient and complete. Extend this, arguments and results, to the general multiple linear regression model of Ex. 3.34.

Ex. 4.25. *Basu's Lemma.* Consider a setting with data Y from some parametric family, indexed by $\theta \in \Theta$, where a suitable T is sufficient and complete.

(a). Assume that a suitable statistic $Z = Z(Y)$ has a distribution not depending on the θ. Show that Z is independent of T. This is called Basu's Lemma. You may follow this path: start writing $\Pr_\theta(Z \in A) = p$, by assumption not depending on θ. Show that $p = \mathrm{E}_\theta\, h(T)$, for all θ, where $h(t) = \Pr_\theta(Z \in A \,|\, T = t)$, another function not depending on θ. Explain that this implies $\Pr_\theta(Z \in A \,|\, T = t) = p$ for all θ. Basu's lemma

(b). Consider the familiar setup with $Y_1, \ldots, Y_n$ being i.i.d. from $\mathrm{N}(\mu, \sigma^2)$. For σ fixed, show that the mean $\bar{Y}$ is sufficient and complete for μ; explain that all statistics Z with distribution not depending on μ must be independent of $\bar{Y}$. Explain hence in particular that $\sum_{i=1}^n (Y_i - \bar{Y})^2$ is independent of $\bar{Y}$. This has been demonstrated in different ways in Ex. 1.48 and 1.49.

(c). In this normal sample setting, argue that the classic t ratio $t = \sqrt{n}(\bar{Y} - \mu)/\widehat{\sigma}$, with $\widehat{\sigma}$ the sample standard deviation, is independent of $(\bar{Y}, \widehat{\sigma})$.

(d). We have seen in Ex. 1.56 that the five-parameter binormal distribution is inside the exponential family class. With $(X_1, Y_1), \ldots, (X_n, Y_n)$ i.i.d. from the binormal, show first that $(\bar{X}, \bar{Y}, \widehat{\sigma}_1, \widehat{\sigma}_2, R_n)$ is complete and sufficient, with R_n the empirical correlation coefficient. Show that R_n is independent of $\bar{X}, \bar{Y}, \widehat{\sigma}_1, \widehat{\sigma}_2$ (xx check this with care xx).

Optimal conditional testing

Ex. 4.26. *Conditional tests.* Suppose in general terms that data y are observed from a model with parameter θ, where the null hypothesis $H_0 \colon \theta \in \Theta_0$ is to be tested, against the alternative that $\theta \notin \Theta_0$. Assume one computes $U = U(y)$ and $V = V(y)$. A *conditional test*, with respect to V, with level α, is then to find a rejection region $R(v)$, using the distribution of U given $V(y) = v$, with conditional tests

$$\Pr_\theta\{U(Y) \in R(v) \,|\, V(Y) = v\} \le \alpha \quad \text{for all } \theta \in \Theta_0.$$

Such tests are natural and important in multiparameter setups, as we shall see, and various constructions succeed in 'reducing the dimensionality' down to the analysis of a one-parameter family, where e.g. Neyman–Pearson more readily applies.

(a). Even when such a test has been constructed in a conditional modus, 'what is unlikely null behaviour of U given that $V = v$', it may of course be translated or paraphrased without the conditioning: one rejects if $U(Y) \in R(V)$. Show that the test also has unconditional level α.

(b). From unconditional to conditional: Conditional tests as above have the form $T(U, V) = I(U \in R(V))$, built to have $\mathrm{E}_\theta\{I(U \in R(v)) \,|\, v\} = \alpha$. Assume now that V is complete at the boundary $\partial\Theta_0$ of the null hypothesis parameter region; see Ex. 4.23. Show that if *any* test $T(U, V)$ has constant level α at this boundary, then it is a conditional test, with this level; $\mathrm{E}_\theta\{T(U, V) \,|\, V = v\} = \alpha$ for all $\theta \in \bar{\Theta}_0$. (It might be useful to check the function $h(v) = \mathrm{E}_\theta\{T(U, V) \,|\, V = v\} - \alpha$.)

Ex. 4.27. *Conditional tests: pairs of exponentials.* Suppose $X \sim \mathrm{Expo}(a)$ and $Y \sim \mathrm{Expo}(a + \delta)$, and that one wishes to test $\delta = 0$, i.e. equal distributions, against $\delta > 0$.

(a). Show that the joint density may be written $a(a + \delta)\exp(-az - \delta y)$, with $z = x + y$. Find the distribution of $Z = X + Y$, and show that the distribution of Y given $Z = z$ has the density

$$g_\delta(y \,|\, z) = \frac{\delta \exp(-\delta y)}{\int_0^z \delta \exp(-\delta y')\,\mathrm{d}y'} = \frac{\delta \exp(-\delta y)}{1 - \exp(-\delta z)} \quad \text{for } 0 \le y \le z.$$

In particular, it does not depend on the θ. For the null hypothesis case of $\delta = 0$, show that $Y \,|\, z$ is uniform on $[0, z]$.

(b). The natural conditional 0.05 level test is then to first compute z, and then to reject if $y \le 0.05\, z$. Show that it indeed has level 0.05, and that is the power optimal test among all conditional tests, using Y given z. Verify that this conditional test is the same as the unconditional test of rejecting when $R = Y/(X + Y) \le 0.05$. Compute the power function of the $T^* = I(Y \le 0.05\, Z)$ test (in the testing function parlance of Ex. 4.7), conditional on z, and unconditionally.

(c). At the boundary of the null, where $\delta = 0$, show that Z is complete. Show hence that *any* test with level 0.05 also must be a conditional on Z 0.05 level test, via Ex. 4.27.

(d). We know that $T^*(y, z) = I(y \le 0.05\, z)$ is the most powerful conditional test with level 0.05; we now wish to extend this statement to T^* actually being the most powerful among *all* tests with level 0.05. For any competing test $T(Y, Z)$ with level 0.05, show, since it must be a z-conditional 0.05 level test, where it cannot beat T^*, that

$$\mathrm{E}_{a,\delta}\{T^*(Y, Z) \,|\, z\} \ge \mathrm{E}_{a,\delta}\{T(Y, Z) \,|\, z\} \quad \text{for all } \delta > 0, z > 0.$$

There is equality, to 0.05, at $\delta = 0$. Show from this that T^* is more powerful than such T, unconditionally; in suitable power function symbols, $\pi_{T^*}(a, \delta) \ge \pi_T(a, \delta)$ for all $\delta > 0$.

(e). Suppose now that there are m independent pairs, $X_i \sim \mathrm{Expo}(a_i)$ and $Y_i \sim \mathrm{Expo}(a_i + \delta)$, with sums $Z_i = X_i + Y_i$; there are hence $m+1$ parameters with $2m$ data points. Show that the optimal test is to reject when $U_m = Y_1 + \cdots + Y_m$ is small, given $z_1, \ldots, z_m$. Explain how the null distribution of U_m can be evaluated via simulations. For an illustration, suppose three pairs (x_i, y_i) are observed: $(0.927, 0.819)$, $(1.479, 0.408)$, $(3.780, 1.311)$. Carry out the test of δ, and compute the p-value.

Ex. 4.28. *Conditional tests: normal.* (xx various situations with distribution of $U \,|\, (V = v)$, followed by natural conditional test. xx) Consider a pair of normals, where interest lies in assessing their difference in means. This may of course be parametrised in different ways, but one natural way is $x \sim \mathrm{N}(\theta, 1)$ and $y \sim \mathrm{N}(\theta + \delta, 1)$. One wishes to test $\delta = 0$ vs. $\delta > 0$, equivalent, of course, to testing equality of the means vs. $\mathrm{E}\, y > \mathrm{E}\, x$.

(a). Show that the joint likelihood can be written

$$\begin{aligned} f(x, y, \theta, \delta) &= (2\pi)^{-1} \exp[-\tfrac{1}{2}\{(x-\theta)^2 + (y-\theta-\delta)^2\}] \\ &= (2\pi)^{-1} \exp\{\theta z + \delta y - \tfrac{1}{2}x^2 - \tfrac{1}{2}y^2 - \tfrac{1}{2}\theta^2 - \tfrac{1}{2}(\theta+\delta)^2\}, \end{aligned}$$

where $z = x + y$, and with the main interaction between parameters and data being in the $\theta z + \delta y$ part.

(b). Show that (y, z) is a binormal, and set up its mean vector and variance matrix. Then use Ex. 1.44, or other algebraic methods, to show that $y \,|\, z \sim \mathrm{N}(\frac{1}{2}(z+\delta), \frac{1}{2})$; in particular, its conditional distribution does not depend on θ.

(c). Through the conditioning on z the testing problem has been reduced from a two-parameter to a one-parameter situation. For $y \,|\, z \sim \mathrm{N}(\frac{1}{2}(z+\delta), \frac{1}{2})$, show that the optimal test is to reject when $y - \frac{1}{2}z > (1/\sqrt{2})c$, with $c = \Phi^{-1}(1-\alpha)$ the standard normal quantile.

(d). Show that the above test, constructed to be optimal in the model for $y \,|\, z$, is equivalent to that of rejecting whem $D = y - x > \sqrt{2}c$. (xx so the conditional test is an ordinary unconditional test in disguise, or vice versa, in this particular situation. the point is the general principle. xx)

(e). (xx Consider m pairs of normal data, of the form $x_i \sim \mathrm{N}(\theta_i, 1)$ and $y_i \sim \mathrm{N}(\theta_i + \delta, 1)$. do the math, with the steps above. log joint density $\sum_{i=1}^m (\theta_i z_i + \delta y_i)$, with $z_i = x_i + y_i$. conditional test, $\sum_{i=1}^k y_i$ big given $z_1, \ldots, z_k$. xx)

(f). (xx something re power. xx)

Ex. 4.29. *Conditional tests: Poisson.* (xx various situations with distribution of $U \,|\, (V = v)$, followed by natural conditional test. xx)

(a). We start with a single pair of Poissons, x with mean θ, y with mean $\theta\gamma$. Show that the joint distribution becomes $\exp(-\theta - \theta\gamma)\theta^{x+y}\gamma^y/(x!\, y!)$. This inspires inspecting the distribution of y given $z = x + y$. Show that $y \,|\, z \sim \mathrm{binom}(z, \gamma/(\gamma+1))$.

(b). To test $\gamma = 1$ against $\gamma > 1$, describe in details the natural conditional test which rejects when y is big, given $z = x + y$.

(c). Next consider independent Poisson pairs x_i, y_i for $i = 1, \ldots, m$, where x_i has mean θ_i and y_i mean $\theta_i\gamma$. The model hence has $m+1$ parameters for the $2m$ observations, with γ the common multiplicative factor. Show that the joint distribution may be written

$$f = \exp\Big[-\sum_{i=1}^m (\theta_i + \theta_i\gamma) + \sum_{i=1}^m \{(x_i + y_i)\log\theta_i + y_i \log\gamma\}\Big] \frac{1}{x_1!\, y_1! \cdots x_m!\, y_m!}.$$

With $z_i = x_i + y_i$, find the distribution of $y_i \,|\, z_i$, and also the distribution of $S = \sum_{i=1}^m y_i$ given $z_1, \ldots, z_m$.

(d). Find the power optimal test for $\gamma = 1$ against $\gamma > 1$, among all those based on S given $z_1, \ldots, z_m$.

(e). (xx more, rounding off. something with limit. point to Ch7 optimal CD. also do $Y \sim \text{Pois}(m_0\theta_0)$ and $Y_1 \sim \text{Pois}(m_1\theta_1)$, with m_0 and m_1 exposure time. With interest being in the ratio parameter $\gamma = \theta_1/\theta_0$, show that $Y_1 \,|\, (Z = z)$ is binomial $(z, m_1\gamma/(m_0 + m_1\gamma))$. xx)

Ex. 4.30. *Conditional tests:* 2×2 *tables.* (xx various situations with distribution of $U \,|\, (V = v)$, followed by natural conditional test. xx)

(a). Consider two binomials $y_0 \sim \text{binom}(m_0, p_0)$ and $y_1 \sim \text{binom}(m_1, p_1)$. The outcomes in such situations are often presented as a two-by-two table,

$$\begin{array}{ll} y_0, & m_0 - y_0 \\ y_1, & m_1 - y_1 \end{array}$$

Consider the so-called logistic parametrisation

$$p_0 = H(\theta) = \frac{\exp(\theta)}{1 + \exp(\theta)} \quad \text{and} \quad p_1 = H(\theta + \gamma) = \frac{\exp(\theta + \gamma)}{1 + \exp(\theta + \gamma)}.$$

Show that $\theta = \log\{p_0/(1 - p_0)\}$ and $\theta + \gamma = \log\{p_1/(1 - p_1)\}$ in terms of the so-called log-odds. Show that the joint distribution can be written

$$f = \binom{m_0}{y_0}\binom{m_1}{y_1} \frac{\exp(\theta(y_0 + y_1))}{\{1 + \exp(\theta)\}^{m_0}} \frac{\exp(\gamma y_1)}{\{1 + \exp(\theta + \gamma)\}^{m_1}}.$$

(b). (xx in view of ... check, calibrate. xx) This inspires reaching inference for γ via the conditional distribution of y_1 given $z = y_0 + y_1$. Show that this distribution becomes

$$g_\gamma(y_1 \,|\, z) = \binom{m_0}{z - y_1}\binom{m_1}{y_1} \exp(\gamma y_1) \Big/ \sum_{y_1' \le \min(m_1, z)} \binom{m_0}{z - y_1'}\binom{m_1}{y_1'} \exp(\gamma y_1')$$

for $y_1 = 0, 1, \ldots, \min(m_1, z)$. In particular, this so-called excentric hypergeometric distribution depends on γ but not θ. We recognise the ordinary hypergeometric for $\gamma = 0$; see Ex. 1.67.

(c). Show that the optimal 0.05 level conditional test for the null hypothesis of equality, $p_0 = p_1$, is to reject when $y_1 > c(z)$, with $c(z)$ the highest number with $\sum_{0 \le y_1 \le c(z)} g_0(y_1 \,|\, z) \le 0.95$.

(d). (xx the power. xx)

(e). (xx to k two-by-two tables, $p_{i,0} = H(\theta_i)$ and $p_{i,1} = H(\theta_i + \gamma)$, $k + 1$ parameters. optimal conditional test for $\sum_{i=1}^k y_{1,i}$ given $z_1, \ldots, z_k$, with $z_i = y_{i,0} + y_{i,1}$. point to Story i.11. xx)

Ex. 4.31. *The t test as an optimal conditional test.* Let $Y_1, \ldots, Y_n$ be i.i.d. from the normal (ξ, σ^2), where we wish to test $\xi = 0$ against $\xi > 0$. The canonical classical test, of level say 0.05, is based on $t = \sqrt{n}\bar{Y}/\sigma$, rejecting if $t \ge t_{n-1,0.95}$, the upper 0.05 point of the t_{n-1} distribution; see also Ex. 3.7 and (3.4). We cannot use the Neyman–Pearson lemma directly to demonstrate optimality of the t test, however. One of several optimality properties may be derived via conditioning.

(a). Write $U = \sqrt{n}\bar{Y}$ and $V = \sum_{i=1}^n Y_i^2$, so that in particular $W = \sum_{i=1}^n (Y_i - \bar{Y})^2 = V - U^2$. Show that the joint density of the data can be written

$$\begin{aligned} f &= \frac{1}{(2\pi)^{n/2}} \frac{1}{\sigma^n} \exp\Big\{-\tfrac{1}{2}\frac{1}{\sigma^2}\sum_{i=1}^n (y_i - \xi)^2\Big\} \\ &= \frac{1}{(2\pi)^{n/2}} \frac{1}{\sigma^n} \exp\Big\{\frac{\xi}{\sigma^2}\sqrt{n}U - \tfrac{1}{2}\frac{1}{\sigma^2}V - \tfrac{1}{2}\frac{\xi^2}{\sigma^2}\Big\}. \end{aligned}$$

Note that the testing problem is equivalent to testing $\lambda = 0$ against $\lambda > 0$, with $\lambda = \xi/\sigma^2$, the mathematics indicating that this is a parameter easier to work with than ξ.

(b). Find the distribution of $U \,|\, (V = v)$, and show in particular that it depends on the parameters only via $\lambda = \xi/\sigma^2$. It is convenient here to work with

$$T = \frac{U}{\{W/(n-1)\}^{1/2}} = (n-1)^{1/2}\frac{U}{(V - U^2)^{1/2}}.$$

(c). Show that the power optimal test, among all tests based on $U \,|\, (V = v)$, is to reject when U is big, say $U \geq c(v)$, with $\Pr_0\{U \geq c(v) \,|\, V = v\} = 0.05$. Then show that this is actually the same as the t test.

(d). Show via arguments used in Ex. 4.27 and in other places above that the t test is not merely optimal among all conditional tests, for given level α, i.e. given V, but among *all* tests with level α.

Ex. 4.32. *Conditional tests: multiparameter exponential models.* In the rather simple situation of Ex. 4.27, with the exponential pair $X \sim \text{Expo}(a)$ and $Y \sim \text{Expo}(a + \delta)$, with sum $Z = X + Y$, we learned that (i) there is a clear level α conditional test for $\delta = 0$ vs. $\delta > 0$, in terms of $Y \,|\, Z$; (ii) that test is uniformly most powerful against all $\delta > 0$, among all conditional tests; and (iii) all other level α competitors are in fact also Z-conditional. Hence the winning test, reject if $Y \leq \alpha Z$, is the uniformly most powerful level α test. – We shall see now that the same arguments essentially go through for the wide class of all exponential families. Consider data Y from a density of the form $f(y, a, b) = \exp\{aU(y) + b^{\mathrm{t}}V(y) - k(a, b)\}h(y)$, as in Ex. 4.20, with one-dimensional U and p-dimensional V. Suppose we need to test $a = a_0$ against $a > a_0$, for some given null hypothesis value a_0.

(a). We have seen in the exercise pointed to that $U \,|\, (V = v)$ has a density depending on a but not b, and that it has an exponential form. Assume for simplicity that the distribution of U is continuous; mild formalistic additional arguments are required if the distribution is discrete. Deduce that there is a most powerful conditional level α test, say $T^*(y) = I\{U > c(V)\}$, with $c(v)$ determined from $\Pr_{a_0}(U > c(v) \,|\, V = v) = \alpha$, and with consequent power function $\pi(a, b) = \Pr_{a,b}\{U > c(V)\}$.

(b). Then consider *any* competing test $T(U, V)$ with level α. From $\mathrm{E}_{a_0,b}\, T(U(Y), V(Y)) = \alpha$, for all b, use completeness in b of $V(Y)$ for fixed a_0 to prove that

$$\mathrm{E}_{a_0,b}\{T(U(Y), V(Y)) \,|\, V(y) = v)\} = \alpha$$

for all v (except perhaps in a region of probability zero), and for all b. Thus the T competitor is also a level α V-conditional test, and we have proved that the conditional test is uniformly most powerful among *all* level α tests.

(c). The theory extends fruitfully to the case of testing $H_0\colon a \le a_0$ against $a > a_0$. Show that the test $T^* = I\{U > c(V)\}$ above, with $c(v)$ determined from $\Pr_{a_0}(U > c(v) \,|\, V = v) = \alpha$ at the boundary, is still of level α. Then show that this test is uniformly most powerful against all competing tests with constant level α at the boundary $a = a_0$; one says that such tests are *unbiased*. This latter very mild limitation is in order for the completeness argument to go through.

(d). (xx briefly about two-sided tests. still based on $U \,|\, (V = v)$. xx)

(e). When $U \,|\, (V = v)$ has a discrete distribution, the arguments still go through, but one cannot expect to find $c(v)$ with e.g. $\Pr_{a_0}\{U > c(v) \,|\, V = v\} = 0.05$. There are two ways out of this mild quandary. The first is to be satisfied with level 0.042, say, if that is how close one comes to 0.050, by appropriate choice of $c(v)$. The other, if one pedantically insists on 0.05, is to finetune $c(v)$ such that $\Pr_{a_0}\{U > c(v) \,|\, V = v\}$ is just below 0.05, and then identify the probability r such that

$$\Pr_{a_0}\{U > c(v) \,|\, V = v\} + r \Pr_{a_0}\{U = c(v) \,|\, V = v\} = 0.05,$$

So one rejects if $U > c(V)$, or, but then with probability r, if $U = v(C)$.

(f). (xx go through the previous exercises about conditional tests, once more, make sure that the tests found there are really uniformly most powerful among all unbiased tests. xx)

Ex. 4.33. *Optimal t testing in linear regression.* Consider the classic linear regression setup with $Y_i \sim \mathrm{N}(a + bx_i, \sigma^2)$ for $i = 1, \ldots, n$, see Ex. 3.33, and for simplicity of presentation here we assume the x_i have been centred, so $\bar{x} = 0$. The three estimators are then $\widehat{a} = \bar{Y}$, $\widehat{b} = \sum_{i=1}^n x_i Y_i / M_n$, with $M_n = \sum_{i=1}^n x_i^2$, and $\widehat{\sigma}^2 = Q_0/(n-2)$, with $Q_0 = \sum_{i=1}^n (Y_i - \widehat{a} - \widehat{b}x_i)^2$.

(a). With $V = \sum_{i=1}^n Y_i^2$, show that $Q_0 = V - n\widehat{a}^2 - M_n\widehat{b}^2$, and that the likelihood may be written

$$\begin{aligned} L &= \frac{1}{\sigma^n} \exp\Big\{-\tfrac{1}{2}\frac{1}{\sigma^2}\sum_{i=1}^n (Y_i - a - bx_i)^2\Big\} \\ &= \frac{1}{\sigma^n} \exp\Big\{-\tfrac{1}{2}\frac{V}{\sigma^2} + \frac{a}{\sigma^2} n\widehat{a} + \frac{b}{\sigma^2} M_n \widehat{b} - \tfrac{1}{2}\frac{na^2 + M_n b^2}{\sigma^2}\Big\}, \end{aligned}$$

which is in the exponential family form, with natural parameters $1/\sigma^2$, a/σ^2, b/σ^2.

(b). Explain from the general exponential family testing theory that the uniformly most powerful test for $b = 0$ against $b > 0$, among all level α tests is to reject when $\widehat{b}$ is large enough, in its conditional distribution given $(V, \widehat{a})$. We need to find $c(V, \widehat{a})$ such that

$$\Pr_{a,\sigma}(\widehat{b} > c(V, \widehat{a}) \,|\, V, \widehat{a}) = \alpha$$

under $b = 0$. But this may be transformed to

$$W = \frac{\widehat{b}}{(V - n\widehat{a}^2)^{1/2}} = \frac{\widehat{b}}{(Q_0 + M_n\widehat{b}^2)^{1/2}} = \frac{\widehat{b}/Q_0^{1/2}}{(1 + M_n\widehat{b}^2/Q_0)^{1/2}} > \frac{c(V, \widehat{a})}{(V - n\widehat{a}^2)^{1/2}}.$$

But W is a smooth increasing function of $t = \widehat{b}/(\widehat{\sigma}/M_n)$, the classic t ratio, which has the t_{n-2} distribution under $b = 0$. Argue that W is therefore independent of $(V, \widehat{a})$, and that the optimal test is to reject when $t > t_0$, the upper α in the t_{n-2} distribtion. You have now shown that the traditional t test is optimal.

(c). Generalise the above to general linear multiple regression models, giving the result that the traditional t tests, for each of the β_j coefficients, are optimal.

Ex. 4.34. *A generalised Poisson distribution.* For a count variable Y, consider the model with point probabilities

$$f(y,\lambda,\gamma) = k(\lambda,\gamma)^{-1}\lambda^y/(y!)^\gamma \quad \text{for } y = 0,1,2,\ldots,$$

where $k(\lambda,\gamma)$ is the normalisation constant $\sum_{y=0}^{\infty}\lambda^y/(y!)^\gamma$. For $\gamma = 1$ we're back to ordinary Pois(λ), with $k(\lambda,1) = \exp(\lambda)$. This two-parameter generalised Poisson model is from Schweder and Hjort (2016, Examples 4.18, 8.16). A regression version model of this type is used in Story iv.7, to assess potential overdispersion in Poisson counts.

(a). Pick some λ, and compute and display curves of the mean $\xi(\lambda,\gamma)$ and the variance-to-mean ratio $\rho(\lambda,\gamma)$, for an interval of γ around 1. Show that this ratio is decreasing in γ; hence $\gamma < 1$ indicates overdispersion and $\gamma > 1$ underdispersion, relative to the Poisson. Also show that the mean of $\log(Y!)$ is decreasing in γ.

(b). Show that the distribution is of the exponential family form, and that the sufficient statistics, after having observed a sample $Y_1,\ldots,Y_n$, is $T = \sum_{i=1}^n Y_i$ and $U = \sum_{i=1}^n \log(Y_i!)$. Show also that the joint distribution of these two must take the form

$$g_n(t,u) = \exp\{t\log\lambda - u\gamma - r_n(\lambda,\gamma)\}h_n(y_1,\ldots,y_n),$$

for appropriate functions r_n and h_n.

(c). For an observed sample $Y_1,\ldots,Y_n$, to test the Poisson assumption, against overdispersion, show that the optimal test is to reject when U is sufficiently small, given $T = t$. In other words, with level the classic 0.05, for example, we reject when $U \le u_0(t)$, where $u_0(t)$ is the 0.95 quantile of the distribution of U given $T = t$, computed at $\gamma = 1$, i.e. under Poisson conditions. (xx this needs more care; distribution of $U \,|\, (T = t)$ needs a formula or two, so we see that U significantly small indicates $\gamma < 1$. xx)

(d). There is no table or simple formula for the distribution of $U \,|\, (T = t)$, but show that it depends on γ, but not λ. Show that under $\gamma = 1$, $(Y_1,\ldots,Y_n) \,|\, (T = t)$ is a multinomial with count t and probabilities $(1/n,\ldots,1/n)$. Explain then how the distribution of $U \,|\, (T = t)$ may be simulated under Poisson conditions.

(e). (xx give them a dataset. nils checks the football matches dataset. decide later if this is a Story or an exercise. xx)

Ex. 4.35. *Testing for correlation.* Consider binormal i.i.d. pairs (X_i,Y_i), with parameters $\xi_1,\xi_2,\sigma_1,\sigma_2,\rho$, as in Ex. 1.56. Suppose we need to test $\rho = 0$ vs. $\rho > 0$.

(a). Argue as in the exercise pointed to that the binormal distribution is inside the exponential family class, with natural parameters

$$\frac{\xi_1}{(1-\rho^2)\sigma_1^2}, \frac{1}{(1-\rho^2)\sigma_1^2}, \frac{\xi_2}{(1-\rho^2)\sigma_2^2}, \frac{1}{(1-\rho^2)\sigma_2^2}, \frac{\rho}{(1-\rho^2)\sigma_1\sigma_2},$$

associated with data functions $X_i, X_i^2, Y_i, Y_i^2, X_iY_i$. Explain that the testing problem is equivalent to testing $\lambda = 0$ vs. $\lambda > 0$, where $\lambda = \rho/\{(1-\rho^2)\sigma_1\sigma_2\}$.

(b). Then argue from the general testing theory for the exponential class that there for given testing level α is a uniformly most powerful test, consisting in rejecting the null provided $A_n = n^{-1}\sum_{i=1}^n X_iY_i$ is big enough, in its conditional null distribution given $(\bar X, n^{-1}\sum_{i=1}^n X_i^2, \bar Y, n^{-1}\sum_{i=1}^n Y_i^2)$. Formally, the rejection threshold c, which depends on these variables, is determined by

$$\Pr_0\{A_n > c \,|\, (\bar X, n^{-1}\sum_{i=1}^n X_i^2, \bar Y, n^{-1}\sum_{i=1}^n Y_i^2)\} = \alpha,$$

with the footscript 0 indicating probability under $\rho = 0$.

(c). Write $\widehat\sigma_1^2 = n^{-1}\sum_{i=1}^n X_i^2 - \bar X^2$ and $\widehat\sigma_2^2 = n^{-1}\sum_{i=1}^n Y_i^2 - \bar Y^2$ for the empirical variances. Explain that in the conditional situation, given values for $\bar X, \bar Y, \widehat\sigma_1, \widehat\sigma_2$, the requirement above can be transformed to $\Pr_0\{R_n > d \,|\, (\bar X, \bar Y, \widehat\sigma_1, \widehat\sigma_2)\} = \alpha$, with R_n the empirical correlation coefficient, see (2.11).

(d). Then explain that R_n is actually independent of $\bar X, \bar Y, \widehat\sigma_1, \widehat\sigma_2$, under the null, so that the optimal test becomes the simpler one, of rejecting if $R_n > d$, where $\Pr_0(R_n > d) = \alpha$.

(e). For this optimal test regime to be specified fully, it remains to find the null distribution of R_n. Explain that this null distribution does not depend on the actual values of $(\xi_1, \xi_2, \sigma_1, \sigma_2)$, which we therefore may take to be $(0, 0, 1, 1)$. For practical purposes one may simulate a high number of R_n and then read off the d quantile; explain also that $\sqrt n R_n \to_d \mathrm N(0,1)$, from Ex. 2.54, so that $z_{1-\alpha}/\sqrt n$ is an approximation to the upper α point. We show however in the following point that $T_n = m^{1/2}R_n/(1-R_n^2)^{1/2}$ has a t_m distribution, under the null, with $m = n-2$ degrees of freedom. With e.g. $n = 25$, find the 0.95 quantile of the R_n distribution.

(f). So let us find the null distribution for R_n. Write $s_x^2 = \sum_{i=1}^n (X_i - \bar X)^2$, $s_y^2 = \sum_{i=1}^n (Y_i - \bar Y)^2$, $s_{xy} = \sum_{i=1}^n (X_i - \bar X)(Y_i - \bar Y)$, so that $R_n = s_{xy}/(s_x s_y)$. We may now use results from linear regression of the Y_i with respect to the X_i. Study in particular the least squares estimators $\widehat a, \widehat b$, minimising $Q(a,b) = \sum_{i=1}^n \{Y_i - a - b(X_i - \bar X)\}^2$, as with Ex. 3.33. Using results from that exercise, show (i) that $\widehat a = \bar Y$ and $\widehat b = s_{xy}/s_x^2$; (ii) that $Q_0 = Q(\widehat a, \widehat b) = s_y^2 - s_x^2\widehat b^2$; and (iii) that $T = \widehat b/\{(Q_0/m)^{1/2}/s_x\} = m^{1/2}s_x\widehat b/Q_0^{1/2}$ has the t_m distribution, with $m = n-2$ degrees of freedom. Show that this implies

$$R_n = \widehat b\frac{s_x}{s_y} = \frac{s_x\widehat b}{(Q_0 + s_x^2\widehat b^2)^{1/2}} = \frac{T/m^{1/2}}{(1 + T^2/m)^{1/2}}.$$

Solve for T to get the $T = m^{1/2}R_n/(1-R_n^2)^{1/2}$, pointed to and used in the previous point. Go on to work out the density of R_n, via the density g_m for the t_m; show that it becomes

$$h_n(r) = g_m\Big(\frac{m^{1/2}r}{(1-r^2)^{1/2}}\Big)\frac{m^{1/2}}{(1-r^2)^{3/2}} = \frac{\Gamma(\frac12(m+1))}{\Gamma(\frac12 m)\sqrt\pi}(1-r^2)^{(m-2)/2}$$

for $r \in (-1,1)$. As a curiosum, check that it is U-shaped for $n = 3$, uniform for $n = 4$, and then bell-shaped for $n \geq 5$.

Ex. 4.36. *Inference for linear multiple regression.* [xx to be done. point back to Ex. 3.34, 3.35, 3.37, and point to one or two stories. xx] in this exercise we show typical

and not so typical inference methods for the linear multiple regression model, using the key results reached in the previous exercise. confidence intervals, tests, also for σ, for a p quantile $F^{-1}(q\,|\,x_0) = x_0^{\rm t}\beta + z_q\sigma$, and delta method for things like $\Pr(Y \le y_0\,|\,x_0)$. and prediction. xx]

(a). (xx typical things first. show that $\widehat\beta_j \sim {\rm N}(\beta_j, k_j^2\sigma^2/n)$, where $k_j^2 = \sigma_n^{j,j}$ the diagonal elements of Σ_n^{-1}. from this show $t_j = (\widehat\beta_j - \beta_j)/(k_j\widehat\sigma/\sqrt{n})$ is a t_{n-p}. then ci for each β_j. and test of $\beta_j = 0$. also ok for $\beta_j - \beta_k$ etc. xx)

(b). (xx inference for σ. xx)

(c). For a given individual, with covariate vector x_0, the outcome Y_0 has the distribution ${\rm N}(x_0^{\rm t}\beta, \sigma^2)$. Consider the inference task for a quantile in this distribution. Show that the q-quantile becomes $\gamma_q = x_0^{\rm t}\beta + z_q\sigma$, with $\Phi(z_q) = q$. With estimator $\widehat\gamma_q = x_0^{\rm t}\widehat\beta + z_q\widehat\sigma$, show that

$$W_q = \frac{\widehat\gamma_q - \gamma_q}{\widehat\sigma} = \frac{x_0^{\rm t}(\widehat\beta - \beta) + z_q(\widehat\sigma - \sigma)}{\widehat\sigma} = \frac{{\rm N}(0, x_0^{\rm t}\Sigma_n^{-1}x_0/n) + z_q\{(\chi^2_m/m)^{1/2} - 1\}}{(\chi^2_m/m)^{1/2}}.$$

(xx round off. the point is that W_q can be simulated. also approximated with a normal. give data example. xx)

(d). (xx prediction, what will Y_0 be, at position x_0. also $\Pr(Y \le y_0\,|\,x_0)$. xx)

Ex. 4.37. *How much of the variance is explained?* (xx the below to be polished and illustrated, and with a clearer link to R^2. exact cc(ρ) to come in Ch. 7. point to illustration in Story **??**. xx) In the linear regression model, the extent to which the covariates influence the outcomes may be assessed in several ways, one of which is to decompose the variance of the outcomes into a covariate part and a 'remaining variability' part. Such assessments relate also to 'signal plus noise' viewpoints; how strong is the signal?

(a). We start out writing the regression model as

$$Y_i = \beta_0 + \beta_1 x_{i,1} + \cdots + \beta_p x_{i,p} + \varepsilon_i = \beta_0 + x_i^{\rm t}\beta + \varepsilon_i \quad \text{for } i = 1, \ldots, n,$$

again with the ε_i seen as i.i.d. ${\rm N}(0, \sigma^2)$, and further assume that the covariates have been centred, having their means subtracted, so that $\sum_{i=1}^n x_{i,j} = 0$ for $j = 1, \ldots, p$. This gives β_0 the interpretation as the overall mean of the Y_i. With $\Sigma_n = (1/n)\sum_{i=1}^n x_i x_i^{\rm t}$ the empirical $p \times p$ variance matrix for the x_i, show that the least squares estimators become

$$\widehat\beta_0 = \bar Y \sim {\rm N}(\beta_0, \sigma^2/n), \quad \widehat\beta = \Sigma_n^{-1}(1/n)\sum_{i=1}^n x_i Y_i \sim {\rm N}_p(\beta, (\sigma^2/n)\Sigma_n^{-1}),$$

and that these two are independent.

(b). Write $\widehat\mu_i = \widehat\beta_0 + x_i^{\rm t}\widehat\beta$ for the model based estimate of the outcome at x_i. For the sum of squared residuals, show that $Q_0 = \sum_{i=1}^n (Y_i - \widehat\mu_i)^2 = \sum_{i=1}^n (Y_i - \bar Y)^2 - n\widehat\beta^{\rm t}\Sigma_n\widehat\beta$. In other words,

$$V_n = \sum_{i=1}^n (Y_i - \bar Y)^2 = \sum_{i=1}^n (Y_i - \widehat\mu_i)^2 + n\widehat\beta^{\rm t}\Sigma_n\widehat\beta,$$

a neat decomposition of the full variability of outcomes as a sum of squared residuals and the covariate part $n\widehat\beta^{\rm t}\Sigma_n\widehat\beta$.

(c). (xx place a little caveat below, regarding interpretation; we need to think about the population being sampled from. xx) Standard themes for the linear regression model are developed with analyses carried out conditional on the covariates. Allow now a change in this narrative, where the x_i are seen as having their own covariate distribution, with mean zero and variance matrix Σ_n. Show that a randomly selected outcome Y_i then has variance $\beta^{\rm t}\Sigma_n\beta+\sigma^2$. Show that the covariate part of the full variability becomes

$$\rho=\frac{\beta^{\rm t}\Sigma_n\beta}{\beta^{\rm t}\Sigma_n\beta+\sigma^2}=\frac{\lambda}{\lambda+1},\quad\text{with }\lambda=\beta^{\rm t}\Sigma_n\beta/\sigma^2.$$

With $\widetilde{\sigma}^2=Q_0/n$ (rather than the unbiased $\widehat{\sigma}^2=Q_0/(n-p-1)$), show that this leads to

$$\widetilde{\rho}=\frac{\widehat{\beta}^{\rm t}\Sigma_n\widehat{\beta}}{\widehat{\beta}^{\rm t}\Sigma_n\widehat{\beta}+\widetilde{\sigma}^2}=\frac{V_n-\sum_{i=1}^n(Y_i-\widehat{\mu}_i)^2}{V_n}=1-\frac{\sum_{i=1}^n(Y_i-\widehat{\mu}_i)^2}{\sum_{i=1}^n(Y_i-\bar{Y})^2}.$$

This is often called the coefficient of determination, or R^2.

(d). To carry out precise inference for ρ, show first that $n\widehat{\beta}^{\rm t}\Sigma_n\widehat{\beta}/\sigma^2\sim\chi^2_p(n\lambda)$; check with Ex. 1.53. Then show that with $\widehat{\lambda}=\widehat{\beta}^{\rm t}\Sigma_n\widehat{\beta}/\widehat{\sigma}^2$, we have

$$F=n\widehat{\lambda}/p=n\widehat{\beta}^{\rm t}\Sigma_n\widehat{\beta}/(p\widehat{\sigma}^2)\sim F(p,m,n\lambda),$$

the noncentral F, see Ex. 1.72, with $m=n-(p+1)$ the degrees of freedom for $\widehat{\sigma}^2$.

(e). Explain how this may be used to set confidence intervals for λ and hence for ρ. You may check with Ex. 7.18 how this may be used to construct full confidence distributions for the fraction ρ of variation explained by the covariates; for an illustration, see Story **??**.

Ex. 4.38. *Testing and confidence for linear combinations of regression parameters.* (xx nils rant, so far; will be cleaned and related to other themes. xx) linear regression, $\widehat{\beta}\sim{\rm N}_p(\beta,(\sigma^2/n)\Sigma_n^{-1})$.

(a). A useful little lemma is as follows. Let x be a vector of length p, with A a symmetric positive definite $p\times p$ matrix. Show that the maximum of $c^{\rm t}x$, under the condition $c^{\rm t}Ac=1$, is $x^{\rm t}A^{-1}x$. You may show this via Lagrange multipliers and a bit of linear algebra.

(b). In the regression problem, show that we for each c have $c^{\rm t}\widehat{\beta}\sim{\rm N}(c^{\rm t}\beta,(\sigma^2/n)c^{\rm t}\Sigma_n^{-1}c)$, and that

$$V_n(c)=\frac{\sqrt{n}(c^{\rm t}\widehat{\beta}-c^{\rm t}\beta)}{\widehat{\sigma}\,c^{\rm t}\Sigma_n^{-1}c}\sim t_m,$$

with $m=n-p$ degrees of freedom. Explain how this can be used to test $\gamma(c)=c^{\rm t}\beta=0$, for any given c, and construct a confidence interval for such a $\gamma(c)$.

(c). Will the hypothesis $\gamma(c)=0$ be rejected, for any linear combination? Show that

$$V_n^*=\max_{\rm all\ }{}_c\,V_n(c)=(\sqrt{n}/\widehat{\sigma})\{(\widehat{\beta}-\beta)^{\rm t}\Sigma_n(\widehat{\beta}-\beta)\}^{1/2},$$

and that $(V_n^*)^2$ has a Fisher distribution with degrees of freedom $(p,n-p)$. (xx then round off. joint confidence for all linear combinations. testing $\beta=0$ via an F test. can also do contrasts, all c with $\sum_{j=1}^p c_j=0$. xx)

Ex. 4.39. *Inference for ratios of standard deviations.* Suppose two independent samples, of sizes n_1 and n_2, come from two populations, with standard deviations σ_1 and σ_2. From the empirical standard deviations $\widehat{\sigma}_1$ and $\widehat{\sigma}_2$, form the ratio $R = \widehat{\sigma}_1/\widehat{\sigma}_2$, to be used for inference about the underlying ratio $\rho = \sigma_1/\sigma_2$.

(a). Suppose first that the two distributions are normal. We then saw in Ex. 1.54 that $R^2 = \rho^2 F$, where $F \sim F_{m_1,m_2}$, an F distribution with degrees of freedom $(m_1, m_2) = (n_1 - 1, n_2 - 1)$. Construct a 95 percent confidence interval for ρ based on this. Also give a 0.05 level test for the equality hypothesis $\sigma_1 = \sigma_2$. this gives c.i. for ρ, and tests for $\rho = 1$. (xx answer: with $\Pr(a \le F \le b) = 0.95$, with (a, b) found from quantiles of the F, we find $[R/b^{1/2}, R/a^{1/2}]$. test: accept if $a < R^2 < b$. xx)

(b). Then consider inference for the ratio ρ outside the assumption of normally distributed data. From Ex. 3.12, find the representation

$$R = \frac{\widehat{\sigma}_1}{\widehat{\sigma}_2} \doteq \frac{\sigma_1}{\sigma_2} \frac{1 + (1/n_1)^{1/2}(\frac{1}{2} + \frac{1}{4}\gamma_{4,1})^{1/2} N_{n_1}}{1 + (1/n_2)^{1/2}(\frac{1}{2} + \frac{1}{4}\gamma_{4,2})^{1/2} N_{n_2}},$$

in terms of the kurtoses $\gamma_{4,1}$ and $\gamma_{4,2}$, where N_{n_1} and N_{n_2} are independent variables tending to standard normals as sample sizes increase. Use the delta method to deduce that $R/\rho \approx_d \mathrm{N}(1, \tau^2)$, with $\tau^2 = (1/n_1)(\frac{1}{2} + \frac{1}{4}\gamma_{4,1}) + (1/n_2)(\frac{1}{2} + \frac{1}{4}\gamma_{4,2})$. For normal data, show that this matches the distributional approximation $F^{1/2} \approx_d \mathrm{N}(1, 1/(2n_1) + 1/(2n_2))$.

(c). Construct an approximate 95 percent confidence interval for σ_1/σ_2, valid also outside normal data. For an application, see the Bach and Reger Story ii.12.

(d). (xx one more point. xx)

Ex. 4.40. *Testing equality across groups.* (xx mention below that $k = 2$ is simple, with pairwise comparisons, t type things. xx) Suppose a parameter has the same interpretation across groups, say θ_j for groups $j = 1, \ldots, k$. associated with estimators, perhaps of different precision. How can we test the null hypothesis of no difference, i.e. $H_0 \colon \theta_1 = \cdots = \theta_k$?

(a). Suppose first that $Y_j \sim \mathrm{N}(\theta_j, \sigma^2)$, with independence, and the same known precision, i.e. σ. Explain from classic results of Ex. 1.48–1.49, that $Z = \sum_{j=1}^k (Y_j - \bar{Y})^2/\sigma^2 \sim \chi^2_{k-1}$, with $\bar{Y} = (1/k)\sum_{j=1}^k Y_j$ the natural estimator for the common parameter under H_0. Explain that in the case of σ not being known, we still have $Z^* = \sum_{j=1}^k (Y_j - \bar{Y})^2/\widehat{\sigma}^2 \to_d \chi^2_{k-1}$, as long as $\widehat{\sigma}$ is consistent, with growing information for each group.

(b). A useful generalisation of the basic results for i.i.d. normal data is as follows. Suppose $Y_j \sim \mathrm{N}(0, 1/a_j)$, with the a_j positive and $a = \sum_{j=1}^k a_j$. Define $Y^* = \sum_{j=1}^k (a_j/a) Y_j$ and then $Z = \sum_{j=1}^k a_j (Y_j - Y^*)^2$. Show that $Z = \sum_{j=1}^k a_j Y_j^2 - a(Y^*)^2$, that it is independent of Y^*, and that its distribution is a χ^2_{k-1}. The classical case, leading to the distribution of the empirical variance and then to the t test, corresponds to all $a_j = 1$.

(c). A natural variation of the previous point is as follows. Assume in general terms that $Y_j \sim \mathrm{N}(\theta, \sigma_j^2)$, with independence, for $j = 1, \ldots, k$. Show that the minimiser of $Q(\theta) = \sum_{j=1}^k (Y_j - \theta)^2/\sigma_j^2$ is $\widehat{\theta} = (\sum_{j=1}^k Y_j/\sigma_j^2)/(\sum_{j=1}^k 1/\sigma_j^2)$, and that $Q_{\min} = Q(\widehat{\theta})$ is χ^2_{k-1}, independent of $\widehat{\theta}$. Demonstrate also that of all unbiased estimators $\theta^* =$

$\sum_{j=1}^k c_j Y_j$ of θ, $\widehat\theta$, with weights proportionao to inverse variances, has the smallest variance. Again, if the σ_j are consistently estimated, as opposed to known, explain that with $\theta^* = (\sum_{j=1}^k Y_j/\widehat\sigma_j^2)/(\sum_{j=1}^k 1/\widehat\sigma_j^2)$, we still have $Q^* = \sum_{j=1}^k (Y_j-\theta^*)^2/\widehat\sigma_j^2$ tending to χ^2_{k-1}.

(d). Consider correlations $\rho_1,\ldots,\rho_k$ for k groups of binormal data, with sample sizes $n_1,\ldots,n_k$, summing to the total n. Via Ex. 2.54, explain that $n_j^{1/2}(\widehat\zeta_j-\zeta_j)\to_d \mathrm{N}(0,1)$, where $\widehat\zeta_j = \frac12\log\{(1+\widehat\rho_j)/(1-\widehat\rho_j)\}$ is Fisher's zeta transform. Deduce that $Q = \sum_{j=1}^k n_j(\widehat\zeta_j-\zeta^*)^2 \to_d \chi^2_{k-1}$, with $\zeta^* = \sum_{j=1}^k (n_j/n)\widehat\zeta_j$. This is accordingly a natural test statistic for $\rho_1=\cdots=\rho_k$.

(e). Situations similar to that of the previous point abound in applied statistics, when needing to compare more groups than two. Suppose $\widehat p_j = Y_j/n_j$ is estimating a binomial p_j, for independent experiments, for $j=1,\ldots,k$. Then form $\widehat p = \sum_{j=1}^k (n_j/n)\widehat p_j$, and show that under $p_1=\cdots=p_k$ conditions, $Q=\sum_{j=1}^k n_j(\widehat p_j-\widehat p)^2 \to_d p(1-p)\chi^2_{k-1}$, using p to denote the common value of the p_j. Construct a test for equality of the p_j.

(f). (xx one more. could test equality of medians in k groups. xx)

Ex. 4.41. *Testing equality across groups, vector parameter case.* In Ex. 4.40 we developed a recipe for testing equality across groups for any one-dimensional parameter. Here we lift those methods to the case where the $\theta_1,\ldots,\theta_k$ in question have dimension $p\geq 2$. We wish again to test equality of these. (xx to be used in Story iii.7. xx)

(a). Assume $Y_j\sim \mathrm{N}_p(\theta_j,\Sigma_j)$ for $j=1,\ldots,k$, with known positive definite variance matrices. Under H_0 of $\theta_1=\cdots=\theta_k$, show that $Q(\theta)=\sum_{j=1}^k (Y_j-\theta)^{\mathrm t}\Sigma_j^{-1}(Y_j-\theta)$ is a χ^2_{kp}, and that it is minimised by $\widehat\theta = A^{-1}\sum_{j=1}^k \Sigma_j^{-1}Y_j$, with $A=\sum_{j=1}^k\Sigma_j^{-1}$. Show further (i) that $\widehat\theta\sim \mathrm{N}_p(\theta,A^{-1})$; (ii) that the full ensemble of $Y_j-\theta$ is independent of $\widehat\theta$; and (iii) that the natural test statistic $Q_0=Q(\widehat\theta)=\sum_{j=1}^k (Y_j-\widehat\theta)^{\mathrm t}\Sigma_j^{-1}(Y_j-\widehat\theta)$ is a $\chi^2_{(k-1)p}$.

(b). This Q_0 may be used as a test statistic for $\theta_1=\cdots=\theta_k$, with an exact chi-square null distribution, provided the Σ_j matrices are known. In various applications these are estimated from data, with $\widehat\Sigma_j$ using data for group j. This leads to $\widehat\theta = \widehat A^{-1}\sum_{j=1}^k \widehat\Sigma_j^{-1}Y_j$, with $\widehat A=\sum_{j=1}^k\widehat\Sigma_j^{-1}$. Show that if these $\widehat\Sigma_j$ are consistent, then $Q^*=Q^*(\widehat\theta)=\sum_{j=1}^k (Y_j-\widehat\theta)^{\mathrm t}\widehat\Sigma_j^{-1}(Y_j-\widehat\theta)$ tends to the $\chi^2_{(k-1)p}$, under the null.

(c). Suppose we have normal datasets with sample sizes n_j, yielding the usual parameter estimates $(\widehat\xi_j,\widehat\sigma_j)$, for groups $j=1,\ldots,k$, From these compute $\widehat\xi=\sum_{j=1}^k (n_j/n)\widehat\xi_j$, the overall mean, and $\widehat\sigma=\prod_{j=1}^k \widehat\sigma_j^{nj/n}$. Show that

$$Q=\sum_{j=1}^k \{n_j(\widehat\xi_j-\widehat\xi)^2+2n_j(\log\widehat\sigma_j-\log\widehat\sigma)^2\}$$

tends to $\chi^2_{2(k-1)}$ under the hypothesis of common (ξ,σ). Write also down the Wilks test, based on attainde log-likelihood maxima. (xx do both, for the mothers and babies story, for the three ethnic groups. xx)

(d). Consider a probability distribution $b=(b_1,\ldots,b_s)$ for suitable outcomes $1,\ldots,s$, and suppose we have multinomial estimators $\widehat b_j=(\widehat b_{j,1},\ldots,\widehat b_{j,s})$ for each of groups

$j = 1, \ldots, r$, resulting from sample sizes $n_1, \ldots, n_k$, with total sum n. In other words, $\widehat{b}_j = (N_{j,1}, \ldots, N_{j,s})/n_j$, for the multinomial experiment for group j. The hypothesis to be tested is that of equal probability distribution across groups. Explain that $\widehat{b}_j \approx_d \mathrm{N}_s(b, \Sigma/n_j)$, where Σ_j has $b_j(1-b_j)$ on the diagonal and $-b_j b_\ell$ outside. As in Story vii.1, write Σ_0 for the $(s-1) \times (s-1)$ submatrix of Σ, needed since Σ does not have full rank, with $\widehat{b}_0$ and $\widehat{b}_{j,0}$ similarly being the shorter $(s-1)$-length vectors, so that the general method above can be applied. Then show that the test statistic Q_0 from the general recipe above can be written

$$\sum_{j=1}^k n_j(\widehat{b}_{j,0} - \widehat{b}_0)^{\mathrm{t}} \widehat{\Sigma}_0^{-1} (\widehat{b}_{j,0} - \widehat{b}_0) = \sum_{j=1}^k n_j \sum_{r=1}^s \frac{(\widehat{b}_{j,r} - \widehat{b}_r)^2}{\widehat{b}_r} = \sum_{j=1}^k \sum_{r=1}^s \frac{(N_{j,r} - E_{j,r})^2}{E_{j,r}},$$

which then tends to $\chi^2_{(k-1)(s-1)}$ under the null. Show that this is actually precisely identical to the Pearson statistic from Story iii.7, for testing independence between the two factors 'group' and 'bin'.

Ex. 4.42. *One-way layout.* (xx a point about $k = 2$ at the end, where F i suppose becomes simply t^2, with usual t test. xx) The most prominent special case of the setup from Ex. 4.40 is that with an i.i.d. normal sample for each of the groups in question, with the same variance for all data. Consider therefore $Y_{j,1}, \ldots, Y_{j,n_j}$ i.i.d. from $\mathrm{N}(\xi_j, \sigma^2)$, for group j, and with overall sample size $n = \sum_{j=1}^k n_j$. This is the so-called normal one-way layout, with the chief hypothesis to be tested is $H_0 \colon \xi_1 = \cdots = \xi_k$.

(a). Let $\bar{Y}_j = (1/n_j) \sum_{r=1}^{n_j} Y_{j,r}$ be the group averages. Explain that we for the overall average have $\bar{Y} = (1/n) \sum_{j=1}^k \sum_{r=1}^{n_j} Y_{j,r} = \sum_{j=1}^k (n_j/n) \bar{Y}_j$. Show then that $Q = \sum_{j=1}^k n_j(\bar{Y}_j - \bar{Y})^2$ has $\mathrm{E}\, Q = (k-1)\sigma^2 + \sum_{j=1}^k n_j(\xi_j - \bar{\xi})^2$, with $\bar{\xi} = \sum_{j=1}^k (n_j/n)\xi_j$. Explain that $Q \sim \sigma^2 \chi^2_{k-1}$, under equality of means, so that the test statistic $W = Q/\widehat{\sigma}^2$ approximately has the χ^2_{k-1} distribution under H_0, as long as $\widehat{\sigma}$ is consistent. This is already a satisfactory answer to the one-way layout testing problem, without further fine-tuning.

(b). In this setup further fine-tuning is available, however. Show that $Q_0 = \sum_{j=1}^k \sum_{r=1}^{n_j} (Y_{j,r} - \bar{Y}_j)^2 \sim \sigma^2 \chi^2_{n-k}$, making $\widehat{\sigma}^2 = Q_0/(n-k)$ the natural unbiased estimator of σ^2. Explain also that Q_0 is independent of Q. Under the null of equal means, then, deduce that

$$F = \frac{Q/(k-1)}{Q_0/(n-k)} = \frac{Q/(k-1)}{\widehat{\sigma}^2} = \frac{\sum_{j=1}^k n_j(\bar{Y}_j - \bar{Y})^2/(k-1)}{\sum_{j=1}^k \sum_{r=1}^{n_j} (Y_{j,r} - \bar{Y}_j)^2/(n-k)}$$

has the F distribution $F_{k-1,n-k}$. Checking whether an observed F is too big, compared to this null distribution, is then a way to assess the hypothesis $\xi_1 = \cdots = \xi_k$. For an illustration, see Story **??**.

(c). (xx a short thing estimating contrasts. xx)

Ex. 4.43. *Testing equality of multinormal means.* (xx a brief thing, used in Story ii.7, can point to ML theory too. establish the following, then apply in two-three settings. xx) Suppose $A \sim \mathrm{N}_p(a, \Sigma_1)$ and $B \sim \mathrm{N}_p(b, \Sigma_2)$ are independent multinormal data vectors, with known variance matrices. How can we test $a = b$? Show that $W = (B-A)^{\mathrm{t}}(\Sigma_1 + \Sigma_2)^{-1}(A-B) \sim \chi^2_p$ under the null.

Ex. 4.44. *The Behrens–Fisher problem.* (xx to come here, making it reasonable simple. one point is that problems go away with reasonable sample sizes. xx)

Further material and leftovers (as of 12-March-2025)

Ex. 4.45. *Non-monotone likelihood ratio.* Consider a simple setup with $Y \sim \mathrm{N}(\theta, \theta^2/m)$, for a known m. (xx point to this being approximate situation in several setups. xx) Suppose one needs to test $\theta = \theta_0$, a known value, against the alternative $\theta > \theta_0$.

(a). For some $\theta_1 > \theta_0$, show that the log-likelihood-ratio can be written

$$R(y, \theta_0, \theta_1) = \log f(y, \theta_1) - \log f(y, \theta_0) = \tfrac{1}{2} max^2 - mby - \log(\theta_1/\theta_0),$$

with $a = 1/\theta_0^2 - 1/\theta_1^2$ and $b = 1/\theta_0 - 1/\theta_1$. Simplify this further to

$$R(y, \theta_0, \theta_1) = \tfrac{1}{2} ma\Big(y - \tfrac{1}{1/\theta_0 + 1/\theta_1}\Big)^2 + \text{const.},$$

(b). Set up the Neyman–Pearson optimal test, at level say 0.05, at this alternative point θ_1. Find the associated optimal power, at θ_1. Compare this power with the simpler test which rejects if $m^{1/2}(y - \theta_0)/\theta_0 > 1.645$.

(c). (xx a few details. so there is no UMP, for $Y \sim \mathrm{N}(\theta, \theta^2/m)$, but there *is* a simple UMP for $Y \sim \mathrm{N}(\theta, \theta/m)$. xx)

(d). (xx a bit more, worth doing, in detail. find the ML estimator $\widehat{\theta}$. explore the Wilks test, $W_m = 2\{\ell(\widehat{\theta}) - \ell(\theta_0)\}$, perhaps calibrated to have null distribution closer to χ^2_1. find its power function. compare with NP optimal in a few positions. xx)

(e). xx

Ex. 4.46. *Two denominators for Wald tests.* (xx this needs work and editing, and might be split in two parts. i need a better prototype example for the 'two denominators' thing than what's here as of 4 april 2021. xx) (xx for $W_n = \widehat{\beta}/D_n$ there are often two choices for the denominator D_n, constructed to be consistent at H_0, or consistent in the wider model. same $\mathrm{N}(0,1)$ limit under H_0, but different powers. need to point to and calibrate wring to monotone likelihood ratio property. we need a good example where this property does not hold. we make a Cauchy exercise too. xx) Consider the simple setup where $\widehat{\theta} \sim \mathrm{N}(\theta, \theta/m)$, with m known, perhaps the sample size behind the estimator $\widehat{\theta}$. We also suppose θ is sufficiently positive so that $\widehat{\theta}$ is also positive with very high probability. One wishes to test the null hypothesis that $\theta = \theta_0$, a known null value, against the alternative that $\theta > \theta_0$.

(a). For a value $\theta_1 > \theta_0$, consider the log-ratio function

$$\begin{aligned} R(y, \theta_1, \theta_0) &= \log\{f(y, \theta_1)/f(y, \theta_0)\} \\ &= -\tfrac{1}{2} m(y - \theta_1)^2/\theta_1 + \tfrac{1}{2} m(y - \theta_0)^2/\theta_0 - \log\theta_1 + \log\theta_0. \end{aligned}$$

Show that R is increasing in x. (xx which implies that $W_{n,0}$ is optimal, on this occasion. need another example to showcase different power behaviour. xx)

(b). Consider the two test ratios

$$W_{m,0} = \frac{\widehat{\theta} - \theta_0}{(\theta_0/m)^{1/2}} \quad \text{and} \quad W_{m,1} = \frac{\widehat{\theta} - \theta_0}{(\widehat{\theta}/m)^{1/2}}.$$

Show that their null distributions both tend to the standard normal, as m grows; in fact $W_{m,0}$ is exactly a $\mathrm{N}(0,1)$ under the null. With a 0.05 level, the tests reject the null if $W_{m,0} \geq z_0$ or $W_{m,1} \geq z_0$, with $z_0 = 1.645$. (xx can ask just a bit more here, a more exact version of the second test, needing then distribution of $(\widehat{\theta}/\theta_0)^{1/2}$. xx)

(c). First consider the local power, at the local alternatives $\theta = \theta_0 + a/m^{1/2}$. Show that both $W_{m,0}$ and $W_{m,1}$ tend to $\mathrm{N}(a/\theta_0, 1)$, and hence that both power functions tend to $\Phi(a/\theta_0 - z_0)$.

(d). Then consider an alternative value, with a non-zero $\delta = \theta - \theta_0$. Show

$$W_{m,0} = \frac{m^{1/2}\{\delta + (\theta/m)^{1/2} Z_m\}}{\theta_0^{1/2}} = (\theta/\theta_0)^{1/2} Z_m + m^{1/2}\delta/\theta_0^{1/2},$$
$$W_{m,1} = \frac{m^{1/2}\{\delta + (\theta/m)^{1/2} Z_m\}}{\widehat{\theta}^{1/2}} = (\theta/\widehat{\theta})^{1/2} Z_m + m^{1/2}\delta/\widehat{\theta}^{1/2}.$$

where $Z_m \sim \mathrm{N}(0,1)$. Argue that this leads to power functions

$$\pi_{m,0}(\theta) = \Pr\{(\theta/\theta_0)^{1/2} N + m^{1/2}(\theta - \theta_0)/\theta_0^{1/2} \geq z_0\},$$
$$\pi_{m,1}(\theta) = \Pr\{N + m^{1/2}(\theta - \theta_0)/\theta^{1/2} \geq z_0\},$$

with N a standard normal; the first power expression is exact, the second an approximation.

(e). xx

Notes and pointers

(xx confidence intervals. testing. connections. power. Neyman–Pearson. point to Lehmann. and to later chapters, Ch. 7 for CDs. point to interplay between modelling, probability calculus, thinking, a bit of philosophy, and practice. xx)

(xx one-way layout: primarily to test equality of means, so 'anova' is arguably a slight misnomer. xx)

Basu's lemma: from Basu (1955), see also Ghosh (2002).

ToDo notes, as of 12-March-2025:

Lots, though the chapter is shaping up. There are lots of 'test $\theta = \theta_0$ against $\theta > \theta_0$ prose in exercises, since it's easiest and cleanest, with NP etc. But we need to say a couple of times that all of this generalises to $\theta \neq \theta_0$ etc.

the Lindqvist and Taraldsen things, for simulating from $U(x) \,|\, \{V(x) = v\}$. do the handball model from CLP.

Do the bioequivalence: test the H_0 that θ is outside $[-\varepsilon, \varepsilon]$ against the alternative that it is inside. Diffent type of tests, and different looking power function.

Do sample size things, and efficiency; $\pi_n(\theta) \doteq \Phi(\sqrt{n}(\theta - \theta_0)/\sigma - z_0)$, with informative derivative $\phi(0)\sqrt{n}/\sigma$ at θ_0. Efficiency things.

Make an example or two, perhaps with Cauchy again, to see that the confidence region might not be an interval.

Point to Cox (1958) and also interviews with him regarding conditional stuff.

I.5

Minimum divergence and maximum likelihood

Consider the joint density of a dataset Y from a parametric model, say $f_{\text{full}}(y, \theta)$. The likelihood function, a fundamental concept in parametric inference, is just this density, but seen as a function of θ, with y fixed at the observed dataset y_{obs}. In this chapter we go through the fundamental likelihood inference methods, in particular with results for the maximum likelihood estimator, the maximiser of the likelihood, and for the attained maximum value itself. It is useful to see this general likelihood theory as a special case of the more flexible machinery of minimum divergence methods. Here statistical divergences, often seen as a distance from one fixed model to a collection of approximation models, lead to empirical divergence functions, then to be minimised to find the best approximation. In fact we develop general theory for minimum divergences first, below, after which likelihod theory is a relatively easy consequence, in that maximum likelihood corresponds to one particular divergence, namely the Kullback–Leibler. The divergence and likelihood methods are practical and versatile, as is demonstrated in several exercises and stories, also for say non-standard regression models. With models outside the most familiar ones, inference analysis essentially flows from being able to programme the log-likelihood function or divergence function. Further material connected to likelihood theory are Cramér–Rao information inequalities, Wilks theorems, influence functions, and certain flexible robustification methods. Crucially, the theory developed does not in general presuppose that the parametric model worked with is correct, as results are established both under and outside the precise model conditions. (xx perhaps there is room for empirical likelihood. xx)

Key words: BHHJ, Cramér–Rao lower bounds, Fisher information matrix, influence function, Kullback–Leibler, least false parameters, log-likelihood, maximum likelihood, minimum divergence, regression models, score function, Wilks tests

In earlier chapters we have met and worked with several methods for estimation parameters in different settings. Sometimes one estimates population parameters directly from data, like the mean, the median, the standard deviation, the skewness, the correlation, the median difference between groups, etc., without necessarily using parametric models. Very frequently, however, the most fruitful data analyses involve fitting some parametric model to the data, as in regression models, where regression coefficients are estimated and assessed to learn how covariates influence the main outcomes. In Chs. 3-4 we have already worked with methods associated with moment fitting, quantile fitting, and least squares for regression models, but the present chapter has a wider aim and applicability,

specifically developing methodology for *minimum divergence function estimators*, and its primary special case, *maximum likelihood estimators*.

In general terms, suppose data y have been modelled via a parametric model, leading to a joint probability density $f_{\rm full}(y,\theta)$, with θ in some relevant parameter region Θ. Then the *likelihood function* is $L(\theta) = f_{\rm full}(y_{\rm obs},\theta)$, studied as a function of the parameters, with y held fixed at the observed $y_{\rm obs}$. The maximum likelihood (ML) estimator is the value $\widehat\theta$ maximising the likelihood, or equivalently the *log-likelihood function* $\ell(\theta) = \log f_{\rm full}(y_{\rm obs},\theta)$. The simplest setup is that of i.i.d. observations $Y_1,\ldots,Y_n$ to be fitted to some parametric $f(y,\theta)$, with θ a parameter vector inside its relevant parameter region. The log-likelihood function is then

the ML estimator

log-likelihood function

$$\ell_n(\theta) = \sum_{i=1}^n \log f(Y_i,\theta). \tag{5.1}$$

The apparatus of minimum divergence function and likelihood estimation carries over to regression models too, where the density of Y_i given a covariate vector x_i is modelled via some $f(y_i \mid x_i,\theta)$.

Generally speaking there are two valid viewpoints when developing the required theory. The first takes *the model to be correct*, so there is a true parameter θ_0 to be estimated, with associated inference. The second takes the parametric model to be a sensible approximation to the real and unknown data-generating mechanism. With i.i.d. data, the model density $f(y,\theta)$ is consequently seen as an approximation to the real underlying density $g(y)$. Estimation and inference then involves the *least false parameter* θ_0, in a suitable sense making $f(y,\theta_0)$ coming as close to the real $g(y)$ as possible.

A generic and powerful statistical idea is to estimate parameters by minimising a relevant distance or divergence from the true mechanism to the data. Our chapter starts with general ways in which to construct relevant *distance functions*, often tied to divergence ideas about the distance from one density to another, which in various settings define the best parameter θ_0 to be minimiser of such functions. This is then followed up by constructing the parameter estimator $\widehat\theta$ as the minimiser of an empirical version of the same distance function. The arguably most important estimation method, the *maximum likelihood method*, is the special case associated with the *Kullback–Leibler divergence* $\mathrm{KL}(g,f_\theta)$. Apart from clarifying the concepts, defining parameters and estimators, in such ways, developing the theory amounts to establishing clear results for the behaviour and performance of these estimators and related data-driven tests, confidence methods, etc. This again involves limiting normality, assessing and estimating variability, and so on. All of this leads to a fruitful and indeed practical general theory, which also in new situations, with models constructed for new purposes, allows the statistician to estimate, to reach relevant confidence statements, test and compare, predict, etc.

In brief, this chapter has these connected parts, with associated correlated groups of exercises.

(i) Motivating and building an apparatus for defining parameters and their estimators via minimum divergence functions, initially for i.i.d. setups. The most important method, the maximum likelihood (ML) method, briefly described above, is a special case, associated with the KL divergence.

(ii) Analysing ML methods involves studying score functions, information functions, and the Fisher information matrix, for general parametric models. This again relates to 'information inequalities', specifically of the Cramér–Rao type, establishing lower bounds

for how small variances can be, with a given model and a given sample size n. The ML methods achieve these lower bounds, under model conditions, for growing n.

(iii) Methods are developed for deriving the basic properties, concerning limiting normality of estimators, chi-squared type results for profiled versions for focus parameters, and quadratic forms for classes of test criteria, via the mechanics of minimisers of random functions.

(iv) Importantly, the i.i.d. setup is then lifted to classes of regression models. Intrigruingly, with attention to certain crucial details, these extensions turn out to be not too strenuous, partly with Lindeberg theorem arguments replacing CLT details for the simpler case. The general likelihood theory in particular then makes it relatively straightforward to define and work with not only multiple linear regression, but logistic regression, Poisson regression, gamma type regressions, and more.

Carrying out data analyses using tools from this chapter is often even surprisingly straightforward, even in new situations with new models. This is partly due the general well-working theory, but also to the modern conveniences of software packages for numerical optimisation, easy calculation of derivatives and second derivatives, and so on, after having programmed the basic empirical distance function to be used. For maximum likelihood methodology, in particular, this is showcased in Story **??**, with a logistic full-data model for how covariates influences birthweights, Story vii.4, concerning models for time-to-failure of machine components, and in Story iv.7, with regression models for the number of bird species on islands outside Ecuador. These applications are also meant to inspire the invention of new parametric models in new situations, perhaps along with context relevant distance functions.

(xx can briefly point to extensive use of likelihood theory in Ch. 7, 10, 11, 12. to be clearer, also using Notes: M-estimators, Z-estimators. the general wondrous recipes around $\widehat\theta \approx_d \mathrm{N}_p(\theta_0, \widehat J_{\rm obs}^{-1})$. make robustness aspects clearer. xx)

Divergences, Kullback–Leibler, likelihoods

Ex. 5.1. *Maximising the log-likelihood.* (xx check that we don't repeat too similar things later on. xx) Here we work through some examples of setting up the log-likelihood function and finding the ML estimator.

(a)**.** Suppose $Y \sim \text{binom}(100, \theta)$ and that you observe $y_{\rm obs} = 22$. Set up the log-likelihood function $\ell(\theta)$ and plot it. Show that its maximiser is $\widehat\theta = 0.22$.

(b)**.** Simulate ten values of Y from the binomial $(100, 0.25)$, and plot the resulting ten log-likelihood functions. Note how they vary, giving different ML estimates of the underlying true $\theta_0 = 0.25$. With your code, experiment a bit with different sizes of the binomial n. For bigger n, the peak is sharper; show this mathematically.

(c)**.** With i.i.d. observations $Y_1, \ldots, Y_n$ from some parametric $f(y, \theta)$, show that the log-likelihood function becomes $\ell_n(\theta) = \sum_{i=1}^n \log f(Y_i, \theta)$. With such data points from the two-parameter normal $\mathrm{N}(\xi, \sigma^2)$ model, show that the log-likelihood function may be written

$$\begin{aligned} \ell_n(\xi, \sigma) &= \sum_{i=1}^n \{-\log\sigma - \tfrac12 (y_i - \xi)^2/\sigma^2 - \tfrac12 \log(2\pi)\} \\ &= -n\log\sigma - \tfrac12 \{Q_0 + n(\xi - \bar Y)^2\}/\sigma^2 - \tfrac12 n\log(2\pi), \end{aligned} \tag{5.2}$$

with $\bar{Y}$ the sample mean and $Q_0 = \sum_{i=1}^n (Y_i - \bar{Y})^2$. Show that the ML estimators are $\widehat{\xi} = \bar{Y}$ and $\widehat{\sigma}^2 = Q_0/n$. Verify also that the log-likelihood maximum value is $\ell_{n,\max} = -n \log \widehat{\sigma} - \frac{1}{2} n\{1 + \log(2\pi)\}$. Compared to the classical empirical variance $Q_0/(n-1)$, see Ex. 1.49, we learn that the ML estimator has a small negative bias, but that the difference is small for moderate to large n.

(d). With data $y_1, \ldots, y_{100}$ from the exponential density $\theta \exp(-\theta y)$, set up the log-likelihood function and find a formula for its maximiser $\widehat{\theta}$.

(e). Simulate 50 points from the uniform $[0, \theta]$, with true value $\theta_0 = 1$. Find and plot the log-likelihood function, and read off the ML estimate $\widehat{\theta}$.

(f). Simulate 100 values y_i from the Gam$(a_0, 1)$ distribution (see Ex. 1.9), with e.g. $a_0 = 3.33$. Set up and plot the log-likelihood function. Set up a little experiment where you keep track of the ML estimators $\widehat{a}$ in repeated experiments from the samme Gam$(a_0, 1)$.

Ex. 5.2. *The likelihood and log-likelihood functions.* Consider the one-parameter model with density $f(y, \theta) = \exp(-\theta y^{1/2})\theta/(2y^{1/2})$ for $y > 0$, and assume $n = 12$ data points have been observed:

```
0.233 0.334 0.067 0.148 0.007 0.639 0.017 0.298 0.030 0.120 0.061 0.063
```

(a). Show that $f(y, \theta)$ indeed is a density, write down the log-likelihood function $\ell_n(\theta)$, and show that it is maximised at $\widehat{\theta} = 1/W_n$, with $W_n = n^{-1} \sum_{i=1}^n y_i^{1/2}$. Find also a formula for the Hessian at the maximum position, $\widehat{J}_{\rm obs} = -\ell_n''(\widehat{\theta})$.

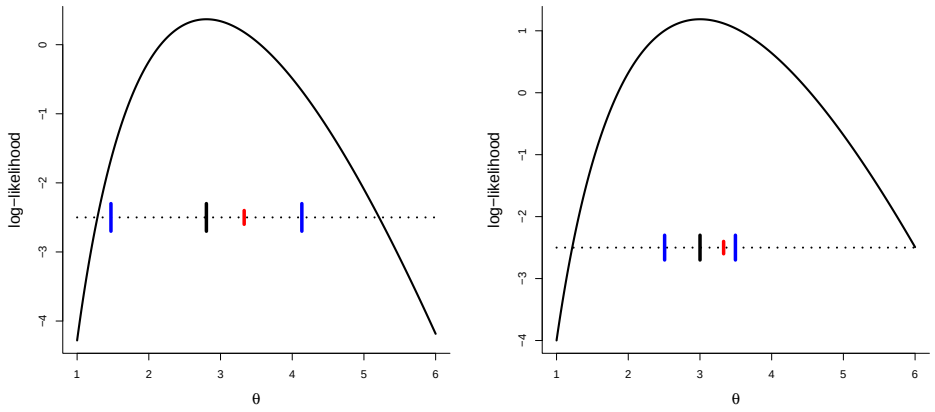


Figure 5.1: *Left panel: The log-likelihood function $\ell_n(\theta)$ for the simple $n = 12$ dataset of Ex. 5.2. The maximum is attained at $\widehat{\theta} = 2.803$. The blue bars indicate the 90 percent confidence interval coming from standard ML estimation theory, see Ex. 5.18, and is $[1.472, 4.134]$. The true value behind the data, $\theta_0 = 3.33$, is indicated as the red bar. Right panel: as for the left, but now with the bigger data set of $n = 100$ datapoints, from the same distribution, where the first 12 are as above. The minus second derivative $\widehat{J}$ at the ML position has increased from 1.527 to 11.103, causing the ML based confidence interval to become considerably tighter, and is $[2.507, 3.494]$.*

(b). Find the limit distribution of $\sqrt{n}(W_n - 1/\theta)$, using the CLT of Ch. 2, and use the delta method to find that the limit distribution of $\sqrt{n}(\widehat{\theta} - \theta)$ is $\mathrm{N}(0, \theta^2)$. You may verify already that this is what comes out of the general ML theory developed below, see e.g. Ex. 5.18.

(c). Anticipating general ML theory to come, check again with Ex. 5.18, it will be seen that $\widehat{\theta} \approx_d \mathrm{N}(\theta, 1/\widehat{J})$, under model conditions. Explain that this leads to approximate confidence intervals of the type $\widehat{\theta} \pm z_0/\widehat{J}^{1/2} = \widehat{\theta}(1 \pm z_0/\sqrt{n})$, with z_0 the appropriate normal quantile. We've actually simulated these few data points above from the model, with true parameter $\theta_0 = 3.33$. Construct a version of Figure 5.1, left panel.

(d). In general the ML estimator might have a complicated distribution (though it is approximately normal, as we have seen here). In this particular model its precise distribution may be worked out, however; show that $\widehat{\theta} \sim \theta\,(2n)/\chi^2_{2n}$. Use this to find a precise 90 confidence interval for θ, and compare to the approximation given above.

(e). To simulate data from this model, show that Y_i is equal in distribution to $(V_i/\theta_0)^2$, with the V_i i.i.d. from the unit exponential. Make a computer programme to simulate n points from the $f(y, \theta_0)$ model, and which then finds the log-likelihood function, the ML estimator, and the approximate 90 percent confidence interval, as above. Run such a programme for say 88 more points, forming a bigger dataset with $n = 100$ datapoints, and comment on what you find. Produce a version of Figure 5.1, right panel. Comment on the main features here, including that $\widehat{J}$ becomes bigger with more data, yielding sharper confidence intervals. We would usually plot the log-likelihood and related aspects, as the confidence curves of Ch. 7 for a shorter range of parameter values than in this right panel, but we here choose to plot using the same range for both $n = 12$ and $n = 100$.

Ex. 5.3. *Minimising L_2 distance.* Suppose i.i.d. data $Y_1, \ldots, Y_n$ come from some data generating density g which we wish to approximate with some parametrically modelled f_θ. The L_2 distance between g and f_θ is

$$D(g, f_\theta) = \int (g - f_\theta)^2 \,\mathrm{d}y = \int f_\theta^2 \,\mathrm{d}y - 2\int g f_\theta \,\mathrm{d}y + a(g),$$

where $a(g)$ does not depend on θ.

(a). Use this to motivate what we may call the minimum L_2 estimator $\widehat{\theta}$, the minimiser of $D_n(\theta) = \int f_\theta^2 \,\mathrm{d}y - 2n^{-1}\sum_{i=1}^n f(Y_i, \theta)$. Operationally, this is easiest when there is a closed-form formula for $\int f_\theta^2 \,\mathrm{d}y$, but numerical minimisation might be carried out even without this. Simulate 100 detapoints from a $\mathrm{Gam}(a, b)$, where you choose (a, b) as you wish, and estimate these using this method.

(b). Carry out a similar simple experiment with 100 datapoints drawn from a normal, i.e. estimate the mean and standard deviation using the minimum L_2 method. Then do the following variation: push one of the 100 datapoints far away from the others, which will cause the traditional estimates of μ and σ to be off; show however that the minimum L_2 estimates are far less affected. – The minimum L_2 method is a special case of the BHHJ method, see Ex. 5.9.

Ex. 5.4. *Estimating means via minimum divergence.* Suppose again i.i.d. data $Y_1, \ldots, Y_n$ stem from some data generating density g. The point to convey here is that key parameters can be seen as minimisers of natural distance functions, or divergences, leading to recipes for estimating them.

(a). For any distribution G, let $\xi(G)$ be the minimiser of $H(\xi) = \mathrm{E}_G\,(Y-\xi)^2$, with estimator the minimiser $\widehat{\xi}$ of the empirical version $H_n(\xi) = n^{-1}\sum_{i=1}^n (Y_i-\xi)^2$. Show, for this simple illustration of minimum divergence function methods, that $\xi(G)$ is the mean $\mathrm{E}_G\,Y$ and that $\widehat{\xi} = \bar{Y}$, the sample mean.

(b). For characterising and then estimating both mean and standard deviation, consider $h(y,\xi,\sigma) = (y-\xi)^2 + \{y^2 - (\xi^2+\sigma^2)\}^2$, and let (ξ,σ) be the minimiser of the distance function $H(\xi,\sigma) = \mathrm{E}_G\,h(Y,\xi,\sigma)$. Show that $\xi(G)$, $\sigma(G)$ are the mean and standard deviation, and work out expressions for $\widehat{\xi}$, $\widehat{\sigma}$, the minimisers of the empirical distance function $H_n(\xi,\sigma) = n^{-1}\sum_{i=1}^n h(Y_i,\xi,\sigma)$.

(c). For positive Y_i data, to be fitted to the $\mathrm{Gam}(a,b)$, consider the function $h(y,a,b) = (y-a/b)^2 + (y^2 - (a/b^2 + a^2/b^2))^2$, constructed in view of a/b and a/b^2 being the formulae for the mean and variance for a Gamma. Explain that the empirical version of the distance function $H(a,b) = \mathrm{E}_G\,h(Y,a,b)$ becomes $H_n(a,b) = n^{-1}\sum_{i=1}^n \{(Y_i-a/b)^2 + (Y_i^2 - (a/b^2+a^2/b^2)\}^2$. Show that the minimum divergence function estimators $(\widehat{a},\widehat{b})$ are equivalent to the moment estimators, worked with in Ex. 3.27, i.e. the solution to the two equations $\bar{Y} = a/b$ and $\widehat{\sigma}^2 = a/b^2$.

Ex. 5.5. *Minimum divergence function estimators: general setup.* We learn from Ex. 5.3–5.4 that classes of estimators may be formed via minimisation of suitable empirical distance functions; these estimate the corresponding minimisers of distance functions operating on the underlying distributions. We call these *minimum divergence function estimators* or *minimum distance estimators*. In general terms, for observations $Y_1,\ldots,Y_n$ from some distribution G, consider a parameter $\theta_0 = \theta(G)$ defined as the minimiser of the function $H(\theta) = \mathrm{E}_G\,h(Y,\theta)$, for a suitable $h(y,\theta)$. It is assumed that θ_0 thus defined, which may also be multidimensional, is the unique minimiser. The empirical version of $H(\theta)$ is $H_n(\theta) = n^{-1}\sum_{i=1}^n h(Y_i,\theta)$, so a natural estimator for θ_0 is $\widehat{\theta} = \mathrm{argmin}(H_n)$. In fact many important estimators are of this or related types, perhaps minimising somewhat more complicated random functions, as we shall see in this chapter. In exercises below we shall develop clear results for how the minimum divergence function estimators behave, under sets of natural assumptions, but the present exercise is meant to illustrate the basic construction via different types of examples.

minimum divergence function estimators

(a). Explain that $H_n(\theta)$ can be written $\int h(y,\theta)\,\mathrm{d}G_n(y)$, with G_n the empirical distribution, having mass $1/n$ at each datapoint; see Ex. 3.9. Explain why $H_n(\theta) \to_{\mathrm{pr}} H(\theta)$ for each θ, and find the limit distribution of $\sqrt{n}\{H_n(\theta) - H(\theta)\}$. What we need, tended to in several exercises to follow, are conditions under which $\widehat{\theta} = \mathrm{argmin}(H_n)$ tends to $\theta_0 = \mathrm{argmin}(H)$, along with a limit distribution.

(b). For one-dimensional Y_i, work through the details of $h(y,\theta) = (y-\theta)^2$. Then consider $h(y,\theta) = [\exp\{c(y-\theta)\} - 1 - c(y-\theta)]/c^2$, with c a balance parameter. Draw 100 datapoints from a normal $\mathrm{N}(\theta,1)$, with θ of your choice, and estimate θ in this minimum H_n fashion, for a few values of the balance parameter c. Show that c close to zero corresponds to the mean.

(c). Let generally $h(y,\theta) = \{r(y)-\theta\}^{\mathrm{t}} V \{r(y)-\theta\}$, for some $r(y) = (r_1(y),\ldots,r_p(y))$ and a symmetric positive definite matrix V. Show that $\theta_0 = \mathrm{E}_G\,r(Y)$ and that $\widehat{\theta} = n^{-1}\sum_{i=1}^n r(Y_i)$.

(d). Consider $h(y,\xi,\tau) = p(\tau) + \frac{1}{2}(y-\xi)^2/\tau^2$, where $p(\tau)$ is a smooth increasing function of $\tau > 0$. Find a recipe for computing the estimates $(\widehat{\xi},\widehat{\tau})$ associated with the distance function $n^{-1}\sum_{i=1}^n h(Y_i,\xi,\tau)$. Check in particular the case of $p(\tau) = \log\tau$.

(e). Consider $h_0(x) = x\arctan x - \frac{1}{2}\log(1+x^2)$, and define θ_0 as the minimiser of $\mathrm{E}_G\, h_0(Y-\theta)$. Show that $\widehat{\theta}$, the minimiser of $H_n(\theta) = n^{-1}\sum_{i=1}^n h_0(Y_i-\theta)$, is also the unique solution to $\sum_{i=1}^n \arctan(Y_i-\theta) = 0$. (xx so connection from minimum divergence estimator to M estimator. round off. xx)

(f). There are connections here to the moment matching estimation method, as worked with in Ex. 3.26. For a one-parameter model first, with $\mathrm{E}\,Y = m(\theta)$, consider $h(y,\theta) = \{y-m(\theta)\}^2$. Show that the implied best parameter is for $m(\theta_0) = \mathrm{E}\,Y$, or $\theta_0 = m^{-1}(\mathrm{E}\,Y)$, and that the minimum divergence function method leads to solving $m(\theta) = \bar{Y}$. Generalise to the case of there being two parameters in the model, and then to the general vector parameter case.

Ex. 5.6. *The Kullback–Leibler divergence and the maximum likelihood method.* With i.i.d. data $Y_1,\ldots,Y_n$ from a density g, to be approximated with a parametric f_θ, the particularly important *maximum likelihood estimation method* is worked with here, seen as the natural cousin to the Kullback–Leibler divergence as a measure of distance from the true density to the parametric approximation. Development in exercises below give a precise description of how the method actually behaves.

the Kullback–Leibler distance

(a). For two densities g and f, defined on a common support, the Kullback–Leibler distance, interpreted to be 'from the first density to the second', is

$$\mathrm{KL}(g,f) = \int g\log\frac{g}{f}\,\mathrm{d}y. \tag{5.3}$$

It is an important concept and tool for communication and information theory, as for probability theory and statistics. The $\log(g/f)$ term will be both positive and negative, in different parts of the domain. Show nevertheless that indeed $\mathrm{KL}(g,f) \geq 0$, perhaps via the Jensen inequality, and that $\mathrm{KL}(g,f) = 0$ only when the two densities are equal a.e. In Ex. 5.7 we learn more details about the KL distance, and look into illustrations, but here the main point is to see its close connection to ML estimation.

maximum likelihood

(b). We now apply the general minimum divergence machinery of Ex. 5.5, with basic function $h(y,\theta) = -\log f(y,\theta)$. Show that this defines $\theta_0 = \theta(G)$ as the minimiser of $\mathrm{KL}(g,f_\theta)$, and that minimising the implied distance function $H_n(\theta)$ is the same as maximising the function $\ell_n(\theta) = \sum_{i=1}^n \log f(Y_i,\theta)$, as in (5.1). This function is sufficiently famous and pervasive to have earned its own name, the *log-likelihood function* (here for i.i.d. observations). As explained above, the ML estimation method, seen here to be equivalent to minimum divergence with the underlying divergence being the Kullback–Leibler, from true density to parametric approximation.

(c). As a simple illustration (to be returned to in Ex. 5.9), generate say $n = 100$ points from the uniform distribution, and use the parametric density $f(y,\theta) = \theta y^{\theta-1}$ on the unit interval. Write down and plot the log-likelihood function, and find the ML estimate. Also show that the implied best parameter value, if the data stem from some g, rather than from the model, is $\theta_0 = 1/\mathrm{E}_g\,\log(1/Y)$.

Ex. 5.7. *The Kullback–Leibler distance: details and illustrations.* (xx repair here. xx) In Ex. 5.6 and other exercises above we have seen that the machinery of maximum likelihood is intimately related to the KL distance $\mathrm{KL}(g,f)=\int g\log(g/f)\,\mathrm{d}y$. Here we work on illustrations to learn more.

(a). For two normal densities, $\mathrm{N}(a,1)$ and $\mathrm{N}(b,1)$, show that the KL distance becomes $\frac12(b-a)^2$. Prove also the somewhat more general result, that with $g\sim\mathrm{N}(\xi_1,\sigma^2)$ and $f\sim\mathrm{N}(\xi_2,\sigma^2)$, the KL distance is $\frac12(\xi_2-\xi_1)^2/\sigma^2$.

(b). Find the KL distance from one Poisson to another.

(c). The KL distance is also perfectly well-defined and meaningful in higher dimension. Show that the KL distance from $\mathrm{N}_p(\xi_1,\Sigma)$ to $\mathrm{N}_p(\xi_2,\Sigma)$ can be expressed as $\frac12\delta^2$, where $\delta=\{(\xi_2-\xi_1)^{\mathrm{t}}\Sigma^{-1}(\xi_2-\xi_1)\}^{1/2}$ is the so-called Mahalanobis distance between the two populations.

the Mahalanobis distance

(d). For several of these examples we find KL distances being symmetric, between the two densities in question, but this is not true in general. Compute the KL distance from $\mathrm{N}(\xi,\sigma_1^2)$ to $\mathrm{N}(\xi,\sigma_2^2)$, and compare to the reciprocal case.

(e). (xx may consider a little reshuffle of exercises. xx) Consider a parametric density $f(y,\theta)$, with score function $u(y,\theta)=\partial\log f(y,\theta)$ and information matrix $J(\theta)=\mathrm{Var}_\theta\,u(Y,\theta)$; see Ex. 5.14 for more on these. Show here that

$$\mathrm{KL}(f(\cdot,\theta),f(\cdot,\theta+\varepsilon))=\tfrac12\varepsilon^{\mathrm{t}}J(\theta)\varepsilon+O(\varepsilon^3).$$

(f). Start from $d(g,f)=-\int g\log\{1+(f/g-1)\}\,\mathrm{d}y$, for densities which are not far from each other, and use Taylor expansion to find

$$\mathrm{KL}(g,f)\approx\tfrac12\int g(f/g-1)^2\,\mathrm{d}y=\tfrac12\Big(\int f^2/g\,\mathrm{d}y-1\Big).$$

(xx some words indicating that the root-KL might have an easier interpretation. xx)

(g) (xx a bit of text, more than a question. xx). As noted the KL distance is not symmetric, so 'distance' has a direction. In various statistical setups it makes sens to interpret $d(g,f)$ as the the distance from 'home density g' to 'approximation candidate f'. As also becoming clear from examples above, it's somehow quadratic in nature, so when numbers are involved, measuring the KL distances, it would typically make more sense to give their square roots, as with $\{d(g,f_\theta)\}^{1/2}$, the degree of closeness of the parametric approximant f_θ to the ground truth g.

Ex. 5.8. *KL approximation.* (xx to be edited. xx) For the following cases the point is to set up a data generating density g, and then check how well a certain parametric family $f(y,\theta)$ does the approximation job. For each case, this tells us how well the ML can do its job, with enough data. For the various cases, find the minimiser, i.e. the best approximation; find the minimum square-root distance $d(g,f(\cdot,\theta_0))^{1/2}$ (since this gives a better picture than on the KL scale itself); and plot the true g alongside the parametric approximant.

(a). Let $g=0.33\,\mathrm{N}(-1,1)+0.67\,\mathrm{N}(1,1)$. Find the best normal approximation.

(b). Let g be a Gamma with parameters $(2.22, 3.33)$. Find the best Weibull approximant, and also the best log-normal approximant. Similarly, start with a Weibull distribution, with parameters say $(3.33, 2.22)$, and find the best Gamma distribution approximation.

(c). Let $g = 0.95\,\text{Expo}(1) + 0.05\,\text{Expo}(0.01)$, which roughly means that about five percent of the data come from a distribution which much higher mean than the mainstream exponential data. Find the best exponential model approximation, and also the best Gamma and Weibull approximations. Display the true g and these three best parametric approximations in the same diagram.

(d). Suppose data really come from $\text{N}(0.333, \sigma_1^2)$, with $\sigma_1 = 1.111$, where a statistician fits the simpler $\text{N}(0, \sigma^2)$ model. First, find out what happens to the ML estimator. Secondly, illustrate 'what goes on' by drawing e.g. ten samples of size $n = 50$ from the true density, and then display the ten versions of $n^{-1}\ell_n(\sigma)$, along with its limit $C(\sigma)$. Comment on your findings.

BHHJ, the minimum divergence process, and its limit

Ex. 5.9. *The BHHJ density power divergence method.* Here we set up the basics for the so-called *density power divergence method*. It involves raising the density function $f(y, \theta)$ to some power a, as one of its ingredients. In the literature it is sometimes called the *BHHJ divergence method*, from its inventors Basu, Harris, Hjort, Jones (Basu et al. (1998), Jones et al. (2001)). We shall see in later exercises that it amounts to a robust modification of the ML method.

(a). For a density g, in what follows to be seen as the true underlying data-generating model, consider measuring the distance to an approximate $f_\theta(y) = f(y, \theta)$ density as

$$d_a(g, f_\theta) = \int \Big\{ f_\theta^{1+a} - \Big(1 + \frac{1}{a}\Big) g f_\theta^a + \frac{1}{a} g^{1+a} \Big\}\, \text{d}y, \tag{5.4}$$

with a a positive tuning parameter. Show that $d_a(g, f_\theta) \geq 0$, and that the distance is zero only when $g = f_\theta$ a.e. Note that this for $a = 1$ is the same as the L_2 distance of Ex. 5.3.

(b). Use Taylor expansion for f_θ^a and g^a for small a, to demonstrate that the integrand in (5.4) may be written

$$f_\theta - g + g \log(g/f_\theta) - a(g - f_\theta) \log f_\theta + \tfrac{1}{2} a g \{(\log g)^2 - (\log f_\theta)^2\} + O(a^2).$$

Hence show that for a small, we have $d_a(g, f_\theta) = \text{KL}(g, f_\theta) + O(a)$, with the Kullback–Leibler distance $\int g \log(g/f_\theta)\, \text{d}y$, assuming that the functions $g \log f_\theta$, $f_\theta \log f_\theta$, $g(\log f_\theta)^2$, $g(\log g)^2$ have finite integrals.

(c). Suppose $Y_1, \ldots, Y_n$ are i.i.d. from some unknown g, and that we wish to estimate θ by making the distance $d_a(g, f_\theta)$ small. The point is now that the third term of d_a does not depend on θ, and that we may accurately estimate the two first, using

$$H_{n,a}(\theta) = \int f(y, \theta)^{1+a}\, \text{d}y - (1 + 1/a)\, n^{-1} \sum_{i=1}^n f(y_i, \theta)^a. \tag{5.5}$$

Show that this is also coming out of the general minimum divergence function apparatus of Ex. 5.5, with $h(y, \theta) = \int f_\theta^{1+a}\, \text{d}y - (1 + 1/a) f(y, \theta)^a$. Show indeed that $H_n(\theta)$ is an unbiased estimator of the two first terms of $d_a(g, f_\theta)$, and give an expression for its variance. We call the minimiser $\widehat{\theta}$ of $H_n(\theta)$ the BHHJ estimator.

the BHHJ estimator

(d). Return to the ML estimator illustration of Ex. 5.6, with $n = 100$ points drawn from the uniform, fitted to the $\theta y^{\theta-1}$ density. For some values of a, compute and plot the $H_{n,a}$ function, finding also the BHHJ estimate $\widehat{\theta}_a$. After having computed this for a grid of a, plot the resulting $\widehat{\theta}_a$ as a function of a, and comment.

Ex. 5.10. *Integrals for BHHJ estimation.* Using the BHHJ method of Ex. 5.9 for estimating the parameters of models $f(y,\theta)$, we are very much helped, algorithmically and numerically, by having formulae for the term $A(a) = \int f_\theta^{1+a}\,\mathrm{d}y$.

(a). For the normal $\mathrm{N}(\xi,\sigma^2)$, show that

$$A(a) = \int f(y,\xi,\sigma)^{1+a}\,\mathrm{d}y = (2\pi)^{-a/2}(1+a)^{-1/2}\sigma^{-a}.$$

Generalise to the multinormal case of $\mathrm{N}_p(\xi,\Sigma)$, with

$$A(a) = \int f(y,\xi,\Sigma)^{1+a}\,\mathrm{d}y = (2\pi)^{-ap/2}(1+a)^{-p/2}|\Sigma|^{-a/2}.$$

(b). For the gamma model with density proportional to $y^{\alpha-1}\exp\{-\beta(y-y_0)\}$ for $y \ge y_0$, we find

$$\int_{y_0}^\infty f(y,\alpha,\beta)^{1+a}\,\mathrm{d}y = \frac{\Gamma(\alpha+a\alpha-a)}{\Gamma(\alpha)^{1+a}}\frac{\beta^a}{(1+a)^{\alpha+a\alpha-a}}.$$

(c). For the log-normal model, where Y is such that $\log Y \sim \mathrm{N}(\xi,\sigma^2)$: Show that

$$\int_0^\infty f(y,\xi,\sigma)^{1+a}\,\mathrm{d}y = (2\pi)^{-a/2}\sigma^{-a}(1+a)^{-1/2}\exp\{-a\xi + \tfrac{1}{2}a^2\sigma^2/(1+a)\}.$$

(d). For the Weibull, with c.d.f. $1-\exp\{-(y/\alpha)^\beta\}$ for $y \ge 0$, show that

$$\int_0^\infty f(y,\alpha,\beta)^{1+a}\,\mathrm{d}y = \Big(\frac{\beta}{\alpha}\Big)^a\frac{\Gamma(1+a-a/\beta)}{(1+a)^{1+a-a/\beta}}.$$

(e). For robust estimation of the three-parameter t distribution, consider the density $f(y,\xi,\sigma,\nu) = g_\nu((y-\xi)/\sigma)(1/\sigma)$ for $Y = \xi + \sigma t_\nu$, with g_ν the t density with ν degrees of freedom. Find

$$\begin{aligned}\int f^{1+a}\,\mathrm{d}y &= \frac{1}{\sigma^a}\int g_\nu^{1+a}\,\mathrm{d}x \\ &= \frac{1}{\sigma^a}\frac{\Gamma(((1+a)\nu+a)/2)}{\Gamma((1+a)(\nu+1)/2)}\Big\{\frac{\Gamma((\nu+1)/2)}{\Gamma(\nu/2)}\Big\}^{1+a}\frac{1}{(\nu\,\pi)^{a/2}}.\end{aligned}$$

Ex. 5.11. *Maximum weighted likelihood estimation.* Another fruitful generalisation of the ML method of Ex. 5.6, in addition to the BHHJ method of Ex. 5.9, starts from adding a weight function to the Kullback–Leibler $\mathrm{KL}(g,f)$ divergence.

(a). Show that $\mathrm{KL}(g,f)$ may be written $\int\{g\log(g/f) - (g-f)\}\,\mathrm{d}y$, that the integrand is nonnegative, and is zero only if $g = f$ a.e. With any nonnegative weight function $w(y)$ on the sample space, deduce that

$$\mathrm{KL}_w(g,f) = \int w(y)\Big[g(y)\{\log\frac{g(y)}{f(y)} - \{g(y)-f(y)\}\Big]\,\mathrm{d}y$$

is a divergence, i.e. nonnegative, and is zero only if $g = f$ on the support set of w. In particular, for g fixed and parametric $f(\cdot,\theta)$, show that $\mathrm{KL}_w(g,f_\theta)$ can be expressed as $\int w(-g\log f_\theta + f_\theta)$ plus other terms not depending on θ.

(b). For data $Y_1, \ldots, Y_n$ from a generating density g, to be approximated with a parametric $f(y, \theta)$, explain that the general minimum divergence function method leads to maximisation of

$$\ell_{n,w}(\theta) = \sum_{i=1}^{n} w(y_i) \log f(Y_i, \theta) - nB(\theta),$$

with $B(\theta) = \int w f_\theta \, \mathrm{d}y$. We call the maximiser $\widehat{\theta}_w$ the maximum weighted likelihood estimator, associated with weight function $w(y)$. The classic ML estimator then corresponds to constant weight function $w(y) = 1$. Properties of $\widehat{\theta}_w$ are derived in Ex. 5.21.

(c). For an illustration, access the birthweight dataset for Oslo children, reported on in Story i.5. For the boys and the girls, fit normal distributions, using maximum weighted likelihood, with weight function $w(y)$ being one on $[2.5, 4.5]$ kg and zero outside. Plot the estimated densities in a diagram, This requires programming a function for optimisation. Compare with curves estimated via full ML.

Ex. 5.12. *Minimum divergence function estimators: a limit process and main heuristics.* After having motivated and worked through particular instances of the minimum divergence function estimators, we now return to the general case, aiming to demonstrate both limiting normality and other associated results, finding recipes for large-sample approximations in the process. The aim of the present exercise is to go through the basic ideas, involving also two main heuristics. These then apply to large classes of estimators and associatead minimum divergence function minima. To make these heuristics precise we need more conditions and details, to which we return in further exercises below. In particular, results derived here and in a few of the following exercises will apply immediately to ML estimators, BHHJ estimators, MWL estimators, via their definitions in Ex. 5.6, 5.9, 5.11.

So $\theta_0 = \theta_0(G)$ is the minimiser of $H(\theta) = \mathrm{E}_G \, h(y, \theta)$, and $\widehat{\theta}$ is its estimator, the minimiser of $H_n(\theta) = \int h(y, \theta) \, \mathrm{d}G_n(y)$, with G_n the empirical distribution for an observed i.i.d. sample $Y_1, \ldots, Y_n$ from G. To present the basic ideas and main heuristics we shall start with these regularity conditions: (i) The true $\theta_0 = \theta_0(G)$ is an inner point in its parameter space inside $\mathbb{R}^p$. (ii) The $h(y, \theta)$ is smooth in θ, in a neighbourhood around θ_0, with at least two derivatives, say $h'(y, \theta)$ (a vector, with components h'_a), $h''(y, \theta)$ (a matrix, with components $h''_{a,b}$). (iii) The matrix $J = \mathrm{E}_G \, h''(Y, \theta_0)$ is finite and positive definite. (iv) The first derivative $h'(Y, \theta_0)$ has finite variance matrix K.

(a). A key idea is to work with the following random function. Write

$$\begin{aligned} A_n(s) &= n\{H_n(\theta_0 + s/\sqrt{n}) - H_n(\theta_0)\} \\ &= \sum_{i=1}^{n} \{h(Y_i, \theta_0 + s/\sqrt{n}) - h(Y_i, \theta_0)\} = U_n^{\mathrm{t}} s + \tfrac{1}{2} s^{\mathrm{t}} J_n s + r_n(s), \end{aligned}$$

with $U_n = (1/\sqrt{n}) \sum_{i=1}^{n} h'(Y_i, \theta_0)$, $J_n = n^{-1} \sum_{i=1}^{n} h''(Y_i, \theta_0)$. Show that $A_n(s)$ is properly defined for all large enough n, that U_n has mean zero and tends to $U \sim \mathrm{N}_p(0, K)$, and that $J_n \to_{\mathrm{pr}} J$.

(b). It is also clear from the Taylor expansion argument that with moderate further regularity, the remainder term $r_n(s)$ will go to zero in probability, for each s. There is then convergence $A_n(s) \to_d A(s) = U^{\mathrm{t}} s + \frac{1}{2} s^{\mathrm{t}} J s$, for each s. Heuristic One is then

to go from $A_n \to_d A$ to $\text{argmin}(A_n) \to_d \text{argmin}(A)$. Explain that this then leads to $\sqrt{n}(\widehat{\theta} - \theta_0) \to_d -J^{-1}U$, which is a $\text{N}_p(0, J^{-1}KJ^{-1})$. We call this limit distribution variance matrix $J^{-1}KJ^{-1}$ the sandwich matrix.

the sandwich matrix

(c). Similarly, Heuristic Two is to go from $A_n \to_d A$ to $\min(A_n) \to_d \min(A)$. Argue that this entails $W_n = 2n\{H_n(\theta_0) - H_n(\widehat{\theta})\} \to_d W = U^{\text{t}}J^{-1}U$. Show for this quadratic form that its mean and variance are $\text{Tr}(J^{-1}K)$ and $2\,\text{Tr}(J^{-1}KJ^{-1}K)$.

(d). (xx give an easy example, to identify J and K. xx)

Ex. 5.13. *Minimum divergence function estimators: from heuristics to proofs.* (xx nils and emil check details carefully. repair and edit. xx) The general setup is as with the previous Ex. 5.12, with minimum divergence function estimator $\widehat{\theta}$ minimising $H_n(\theta) = n^{-1}\sum_{i=1}^n h(Y_i, \theta)$, with observations coming from G. In addition to regulatity conditions (i)–(iv) given there, we postulate (v), that also third derivatives of $h(y, \theta)$ exist, say $h'''_{a,b,c}(y, \theta)$, and that these have finite means in a neighbourhood around θ_0. We use the process $A_n(s) = n\{H_n(\theta_0 + s/\sqrt{n}) - H_n(\theta_0)\}$ from the previous exercise.

(a). Explain first that the minimiser of A_n is $\alpha_n = \sqrt{n}(\widehat{\theta} - \theta_0)$, where we shall also study the overall minimum $A_{n,\min} = A_n(\alpha_n)$ below. Let $B_n(s) = U_n^{\text{t}}s + \frac{1}{2}s^{\text{t}}J_n s$ be the quadratic approximation to A_n, with minimiser β_n and overall minimum $B_{n,\min} = \min\{B_n(s)\colon \text{all } s\}$. Show that

$$\begin{aligned} \beta_n &= -J_n^{-1}U_n \to_d -J^{-1}U \sim \text{N}_p(0, J^{-1}KJ^{-1}), \\ B_{n,\min} &= -\tfrac{1}{2}U_n^{\text{t}}J_n^{-1}U_n \to_d -\tfrac{1}{2}U^{\text{t}}J^{-1}U. \end{aligned}$$

So things are simple and clear for the quadratic approximation B_n; we need to show that the same results obtain for the real thing, the A_n.

(b). Supposing $|r_n(s)| \le \delta$ for all s in a subset S, show from $A_n = B_n + r_n$ that

$$|\min_{s\in S} A_n(s) - \min_{s\in S} B_n(s)| \le \delta.$$

We next establish that α_n cannot be far away. Show that when $\|s\| \ge cn^{1/8}$, $B_n(s) \ge D_n n^{1/4}$, with D_n positive and bounded in probability; and that when $\|s\| \le cn^{1/8}$, then $|r_n(s)| \le E_n/n^{1/8}$, with E_n also bounded in probability. Show from this that $\alpha_n = O_{\text{pr}}(n^{1/8})$, and that $A_{n,\min} - B_{n,\min} \to_{\text{pr}} 0$.

(c). Then consider B_n a certain distance away from the minimum. For given small ε, show that for v with $\|v\| \ge \varepsilon$, we have

$$B_n(\beta_n + v) = B_{n,\min} + \tfrac{1}{2}v^{\text{t}}J_n v \ge B_{n,\min} + \tfrac{1}{2}j_n\varepsilon^2,$$

where j_n is the smallest eigenvalue of J_n. Show then that the event Ω_n, where $A_n(\beta_n + v) \ge B_{n,\min} + \frac{1}{4}j_n\varepsilon^2$ for all v with $\|v\| \ge \varepsilon$, must have $\Pr(\Omega_n) \to 1$. Prove from these established statements that $\alpha_n - \beta_n$ must tend to zero in probability. Conclude that

$$\begin{aligned} \sqrt{n}(\widehat{\theta} - \theta_0) &\to_d -J^{-1}U \sim \text{N}_p(0, J^{-1}KJ^{-1}), \\ 2n\{H_n(\theta_0) - H_n(\widehat{\theta})\} &\to_d W = U^{\text{t}}J^{-1}U. \end{aligned} \tag{5.6}$$

Score function, Fisher information matrix, Cramér–Rao lower bounds

Ex. 5.14. *Score functions and the Fisher information matrix.* Consider a parametric model with density $f(y,\theta)$ with respect to some measure μ, where $\theta = (\theta_1,\ldots,\theta_p)^{\mathrm{t}}$, the parameter of the model is contained in some open parameter space $\Theta \subset \mathbb{R}^p$. Introduce score function

$$u(y,\theta) = \partial \log f(y,\theta)/\partial\theta \quad \text{and} \quad i(y,\theta) = \partial^2 \log f(y,\theta)/\partial\theta\,\partial\theta^{\mathrm{t}},$$

called the *score function*, and is a p-dimensional column vector , and the *information function*, a $p\times p$ matrix. For the theory developed in this exercise, these partial derivatives are assumed to exist and, they must be continuous in θ. We also assume that the *support* for the distribution, the smallest closed set for which the density is positive, does not depend on θ. Cases falling outside such assumptions are, for example, the uniform distribution on an unknown interval $[0,\theta]$.

(a). Assume that $(\partial/\partial\theta)\int f(y,\theta)\,\mathrm{d}\mu(y) = \int (\partial/\partial\theta) f(y,\theta)\,\mathrm{d}\mu(y)$. Show that the score function has mean zero, first Bartlett identity

$$\mathrm{E}_\theta\, u(Y,\theta) = \int f(y,\theta) u(y,\theta)\,\mathrm{d}\mu(y) = 0.$$

Let next

$$K(\theta) = \mathrm{Var}_\theta\, u(Y,\theta) \quad \text{and} \quad J(\theta) = -\mathrm{E}_\theta\, i(Y,\theta),$$

Assume also that $(\partial^2/\partial\theta\partial\theta^{\mathrm{t}})\int f(y,\theta)\,\mathrm{d}\mu(y) = \int (\partial^2/\partial\theta\partial\theta^{\mathrm{t}}) f(y,\theta)\,\mathrm{d}\mu(y)$, and show that indeed second Bartlett identity

$$K(\theta) = J(\theta).$$

This matrix is *the Fisher information matrix* for the model. It provides a measure of how much information about a parameter a data set provides. Note that the calculation of both $J(\theta)$ and $K(\theta)$ is taking place under the assumption that the model is actually correct.

(b). Show that if $\theta \mapsto (\partial/\partial\theta) f(y,\theta)$ is continuous for each y, and each element of $(\partial/\mathrm{d}\partial) f(y,\theta)$ is bounded by an integrable function not depending on θ, then the first partial derivative can be passed under the integral sign in $\int f(y,\theta)\,\mathrm{d}\mu(y)$. Deduce conditions under which we may pass the second partial derivative under the same integral sign.

(c). For the exponential model, with density $\theta\exp(-\theta y)$, find the score function, and compute the Fisher information function in two ways. The second (derivative) way of computing the Fisher information here was quite simple.

(d). Consider the general exponential family, with its natural parametrisation $f(y,\theta) = \exp\{\theta^{\mathrm{t}} T(y) - k(\theta)\}h(y)$, see Ex. 1.55. Explain that the score function becomes $T(y) - k'(\theta)$, with Fisher matrix $J(\theta) = k''(\theta)$, that of the second order derivatives of $k(\theta)$.

(e). For the normal $\mathrm{N}(\xi,\sigma^2)$ model, show that the score function can be expressed as

$$u(y,\xi,\sigma) = \begin{pmatrix} \frac{1}{\sigma}(y-\xi)/\sigma \\ \frac{1}{\sigma}\{(y-\xi)^2/\sigma^2 - 1\} \end{pmatrix} = \frac{1}{\sigma}\begin{pmatrix} z \\ z^2-1 \end{pmatrix},$$

writing $z = (y - \xi)/\sigma$, which is a standard normal when y comes from the model. Demonstrate that the Fisher information matrix becomes

$$J(\xi, \sigma) = \mathrm{Var}_{\xi,\sigma}\, u(Y, \xi, \sigma) = \begin{pmatrix} 1/\sigma^2 & 0 \\ 0 & 2/\sigma^2 \end{pmatrix}.$$

(f). (xx Check with a few more of your favourite parametric models, where you find the score function and the information function, and where then formulae for both $J(\theta)$ and the variance matrix $K(\theta)$ of the score function, verifying that they are the same. ask for poisson, for binomial with parametrisation $p = \exp(\theta)/\{1 + \exp(\theta)\}$, geometric. xx)

(g). (xx more care here, since we do CR a bit later. xx) If Y has the uniform distribution on $[0, \theta]$, which of the regularity conditions listed above fail? In this situation, one might try to define the Fisher information to be $1/\theta^2$. Assuming that this is indeed the Fisher information, use the Cramér–Rao lower bound to derive a contradiction.

Ex. 5.15. *Cramér–Rao lower bounds for estimators.* A certain basic and classic inequality provides a lower bound for the variance of any unbiased estimator of a given parameter. There are various versions and generalisations, some of which we go through here. Inequalities of the type encountered here are sometimes called 'information inequalities', as they may be used to define and analyse how much information there can be in a finite set of data. We also discover clear links to log-likelihoods and ML estimation; the behaviour of ML estimators for growing sample size, described in Ex. 5.18, exactly matches the Cramér–Rao lower bound for variances.

(a). To begin simply, suppose Y is an observation from the density $f(y, \theta)$, assumed smooth in its one-dimensional parameter. Let $u(y, \theta) = \partial \log f(y, \theta)/\partial\theta$ be the score function, with finite variance, by definition equal to the Fisher information, $J(\theta) = \int f(y, \theta) u(y, \theta)^2 \, \mathrm{d}y$. Let $T = T(Y)$ be any estimator unbiased for θ, and assume that $f(y, \theta)$ satisfies the conditions of Ex. A.12(g). From $\mathrm{E}_\theta\, T = \int T(y) f(y, \theta) \, \mathrm{d}y = \theta$, use the conditions on $f(y, \theta)$ to deduce that $(\mathrm{d}/\mathrm{d}\theta) \int T(y) f(y, \theta) u(y, \theta) \, \mathrm{d}y = 1$. Show that $\mathrm{cov}_\theta\,(T, u(Y, \theta)) = 1$, and, consequently $1 \le \mathrm{Var}_\theta\, T \,\mathrm{Var}\, u(Y, \theta)$, from which the we get one-dimensional classic Cramér–Rao inequality

$$\mathrm{Var}_\theta\, T \ge 1/J(\theta).$$

(b). It is easy to generalise the above to the more interesting case of having more observation than one. Suppose $Y_1, \ldots, Y_n$ are i.i.d. from the parametric model $f(y, \theta)$. Show that the arguments above still hold, essentially since $Y = (Y_1, \ldots, Y_n)$ can be considered a single datum from the model with joint density $f(y_1, \theta) \cdots f(y_n, \theta)$. Show that the score function now becomes $u(y_1, \ldots, y_n, \theta) = \sum_{i=1}^n u_0(y_i, \theta)$, writing for emphasis $u_0(y, \theta) = \partial \log f(y, \theta)/\partial\theta$ for the score function for a single observation. Deduce that the combined Fisher information for the full sample is $J_n = \mathrm{Var}_\theta\, u(Y_1, \ldots, Y_n) = nJ_0$, with $J_0 = \mathrm{Var}_\theta\, u_0(Y_i, \theta)$ the information in a single observation.

(c). Show from this that if $T = T(Y_1, \ldots, Y_n)$ is an unbiased estimator for θ, then

Cramér–Rao lower bound

$$\mathrm{Var}_\theta\, T \ge \frac{1}{nJ_0(\theta)} = \frac{1/J_0(\theta)}{n}.$$

This says that there is a clear limit to how well one might estimate a parameter in a model, with n observations. If you're not entirely satisfied with $\mathrm{Var}_\theta\, T = 0.10$, say, and wish for variance 0.05 instead, then shell out more money to get twice as many observations.

(d). Show more generally that if $T = T(Y_1, \ldots, Y_n)$ is an estimator for θ, with mean $\mathrm{E}_\theta\, T = \theta + b(\theta)$, i.e. with a certain bias $b(\theta)$, then

$$\mathrm{Var}_\theta\, T \geq \frac{1}{n}\frac{\{1 + b'(\theta)\}^2}{J_0(\theta)}.$$

In particular, show that there's a lower bound on the mean squared error for *any* estimator (i.e. not merely the unbiased ones):

$$\mathrm{mse}(\theta) = \mathrm{E}\,\{T(Y_1, \ldots, Y_n) - \theta\}^2 \geq n^{-1}\{1 + b'(\theta)\}^2/J_0(\theta) + b(\theta)^2.$$

(e). Go through the following examples, in each case finding the score function, the information $J_0(\theta)$, and the lower bound for any unbiased estimator of the model parameter. (i) y is binomial (n, θ). (ii) y is Poisson θ. (iii) y is normal (θ, σ^2), with σ known. (iv) y is normal (θ, σ^2), with θ known, and σ to be estimated. Comment on the implications of your findings.

Ex. 5.16. *Cramér–Rao bounds for the multidimensional case.* In generalisation of the above situation to the case of multiparameter models, assume first that y is a single observation from the model $f(y, \theta)$, with $\theta = (\theta_1, \ldots, \theta_p)$ of dimension p. Let $u_0(y, \theta) = \partial \log f(y, \theta)/\partial\theta$ be the score function, for such a single y, with the $p \times p$ Fisher information matrix $J_0(\theta) = \mathrm{Var}_\theta\, u_0(Y, \theta)$ assumed positive definite.

(a). For a symmetric $p \times p$ matrix A we write $A \geq 0$ provided it is nonnegative definite, i.e. that $c^{\mathrm{t}} A c \geq 0$ for all c. Show that the covariance matrix of a random vector is necessarily nonnegative, and that $A \geq 0$ is is equivalent to its eigenvalues being nonnegative. Explain that $a_{ii} \geq 0$, for the diagonal elements, but that we may still have $a_{i,j} < 0$ for some off-diagonal elements. With two symmetric matrices, we write $A \geq B$ if $A - B \geq 0$, which is the ordering of variance matrices we use below.

(b). Assume that $T = T(Y)$ is an unbiased estimator of θ, which also means that $\mathrm{E}_\theta\, T_j(Y) = \theta_j$ for each component j. With I_p the identity matrix of size $p \times p$, deduce from $\mathrm{E}_\theta\, T = \int T(y) f(y, \theta)\, \mathrm{d}y$ that

$$(\partial/\partial\theta)\, \mathrm{E}_\theta\, T = \int T(y) f(y, \theta) u_0(y, \theta)^{\mathrm{t}}\, \mathrm{d}y = I_p.$$

(c). Then work out that

$$\mathrm{Var}_\theta\, \{T - J_0(\theta)^{-1} u_0(Y, \theta)\} = \mathrm{E}_\theta\, \{T - \theta - J_0(\theta)^{-1} u_0(Y, \theta)\}\{T - \theta - J_0(\theta)^{-1} u_0(Y, \theta)\}^{\mathrm{t}}$$

can be expressed as $\mathrm{Var}_\theta\, T - J_0(\theta)^{-1}$. We have then shown a multidimensional version of the Cramér–Rao inequality, that $\mathrm{Var}_\theta\, T \geq J_0(\theta)^{-1}$.

(d). Generalise the above to the case of n i.i.d. observations $Y_1, \ldots, Y_n$ from the model. Show that the information matrix for the full data set becomes

$$J_n(\theta) = \mathrm{Var}_\theta\, \frac{\partial \log\{f(Y_1, \theta) \cdots f(Y_n, \theta)\}}{\partial\theta} = n J_0(\theta),$$

Cramér–Rao lower bound, matrix case

and that for *any* unbiased estimator $T = T(Y_1, \ldots, Y_n)$ of θ, we must have

$$\mathrm{Var}_\theta\, T \geq \{n J_0(\theta)\}^{-1} = n^{-1} J_0(\theta)^{-1}.$$

More generally, if independent observations $Y_1, \ldots, Y_n$ come from densities $f_1(y, \theta), \ldots, f_n(y, \theta)$, with Fisher information matrices $J_1(\theta), \ldots, J_n(\theta)$, show that any unbiased estimator T of θ has $\mathrm{Var}_\theta\, T \geq \{J_1(\theta) + \cdots + J_n(\theta)\}^{-1}$.

(e). Suppose now that $\mu = \mu(\theta_1, \ldots, \theta_p)$ is a one-dimensional parameter in focus, and that $T = T(Y)$ is an unbiased estimator. With $c(\theta) = \partial\mu(\theta)/\partial\theta$ the $p \times 1$ gradient, show that

$$\mathrm{Var}_\theta \{T - c(\theta)^{\mathrm{t}} J(\theta)^{-1} u_0(Y)\} = \mathrm{Var}_\theta T - c(\theta)^{\mathrm{t}} J(\theta)^{-1} c(\theta),$$

and conclude the lower bound $\mathrm{Var}_\theta T \geq \lambda(\theta) = c(\theta)^{\mathrm{t}} J(\theta)^{-1} c(\theta)$. Generalise to the case of $Y_1, \ldots, Y_n$ being i.i.d. from $f(y, \theta)$: if $T_n = T(Y_1, \ldots, Y_n)$ is unbiased for μ, show that $\mathrm{Var}_\theta T \geq \lambda(\theta)/n$. We learn also a bit more from these arguments. Show that

$$\mathrm{Var}_\theta \Big\{ T_n - n^{-1} \sum_{i=1}^n c(\theta)^{\mathrm{t}} J(\theta)^{-1} u_0(Y_i, \theta) \Big\} = \mathrm{Var}_\theta T_n - \lambda(\theta)/n,$$

so T_n having a variance coming close to the lower bound means $T_n - \mu$ being equal to or close to the random variable $n^{-1} \sum_{i=1}^n c(\theta)^{\mathrm{t}} J(\theta)^{-1} u_0(Y_i, \theta)$.

(f). (xx something here, tying these matters to ML. consider $T_n^* = \theta + n^{-1} J(\theta)^{-1} \sum_{i=1}^n u_0(Y_i, \theta)$. show that this T_n^* is unbiased and achieves the Cramér–Rao lower bound. this happens in exponential families, give pointer. the point is then that ML in the general case is close to this with growing n. xx)

(g). Also other estimators for θ deserve to be studied, even when they are not exactly unbiased. We start with a single observation Y from $f(y, \theta)$, with score function $u_0(y, \theta)$ as above, and then generalise to n observations afterwards. Assume therefore that $T = T(Y)$ is such that

$$\mathrm{E}_\theta T = \int T(y) f(y, \theta)\, \mathrm{d}y = \theta + b(\theta) = \begin{pmatrix} \theta_1 + b_1(\theta) \\ \vdots \\ \theta_p + b_p(\theta) \end{pmatrix},$$

for suitable bias functions $b_1(\theta), \ldots, b_p(\theta)$, perhaps not far from zero. Show that

$$(\partial/\partial\theta)\, \mathrm{E}_\theta T = \begin{pmatrix} 1 + \partial b_1(\theta)/\partial\theta \\ \vdots \\ 1 + \partial b_p(\theta)/\partial\theta \end{pmatrix} = \int T(y) f(y, \theta) u_0(y, \theta)^{\mathrm{t}}\, \mathrm{d}y.$$

Then work with $\mathrm{Var}_\theta [T - \{I_p + b'(\theta)\} J_0(\theta)^{-1} u_0(Y, \theta)]$ to demonstrate that

$$\mathrm{Var}_\theta T \geq \{I_p + b'(\theta)\} J_0(\theta)^{-1} \{I + b'(\theta)\}^{\mathrm{t}}.$$

(h). Generalise to the case of n i.i.d. observations, to reach

$$\mathrm{Var}_\theta T \geq n^{-1} \{I_p + b'(\theta)\} J_0(\theta)^{-1} \{I + b'(\theta)\}^{\mathrm{t}}.$$

(i). (xx a bit more to round it off. an example or two. CR lower bound not always attained, in some models only for growing n, but that's ok. i make a separate point that the arguments also lead to bounds of type

$$\mathrm{Var}_\theta T \geq \{J_1(\theta) + \cdots + J_n(\theta)\}^{-1},$$

in cases with different situations or types of information sources, for the same θ. tie it all to the large-sample ML results. xx)

Ex. 5.17. *Fisher information and parameter transformations.* Consider a parametric model with density $f(y,\theta)$, with θ of length p, say. We have seen the central role the Fisher information matrix $J(\theta)$ plays, with more to come for the likelihood machinery below. Here we look into reparametrisation, going from $\theta = (\theta_1,\ldots,\theta_p)$ to say $\phi = (\phi_1,\ldots,\phi_p)$, with smooth 1-1 transformations $\theta = h(\phi)$, or $\phi = h^{-1}(\theta)$. It is assumed that the $p \times p$ matrix $G(\phi) = \partial h(\phi)/\partial\phi$, with rows $\partial h_j(\phi)/\partial\phi$, exists and is non-zero throughout the parameter space.

(a). Let $u(y,\theta) = \partial \log f(y,\theta)/\partial\theta$ be the score function with the θ parametrisation. Show that for the ϕ parametrisation, we have

$$u^*(y,\phi) = \frac{\partial \log f(y,h(\phi))}{\partial\phi} = \begin{pmatrix} u_1(y,h(\phi))^{\rm t} g_1(\phi) \\ \vdots \\ u_p(y,h(\phi))^{\rm t} g_p(\phi) \end{pmatrix} = G(\phi)u(y,h(\phi)).$$

(b). Show from this that the Fisher information matrix for the ϕ parametrisation becomes $J^*(\phi) = G(\phi)J(h(\theta))G(\phi)^{\rm t}$. Deduce also that the important inverse can be written $J^*(\phi)^{-1} = (G(\phi)^{\rm t})^{-1}J(h(\theta))^{-1}G(\phi)^{-1}$.

(c). As illustrations, consider the binomial (n,θ), with reparametrisations (i) $\phi_1 = \theta^{1/2}$; (ii) $\phi_2 = \log(\theta/(1-\theta))$, so that $\theta = \exp(\phi)/\{1+\exp(\theta)\}$; (iii) $\phi = 2\arcsin(\theta^{1/2})$. For each of these, find expressions for the Fisher information $J^*(\phi)$.

(d). Consider the normal (ξ,σ^2). Find first the Fisher information matrix for the (ξ,σ) parametrisation, then for the (ξ,λ) parametrisation, where $\lambda = \sigma^2$.

Ex. 5.18. *Maximum likelihood estimators.* Thanks to the general efforts in exercises above, in particular 5.12–5.13, we may already learn the basic properties of the most important of all estimation methods, namely the ML method, introduced in Ex. 5.6. We do return to further details, extensions, results, illustrations, applications in exercises to come. The basic setup, to be generalised later, is that of observations $Y_1,\ldots,Y_n$ being i.i.d. from some $g(y)$, to be fitted to a parametric $f_\theta(y) = f(y,\theta)$, and with $\widehat\theta$ the maximiser of $\ell_n(\theta) = \sum_{i=1}^n \log f(Y_i,\theta)$. The log-density derivatives $u(y,\theta)$ and $i(y,\theta)$ are those studied in Ex. 5.14. There is a short list of regularity conditions to secure results reached below, inherited from those called (i)-(v) in Ex. 5.13.

(a). Arguably the first natural question, for any estimation method, is what it aims for. Show that $\widehat\theta \to_{\rm pr} \theta_0$, the minimiser of $\mathrm{KL}(g,f_\theta)$, assumed here to be unique and an inner point in the parameter space. So we've uncovered what goes on in the mindset of the ML operator; it aims for this *least false parameter*, the θ_0 minimising the KL distance from the truth to the model approximation.

the least false parameter value

(b). With p the dimension of θ, show that the two crucial $p \times p$ matrices J and K, from the general treatment of Ex. 5.13, become

$$J = -\mathrm{E}_G\, i(Y,\theta_0) \quad \text{and} \quad K = \mathrm{Var}_G\, u(Y,\theta_0).$$

We have seen in Ex. 5.14 that these are equal under model conditions, i.e. that $g(y) = f(y,\theta_0)$ for all y.

(c)**.** Via efforts in previous exercises, show (i) that $U_n = n^{-1/2}\sum_{i=1}^n u(Y_i,\theta_0) \to_d U \sim \mathrm{N}_p(0,K)$; (ii) that $J_n = -n^{-1}\sum_{i=1}^n i(Y_i,\theta_0) \to_{\mathrm{pr}} J$; and the basic log-likelihood process convergence result

$$\begin{aligned} A_n(s) &= \ell_n(\theta_0 + s/\sqrt{n}) - \ell_n(\theta_0) \\ &= U_n^{\mathrm{t}} s - \tfrac{1}{2} s^{\mathrm{t}} J_n s + r_n(s) \to_d A(s) = U^{\mathrm{t}} s - \tfrac{1}{2} s^{\mathrm{t}} J s, \end{aligned} \tag{5.7}$$

for each s. This drives much of the large-sample results for likelihood inference, including the important Wilks theorems we return to in Ex. 5.29, along with further generalisations for e.g. regression models. Show also that

$$B_n(s) = \ell_n(\widehat{\theta} + s/\sqrt{n}) - \ell_n(\widehat{\theta}) = -\tfrac{1}{2} s^{\mathrm{t}} \widehat{J}_n s + r_n'(s) \to_d B(s) = -\tfrac{1}{2} s^{\mathrm{t}} J s,$$

for each s, i.e. $r_n'(s) \to_{\mathrm{pr}} 0$, where $\widehat{J}_n = -n^{-1}\partial^2 \ell_n(\widehat{\theta})/\partial\theta\,\partial\theta^{\mathrm{t}}$ is the *normalised observed Fisher information matrix*.

(d)**.** Let $\ell_{n,\max} = \ell_n(\widehat{\theta})$ be the maximised log-likelihood. From (5.6), deduce that

$$\begin{aligned} \sqrt{n}(\widehat{\theta} - \theta_0) &\to_d J^{-1}U \sim \mathrm{N}_p(0, J^{-1}KJ^{-1}), \\ 2\{\ell_{n,\max} - \ell_n(\theta_0)\} &\to_d W = U^{\mathrm{t}} J^{-1} U, \end{aligned}$$

in which $U \sim \mathrm{N}_p(0,K)$. Remarkably, we have been able to reach these general results for ML estimators without going into special cases, and without caring about there being explicit formulae for these or not.

(e)**.** As simple corollaries, we are also reaching the important and very frequently used consequences for ML estimation in general smooth parametric models, under model conditions: with $J = J(\theta_0)$ the Fisher information matrix, at the underlying true parameter value, show that

$$\sqrt{n}(\widehat{\theta} - \theta_0) \to_d \mathrm{N}_p(0, J(\theta_0)^{-1}) \quad \text{and} \quad 2\{\ell_{n,\max} - \ell_n(\theta_0)\} \to_d \chi^2_p. \tag{5.8}$$

Ex. 5.19. *The J, K, and sandwich matrices for the BHHJ method.* (xx place inside prose here that we also need estimation of these matrices, to which we return in Ex. 5.28. xx) When fitting a parametric family $f(y,\theta)$ to observations, we have seen in Ex. 5.12-5.13 that the behaviour of minimum divergence function estimators is characteriserd by (i) the implied least false parameter $\theta_0 = \theta_0(G)$ and then (ii) the crucial matrices J and K, leading to the sandwich matrix $\Sigma = J^{-1}KJ^{-1}$. As above G, with density g, is the data generating mechanism. Here we consider the BHHJ estimation method Ex. 5.9, with the estimator $\widehat{\theta}$ minimising $H_{n,a}(\theta)$ of (5.5). Again, a is a positive finetuning parameter, and for small a the method is close to ML estimation. For the following points, we assume regularity conditions (i)-(v) are in place, from Ex. 5.12-5.13, and in particular that there is a unique minimiser $\theta_0 = \theta_{0,a} = \theta_{0,a}(G)$ of the $d_a(g, f_\theta)$ of (5.4).

(a)**.** Explain that the theory developed in these earlier exercises implies that $\widehat{\theta} \to_{\mathrm{pr}} \theta_0$. In the following points, write for simplicity $f_0(y) = f(y,\theta_0)$, $u_0(y) = u(y,\theta_0)$, $i_0(y) = i(y,\theta_0)$. Show also that θ_0 is characterised as the solution to $\int f_\theta^{1+a} u_\theta \,\mathrm{d}y = \int g f_\theta^a u_\theta \,\mathrm{d}y$. For this vector, evaluated at θ_0, we write

$$\xi_a = \int g f_0^a u_0 \,\mathrm{d}y = \int f_0^{1+a} u_0 \,\mathrm{d}y.$$

(b). As seen in Ex. 5.9, the BHHJ method corresponds to using $h(y,\theta) = \int f_\theta^{1+a}\,\mathrm{d}y - (1+1/a)f(y,\theta)^a$ in the general minimum divergence function setup. Find the derivative $h'(y,\theta)$ and deduce that

$$K_a = (1+a)^2\Big\{\int g f_0^{2a} u_0 u_0^{\mathrm{t}} - \xi_a \xi_a^{\mathrm{t}}\Big\},$$

(xx find K_a. with simplified at model. xx)

(c). (xx check this carefully. also: when estimating J_a, we don't need such a formula, as we get it as the Hessian of the minimisation. xx) Find then an expression for the second derivative $h''(y,\theta)$, and derive a formula for the J matrix:

$$J_a = (1+a)\int \{f_0^{1+a} u_0 u_0^{\mathrm{t}} + (f_0^{1+a} - g f_0^a)(i_0 + a u_0 u_0^{\mathrm{t}})\}\,\mathrm{d}y.$$

Ex. 5.20. *BHHJ precision under model conditions.* The BHHJ method is designed to do well when the parametric model used is not perfect. We may however also study its precision when the model actually holds. For matrices studied in Ex. 5.19 there are simplified expressions under model conditions, i.e. when $g(y) = f(y,\theta_0)$ for all y.

(a). Under such model conditions, and with f_0, u_0 notation from Ex. 5.19, show that

$$J_a = (1+a)\int f_0^{1+a} u_0 u_0^{\mathrm{t}}\,\mathrm{d}y, \quad K_a = (1+a)^2\Big(\int f_0^{1+2a} u_0 u_0^{\mathrm{t}}\,\mathrm{d}y - \xi_a \xi_a^{\mathrm{t}}\Big),$$

with a consequent simplified sandwich matrix $\Sigma_a = J_a^{-1} K_a J_a^{-1}$. Show that for $a \to 0$, $\Sigma_a \to J_0^{-1}$, with $J_0 = J(\theta_0)$ the Fisher information matrix for the model.

(b). We may now check the loss of efficiency of the BHHJ method, with tuning parameter a, compared to the ML method, under model conditions, by comparing $J_a^{-1} K_a J_a^{-1}$ to the inverse Fisher information matrix. Carry out the relevant computations for estimation θ in the exponential $\theta \exp(-\theta y)$ model. For the case of the $\mathrm{N}(\xi, \sigma^2)$ model, check with Story vii.5. For small a, there is significant gain in robustness at a low cost in efficiency.

(c). For the following, assume

$$A = \int f_0 \log f_0\, u_0 u_0^{\mathrm{t}}\,\mathrm{d}y, \quad B = \int f_0 (\log f_0)^2\, u_0 u_0^{\mathrm{t}}\,\mathrm{d}y, \quad c = \int f_0 \log f_0\, u_0\,\mathrm{d}y$$

are finite; A and B are $p \times p$ matrices and c a $p \times 1$ vector. Below we carry out Taylor expansions to order a^2, and use $\doteq$ to indicate the consequent approximations for small a. Show that to this order

$$J_a \doteq (1+a)(J_0 + aA + \tfrac{1}{2}a^2 B), \quad K_a \doteq (1+a)^2\{J_0 + 2aA + a^2(2B - cc^{\mathrm{t}})\}.$$

As a useful interlude, consider $p \times p$ symmetric matrices M and ε, where M is positive definite and ε is smaller, in the technical sense that the eigenvalues of $M^{-1}\varepsilon$ are smaller than 1 in absolute value. Show that

$$(M+\varepsilon)^{-1} = M^{-1} - M^{-1}\varepsilon M^{-1} + M^{-1}\varepsilon M^{-1}\varepsilon M^{-1} + \cdots.$$

Use this to reach

$$(J_0 + aA + \tfrac{1}{2}a^2 B)^{-1} \doteq J_0^{-1} - aJ_0^{-1}AJ_0^{-1} + a^2(J_0^{-1}AJ_0^{-1}AJ_0^{-1} - \tfrac{1}{2}J_0^{-1}BJ_0^{-1}).$$

It is then a matter of bureaucratic algebra to reach an informative approximation to the sandwich matrix

$$\Sigma_a \doteq (J_0 + aA + \tfrac{1}{2}a^2B)^{-1}\{J_0 + 2aA + a^2(2B - cc^{\mathrm{t}})\}(J_0 + aA + \tfrac{1}{2}a^2B)^{-1}.$$

Show that this leads to $\Sigma_a \doteq J_0^{-1} + a^2D$, with $D = J_0^{-1}(B - cc^{\mathrm{t}} - AJ_0^{-1}A)J_0^{-1}$. This indicates indeed that for small a, the efficiency loss is small.

Ex. 5.21. *The J, K, and sandwich matrices for the maximum weighted likelihood method.* (xx check all this. xx) Just as we for the BHHJ estimation method worked out the basics for J_a, K_a, $J_a^{-1}K_aJ_a^{-1}$ in Ex. 5.19, we here tend to their parallels for the maximum weighted likelihood method of Ex. 5.11. For a parametric model $f(y,\theta)$, and with a given weight function $w(y)$, it consists in maximising the weighted log-likelihood function $\ell_{n,w}(\theta) = \sum_{i=1}^n w(y_i)\log f(Y_i,\theta) - nB(\theta)$, with $B(\theta) = \int w(y)f(y,\theta)\,\mathrm{d}y$. We assume the regularity conditions of Ex. 5.13 are in force.

(a). When observations $Y_1,\ldots,Y_n$ stem from G, with density g, show that the MWL estimator $\widehat{\theta}$ tends to $\theta_0 = \theta_0(G)$, the minimiser of the weighted KL divergence $\mathrm{KL}_w(g,f_\theta)$. When g is not inside the parametric family, this least false parameter value depends also on the weight function. Show also that θ_0 is characterised as the solution to $\int w(y)\{g(y)-f(y,\theta_0)\}u(y,\theta_0)\,\mathrm{d}y = 0$. As for Ex. 5.19 we write f_0,u_0,i_0 for the functions $f(y,\theta_0),u(y,\theta_0),i(y,\theta_0)$, and $\xi_0 = \int wgu_0\,\mathrm{d}y = \int wf_0u_0\,\mathrm{d}y$.

(b). Explain that maximising $\ell_{n,w}(\theta)$ is the same as minimising $n^{-1}\sum_{i=1}^n h(Y_i,\theta)$, with $h(y,\theta) = w(y)\log f(y,\theta) - B(\theta)$, or solving $n^{-1}\sum_{i=1}^n h'(Y_i,\theta) = 0$, with $h'(y,\theta) = w(y)u(y,\theta) - \xi(\theta)$, where $\xi(\theta) = \partial B(\theta)/\partial\theta = \int w(y)f(y,\theta)u(y,\theta)\,\mathrm{d}y$. Show that

$$K_w = \int w^2gu_0u_0^{\mathrm{t}}\,\mathrm{d}y - \xi_0\xi_0^{\mathrm{t}}.$$

Under model conditions, we have $\xi_0 = 0$, and the above simplies to $K_w = \int w^2f_0u_0u_0^{\mathrm{t}}\,\mathrm{d}y$.

(c). Then show that $J_w = -\int wgi_0\,\mathrm{d}y + \partial^2B(\theta_0)/\partial\theta\,\partial\theta^{\mathrm{t}}$. Under model conditions, $J_w = -\int wf_0i_0\,\mathrm{d}y$.

(d). (xx find sandwich, and compare with J^{-1} under model. xx)

Ex. 5.22. *Maximum weighted likelihood for multinomial models.* The weighted likelihood ideas of Ex. 5.11, 5.21 can also be used to estimate parameters in models for multinomial probabilities, allowing different weights of importance for different outcomes. (xx pointer to Story ii.8. xx)

(a). Show first that $p\log p/q-(p-q)$ is always nonnegative, for p,q in $(0,1)$. Now consider $(Y_1,\ldots,Y_k)$ being multinomial with count n and probability vector $p = (p_1,\ldots,p_k)$, as with Ex. 1.5. Then

$$d_w(p,p_\theta) = \sum_{j=1}^k w_j\{p_j\log(p_j/p_{j,\theta}) - (p_j - p_{j,\theta})\},$$

is a proper nonnegative divergence, from p to some p_θ, as long as the weights $w_1,\ldots,w_k$ are nonnegative. With equal weights, show that this is the Kullback–Leibler divergence $\mathrm{KL}(p,p_\theta)$. If $p_\theta = (p_1(\theta),\ldots,p_k(\theta))$ is some postulated model, with θ of dimension r,

say, show that the minimum divergence method amounts to maximising the weighted and modified log-likelihood

$$\ell_{n,w}(\theta) = n\sum_{j=1}^{k} w_j\{\widehat{p}_j \log p_j(\theta) - p_j(\theta)\} = \sum_{j=1}^{k} w_j Y_j \log p_j(\theta) - n\sum_{j=1}^{k} w_j p_j(\theta),$$

with $\widehat{p}_j = Y_j/n$. This generalises the usual ML method, which corresponds to equal weights.

(b). Let $p = (p_1, \ldots, p_k)$ denote the true probability vector. Explain that the weighted likelihood estimator $\widehat{\theta}_w$, maximising $\ell_{n,w}(\theta)$, tends in probability to the least false parameter maximising $\sum_{j=1}^{k} w_j\{p_j \log p_j(\theta) - p_j(\theta)\}$, i.e. to the minimiser of the weighted KL distance $d_w(p, p_\theta)$. Writing $u_j(\theta) = \partial \log p_j(\theta)/\partial\theta$, show that the least false parameter is also the solution to $\sum_{j=1}^{k} w_j\{p_j - p_j(\theta)\}u_j(\theta) = 0$, assumed here to be unique. Show furthermore that

$$U_n = n^{-1/2}\frac{\partial \ell_{n,w}(\theta_0)}{\partial\theta} = \sum_{j=1}^{k} w_j\sqrt{n}\{\widehat{p}_j - p_j(\theta_0)\}u_j(\theta_0) \to_d U = \sum_{j=1}^{k} w_j Z_j u_j(\theta_0),$$

using results and notation from the multinomial CLT worked with in Ex. 2.49. Deduce that $U \sim \mathrm{N}_r(0, K_w)$, with

$$K_w = \sum_{j=1}^{k} w_j^2 p_j u_j(\theta_0)u_j(\theta_0)^{\mathrm{t}} - \xi_0\xi_0^{\mathrm{t}}, \quad \text{with } \xi_0 = \sum_{j=1}^{k} w_j p_j u_j(\theta_0).$$

(c). Next demonstrate that the normalised Hessian matrix $\widehat{J}_{n,w} = -n^{-1}\partial^2\ell_{n,w}(\widehat{\theta})/\partial\theta\partial\theta^{\mathrm{t}}$ tends in probability to a well-defined J_w. In fact, writing $i_j(\theta) = \partial^2 \log p_j(\theta)/\partial\theta\partial\theta^{\mathrm{t}}$, show that $J_w = J_w(\theta_0)$, where

$$\begin{aligned} J_w(\theta) &= -\frac{\partial^2}{\partial\theta\partial\theta^{\mathrm{t}}}\sum_{j=1}^{k} w_j\{p_j \log p_j(\theta) - p_j(\theta)\} \\ &= \sum_{j=1}^{k} w_j\big[p_j(\theta)u_j(\theta)u_j(\theta)^{\mathrm{t}} - \{p_j - p_j(\theta)\}i_j(\theta)\big], \end{aligned}$$

with the expression simplifying under model conditions. Use this in conjunction with the general minimum divergence theory to establish that $\sqrt{n}(\widehat{\theta}_w - \theta_0) \to_d J_w^{-1}U \sim \mathrm{N}_r(0, \Sigma_w)$, with the sandwich matrix $\Sigma_w = J_w^{-1}K_w J_w^{-1}$.

(d). (xx a simple example here, with pointer to Story ii.8. xx)

Ex. 5.23. *Completing increasingly simpler tasks.* In a certain game of learning a player needs to complete tasks $1, 2, \ldots, n$, which become increasingly simpler with each passing of a new level. Assume that the time needed to complete these tasks are $V_1, \ldots, V_n$, with these being independent with $V_i \sim \mathrm{Expo}(i/\theta)$, where θ is an unknown parameter. – For the following questions, you may encounter the partial sums

$$a_n = 1 + 1/2 + 1/3 + \cdots + 1/n, \quad b_n = 1 + 1/2^2 + 1/3^2 + \cdots + 1/n^2.$$

Here the first is slowly divergent, with $a_n \doteq \log n + 0.5772$, and the second is convergent, with $b_n \to \pi^2/6$, as shown by Euler in 1734, bringing him instant world fame.

(a). Find expressions for the mean and variance of $T_n = V_1 + \cdots + V_n$, the time it takes the player to complete all tasks. In particular, show that T_n has mean $a_n\theta$. Put up the unbiased estimator based on T_n, say $\widehat{\theta}$. Find its variance, and show that the estimator is consistent.

(b). Then work out a formula for the log-likelihood function, based on having observed not merely the total time T_n, but the individual waiting times $V_1, \ldots, V_n$. Find the maximum likelihood estimator, say θ^*. Show that also this estimator θ^* is unbiased, and compare its variance to that of $\widehat{\theta}$. Find also the Cramér–Rao lower bound for variances of unbiased estimators for θ, and comment.

(c). Assume the game goes on, up to level $2n$, and consider the time a player needs to pass the last half of these levels, i.e. $T_n^* = T_{2n} - T_n$. Show that T_n^* tends in probability to a certain limit as n grows.

Ex. 5.24. *Profiling quadratic functions.* Consider some quadratic function $A(s) = \frac{1}{2}s^{\mathrm{t}}Js$, with J symmetric and positive definite of dimension $p \times p$. It is useful to sort out minima of A under different types of side constraints.

(a). We start out examining the minimum of $A(s)$ over all s with $c^{\mathrm{t}}s = x$, for some given vector c and level x. The Lagrange multiplier way of solving such a problem is to minimise the function $\frac{1}{2}s^{\mathrm{t}}Js - \lambda(c^{\mathrm{t}}s - x)$, with no constraint on s, and in the process find the λ agreeing with the constraint. Taking derivatives, show that minimum occurs for $s_0 = \lambda J^{-1}c$, leading to $c^{\mathrm{t}}s_0 = \lambda c^{\mathrm{t}}J^{-1}c$, which should then equal x. Explain that this leads to minimiser $s_0 = x/(c^{\mathrm{t}}J^{-1}c)$ and attained minimum $A_{\min} = \frac{1}{2}x^2/(c^{\mathrm{t}}J^{-1}c)$.

(b). Generalising the above, from a scalar to a vector, consider a $p_0 \times p$ matrix C and an x of dimension p_0, where it is assumed that $CJ^{-1}C^{\mathrm{t}}$ has full rank. The task is to work out the minimiser and minimum of $\frac{1}{2}s^{\mathrm{t}}Js$ over all s with $Cs = x$. Work out that the minimiser is

$$s_0 = J^{-1}C^{\mathrm{t}}\lambda = J^{-1}C^{\mathrm{t}}(CJ^{-1}C^{\mathrm{t}})^{-1}x \text{ with } A_{\min} = \tfrac{1}{2}x^{\mathrm{t}}(CJ^{-1}C^{\mathrm{t}})^{-1}x.$$

(xx we should find the following from this; round off. xx) with $\mu = k(\theta) = (\mu_1, \ldots, \mu_{p_0})$,

$$n(H_{n,\min,\mathrm{narr}} - H_{n,\min,\mathrm{wide}}) \to_d \tfrac{1}{2}U^{\mathrm{t}}J^{-1}C(C^{\mathrm{t}}J^{-1}C)^{-1}C^{\mathrm{t}}J^{-1}U.$$

Ex. 5.25. *Profiling a minimum divergence function, I.* For data $Y_1, \ldots, Y_n$ from some distribution G, we have in Ex. 5.5 considered estimating a parameter $\theta_0 = \mathrm{argmin}(H)$, for $H(\theta) = \mathrm{E}_G\, h(Y, \theta)$, by minimising the distance function $H_n(\theta) = n^{-1}\sum_{i=1}^n h(Y_i, \theta)$. For a focus parameter $\mu = k(\theta)$, a smooth function of $\theta = (\theta_1, \ldots, \theta_p)$, it is useful to work with the associated profile function

profiling a distance function

$$H_{n,\mathrm{prof}}(\mu) = \min\{H_n(\theta)\colon k(\theta) = \mu\}.$$

(a). As an introductory illustration, consider estimating the parameters (a, b) of a Gamma distribution via minimum L_2, as in Ex. 5.3. Simulate 100 datapoints from a gamma; compute $(\widehat{a}, \widehat{b})$; and compute and display also the profile function $H_{n,\mathrm{prof}}(\mu)$ for the mean parameter $\mu = a/b$.

(b). From the full model, with ensuing minimum divergence function estimator $\widehat{\theta}$, show that the consequent $\widehat{\mu} = k(\widehat{\theta})$ becomes normal. With setup and notation as in Ex. 5.13,

prove indeed that $\sqrt{n}(\widehat{\mu}-\mu_0) \to_d c^{\rm t}J^{-1}U$, with $c=\partial k(\theta_0)/\partial\theta$, and that the limit is a zero-mean normal with variance $c^{\rm t}J^{-1}KJ^{-1}c$. Show also that this $\widehat{\mu}$ is identical to the minimiser of the profile function.

(c)**.** In other words we already know the basic story for any focus parameter estimator $\widehat{\mu}$, thanks to the delta method. It is however fruiful to work with representations and approximations stemming from examining the associated profile function. With methods from Ex. 5.13, show that

$$n\{H_n(\widehat{\theta}+s/\sqrt{n})-H_n(\widehat{\theta})\} = \tfrac{1}{2}s^{\rm t}J_ns+r_n(s), \quad \text{with } r_n(s)=O_{\rm pr}(\|s\|^3/\sqrt{n}).$$

For the profiling, therefore, we must minimise this expression over all s such that $k(\theta)=k(\widehat{\theta}+s/\sqrt{n})=\mu$. With $k(\theta)=\widehat{\mu}+c_n^{\rm t}s/\sqrt{n}+O_{\rm pr}(\|s\|^2/n)$, here writing $c_n=\partial k(\widehat{\theta})/\partial\theta$, the essence is to minimise $\frac{1}{2}s^{\rm t}J_ns$ under $c_n^{\rm t}s=\sqrt{n}(\mu-\widehat{\mu})=x_n$, say. Appealing to Ex. 5.24, show that this minimum becomes $\frac{1}{2}x_n^2/c_n^{\rm t}J_n^{-1}c_n=n(\widehat{\mu}-\mu)^2/c_n^{\rm t}J_n^{-1}c_n$. Fill in more details to prove that with $\mu_0=k(\theta_0)$ the true parameter in question,

$$2n\{H_{n,{\rm prof}}(\mu_0)-H_{n,{\rm prof}}(\widehat{\mu})\} = \frac{n(\widehat{\mu}-\mu_0)^2}{c_n^{\rm t}J_n^{-1}c_n}+o_{\rm pr}(1) \to_d \frac{(c^{\rm t}J^{-1}U)^2}{c^{\rm t}J^{-1}c} \sim \kappa\chi_1^2,$$

with $\kappa=c^{\rm t}J^{-1}KJ^{-1}c/c^{\rm t}J^{-1}c$.

(d)**.** (xx explain that this often leads to better approximations than the direct limiting normality thing. then an illustration of this. and pointer to Wilks. and pointer to regression versions of these methods and results; the Y_i need not at all be i.i.d. xx)

Ex. 5.26. *Profiling a distance function, II.* In Ex. 5.5, 5.13, 5.25 we have considered parameters defined as minimisers of functions $H(\alpha)=\mathrm{E}_G\,h(Y,\alpha)$, and developed the basic theory for the associated minimum divergence function estimators. We now consider situations where some of the components of the argmin(H) parameter are specified. Such occur when one tests for lower-dimensional structure, etc. This invites setting up the following framework, with a wide model having $\alpha=(\theta,\gamma)$ of length $p+q$, and the narrow model considered has (θ,γ_0), with θ unknown but $\gamma=\gamma_0$ fixed. Estimators $(\widehat{\theta},\widehat{\gamma})$ in the wide model minimise $H_n(\theta,\gamma)=\int h(y,\theta,\gamma)\,{\rm d}G_n(y)$ whereas $\widetilde{\theta}$ for the narrow model minimises $H_n(\theta,\gamma_0)=\int H(y,\theta,\gamma_0)\,{\rm d}G_n(y)$. The theory developed in the previous exercises mentioned holds for the wide and the narrow models, separately, and below we postulate that the regularity conditions (i)–(v) put up in Ex. 5.13 are in force. Efforts of linear and matrix algebra are required in order to handle these models jointly, however. Define therefore

$$J_{\rm wide}=\mathrm{E}_G\frac{\partial^2 H(Y,\theta_0,\gamma_0)}{\partial\alpha\partial\alpha^{\rm t}}=\begin{pmatrix}J_{00}&J_{01}\\J_{10}&J_{11}\end{pmatrix} \quad \text{with inverse} \quad J_{\rm wide}^{-1}=\begin{pmatrix}J^{00}&J^{01}\\J^{10}&J^{11}\end{pmatrix},$$

where $J_{00}=J_{\rm narr}$ is of size $p\times p$, etc. There is similarly a $(p+q)\times(p+q)$ matrix $K_{\rm wide}$, with submatrices $K_{00},K_{01},K_{10},K_{11}$, the variance matrix of $U=(U_0^{\rm t},U_1^{\rm t})^{\rm t}$, the first derivative $\partial H(Y,\theta_0,\gamma_0)/\partial\alpha$.

(a)**.** The following developements are under the $\gamma=\gamma_0$ constraint, so $\alpha_0=(\theta_0,\gamma_0)$ is the true parameter, determined by the distribution G. Argue that

$$\begin{pmatrix}\sqrt{n}(\widehat{\theta}-\theta_0)\\ \sqrt{n}(\widehat{\gamma}-\gamma_0)\end{pmatrix} \to_d -J_{\rm wide}^{-1}\begin{pmatrix}U_0\\U_1\end{pmatrix}, \quad \sqrt{n}(\widetilde{\theta}-\theta_0)\to_d -J_{00}^{-1}U_0.$$

Show also, again using results reached earlier, that

$$n\{H_{n,\text{wide}} - H_n(\theta_0,\gamma_0)\} \to_d -\tfrac{1}{2}U^{\mathrm{t}}J_{\text{wide}}^{-1}U,$$
$$n\{H_{n,\text{narr}} - H_n(\theta_0,\gamma_0)\} \to_d -\tfrac{1}{2}U_0^{\mathrm{t}}J_{00}^{-1}U_0,$$

with $H_{n,\text{wide}} = H_n(\widehat{\theta},\widehat{\gamma})$ and $H_{n,\text{narr}} = H_n(\widetilde{\theta},\gamma_0)$. Deduce that

$$W_n = 2n(H_{n,\text{narr}} - H_{n,\text{wide}}) \to_d W = U^{\mathrm{t}}J_{\text{wide}}^{-1}U - U_0^{\mathrm{t}}J_{00}^{-1}U_0. \tag{5.9}$$

Show that this limit variable has mean $\mathrm{Tr}(J_{\text{wide}}^{-1}K_{\text{wide}}) - \mathrm{Tr}(J_{00}^{-1}K_{00})$.

(b). (xx clean and simplify this. xx) We tend to a few matrix and submatrix identities here, as they come in handy for some of the technical arguments below. The $q\times q$ matrix J^{11} has an important role, here and on later occasions (as with the FIC in Ch. 11). Show that

$$Q = J^{11} = (J_{11} - J_{10}J_{00}^{-1}J_{01})^{-1}. \tag{5.10}$$

Similarly, we have $J^{00} = J_{00}^{-1} + J_{00}^{-1}J_{01}QJ_{10}J_{00}^{-1}$. Show also that $J^{00} - J_{00}^{-1} = J^{01}J_{10}J_{00}^{-1}$, $J^{10} = -QJ_{10}J_{00}^{-1}$.

(c). We now use the structure of K_{wide} to transform (U_0,U_1) to (U_0,V), with $V = U_1 - K_{10}K_{00}^{-1}U_0$, the point being that U_0 and V become independent. Work through the details of

$$\begin{pmatrix}U_0\\V\end{pmatrix} = \begin{pmatrix}U_0\\U_1 - K_{10}K_{00}^{-1}U_0\end{pmatrix} \sim \mathrm{N}_{p+q}(0, \begin{pmatrix}K_{00} & 0\\0 & K_{11} - K_{10}K_{00}^{-1}K_{01}\end{pmatrix}).$$

Show also that the variance of V is the same as $(K^{11})^{-1}$ (xx check with care xx).

(d). With this transformation, work out the following formula for W, in terms of the independent U_0 and V (xx check all this xx):

$$\begin{aligned}W &= U_0^{\mathrm{t}}(J^{00} - J_{00}^{-1})U_0 + (V + K_{10}K_{00}^{-1}U_0)^{\mathrm{t}}Q(V + K_{10}K_{00}^{-1}U_0)\\ &\quad + 2\,U_0^{\mathrm{t}}J^{01}(V + K_{10}K_{00}^{-1}U_0)\\ &= V^{\mathrm{t}}QV + U_0^{\mathrm{t}}(J^{00} - J_{00}^{-1} + K_{00}^{-1}K_{01}QK_{10}K_{00}^{-1} + 2\,J^{01}K_{10}K_{00}^{-1})U_0\\ &\quad + U_0^{\mathrm{t}}(J^{01} + K_{00}^{-1}K_{01}Q)V + V^{\mathrm{t}}(J^{10} + QK_{10}K_{00}^{-1})U_0.\end{aligned}$$

(e). There are additional informative and insightful representations of the W above. Start by showing $\sqrt{n}(\widehat{\gamma} - \gamma_0) \to_d -Z$, where

$$\begin{aligned}Z &= J^{10}U_0 + J^{11}U_1\\ &= J^{10}U_0 + Q(V + K_{10}K_{00}^{-1}U_0) = QV + (J^{10} + QK_{10}K_{00}^{-1})U_0,\end{aligned}$$

Show that $Z \sim \mathrm{N}_q(0,\Sigma_{11})$, with $\Sigma = J^{-1}KJ^{-1}$ the sandwich matrix. The point is now to demonstrate that W above is identical to $W' = Z^{\mathrm{t}}Q^{-1}Z$. Verify first that its mean $\mathrm{Tr}(Q^{-1}\Sigma_{11})$ is identical to the formula found above for $\mathrm{E}\,W$. Work out that

$$\begin{aligned}W' &= [QV + (J^{10} + QK_{10}K_{00}^{-1})U_0]^{\mathrm{t}}Q^{-1}[QV + (J^{10} + QK_{10}K_{00}^{-1})U_0]\\ &= V^{\mathrm{t}}QV + U_0^{\mathrm{t}}(J^{01} + K_{00}^{-1}K_{01}Q)Q^{-1}(J^{10} + QK_{10}K_{00}^{-1})U_0\\ &\quad + U_0^{\mathrm{t}}(J^{01} + K_{00}^{-1}K_{01}Q)V + V^{\mathrm{t}}(J^{10} + QK_{10}K_{00}^{-1})U_0.\end{aligned}$$

Prove $W = W'$ by by checking the separate terms. One needs to verify that $A = A'$, in

$$\begin{aligned} A &= J^{00} - J_{00}^{-1} + K_{00}^{-1}K_{01}QK_{10}K_{00}^{-1} + J^{01}K_{10}K_{00}^{-1} + K_{00}^{-1}K_{01}J^{10}, \\ A' &= (J^{01} + K_{00}^{-1}K_{01}Q)Q^{-1}(J^{10} + QK_{10}K_{00}^{-1}). \end{aligned}$$

(xx nils cleans and checks all of this. xx)

(f). For the special case $J = K$, which we meet for ML estimation under model conditions, show that $W \sim \chi^2_q$. For the case $K = cJ$, which turns up in certain overdispersion setups, show that $W \sim c\chi^2_q$.

(g). (xx give an example. can simulate from limit distribution. clarify connections to the case of narrow model $p-1$, wide model p, i.e. profiling over a 1-dimensional $\mu = k(\theta)$. xx)

Ex. 5.27. *ML and minimum divergence function estimators in practice.* In previous exercises we have learned that $\widehat\theta$, the minimiser of $H_n(\theta) = n^{-1}\sum_{i=1}^n h(Y_i,\theta)$, is a natural estimator for θ_0, the minimiser of $\mathrm{E}_G\, h(Y,\theta)$, and that its distribution approaches normality. In order to use such results in practice, for testing, setting confidence intervals, etc., we need to estimate the two crucal matrices $J = \mathrm{E}_G\, h''(Y,\theta_0)$ and $K = \mathrm{Var}_G\, h'(Y,\theta_0)$. These comments apply in particular to ML estimation, with $J = -\mathrm{E}_G\, i(Y,\theta_0)$ and $K = \mathrm{Var}_G\, u(Y,\theta_0)$, see Ex. 5.18.

(a). Consider first, in general terms, some function $p(y,\theta)$ with finite mean in a neighbourhood of the true θ_0, with $\widehat\theta$ an estimator of θ_0. Explain first that $p_n = n^{-1}\sum_{i=1}^n p(Y_i,\theta_0)$ tends to $p_0 = \mathrm{E}_G\, p(Y,\theta_0)$. The best we may do for estimating p_0 in practice is $\widehat p_n = n^{-1}\sum_{i=1}^n p(Y_i,\widehat\theta)$. Show that if $|p(y,\theta_0+\varepsilon) - p(y,\theta_0)| \le M(y)\|\varepsilon\|$, for all small $\|\varepsilon\|$, for some function $M(y)$ with finite mean, then indeed $\widehat p_n \to_{\mathrm{pr}} p_0$.

(b). Give conditions under which the natural estimators

$$\widehat J = n^{-1}\sum_{i=1}^n h''(Y_i,\widehat\theta) \quad\text{and}\quad \widehat K = n^{-1}\sum_{i=1}^n h'(Y_i,\widehat\theta)h'(Y_i,\widehat\theta)^{\mathrm t}$$

are consistent for J and K. Deduce that when such hold, the empirical sandwich matrix $\widehat\Sigma = \widehat J^{-1}\widehat K\widehat J^{-1}$ is consistent for $\Sigma = J^{-1}KJ^{-1}$. Note also that $\widehat J = H_n''(\widehat\theta)$ is the Hessian matrix of the criterion function H_n, often computed directly when using numerical minimisation methods for finding $\widehat\theta$ in the first place. (xx nils emil small note: where do we say that $A_n \to_{\mathrm{pr}} A$ and $B_n \to_{\mathrm{pr}} B$ for matrices implies $A_nB_n \to_{\mathrm{pr}} AB$, etc.? xx)

(c). Explain how confidence intervals for the components of θ may be read off from this. More generally, for any focus parameter $\mu = k(\theta)$, with estimator $\widehat\mu = k(\widehat\theta)$, show that $\widehat\mu \pm 1.96\,\widehat\kappa/\sqrt n$ is an approximate 95 percent interval for μ, where $\widehat\kappa^2 = \widehat c^{\mathrm t}\widehat J^{-1}\widehat K\widehat J^{-1}\widehat c$, and $\widehat c = \partial k(\widehat\theta)/\partial\theta$.

(d). The case of ML estimation is a special case of the setup above, see Ex. 5.18. Show that the estimated matrices become

$$\widehat J = -n^{-1}\sum_{i=1}^n \frac{\partial^2 \log f(Y_i,\widehat\theta)}{\partial\theta\partial\theta^{\mathrm t}} \quad\text{and}\quad \widehat K = n^{-1}\sum_{i=1}^n \widehat u_i\widehat u_i^{\mathrm t},$$

where $\widehat u_i = u(Y_i,\widehat\theta)$. Note also that $\widehat J = -n^{-1}\ell_n''(\widehat\theta)$, the normalised Hessian matrix of the log-likelihood function computed at the ML position.

(e)**.** To illustrate how the above machinery works in practice, simulate 100 points from the standard normal, and then estimate the two normal parameters (ξ,σ) via minimum L_2, as in Ex. 5.5. Explain that this means minimising the empirical distance function

$$H_n(\xi,\sigma)=\int f(y,\xi,\sigma)^2\,\mathrm{d}y-\frac{2}{n}\sum_{i=1}^n f(y_i,\xi,\sigma)=\frac{\frac12/\pi^{1/2}}{\sigma}-\frac{2}{n}\sum_{i=1}^n\frac{1}{\sigma}\phi\Big(\frac{y_i-\xi}{\sigma}\Big).$$

Carry out this minimisation using e.g. `nlm` in R, a non-linear minimisation algorithm, which finds both $(\widehat\xi,\widehat\sigma)$ and the Hessian $\widehat J$. Compute also $\widehat K$, and find confidence intervals for ξ, for σ, and for $p(y_0)=\Pr(Y\ge y_0)$, with say $y_0=1.00$.

(f)**.** Change one or two of your simulated datapoints to somewhat far-off values, e.g. $y_{99}=d$ and $y_{100}=d$, with $d=5.00$ (which indeed is really far off for the standard normal). Observe what then happens to the ordinary ML estimators, and compare with what happens with the minimum L_2 estimators. The point is the the minimum L_2 method is much more robust than the ML method.

(g)**.** For the ML method, we know from the Bartlett identity that the two matrices J and K are equal under model conditions. For various models we would then have two different estimators, with the same aim, and perhaps different precision; this does not matter for the first-order large-sample theory, since consistency is what matter for $\widehat J$ and $\widehat K$, and specifically for the estimated sandwich matrix $\widehat\Sigma=\widehat J^{-1}\widehat K\widehat J^{-1}$. For illustration, consider Y_i from the Pois(θ) model. Show that $J=1/\theta_0$ and $K=\tau_0^2/\theta_0^2$, in terms of the true mean and variance θ_0 and τ_0^2 of the underlying distribution. The recipes above lead to $\widehat J=1/\bar Y$ and $\widehat K=V_n/\bar Y^2$, in terms of sample mean and sample variance. Under model conditions, both aim for $1/\theta_0$. Show also that the consequent estimated sandwich becomes $\widehat\Sigma=V_n$. The general recipes hence lead to two somewhat different confidence intervals for θ, namely $\widehat\theta\pm z_0\widehat\kappa/\sqrt n$, with either $\widehat\kappa=\bar Y^{1/2}$ of $\widehat\kappa=V_n^{1/2}$. Argue that both are valid, with both aiming for the same quantity under Poisson conditions, whereas the second option might be preferred if there might be overdispersion.

(h)**.** (xx profiling too, to illustrate, for $p(y_0)$. again handled by the general theory. intervals need $\kappa=c^{\mathrm t}J^{-1}KJ^{-1}c/c^{\mathrm t}J^{-1}c$. xx)

Ex. 5.28. *BHHJ analysis in practice.* We have seen how the behaviour of the model robust BHHJ estimators, defined in Ex. 5.9, is described via the two crucial matrices J_a and K_a worked with in Ex. 5.19. Here we go into the details of their estimation. The setting is having i.i.d. data $Y_1,\ldots,Y_n$ from some density g, fitted to a parametric $f(y,\theta)$ by minimising $H_{n,a}(\theta)$ of (5.5). We let $\widehat\theta$ be this BHHJ estimator, computed for the given balance parameter a. Below, we write $\widehat u_i=u(Y_i,\widehat\theta)$ and $\widehat f_i=f(Y_i,\widehat\theta)$.

(a)**.** Show that $\widehat\xi_a=n^{-1}\sum_{i=1}^n\widehat f_i^a\widehat u_i$ is equal to $\int f(y,\widehat\theta)^{1+a}u(y,\widehat\theta)\,\mathrm{d}y$, and that it is consistent for the ξ_a defined in Ex. 5.19.

(b)**.** Show next that $\widehat J_a=H''_{n,a}(\widehat\theta)$, the Hessian matrix associated with minimisation of the criterion function, is consistent for the J_a matrix.

(c)**.** Show that

$$\widehat K_a=(1+a)^2\Big\{n^{-1}\sum_{i=1}^n\widehat f_i^{2a}\widehat u_i\widehat u_i^{\mathrm t}-\widehat\xi_a\widehat\xi_a^{\mathrm t}\Big\}=(1+a)^2\,n^{-1}\sum_{i=1}^n(\widehat f_i^a\widehat u_i-\widehat\xi_a)(\widehat f_i^a\widehat u_i-\widehat\xi_a)^{\mathrm t}$$

is consistent for K_a.

Ex. 5.29. *Wilks theorems.* (xx repair and round off. check if these two exercises may be pushed upwards. link to Ch7 things. log-likelihood profiling, deviance functions, Wilks theorems. setup as in Ex. 5.18. we harvest from earlier profiling efforts. xx)

(a). Consider a one-dimensional focus parameter $\mu = k(\theta)$, and the consequent profile log-likelihood function

$$\ell_{n,\text{prof}}(\mu) = \max\{\ell_n(\theta)\colon k(\theta) = \mu\}.$$

Show first that its maximum is reached for $\widehat{\mu} = k(\widehat{\theta})$, so this $\widehat{\mu}$ is also maximising the profile function, and is rightly the ML estimator of μ. With $\mu_0 = k(\theta_0)$, the least false parameter value for μ, write also $c = \partial k(\theta_0)/\partial\theta$ for the gradient vector. Show from Ex. 5.25 that

the deviance function

$$D_n(\mu_0) = 2\{\ell_{n,\max} - \ell_{n,\text{prof}}(\mu_0)\} \to_d (c^{\mathrm{t}}J^{-1}U)^2/c^{\mathrm{t}}J^{-1}c \sim \kappa\chi^2_1,$$

with $\kappa = c^{\mathrm{t}}J^{-1}KJ^{-1}/c^{\mathrm{t}}J^{-1}c$. This $D_n(\mu)$, when computed for a range of μ values, is called the *deviance function* for that parameter. Under model conditions, $D_n(\mu_0) \to_d \chi^2_1$, which is one of several so-called Wilks theorems.

(b). Explain that this Wilks theorem may be used rather directly, without the need to estimate the J matrix, to construct confidence intervals for the focus parameter $\mu = \mu(\theta)$, assuming that the $f(y,\theta)$ model holds. With $C_n(\alpha) = \{\mu\colon \Gamma_1(D_n(\mu)) \le \alpha\}$, writing $\Gamma_1(\cdot)$ for the c.d.f. of the χ^2_1, show in fact that $\Pr_\theta\{\mu \in C_n(\alpha)\} \to \alpha$, for any desired confidence level α. This is a key method for constructing full confidence distributions, the core topic of Ch. 7.

(c). To formulate and prove a Wilks theorem for testing a submodel within a bigger model, via ML estimation, consider $Y_1,\ldots,Y_n$ i.i.d. from a parametric model of the form $f(y,\alpha)$, with $\alpha = (\theta,\gamma)$. We call this the wide model, of dimension $p+q$, with p and q the dimensions of θ and γ, and then study the narrow model, of dimension p, corresponding to $\gamma = \gamma_0$ for some fixed γ_0. Let $\ell_{n,\max,\text{wide}}$ and $\ell_{n,\max,\text{narr}}$ be the log-likelihood maxima for the wide and the narrow models. Show now via Ex. 5.26 that if the narrow model holds, with data arising from $f(y,\theta_0,\gamma_0)$ for some θ_0, that

$$W_n = 2(\ell_{n,\max,\text{wide}} - \ell_{n,\max,\text{narr}}) \to_d W = Z^{\mathrm{t}}Q^{-1}Z.$$

Here $Z \sim \mathrm{N}_q(0,Q)$, with $Q = J^{11}$ the lower right $q\times q$ submatrix of J^{-1}, as in (5.10). Under the narrow model conditions, then, $W_n \to_d \chi^2_q$, another Wilks theorem.

(d). (xx choose a simple illustration, one covered in score function exercise above. may point to a couple of stories. xx)

Ex. 5.30. *Wilks theorems for the normal model.* Consider the familiar normal situation with $Y_1,\ldots,Y_n$ independent from $\mathrm{N}(\xi,\sigma^2)$. The log-likelihood function is $\ell_n(\xi,\sigma)$, as in (5.2). Here we work with profiled consequences, also to illustrate the Wilks theorem of Ex. 5.29.

(a). For σ in focus, show that

$$\ell_{n,\text{prof}}(\sigma) = -n\log\sigma - \tfrac12 Q_0/\sigma^2 - \tfrac12 n\log(2\pi),$$

with $Q_0 = \sum_{i=1}^n (Y_i - \bar{Y})^2$. From this show that the deviance function may be written

$$D_n(\sigma) = 2\{\ell_{n,\max} - \ell_{n,\text{prof}}(\sigma)\} = nh(\widehat{\sigma}^2/\sigma^2),$$

with $h(x) = x - \log x - 1$. The Wilks theorem implies that $D_n(\sigma) \to_d \chi^2_1$, under model conditions; also verify this via a two-term Taylor expansion for $h(x)$ around $x = 1$.

(b). Then with attention to ξ, show that

$$\ell_{n,\text{prof}}(\xi) = -n \log \widehat{\sigma}(\xi) - \tfrac{1}{2} n\{1 + \log(2\pi)\},$$

with $\widehat{\sigma}(\xi)^2 = n^{-1}\{Q_0 + n(\xi - \bar{Y})^2\}$. From this show that

$$D_n(\xi) = 2\{\ell_{n,\max} - \ell_{n,\text{prof}}(\xi)\} = n \log(1 + n^{-1} K_n(\xi)), \quad \text{with } K_n(\xi) = n(\bar{Y} - \xi)^2/\widehat{\sigma}^2.$$

Verify directly what the Wilks theorem says here, that $D_n(\xi) \to_d \chi^2_1$.

(c). The Wilks theorem gives in general terms the recipe $\{\mu\colon D_n(\mu) \le 1.96^2\}$ for an approximate 95 percent confidence interval. With some simulated data, with sample sizes say $10, 30, 50$, check what this recipe gives for forming intervals for ξ and for σ, compared to the exact ones. (xx can also say, do this, exact and Wilks, for the $x_{i,E}$ data of Story ii.2, $n = 64$ word-length averages. xx)

(d). Importantly, the Wilks theorem with consequent confidence recipe can be used for *any* smooth function of the parameters here, i.e. not merely for ξ and for σ. With some simulated data, illustrate this, by computing the log-likelihood profile function and the deviance function, for parameters $\gamma = \xi + 1.96\,\sigma$ and $p = 1 - \Phi((y_0 - \xi)/\sigma)$, for suitable threshold y_0.

Ex. 5.31. *log-likelihood and ML for the binomial, trinomial, multinomial.* Working through the likelihood mechanics of the binomial, trinomial, multinomial models provide good illustrations of the developed methodology. We already know the basic large-sample behaviour for the natural estimators in these models, via Ex. 2.49, but here we connect such results to the general likelihood theory. See also Story vii.1.

(a). For $X \sim \text{binom}(n, p)$, a sum of independent Bernoulli variables, show that its log-likelihood function is $\ell_n(p) = X \log p + (n - X) \log(1 - p)$, with maximiser $\widehat{p} = X/n$. Show that $nJ = n/\{p(1-p)\}$, here with J defined relative to a single Bernoulli variable. Explain that we may hence read off the limit distribution of $\sqrt{n}(\widehat{p} - p)$ being $\text{N}(0, p(1-p))$, without necessarily even knowing the X/n formula, or about the CLT for binomials. One may even argue the other way, starting with the X/n formula: since we know the limit is $\text{N}(0, p(1-p))$, via the CLT, we must have $J^{-1} = p(1-p)$, hence J must be $1/\{p(1-p)\}$.

(b). For further illustration of the binomial likelihood mechanics, draw the log-likelihood function $\ell_n(p)$, computing both its maximiser, its maximum, and its second derivative at the maximum, for cases (i) $n = 20, y = 4$, (ii) $n = 40, y = 8$, (iii) $n = 200, y = 40$. Note how $\widehat{J}_{\text{obs}} = -\ell_n''(\widehat{p})$ becomes bigger and the implied estimated standard deviance $1/\widehat{J}_{\text{obs}}^{1/2}$ becomes smaller.

(c). Let (X, Y) be trinomial (n, p, q), and use $r = 1 - p - q$ and $Z = n - X - Y$. Show that the log-likelihood function becomes $X \log p + Y \log q + Z \log(1 - p - q)$, and for the Fisher information matrix that

$$J = \begin{pmatrix} 1/p + 1/r & 1/r \\ 1/r & 1/q + 1/r \end{pmatrix} \quad \text{with} \quad J^{-1} = \begin{pmatrix} p(1-p) & -pq \\ -pq & q(1-q) \end{pmatrix}.$$

Conclude, even before finding or perhaps without caring that there are clear formulae $\widehat{p} = X/n$ and $\widehat{q} = Y/n$ for the ML estimators, that $(\sqrt{n}(\widehat{p}-p), \sqrt{n}(\widehat{q}-q))$ tends to the binormal zero-mean distribution with the J^{-1} covariance matrix. This may of course also be shown, for $(X/n, Y/n)$, without knowing that these are ML estimators.

(d). Generalise the above to the full multinomial model, with data $(X_1, \ldots, X_k)$ with sum n and probabilities $p_1, \ldots, p_k$ summing to 1 over k boxes. The model has $k-1$ parameters, since p_k is known when $(p_1, \ldots, p_{k-1})$ is known. Find the J and J^{-1} matrices, of size $(k-1) \times (k-1)$.

Ex. 5.32. *The maximum likelihood estimator: examples.* Here we work through some examples, where the task is to set up the log-likelihood function and if feasible also explicit formulae for the ML estimators.

(a). With $Y_1, \ldots, Y_n$ from the normal model $\mathrm{N}(\mu, \sigma^2)$, write down the log-likelihood function. Find the ML estimator for σ when μ is a known value, and find also the ML estimators $(\widehat{\mu}, \widehat{\sigma})$ in the case where both parameters are unknown.

(b). Suppose $Y_i \sim \mathrm{Pois}(w_i\theta)$, with known exposure times w_i, and that the observations are independent, for $i = 1, \ldots, n$. Write down the log-likelihood function, find the ML estimator, and find its mean and variance.

(c). (xx one or two more, with explicit formulae for ML. xx)

(d). Let $Y_1, \ldots, Y_n$ be i.i.d. from the uniform model on $[0, \theta]$, with θ the unknown end-point. Set up the likelihood function and find the ML estimator.

Ex. 5.33. *Maximum likelihood for the Beta and Gamma models.* Consider the Beta and Gamma two-parameter models, with densities

$$\mathrm{be}(y, a, b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} y^{a-1}(1-y)^{b-1} \text{ and } g(y, a, b) = \frac{b^a}{\Gamma(a)} y^{a-1} \exp(-by),$$

for $y \in (0,1)$ and $y > 0$, respectively. The task is in each case to estimate the parameters based on an i.i.d. sample $Y_1, \ldots, Y_n$.

(a). We start with the Beta distribution, see Ex. 1.19, where we in particular have found formulae for the mean and variance in terms of (a, b). With empirical mean and variance $\bar{y}$ and $\widehat{\sigma}^2$, show how (a, b) can be fitted by solving the two equations $\bar{y} = \mathrm{E}\, Y$ and $\widehat{\sigma}^2 = \mathrm{Var}\, Y$. With solutions $(\widehat{a}_m, \widehat{b}_m)$ for these moment estimators, and assuming the Beta model is correct, explain how you can find limit distributions of $\sqrt{n}(\widehat{a}_m - a, \widehat{b}_m - b)$.

(b). Write down the log-likelihood function, say $\ell_n(a, b)$. Show that the ML estimators $(\widehat{a}, \widehat{b})$ are the solutions to the two equations

$$n^{-1} \sum_{i=1}^{n} \log Y_i = \psi(a) - \psi(a+b), \quad n^{-1} \sum_{i=1}^{n} \log(1-Y_i) = \psi(b) - \psi(a+b),$$

where $\psi(x) = \partial \log \Gamma(x)/\partial x$ is the so-called digamma function. There are no explicit formulae here, but the two equations may be easily solved numerically. Explain how the limit distribution for $\sqrt{n}(\widehat{a} - a, \widehat{b} - b)$ may be found, via the two-dimensional central limit theorem. – Here it turns out (i) that the ML estimators are more precise than the moment estimators, and (ii) that finding the limit distribution is rather easier via the general results about ML behaviour, already sorted out in Ex. 5.18.

(c). Then turn attention to the two-parameter Gamma model, where arguments and results will be similar. Show that the mean and variance are a/b and a/b^2, and find moment estimators $\widehat{a}_m, \widehat{b}_m$ based on this. Find the limit distribution of $\sqrt{n}(\widehat{a}_m - a, \widehat{b}_m - b)$ using Ex. 2.68.

(d). Then write down the log-likelihood function, take derivatives, and show that the ML estimators $(\widehat{a}, \widehat{b})$ are the solutions to the two equations

$$\bar{Y} = a/b, \quad n^{-1}\sum_{i=1}^{n} \log Y_i = \psi(a) - \log b.$$

Explain how the limit distribution of $\sqrt{n}(\widehat{a} - a, \widehat{b} - b)$ may be obtained. Again, this is rather easier via the general recipes worked out in Ex. 5.18; also, we shall again find that ML estimators under model conditions are more precise than the moment estimators.

Ex. 5.34. *Log-linear mixing of densities.* How far have we moved, from A to B? Suppose i.i.d. data $Y_1, \ldots, Y_n$ come from a density which is percieved of as being 'betweeen' given densities f_A and f_B. There are various ways of creating probabilistic bridges from A to B, via which one then can estimate the position of the current density, e.g. to assess whether it is close to A or to B. One such model takes $f(y, \lambda) = f_A(y)^{1-\lambda} f_B(y)^{\lambda}/R(\lambda)$, with $R(\lambda)$ the normalisation constant $\int f_A^{1-\lambda} f_B^{\lambda}\, \mathrm{d}y$. This creates such a bridge, with $\lambda \in [0, 1]$, and with A and B corresponding to $\lambda = 0$ and $\lambda = 1$.

(a). Explain that $R(0) = R(1) = 1$, and that the derivative of $R(\lambda)$, at the endpoints 0 and 1, can be written

$$R'(0) = -\mathrm{KL}(f_A, f_B), \quad R'(1) = \mathrm{KL}(f_B, f_A).$$

in terms of the Kullback–Leibler distances discussed in Ex. 5.7 and later on. In particular, the $R(\lambda)$ has negative derivative at the start and positive derivative at the end.

(b). For a concrete illustration, suppose f_A is $\mathrm{N}(0, 1)$ and f_B is $\mathrm{N}(c, 1)$, with c positive. Show that $R(\lambda) = \exp\{-\frac{1}{2}c^2\lambda(1-\lambda)\}$. From this explain that the score function is $u(y, \lambda) = cy - \lambda c^2$ and that the Fisher information becomes the constant c^2. Show indeed that $f(y, \lambda)$ is $\mathrm{N}(\lambda c, 1)$, that the ML estimator becomes $\widehat{\lambda} = \bar{Y}/c$, truncated to $[0, 1]$. From general ML theory, explain that if the real λ is inside $(0, 1)$, then $\sqrt{n}(\widehat{\lambda} - \lambda) \to_d \mathrm{N}(0, 1/c^2)$. For c small the variance will be very big, signifying that it is a hard task to estimate the balance parameter when f_A and f_B are close.

(c). Returning to the general setup, write $S(y) = \log f_B(y) - \log f_A(y)$. Show that the log-likelihood function becomes

$$\ell_n(\lambda) = n\{\lambda \bar{S}_n - \log R(\lambda)\}, \quad \text{with } \bar{S}_n = n^{-1}\sum_{i=1}^{n} S(y_i).$$

Explain that the ML estimator $\widehat{\lambda}$ hence is the solution to $\bar{S}_n = R'(\lambda)/R(\lambda)$. (xx a bit more. and example. limiting normality. evaluate for $\lambda = 0$ and $\lambda = 1$. xx)

(d). Bayes posterior: $\pi(\lambda) \exp[n\{\lambda S_n - \log R(\lambda)\}]$.

(e). (xx there is perhaps hope for the semiparametric construction $\widehat{f}(y) = f_0(y)^{1-\lambda} f_n(y)^{\lambda}/R_n(\lambda)$, resembling the Hjort–Glad estimator, see Hjort and Glad (1995). xx)

Convex processes, log-concave likelihoods

Ex. 5.35. *Minimisers of convex processes, I.* We have seen useful constructions, methods, and results for minimum divergence function estimators in Ex. 5.5 and 5.25, in particular when applied to ML estimators, as with the general apparatus of Ex. 5.18. There are issues worth refining and generalising, however. The regularity conditions required for the Taylor expansion based arguments to go fully through are a bit cumbersome, and there are important constructions where the distance function $h(y,\theta)$ in $H(\theta) = \mathrm{E}_G\, h(Y,\theta)$ is not smooth. Here we give the basics for how matters simplify, with weaker conditions, if the distance function is convex.

(a). From pointwise to uniform: Suppose $A_n(s)$ is a sequence of convex random functions defined on an open convex set $\mathcal{S}$ of $\mathbb{R}^p$, which convergences in probability to some $A(s)$, for each $s \in \mathcal{S}$. Show that the convergence is automatically uniform; $\max_{s\in\mathcal{S}} |A_n(s) - A(s)| \to_{\mathrm{pr}} 0$.

(b). Nearness of argmins: Suppose $A_n(s)$ is convex and is approximated by $B_n(s)$. Let α_n and β_n be the argmins of A_n and B_n. Then there is a probabilistic bound on how far these minimisers can be from each other: show that

$$\Pr(\|\alpha_n - \beta_n\| \ge \delta) \le \Pr\{\Delta_n(\delta) \ge \tfrac12 h_n(\delta)\},$$

in which

$$\Delta_n(\delta) = \sup_{\|s-\beta_n\|\le\delta} |A_n(s) - B_n(s)| \quad \text{and} \quad h_n(\delta) = \inf_{\|s-\beta_n\|=\delta} B_n(s) - B_n(\beta_n).$$

(c). Basic corollary: Suppose $A_n(s)$ is convex and can be represented as $\frac12 s^{\mathrm{t}} J s + U_n^{\mathrm{t}} s + C_n + r_n(s)$, where J is symmetric and positive definite, U_n is stochastically bounded, C_n is arbitrary, and $r_n(s) \to_{\mathrm{pr}} 0$ for each s. For the approximation $B_n(s) = \frac12 s^{\mathrm{t}} J s + U_n^{\mathrm{t}} s + C_n$, show that $\beta_n = -J^{-1}U_n$ is its argmin. Then demonstrate that their minimisers as well as their minima are close. Specifically, show (i) that $\alpha_n - \beta_n \to_{\mathrm{pr}} 0$; and (ii) that $A_{n,\min} - B_{n,\min} \to_{\mathrm{pr}} 0$.

(d). Show that if in addition $U_n \to_d U$, then $\alpha_n \to_d -J^{-1}U$, and that $B_{n,\min} - C_n$ as well as $A_{n,\min} - C_n$ tend to $-\frac12 U^{\mathrm{t}} J^{-1} U$. These two statements are what we worked hard for in Ex. 5.13, 5.26 see (5.6)–(5.9), now obtained in a simpler fashion and with weaker smoothness assumptions, though bought with the extra convexity condition.

(e). Prove the following modest but useful generalisation of the above: the statements continue to hold if a random matrix J_n replaces V, provided $J_n \to_{\mathrm{pr}} J$.

Ex. 5.36. *Minimisers of convex processes, II.* The framework worked with now is as in Ex. 5.5 and 5.25, with $Y_1, \ldots, Y_n$ being i.i.d. from some G, a possibly multidimensional parameter θ_0 defined as the minimiser of $H(\theta) = \int h(y,\theta)\,\mathrm{d}G(y)$, with estimator $\widehat\theta$ the minimiser of the distance function $H_n(\theta) = \int h(y,\theta)\,\mathrm{d}G_n(y) = n^{-1}\sum_{i=1}^n h(Y_i,\theta)$. Here we put in one more condition, however, that $h(y,\theta)$ is convex in θ. The point is that this both simplifies various technical arguments, via methods of Ex. 5.35, and allows for nonsmooth distance functions.

(a). With $h(y,\mu) = |y-\mu|$, show that $\mu_0 = \mathrm{med}(G)$, the median, with $\widehat\mu = M_n$, the sample median. More generally, for some $q \in (0,1)$, consider

$$h_q(y,\mu) = q(y-\mu)_+ + (1-q)(\mu-y)_+ = \begin{cases} q(y-\mu) & \text{if } y \ge \mu, \\ (1-q)(\mu-y) & \text{if } y \le \mu. \end{cases}$$

Show that $\mu_0 = G^{-1}(q)$, the q quantile.

(b). For an $\alpha \geq 1$, consider the parameter θ_0 being the minimiser of $\mathrm{E}_G\,|Y-\theta|^\alpha$. Show that the distance function indeed is convex, and that the special cases $\alpha = 1$ and $\alpha = 2$ correspond to the median and the mean, respectively.

(c). We now work through regularity conditions ensuring control over the behaviour of such estimators. Part of the point is that we avoid needing smooth derivatives in θ. Suppose that

$$h(y,\theta_0+\varepsilon) - h(y,\theta_0) = D(y)^{\mathrm{t}}\varepsilon + R(y,\varepsilon),$$

for a $D(y)$ with mean zero under G, and that

$$\mathrm{E}\,\{h(Y,\theta_0+\varepsilon) - h(Y,\theta_0)\} = \mathrm{E}\,R(Y,\varepsilon) = \tfrac{1}{2}\varepsilon^{\mathrm{t}}J\varepsilon + o(\|\varepsilon\|^2) \quad \text{as } \varepsilon \to 0$$

for a positive definite J. Assume furthermore that the variance matrix $K = \mathrm{Var}_G\, D(Y)$ is finite and that $\mathrm{Var}\,R(Y,\varepsilon) = o(\|\varepsilon\|^2)$. Show that $\sqrt{n}(\widehat{\theta}-\theta_0) = -J^{-1}(1/\sqrt{n})\sum_{i=1}^n D(Y_i) + o_{\mathrm{pr}}(1)$. In particular, it tends to $\mathrm{N}_p(0, J^{-1}KJ^{-1})$. Show also that

$$W_n(\theta_0) = 2n\{H_n(\theta_0) - H_n(\widehat{\theta})\} \to_d U^{\mathrm{t}}J^{-1}U,$$

and explain how this may be used to find a confidence region for θ_0.

(d). The median: Suppose $Y_1,\ldots,Y_n$ are i.i.d. from a distribution G with a density g positive at the median μ. For the median distance function $|y-\mu|$, show that

$$|y-(\mu+\varepsilon)| - |y-\mu| = D(y)\varepsilon + R(y,\varepsilon),$$

with $D(y) = I(y \leq \mu) - I(y > \mu)$ and

$$R(y,\varepsilon) = \begin{cases} 2\{\varepsilon - (y-\mu)\}\, I(\mu \leq y \leq \mu+\varepsilon) & \text{if } \varepsilon > 0, \\ 2\{(y-\mu)-\varepsilon\}\, I(\mu+\varepsilon \leq y \leq \mu) & \text{if } \varepsilon < 0, \end{cases}$$

with $R(y,0) = 0$. Verify from this that $\mathrm{E}\,R(Y,\varepsilon) = g(\mu)\varepsilon^2 + o(\varepsilon^2)$ and $\mathrm{E}\,R(Y,\varepsilon)^2 = (4/3)g(\mu)|\varepsilon|^3 + o(|\varepsilon|^3)$. Deduce that $\sqrt{n}(M_n-\mu) \to_d \mathrm{N}(0, 1/\{4g(\mu)^2\})$.

(e). Generalise to the case of $\mu_q = G^{-1}(q)$, with $Q_{n,q} = G_n^{-1}(q)$ the empirical q quantile. Show in fact that

$$Z_n(q) = \sqrt{n}(Q_{n,q}-\mu_q) = -g(\mu_q)^{-1}\sqrt{n}\{G_n(\mu_q)-q\} + \varepsilon_n(q),$$

where $\varepsilon_n(q) \to_{\mathrm{pr}} 0$ for each q. Derive from this that the limit is a $\mathrm{N}(0, q(1-q)/g(\mu_q)^2)$. We saw this in Ex. 3.21, but here the technique is distinctly different, and we find an interesting representation in terms of the empirical $G_n(\mu_q)$.

(f). Let ξ_α be the minimiser of $\mathrm{E}\,|Y_i-\xi|^\alpha$, with estimator $M_{n,\alpha}$ minimising $\sum_{i=1}^n |Y_i-\xi|^\alpha$. Show that

$$\sqrt{n}(M_{n,\alpha}-\xi_\alpha) \to_d \mathrm{N}(0,\tau_\alpha^2) \quad \text{with } \tau_\alpha^2 = \frac{\mathrm{E}\,|Y-\xi_\alpha|^{2(\alpha-1)}}{\{(\alpha-1)\,\mathrm{E}\,|Y-\xi_\alpha|^{\alpha-2}\}^2}.$$

Explain that if the distribution is symmetric, then the ξ_α is the same for each α. Compute and display τ_α, for $\alpha \in [1,2]$, i.e. the range from median to mean, for the normal and for the Laplace,

Ex. 5.37. *Maximum likelihood asymptotics with log-concave likelihood.* (xx edit with care. regularity condition on R. point to expofamily in later exercise. xx) Suppose $Y_1, \ldots, Y_n$ are i.i.d. from a density g, modelled as $f(y, \theta)$, where $\log f(y, \theta)$ is concave in θ for each y; the $\theta = (\theta_1, \ldots, \theta_p)$ is allowed to be multidimensional. Let $\ell_n(\theta)$ be the log-likelihood, with ML estimator $\widehat{\theta}$. From efforts of Ex. 5.18 we know $\widehat{\theta}$ is consistent for the least false parameter θ_0, the KL minimiser, assumed here to be an inner point in the parameter space. log-concave density

(a). Show first that ℓ_n is concave. To start with mild conditions, assume merely that

$$\log f(y, \theta_0 + h) - \log f(y, \theta_0) = h^{\mathrm{t}} D(y) + R(y, h), \tag{5.11}$$

for a $D(y)$ with mean zero, and that $\mathrm{E}\, R(Y, h) = \frac{1}{2} h^{\mathrm{t}} J h + o(\|h\|^2)$ as $h \to 0$ for some positive definite J, along with $\mathrm{Var}\, R(Y, \varepsilon) = o(\|h\|^2)$. Show via the convexity driven methods of Ex. 5.35 and 5.36 that $\sqrt{n}(\widehat{\theta} - \theta_0) = J^{-1} U_n + o_{\mathrm{pr}}(1)$, where $U_n = (1/\sqrt{n}) \sum_{i=1}^n u(Y_i, \theta_0)$. Here $U_n \to_d U \sim \mathrm{N}_p(0, K)$, with $K = \mathrm{Var}_g\, u(Y, \theta_0)$ assumed finite. Deduce the two fundamental results

$$\begin{aligned} Z_n &= \sqrt{n}(\widehat{\theta} - \theta_0) \to_d \mathrm{N}_p(0, J^{-1} K J^{-1}), \\ W_n &= 2\{\ell_{n,\max} - \ell_n(\theta_0)\} \to_d U^{\mathrm{t}} J^{-1} U, \end{aligned} \tag{5.12}$$

with simplifications under model conditions. Show also that consistency of the ML estimator for θ_0 is a simple consequence of the first statement. We have seen such results before, see Ex. 5.18 and in particular (5.8); the point is that they have now been derived more simply and under weaker conditions, as long as there is log-density concavity.

(b). Work through the details for the case of the Laplace distribution with $f(y, \theta) = \frac{1}{2} \exp(-|y - \theta|)$. The point is that with log-concavity, we reach the required results of (5.12) even without needing full smoothness in the parameter; here the score function is not defined for all values, etc. Usually, though, the $D(y)$ is the score function $u(y, \theta_0) = \partial \log f(y, \theta_0)/\partial\theta$ and $J = J(\theta_0)$ is the variance matrix of this score function, i.e. the Fisher information matrix, evaluated at the true position in the parameter space.

(c). (xx a couple of illustrations. for each, find the limit distribution of $\sqrt{n}(\widehat{\theta} - \theta_0)$. the poisson. gamma. beta. normal. also something like $Y_i \sim \mathrm{Pois}(e_i \theta)$, with different exposures w_i, the point is non-i.i.d. nils does this after mild extra editing for Ex. 5.35 and 5.36, with a bit of Lindeberg too. xx)

(d). (xx to polish. xx) something nontrivial. can do

$$f(y, \theta) = (1/k) \exp\{(y - \theta) \arctan(y - \theta)\}/\{1 + (y - \theta)^2\}^{1/2},$$

which has log-density $(y - \theta) \arctan(y - \theta) - \frac{1}{2} \log\{1 + (y - \theta)^2\}$ and nice score function $\arctan(y - \theta)$. something nontrivial II. and $f = (1/k) \exp(-|y - \theta|^{1.5})$.

Ex. 5.38. *Differentiability in quadratic mean.* The key to proving asymptotic normality of the ML estimator for the log-concave densities in Ex. 5.37 was the assumption that $\log f(y, \theta_0 + h) - \log f(y, \theta_0) = h^{\mathrm{t}} D(y) + R(y, h)$, where $D(y)$ and $R(y, h)$ satisfy the conditions specified in (a) of that exercise. This raises the question of what conditions are needed for such an expansion of the log-likelihood ratio to hold. As above $u(y, \theta) = \partial \log f(y, \theta)/\partial\theta$, and $i(y, \theta) = -\partial^2 \log f(y, \theta)/\partial\theta\partial\theta^{\mathrm{t}}$, and $J = -\mathrm{E}\, i(Y, \theta_0)$ and $K = \mathrm{Var}\, u(Y, \theta_0)$, where θ_0 denotes the true parameter value.

(a). Assume that $\theta \mapsto \log f(y,\theta)$ is two times continuously differentiable for every y; that $\mathrm{E}\, u(Y,\theta_0) = 0$ and that every element of $-\partial^2 \log f(y,\theta)/\partial\theta\partial\theta^{\mathrm{t}}$ is bounded by an integreable function $k(y)$ not depending on θ. Show that $\log f(y,\theta_0+h) - \log f(y,\theta_0) = h^{\mathrm{t}}u(y,\theta_0) + R(y,h)$, where $R(y,h)$ satisfies the conditons of Ex. 5.37(a) .

(b). With $Y_1,\ldots,Y_n$ an i.i.d. sample from $f(y,\theta_0)$, show that the conditions in (a) entails the $A_n(h) = \sum_{i=1}^n \log f(Y_i,\theta_0+h/\sqrt{n})/f(Y_i,\theta_0) = \ell_n(\theta_0+h/\sqrt{n}) - \ell_n(\theta_0)$ can be expressed as

$$A_n(h) = h^{\mathrm{t}}U_n - \tfrac{1}{2}h^t Jh + o_{\mathrm{pr}}(1). \tag{5.13}$$

where $U_n = n^{-1/2}\sum_{i=1}^n u(Y_i,\theta_0)$

(c). The representation in (5.13) can be derived under a weaker condition that only involves the first derivative of $\log f(y,\theta)$, namely differentiability in quadratic mean. Let $\mathcal{P} = \{P_\theta \colon \theta \in \Theta\}$ be a family of distributions, all with densities $f(y,\theta)$ with respect to some dominating measure μ. We say that $\mathcal{P}$ is differentiable in quadratic mean at θ_0 if

differentiability in quadratic mean

$$\int \bigl(f(y,\theta_0+h)^{1/2} - f(y,\theta_0)^{1/2} - \tfrac{1}{2}h^{\mathrm{t}}u(y,\theta_0)f(y,\theta_0)^{1/2}\bigr)^2 \,\mathrm{d}\mu(y) = o(\|h\|^2), \tag{5.14}$$

as $h \to 0$. In this and the following five exercises we will see that this differentiability in quadratic mean assumption entails that the log-likelihood ratio $A_n(h)$ has the representation given in (5.13). Show that $\mathrm{E}\, u(Y,\theta_0) = 0$ and that K is finite.

(d). Define $W_{n,i} = 2\bigl\{f(y,\theta+n^{-1/2}h)^{1/2}/f(y,\theta)^{1/2} - 1\bigr)$, then verify that $A_n(h) = 2\sum_{i=1}^n \log(\frac{1}{2}W_{n,i}+1)$, and show that by a Taylor expansion with the remainder on mean value form

$$A_n(h) = \sum_{i=1}^n W_{n,i} - \tfrac{1}{4}\sum_{i=1}^n W_{n,i}^2 + \tfrac{1}{4}\sum_{i=1}^n W_{n,i}^2 R(W_{n,i}),$$

where $R(2w) = 1 - \{1+\tilde{w}(w)\}^{-2}$, and $\tilde{w}(w)$ is a point between 0 and w.

(e). We start by establishing some useful facts. Show that $n\,\mathrm{E}\, W_{n,i} = -\frac{1}{4}h^{\mathrm{t}}Kh$, that $\mathrm{E}\, W_{n,i}^2 = -4\,\mathrm{E}\, W_{n,i}$, and that $\mathrm{E}\,(W_{n,i} - n^{-1/2}h^{\mathrm{t}}u(Y_i,\theta_0))^2 \le 4D(h/\sqrt{n})$, where $D(h)$ is the expression on the left in (5.14). Deduce that

$$\mathrm{E}\,|W_{n,i}^2 - n^{-1}(h^{\mathrm{t}}u(Y_i,\theta_0))^2| \le n^{-1/2}4\{D(h/\sqrt{n})h^{\mathrm{t}}Kh\}^{1/2}.$$

(f). Define the mean zero random variables $\xi_{n,i} = W_{n,i} - n^{-1/2}h^{\mathrm{t}}u(Y_i,\theta_0) - \mathrm{E}\, W_{n,i}$, and verify, using the results from (e), that

$$\sum_{i=1}^n W_{n,i} = h^{\mathrm{t}}U_n - \tfrac{1}{4}h^{\mathrm{t}}Kh + \sum_{i=1}^n \xi_{n,i}.$$

Show also that $\sum_{i=1}^n \xi_{n,i} = o_{\mathrm{pr}}(1)$.

(g). Define the random variables $\eta_{n,i} = W_{n,i}^2 - (n^{-1/2}h^{\mathrm{t}}u(Y_i,\theta_0))^2$, as well as the mean zero random variables $\zeta_{n,i} = h^{\mathrm{t}}Kh - (h^{\mathrm{t}}u(Y_i,\theta_0))^2$. Use the facts deduced in (e), along with the law of large numbers, to show that

$$\sum_{i=1}^n W_{n,i}^2 = h^{\mathrm{t}}Kh - \sum_{i=1}^n \eta_{n,i} + n^{-1}\sum_{i=1}^n \zeta_{n,i} = h^{\mathrm{t}}Kh + o_{\mathrm{pr}}(1).$$

This means that at this point we have $A_n(h) = h^{\mathrm{t}}U_n - \frac{1}{2}h^{\mathrm{t}}Kh + \sum_{i=1}^n W_{n,i}^2 R(W_{n,i}) + o_{\mathrm{pr}}(1)$. It remains to show that $\sum_{i=1}^n W_{n,i}^2 R(W_{n,i}) = o_{\mathrm{pr}}(1)$.

(h). Add and subtract $n^{-1}(h^{\rm t}u(Y_1,\theta_0))^2$ from $W_{n,1}^2$ to show that

$$\Pr(W_{n,1}^2\geq 2\varepsilon^2)\leq \frac{\|h\|}{n\varepsilon^2}\mathrm{E}\,\|u(Y_1,\theta_0)\|^2 I(\|u(Y_1,\theta_0)\|\geq \frac{\sqrt{n}\varepsilon}{\|h\|})+\frac{4}{\varepsilon^2\sqrt{n}}\{D(h/\sqrt{n})h^{\rm t}Kh\}^{1/2},$$

for each $\varepsilon>0$; and deduce from this expression that $\max_{1\leq i\leq n}|W_{n,i}|\to_{\rm pr}0$ as $n\to\infty$. Finally, show that $\max_{1\leq i\leq n}|W_{n,i}|=o_{\rm pr}(1)$ implies that $\sum_{i=1}^n W_{n,i}^2R(W_{n,i})=o_{\rm pr}(1)$.

(i). [xx an exercise with conditions ensuring that a family of distributions is dqm xx]

Ex. 5.39. *The exponential family class and ML.* The ML and log-likelihood machinery works particularly well for the exponential family class of models, see Ex. 1.55, 4.20, **??**. This is due the log-linear structure for parameters and sufficient statistics, and also to the consequent log-concavity of densities. Consider i.i.d. data $Y_1,\ldots,Y_n$ from a density g, modelled with the generic parametric $f(y,\theta)=\exp\{\theta^{\rm t}T(y)-k(\theta)\}h(y)$, with $\theta^{\rm t}T(y)=\theta_1T_1(y)+\cdots+\theta_pT_p(y)$ involving basic data functions $T_1,\ldots,T_p$.

(a). Show that the score function is $u(y,\theta)=T(y)-\xi(\theta)$, with $\xi(\theta)=\partial k(\theta)/\partial\theta$. Also, the information function becomes $i(y,\theta)=-J(\theta)$, with $J(\theta)=\partial^2k(\theta)/\partial\theta\,\partial\theta^{\rm t}$.

(b). Under model conditions, show that $\mathrm{E}_\theta\,T(Y)=\xi(\theta)$ and $\mathrm{Var}_\theta\,T(Y)=J(\theta)$. Outside model conditions, with g not belonging to the f_θ, assume merely that $T(Y)$ has true mean ξ_0 and true variance K. Show that the least false parameter θ_0 minimising $\mathrm{KL}(g,f(\cdot,\theta))$ is characterised by $\xi_0=\xi(\theta_0)$. This may also be written $\theta_0=M(\xi_0)$, with M the inverse map. Write also $J=J(\theta_0)$, the Fisher information matrix computed at the least false parameter position.

(c). Show that the log-likelihood function becomes $\ell_n(\theta)=n\{\theta^{\rm t}\bar T-k(\theta)\}$, in terms of sample averages $\bar T=(\bar T_1,\ldots,\bar T_p)^{\rm t}$, and that it is concave. Deduce that the ML estimator has $\xi(\widehat\theta)=\bar T$, of $\widehat\theta=M(\bar T)$.

(d). We have $U_n=\sqrt{n}(\bar T-\xi_0)\to_d U\sim \mathrm{N}_p(0,K)$, by the CLT. For the basic log-likelihood function process, show that

$$\begin{aligned}A_n(s)&=\ell_n(\theta_0+s/\sqrt{n})-\ell_n(\theta_0)=\sqrt{n}\bar T^{\rm t}s-n\{k(\theta_0+s/\sqrt{n})-k(\theta_0)\}\\&=U_n^{\rm t}s-\tfrac12 s^{\rm t}J(\theta_0)s+o(\|s\|^3/\sqrt{n})\to_d A(s)=U^{\rm t}s-\tfrac12 s^{\rm t}Js.\end{aligned}$$

This is as with (5.7), but now established rather simply and directly, with a minimum of regularity conditions. By appealing to Ex. 5.12, 5.13, deduce also that

$$\sqrt{n}(\widehat\theta-\theta_0)\to_d \mathrm{N}(0,J^{-1}KJ^{-1}),\quad 2\{\ell_{n,\max}-\ell_n(\theta_0)\}\to_d W=U^{\rm t}J^{-1}U.$$

Under model conditions, the sandwich matrix is $J(\theta_0)^{-1}$ and $W\sim\chi^2_p$. Also, the Wilks theorems of Ex. 5.29 hold, for focus parameters and for testing submodels.

(e). (xx a bit more. xx) From $\widehat\theta_j=M_j(\bar T)$, for components $M_1,\ldots,M_p$ of the mapping $\theta_0=M(\xi_0)$, write

$$\widehat\theta_j=M_j(\xi_0)+M_j'(\xi_0)^{\rm t}(\bar T-\xi_0)+\tfrac12(\bar T-\xi_0)^{\rm t}M_j''(\xi_0)(\bar T-\xi_0)+O_{\rm pr}(1/n^{3/2}).$$

Show that this implies $\mathrm{E}\,\widehat\theta_j=\theta_{0,j}+\frac12c_j/n+o(1/n)$, with $c_j=\mathrm{Tr}(M_j''(\xi_0)K)$.

(f). (xx a simple example here. xx)

Ex. 5.40. *ML asymptotics under model conditions: applications.* (xx cleaning required, and more of the directly useful up front; we estimate parameters and have confidence, almost automatically, via normality and delta method. xx) Results reached in Ex. 5.18 are central in applied statistics. The versatile ML machinery allows the statistician to construct good estimators for even complicated functions of parameters in new models, and to supplement these estimators with confidence intervals, tests, etc. We will also see that the general ML asymptotics results may be used to verify what we already knew, so to speak, regarding estimators in the more familiar models. (xx check all this to make sure that we don't become repetitive. xx)

(a). Let Y be binomial (n, p). Even before you find a formula for the ML estimator for p, show that $\sqrt{n}(\widehat{p} - p) \to \mathrm{N}(0, p(1-p))$. By all means, show also that $\widehat{p} = Y/n$.

(b). In a similar vein, study the classic case of $Y_1, \ldots, Y_n$ being i.i.d. from the normal (ξ, σ^2). Using the Fisher information matrix found in Ex. 5.14, show that the ML estimators $\widehat{\xi}$ and $\widehat{\sigma}$ must be independent, in the limit, with approximate distributions $\mathrm{N}(\xi, \sigma^2/n)$ and $\mathrm{N}(\sigma, \frac{1}{2}\sigma^2)$. Remarkably, these results follow from the ML apparatus even without or before knowing any formulae for the estimators, and without or before knowing any finite-sample theory for these. As we know (xx crossref here xx) there is *exact* independence here, and the distribution for $\widehat{\xi}$ is exactly correct, for each n.

(c). (xx something with $\mathrm{Gam}(a, b)$, and approximate distribution for $\widehat{\mu}$, estimator for the median $\mu = \mu(a, b)$. illustrate also Wilks. which can be used without explicit formulae. xx)

(d). (xx the Weibull $F(t) = 1 - \exp\{-(t/a)^b\}$. perhaps an earlier exercise where we find $J(a, b)$. xx)

(e). Consider random i.i.d. pairs (X_i, Y_i) from the standardised binormal distribution with zero means, unit variances, and correlation ρ. Set up the log-likelihood function $\ell_n(\rho)$, show that the Fisher information becomes $J(\rho) = (1+\rho^2)/(1-\rho^2)^2$, and find the limiting distribution of the ML estimator. How much better is the ML estimator compared to the usual empirical correlation coefficient R_n? (xx kladd to be pushed to solutions follows. xx)

$$\log f(x, y, \rho) = -\tfrac{1}{2}\log(1-\rho^2) - \tfrac{1}{2}(x^2 + y^2 - 2\rho xy)/(1-\rho^2),$$

score function

$$u(x, y, \rho) = \frac{1}{(1-\rho^2)^2}\{\rho - \rho^2 + (1+\rho^2)xy - \rho(x^2+y^2)\},$$

where we may check that the mean is zero. We find $\mathrm{Var}\, XY = 1 + \rho^2$, $\mathrm{Var}\,(X^2 + Y^2) = 4(1+\rho^2)$, $\mathrm{cov}(XY, X^2+Y^2) = 4\rho$, and this leads to Fisher information $J(\rho) = (1+\rho^2)/(1-\rho^2)^2$.

Ex. 5.41. *Examples of agnostic ML operations.* It is useful to go through a list of special cases, to see how the agnostic ML theory pans out in practice. Note that convergence to the normal $\mathrm{N}_p(0, J^{-1}KJ^{-1})$ takes place in general, model after model after model (including those you might invent next week), without any need for working with explicit formulae for the ML estimators etc.

(a). For the exponential model $\theta\exp(-\theta y)$, show that the score function is $u(y,\theta) = 1/\theta - y$, that its least false parameter value is $\theta_0 = 1/\xi_0$, in termas of the true mean $\xi_0 = \mathrm{E}\,Y$. Show that $\sqrt{n}(\widehat{\theta} - \theta_0)$ has limit distribution $\mathrm{N}(0, \sigma_0^2\theta_0^4)$, where σ_0^2 is the true variance. Show that this generalises the 'usual result' derived under model conditions.

(b). Then do the normal: assume data follow some density g, and the normal $\mathrm{N}(\xi,\sigma^2)$ model is used. We already know that the least false parameters are ξ_0 and σ_0, the true mean and standard deviation (i.e. even if g is far from the normal). Assume that the fourth moment is finite, so that skew $= \mathrm{E}\,Z^3$ and kurt $= \mathrm{E}\,Z^4 - 3$ are finite, with $Z = (Y - \mathrm{E}\,Y)/\mathrm{sd}(Y) = (Y-\xi_0)/\sigma_0$. Working with the score function, and the second order derivatives, show that

$$J = \frac{1}{\sigma_0^2}\begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix} \quad \text{and} \quad K = \frac{1}{\sigma_0^2}\begin{pmatrix} 1 & \gamma_3 \\ \gamma_3 & 2+\gamma_4 \end{pmatrix}.$$

(c). For the ML estimators $\widehat{\xi}$ and $\widehat{\sigma}$, show from this that

$$\begin{pmatrix} \sqrt{n}(\widehat{\xi} - \xi) \\ \sqrt{n}(\widehat{\sigma} - \sigma) \end{pmatrix} \to_d \mathrm{N}_2\Big(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \sigma^2 \begin{pmatrix} 1 & \frac{1}{2}\gamma_3 \\ \frac{1}{2}\gamma_3 & \frac{1}{2} + \frac{1}{4}\gamma_4 \end{pmatrix}\Big).$$

Note that this is a 'rediscovery' of what we found in Ex. 2.68 and 2.73, but here we managed to find the limit distribution fully without knowing (or caring) about the exact expressions for the ML estimators.

(d) (xx one more case to come here. xx).

Ex. 5.42. *An average power optimality property.* (xx we shall see how this pans out, and how it can be best told. make connection to BIC of Ch. 11. xx) Suppose $Y_1,\ldots,Y_n$ are i.i.d. from a smooth parametric model $f(y,\theta)$, where we need to test $\theta = \theta_0$, a given value, against $\theta \neq \theta_0$. In general there is no uniformly most powerful test. We have seen in Ex. 4.16, however, that there is a well-defined test maximising the weighted average power $\bar{\pi}_n = \int \pi_n(\theta)\,\mathrm{d}w(\theta)$, with $\pi_n(\theta)$ the power at position θ and $\mathrm{d}w(\theta)$ a given probability measure on the alternative region, here $\theta \neq \theta_0$. This optimal strategy is to use the Neyman–Pearson Lemma for the marginal density $\bar{f}(y_1,\ldots,y_n) = \int \prod_{i=1}^n f(y_i,\theta)\,\mathrm{d}w(\theta)$.

(a). Let $\ell_n(\theta)$ be the log-likelihood function, with $\widehat{\theta}$ the ML estimator, and write also f_0 for the model at the null value θ_0. Show that the Neyman–Pearson ratio can be expressed as

$$\begin{aligned} R_n = \frac{\bar{f}(y_1,\ldots,y_n)}{f_0(y_1,\ldots,y_n)} &= \int \exp\{\ell_n(\theta) - \ell_n(\theta_0)\}\,\mathrm{d}w(\theta) \\ &= \exp\{\ell_n(\widehat{\theta}) - \ell_n(\theta_0)\} \int \exp\{\ell_n(\theta) - \ell_n(\widehat{\theta})\}\,\mathrm{d}w(\theta). \end{aligned}$$

(b). For θ close to $\widehat{\theta}$, use Taylor expansion to get

$$\ell_n(\theta) - \ell_n(\widehat{\theta}) \doteq -\tfrac{1}{2}n(\theta - \widehat{\theta})^{\mathrm{t}} J_n (\theta - \widehat{\theta}),$$

with $J_n = -n^{-1}\partial^2\ell_n(\widehat{\theta})/\partial\theta\,\partial\theta^{\mathrm{t}}$ the normalised Hessian matrix at the max point.

(c). The optimal test consists in rejecting when R_n is above its null distribution threshold. Show that the above leads to

$$\begin{aligned} R_n &\doteq \exp(\tfrac{1}{2}D_n)(2\pi)^{p/2}|nJ_n|^{1/2}w(\widehat{\theta}), \\ 2\log R_n &\doteq D_n - p\log n + \log|J_n| + p\log(2\pi) + 2\log w(\widehat{\theta}). \end{aligned}$$

Here $D_n = 2\{\ell_{n,\max} - \ell_n(\theta_0)\}$ is the Wilks or log-likelihood-ratio test statistic, with its χ^2_p limiting null distribution.

(d). Conclude that the D_n test is an approximation to the maximum averaged power test, almost regardless of the weighting measure.

(e). (xx round this off. example. xx)

Divergences and likelihoods in regression models

Ex. 5.43. *Extending theory and methods to regression setups, I.* Above we have dealt with likelihood methods, involving ML estimation, limit distributions under and outside model conditions, the Wilks theorem for profiled log-likelihoods, broadly valid for all smooth parametric models, etc. – but after all under simple i.i.d. conditions. Crucially, most of these concepts, methods, and results extend to classes of general regression models. Here we go through the various steps to see how the scene broadens and to learn the appropriate extensions for concepts, techniques, and results.

Consider in general terms regression data of the form (x_i, Y_i), with x_i a covariate vector, of length say p, thought to influence the main outcome Y_i. We assume here that the Y_i are independent given the covariates. Let $f(y_i \,|\, x_i, \theta)$ be a suitable density for y_i given x_i, with score function $u(y_i \,|\, x_i, \theta) = \partial \log f(y_i \,|\, x_i, \theta)/\partial\theta$ and information function $i(y_i \,|\, x_i, \theta) = \partial^2 \log f(y_i \,|\, x_i, \theta)/\partial\theta\,\partial\theta^{\mathrm{t}}$. The θ could comprise both regression coefficients and parameters describing the shape of the distributions. In this exercise we assume that the model holds, with θ_0 denoting the true parameter, an inner point in the parameter spoace. We also postulate *the ergodic condition*, that averages over covariates stabilise, with increasing sample size; formally, for each bounded $h(x)$, there is a well-defined limit $h_0 = \int h(x)\,\mathrm{d}R(x)$ for $n^{-1}\sum_{i=1}^n h(x_i)$, for an appropriate distribution R on the covariate space. For theory and applications, we do not need to model this R, or take it explicitly into account, beyond postulating its existence.

ergodic conditions

(a). First of all, there is a log-likelihood function, also in these regression setups, $\ell_n(\theta) = \sum_{i=1}^n \log f(y_i \,|\, x_i, \theta)$. The ML estimator $\widehat{\theta}$ is its maximiser, satisfying also $U_n(\widehat{\theta}) = 0$, with $U_n(\theta) = \sum_{i=1}^n u(y_i \,|\, x_i, \theta)$. Secondly, to extend theory and results for the i.i.d. case, see Ex. 5.18, we need to understand $U_n = n^{-1/2}U_n(\theta_0) = n^{-1/2}\sum_{i=1}^n u(Y_i \,|\, x_i, \theta_0)$. Show that it has mean zero and variance matrix $J_n = n^{-1}\sum_{i=1}^n J(x_i)$, where $J(x_i) = \mathrm{Var}_{\theta_0} u(Y_i \,|\, x_i, \theta_0)$. Under ergodic assumptions, there is convergence $J_n \to J$, say. Give Lindeberg type conditions under which $U_n \to_d U \sim \mathrm{N}_p(0, J)$.

(b). Extend techniques from Ex. 5.18 to deduce the natural parallel to (5.7), in this general regression setting, that

$$A_n(s) = \ell_n(\theta_0 + s/\sqrt{n}) - \ell_n(\theta_0) = U_n^{\mathrm{t}}s - \tfrac{1}{2}J_n s + r_n(s) \to_d A(s) = U^{\mathrm{t}}s - \tfrac{1}{2}s^{\mathrm{t}}Js,$$

Set up clear but mild regularity conditions, as in Ex. 5.13, to secure the required $r_n(s) \to_{\rm pr} 0$ for each s. Explain as in Ex. 5.18 that this leads to the two fundamental results

$$\sqrt{n}(\widehat{\theta}-\theta_0) \to_d {\rm N}_p(0, J^{-1}) \quad \text{and} \quad 2\{\ell_{n,\max} - \ell_n(\theta_0)\} \to_d \chi^2_p,$$

under these Lindeberg type conditions. Show also that the observed Fisher information matrix

observed Fisher information matrix

$$\widehat{J}_{n,\rm full} = -\partial^2 \ell_n(\widehat{\theta})/\partial\theta\,\partial\theta^{\rm t}, \tag{5.15}$$

i.e. the Hessian mastrix associated with the maximisation of the log-likelihood, satisfies $\widehat{J}_n = n^{-1}\widehat{J}_{\rm full} \to_{\rm pr} J$. Deduce from this that $\widehat{\theta} \approx_d {\rm N}_p(\theta_0, \widehat{J}^{-1}_{\rm full})$. (xx need to point to Ch2 thing with weak LLN for averages of non-i.i.d. xx)

(c). (xx rounding this off. also log-likelihood profiling and deviance and χ^2_1 limit and CDs. the point is that i.i.d. results all extend to the broad regression cases. point also to exercise below with outside the model. xx)

Ex. 5.44. *Linear regression revisited.* (xx edit and clean. xx) Consider the linear regression model of Ex. 3.34, with $Y_i \,|\, x_i \sim {\rm N}(x_i^{\rm t}\beta, \sigma^2)$, for which exact finite-sample theory has been well developed. We now take another look at this classical model, with the general likelihood tools.

(a). For the log-likelihood, show that $\ell_n(\beta,\sigma) = -n\log\sigma - \frac12 Q(\beta)/\sigma^2 - \frac12 n\log(2\pi)$, with $Q(\beta) = \sum_{i=1}^n (y_i - x_i^{\rm t}\beta)^2$ Show that the ML estimator for β is the least squares estimator $\widehat{\beta} = \Sigma_n^{-1} n^{-1} \sum_{i=1}^n x_i Y_i$, with $\Sigma_n = n^{-1}\sum_{i=1}^n x_i x_i^{\rm t}$, see the exercise mentioned, and that $\widehat{\sigma} = (Q_0/n)^{1/2}$, with $Q_0 = Q(\widehat{\beta})$ the minimum of $Q(\beta)$.

(b). Show that the score function becomes

$$u(y_i \,|\, x_i, \beta, \sigma) = \begin{pmatrix} (1/\sigma^2)(y_i - x_i^{\rm t}\beta)x_i \\ -1/\sigma + (1/\sigma^3)(y_i - x_i^{\rm t}\beta)^2 \end{pmatrix} = \begin{pmatrix} (1/\sigma)\varepsilon_i x_i \\ (1/\sigma)(\varepsilon_i^2 - 1) \end{pmatrix}$$

in terms of $\varepsilon_i = (y_i - x_i^{\rm t}\beta)/\sigma$, which are independent standard normals under the model. With (β_0, σ_0) the true parameters, deduce that the $(p+1)\times(p+1)$ Fisher information matrix becomes

$$J_n = n^{-1}\sum_{i=1}^n {\rm Var}_{\beta_0,\sigma_0} u(Y_i \,|\, x_i, \beta_0, \sigma_0) = (1/\sigma_0^2)\,{\rm diag}(\Sigma_n, 2),$$

with $\Sigma_n = n^{-1}\sum_{i=1}^n x_i x_i^{\rm t} = n^{-1}X^{\rm t}X$. Show also that the observed Fisher information matrix becomes $\widehat{J}_{n,\rm full} = (n/\widehat{\sigma}^2)\,{\rm diag}(\Sigma_n, 2)$.

(c). (xx then on to what likelihood theory implies for $\widehat{\beta}$ and $\widehat{\sigma}$. the J_n and $\widehat{J}_n$. a t_{n-p} vs. approximate normality. we reproduce the $\widehat{\beta}$ distribution, and come close for $\widehat{\sigma}^2$. xx)

Ex. 5.45. *Logistic regression.* Consider binary outcome data, where the values 0-1 for Y_i are influenced by a covariate vector x_i, of dimension say p. The logistic regression model takes the probabilities to be

logistic regression

$$p_i = \Pr(Y_i = 1 \,|\, x_i) = H(x_i^{\rm t}\beta) = \frac{\exp(x_i^{\rm t}\beta)}{1+\exp(x_i^{\rm t}\beta)} \quad \text{for } i = 1,\ldots,n, \tag{5.16}$$

with $H(u) = \exp(u)/\{1+\exp(u)\}$ the logistic transform, studied in Ex. 1.62. One may interpret the model via underlying or latent i.i.d. variables Z_i, having the logistic distribution, with the individuals having different thresholds; $p_i = \Pr(Z_i \le x_i^{\rm t}\beta)$ says that outcomes are '1' for those individuals whose Z_i fall to the left of the threshold.

(a). Show that $H(u) = p$ means $u = H^{-1}(p) = \log\{p/(1-p)\}$, so that the model can be represented as $\log\{p_i/(1-p_i)\} = x_i^{\rm t}\beta$.

(b). Show that $\Pr(Y_i = y \,|\, x_i) = p_i^y(1-p_i)^{1-y}$, for the two outcomes, and deduce that the log-likelihood function can be written

$$\ell_n(\beta) = \sum_{i=1}^n \{y_i \log p_i + (1-y_i)\log(1-p_i)\} = \sum_{i=1}^n \big[y_i x_i^{\rm t}\beta - \log\{1+\exp(x_i^{\rm t}\beta)\}\big].$$

Show from this that the estimation equation, giving rise to the ML estimator $\widehat\beta$, is $\sum_{i=1}^n (y_i - p_i)x_i = 0$, and that

$$J_{n,\rm full}(\beta) = -\frac{\partial^2 \ell_n(\beta)}{\partial\beta\partial\beta^{\rm t}} = \sum_{i=1}^n p_i(1-p_i)x_i x_i^{\rm t} = \sum_{i=1}^n H(x_i^{\rm t}\beta)\{1-H(x_i^{\rm t}\beta)\}x_i x_i^{\rm t}.$$

This matrix is assumed here to be positive definite, which in particular requires $n \ge p$. Explain that the log-likelihood function is concave, with $\widehat\beta$ the unique maximiser.

(c). Show that, under model conditions, $\widehat\beta \approx_d {\rm N}_p(\beta, \widehat J_{n,\rm full}^{-1})$, where $\widehat J_{n,\rm full} = J_{n,\rm full}(\widehat\beta)$ is the observed Fisher information matrix (the Hessian matrix of minus the normalised log-likelihood function, at the ML position).

(d). Consider an individual, perhaps outside the dataset, with covariate vector x_0. Show that $x_0^{\rm t}\widehat\beta$ is approximately a normal $(x_0^{\rm t}\beta, x_0^{\rm t}\widehat J_{n,\rm full}^{-1}x_0)$, and use this to construct a confidence interval for $p(x_0) = \Pr(Y_0 = 1 \,|\, x_0)$.

(e). (xx repair a bit here. xx) Consider the important special case of a single x_i recorded for Y_i, where we write the model equation as $p_i = H(a+bx_i)$, corresponding to 2-size vectors $(1,x_i)^{\rm t}$ in the more general notation used above. Show that $(\widehat a, \widehat b)$ are the solutions to $\sum_{i=1}^n (y_i - p_i) = 0$ and $\sum_{i=1}^n (y_i - p_i)x_i = 0$, and that

$$J_n(a,b) = \sum_{i=1}^n \frac{\exp(a+bx_i)}{\{1+\exp(a+bx_i)\}^2}\begin{pmatrix} 1 & x_i \\ x_i & x_i^2\end{pmatrix}.$$

(f). The logistic regression model uses the H transform, which is symmetric around zero. A more flexible model takes $p_i = H(x_i^{\rm t}\beta)^\kappa$, allowing skewness in the underlying distribution for the latent variables. Write down the log-likelihood function, explain that standard theory applies for the ML estimator $(\widehat\beta, \widehat\kappa)$. (xx point perhaps to illustration, Polish girls story, where κ is significantly larger than 1. xx)

Ex. 5.46. *Probit regression.* Consider again data (x_i, Y_i), with covariates x_i and 0-1 outcomes Y_i. For the logistic regression of Ex. 5.45 we modelled the 1-probabilities as $p_i = \Pr(Z_i \le x_i^{\rm t}\beta) = H(x_i^{\rm t}\beta)$, with H the logistic c.d.f. Clearly there are alternatives, and the more famous one among these is *probit regression*, which takes the Z_i to be standard normal, i.e. $p_i = \Phi(x_i^{\rm t}\beta)$.

probit regression, from probability of unit

(a). Explain that the log-likelihood becomes $\ell_n(\beta) = \sum_{i=1}^n [y_i \log p_i(\beta) + (1-y_i)\log\{1-p_i(\beta)\}]$. The ML estimator $\widehat\beta$ is its maximiser.

(b). Explain also that standard regression likelihood theory applies, with $\widehat{\beta} \approx_d \mathrm{N}_p(\beta, \widehat{J}_n^{-1})$, and give a recipe for computing the variance matrix.

(c). To compare analyses from logistic and probit regressions, we may scale the normal above to have the same variance $\tau^2 = \pi^2/3$ as the logistic. Explain that this leads to an equivalent model, though with scaled parameters, namely $p_i = \Pr(Z_i' \le x_i^{\mathrm{t}}\beta) = \Phi(x_i^{\mathrm{t}}\beta/\tau)$.

(d). Simulate a simple dataset of size $n = 1000$ pairs, with x_i taken $\mathrm{N}(0, \kappa^2)$ with $\kappa = 2$ and true probabilities $p_i = H(a_0 + b_0 x_i)$ with $a_0 = -0.33$ and $b_0 = 0.77$. The litt point with $\kappa = 2$ is to secure a fair range for the x_i, in view of the probabilities. (i) Estimate (a, b) in the logistic regression model $p_i = H(a + bx_i)$; and then (ii) estimate (a, b) in the probit model $p_i = \Phi((a + bx_i)/\tau)$. Draw the two estimated regression curves, say $H(\widehat{a} + \widehat{b}x)$ and $\Phi((\widetilde{a} + \widetilde{b}x)/\tau)$. Comment on what you find.

Ex. 5.47. *Poisson regression.* Consider independent count data $y_1, \ldots, y_n$, influenced by covariate vectors $x_1, \ldots, x_n$. The Poisson regression model, in its standard form, takes $Y_i \sim \mathrm{Pois}(\mu_i)$, with $\mu_i = \exp(x_i^{\mathrm{t}}\beta)$.

(a). Show that the log-likelihood function becomes

$$\ell_n(\beta) = \sum_{i=1}^n \{-\mu_i + y_i \log(\mu_i)\} = \sum_{i=1}^n \{y_i x_i^{\mathrm{t}}\beta - \exp(x_i^{\mathrm{t}}\beta)\},$$

and that the equations $\sum_{i=1}^n \{y_i - \mu_i(\beta)\}x_i = 0$ define the ML estimators.

(b). Show that $-\partial^2 \ell_n(\beta)/\partial\beta\,\partial\beta^{\mathrm{t}} = \sum_{i=1}^n \mu_i x_i x_i^{\mathrm{t}}$, leading to the observed Fisher information matrix $\widehat{J}_{n,\mathrm{full}} = \sum_{i=1}^n \widehat{\mu}_i x_i x_i^{\mathrm{t}}$, where $\widehat{\mu}_i = \exp(x_i^{\mathrm{t}}\widehat{\beta})$.

(c). For a case with covariate vector x_0, estimate the associated expected count $\mu_0 = \exp(x_0^{\mathrm{t}}\beta)$, and construct a confidence interval.

(d). Sometimes there is overdispersion, compared to how the counts y_i should behave under Poisson conditions. Such overdispersion could e.g. reflect 'hidden covariates' not taken on board in the model. A method handling overdispersion is to take $Y_i \,|\, \mu \sim \mathrm{Pois}(\mu_i)$, but modelling potential extra randomness via $\mu_i \sim \mathrm{Gam}(\exp(x_i^{\mathrm{t}}\beta)/\tau, 1/\tau)$. This is an application of Poisson-gamma mixtures from Ex. 1.27. Show that μ_i has mean $\exp(x_i^{\mathrm{t}}\beta)$ and variance $\tau \exp(x_i^{\mathrm{t}}\beta)$; τ small means getting back to ordinary Poisson regression. Show altso that Y_i has mean $\exp(x_i^{\mathrm{t}}\beta)$ variance $(1+\tau)\exp(x_i^{\mathrm{t}}\beta)$. Furthermore, work out that an expression for the log-likelihood function $\ell_n(\beta, \tau)$ is

$$\sum_{i=1}^n \{(a_i/\tau)\log(1/\tau) - \log\Gamma(a_i/\tau) + \log\Gamma(a_i/\tau + y_i) - (a_i/\tau + y_i)\log(1/\tau + 1)\},$$

in which $a_i = \exp(x_i^{\mathrm{t}}\beta)$. In applications it is often fruitful to profile out the β, and studying $\ell_{n,\mathrm{prof}}(\tau) = \max_{\mathrm{all}\ \beta} \ell_n(\beta, \tau)$. (xx pointer to Story iv.7. xx)

Ex. 5.48. *Exponential regression.* (xx to be polished. xx) Consider regression data (x_i, y_i) for $i = 1, \ldots, n$, for nonnegative outcomes y_i, which we now model as $Y_i \sim \mathrm{Expo}(\theta_i)$ with $\theta_i = \exp(x_i^{\mathrm{t}}\beta)$.

(a). Show that the log-likelihood may be written $\ell_n(\beta) = \sum_{i=1}^n \{x_i^{\mathrm{t}}\beta - \exp(x_i^{\mathrm{t}}\beta)y_i\}$, and the the ML estimator $\widehat{\beta}$ is the solution to $\sum_{i=1}^n x_i\{1 - \exp(x_i^{\mathrm{t}}\beta)y_i\} = 0$.

(b). Show next that the normalised expected Fisher information matrix takes the simple form $J_n = n^{-1}\sum_{i=1}^n x_i x_i^{\rm t}$. Explain from this that $\widehat\beta \approx_d {\rm N}_p(\beta, n^{-1}J_n^{-1})$, under model conditions.

(c). Conservative analysis, i.e. not necessarily trusting model conditions, leads to a more general approximation for the variance matrix, with $n^{-1}J_n^{-1}K_nJ_n^{-1}$. Identify the K_n and give a consistent estimator.

(d). (xx perhaps pointed to illustration in a story. xx) can also do hazard regression with $h_i(t) = h_0(t)\exp(x_i^{\rm t}\beta)$. then with $h_i(t) = h_0(t,\theta)\exp(x_i^{\rm t}\beta)$.

Ex. 5.49. *GLM regression.* (xx to be polished. xx) basic expofamily model $f(y,\theta,\kappa) = \exp\{\theta T(y) + \kappa U(y) - k(\theta,\kappa)\}h(y)$, here with one-dimensional θ,κ. now regression with $\theta_i = x_i^{\rm t}\beta$.

$$\ell_n(\beta,\kappa) = \sum_{i=1}^n \{x_i^{\rm t}\beta T(y_i) + \kappa U(y_i) - k(x_i^{\rm t}\beta,\kappa)\}.$$

then ML $\widehat\beta,\widehat\kappa$, and more.

(a).

(b). (xx example. xx)

Ex. 5.50. *Wald tests.* (xx something here, re Wald tests, used e.g. in regression models. p-value. more on power for the two variations with two nevnere. point back to Ex. 4.5, and to a couple of stories where we use this tool. xx)

Ex. 5.51. *A heteroscedastic linear regression model.* (xx edit and clean. xx) In various linear regression type applications for (x_i,y_i) data the linear mean assumption can be reasonable, whereas the variance might not be taken constant across covariates. Consider therefore the model with independent $Y_i\,|\,(x_i,w_i) \sim {\rm N}(x_i^{\rm t}\beta,\sigma_i^2)$, for $i=1,\ldots,n$, with covariate vectors x_i of length p and variance related covariates w_i of length q, influencing $\sigma_i = \sigma\exp(\gamma^{\rm t}w_i)$. These w_i could be a subset of the x_i or functions thereof. It is convenient to normalise these such that $\bar w = n^{-1}\sum_{i=1}^n w_i = 0$, which also means that σ is the standard deviation for an average individual, with w_i equal to $\bar w$.

(a). (xx log-likelihood. score function. J_n and $\widehat J_n$. approximations. pointers. xx) Show that the log-likelihood function can be written

$$\ell_n(\beta,\gamma) = -n\log\sigma - \tfrac12(1/\sigma^2)Q(\beta,\gamma), \quad \text{with } Q(\beta,\gamma) = \sum_{i=1}^n \frac{(y_i - x_i^{\rm t}\beta)^2}{\exp(2w_i^{\rm t}\gamma)}.$$

Show that minimising $Q(\beta,\gamma)$ over β, for fixed γ, takes place for

$$\widehat\beta(\gamma) = \Big\{\sum_{i=1}^n \frac{x_ix_i^{\rm t}}{\exp(2\gamma^{\rm t}w_i)}\Big\}^{-1}\sum_{i=1}^n \frac{x_iy_i}{\exp(2\gamma^{\rm t}w_i)}.$$

Demonstrate that this leads to the profiled log-likelihood $\ell_{n,{\rm prof}}(\gamma) = -n\log\widehat\sigma(\gamma) - \frac12 n$, where $\widehat\sigma(\gamma)^2 = Q_0(\gamma)/n$, with $Q_0(\gamma) = Q(\widehat\beta(\gamma),\gamma)$ the minimum sum of squares. Deduce from this that a recipe for finding the ML estimators consists in (i) minimising $Q_0(\gamma)$ over γ, yielding $\widehat\gamma$; (ii) reading off $\widehat\beta = \widehat\beta(\widehat\gamma)$ and $\widehat\sigma = \widehat\sigma(\widehat\gamma)$.

(b). (xx calibrate this with Wilks things. xx) Is it worthwhile, turning from classic linear regression, to include the extra layer of variance heterogeneity sophistication? Show that the log-likelihood-ratio test becomes that of comparing $D_n = 2n \log\{\widehat{\sigma}(0)/\widehat{\sigma}(\widehat{\gamma})\}$ to the χ^2_q, in which $\widehat{\sigma}(0)^2 = Q_0(0)/n$ is the standard estimator for σ^2 under variance constancy.

(c). For the $p+q+1$-parameter model, with parameters β, γ, σ, show that the score function becomes

$$u(y_i \,|\, x_i) = \begin{pmatrix} (1/\sigma^2)(y_i - x_i^{\mathrm{t}}\beta)x_i/\exp(2\gamma^{\mathrm{t}}w_i) \\ -w_i + (1/\sigma^2)(y_i - x_i^{\mathrm{t}}\beta)^2 w_i/\exp(2\gamma^{\mathrm{t}}w_i) \\ -1/\sigma + (1/\sigma^3)(y_i - x_i^{\mathrm{t}}\beta)^2/\exp(2\gamma^{\mathrm{t}}w_i) \end{pmatrix} = \begin{pmatrix} (1/\sigma)\varepsilon_i x_i/\exp(\gamma^{\mathrm{t}}w_i) \\ (1/\sigma)(\varepsilon_i^2 - 1)w_i \\ (1/\sigma)(\varepsilon_i^2 - 1) \end{pmatrix},$$

in terms of $\varepsilon_i = (y_i - x_i^{\mathrm{t}}\beta)/\{\sigma \exp(\gamma^{\mathrm{t}}w_i)\}$. Show from this that the normalised Fisher information matrix becomes $J_n = (1/\sigma_0^2)\mathrm{diag}(\Sigma_n(\gamma_0), M_n, 2)$, at the true parameters $(\beta_0, \gamma_0, \sigma_0)$, in terms of $\Sigma_n(\gamma) = n^{-1}\sum_{i=1}^n x_i x_i^{\mathrm{t}}/\exp(2\gamma^{\mathrm{t}}w_i)$ and $M_n = n^{-1}\sum_{i=1}^n w_i w_i^{\mathrm{t}}$.

(d). (xx check if $\widehat{J}_{n,\mathrm{full}}$ has these off-diagonal zeroes, or if it only holds for the information calculus. spell out nice behaviour for ML estimators. $\widehat{\gamma} \approx_d \mathrm{N}_q(\gamma, (\sigma^2/n)M_n^{-1})$. xx)

(e). (xx round off. confidence for $\mu(x_0, w_0)$, $x_0^{\mathrm{t}}\widehat{\beta} \pm 1.96\,\widehat{\sigma}\exp(\widehat{\gamma}^{\mathrm{t}}w_0)$. pointer to Story iv.4. xx)

Ex. 5.52. *Nonlinear regression.* (xx calibrate this with both classic linear regression and what we've said with genereal regression models. xx) Consider in general terms the model with independent $y_i \sim \mathrm{N}(m_i(\beta), \sigma^2)$ for $i = 1, \ldots, n$, wherethe means $m_i(\beta)$ are perhaps nonlinear functions of an appropriate vector prarameter β, involving also covariates.

(a). Show that the log-likelihood function becomes $-n\log\sigma - \frac{1}{2}Q_n(\beta)/\sigma^2$, with $Q_n(\beta) = \sum_{i=1}^n \{y_i - m_i(\beta)\}^2$. Show that the ML estimator for β is the minimiser of Q_n, and that $\widehat{\sigma}^2 = Q_n(\widehat{\beta})/n$.

(b). Show that the normalised Fisher information matrix becomes

$$J_n(\beta, \sigma) = \frac{1}{\sigma^2}\begin{pmatrix} \Sigma_n & 0 \\ 0 & 2 \end{pmatrix}, \quad \text{with } \Sigma_n = n^{-1}\sum_{i=1}^n m_i^*(\beta) m_i^*(\beta)^{\mathrm{t}},$$

in which $m_i^*(\beta) = \partial m_i(\beta)/\partial\beta$.

(c). Deduce that $\sqrt{n}(\widehat{\beta} - \beta) \to_d \mathrm{N}(0, \sigma^2\Sigma^{-1})$, under Lindeberg type conditions, where Σ is the limit of Σ_n. (xx more here. also the case with different normalisation, for the time series cyclic thing. xx)

Ex. 5.53. *We can do things.* The spirit of this exercise is to see that the log-likelihood machinery is useful, versatile, flexible, and not too hard to use in new situations, outside ordinary textbook terrain.

(a). One records the number of a certain event, say per week, over a time period. Suppose most of these counts are Poisson-like, with some parameter θ, but that a fraction comes from another Poisson with a higher parameter. A model for such data is that Y_i stems from the mixture distribution $(1-p)\,\mathrm{Pois}(\theta) + p\,\mathrm{Pois}(c\theta)$, with $c > 1$. (i) Find the mean and variance for this distribution. (ii) Generate such a dataset, say $y_1, \ldots, y_n$ with $n = 250$, $\theta = 10$, $c\theta = 20$, $p = 0.10$. Taking first $c = 2$ known, estimate the parameters

(p, θ), along with confidence intervals. (iii) Using the same data, now with c unknown, estimate all three parameters, with confidence intervals. Briefly investigate how much is earned in precision for estimating (p, θ) when c is known compared to it being unknown.

(b). Suppose data Y_i^0 are generated according to some normal (ξ, σ^2), but that we are only able to see those falling inside some observation window $[a, b]$. Write down the log-likelihood function, for the observed data Y_i. To see how it works, generate say $n_0 = 1000$ points from the standard normal, keep the n amongst these that fall inside $[1.0, 3.0]$, and estimate the two parameters from these. Find also approximate confidence intervals for the two parameters.

(c). For two distributions with densities f_1, f_2 and c.d.f.s F_1, F_2, consider the bivariate density

$$f(x, y, a) = f_1(x) f_2(y)[1 + a\{F_1(x) - \tfrac{1}{2}\}\{F_2(y) - \tfrac{1}{2}\}].$$

What is the range of the dependence parameter a? Show that the marginal distributions are f_1, f_2. Generate first binormal data with $\mathrm{N}(0, 1)$ marginals, and estimate a, with confidence interval for a. Then generate another binormal dataset, and estimate the five parameters $(\xi_1, \xi_2, \sigma_1, \sigma_2, a)$, along with confidence intervals.

(d). Tired of normality? Then generalise it. With $\mathrm{Be}_{a,b}(x)$ the c.d.f. for the $\mathrm{Beta}(a, b)$ distribution, consider the four-parameter model with c.d.f. $F(y) = \mathrm{Be}_{a,b}(\Phi((y - \xi)/\sigma))$. Write up an expression for the log-likelihood function. Generate first $n = 500$ datapoints from the standard normal, and check if the ML metod succeeds in coming close to the true parameters $(0, 1, 1, 1)$. Argue that you here have a test for normality, and spell out the ingredients. Then simulate another dataset, from $(a, b) = (1.23, 2.34)$, compute ML estimates and their estimated standard deviations.

Ex. 5.54. *Extending theory and methods to regression setups, II.* (xx smooth and polish. regression outside model conditions. careful with notion, point to previous for u, i. then $\theta_{0,n}$, J_n, K_n, sandwich. xx) For i.i.d. setups we have seen how the behaviour of ML estimators can be accurately described also when the data-generating density g is not inside the parametric model f_θ, via a clearly defined least false parameter θ_0 and the two matrices J and K, see Ex. 5.18. Under model conditions, the two are equal. Here we extend these notions and results to regression models. The setting is partly as in Ex. 5.43, with a parametric model $f(y_i \,|\, x_i, \theta)$ for $Y_i \,|\, x_i$, but notions and results need to be lifted to the agnostic state of affairs when the real density $g(y_i \,|\, x_i)$ is not necessarily inside the model.

(a). For given x, there is a KL distance

$$\mathrm{KL}_x(g(\cdot \,|\, x), f(\cdot \,|\, x, \theta)) = \int g(y \,|\, x) \log \frac{g(y \,|\, x)}{f(y \,|\, x, \theta)} \, \mathrm{d}y.$$

Consider the overall KL distance, weighted over the covariate vectors in the sample,

$$\mathrm{KL} = n^{-1} \sum_{i=1}^n \mathrm{KL}_{x_i} = \int \int g(y \,|\, x) \log \frac{g(y \,|\, x)}{f(y \,|\, x, \theta)} \, \mathrm{d}y \, \mathrm{d}R_n(x),$$

with R_n the empirical distribution over these covariate vectors. Show that this is a divergence, i.e. nonnegative and equal to zero only if $g(y \,|\, x) = f(y \,|\, x, \theta)$ for all y and

all $x_1, \ldots, x_n$. Let $\theta_{0,n}$ be the least false parameter, minimising this KL, for the given covariate vectors, and show that this is the same as the maximiser of

$$M_n(\theta) = n^{-1} \sum_{i=1}^{n} \int g(y \,|\, x_i) \log f(y \,|\, x_i, \theta) \, \mathrm{d}y.$$

Explain also that $M_n(\theta) \to M(\theta)$, under ergodic conditions, with this limit involving the distribution R over covariate space, using $\int h \, \mathrm{d}R_n \to \int h \, \mathrm{d}R$. Letting θ_0 be the maximiser of this limit $M(\theta)$, show also that $\theta_{0,n} \to \theta_0$, under weak conditions.

(b). The following arguments support the notion that the ML estimator $\widehat{\theta}$ aims at this least false $\theta_{0,n}$. With $\ell_n(\theta) = \sum_{i=1}^{n} \log f(y_i \,|\, x_i, \theta)$ the log-likelihood function, explain that $n^{-1}\ell_n(\theta)$ has mean $M_n(\theta)$. the mean of $n^{-1}\ell_n(\theta)$ is $M_n(\theta)$, and that its variance goes to zero. Now show that $\widehat{\theta} - \theta_{0,n} \to_{\mathrm{pr}} 0$.

(c). As in the earlier and simpler case of (5.7), it is fruitful to work with the log-likelihood function process. Show that

$$A_n(s) = \ell_n(\theta_{0,n} + s/\sqrt{n}) - \ell_n(\theta_{0,n}) = U_n^{\mathrm{t}} s - \tfrac{1}{2} s^{\mathrm{t}} J_n s + r_n(s),$$

now with $U_n = n^{-1/2} \sum_{i=1}^{n} u(Y_i \,|\, x_i, \theta_{0,n})$ and $J_n = -n^{-1} \sum_{i=1}^{n} i(Y_i \,|\, x_i, \theta_{0,n})$. Let $K(x_i) = \mathrm{Var}\, u(Y_i \,|\, x_i, \theta_{0,n})$. Under Lindeberg type conditions, show that $U_n \to_d U \sim \mathrm{N}_p(0, K)$, with K the limit of $n^{-1} \sum_{i=1}^{n} K(x_i)$. Explain that $A_n(s) \to_d A(s) = U^{\mathrm{t}} s - \frac{1}{2} s^{\mathrm{t}} J s$.

(d). (xx then read off limits. proof of pudding lies in eating, i.e. applications, below. estimates $\widehat{J}$ and $\widehat{K}$. xx) $\widehat{J} = -n^{-1} \partial^2 \ell_n(\widehat{\theta}) / \partial\theta \, \partial\theta^{\mathrm{t}}$, for estimating K, use the estimated scores $\widehat{u}_i = u(Y_i \,|\, x_i, \widehat{\theta})$, noting that these sum to zero. then use $\widehat{K} = n^{-1} \sum_{i=1}^{n} \widehat{u}_i \widehat{u}_i^{\mathrm{t}}$. sandwich $\widehat{\Sigma} = \widehat{J}^{-1} \widehat{K} \widehat{J}^{-1}$.

Ex. 5.55. *Linear regression: agnostic analysis.* (xx to be repaired and cleaned and calibrated with next. xx) When a regression model is used, without being fully correct, the general theory of Ex. 5.54 explains (i) how the ML estimators implicitly aim for the best parametric approximation, defined via Kullback–Leibler divergences weighted over the available covariate vectors, and (ii) how the approximate distributions are affected. Here we consider such consequences for the classic linear normal regression model. For a dataset of pairs (x_i, Y_i), assume merely that $\mathrm{E}\, Y_i = m(x_i)$, for some smooth $m(x)$ function, with i.i.d. errors ε_i with mean zero and variance σ_0^2. The linear model approximates this via the standard $Y_i \,|\, x_i \sim \mathrm{N}(x_i^{\mathrm{t}}\beta, \sigma^2)$, for which we have found the ML estimators $\widehat{\beta}$ and $\widehat{\sigma}$ in Ex. 5.44.

(a). Explain that minimising the KL distance, as per Ex. 5.54, is the same as maximising

$$n^{-1} \sum_{i=1}^{n} \{-\log\sigma - \tfrac{1}{2} \frac{1}{\sigma^2} \mathrm{E}\,(Y_i - x_i^{\mathrm{t}}\beta)^2\} = -\log\sigma - \tfrac{1}{2} \frac{1}{\sigma^2} n^{-1} \sum_{i=1}^{n} [\{m(x_i) - x_i^{\mathrm{t}}\beta\}^2 + \sigma_0^2].$$

Argue from this that the least false parameter $\beta_{0,n}$ is the one minimising $Q_n(\beta) = n^{-1} \sum_{i=1}^{n} \{m(x_i) - x_i^{\mathrm{t}}\beta\}^2$, and show that this means $\beta_{0,n} = \Sigma_n^{-1} n^{-1} \sum_{i=1}^{n} m(x_i) x_i$, where $\Sigma_n = n^{-1} \sum_{i=1}^{n} x_i x_i^{\mathrm{t}}$. If the model is perfect, there is a true β_{true} for which $\mathrm{E}\,(Y_i \,|\, x_i) = x_i^{\mathrm{t}} \beta_{\mathrm{true}}$, in which case $\beta_{0,n} = \beta_{\mathrm{true}}$. Show next that the least false $\sigma_{0,n}$ is determined by $\sigma_{0,n}^2 = Q_{0,n} + \sigma_0^2$, where $Q_{0,n}$ is the minimum of $Q_n(\beta)$. If the linear mean is a good

model, the $Q_{0,n}$ is small and the $\sigma_{0,n}$ aimed at by the ML estimator $\widehat{\sigma}$ is not much bigger than σ_0. If the linear mean is not a good approximation to the real $m(x)$, then the $\widehat{\sigma}$ implicitly picks up both data variability and the difference between $x_i^{\mathrm{t}}\beta$ and the real $m(x_i)$.

(b). For the two matrices J and K determining the behaviour of the ML estimators we have found in Ex. 5.44 that $\widehat{J} = (1/\widehat{\sigma}^2)\,\mathrm{diag}(\Sigma_n, 2)$, and that the score vectors, computed at the least false values, can be written $u(y_i \,|\, x_i, \theta_0) = (1/\sigma_{0,n})(x_i z_i, z_i^2 - 1)^{\mathrm{t}}$, in which $z_i = (y_i - x_i^{\mathrm{t}}\beta_{0,n})/\sigma_{0,n}$. Note that we for the estimated standardised residuals $\widehat{z}_i = (y_i - x_i^{\mathrm{t}}\widehat{\beta})/\widehat{\sigma}$ have $\sum_{i=1}^n \widehat{z}_i x_i = 0$ and $n^{-1}\sum_{i=1}^n \widehat{z}_i^2 = 1$. For the K matrix, and with notation from Ex. 5.54, show that

$$n^{-1}\sum_{i=1}^n K(x_i) = n^{-1}\sum_{i=1}^n \frac{1}{\sigma_{0,n}^2}\mathrm{Var}\begin{pmatrix} z_i x_i \\ z_i^2 - 1\end{pmatrix}.$$

(xx round this off. two substories, depending on what is assumed for z_i. with $y_i = m(x_i) + \sigma_0\varepsilon_i$, the z_i are not i.i.d., but we may cope with that too, via the right $\widehat{K}$. simplification if we postulate $m(x_i) = x_i^{\mathrm{t}}\beta_0$, then involving only $\widehat{\gamma}_3$ and $\widehat{\gamma}_4$ for the $\widehat{z}_i$. xx) a somewhat messy model-robust approach would use

$$\widehat{K} = \frac{1}{\widehat{\sigma}^2} n^{-1}\sum_{i=1}^n \begin{pmatrix} \widehat{z}_i^2 x_i x_i^{\mathrm{t}} & \widehat{z}_i^3 x_i \\ \widehat{z}_i^3 x_i^{\mathrm{t}} & \widehat{z}_i^4 - 1\end{pmatrix}.$$

(c).

Ex. 5.56. *Logistic regression, Poisson regression: agnostic analysis.* (xx to be repaired and cleaned. xx) When the regression model used does not necessarily hold up, ML estimators aim for the relevant least false parameters, the general theory of Ex. 5.54 shows that $\widehat{\theta} \approx_d \mathrm{N}(\theta_{0,n}, \widehat{\Sigma}/n)$, for the relevant least false parameter and with estimated sandwich matrix $\widehat{\Sigma}_n = \widehat{J}^{-1}\widehat{K}\widehat{J}^{-1}$. Under model conditions, matters simplify, the underlying J_n and K_n matrices are equal, and we have the standard result $\widehat{\theta} \approx_d \mathrm{N}(\theta_0, \widehat{J}^{-1}/n)$, used extensively in statistical software packages for a range of regression models. Here we check what the model agnostic setup leads to, for securing model robust inference, for standard logistic and Poission regression models; see also the agnostic analysis for the linear regression model in Ex. 5.55.

(a). Consider logistic regression (xx point back xx), with $p_i = H(x_i^{\mathrm{t}}\beta)$: show that $u(y_i \,|\, x_i, \beta) = (y_i - p_i)x_i$, and that

$$\widehat{J}_n = n^{-1}\sum_{i=1}^n \widehat{p}_i(1-\widehat{p}_i)x_i x_i^{\mathrm{t}}, \quad \widehat{K}_n = n^{-1}\sum_{i=1}^n (y_i - \widehat{p}_i)^2 x_i x_i^{\mathrm{t}}.$$

(b). Then look at Poisson regression (xx point back xx), with $\mu_i = \exp(x_i^{\mathrm{t}}\beta)$: show that $u(y_i \,|\, x_i, \beta) = (y_i - \mu_i)x_i$, with

$$\widehat{J}_n = n^{-1}\sum_{i=1}^n \widehat{\mu}_i x_i x_i^{\mathrm{t}}, \quad \widehat{K}_n = n^{-1}\sum_{i=1}^n (y_i - \widehat{\mu}_i)^2 x_i x_i^{\mathrm{t}}.$$

If there is overdispersion, with $(y_i - \widehat{\mu}_i)^2$ tending to be bigger than $\widehat{\mu}_i$, this is picked up here, with sandwich matrix bigger than $\widehat{J}^{-1}$. The potential error of applying straightforward Poisson regression, perhaps via standard packages, is that variability is underestimated, with confidence intervals becoming too narrow. For an illustration of this, see Story iv.7.

(c). With the Poisson-gamma overdispersion regression model studied in Ex. 5.47, where $Y_i \,|\, x_i$ has mean $\exp(x_i^{\mathrm{t}}\beta)$ and variance $(1+\tau)\exp(x_i^{\mathrm{t}}\beta)$, show that $\widehat{K}_n \to_{\mathrm{pr}} (1+\tau)J$, in terms of the limit J of $\widehat{J}_n$. Deduce that if $\widehat{\beta}$ is the ML estimator computed for the standard Poisson regression model, then $\sqrt{n}(\widehat{\beta} - \beta_{0,n}) \to_d \mathrm{N}_p(0, (1+\tau)J^{-1})$.

Influence functions

Ex. 5.57. *Influence functions.* For a distribution function F, consider some associated parameter, say $\theta = T(F)$, with T the appropriate functional mapping the distribution to the parameter value in question. Examples include the mean, the standard deviation, the skewness, the interquartiale range, a threshold probability. The *influence function* for $\theta = T(F)$ is a very useful quantity, as we shall see. It is defined as

$$\mathrm{IF}(F, y) = \lim_{\varepsilon \to 0} (1/\varepsilon)\{T((1-\varepsilon)F + \varepsilon\delta(y)) - T(F)\}. \tag{5.17}$$

Here $\delta(y)$ is the measure putting full mass 1 at the point y, and $(1-\varepsilon)F + \varepsilon\delta(y)$ the consequent mixture distribution. A variable Y_ε drawn from this mixture is from F with probability $1-\varepsilon$ and is equal to y with probability ε.

(a). Consider $\theta(F) = \mathrm{E}_G\, h(Y) = \int h(y)\,\mathrm{d}F(y)$, the mean of $h(Y)$. Show that $\mathrm{IF}(F, y) = h(y) - \theta(F)$. In particular, the influence is bounded when h is, but unbounded e.g. in the case of the plain mean $h(y) = y$, which signifies a potential lack of robustness of this mean parameter functional $\theta = \mathrm{E}_G\, Y$.

(b). Then consider the class of smooth functions of means. For mean type parameters $\gamma_1 = \mathrm{E}_G\, h_1(Y), \ldots, \gamma_k = \mathrm{E}_G\, h_k(Y)$, let $\theta = T(F) = A(\gamma_1(F), \ldots, \gamma_k(F))$, where $A(u_1, \ldots, u_k)$ is smooth in a neighbourhood of $(\gamma_1(F), \ldots, \gamma_k(F))$. Show that this parameter has influence function

$$\begin{aligned} \mathrm{IF}(F, y) &= c_1(F)\,\mathrm{IF}_{\gamma_1}(F, y) + \cdots + c_k(F)\,\mathrm{IF}_{\gamma_k}(F, y) \\ &= c_1(F)\{h_1(y) - \gamma_1(F)\} + \cdots + c_k(F)\{h_k(y) - \gamma_k(F)\}, \end{aligned}$$

with $c_j(F)$ is the partial derivative $\partial A(u_1, \ldots, u_k)/\partial u_j$, evaluated at $(\gamma_1(F), \ldots, \gamma_k(F))$.

(c). Writing $\mu_F = \mathrm{E}_G\, Y$ for the mean, show for the variance parameter $\sigma_F^2 = \mathrm{E}_G\, Y^2 - \mu_F^2$ that its influence function becomes

$$\mathrm{IF}(F, y) = -2\,\mu_F(y - \mu_F) + y^2 - \mathrm{E}_G\, Y^2 = (y - \mu_F)^2 - \sigma_F^2.$$

Then for the standard deviation parameter $\sigma(F)$ itself, show that its influence function becomes

$$\mathrm{IF}_\sigma(F, y) = \tfrac{1}{2}(1/\sigma_F)\{(y - \mu_F)^2 - \sigma_F^2\}.$$

(d). For a given parametric family $f(y, \theta)$, consider the ML functional $T(F)$, mapping a given F with density f to the least parameter value $\theta_0 = \theta_0(F)$, the minimiser of the information distance $\mathrm{KL}(f, f(\cdot, \theta))$, or the maximiser of $\int \log f(y, \theta)\,\mathrm{d}F(y)$. With F_n the empirical distribution of the data, placing probability $1/n$ on each data point, cf. Ex. 3.9, Show that $T(F_n)$ is the ML estimator, and that its influence function becomes $\mathrm{IF}(F, y) = J^{-1}u(y, \theta_0)$.

Ex. 5.58. *Influence for quantiles.* Assume the c.d.f. F has a smooth density f.

(a). Consider $\mu = F^{-1}(\frac12)$, the median. Show that the influence function $\mathrm{IF}(F,y)$ becomes $-\frac12/f(\mu)$ for $y \le \mu$ and $\frac12/f(\mu)$ for $y > \mu$. Verify that the influence function has mean zero and variance $\frac14/f(\mu)^2$.

(b). More generally, let $\mu_q = F^{-1}(q)$, for some quantile level $q \in (0,1)$. Show that $\mathrm{IF}(F,y)$ is $-(1-q)/f(\mu_q)$ for $y \le \mu_q$ and $q/f(\mu_q)$ for $y > \mu_q$. Comment on the fact that these are bounded.

(c). Then consider a smooth function of quantiles, say $\gamma(F) = A(Q_1(F),\ldots,Q_k(F))$, where $Q_j(F) = F^{-1}(q_j)$. Show that its influence function is $\mathrm{IF}(F,y) = c_1(F)\,\mathrm{IF}_1(F,y) + \cdots + c_k(F)\,\mathrm{IF}_k(F,y)$, in terms of $\mathrm{IF}_j(F,y)$ the influence function of $F^{-1}(q_j)$ calculated as above, and where $c_j(F) = \partial A(Q_1,\ldots,Q_k)/\partial Q_j$ computed at $(F^{-1}(q_1),\ldots,F^{-1}(q_k))$. As an illustration, find and graph the influence function for $\gamma(F) = F^{-1}(0.90) - F^{-1}(0.10)$.

Ex. 5.59. *An estimator represented via its influence function.* Consider an i.i.d. sequence $Y_1, Y_2, \ldots$ from F, with $\theta = T(F)$ a parameter of interest. It may be estimated nonparametrically using $\widehat\theta = T(F_n)$, with F_n the empirical distribution. Here we work towards a representation of $\widehat\theta - \theta = T(F_n) - T(F)$ in terms of the influence function.

(a). Consider the case of $\theta = A(\gamma(F))$, where $\gamma(F) = \mathrm{E}_G\, h(Y) = \int h\,\mathrm{d}F$. Show that $\widehat\theta = T(F_n)$ is equal to $A(\bar h)$, with $\bar h = \int h\,\mathrm{d}F_n = n^{-1}\sum_{i=1}^n h(y_i)$. Assuming $A(u)$ smooth, with two derivatives, show that

$$\widehat\theta = A(\gamma_0) + A'(\gamma_0)(\bar h - \gamma_0) + \tfrac12 A''(\gamma_0)(\bar h - \gamma_0)^2 + o_{\mathrm{pr}}(1/n),$$

with $\gamma_0 = \gamma(F)$. Deduce that $\mathrm{E}\,\widehat\theta = \theta + \frac12 A''(\gamma_0)\tau^2/n + o(1/n)$, in terms of $\tau^2 = \mathrm{Var}\, h(Y_i)$, and that $\widehat\theta - \theta = n^{-1}\sum_{i=1}^n \mathrm{IF}(F,y_i) + b/n + o_{\mathrm{pr}}(1/\sqrt n)$, where $b = \frac12 A''(\gamma_0)\tau^2$.

(b). Generalise to the case of $T(F) = A(\gamma_1(F),\ldots,\gamma_p(F))$ being a smooth function of several means, as studied also in Ex. 5.57, with $\gamma_j(F) = \int h_j\,\mathrm{d}F$. Show that $\mathrm{E}\,\widehat\theta = \theta + b/n + o(1/n)$, with $b = \frac12 \mathrm{Tr}(A''(\gamma_0)K)$, with K the variance matrix of $(h_1(Y_i),\ldots,h_p(Y_i))^{\mathrm t}$, and $A''(\gamma_0)$ the second order derivative matrix of A, computed at γ_0. Show that

$$\widehat\theta - \theta = T(F_n) - T(F) = n^{-1}\sum_{i=1}^n \mathrm{IF}(F,Y_i) + b/n + \varepsilon_n, \tag{5.18}$$

with $\varepsilon_n = o_{\mathrm{pr}}(1/\sqrt n)$, i.e. small enough to have $\sqrt n\,\varepsilon_n \to_{\mathrm{pr}} 0$. Show

(c). The powerful representation (5.18) actually holds quite generally, as long as $T(F)$ is a moderately smooth functional (xx find refs, Shao (1991), Jullum and Hjort (2017) xx), though with no easy general formula for the bias component b/n. Deduce that $\sqrt n(\widehat\theta - \theta)$ has the limit distribution $\mathrm{N}(0,\kappa^2)$, with κ^2 the variance of $\mathrm{IF}(Y_i,\theta)$, and with the bias part b/n disappearing in this normal limit.

(d). Use the above to find the limit distribution of $\sqrt n(\widehat\sigma - \sigma)$. This gives a new and partly simpler proof of things proved in Ex. 2.72.

Ex. 5.60. *Influence functions for BHHJ and for weighted likelihood estimators.* For some smooth parametric model $f_\theta(y) = f(y,\theta)$, consider the BHHJ estimation method of Ex. 5.9 and 5.19. The setup involves the data-generating density g, and for the given tuning parameter a, the defining parameter is $\theta_0 = \theta_{0,a}$, the minimiser of $\int f_\theta^{1+a}\,\mathrm{d}y - (1+1/a)\int g f_\theta^a\,\mathrm{d}y$, also the solution to $\int (f_\theta^{1+a} - g f_\theta^a)u_\theta\,\mathrm{d}y = 0$. Here $u_\theta(y) = u(y,\theta)$ is the score function, and we will also need the information function $i_\theta(y) = i(y,\theta)$ below. As in Ex. 5.19 we use f_0, u_0, i_0 for these functions at θ_0.

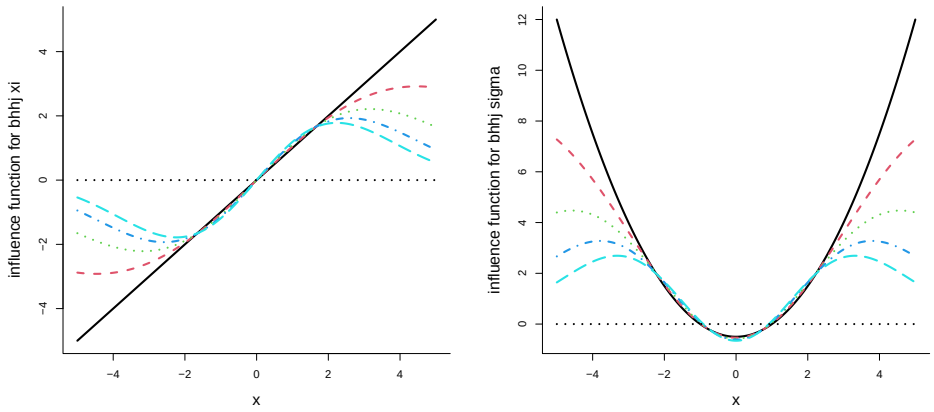


Figure 5.2: *Influence functions for the BHHJ estimator, for the normal* (ξ, σ^2) *model, for values* $a = 0, 0.05, 0.10, 0.15, 0.20$*; the full black curves are for the ML case* $a = 0$*. Left panel: for* $\widehat{\xi}_a$*; right panel: for* $\widehat{\sigma}_a$*.*

(a)**.** Explain that to find the influence function $\mathrm{IF}_a(G, y)$, we need for a small ε to assess the solution $\theta_\varepsilon = \theta_0 + \delta$ to the equation

$$\int f_{\theta_0+\delta}^{1+a} u_{\theta_0+\delta} \,\mathrm{d}y = (1-\varepsilon) \int g f_{\theta_0+\delta}^{a} u_{\theta_0+\delta} \,\mathrm{d}y + \varepsilon f(y, \theta_0+\delta)^a u(y, \theta_0+\delta).$$

The programme is now for the given small ε to carry out first order Taylor expansion for δ small and then solve for δ. Some details are as follows: show first

$$f_{\theta_0+\delta}^{a} \doteq f_0^a(1 + a\delta^{\mathrm{t}} u_0), \quad f_{\theta_0+\delta}^{1+a} \doteq f_0^{1+a}\{1 + (1+a)\delta^{\mathrm{t}} u_0\}, \quad u_{\theta_0+\delta} \doteq u_0 + i_0\delta.$$

Again with $\xi_a = \int f_0^{1+a} u_0 \,\mathrm{d}u = \int g f_0^a u_0 \,\mathrm{d}y$, use this to write the left and right hand sides of the equation as

$$\begin{aligned} L &\doteq \xi_a + \Big\{(1+a)\int f_0^{1+a} u_0 u_0^{\mathrm{t}} \,\mathrm{d}y + \int f_0^{1+a} i_0 \,\mathrm{d}y\Big\}\delta, \\ R &\doteq \xi_a + \Big\{\int g f_0^a (u_0 u_0^{\mathrm{t}} + i_0) \,\mathrm{d}y\Big\}\delta + \varepsilon\{f(y,\theta_0)^a u(y,\theta_0) - \xi_a\}. \end{aligned}$$

Conclude that $\mathrm{IF}_a(G, y) = M_a^{-1}\{f_0(y)^a u_0(y) - \xi_a\}$, with the matrix

$$M_a = (1+a)\int f_0^{1+a} u_0 u_0^{\mathrm{t}} \,\mathrm{d}y + \int f_0^{1+a} i_0 \,\mathrm{d}y - \int g f_0^a (u_0 u_0^{\mathrm{t}} + i_0) \,\mathrm{d}y.$$

Check with the J_a matrix of Ex. 5.19 that in fact $M_a = J_a/(1+a)$, so that

$$\mathrm{IF}_a(G, y) = (1+a) J_a^{-1}\{f(y,\theta_0)^a u(y,\theta_0) - \xi_a\}.$$

Explain that as $a \to 0$, the limit is $J_0^{-1} u(y, \theta_0)$, the influence function for the ML estimation method. Show also that under conditions, M_a simplifies to $\int f_0^{1+a} u_0 u_0^{\mathrm{t}} \,\mathrm{d}y$.

(b)**.** Consider the case of the normal $\mathrm{N}(\xi, \sigma^2)$. Under model conditions, find formulae for and graph the influence functions $\mathrm{IF}_{1,a}(y)$ and $\mathrm{IF}_{2,a}(y)$, for the ξ and the σ, at the

standard position $(\xi,\sigma)=(0,1)$, for a equal to 0, 0.05, 0.10, 0.15, 0.20. Construct a version of Figure 5.2. Note that the BHHJ influence functions mimic those for the ML (i.e. for $a=0$) for the main expected data range, but that they sensibly redescend down to zero for values far away, via density downweighting. Show in particular that the influence functions are bounded, for each $a>0$. For illustration of how the BHHJ successfully deals with outliers (without needing to identify them), and for seeing that very little efficiency is lost for small a, see Story vii.5. Formulae needed here include these, which you should prove, in addition to $\xi_{1,a}=0$:

$$\begin{aligned}\xi_{2,a} &= -\sigma^{-a}(2\pi)^{-a/2}a/(1+a)^{3/2},\\ M_{11,a} &= (1/\sigma^2)\sigma^{-a}(2\pi)^{-a/2}/(1+a)^{3/2},\\ M_{22,a} &= (1/\sigma^2)\sigma^{-a}(2\pi)^{-a/2}\{3/(1+a)^2-2/(1+a)+1\}.\end{aligned}$$

(c)**.** Carry out similar analysis for the Gamma (α,β) model.

(d)**.** Consider then the MWL method of Ex. 5.11 and 5.21. For a given weight function $w(y)$, the parameter aimed at for the MWL estimator is the maximiser $\theta_0=\theta_{0,w}$ of $\int w(g\log f_\theta - f_\theta)$. We need $J_w=\int wf_0u_0u_0^{\rm t}\,{\rm d}y+\int w(f_0-g)i_0\,{\rm d}y$ and $\xi_w=\int wf_0u_0\,{\rm d}y=\int wgu_0\,{\rm d}y$, computed at θ_0. Show that the influence function becomes

$$\mathrm{IF}(G,y)=J_w^{-1}\{w(y)u(y,\theta_0)-\xi_w\}.$$

Further material and leftovers (as of 12-March-2025)

Ex. 5.61. *Likelihood inference for binormal correlation.* (xx to be calibrated with what is in Ch2 for R_n. finding J directly from derivatives of log-likelihood is somewhat messy, but here we may find it via other means. xx) Let (X_i,Y_i) be binormal i.i.d. pairs with zero means, standard deviations σ_1,σ_2, and correlation ρ. We have studied the behaviour of the empirical correlation coefficient R_n in Ex. 2.54, and shall now revisit that analysis, via likelihoods.

(a)**.** Show that the log-likelihood function may be written

$$\ell_n(\sigma_1,\sigma_2,\rho)=\sum_{i=1}^n\Big\{-\tfrac12\log(1-\rho^2)-\log\sigma_1-\log\sigma_2-\tfrac12\frac{1}{1-\rho^2}\Big(\frac{X_i^2}{\sigma_1^2}+\frac{Y_i^2}{\sigma_2^2}-2\rho\frac{X_i}{\sigma_1}\frac{Y_i}{\sigma_2}\Big)\Big\}.$$

From this establish that the ML estimators are $\widehat\sigma_1^2=n^{-1}\sum_{i=1}^n X_i^2$, $\widehat\sigma_2^2=n^{-1}\sum_{i=1}^n Y_i^2$, $\widehat\rho=n^{-1}\sum_{i=1}^n X_iY_i/(\widehat\sigma_1\widehat\sigma_2)$.

(b)**.** Explain that we may write $X_i=\sigma_1U_i$ and $Y_i=\sigma_2V_i$, where the (U_i,V_i) have standard deviations 1. Use the CLT and formulae for product moments to argue that

$$\sqrt n\begin{pmatrix}n^{-1}\sum_{i=1}^n U_i^2-1\\ n^{-1}\sum_{i=1}^n V_i^2-1\\ n^{-1}\sum_{i=1}^n U_iV_i-\rho\end{pmatrix}\to_d\begin{pmatrix}A\\B\\C\end{pmatrix}\sim\mathrm{N}_3(0,\Sigma),\ \text{with }\Sigma=\begin{pmatrix}2&2\rho^2&2\rho\\2\rho^2&2&2\rho\\2\rho&2\rho&1+\rho^2\end{pmatrix}.$$

Use the delta method to demonstrate that

$$\begin{pmatrix}\sqrt n(\widehat\sigma_1-\sigma_1)\\ \sqrt n(\widehat\sigma_2-\sigma_2)\\ \sqrt n(R_n-\rho)\end{pmatrix}\to_d\begin{pmatrix}\frac12\sigma_1A\\ \frac12\sigma_2B\\ C-\frac12\rho A-\frac12\rho B\end{pmatrix}.$$

(c). Then find the variance matrix for this limit distribution,

$$M = \begin{pmatrix} \frac{1}{2}\sigma_1^2 & \frac{1}{2}\sigma_1\sigma_2\rho^2 & \frac{1}{2}\sigma_1(1-\rho+\rho^2-\rho^3) \\ \frac{1}{2}\sigma_1\sigma_2\rho^2 & \frac{1}{2}\sigma_2^2 & \frac{1}{2}\sigma_2(1-\rho+\rho^2-\rho^3) \\ \frac{1}{2}\sigma_1(1-\rho+\rho^2-\rho^3) & \frac{1}{2}\sigma_2(1-\rho+\rho^2-\rho^3) & (1-\rho^2)^2 \end{pmatrix}.$$

Argue from general ML theory that $M = J^{-1}$, with the information matrix $J = J(\sigma_1, \sigma_2, \rho)$. In particular, as established in Ex. 2.54, $\sqrt{n}(\widehat{\rho} - \rho) \to_d \mathrm{N}(0, (1-\rho^2)^2)$. Find also the J matrix from M.

(d). Generalise the above to the full five-parameter binormal model, say $(\xi_1, \xi_2, \sigma_1, \sigma_2, \rho)$. Find ML estimators for these, and the 5×5 information matrix.

Ex. 5.62. *Estimating the size of the hidden box.* Suppose a multinomial experiment has been carried out, involving N individuals, sorted in boxes $0, 1, \ldots, k$, with probabilities $p_0, p_1, \ldots, p_k$ summing to one. The full set of observations should have been $(N_0, N_1, \ldots, N_k)$, counting the numbers in these boxes, but here we have in mind situations where the zero-box count N_0 has not been observed. The task is to estimate N_0, and hence $N = N_0 + R$, with $R = N_1 + \cdots + N_k$ the total observed count in boxes $1, \ldots, k$. The setup of Ex. 2.87, involving boxes 00, 01, 10, 11, is a special case. Methods developed below will be used in Story iii.14 to estimate how many persons were killed in Guatemala, 1978–1996, based on a Venn diagram constructed from three lists of murdered people.

(a). Assume a smooth parametric model is used for the probabilities, with $p_j = p_j(\theta)$ for $j = 0, 1, \ldots, k$, with θ of dimension $r \leq k-1$, admitting two continuous derivatives. Show that the log-likelihood may be written

$$\ell(N, \theta) = \log(N!) - \log((N-R)!) + (N-R)\log p_0(\theta) + \sum_{j=1}^{k} N_j \log p_j(\theta).$$

Explain that this also leads to a profiled log-likelihood in N,

$$\ell_{\mathrm{prof}}(N) = \max\{\ell(N, \theta) \colon \text{all } \theta\} = \ell(N, \widehat{\theta}_N).$$

Maximising this function gives the ML estimators $(\widehat{N}, \widehat{\theta})$.

(b). We do assume that the parametric model holds, with true parameter θ_0 an inner point in the parameter space; write also N_0 for the true N. Large-sample results will be reached below in the framework of growing N_0. As is seen for the special case of Ex. 2.87, $\widehat{N}$ and $\widehat{\theta}$ will converge at different rates, so it is helpful to write $N = N_0\rho$ and $\widehat{\rho} = \widehat{N}/N_0$. Explain that this leads to the approximate log-likelihood $\ell(\rho, \theta)$ equal to

$$N_0\rho \log(N_0\rho) - (N_0\rho - R)\log(N_0\rho - R) + (N_0\rho - R)\log p_0(\theta) + \sum_{j=1}^{k} N_j \log p_j(\theta).$$

(c). From basic multinormal limits for multinomials, as with Ex. 2.49, show that there is joint convergence $N_0^{1/2}\{N_j/N - p_j(\theta_0)\} \to_d A_j$, a zero-mean multinormal vector with $\mathrm{cov}(A_j, A_\ell) = p_j(\theta_0)\{\delta_{j,\ell} - p_\ell(\theta_0)\}$; note that $A_0 + \cdots + A_k = 0$. With $\widetilde{p}_0 = 1 - R/N_0$, show that $N_0^{1/2}\{\widetilde{p}_0 - p_0(\theta_0)\} \to_d A_0$.

(d). Letting $\psi_j(\theta) = \partial \log p_j(\theta)/\partial\theta$, explain that $\sum_{j=0}^k p_j(\theta)\psi_j(\theta) = 0$. Below, write for crispness p_j and ψ_j for $p_j(\theta_0)$ and $\psi_j(\theta_0)$. Put up formuale for the first-order partial derivatives of the log-likelihood $\ell(\rho,\theta)$. Show that when evaluated at (N_0,θ_0), these become

$$\begin{aligned} U_0 &= N_0(-\log\widetilde{p}_0 + \log p_0), \\ V_0 &= N_0\Big\{(\widetilde{p}_0 - p_0)\psi_0 + \sum_{j=1}^k (N_j/N_0 - p_j)\psi_j\Big\}. \end{aligned}$$

(e). A key quantity is $M = \sum_{j=0}^k p_j\psi_j\psi_j^{\mathrm{t}}$. Show that there is convergence in distribution

$$N_0^{-1/2}\begin{pmatrix} U_0 \\ V_0\end{pmatrix} \to_d \begin{pmatrix} U \\ V\end{pmatrix} = \begin{pmatrix} -(1/p_0)A_0 \\ \sum_{j=0}^k A_j\psi_j\end{pmatrix} \sim \mathrm{N}_{r+1}(0,J),$$

with

$$J = \begin{pmatrix} (1-p_0)/p_0 & -\psi_0^{\mathrm{t}} \\ -\psi_0 & M\end{pmatrix}.$$

This is the Fisher information matrix for this reparametrised model, with parameters (ρ,θ), evaluated at $(1,\theta_0)$. Verify via derivatives and limits that J is also the limit in probability of $-N_0^{-1}H_0$, where H_0 is the $(r+1)\times(r+1)$ Hessian matrix of second order partial derivatives of $\ell(\rho,\theta)$, evaluated at $(1,\theta_0)$.

(f). Show from this that $N_0^{1/2}(\widehat{N}/N_0 - 1) = (\widehat{N} - N_0)/N_0^{1/2}$ tends in distribution to $\mathrm{N}(0,\tau^2)$, with

$$\tau^2 = \frac{1}{(1-p_0)/p_0 - \psi_0^{\mathrm{t}}M^{-1}\psi_0} = \frac{p_0}{1-p_0-p_0\delta}, \quad \text{where } \delta = \psi_0^{\mathrm{t}}M^{-1}\psi_0.$$

Though the main concern might be estimating N_0 and $N = N_0 + R$, show also that $N_0^{1/2}(\widehat{\theta} - \theta_0) \to_d \mathrm{N}_r(0,K)$, with

$$K = (M - (1-p_0)\psi_0\psi_0^{\mathrm{t}}/p_0)^{-1} = M^{-1} + dM^{-1}\psi_0\psi_0^{\mathrm{t}}M^{-1},$$

where we write $d = r_0/(1-r_0\delta) = p_0/(1-p_0-p_0\delta)$ and $r_0 = p_0/(1-p_0)$. Explain that had N_0 been observed too, the result would have beem $\mathrm{N}_r(0,M^{-1})$, so $dM^{-1}\psi_0\psi_0^{\mathrm{t}}M^{-1}$ is the extra amount of uncertainty when estimating θ, caused by N_0 not being counted.

(g). Show via Wilks theory for profiled log-likelihood functions, as in Ex. 5.29, that

$$D(N_0) = 2\{\ell_{\mathrm{prof,max}} - \ell_{\mathrm{prof}}(N_0)\} \to_d \chi_1^2.$$

This leads to confidence intervals, e.g. $\{N_0\colon D(N_0) \le 1.96^2\}$, and a full confidence curve $\mathrm{cc}(N) = \Gamma_1(D(N_0))$, as per Ch. 7, which we can read off directly from the profiled log-likelihood, without needing the perhaps cumbersome estimation of the J matrix.

(h). For the above results we are assuming that the parametric model $p_j(\theta)$ holds, also for the null box, which cannot really be verified or tested from $N_1,\ldots,N_k$ data alone, without other assumptions again, or perhaps follow-up data. Show in general terms that the likelihood can be written

$$\begin{aligned} L(N,\theta) &= \frac{N!}{N_0!\,N_1!\cdots N_k!}p_0^{N_0}p_1^{N_1}\cdots p_k^{N_k} \\ &= \frac{N!}{R!\,(N-R)!}p_0^{N-R}(1-p_0)^R\frac{R!}{N_1!\cdots N_k!}q_1^{N_1}\cdots q_k^{N_k}, \end{aligned}$$

splitting into a binomial $R \sim \text{binom}(N, 1-p_0)$ and a multinomial $(R, q_1, \ldots, q_k)$ for what is actually observed, with $q_j = p_j/(1-p_0)$. We may infer the q_j from the p_j but not p_0 from $q_1, \ldots, q_k$. Show also the ML estimator $\widehat{\theta}$ can be found from the multinomial $(N_1, \ldots, N_k)$ given R, followed by $\widehat{N} = R/\{1 - p_0(\widehat{\theta})\}$.

Ex. 5.63. *Capture-recapture.* (xx nils thinks this exercise goes away, with the material taken care of inside Story iii.13. xx) The methods developed in Ex. 5.62 for hidden boxes in multinomial setups apply in particular to the estimation of closed populations, via capture-recapture. The case dealt with in Ex. 2.87, involving boxes 00, 01, 10, 11, leading to the Petersen (1896) estimator, is a special case, and we now apply better and more versatile methods. The fish population is $\{1, \ldots, N\}$, and the task is to estimate N based on capture, mark, release, recapture.

(a). With capture probabiities p and q for two independent rounds, this leads to a Venn diagram of $N_{0,0}, N_{0,1}, N_{1,0}, N_{1,1}$, with $N_{1,1}$ the number of fish caught in both rounds, etc. The number $N_{0,0}$ of fish caught in neither of the two rounds is not observed. Show that the this amounts to a mutlinomial setup with probabilities $p_{0,0} = (1-p)(1-q)$, $p_{0,1} = (1-p)q$, $p_{1,0} = p(1-q)$, $p_{1,1} = pq$. Show that the log-likelihood function may be written

$$\begin{aligned}\ell(N, p, q) \;&=\; \log(N!) - \log((N-R)!) + (N-R)\log p_{0,0} \\ &\quad + N_{0,1}\log p_{0,1} + N_{1,0}\log p_{1,0} + N_{1,1}\log p_{1,1}.\end{aligned}$$

Simplify and find an expression for the profiled $\ell_{\text{prof}}(N) = \ell(N, \widehat{p}_N, \widehat{q}_N)$.

(b). The theory of Ex. 5.62 applies, both for reading off confidence intervals from $D(N) = 2\{\ell_{\text{prof}}(\widehat{N}) - \ell_{\text{prof}}(N)\}$ and for assessing the variance of $\widehat{N}$. To this end, find the four log-probability derivatives $\psi_{i,j}(p, q)$, verify that $\sum_{i,j} p_{i,j}\psi_{i,j} = 0$, and show that $M = \sum_{i,j} p_{i,j}\psi_{i,j}\psi_{i,j}^{\mathrm{t}}$ is the diagonal matrix with elements $1/(p(1-p)), 1/(q(1-q))$. From this find also

$$\delta = \psi_{0,0}^{\mathrm{t}} M^{-1} \psi_{0,0} = p/(1-p) + q/(1-q),$$

and show that the limit variance for $N^{1/2}(\widehat{N}/N - 1)$ is $\tau^2 = (1-p)(1-q)/(pq)$. This is indeed equal to the limit variance found for $N^{1/2}(N^*/N - 1)$ in Ex. 2.87 for the Petersen estimator N^*. This is no coincidence, show that N^* and $\widehat{N}$ is at most 1 apart. Note that these estimates and confidence intervals are found not only without knowing $N_{0,0}$, but also not knowing the probabilites p and q of fish being caught in the two rounds.

(c). We now go on to three rounds of capture-mark-release, with probabilities p, q, r. In the multinomial setup, with $2^3 = 8$ counts $N_{i,j,k}$, show that $p_{0,0,0} = (1-p)(1-q)(1-r)$, and so on, up to $p_{1,1,1} = pqr$ the probability of a fish being caught in each of the threee rounds. Again, the theory of Ex. 5.62 applies, so it is a matter of working out the details for this setup. Find an explicit formula for the profiled log-likelihood function $\ell_{\text{prof}}(N) = \ell(N, \widehat{p}_N, \widehat{q}_N, \widehat{r}_N)$, amenable for implementation and for reading off confidence intervals for N via the Wilks theorem. Set up formulae for the eight $\psi_{i,j,\ell}$; verify that $\sum_{i,j,\ell} p_{i,j,\ell}\psi_{i,j,\ell} = 0$; and show that $M = \sum_{i,j,\ell} p_{i,j,\ell}\psi_{i,j,\ell}\psi_{i,j,\ell}^{\mathrm{t}}$ is diagonal, with elements $\{p(1-p)\}^{-1}$, $\{q(1-q)\}^{-1}$, $\{r(1-r)\}^{-1}$. Explain that also leads to

$$\delta = \psi_{0,0,0}^{\mathrm{t}} M^{-1} \psi_{0,0,0} = p/(1-p) + q/(1-q) + r/(1-r).$$

Conclude that $(\widehat{N}-N)/N^{1/2} \to_d \mathrm{N}(0,\tau^2)$, with

$$\tau^2 = \frac{(1-p)(1-q)(1-r)}{1-(1-p)(1-q)(1-r)(1+\delta)} = \frac{(1-p)(1-q)(1-r)}{pq+pr+qr-2pqr}.$$

Discuss how the variance of $\widehat{N}$ is influenced by say small, moderate, and higher values of p,q,r. In various applications, from fishery sciences to estimating the number of bank acccount cheaters, these probabilities might tend to be small.

(d). Generalise the above to the case of four independent rounds of capture-mark-release.

Ex. 5.64. *Hardy-Weinberg when seeing only two of the three categories.* (xx perhaps a Simulated Story. xx) We visit Hardy-Weinberg island, in clear equilibritum, where individuals may be sorted into categories 0, 1, 2, with probabilities $(1-\theta)^2, 2\theta(1-\theta), \theta^2$. Suppose however that when accountants arrive, all the zero box people go into hiding, so from the trinomial (N_0,N_1,N_2) only N_1,N_2 are counted. How can we estimate N_0, and the size of the population $N=N_0+N_1+N_2$ investigated?

(a). Show that given $R=N_1+N_2$, N_2 is a binomial (R,q_2), with $q_2=p_1/(p_1+p_2)=\theta/(2-\theta)$. Show that the ML estimators are

$$\widehat{\theta} = \frac{2N_2}{N_1+2N_2}, \quad \widehat{N} = \frac{R}{1-(1-\widehat{\theta})^2} = \frac{R}{1-\{N_1/(N_1+2N_2)\}^2}.$$

(b). We may find confidence intervals via the methods of Ex. 5.62. Explain that the log-likelihood is

$$\begin{aligned}\ell(N,\theta) &= \log\frac{(N!)}{(N-R)!} + 2(N-R)\log(1-\theta) + N_1\{\log\theta+\log(1-\theta)\} + 2N_2\log\theta \\ &= \log(N!) - \log((N-R)!) + (N_1+2N_2)\log\theta + (2(N-R)+N_1)\log(1-\theta),\end{aligned}$$

and that this leads to the profiled log-likelihood

$$\ell_{\mathrm{prof}}(N) = \log(N!) - \log((N-R)!) + (N_1+2N_2)\log\widehat{\theta}_N + (2(N-R)+N_1)\log(1-\widehat{\theta}_N),$$

in which $\widehat{\theta}_N = (N_1+2N_2)/N$.

(c). To find the limiting variance τ^2 of $(\widehat{N}-N)/N^{1/2}$, we need $\psi_0 = -2/(1-\theta)$, $\psi_1 = 1/\theta - 1/(1-\theta)$, $\psi_2 = 2/\theta$. Verify that $\sum_{0,1,2} p_j\psi_j = 0$, and show that

$$M = \sum_{0,1,2} p_j\psi_j^2 = 8 + 8\frac{(\theta-\frac12)^2}{\theta(1-\theta)} = \frac{2}{\theta(1-\theta)}.$$

Derive from this that $\delta = \psi_0^2/M = 2\theta/(1-\theta)$, and that the formula for τ^2 of Ex. 5.62 simplifies to $(1-\theta)^2/\theta^2$. It is hence considerably harder to estimate N with good precision if θ is small, with a high proportion N_0/N in hiding.

(d). (xx round off. variance of $\widehat{N}/N$. may ask for a few cc(N). xx)

Ex. 5.65. *Finding the ML estimate via simulation.* We have seen in Ex. 5.39 that if Y comes from a model of the form $f(y,\theta) = \exp\{\theta^{\mathrm{t}}T(y)-k(\theta)\}h(y)$, then the ML estimate is the solution to $\xi(\theta) = \mathrm{E}_\theta\, T(Y) = t_{\mathrm{obs}}$, where $t_{\mathrm{obs}} = T(y_{\mathrm{obs}})$. To find $\widehat{\theta}$, for a given dataset, there are various options. If the equation is easy to work with, one solves it; in many other situations, one may programme the log-likelihood function, and throw it to an optimiser. There is a third option, however, described now.

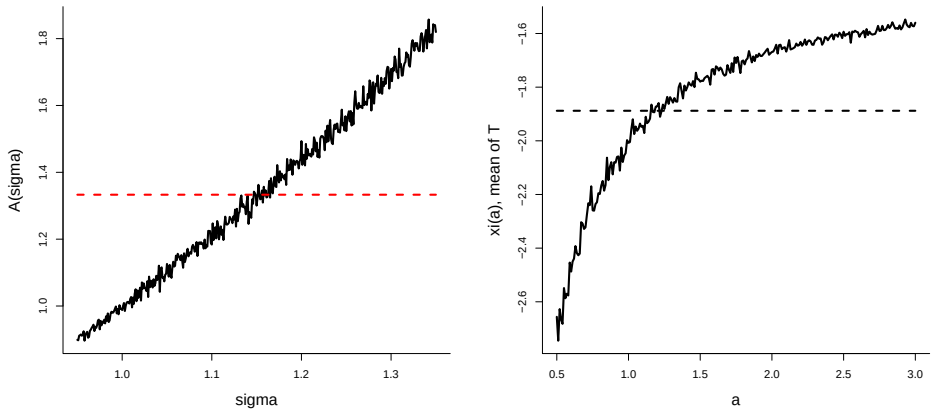


Figure 5.3: *Left panel: the estimated* $A(\sigma) = \mathrm{E}_\theta\, T_n$*, from* 10^3 *simulated outcomes of* T_n *at each* σ*, to see when it is equal to* $t_{n,\mathrm{obs}} = 1.333$*. Right panel: for the Beta-extended normal model, the estimated* $\xi(a) = \mathrm{E}_a\, T_n$*, from* 10^3 *simulated outcomes of* T_n *at each* a*, to see when it crosses* $t_{n,\mathrm{obs}} = -1.888$*.*

(a)**.** Suppose that there is no easy way of solving $\xi(\theta) = t_{\mathrm{obs}}$, and that the log-likelihood function cannot be worked with, perhaps because the $k(\theta)$ is too complicated or impossible to compute. Suppose however that one may *simulate outcomes* for each θ. With a high number B of outcomes Y^*, for a given θ, explain why $\widehat{\xi}(\theta) = (1/B)\sum_{j=1}^B T(Y_j^*)$ is a good (but simulation based) estimate of $\xi(\theta)$. Explain also how one can assess its variance. Argue furthermore that this may be used to solve $\widehat{\xi}(\theta) = t_{\mathrm{obs}}$, as a simulation based estimate of the real ML estimate.

(b)**.** Work through the following simple illustration of this principle. Let $Y_1, \ldots, Y_n$ be i.i.d. $\mathrm{N}(0, \sigma^2)$. Show that the log of the joint likelihood is $-n \log \sigma - \frac{1}{2} n T_n/\sigma^2$, with $T_n = n^{-1}\sum_{i=1}^n y_i^2$, and that the ML estimator admits the exact formula $\widehat{\sigma} = T_n^{1/2}$. Here we pretend that we are not clever enough to do this bit of mathematics, and that we neither are clever enough to programme the log-likelihood function – so we tend to simulations instead. Assume $n = 10$ and that we observe $t_{n,\mathrm{obs}} = 1.333$. Simulate say 10^3 outcomes of T_n for each σ, to compute $A(\sigma) = \mathrm{E}_\sigma\, T_n$. Construct a version of Figure 5.3, left panel. Read off the consequent simulation based estimate σ^*, and compare it to the ecact ML estimate $\widehat{\sigma}_{\mathrm{obs}}$.

(c)**.** Consider the model with c.d.f. $F(y, a) = \mathrm{Be}(\Phi(x), a, a)$, with a an unknown parameter; the standard normal corresponds to $a = 1$. Show that the density may be written

$$f(y, a) = \mathrm{be}(\Phi(y), a, a)\phi(y) = \frac{\Gamma(2a)}{\Gamma(a)^2}[\Phi(y)\{1 - \Phi(y)\}]^{a-1}\phi(y),$$

and that this is of exponential family form with basic data function $T(y) = \log \Phi(y) + \log\{1 - \Phi(y)\}$. Explain that the ML estimate $\widehat{a}$ therefore is the solution to $\xi(a) = \bar{T} = n^{-1}\sum_{i=1}^n T(y_i)$, with $\xi(a) = \mathrm{E}_a\, T(Y_i)$. Assume $\bar{T}_n$ has been observed to be -1.888. Give a version of Figure 5.3, right panel, and estimate the real ML estimate $\widehat{a}$. (xx more comments here. can still do this numerically here. in other situations we can't write

down the log-likelihood function. can zoom in and increase simulation sample size etc. note we found ML value without knowint more than simulating T. xx)

(d). (xx one more example, to make it interesting. point to magis squares Story vi.3, where we have a subquestion on estimating λ from an MCMC run. xx)

Ex. 5.66. *Two-group classification.* (xx a little rant, perhaps to be placed in Ch. 16, or in Ch. 6, but it has relevance to binary regression. xx) Two groups, probabilities π_0 and π_1, with feature vectors having densities f_0 and f_1 from 0 and 1. In this exercise we assume that these two class densities are known; in practice they need to be estimated from training data.

(a). Suppose an object comes from group g, and our decision after seeing its x is $\widehat{g}$. Let the loss function be of symmetric 0-1 type, $L(g,\widehat{g}) = I(\widehat{g} \neq g)$. Show that the optimal Bayes solution is to claim '1' if $p(x) \geq \frac{1}{2}$ and '0' if $p(x) < \frac{1}{2}$, where

$$p(x) = \Pr(g = 1 \,|\, x) = \frac{\pi_1 f_1(x)}{\pi_0 f_0(x) + \pi_1 f_1(x)}.$$

Show that this is equivalent to claiming '1' if $R(x) = \pi_1 f_1(x)/\{\pi_0 f_0(x)\} \geq 1$.

(b). Show that the two group-conditional error rates, using this optimal classifier, is

$$\begin{aligned} \mathrm{err}_0 &= \Pr(\text{saying } 1 \,|\, g = 0) = \int_A f_0 \,\mathrm{d}x, \\ \mathrm{err}_1 &= \Pr(\text{saying } 0 \,|\, g = 1) = \int_{A^c} f_1 \,\mathrm{d}x, \end{aligned}$$

with A the set $\{x\colon R(x) \geq 1\}$. Deduce that the total optimal error rate, seen if this classifier is used for a high number of cases, is

$$\mathrm{err} = \pi_0 \,\mathrm{err}_0 + \pi_1 \,\mathrm{err}_1 = \int_A \pi_0 f_0 \,\mathrm{d}x + \int_{A^c} \pi_1 f_1 \,\mathrm{d}x.$$

(c). Then work out that

$$I = \int |\pi_1 f_1 - \pi_0 f_0| \,\mathrm{d}x = 2 \int_A (\pi_1 f_1 - \pi_0 f_0) \,\mathrm{d}x - (\pi_1 - \pi_0),$$

and from this that

$$\mathrm{err} = \pi_1 - \int_A (\pi_1 f_1 - \pi_0 f_0) \,\mathrm{d}x = \tfrac{1}{2} - \tfrac{1}{2} I.$$

In other words, $\frac{1}{2} + \frac{1}{2} \int |\pi_1 f_1 - \pi_0 f_0| \,\mathrm{d}x$ is the success rate, $\Pr(\widehat{g} = g)$. In particular, for the balanced case of $\pi_0 = \pi_1 = \frac{1}{2}$, show that $\mathrm{err} = \frac{1}{2} - \frac{1}{4} \int |f_1 - f_0| \,\mathrm{d}x$, with success rate $\frac{1}{2} + \frac{1}{4} \int |f_1 - f_0| \,\mathrm{d}x$.

(d). (xx more. Mahalanobis distance. mild but important extension to losses $L_{0,1}$ and $L_{1,0}$. $\mathrm{E}\,\{L(g,0) \,|\, x\} = L_{1,0}(1 - p(x))$ and $\mathrm{E}\,\{L(g,1) \,|\, x\} = L_{0,1} p(x)$. tidy up here. wish to see methods like 'say 1 if $p(x) \geq 0.90$'. xx)

(e). xx

Ex. 5.67. *Estimating in a Weibull model.* Suppose $Y_1, \ldots, Y_n$ are i.i.d. from the one-parameter model with $\Pr(Y_i \leq y) = 1 - \exp(-y^\theta)$ for $y \geq 0$.

(a). Set up a formula for the log-likelihood function. Find the limit distribution for $\sqrt{n}(\widehat{\theta}-\theta)$. (xx this needs: $f(y,\theta)=\exp(-y^\theta)\theta y^{\theta-1}$, with $\log -y^\theta+\log\theta+(\theta-1)\log y$, then score function

$$u(y,\theta)=1/\theta+\log y-y^\theta\log y=(1/\theta)(1+\log v-v\log v),$$

in terms of $v=y^\theta$, a unit exponential. xx)

(b). (xx Analyse the log-likelihood-ratio function. no simple sufficient statistic. check wilks approximation. confidence interval, tests. xx)

(c). xx

Ex. 5.68. *Estimation and testing in the Cauchy location model.* (xx need repair and cleaning. xx) Consider the Cauchy density model $f(y,\theta)=(1/\pi)\{1+(y-\theta)^2\}^{-1}$; see Ex. 1.18.

(a). Show that the score function becomes $\partial\log f(y,\theta)/\partial\theta=2(y-\theta)/\{1+(y-\theta)^2\}$, and hence that the Fisher information can be expressed as $J(\theta)=\mathrm{Var}\,\{2X/(1+X^2)\}$, with X a unit Cauchy, with density $(1/\pi)(1+x^2)^{-1}$. To find $J(\theta)$, use Ex. 1.18 to see that X can be represented as U/V, with U and V independent standard normals. Write $U=R\cos\theta$ and $V=R\sin\theta$, to get to $2X/(1+X^2)=2\cos\theta\sin\theta$, with θ uniform on $[0,2\pi]$. Then show that $J(\theta)=\frac12$. With i.i.d. data $Y_1,\ldots,Y_n$ from the Cauchy model, therefore, show for the ML estimator that $\sqrt{n}(\widehat{\theta}_{\rm ml}-\theta)\to_d \mathrm{N}(0,\tau^2)$, with $\tau=\sqrt{2}$.

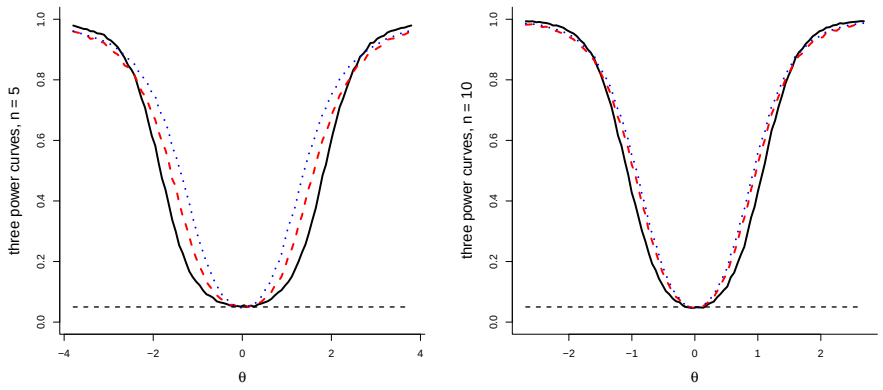


Figure 5.4: *Power functions for three tests for the Cauchy location model, based on A_n (black, full), B_n (red, slanted), D_n (blue, dotted) of Ex. 5.68. The sample sizes are $n=5$ (left panel) and $n=10$ (right panel). Note that the θ scale is different in the two panels; from $n=5$ to $n=10$, the three powers have increased.*

(b). Let $X_1,\ldots,X_n$ be i.i.d. from the Cauchy location model, and suppose one needs to test θ_0 vs. $\theta\neq0$. For the $\widehat{\theta}_{\rm ml}$, we already know its limit distribution, but numerical care is needed to maximise the log-likelihood function $\sum_{i=1}^n[-\log\{1+(x_i-\theta)^2\}]$. Show, via concrete data examples of size e.g. $n=5$, that the log-likelihood function can have more than one maximum, i.e. several bumps. (xx calibrate here, where is what. xx)

(c). For testing the null, at level say 0.05, three choices are as follows: (i) Compute $A_n = \sqrt{n}\widehat{\theta}/\tau$, and reject when $|A_n| \geq a_n$, the 0.95 point in the null distribution of $|A_n|$. (ii) Compute $B_n = \sqrt{n}\widehat{\theta}/\widehat{\tau}$, and reject when $|B_n| \geq a_n$, the 0.95 point in the null distribution of $|B_n|$; here $\widehat{\tau}^2 = -n^{-1}\partial^2\ell_n(\widehat{\theta})/\partial\theta$ is the normalised Hessian associated with the maximisation procedure for finding $\widehat{\theta}$. (iii) Compute $D_n = 2\{\ell_{n,\max} - \ell_n(0)\}$, and reject when $D_n \geq d_n$, the 0.95 point in the null distribution of D_n. Show that $a_n \to 1.96$, $b_n \to 1.96$, $d_n \to \Gamma_1^{-1}(0.95) = 1.96^2$, as n increases. For our power analysis below we have used exact non-asymptotic values of a_n, b_n, d_n, however, computed via simulations. (xx a bit more: need perhaps $n \geq 30$ for asymptotics to have kicked in properly. xx)

(d). We may now compute the associated power functions

$$\pi_{A,n}(\theta) = \Pr_\theta(|A_n| \geq a_n),\ \pi_{B,n}(\theta) = \Pr_\theta(|B_n| \geq b_n),\ \pi_{D,n}(\theta) = \Pr_\theta(D_n \geq d_n),$$

via simulation. Produce a version of Figure 5.4, where we have used $n = 5$ and $n = 10$. With these two n, and with 10^5 simulations for the null distributions, we arrived at $2.959, 2.387$ for a_n, $2.981, 2.291$ for b_n, $2.144^2, 2.012^2$ for d_n.

(e). (xx round off, some comments. moral: direct $\sqrt{n}\widehat{\theta}/\tau$ loses to the less direct $\sqrt{n}\widehat{\theta}/\widehat{\tau}$, even though τ is a known quantity, for the main area around the null value. but it does win, over the B_n and D_n, at a certain distance away from the null, where the power is high for all competitors. deviance test a bit better than the others, for the main neighbourhood. but differences are small. also: more equal powers with $n = 10$ than with $n = 5$, etc. xx)

(f).

Ex. 5.69. *Gamma regression.* (xx a few versions of gamma regression. also 'doubly linear', in two parameters. this could mean $Y_i \sim \text{Gam}(a_i, b_i)$ with $a_i = \exp(x_i^{\rm t}\beta)$ and $b_i = \exp(z_i^{\rm t}\gamma)$. xx)

Ex. 5.70. *Estimating cycle length.* (xx belongs in Ch. 12, depending on what else is included there. xx) Consider the model with mean structure $m_i = a\cos(2\pi i/\omega + \phi)$, common in time series models, with ω related to the length of a cycle. Assume first that a, ϕ are known. Work out the asymptotics for the ML estimator $\widehat{\omega}$. (xx answer: it is surprisingly sharp. xx) Then assume all three parameters need to be estimated. (xx again: ω can be estimated with high precision i think. xx)

(a). For the case of known a, ϕ, and cycle length ω estimated by minimising $Q_n(\omega) = \sum_{i=1}^n \{y_i - a\cos(2\pi i/\omega\phi)\}^2$, show that

$$n^{3/2}(\widehat{\omega} - \omega) \to_d \frac{\sqrt{6}}{2\pi}\omega^2\frac{\sigma}{a}\text{N}(0,1).$$

Ex. 5.71. *Estimation for a triangular distribution.* Consider the triangular distribution on the unit interval, with mode parameter θ in $(0,1)$. We define this by the density $f(x,\theta)$ starting and ending at zero, increasing linearly from 0 to θ and then decreasing linearly from θ to 1. (xx can it really be, from Nguyen and McLachlan 2016, 2017, that numerical schemes are investigated for finding ML, but not theory. surely there is limiting normality here, for us to find and check. easy for moment estimator. xx)

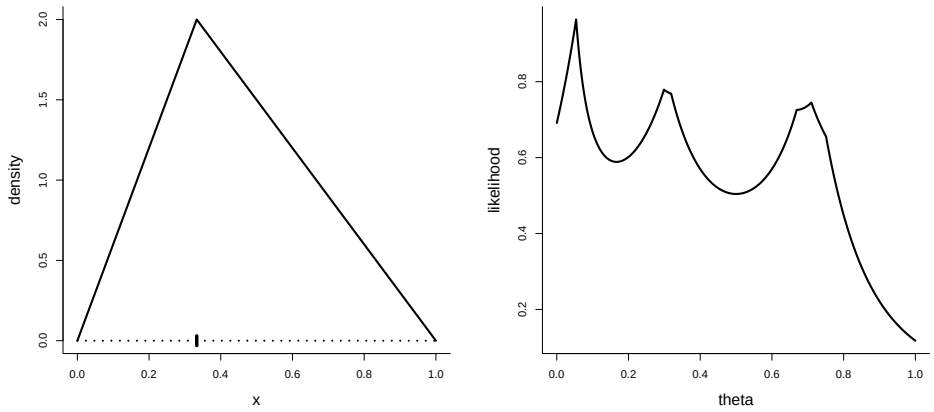


Figure 5.5: *Left panel: Triangular density $f(x, \theta_0)$, with mode at $\theta_0 = 0.333$. Right panel: The likelihood function, with data points* `0.054, 0.299, 0.319, 0.669, 0.710, 0.751`.

(a). Writing $f(x, \theta) = ax$ for $x \le \theta$, show that we then must have $f(x, \theta) = a\theta - b(x - \theta)$ for $x \ge \theta$, and that $b = a\theta/(1 - \theta)$. Derive an expression for the c.d.f. $F(x, \theta)$; use this to identify a as determined by θ; and deduce that

$$f(x, \theta) = \begin{cases} 2x/\theta & \text{if } x \le \theta, \\ 2 - 2(x - \theta)/(1 - \theta) = 2(1 - x)/(1 - \theta) & \text{if } x > \theta. \end{cases}$$

Note that the mode is 2, and that the two triangles have probability masses θ and $1 - \theta$. Show that (xx check this xx) X has mean $(1 + \theta)/3$ and variance $(1 + \theta^2 - \theta)/18$.

(b). Show furthermore that the c.d.f. becomes

$$F(x, \theta) = \begin{cases} x^2/\theta & \text{if } x \le \theta, \\ \theta + 2(x - \theta) - (x - \theta)^2/(1 - \theta) = 1 - (1 - x)^2/(1 - \theta) & \text{if } x > \theta. \end{cases}$$

Show that solving $F(x, \theta) = u$ leads to

$$x = \begin{cases} (\theta u)^{1/2} & \text{if } u \le \theta, \\ 1 - (1 - \theta)\{1 - (u - \theta)/(1 - \theta)\}^{1/2} & \text{if } u > \theta. \end{cases}$$

Use this to generate 1000 data points from the triangular distribution, with say $\theta = 0.345$, and check with the histogram that the recipe works.

(c). (xx moment estimator: $3\bar{X} - 1$, with variance. xx)

(d). (xx then ML estimator. log-likelihood function non-smooth. xx) Show that the likelihood function may be written

$$L_n(\theta) = 2^n \prod_{x_i \le \theta} \frac{x_i}{\theta} \prod_{x_i > \theta} \frac{1 - x_i}{1 - \theta}.$$

Construct a version of Figure 5.5, for a case with the small dataset `0.054, 0.299, 0.319, 0.669, 0.710, 0.751`. To work with this non-smooth likelihood, consider the log-density

$$\log f(x_i,\theta) = (\log x_i - \log\theta)I(x_i \le \theta) + \{\log(1-x_i) - \log(1-\theta)\}I(x_i > \theta),$$

seen as a function of θ for given x_i. We need to work with the log-density difference

$$\begin{aligned} E_i = \log\frac{f(x_i,\theta_0+\varepsilon)}{f(x_i,\theta_0)} &= \log\frac{x_i}{\theta_0+\varepsilon}I(x_i \le \theta_0+\varepsilon) + \log\frac{1-x_i}{1-\theta_0-\varepsilon}I(x_i > \theta_0+\varepsilon) \\ &\quad - \log\frac{x_i}{\theta_0}I(x_i \le \theta_0) - \log\frac{1-x_i}{1-\theta_0}I(x_i > \theta_0) \end{aligned}$$

seen as a function of θ for given x_i. Letting $D(x_i) = -1/\theta_0$ for $x_i \le \theta_0$ and $1/(1-\theta_0)$ for $x_i > \theta_0$, show that $D(X_i)$ has mean zero and variance $J = 1/\{\theta_0(1-\theta_0)\}$. Then show that (xx the following to be made clearer xx)

$$E_i = D(x_i)\varepsilon + R(x_i,\varepsilon)\varepsilon^2,$$

where $R(x_i,\varepsilon) = \frac12/\theta_0^2$ for $x_i \le \theta_0$ and $\frac12/(1-\theta_0)^2$ for $x_i > \theta_0$. (xx then a bit more for showing what we need to show. xx)

(e). show that the log-likelihood function may be expressed as

$$\ell_n(\theta) = \sum_{x_i\le\theta}(\log x_i - \log\theta) + \sum_{x_i>\theta}\{\log(1-x_i) - \log(1-\theta)\} = n\{A_n(\theta) + B_n(\theta)\},$$

omitting a $\log 2$ constant, where $A_n(\theta) = n^{-1}\sum_{i=1}^n \log(x_i/\theta)\, I(x_i \le \theta)$ and $B_n(\theta) = n^{-1}\sum_{i=1}^n \log((1-x_i)/(1-\theta))\, I(x_i > \theta)$. (xx should now find clear limit $L(\theta,\theta_0)$ of $n^{-1}\ell_n(\theta)$, when data stem from θ_0, and then show that L is maximised for $\theta = \theta_0$. (xx more via nilswork24 com13* which verifies that ML works. need Hjort Pollard to prove normal limit. xx)

Notes and pointers

(xx some pointers to more, including Empirical Likelihood, Hjort et al. (2009, 2018). xx)

(xx Bhattacharyya, 1946, lower bounds more general than CR. xx)

(xx make sure we have something on board with BHHJ for regression models, placed after ML for regression. xx)

[xx CR bound: In its simplest form, the inequality goes back to Cramér (1946) and Rao (1945). xx]

(xx M-estimators, Z-estimators, Huber, minimum divergence, point to Basu et al. (2011), more on minimum divergence. two-stage estimators, WalkerHjort24. xx)

[xx least false: a term invented by Hjort, Hjort believes, see Hjort (1986b, 1992), and now used somewhat frequently in the literature. xx]

Read more about risk functions in DeGroot (1970).

For Ex. 5.11, point to Hjort and Jones (1996), Schweder and Hjort (2016), and WalkerHjort24.

[xx check and calibrate what's here and what's in Ch. 7, regarding CD things. xx]

(xx we see that ML matches CR bounds, under model conditions. point to Hajek convolution theorem, and other characterisations. perhaps Hodges superefficient thing too. xx)

(xx point to Hjort (2008), re ML and least false etc. xx)

I.6

Bayesian inference and computation

In frequentist parametric inference, there is a fixed underlying true parameter value, say θ_0, and methods aim at estimating this value, perhaps along with confidence regions or testing. Bayesian inference is radically different, conceptually and operationally. It starts with a prior distribution for the model parameter θ, and proceeds via Bayes theorems to produce the posterior distribution, of the full θ or of relevant focus parameters. Thus 'not knowing θ well' is expressed in terms of probability distributions. This chapter goes through these concepts and operations, including also computational schemes to simulate outcomes from the posterior distributions. One prominent class of such schemes amounts to setting up a Markov Chain Monte Carlo algorithm for the given problem, where the stationary distribution for the chain is precisely the required posterior. This makes Bayesian inference possible in a host of complicated setups, without needing to rely on mathematically feasible formulae.

Key words: Bayes solutions, Bernshteĭn–von Mises theorems, conjugate priors, Jeffreys prior, loss functions, MCMC, prior to posterior distributions

The Bayesian paradigm is to formulate uncertainty about model parameters through probability distributions. If the pre-data uncertainty is a prior density $\pi(\theta)$, this is updated to the post-data posterior density $\pi(\theta \,|\, \text{data})$, via the Bayes theorems.

Consider for illustration and clarification the classical coin flipping experiment, with θ the probability of 'head up'. With n independent flips we have $Y \sim \text{binom}(n, \theta)$. The frequentist postulates that there is an underlying true θ_0, uses perhaps the estimator $\widehat{\theta} = Y/n$, reaches the 95 percent interval $I_n = \widehat{\theta} \pm 1.96\,\{\widehat{\theta}(1-\widehat{\theta})\}/\sqrt{n}$, etc. The key property here is $\Pr_{\theta_0}(\theta_0 \in I_n) = 0.95$; so I_n is a random interval, covering the true θ_0 in 95 percent of actual cases. The Bayesian viewpoint is strikingly different, starting with a prior density $\pi(\theta)$ to reflect what might be considered understanding of θ before the first flip. Post flipping, the Bayesian has reached $\pi(\theta \,|\, y) \propto \pi(\theta) f(y, n, \theta)$, with the binomial likelihood. This may e.g. be used to construct a 95 percent posterior interval J_n for θ, with $\Pr(\theta \in J_n \,|\, y) = 0.95$. The Bayesian is then not interested in 'independent repeated experiments', but just in the data at hand. She is also allowed the statistical luxury of putting prior knowledge into the analysis; if it can be considered known that θ must be close to 0.50, with values outside $[0.40, 0.60]$ less likely than 2 percent, that can effectively be utilised in Bayesian analysis, but not so easily in frequentist analysis.

In this chapter we go through the basics of such constructions and methods, conceptually and operationally. We also uncover conditions under which the frequentist and

Bayesian might actually (approximately) agree, in their final inference statements. The two 95 percent intervals I_n and J_n in the previous paragraph will e.g. tend to be very similar, at least with increasing n.

An attractive feature of Bayesian analysis is that the answer, to a sufficiently well-posed inference question, is crystal clear, without having to study competing methods, carrying out performance and comparison analyses, etc. Essentially, if you give a Bayesian (i) data, (ii) a model, (iii) a list of possible actions, (iv) a prior for the unknown parameters, and (v) a loss function, there is a Master Recipe for the very best action.

Modern Bayesian statistics has flourished since around 1980, partly through computer power and algorithms, making calculations possible that would have been too hard for previous generations. The operational goal is often to be able to generate samples from the posterior distribution, and we give methods for accomplishing this, including Markov Chain Monte Carlo (MCMC).

The Bayesian Master Recipe, with examples and applications

Ex. 6.1. *Poisson data with gamma priors.* This exercise illustrates the basic prior to posterior updating mechanism in a simple Poisson setting. Suppose $Y_1, Y_2, \ldots$ are i.i.d. Poisson with unknown mean θ.

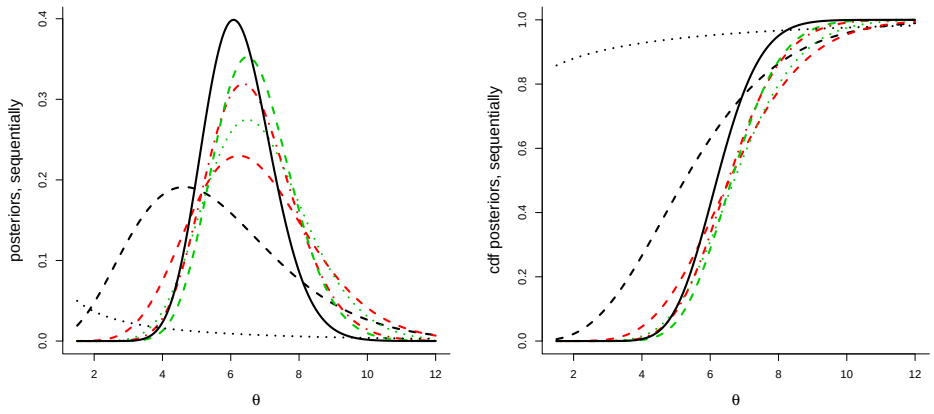


Figure 6.1: *Seven curves are displayed, corresponding to the* Gam$(0.1, 0.1)$ *intial prior for the Poisson parameter* θ*, along with the six first updates following each of the observations 6, 8, 7, 6, 7, 4, 11, 8, 6, 3. The distributions become tighter as more data come in. The smooth black curve is after six data points, the most tight distribution so far. Left panel: the densities; right panel: the c.d.f.s.*

(a)**.** Recall definition and properties of the Gamma distribution from Ex. 1.9. In the present Bayesian context, let $\theta \sim \text{Gam}(a, b)$. The prior mean and variance are $a/b = \theta_0$ and $a/b^2 = \theta_0/b$. In particular, low and high values of b signify high and low variability, respectively. Explain how (a, b) may be set from values of prior mean and prior variance. To exemplify, if these are $(5.5, 7.7)$, find (a, b).

(b)**.** With a single observation Y, which is Pois(θ) given θ, show that $\theta \,|\, y \sim \text{Gam}(a + y, b + 1)$.

(c). Then suppose there are repeated observations $y_1, \ldots, y_n$, being i.i.d. $\sim$ Pois(θ) for given θ. Use the above result repeatedly, e.g. interpreting $p(\theta \,|\, y_1)$ as the new prior before observing y_2, etc., to show that $\theta \,|\, y_1, \ldots, y_n \sim \text{Gam}(a + y_1 + \cdots + y_n, b + n)$. Also derive this result directly, i.e. without necessarily thinking about the data having emerged sequentially.

(d). Suppose the prior used is a rather flat Gam(0.1, 0.1) and that the Poisson data are 6, 8, 7, 6, 7, 4, 11, 8, 6, 3. Reconstruct a version of Figure 6.1 in your computer, plotting the six first posterior densities $p(\theta \,|\, \text{data}_j)$ (left panel), where data_j is $y_1, \ldots, y_j$, along with the prior density; in the right panel we have the corresponding posterior cumulatives. Complement with another figure also including updated densities 7, 8, 9, 10, for the four last observations, and comment. Also compute the ten Bayes estimates $\widehat{\theta}_j = \text{E}(\theta \,|\, \text{data}_j)$ and the posterior standard deviations, for $j = 1, \ldots, 10$.

(e). The mathematics turned out to be rather uncomplicated in this situation, since the Gamma continuous density matches the Poisson discrete density so nicely. Suppose instead that the initial prior for θ is a uniform over $[5.0, 11.0]$. Try to compute posterior distributions, Bayes estimates and posterior standard deviations also in this case, and compare with what you found above. Note also that we in this exercise cared about and managed to reach answers for the posterior distributions $\theta \,|\, y_1, \ldots, y_n$, without needing to deal with the also implied marginal distribution for $(y_1, \ldots, y_m)$. These are not Poisson any longer, with the variability in θ taken into account; see Ex. 6.24.

Ex. 6.2. *The Bayesian Master Recipe.* The general setup is as follows. We have data y_{obs}, seen as the outcome of a random Y over the sample space $\mathcal{Y}$, generated from a model $f_\theta(y) = f(y, \theta)$, with θ having a prior distribution $\pi(\theta)\,\text{d}\theta$ over its parameter space Θ. There is a decision a to be reached, with a belonging to an appropriate action space $\mathcal{A}$, along with a loss function $L(\theta, a)$, a measure of the consequences of decision a if the truth is θ. Whereas the frequentist can attempt different methods for deciding $\widehat{a} = \widehat{a}(Y)$, then compare risk functions, etc., there is a unique optimal strategy for the Bayesian.

(a). Show that the posterior density of θ, that is, the distribution of the parameter given the data, takes the form

$$\pi(\theta \,|\, y) = f_\theta(y)\pi(\theta)/m(y),$$

where $m(y)$ is the required integration constant $\int_\Theta f_\theta(y)\pi(\theta)\,\text{d}\theta$. This is *Bayes' theorem*, and we typically write $\pi(\theta \,|\, y) \propto \pi(\theta) f_\theta(y)$, which reads 'posterior is proportional to prior times likelihood'. Show also that the *marginal distribution* of the data y is $m(y)$.

(b). A decision $a = \widehat{a}(y)$ needs to be reached, as a function of the data. The *Bayes risk* for such a decision is the associated expected loss, $\text{BR}(a, \pi) = \text{E}\, L(\theta, \widehat{a}(Y))$, involving randomness on two levels; θ has a prior, and $\widehat{a}(Y) \,|\, \theta$ is random. Show that it may be expressed in two informative ways:

$$\text{BR}(a, \pi) = \begin{cases} \text{E}\,\{\text{E}\, L(\theta, \widehat{a}(Y)) \,|\, \theta\} = \int_\Theta R(\theta, \widehat{a})\pi(\theta)\,\text{d}\theta, \\ \text{E}\,\{\text{E}\, L(\theta, \widehat{a}(Y)) \,|\, Y\} = \int_{\mathcal{Y}} \{\int_\Theta L(\widehat{a}(y), \theta)\pi(\theta \,|\, y)\,\text{d}\theta\}\, m(y)\,\text{d}y. \end{cases} \tag{6.1}$$

The first expression involves the frequentist risk function $R(\theta, \widehat{a}) = \text{E}_\theta\, L(\theta, \widehat{a}(Y))$, then averaged with respect to the prior. The 'inner expectation' of the second expression is $\text{E}_\pi\,\{L(\theta, \widehat{a}(y)) \,|\, Y = y\}$, that is, the expected loss given data.

(c). Show then that the optimal Bayes strategy, the one minimising the Bayes risk, is achieved by using

the Bayesian Master Recipe

$$\widehat{a} = \operatorname{argmin} g = \text{the value minimising } g,$$

where $g = g(a) = \mathrm{E}_\theta\{L(\theta, a) \,|\, y_{\mathrm{obs}}\}$ is the expected posterior loss. The function g is evaluated and minimised over all a, for the given data $y = y_{\mathrm{obs}}$. This is the Bayes recipe. Note that using the recipe in practice only concerns the observed data y_{obs}, and that one does not need to evaluate its risk function. The *minimum Bayes risk* $\mathrm{MBR}(\pi) = \mathrm{BR}(\widehat{a}, \pi)$ is the resulting Bayes risk when using this optimal strategy.

minimum Bayes risk

(d). Above results have been presented and reached in terms of prior and posterior densities, $\pi(\theta)$ and $\pi(\theta \,|\, y)$, partly for notational convenience. Show that the arguments go through also for more general priors; these may in particular be mixtures over continuous and discrete measures.

Ex. 6.3. *Some loss functions and their associated Bayes rules.* The Master Recipe of Ex. 6.2 is completely general, and can be applied in new and complicated situations, as long as we have data, a model, a prior for the unknowns, and a loss function. In the Bayesian setup finding or evaluation of the posterior distribution of the parameters is always important, carrying separate weight, but if clear decisions are needed one needs also the loss function, say $L(\theta, a)$. Here we go through a short list of commonly used loss functions.

(a). For estimating a one-dimensional θ, with squared error loss $L(\theta, a) = (a - \theta)^2$, show that the Bayes estimator is $\widehat{\theta}_B = \mathrm{E}(\theta \,|\, y)$, the posterior mean. Show also that the minimum Bayes risk, with this loss function, can be expressed as $\mathrm{MBR}(\pi) = \mathrm{E}\,\mathrm{Var}(\theta \,|\, Y)$.

(b). If the loss function is $L(\theta, a) = w(\theta)(a - \theta)^2$, show that the Bayes estimator is

$$\widehat{\theta}_B = \mathrm{E}\{w(\theta)\theta \,|\, \mathrm{data}\}/\mathrm{E}\{w(\theta) \,|\, \mathrm{data}\}.$$

In particular, when estimating a positive θ using loss $(a - \theta)^2/\theta$, show that the Bayes estimator is $1/\mathrm{E}\{(1/\theta) \,|\, \mathrm{data}\}$, the inverse of the inverse mean.

(c). Consider the natural absolute loss function, $L(\theta, a) = |a - \theta|$. Show that the Bayes solution becomes the posterior median, i.e. $\widehat{\theta}_B = G^{-1}(\frac{1}{2} \,|\, \mathrm{data})$, where $G(\theta \,|\, \mathrm{data})$ is the posterior c.d.f.

(d). A differently spirited loss function is $L_\varepsilon(\theta, a) = I(|a - \theta| > \varepsilon)$, for some small ε; zero loss if you're very close, but then a fixed positive loss if you err with more than ε. Show that the Bayes solution is the a which makes the probability $\Pr(\theta \in [a - \varepsilon, a + \varepsilon] \,|\, \mathrm{data})$ as big as possible. Explain that this for small ε means finding the mode of the posterior density.

(e). Suppose one needs the joint estimation of several parameters, say all of $\theta = (\theta_1, \ldots, \theta_p)$, via the loss function $L(\theta, a) = (a - \theta)^{\mathrm{t}} M (a - \theta)$, for an appropriate full-rank symmetric matrix M. Show that the Bayes solution again is the posterior mean, but now for the full vector, i.e. $\mathrm{E}(\theta \,|\, \mathrm{data})$. In particular, the Bayes solution does not depend on the M matrix, though the actual posterior expected loss, and the Bayes risk, do.

(f). In the previous subquestions the framework has been that of estimating a one-dimensional θ. Check that you understand how these results and insights, for the Bayes solutions, change when the situation is changed to that of estimating a *focus parameter*, say $\mu = g(\theta_1, \ldots, \theta_p)$, a function of the full model parameter.

Ex. 6.4. *How many streetcars in San Francisco?* The streetcars in this city are numbered $1, \ldots, N$. You observe $Y = 203$ and wonder what N is.

(a). Supposing Y has the uniform distribution on $1, \ldots, N$, set up the likelihood, and identify the ML estimate. With the prior $\pi(N)$ proportional to $1/N$, say for $N = 1, 2, \ldots, N_{\max}$ for a suitably high $N_{\max}$, find the posterior distribution, along with its mean and median.

(b). Suppose you after having seen no. 203 also see nos. 157, 222. Update your posterior, and again find the mean and median, qua updated estimates.

(c). The Bayesian setup allows any choice for the prior, so think for a minute and construct your own $\pi_{\text{you}}(N)$. Do the updating, for the data 203, 157, 222, and compare with the posterior found above. What is your best estimate of N, if your loss function is of 0-1 type, with $L(N, \widehat{N})$ being 0 if $\widehat{N} = N$ and 1 if $\widehat{N} \neq N$?

Ex. 6.5. *The linex loss function.* When estimating a one-dimensional θ with a $\widetilde{\theta}$, the most traditional loss function is that of squared error, $(\widetilde{\theta} - \theta)^2$, which in particular is symmetric, treating over- and underestimation as equally important. A more flexible loss function is the so-called linex loss, with

$$L_c(\theta, \widetilde{\theta}) = \exp\{c(\widetilde{\theta} - \theta)\} - 1 - c(\widetilde{\theta} - \theta).$$

The c is fine-tuning loss parameter, for the statistician to set, balancing over- against underestimation. Note that both positive and negative values of c are allowed here. There is more on the linex loss in Ex. 8.18, 8.19, but here we deal with the basics.

(a). Show that the L_c is always nonnegative. Show that $c > 0$ means penalising overestimation more than underestimation, and vice versa for $c < 0$. For small $|c|$, show that $L_c(\theta, \widetilde{\theta}) \doteq \frac{1}{2}c^2(\widetilde{\theta} - \theta)^2$, getting back to squared error loss. The constant in front is immaterial for evaluating and comparing loss and risk, and one may use $L_c^*(\theta, \widetilde{\theta}) = L_c(\theta, \widetilde{\theta})/(\frac{1}{2}c^2)$ to have a smoother transition to the $c = 0$ case of squared error loss.

(b). Show that the expected loss given data can be expressed as

$$\begin{aligned} \mathrm{E}\,\{L_c(\theta, t) \,|\, \text{data}\} &= \mathrm{E}\,[\exp\{c(t - \theta)\} - 1 - c(t - \theta) \,|\, \text{data}] \\ &= \exp(ct)M(-c) - 1 - c(t - \widehat{\xi}), \end{aligned}$$

where $\widehat{\xi}$ is the posterior mean and $M(-c) = \mathrm{E}\,\{\exp(-c\theta) \,|\, \text{data}\}$, the m.g.f. of θ given data, computed at $-c$.

(c). Show that this is minimised for the t_0 where $\exp(ct_0)M(-c) = 1$, or $ct_0 + \log M(-c) = 0$, so that the Bayes estimator becomes $\widehat{\theta}_B = -(1/c)\log M(-c)$. This may be computed numerically, perhaps by simulation, in cases where no clear formula exists for $M(-c)$. Show also that the expected posterior loss, using the Bayes solution, is

$$\min \mathrm{E}\,\{L_c(\theta, t) \,|\, \text{data}\} = -c(t_0 - \widehat{\xi}) = \log M(-c) + c\widehat{\xi}.$$

(d). Using approximations for m.g.f.s close to zero to show that $M(-c) \doteq 1 - c\widehat{\xi} + \frac{1}{2}c^2(\widehat{\xi}^2 + \widehat{\sigma})$, with $\widehat{\xi}$ and $\widehat{\sigma}^2$ the posterior mean and variance. Deduce that $\widehat{\theta}_B \doteq \widehat{\xi} - \frac{1}{2}c\widehat{\sigma}^2$.

(e). In situations where the posterior is based on a sample of size n, the posterior mean $\widehat{\xi}_n$ stays stable wheras the posterior variance $\widehat{\sigma}^2$ goes down with speed $1/n$, i.e. as $\widehat{\sigma}_0^2/n$, for the relevant $\widehat{\sigma}_0^2$. In such cases, $\widehat{\xi}_n - \frac{1}{2}c\widehat{\sigma}_0^2/n$ becomes the approximation to the Bayes linex estimator $\widehat{\theta}_B$. Find in fact the exact Bayes linex estimator, for the case of $Y_1, \ldots, Y_n$ being i.i.d. $\mathrm{N}(\theta, 1)$, with a $\mathrm{N}(0, \tau^2)$ prior for θ; here you may very well work from scratch, though it is useful to see the updating results from Ex. 6.16(b).

Ex. 6.6. *A Bayesian take on hypothesis testing.* Assume the model parameter θ is either in Θ_0, which we may call the null hypothesis, or not, i.e. in its complement Θ_0^c. Suppose also that the statistician needs to make a decision, either to reject the null, or to accept it. This is the basic framework of hypothesis testing, see Ch. 4, but we now consider the problem from a Bayesian viewpoint.

(a). The decision space is $\{\text{accept}, \text{reject}\}$. For the loss function, take $L(\theta, \text{accept})$ equal to 0 or L_0, if θ is inside or outside Θ_0, and $L(\theta, \text{reject})$ equal to 0 or L_1, if θ is outside or inside Θ_0. Show that

$$\mathrm{E}\{L(\theta, \text{accept}) \,|\, \text{data}\} = L_0\, p(\text{data}), \quad \mathrm{E}\{L(\theta, \text{reject}) \,|\, \text{data}\} = L_1\{1 - p(\text{data})\},$$

where $p(\text{data}) = \Pr(\theta \in \Theta_0^c \,|\, \text{data})$, the probability that the null is wrong, as measured by the Bayesian posterior distribution.

(b). Deduce that one should reject the null when the probability $p(\text{data})$ for its falseness is sufficiently overwhelming, namely when $p(\text{data}) \geq L_1/(L_0 + L_1)$. – If this threshold is 0.95, for example, show that this corresponds to $L_1/L_0 = 19$. Briefly discuss ways of assigning losses L_0 and L_1.

(c). Suppose now there is a 'doubt option', refraining from deciding on the either-or dichotomy. The setting is as above, with losses L_0 and L_1, but the action list is $\{\text{accept}, \text{reject}, \text{doubt}\}$, with a fixed loss L_d for deciding to be in doubt. Explain that the expected losses given the data are $L_0 p(\text{data})$, $L_1\{1 - p(\text{data})\}$, L_d, and spell out what the Bayes solution is.

Ex. 6.7. *Which subset does the model parameter belong to?* Consider a setup with data from a model with model parameter θ inside its region Θ. Suppose you need to take one of five different possible decisions, $D_1, \ldots, D_5$, and that these are related to where the underlying parameter θ is positioned; if $\theta \in \Theta_j$ the best decision would be D_j, for $j = 1, \ldots, 5$. Here the Θ_j are disjoint and their union is the full parameter region.

(a). Suppose the loss function $L(\theta, D_j)$ is 0 if $\theta \in \Theta_j$ and 100 if $\theta \notin \Theta_j$. Show that $\mathrm{E}\{L(\theta, D_j) \,|\, \text{data}\} = 100\{1 - p_j(\text{data})\}$, where $p_j(\text{data}) = \Pr(\theta \in \Theta_j \,|\, \text{data})$. Hence show that the optimal Bayes strategy is to take the decision associated with the highest posterior probability $p_j(\text{data})$.

(b). Assume there in addition is a 'doubt option', associated with a doubt cost $L_d = 10$; this could e.g. mean planning for getting further data. With decision space $\{D_1, \ldots, D_5, \text{doubt}\}$, what is now the Bayesian strategy?

(c). Generalise the previous setup, and results, to the case where the costs associated with reaching the wrong decision are not equally balanced, say $L(\theta, D_j) = c_{i,j}$, if $\theta \in \Theta_i$, for $i = 1, \ldots, 5$, with $c_{i,i} = 0$ but the other $c_{i,j}$ positive.

Binomial and beta, multinomial and Dirichlet

Ex. 6.8. *The binomial-beta setup.* Let Y given θ be a binomial (n,θ), and for θ take a Beta(a,b) prior, see Ex. 1.22. There we worked with the marginal distribution of Y, and looked at certain properties, but here our aims are Bayesian.

(a)**.** Show that $\theta \,|\, y \sim \text{Beta}(a+y, b+n-y)$. – This is the main and always crucial updating step, getting from the prior to the posterior. In the present case the step is an easy one, since there is only one unknown parameter, and since the product of the prior and the likelihood takes an easy form. Give a description of the posterior also for the not quite so standard case where the prior for θ is uniform on $[0.30, 0.70]$.

(b)**.** Going back to the Beta(a,b) prior again, show that the Bayes estimator, under squared error loss, is $\widehat{\theta}_B = (1-w_n)\theta_0 + w_n y/n$, where $\theta_0 = a/(a+b)$ is the prior mean and $w_n = n/(a+b+n)$. For the case of a uniform prior, show that this leads to $(y+1)/(n+2)$. Compute the risk functions $r(\theta) = \text{E}_\theta\,(\widehat{\theta}-\theta)^2$, for the classic frequentist Y/n and for the this $(Y+1)/(n+2)$, and find the interval where the latter is better than the former.

(c)**.** It is helpful to parametrise (a,b) as $(c\theta_0, c(1-\theta_0))$, with prior mean θ_0 and prior variance $\theta_0(1-\theta_0)/(c+1)$. Explain that for the Bayes estimator and the posterior variance we then have

$$\widehat{\theta}_B = \text{E}\,(\theta \,|\, y) = \frac{c\theta_0 + y}{c+n}, \quad \text{Var}\,(\theta \,|\, y) = \frac{\widehat{\theta}_B(1-\widehat{\theta}_B)}{n+c+1}.$$

Use results from Ex. 1.22 to show that with its marginal distribution, $\text{Var}\, Y = n\theta_0(1-\theta_0)(c+n)/(c+1)$.

(d)**.** Then use $\text{Var}\,\theta = \text{E}\,\text{Var}\,(\theta \,|\, Y) + \text{Var}\,\text{E}\,(\theta \,|\, Y)$ to find that with the Beta $(c\theta_0, c(1-\theta_0))$ prior, the minimum Bayes risk is $\text{MBR} = \theta_0(1-\theta_0)c/\{(c+1)(c+n)\}$. Show that the maximum of the minimum Bayes risk happens for the prior Beta$(\frac12\sqrt{n}, \frac12\sqrt{n})$, with value $\frac14/(\sqrt{n}+1)^2$, and consequent Bayes estimator

$$\widehat{\theta}_B = \frac{\frac12\sqrt{n}+Y}{\sqrt{n}+n} = \frac{1}{\sqrt{n}+1}\frac12 + \frac{\sqrt{n}}{\sqrt{n}+1}\frac{Y}{n}. \tag{6.2}$$

This is the *least favourable prior*; explain that if Nature chooses θ from this distribution, then this is the worst of all Beta scenarios for the statistician. The $\widehat{\theta}_B$ estimator has a corresponding *minimax* property, no competing estimator has a smaller maximum risk; see material on this in Ch. 8.

(e)**.** It is useful to see the minimum Bayes risk in two ways, via (6.1), i.e. also via the risk function of the Bayes solution. Show that this risk function becomes

$$r_B(\theta) = \text{E}_\theta\,(\widehat{\theta}_B - \theta)^2 = \frac{1}{(c+n)^2}\{n\theta(1-\theta) + c^2(\theta-\theta_0)^2\}.$$

Then take the mean of this with respect to the prior to find the MBR formula above again. For the least favourable prior, show that the risk function is constant, equal to $\frac14/(\sqrt{n}+1)^2$.

(f). If Y is from the binomial $(n, \theta_{\text{true}})$ model, show that Y/n and the Bayes estimator $\widehat{\theta}_B$ are large-sample equivalent, with $\sqrt{n}(Y/n - \widehat{\theta}_B) \to_{\text{pr}} 0$, as long as $c = a + b$ stays fixed or bounded. Deduce that they have the same limit distribution. This does not hold for the least favourable prior, however, as its parameters $(\frac{1}{2}\sqrt{n}, \frac{1}{2}\sqrt{n})$ then move with sample size; find the limit distribution of $\sqrt{n}(\widehat{\theta}_B - \theta_{\text{true}})$ for this case.

Ex. 6.9. *Binomial-beta with weighted loss function.* One's best decisions depend on how the consequences are to be evaluated. Consider as in Ex. 6.8 the $Y \sim \text{binom}(n, \theta)$ situation, but now with the weighted loss function $L(\theta, \widehat{\theta}) = (\widehat{\theta} - \theta)^2/\{\theta(1-\theta)\}$.

(a). Discuss briefly what such a loss function entails, compared to the traditional squared error loss, regarding situations where the θ probabilities are close to the boundaries. For the ML estimator Y/n, show that the risk function is the constant $1/n$.

(b). With θ having any prior, show that the Bayes estimate is

$$\widehat{\theta}_B = \text{E}\Big(\frac{1}{1-\theta}\,|\,y\Big)/\text{E}\Big(\frac{1}{\theta(1-\theta)}\,|\,y\Big).$$

If θ has a Beta(a, b) disribution, show that the formula for $\text{E}\,\{1/(1-\theta)\}/\text{E}\,1/\{\theta(1-\theta)\}$ becomes $(a-1)/(a+b-2)$, provided $a > 1$ and $a + b > 2$. With the uniform prior, show that this again leads to $\widehat{\theta}_B = y/n$, first for cases $y = 1, \ldots, n-1$, then also for cases $y = 0$ and $y = n$, where some extra care is required for the two boundary cases.

Ex. 6.10. *Bayes estimators are biased.* Excpet for a few particular cases, perhaps involving special loss functions, the Bayes estimator is not unbiased; it stems from a prior, which is making its inherently biased presence felt. To be specific, suppose a one-dimensional θ is to be estimated with quadratic loss $(a-\theta)^2$, based on observations Y, and that one starts with a proper prior for the θ. We assume, for the arguments to go through smoothly, that θ has a finite prior variance.

(a). We have seen in Ex. 6.3 that the Bayes estimator is $\widehat{\theta}_B(y) = \text{E}\,(\theta\,|\,y)$. Suppose it is unbiased, $\text{E}_\theta\,\widehat{\theta}_B(Y) = \theta$ for all θ. Go through the details of the following arguments, where the expectation operator refers to the joint distribution of (θ, Y):

$$\text{E}\,(\widehat{\theta}_B - \theta)^2 = \text{E}\,\theta^2 - \text{E}\,\theta\widehat{\theta}_B - \text{E}\,\widehat{\theta}_B\theta + \text{E}\,\widehat{\theta}_B^2 = \text{E}\,\theta^2 - \text{E}\,\theta^2 - \text{E}\,\widehat{\theta}_B^2 + \text{E}\,\widehat{\theta}_B^2 = 0.$$

Conclude that unbiasedness can only take place in the trivial case where $\widehat{\theta}_B = \theta$ with probability 1, which would mean θ is known.

(b). There are various situations where the posterior mean estimator is close to being unbiased, but by the argument above there will be some remaining bias. In the binomial setup of Ex. 6.8, the Bayes estimator is $(c\theta_0 + Y)/(c+n)$, which becomes the familiar unbiased Y/n for $c \to 0$; this limiting case, corresponding to a prior proportional to $\theta^{-1}(1-\theta)^{-1}$, is not a proper prior (finite integral), however. We learn that posterior mean estimators computed under certain improper priors (infinite integral) might be unbiased. Note that the arguments used in (a) involve double expectation rules, which presuppose a proper prior.

(c). Explain why deriving the Bayes estimator Y/n in Ex. 6.9 does not contradict what is found above.

Ex. 6.11. *The multinomial-Dirichlet setup.* Here we extend the setup and result of binomial-Beta, to the case of three or more categories. We start with $Y = (Y_1, \ldots, Y_k)$ which for given $p = (p_1, \ldots, p_k)$ is a multinomial $(n, p_1, \ldots, p_k)$. For p we take the $\text{Dir}(a_1, \ldots, a_k)$ prior. For details regarding the multinomial and the Dirichlet, see Ex. 1.5 and 1.20.

(a). Show the important and useful result that $(p_1, \ldots, p_k) \,|\, (y_1, \ldots, y_k) \sim \text{Dir}(a_1 + y_1, \ldots, a_k + y_k)$.

(b). Show that the Bayes estimator under squared error loss becomes

$$\widehat{p}_{i,B} = \text{E}\,(p_i \,|\, \text{data}) = \frac{a_i + y_i}{a + n} = (1 - w_n)p_{0,i} + w_n(y_i/n)$$

for $i = 1, \ldots, k$, with prior means $p_{0,i} = a_i/a$, with $a = a_1 + \cdots + a_k$, and weight $w_n = n/(a + n)$. Find also the posterior variance and posterior correlation between p_i and p_j.

(c). Find the Bayes estimators, under quadratic loss, for the geometric mean parameters $(p_1 p_2)^{1/2}$ and $(p_1 \cdots p_k)^{1/k}$.

Ex. 6.12. *Minimax estimation of multinomial parameters.* We shall consider estimation in the general multinomial model, but for concreteness start with the trinomial case. Consider therefore (X, Y, Z) trinomial (n, p, q, r), with $p + q + r = 1$. We work with the loss function

$$L((p, q, r), (a, b, c)) = (a - p)^2 + (b - q)^2 + (c - r)^2$$

for estimating (p, q, r) jointly.

(a). Show that using a Dirichlet (cp_0, cq_0, cr_0) for (p, q, r) leads to Bayes estimators $\widehat{p} = (cp_0 + x)/(c + n)$, $\widehat{q} = (cq_0 + y)/(c + n)$, $\widehat{r} = (cr_0 + z)/(c + n)$, and that the associated risk function becomes

$$\frac{1}{(c+n)^2}\{np(1-p) + c^2(p - p_0)^2 + nq(1-q) + c^2(q - q_0)^2 + nr(1-r) + c^2(r - r_0)^2\}.$$

(b). Show that the minimum Bayes risk, for such a prior, becomes

$$\text{MBR} = \frac{c}{(c+1)(c+n)}\{p_0(1-p_0) + q_0(1-q_0) + r_0(1-r_0)\}.$$

Show further that this is maximised over c, for $c = \sqrt{n}$, and then over all prior means for $(p_0, q_0, r_0) = (1/3, 1/3, 1/3)$. Verify that the maximised minimum Bayes risk becomes $(2/3)/(\sqrt{n}+1)^2$. This invites using the $\text{Dir}(\sqrt{n}(1/3, 1/3, 1/3))$ prior, with ensuing estimators $(\widehat{p}, \widehat{q}, \widehat{r})$ equal to

$$\frac{1/3}{\sqrt{n}+1} + \frac{\sqrt{n}}{\sqrt{n}+1}\frac{x}{n}, \frac{1/3}{\sqrt{n}+1} + \frac{\sqrt{n}}{\sqrt{n}+1}\frac{y}{n}, \frac{1/3}{\sqrt{n}+1} + \frac{\sqrt{n}}{\sqrt{n}+1}\frac{z}{n}.$$

(c). Show that this estimation scheme has constant risk function, equal to $(2/3)/(\sqrt{n}+1)^2$, i.e. identical to the maximised minimum Bayes risk. As for the binomial parallel (6.2), this $(\widehat{p}, \widehat{q}, \widehat{r})$ method has a nice minimax property, see Ex. 8.10.

(d). Generalise the above, from the trinomial to the general multinomial setting.

Ex. 6.13. *Gott würfelt nicht.* For the multinomial-Dirichlet setup of Ex. 6.11, we reached the posterior characterisation $p \,|\, \text{data} \sim \text{Dir}(a_1+y_1,\ldots,a_k+y_k)$. The importance of this lies in the easy usefulness of simulations, where the posterior distribution of any functions of $(p_1,\ldots,p_k)$ may be read off.

(a)**.** Explain how you may simulate e.g. 10^5 vectors $p=(p_1,\ldots,p_k)$ from the posterior distribution, using the characterisation from Ex. 1.20. Concretely, show that one may use $p_1=G_1/G,\ldots,p_k=G_k/G$, with independent $G_1\sim\text{Gam}(a_1+y_1),\ldots,G_k\sim\text{Gam}(a_k+y_k)$, and sum $G=G_1+\cdots+G_k$.

(b)**.** Suppose you throw a certain and perhaps not entirely standard die 30 times and have counts $(2,5,3,7,5,8)$ of outcomes 1, 2, 3, 4, 5, 6. Use either of the priors (i) 'flat', $\text{Dir}(1,1,1,1,1,1)$; (ii) 'symmetric and more confident', $\text{Dir}(3,3,3,3,3,3)$; (iii) 'unwilling to guess', $\text{Dir}(0.1,0.1,0.1,0.1,0.1,0.1)$, for the probabilities $(p_1,\ldots,p_6)$, to assess the posterior distribution of each of the following quantities:

$$\alpha=\frac{p_6}{p_1},\ \beta=(1/6)\sum_{j=1}^6(p_j-1/6)^2,\ \gamma=(1/6)\sum_{j=1}^6|p_j-1/6|,\ \delta=\frac{(p_4p_5p_6)^{1/3}}{(p_1p_2p_3)^{1/3}}.$$

For each of $\alpha,\beta,\gamma,\delta$, and for each of the priors, give the 0.05, 0.50, 0.95 quantile points, from 10^5 simulations from the posterior distributions. You should also plot the posterior densities, for each of the four quantities noting the extent to which the prior influences the results.

(c)**.** For the case of $\alpha=p_6/p_1$, exact numerical simulation is possible, without simulation. Do this, and compare with the answers reached via simulation.

(d)**.** The above priors are slightly artificial in this context, since they do not allow the explicit possibility that the die in question is plain boring utterly simply a correct one, i.e. that $p=p_0=(1/6,\ldots,1/6)$. The priors used hence do not give us the possibility to admit that perhaps $\rho=1$, $\alpha=0$, $\beta=0$, $\gamma=1$, after all. This motivates using a mixture prior which allows a positive chance for $p=p_0$. Redo therefore the Bayesian analysis above, with the same $(2,5,3,7,5,8)$ data, for the prior $\frac12\,\delta(p_0)+\frac12\,\text{Dir}(1,1,1,1,1,1)$. Here $\delta(p_0)$ is the 'degenerate prior' that puts unit point mass at position p_0. Compute in particular the posterior probability that $p=p_0$, and display the posterior distributions of ρ,α,β,γ.

Sampling from the posterior via Markov Chain Monte Carlo

Ex. 6.14. *MCMC, I: simulating from a given distribution.* Suppose we need to simulat realisations from a given and possibly complicated density $f(x)$ on some domain. The Metropolis algorithm works as follows. The task is to form a long chain $x_1,x_2,\ldots$ in your computer, where its limiting distribution is precisely the given f. After having generated x_{old}, decide on x_{new} by (i) deciding on a proposal, x_{prop}, drawn from a suitable distribution symmetric in $(x_{\text{old}},x_{\text{prop}})$, and then (ii) accepting this proposal with probability $p_{\text{accept}}=\min(1,f(x_{\text{prop}})/f(x_{\text{old}}))$. It is to be noted that any factors of proportionality, as with $f(x)=cf_0(x)$ with a perhaps complicated c, cancel out and disappear. In algorithmic terms,

$$x_{\text{new}}=(1-\text{ok})\,x_{\text{old}}+\text{ok}\,x_{\text{new}},\quad\text{where ok}=I(\text{accept})=I(u\le p_{\text{accept}}),$$

writing u for a simulated uniform. Markov chain theory (xx point to Ch. 12 xx) secures that this scheme works, in the sense that the chain converges in distribution to that of π; here, handling distributions with densities, $\Pr(x_n \in A) \to \int_A f(x)\,\mathrm{d}x$ for all sets A.

(a). Set up such a scheme to generate outcomes from the standard normal (ignoring the existence of simpler direct algorithms). Start at any x_0, then draw proposals $x_{\rm prop} \sim \text{unif}[x_{\rm old} - a, x_{\rm old} + a]$, with acceptance probabilities set up via the general scheme above. Run the chain for a suitably long time and check with a fine histogram that the distribution matches the normal. Keep track of the acceptance probabilities and the overall acceptance rate. The theory works and says that the chain will converge to the standard normal, for any positive fine-turning parameter a; explain however in which ways too small or too large values of a will be ineffective.

(b). Consider a somewhat harder challenge, simulating realisations from the density f of $0.05\,\mathrm{N}(-2,1) + 0.90\,\mathrm{N}(0,1) + 0.05\,\mathrm{N}(2,1)$ via MCMC. Set up a chain that converges to this f in distribution; check that the output produces a fine histogram matching this f.

(c). The algorithm works also in higher dimension. Set up a chain that produces outcomes (x,y,z) from the model on the unit cube with density $f(x,y,z) = 1 + \theta(x - \frac12)(y - \frac12)(z - \frac12)$, with $\theta = 3.45$. From the output, read off the correlations, and the probability that $Z \geq (XY)^{1/2}$.

(d). With such MCMC output one typically reads off key quanties, like means, quantiles, probabilities for certain sets, etc. The key here is that sample averages converge, even though the chains being sampled do not produce i.i.d. output. Check with Ex. 3.42 to explain that with such a chain x_j, with $x_m \to_d x$ in distribution, we have $m^{-1}\sum_{j=1}^m h(x_j)$ converging to $h_0 = \int h(x) f(x)\,\mathrm{d}x$, for each function h, though perhaps with higher sampling variability than if the x_j had been pure i.i.d. Explain also that a histogram, say based on 10^4 values of x_j, is a bona fide estimator for the underlying density $f(x)$. These matters of autocorrelated chain output do not concern us much as long as computers give us long chains.

Ex. 6.15. *MCMC, II: simulating from the posterior.* The Metropolis algorithm of Ex. 6.14 is generic, and finds extensive use in Bayesian statistics, where the object is to sample from the posterior distribution, given the data at hand.

(a). With the posterior density written as $\pi(\theta \,|\, \text{data}) = c\,\pi(\theta)L(\theta)$, in terms of prior times likelihood and a proportionaly constant c, show that the Metropolis algorithm can be described as follows. Form a long Markov chain $\theta_1, \theta_2, \ldots$ in your computer, starting anywhere in the parameter region. After having generated $\theta_{\rm old}$, (i) get hold of a proposal $\theta_{\rm prop}$, from a proposal distribution symmetric in $(\theta_{\rm old}, \theta_{\rm prop})$, and then (ii) accept it with probability $p_{\rm accept} = \min(1, r)$, with ratio

$$r = \frac{\pi(\theta_{\rm prop} \,|\, \text{data})}{\pi(\theta_{\rm old} \,|\, \text{data})} = \frac{\pi(\theta_{\rm prop})L(\theta_{\rm prop})}{\pi(\theta_{\rm old})L(\theta_{\rm old})} = \frac{\pi(\theta_{\rm prop})}{\pi(\theta_{\rm old})} \exp\{\ell(\theta_{\rm prop}) - \ell(\theta_{\rm old})\}.$$

Note that the proportionality constant c has disappeared, so for both the prior and the likelihood we merely need to give these up to a proportionality factor.

(b). In the Poisson-gamma setup of Ex. 6.1 we could find the posterior distribution directly. Suppose however that the prior is outside the gammas, say uniform on $[4.0, 8.0]$. Set up an MCMC to sample from the posterior, given data 6, 8, 7, 6, 7, 4, 11. Compute also the mean, median, standard deviation.

(c). As an illustration, generate first $y_1, \ldots, y_n$ from the uniform distribution on the unit interval. For these data, consider the log-linear density model $f(y, \theta_1, \theta_2) = \exp\{\theta_1(y - \frac{1}{2}) + \theta_2(y - \frac{1}{2})^2 - k(\theta_1, \theta_2)\}$, briefly studied in Ex. 1.57, and use independent zero-mean priors for θ_1, θ_2, with standard deviations $1.0, 0.5$. Set up an MCMC for the posterior of (θ_1, θ_2). Plot say 25 densities drawn from the posteriori model in this way.

Bayesian analysis for normal models

Ex. 6.16. *The normal prior and posterior with normal data.* Here we go through the basic steps and results for situations with normal data and normal priors for unknown mean parameters. More elaborate constructions and technical issues are needed when there in addition are unknown parameters in the variance and covariance structure, to be pursued in Ex. 6.18, 6.25.

(a). There are things to think through and to learn from, by working through this very simple setup first. (i) For a single observation Y assume it comes from the $\mathrm{N}(\xi, \sigma^2)$, and take σ as known; (ii) for the unknown mean ξ assume it comes from the prior $\mathrm{N}(\xi_0, \tau_0^2)$, with specified prior parameters ξ_0, τ_0. Show that this leads to a binormal joint distribution for parameter and observation,

$$\begin{pmatrix} \xi \\ Y \end{pmatrix} \sim \mathrm{N}_2\Big(\begin{pmatrix} \xi_0 \\ \xi_0 \end{pmatrix}, \begin{pmatrix} \tau_0^2 & \tau_0^2 \\ \tau_0^2 & \tau_0^2 + \sigma^2 \end{pmatrix}\Big).$$

(b). Use general conditioning results from Ex. 1.44 to infer that

$$\xi \,|\, y \sim \mathrm{N}(\xi_0 + w(y - \xi_0), w\sigma^2), \quad \text{with } w = \tau^2/(\tau^2 + \sigma^2).$$

So w and $1 - w$ are the weights given to the data-based estimate y and the prior guess ξ_0, respectively. Also, w is the reduction factor with which the variance of the prior-free estimator Y, from σ^2 to $w\sigma^2$.

(c). An easy but important extension is to the case of a full sample $Y_1, \ldots, Y_n$ from the $\mathrm{N}(\xi, \sigma^2)$ distribution, independent given the ξ, again with σ taken known and the normal prior $\mathrm{N}(\xi_0, \tau_0^2)$ for the unknown mean. Show that $\bar{Y} = (1/n)\sum_{i=1}^n Y_i$ is sufficient (xx xref here xx), and that

$$\begin{pmatrix} \xi \\ \bar{Y} \end{pmatrix} \sim \mathrm{N}_2\Big(\begin{pmatrix} \xi_0 \\ \xi_0 \end{pmatrix}, \begin{pmatrix} \tau_0^2 & \tau_0^2 \\ \tau_0^2 & \tau_0^2 + \sigma^2/n \end{pmatrix}\Big).$$

Show from this that

$$\xi \,|\, \text{data} \sim \mathrm{N}(\xi_0 + w_n(\bar{y} - \xi_0), w_n\sigma^2/n), \quad \text{with } w_n = \frac{\tau^2}{\tau^2 + \sigma^2/n} = \frac{n\tau^2}{n\tau^2 + \sigma^2}.$$

Again, w_n is both the weight given to the data-based estimate and the factor with which the neutral estimator's variance is reduced, from σ^2/n to $w_n\sigma^2/n$. Note that $w_n \to 1$; discuss how this may be seen as 'the data wash out the prior'.

(d). Discuss the case of a 'flat prior', where τ is taken large, for $\xi \sim \mathrm{N}(\xi_0, \tau^2)$. Show indeed that this leads to the intuitively plausibe $\xi \,|\, \text{data} \sim \mathrm{N}(\bar{y}, \sigma^2/n)$.

(e). In addition to having a coherent updating machine, changing the prior to the posterior, for each new data point, the Bayesian structure implies positive dependence among the observations. From $\mathrm{E}\,(Y_i \,|\, \xi) = \xi$, $\mathrm{Var}\,(Y_i \,|\, \xi) = \sigma^2$, $\mathrm{E}\,(Y_i Y_j \,|\, \xi) = \xi^2$, show that

$$\mathrm{E}\, Y_i = \xi_0, \quad \mathrm{Var}\, Y_i = \sigma^2 + \tau_0^2, \quad \mathrm{cov}(Y_i, Y_j) = \tau_0^2, \quad \mathrm{corr}(Y_i, Y_j) = \frac{\tau_0^2}{\sigma^2 + \tau_0^2},$$

Show also that $\mathrm{Var}\,\bar{Y}_n = \sigma^2/n + \tau_0^2$, and discuss what this means for large n.

(f). Prove first the convenient formula

$$v(\xi - \xi_0)^2 + n(\xi - \bar{y})^2 = (v+n)(\xi - \xi^*)^2 + d_n(\bar{y} - \xi_0)^2,$$

where

$$\xi^* = \frac{v\xi_0 + n\bar{y}}{v+n} \quad \text{and} \quad d_n = \frac{vn}{v+n} = (v^{-1} + n^{-1})^{-1},$$

which also may be written and interpreted via $1/d_n = 1/v + 1/n$.

(g). Show that with any prior $p(\xi)$ for ξ, the marginal density of $(Y_1, \ldots, Y_n)$ can be written

$$\bar{f}(y_1, \ldots, y_n) = \int (2\pi)^{-n/2} \sigma^{-n} \exp\Big\{-\tfrac{1}{2}\frac{1}{\sigma^2}\sum_{i=1}^n (y_i - \xi)^2\Big\} p(\xi)\, \mathrm{d}\xi.$$

check all of this with care

For the case of $\mathrm{N}(\xi_0, \tau_0^2)$ worked with above, let first $Q_0 = \sum_{i=1}^n (y_i - \bar{y})^2$, and verify that $\sum_{i=1}^n (y_i - \xi)^2 = Q_0 + n(\xi - \bar{y})^2$. Writing for mathematical convenience $1/\tau^2 = v/\sigma^2$, and $1/d_n = 1/n + 1/v$, show that

$$\bar{f}(y_1, \ldots, y_n) = \frac{1}{(2\pi)^{n/2}} \frac{1}{\sigma^n} \exp(-\tfrac{1}{2} Q_0/\sigma^2) \Big(\frac{v}{v+n}\Big)^{1/2} \exp\{-\tfrac{1}{2} d_n(\bar{y} - \xi_0)^2/\sigma^2\}.$$

(h). With the marginal density seen as a function of the two fine-tuning parameters (ξ_0, τ_0^2), find the maximum marginal likelihood estimators $\widehat{\xi}_0$ and $\widehat{\sigma}$.

(i). We record two matrix identities here, as they will come in handy both here and on later occasions. For a square and invertible A, show that

$$(A + xx^{\mathrm{t}})^{-1} = A^{-1} - cA^{-1}xx^{\mathrm{t}}A^{-1}, \quad \text{with } c = 1/(1 + x^{\mathrm{t}}A^{-1}x);$$

also, that $|A + xx^{\mathrm{t}}| = |A|(1 + x^{\mathrm{t}}A^{-1}x)$.

(j). Argue directly that $Y = (Y_1, \ldots, Y_n)^{\mathrm{t}}$ must be multinormal, with $Y \sim \mathrm{N}_n(\xi_0 \mathbf{1}, \sigma^2 I_n + \tau_0^2 \mathbf{1}\mathbf{1}^{\mathrm{t}})$, with $\mathbf{1}$ the vector $(1, \ldots, 1)^{\mathrm{t}}$; the variance matrix has $\sigma^2 + \tau_0^2$ on the diagonal and τ_0^2 outside. Show that this agrees with the marginal density formula reached above.

Ex. 6.17. *When underestimation doesn't matter so much.* (xx nils rant; we include if turning out well. xx) Suppose θ is to be estimated, where underestimation is less important than overestimation. One version of this is to use the loss function $L(\theta, a) = L_0(a - \theta)$, with $L_0(u) = u^2$ for $u \geq -c$, but then $L_0(u) = c^2$, for all $u \leq -c$, for a relevant negative threshold $-c$. Thus a big error to the left doesn't matter more than a moderate error to the left.

(a). For any random θ, with density f, c.d.f. F, $E_1(x) = \int_{-\infty}^x \theta f(\theta)\,\mathrm{d}\theta$, $E_2(x) = \int_{-\infty}^x \theta^2 f(\theta)\,\mathrm{d}\theta$, show that

$$\begin{aligned} R(a) = \mathrm{E}\, L_0(a-\theta) &= \int_{-\infty}^{a+c} (a-\theta)^2 f(\theta)\,\mathrm{d}\theta + c^2 \int_{a+c}^{\infty} f(\theta)\,\mathrm{d}\theta \\ &= a^2F(a+c) - 2aE_1(a+c) + E_2(a+c) + c^2(1-F(a+c)), \end{aligned}$$

and that its derivative simplifies to

$$R'(a) = 2aF(a+c) - 2E_1(a+c).$$

Check that the solution a_0 to setting this to zero is equal to the mean of θ in the case of $-c$ being far to the left.

(b). Explain that the Bayes estimator, for a given dataset, is the solution $\widehat{a}$ to the equation $aF(a+c\,|\,\text{data}) = \mathrm{E}_1(a+c\,|\,\text{data})$, featuring the posterior c.d.f. $\Pr(\theta \le a + c\,|\,\text{data})$ and $E_1(a+c\,|\,\text{data}) = \int_{-\infty}^{a+c} v\pi(v\,|\,\text{data})\,\mathrm{d}v$.

(c). For the case of $y_1, \ldots, y_n$ i.i.d. from the $\mathrm{N}(\xi, 1)$, with a flat prior for ξ, show first that $\xi\,|\,\text{data} \sim \mathrm{N}(\bar{y}, 1/n)$, which may be represented as $\xi = \bar{y} + N/\sqrt{n}$, with a standard normal N. Show then that the equation for finding the Bayes estimator is

$$(\widehat{a} - \bar{y})\Phi(\sqrt{n}(\widehat{a} - \bar{y} + c)) = -\phi(\sqrt{n}(\widehat{a} - \bar{y} + c))/\sqrt{n}.$$

Explain that this means that the Bayes estimator pushes the ML estimator $\bar{y}$ to the left, with an amount diminishing with sample size n.

(d). (xx round off. an illustration. xx)

Ex. 6.18. *The gamma-normal prior and posterior.* Let data $y_1, \ldots, y_n$ for given parameters ξ and σ be i.i.d. $\mathrm{N}(\xi, \sigma^2)$. We have seen in Ex. 6.16 that when σ may be taken as a known quantity, the canonical class of priors for ξ is the normal one. When both parameters are unknown, however, as in most practical encounters, a more elaborate analysis is called for.

(a). Show that the likelihood function may be written as being proportional to

$$L_n(\xi, \sigma) = \exp\Big[-n\log\sigma - \tfrac{1}{2}\frac{1}{\sigma^2}\{Q_0 + n(\xi - \bar{y})^2\}\Big],$$

where $\bar{y} = (1/n)\sum_{i=1}^n y_i$ and $Q_0 = \sum_{i=1}^n (y_i - \bar{y})^2$.

(b). With *any* given prior $p(\xi, \sigma)$, explain how you may set up a Metropolis type MCMC to draw samples from the posterior distribution. Try this out in practice, using the prior that takes ξ and $\log\sigma$ independent and uniform on say $[-5, 5]$ and $[-10, 10]$, with data that you simulate for the occasion from a $\mathrm{N}(2.345, 1.234^2)$, with $n = 25$. Note that this approach does not need more mathematical algebra as such, apart from the likelihood function above (along with basic programming skills).

(c). There is however a popular and convenient conjugate class of priors for which posterior distributions become particularly clear, with the appropriate algebraic efforts. These in particular involve placing a Gamma prior on the inverse variance $\lambda = 1/\sigma^2$. Say that (λ, ξ) has the gamma-normal distribution with parameters (a, b, ξ_0, v), and write this as

$(\lambda,\xi) \sim \mathrm{GN}(a,b,\xi_0,v)$, provided $\lambda = 1/\sigma^2 \sim \mathrm{Gam}(a,b)$ and $\xi \,|\, \sigma \sim \mathrm{N}(\xi_0, \sigma^2/v)$. Show that the prior can be expressed as

$$p(\lambda,\xi) \propto \lambda^{a-1}\lambda^{1/2}\exp\big[-\lambda\{b+\tfrac{1}{2}v(\xi-\xi_0)^2\}\big].$$

What is the unconditional prior variance of ξ?

(d). Using the identity from Ex. 6.16(f), show that if the prior is $(\lambda,\xi) \sim \mathrm{GN}(a,b,\xi_0,v)$, then

$$(\lambda,\xi) \,|\, \text{data} \sim \mathrm{GN}(a+\tfrac{1}{2}n, b+\tfrac{1}{2}Q_0+\tfrac{1}{2}d_n(\bar y-\xi_0)^2, \xi^*, v+n).$$

(e). The special case of a 'flat prior' for ξ, corresponding to letting $v \to 0$ above, is particularly easy to deal with. Show that then

$$(\lambda,\xi) \,|\, \text{data} \sim \mathrm{GN}(a+\tfrac{1}{2}n, b+\tfrac{1}{2}Q_0, \bar y, n).$$

Find the posterior mean of σ^2 under this prior.

(f). For an illustration, consider the cigarette consumption data $x_1,\ldots,x_n$, from 2.B, for $n = 44$ US states (actually, 43 states plus the District of Columbia), from the early 1960ies, taken here to form a sample from $\mathrm{N}(\xi,\sigma^2)$. The x_i is the consumption in state i, in hundreds per year, ranging then from 14.0 to 42.4 (which translates to from 3.84 per day to 11.61 per day, in the adult population). Fix the prior that takes (i) $\lambda = 1/\sigma^2 \sim \mathrm{Gam}(a,b)$, with (a,b) taken to correspond to 0.10 and 0.90 prior quantiles for σ being 1.0 and 10.0; (ii) given σ, ξ is a $\mathrm{N}(20.0, \sigma^2/v)$ with $v = 0.50$. Find the precise posterior distribution for (λ,ξ). Give posterior means and 90 percent credibility intervals for ξ, σ, and for the probability $p = \Pr(X \ge 40)$.

(g). For the same data and setup as in the previous point, carry out a second Bayesian analysis, with the simpler prior that corresponds to $v \to 0$. Comment briefly on the differences in results.

Ex. 6.19. *The gamma-multinormal prior for linear regression models.* The aim of the present exercise is to generalise the gamma-normal conjugate prior class above to the linear-normal regression model. The model is the very classical one of Ex. 3.34, where

$$y_i = x_{i,1}\beta_1 + \cdots + x_{i,k}\beta_k + \varepsilon_i = x_i^{\mathrm{t}}\beta + \varepsilon_i \quad \text{for } i = 1,\ldots,n,$$

with the ε_i taken i.i.d. $\mathrm{N}(0,\sigma^2)$. Write X for the $n \times k$ matrix of covariates (explanatory variables), with $x_i = (x_{i,1},\ldots,x_{i,k})$ as its ith row, and use y and ε to indicate the vectors of y_i and ε_i. Then $y = X\beta + \varepsilon \sim \mathrm{N}_n(X\beta, \sigma^2 I_n)$ is a concise way to write the full model.

(a). Show that the likelihood function may be written as being proportional to

$$L_n(\beta,\sigma) = \sigma^{-n}\exp\Big[-\tfrac{1}{2}\frac{1}{\sigma^2}\{Q_0 + n(\beta-\widehat\beta)^{\mathrm{t}}M_n(\beta-\widehat\beta)\}\Big],$$

in which

$$\Sigma_n = (1/n)X^{\mathrm{t}}X = n^{-1}\sum_{i=1}^n x_i x_i^{\mathrm{t}} \quad \text{and} \quad \widehat\beta = (X^{\mathrm{t}}X)^{-1}X^{\mathrm{t}}y = \Sigma_n^{-1}n^{-1}\sum_{i=1}^n x_i y_i.$$

Also,

$$Q(\beta) = \|y - X\beta\|^2 = Q_0 + n(\beta - \widehat{\beta})^{\mathrm{t}}\Sigma_n(\beta - \widehat{\beta}),$$

with $Q_0 = \sum_{i=1}^n (y_i - x_i^{\mathrm{t}}\widehat{\beta})^2$ the minimum value of Q over all β. Note that $\widehat{\beta}$ is the classical least squares estimator (and the ML estimator), which in the frequentist framework is unbiased with variance matrix equal to $\sigma^2(X^{\mathrm{t}}X)^{-1} = (\sigma^2/n)\Sigma_n^{-1}$. This is the basis of all classical methods related to the widely popular linear regression model.

(b). Let $p(\beta, \sigma)$ be *any* prior for the $(k+1)$-dimensional parameter of the model. Set up formulae for a Metropolis type MCMC algorithm for drawing samples from the posterior distribution of (β, σ).

(c). In spite of the possibility of solving problems via MCMC (or perhaps acceptance-rejection sampling), as with the previous exercise it is very much worthwhile setting up explicit formulae for the case of a certain canonical prior class. Write $(\lambda, \beta) \sim \mathrm{GN}_k(a, b, \beta_0, \Sigma_0)$ to indicate the gamma-normal prior where

$$\lambda = 1/\sigma^2 \sim \mathrm{Gam}(a, b) \quad \text{and} \quad \beta \,|\, \sigma \sim \mathrm{N}_k(\beta_0, \sigma^2\Sigma_0^{-1}).$$

Show that this prior may be expressed as

$$p(\lambda, \beta) \propto \lambda^{a-1}\lambda^{k/2} \exp\bigl[-\lambda\{b + \tfrac{1}{2}(\beta - \beta_0)^{\mathrm{t}}\Sigma_0(\beta - \beta_0)\}\bigr].$$

(d). When multiplying the prior with the likelihood it is convenient to use the following linear algebra identity about quadratic forms, which you should prove first. For symmetric and invertible matrices A and B, and for any vectors a, b, x of the appropriate dimension,

$$(x-a)^{\mathrm{t}}A(x-a) + (x-b)^{\mathrm{t}}B(x-b) = (x-\xi)^{\mathrm{t}}(A+B)(x-\xi) + (b-a)^{\mathrm{t}}D(b-a),$$

where $\xi = (A+B)^{-1}(Aa + Bb)$ (a weighted average of a and b) and D is a matrix for which several equivalent formulae may be used:

$$\begin{aligned} D &= A(A+B)^{-1}B = B(A+B)^{-1}A \\ &= A - A(A+B)^{-1}A = B - B(A+B)^{-1}B = (A^{-1} + B^{-1})^{-1}. \end{aligned}$$

(e). Prove that if (λ, β) has the $\mathrm{GN}_k(a, b, \beta_0, \Sigma_0)$ prior, then

$$(\lambda, \beta) \,|\, \text{data} \sim \mathrm{GN}_k(a + \tfrac{1}{2}n, b + \tfrac{1}{2}Q_0 + \tfrac{1}{2}(\widehat{\beta} - \beta_0)^{\mathrm{t}}D_n(\widehat{\beta} - \beta_0), \beta^*, \Sigma_0 + n\Sigma_n),$$

where

$$\beta^* = (\Sigma_0 + n\Sigma_n)^{-1}(\Sigma_0\beta_0 + n\Sigma_n\widehat{\beta}) \quad \text{and} \quad D_n = \Sigma_0(\Sigma_0 + n\Sigma_n)^{-1}n\Sigma_n.$$

This characterisation makes it easy to simulate a large number of (β, σ) from the posterior distribution and hence to carry out Bayesian inference for any parameter of quantity of interest.

(f). Note the algebraic simplifications that result when the Σ_0 in the prior is chosen as being proportional to the covariate sample variance matrix, i.e. $\Sigma_0 = c_0\Sigma_n$. Show that then

$$\beta^* = (c_0\beta_0 + n\widehat{\beta})/(c_0 + n) \quad \text{and} \quad D_n = c_0 n/(c_0 + n).$$

In this connection c_0 has a natural interpretation as 'prior sample size'.

(g). A special case of the above, leading to simpler results, is that where β has a flat, non-informative prior, corresponding to very large prior variances, i.e. to $\Sigma_0 \to 0$. Show that with such a prior,

$$(\lambda, \beta) \,|\, \text{data} \sim \text{GN}_k(a + \tfrac{1}{2}n, b + \tfrac{1}{2}Q_0, \widehat{\beta}, n\Sigma_n).$$

The prior is improper (infinite integral), but the posterior is proper as long as $\widehat{\beta}$ exists, which requires $X^{\rm t}X$ to have full rank, which again means at least k linearly independent covariate vectors, and, in particular, $n \geq k$.

(h). (xx repair this. do bladder cancer rates, since linear is ok, whereas lung cancer is more quadratic, done in Story i.16. xx) Go again to the dataset 2.B, for illustration and for flexing your operational muscles. For y use the bladder cancer column of deaths per 100,000 inhabitants and for x use the number of cigarettes sold per capita. Your task is to carry out Bayesian analysis within the linear regression model $y_i = \beta_0 + \beta_1(x_i - \bar{x}) + \varepsilon_i$ for $i = 1, \ldots, 44$, with ε_i taken i.i.d. $\text{N}(0, \sigma^2)$. Specifically, we wish point estimates along with 95 percent credibility intervals for (i) each of the three parameters β_0, β_1, σ; (ii) the probability that $y \geq 25.0$, for a state with cigarette consumption $x = 35.0$; (iii) the bladder cancer death rates y_{45} and y_{46}, per 100,000 inhabitants, for states with cigarette consumption rates $x_{45} = 10.0$ (low) and $x_{46} = 50.0$ (high). You are to carry out such inference with two priors (xx nils repairs this xx): (1) First, the informative one which takes $1/\sigma^2$ a gamma with 0.10 and 0.90 quantiles for σ equal to 1.0 and 5.0, and β_0 and β_1 as independent normals $(5.0, 0.5\,\sigma^2)$ and $(0.0, (2.0\,\sigma)^2)$, given σ. (2) Then, the simpler and partly non-informative one that takes a flat prior for (β_0, β_1) and the less informative one for σ that uses 0.10 and 0.90 prior quantiles 0.5 and 10.0. Finally, compare your results from those arrived at using classical frequentist methods.

Ex. 6.20. *Mixture priors.* Suppose data Y come from a model $f(y, \theta)$, where different priors $\pi_1(\theta), \ldots, \pi_k(\theta)$ can be used, each leading to posterior distributions $\pi_1(\theta \,|\, y), \ldots, \pi_k(\theta \,|\, y)$.

(a). For each of these possible priors (and hence possible posteriors), show that there is a representation $f_j(y, \theta) = \pi_j(\theta) f(y, \theta) = \pi_j(\theta \,|\, y) \bar{f}_j(y)$, where $\bar{f}_j(y) = \int f(y, \theta) \pi_j(\theta)\, {\rm d}\theta$ is the marginal density of Y, associated with the $\pi_j(\theta)$ prior.

(b). Suppose now that a full mixture prior is assigned to θ, of the type $\pi(\theta) = p_1\pi_1(\theta) + \cdots + p_k\pi_k(\theta)$, with probabilities $p_1, \ldots, p_k$ summing to 1. Explain that this can be interpreted as θ being drawn from prior j with probability p_j. Show also that the marginal distribution of Y can be expressed as $\bar{f}(y) = \sum_{j=1}^k p_j \bar{f}_j(y)$.

(c). Then show that the posterior distribution for θ becomes

$$\pi(\theta \,|\, y) = p_1^* \pi_1(\theta \,|\, y) + \cdots + p_k^* \pi_k(\theta \,|\, y),$$

with revised prior probabilities $p_j^* = p_j \bar{f}_j(y) / \sum_{j'=1}^k p_{j'} \bar{f}_{j'}(y)$ for the different types of priors.

(d). Suppose $Y \sim \text{binom}(n, \theta)$, and that the prior used for θ is $0.15\,\text{Beta}(2, 10) + 0.70\,\text{Beta}(15, 15) + 0.15\,\text{Beta}(10, 2)$. Draw this prior in a plot. With $n = 100$, compute the posterior probabilities p_1^*, p_2^*, p_3^*, and draw the posterior distribution for θ, along with the prior, for each of the cases $y = 12$, $y = 48$, $y = 91$.

(e). (xx revisit the Gott würfelt nicht, Ex. 6.13. do a mixture prior, perhaps $0.50\,\mathrm{Dir}(1,1,1,1,1,1)+0.50\,\mathrm{Dir}(s,s,s,s,s,s)$, where s is quite big, reflecting the possibility that the die is perfectly fair with probabilities equal to or very close to $(1/6,1/6,1/6,1/6,1/6,1/6)$. xx)

(f). Generalise the above to the situation where $\pi(\theta)=\int \pi_\alpha(\theta)\,\mathrm{d}G(\alpha)$ is a mixture of $\pi_\alpha(\theta)$ priors, with $\mathrm{d}G(\alpha)$ a fully general probability measure over the space of hyperparameter α. The α coud be a parameter belonging to a finite set, matching the setup above, or a full continuous mixture. Show that the posterior can be represented as $\pi(\theta\,|\,y)=\int \pi_\alpha(\theta\,|\,y)\,\mathrm{d}G(\alpha\,|\,\mathrm{data})$, where $\mathrm{d}G(\alpha\,|\,\mathrm{data})$ is the posterior for the hyperparameter, and $\pi_\alpha(\theta\,|\,y)$ is the posterior for θ in the setup where α is fixed and known.

The Jeffreys prior

Ex. 6.21. *The Jeffreys prior: the basics.* Consider a smooth parametric model $f(y,\theta)$, with score function $u(y,\theta)=\partial\log f(y,\theta)/\partial\theta$ and finite Fisher information matrix $J(\theta)=\mathrm{Var}_\theta\,u(Y,\theta)$; see Ex. 5.14. The *Jeffreys prior* is defined as $\pi_{\mathrm{jeff}}(\theta)\propto|J(\theta)|^{1/2}$. If this has a finite integral, over the parameter domain, the proportionality constant is fixed to make the integral 1, but we explicitly allow also improper priors here, i.e. positive functions with infinite integral. It will be seen that in most cases, the passing from $\pi_{\mathrm{jeff}}(\theta)$ to $\pi_{\mathrm{jeff}}(\theta\,|\,\mathrm{data})$ via the Bayes theorem gives a proper posterior.

(a). Consider a smooth 1-1 reparametrisation, from $\theta=(\theta_1,\ldots,\theta_p)$ to $\phi=(\phi_1,\ldots,\phi_p)$, with $\theta=h(\phi)$ and $\phi=h^{-1}(\theta)$. Use Ex. 1.12 to show that if θ has the Jeffreys prior, then ϕ has density $\pi^*(\phi)=\pi_{\mathrm{jeff}}(h(\phi))|G(\phi)|$ with the gradient matrix $G(\phi)=\partial h(\phi)/\partial\phi$.

(b). We have seen in Ex. 5.17 that the Fisher information matrix on the ϕ scale is $J^*(\phi)=G(\phi)J(h(\phi))G(\phi)^{\mathrm{t}}$. Deduce that the Jeffreys prior constructed on the ϕ scale becomes $\pi^*_{\mathrm{jeff}}(\phi)=\pi_{\mathrm{jeff}}(h(\phi))|G(\phi)|$. But this is the same as $\pi^*(\phi)$ above. Explain that this demonstrates that the Jeffreys prior is invariant under smooth parameter transformations. Three Bayesian statisticians, with the same data and the same model, will reach the same conclusions, even if they use three different parametrisations, as long as they use the Jeffreys prior construction in their respective model formulations.

Ex. 6.22. *The Jeffreys prior in certain models.* (xx to come, examples, illustrations, some models. xx)

(a). For the binomial (n,p) model, show that the Jeffreys prior is the Beta$(\frac12,\frac12)$. Compare this with the geometric, with $f(y,p)=(1-p)^{y-1}p$, for which you should find $\pi_{\mathrm{jeff}}(p)\propto p^{-1}(1-p)^{-1/2}$, which can be thought of as a Beta$(\varepsilon,\frac12)$ with a very small ε. Show that this prior is improper, but that the posterior becomes proper as long as there are $n\geq 1$ observations.

(b). Then consider the trinomial (X,Y,Z), with probabilities proportional to $p^xq^y(1-p-q)^z$, with $x+y+z=n$. Find the 2×2 Fisher information matrix and its determinant, and show that the Jeffreys prior is proportional to $p^{-1/2}q^{-1/2}(1-p-q)^{-1/2}$, which is the Dirichlet $(\frac12,\frac12,\frac12)$ for (p,q,r). Generalise to the multinomial $(n,p_1,\ldots,p_k)$ case, with $p_1+\cdots+p_k=1$.

(c). For the normal ...

(d). For the Poisson with parameter θ, show that the Jeffreys prior is $1/\theta^{1/2}$.

(e). For the Gamma (a, b), show that the Jeffreys prior takes the form $\pi(a, b) = \pi_1(a)\pi_2(b)$, with $\pi_1(a) \propto \{a\psi'(a) - 1\}^{1/2}$ and $\pi_2(b) \propto 1/b$.

Ex. 6.23. *A simple model for deviations from uniformity.* This exercise illustrates how we can carry out Bayesian analysis for almost any given one-parameter model, via simple numerical techniques; bigger models need bigger tools, as we come back to (xx where xx). Consider the model $f(y, \theta) = 1 + \theta(y - \frac{1}{2})$ for $y \in [0, 1]$.

(a). Show that this indeed defines a bona fide model, for $\theta \in [-2, 2]$, and with c.d.f. $F(y, \theta) = y + \frac{1}{2}\theta(y^2 - y)$. Show that the Fisher information is

$$J(\theta) = \int_{-1/2}^{1/2} \frac{x^2}{1 + \theta x}\,\mathrm{d}x.$$

(xx perhaps more, a formula. xx) Compute and display the Jeffreys prior.

(b). Take $\theta_{\rm true} = 0.333$, simulate say $n = 100$ points from the model, and give a graph for the log-likelihood function. Compute the ML and an approximate 90 percent interval for θ via the methods of Chapter 5. Then, with a uniform prior on $[-2, 2]$, compute and display the posterior distribution for θ.

(c). Using a fine grid, e.g. with grid length 0.0001, sample say 10^5 points from the posterior distribution. From these, provide $0.05, 0.50, 0.95$ quantiles. (xx just a bit more; round off; point away. xx)

Marginal distributions and Bernshteĭn–von Mises theorems

Ex. 6.24. *The marginal distribution.* Suppose we have data $y_1, \ldots, y_n$ from a model $f(y, \theta)$, with a prior $\pi(\theta)$. Most of the time Bayesians care about the posterior distribution, but on occasion, also in connection with bigger setups, one needs the marginal distribution, which is $\bar f(y_1, \ldots, y_n) = \int L_n(\theta)\pi(\theta)\,\mathrm{d}\theta$, in terms of the likelihood function $L_n(\theta)$. In various setups there will be a clear formula for this $\bar f$, as seen below; see also Ex. 6.26 for a very useful approximation method for more complex cases.

(a). Let $y_1, \ldots, y_n$ be independent Bernoulli variables with $\Pr(y_i = 1 \,|\, \theta) = \theta$, and let $\theta \sim \mathrm{Beta}(a, b)$. Writing $z = \sum_{i=1}^n y_i$ for the number of 1s, show that

$$\bar f(y_1, \ldots, y_n) = \frac{\Gamma(a + b)}{\Gamma(a)\Gamma(b)} \frac{\Gamma(a + z)\Gamma(b + n - z)}{\Gamma(a + b + n)}.$$

(b). Let then $y_1, \ldots, y_n$ be independent $\mathrm{Pois}(\theta)$, with a $\mathrm{Gam}(a, b)$ prior, as with Ex. 6.1. Show that

$$\bar f(y_1, \ldots, y_n) = \frac{b^a}{\Gamma(a)} \frac{\Gamma(a + n\bar y)}{(b + n)^{a+n\bar y}} \frac{1}{y_1! \cdots y_n!}.$$

(c). Consider i.i.d. data $y_i \sim \mathrm{N}(\xi, \sigma^2)$, with known σ and a normal prior $\xi \sim \mathrm{N}(\xi_0, \sigma_0^2)$ for ξ. Find the marginal distribution (xx give a formula here xx).

(d). (xx do also $\mathrm{N}(x_i^{\rm t}\beta, \sigma^2)$ with $\beta \sim \mathrm{N}(\beta_0, \Sigma_0)$. find the marginal. check with other exercises. xx)

(e). (xx then also for the gamma-normal; give a formula for $\bar f(y_1, \ldots, y_n)$. xx)

(f). (xx something to point to empirical Bayes. calibrate carefully with loss-risk Ch. 8. can already point to Stein things. and to mixtures, where these $\bar{f}(y)$ turn up as ingredients. xx)

Ex. 6.25. *The gamma-normal induced marginal model.* (xx edit intro sentences. $Y_1, \ldots, Y_n$ are i.i.d. from the $\mathrm{N}(\xi, \sigma^2)$, given these two parameters. xx) For the direct Bayesian use one only needs the prior to posterior computation, in this case from the initial $\mathrm{GN}(a, b, \xi_0, v)$ to the updated GN given in Ex. 6.18, and one somehow bypasses the marginal density $\bar{f}(y_1, \ldots, y_n)$ of the data, the likelihood with the parameters (ξ, σ) integrated out according to the prior. On occasion this marginal distribution is important, however, and also finds use as a model in its own right, for positively dependent data.

(a). (xx first things with $\xi \sim \mathrm{N}(\xi_0, \tau_0^2)$, known σ. two ways of computing and seeing $\bar{f}(y_1, \ldots, y_n)$. do marginal moments and correlations. xx)

(b). xx

(c). Then the $\mathrm{GN}(a, b, \xi_0, v)$ gamma-normal prior for (λ, ξ), as with Ex. 6.18. Show first that the likelihood times the prior, $L_n(\lambda, \xi)p(\lambda, \xi)$, can be expressed as

$$\frac{\lambda^{n/2}}{(2\pi)^{n/2}} \exp[-\tfrac{1}{2}\lambda\{Q_0 + n(\xi - \bar{y})^2\}] \frac{b^a}{\Gamma(a)} \lambda^{a-1}(v\lambda)^{1/2} \exp[-\lambda\{b + \tfrac{1}{2}v(\xi - \xi_0)^2\}].$$

Then integrate out the ξ to get

$$\frac{1}{(2\pi)^{n/2}} \frac{b^a}{\Gamma(a)} \lambda^{a+n/2-1} \Big(\frac{v}{v+n}\Big)^{1/2} \exp\{-\lambda(b + \tfrac{1}{2}Q_0 + \tfrac{1}{2}d_n(\bar{y} - \xi_0)^2\},$$

with $d_n = (1/v + 1/n)^{-1}$ as per Ex. 6.18. Show then that this leads to the marginal density being

$$\bar{f}(y_1, \ldots, y_n) = (2\pi)^{-n/2} \Big(\frac{v}{v+n}\Big)^{1/2} \frac{b^a}{\Gamma(a)} \frac{\Gamma(a + n/2)}{\{b + \frac{1}{2}Q_0 + \frac{1}{2}d_n(\bar{y} - \xi_0)^2\}^{a+n/2}}.$$

Ex. 6.26. *Approximating the marginal distribution.* In the setup of Ex. 6.24, we go through a useful type of Laplace approximation for the marginal.

(a). Writing as usual $\ell_n(\theta)$ for the log-likelihood, with maximum value $\ell_{n,\max} = \ell_n(\widehat{\theta})$, in terms of the ML estimator, show that

$$\bar{f}(y_1, \ldots, y_n) = \exp(\ell_{n,\max}) \int \exp\{\ell_n(\theta) - \ell_n(\widehat{\theta})\}\pi(\theta)\,\mathrm{d}\theta.$$

With $J_n = -(1/n)\partial^2 \ell_n(\widehat{\theta})/\partial\theta\partial\theta^{\mathrm{t}}$ the normalised observed information matrix, of dimension say $p \times p$, show that the marginal can be approximated with

$$\begin{aligned}
\bar{f} &\doteq \exp(\ell_{n,\max}) \int \exp\{-\tfrac{1}{2}n(\theta - \widehat{\theta})^{\mathrm{t}} J_n (\theta - \widehat{\theta})\}\pi(\theta)\,\mathrm{d}\theta \\
&= \exp(\ell_{n,\max}) \int \exp(-\tfrac{1}{2}s^{\mathrm{t}} J_n s)\pi(\widehat{\theta} + s/\sqrt{n})\,\mathrm{d}s/n^{p/2} \\
&\doteq \exp(\ell_{n,\max})\pi(\widehat{\theta})(2\pi)^{p/2}|J_n|^{-1/2}/n^{p/2}.
\end{aligned}$$

(b). (xx a couple of things here. check how successful the approximation is in two setups. the formula

$$\log \bar{f} \doteq \ell_{n,\max} - \tfrac{1}{2}p\log n + \log \pi(\widehat{\theta}) - \tfrac{1}{2}\log|J_n| + \tfrac{1}{2}p\log(2\pi),$$

with its two leading terms, lead to the BIC in Ch. 11. xx)

Ex. 6.27. *Bernshteĭn–von Mises approximations.* Suppose observations $Y_1, \ldots, Y_n$ are i.i.d. from a density $f(y,\theta)$, with $\pi(\theta)$ a prior for the model parameter, of dimension say p. The posterior density can of course be quite complicated, perhaps necessitating numerical efforts, or simulation, for its evaluation. Remarkably, there are generic and simple normal approximations, however.

(a). Show that the posterior density $\pi_n(\theta) = \pi(\theta \,|\, \text{data})$ is proportional to $\pi(\theta)\exp\{\ell_n(\theta)\}$, with $\ell_n(\theta) = \sum_{i=1}^n \log f(y_i,\theta)$ the log-likelihood function.

(b). Let $\widehat{\theta}$ be the ML estimator, and $J_n = -(1/n)\partial^2\ell_n(\widehat{\theta})/\partial\theta\partial\theta^{\mathrm{t}}$ the normalised observed information, as per likelihood theory of Ch. 5. Show that the density of $Z_n = \sqrt{n}(\theta - \widehat{\theta})$ is $g_n(z) = \pi_n(\widehat{\theta} + z/\sqrt{n})(1/n^{p/2})$, and that it can be approximated as

$$g_n(z) \propto \pi(\widehat{\theta} + z/\sqrt{n})\exp\{\ell_n(\widehat{\theta} + z/\sqrt{n}) - \ell_n(\widehat{\theta})\} \doteq \pi(\widehat{\theta} + z/\sqrt{n})\exp(-\tfrac{1}{2}z^{\mathrm{t}}J_n z).$$

(c). Suppose then that the data really are i.i.d. from the model, with an underlying θ_{true}. In particular, then $\widehat{\theta} \to_{\text{pr}} \theta_{\text{true}}$ and $J_n \to_{\text{pr}} J = J(\theta_{\text{true}})$, by the general likelihood theory of Ch. 5. If $\pi(\theta)$ is continuous in a neighbourhood around θ_{true}, show that $g_n(z)$ tends to the density of a $\mathrm{N}_p(0, J^{-1})$, in probability. In concrete terms, show

$$D_n = \int |g_n(z) - \phi_p(z, 0, J^{-1})|\,\mathrm{d}z \to_{\text{pr}} 0.$$

This is one of several versions of *Bernshteĭn–von Mises theorems*. These are Bayesian mirror versions of the classical maximum likelihood asymptotics results in the frequentist camp:

sometimes called the Bayesian Central Limit Theorem

$$\begin{aligned}\sqrt{n}(\widehat{\theta} - \theta_{\text{true}}) &\to_d \mathrm{N}(0, J^{-1}),\\ \sqrt{n}(\theta - \widehat{\theta})\,|\,\text{data} &\to_d \mathrm{N}(0, J^{-1}),\ \text{in probability.}\end{aligned}$$

(d). Check two clear situations in detail, comparing the exact posterior density $\pi(\theta \,|\, \text{data})$ with the normal approximation: (i) where $Y \,|\, \theta \sim \text{binom}(n,\theta)$, and $\theta \sim \text{Beta}(a_0, b_0)$; (ii) where $Y_1, \ldots, Y_n \,|\, \theta$ are i.i.d. Pois(θ), and $\theta \sim \text{Gam}(a_0, b_0)$. Choose n and (a_0, b_0), and also the true θ_{true}, for your brief investigations.

(e). (xx just a bit more. lazy Bayesian. prior disappears. different Bayesians agree with each other, and also with the frequentist. xx)

Ex. 6.28. *Multiparameter inference, I.* Consider independent $Y_i \sim \mathrm{N}(\theta_i, 1)$ for $i = 1, \ldots, k$, where the task is to estimate the ensemble of means, with loss function $L = \sum_{i=1}^k (a_i - \theta_i)^2$.

(a). The standard estimator is Y_i, simply; show e.g. that this is the ML estimator. Show that the risk function $R(\theta, Y)$ is constant, equal to k. Show also that with this loss function, the Bayes solution, after having started with some prior for $(\theta_1, \ldots, \theta_k)$, are the posterior means, $\mathrm{E}\,(\theta_i \,|\, \text{data})$.

(b). Then let θ_i have the prior $\mathrm{N}(0,\tau^2)$, independently of the others, for $i=1,\ldots,k$. Show as with 6.16 that $\theta_i\,|\,\text{data}\sim\mathrm{N}(\rho y_i,\rho^2)$, with $\rho=\tau^2/(\tau^2+1)$. For the Bayes estimator $\widehat{\theta}_i(\tau)=\rho Y_i$, show that the risk function is $\sum_{i=1}^k\{\rho^2+(1-\rho)^2\theta_i^2\}$ and that the Bayes risk is $k\rho$.

(c). The ρY_i strategy requires of the Bayesian that τ is specified. When this is a too tall order, a natural method is to use the data to estimate ρ, leading to the *empirical Bayes strategy* $\theta_i^*=\widehat{\theta}_i(\widehat{\rho})=\widehat{\rho}Y_i$. With the $\theta_i\sim\mathrm{N}(0,\tau^2)$ prior, explain that Y_i, marginally, is a $\mathrm{N}(0,\tau^2+1)$, and that $V=\sum_{i=1}^k Y_i^2\sim(\tau^2+1)\chi^2_k$. Since V/k estimates τ^2+1, argue that $\widehat{\rho}=\max(1-k/V,0)$, and relatives, are sensible estimators of ρ. Technically, this is not a bona fide Bayes method, as $\widehat{\rho}Y_i$ is not the posterior mean stemming from a prior on $(\theta_1,\ldots,\theta_k)$, but its frequentist performance may be excellent. In particular, versions of this empirical Bayes scheme will succeed in having risk functions below the standard level risk level k, in all of the parameter space (provided $k\geq 3$, actually). This is the so-called Stein phenomenon; see connections from the present exercise to Ex. **??**. empirical Bayes

(d). A full Bayesian solution, without resorting to the half-Bayes method above, is to give a prior for τ, or equivalently the ρ, along with $\theta_i\,|\,\tau\sim\mathrm{N}(0,\tau^2)$. Show that the joint density of the $2k+1$ quantities in the game, i.e. $\rho,\theta_1,\ldots,\theta_k,y_1,\ldots,y_k$, may be written

$$\pi(\rho)\prod_{i=1}^k\Big\{\frac{1}{\tau}\phi\Big(\frac{\theta_i}{\tau}\Big)\phi(y_i-\theta_i)\Big\}=\pi(\rho)\prod_{i=1}^k\Big\{\frac{1}{\rho^{1/2}}\phi\Big(\frac{\theta_i-\rho y_i}{\rho^{1/2}}\Big)(1-\rho)^{1/2}\phi((1-\rho)^{1/2}y_i)\Big\}.$$

Integrate out the θ_i to show that $(\rho,y_1,\ldots,y_k)$ follow the joint density proportional to $\pi(\rho)(1-\rho)^{k/2}\exp\{-\frac12(1-\rho)v)\}$. Show that the Jeffreys prior is proportional to $1/(1-\rho)$, with which the posterior for ρ becomes $\pi(\rho\,|\,\text{data})\propto(1-\rho)^{k/2-1}\exp\{-\frac12(1-\rho)v\}$ on the unit interval.

(e). Explain that posterior simulation may be carried out via (i) drawing ρ from its posterior, e.g. via one-dimensional MCMC, and then (ii) drawing $\theta_1,\ldots,\theta_k$ independently from their $\mathrm{N}(\rho y_i,\rho)$. For the loss function L given above the Bayes solutions are the posterior means; show that these are $\widetilde{\theta}_i=\widetilde{\rho}y_i$, with $\widetilde{\rho}=\mathrm{E}\,(\rho\,|\,\text{data})$. This can easily be computed numerically.

(f). As long as one may carry out posterior simulations, one may also find Bayes solutions with respecte to more complicated loss functions than the quadratic one. For illustration, suppose the observations are

```
-1.277, -1.254, -0.791, -0.148,  0.178, 0.294,  0.847,  1.229,  1.720,  1.749,
```

and set up a scheme to simulate say 10^5 realisations of $(\theta_1,\ldots,\theta_k)$ given these datapoints, from the prior described above. Then find the Bayes estimators with the loss function $\sum_{i=1}^k|\widehat{\theta}_i-\theta_i|$. Check what happens if data point no. 10 is pushed some steps and then farther to the right.

Ex. 6.29. *Multiparameter inference, II.* (xx nils rant. need polish and finish. xx) the natural real thing: independent $Y_i\sim\mathrm{N}(\theta_i,\sigma^2)$ for $i=1,\ldots,k$, with unknown means and the same unknown variance. the object is estimating all the means, with loss $L=\sum_{i=1}^k(\widehat{\theta}_i-\theta_i)^2/\sigma^2$. direct estimator is Y_i, with risk k. with any prior construction, the Bayes solution is $\mathrm{E}\,(\theta_i\,|\,\text{data})$. natural prior: $\theta_i\sim\mathrm{N}(\theta_0,\tau^2)$. then

$$\theta_i\,|\,y_i\sim\mathrm{N}((1-\rho)\theta_0+\rho y_i,\rho\sigma^2),\quad\text{with }\rho=\tau^2/(\sigma^2+\tau^2).$$

here we are at a higher level of complexity compared to Ex. 6.29; there we knew θ_0 and σ, and could concentrate on a prior for ρ. for empirical Bayes, note that marginally, $Y_i \sim \mathrm{N}(\theta_0, \sigma^2+\tau^2)$. the sample mean is $\widehat{\theta}_0 = \bar{Y}$, and $V = \sum_{i=1}^k (Y_i-\bar{Y})^2 \sim (\sigma^2+\tau^2)\chi^2_{k-1}$, leading to an estimator for $\sigma^2+\tau^2$. the Bayes estimator is

$$\widehat{\theta}_i) = \mathrm{E}\,\{(1-\rho)\theta_0 + \rho y_i \,|\, \mathrm{data}\}.$$

we need posterior for (θ_0, ρ), but perhaps need a prior for the full triple (σ, θ_0, τ). we may write up

$$\begin{aligned} f &= \pi(\sigma,\theta_0,\tau) \prod_{i=1}^k \Big\{ \frac{1}{\tau}\phi\Big(\frac{\theta_i-\theta_0}{\tau}\Big)\frac{1}{\sigma}\phi\Big(\frac{y_i-\theta_i}{\sigma}\Big)\Big\} \\ &= \pi(\sigma,\theta_0,\tau) \prod_{i=1}^k \Big\{ \frac{1}{\rho^{1/2}\sigma}\phi\Big(\frac{\theta_i-(1-\rho)\theta_i-\rho y_i}{\rho^{1/2}\sigma}\Big)\frac{1}{(\sigma^2+\tau^2)^{1/2}}\phi\Big(\frac{y_i-\theta_0}{(\sigma^2+\tau^2)^{1/2}}\Big)\Big\}. \end{aligned}$$

integrating out the θ_i we have

$$\pi(\sigma,\theta_0,\tau) \prod_{i=1}^k \frac{1}{(\sigma^2+\tau^2)^{1/2}}\phi\Big(\frac{y_i-\theta_0}{(\sigma^2+\tau^2)^{1/2}}\Big).$$

Notes and pointers

(xx mention Varian (1975); Zellner (1986); Claeskens and Hjort (2008a) for the linex loss; but this should perhaps be in Ch. 8. mention Stigler (1983), who uses Bayes's Theorem to help him spot a candidate for an earlier discoverer of Bayes' Theorem. mention MCMC revolution since the 1990ies. point to Story ii.5. xx)

(xx on the Jeffreys prior. the argument used is slightly shaky since the transformation formula is used for possibly improper densities. point to Bernardo-Berger. can push in the nils-emil combo prior; informative on some components, non-informative on the rest. Cencov papers re Jeffreys being only solution to invariance demand. xx)

(xx hierarchical Bayes, with or without empirical Bayes. point to Efron etc. xx)

I.7

Confidence distributions, confidence curves, combining information sources

With μ a focus parameter, a function of the full parameter vector θ, the Bayesian setup gives a posterior distribution. This requires the conceptually and practically difficult task of defining a prior for the full θ, however. Confidence distributions (CDs) are a frequentist parallel, yielding post-data distributions for such focus parameters, without any prior. In this chapter we develop theory for CDs and confidence curves, and also find ways of combining CDs across different information sources. Constructing and computing CDs is not an easy or automatic task, but we develop and illustrate several recipes. For the exponential family class, we derive optimal CDs, with their own clear recipes.

Key words: boundary constraints, combining CDs, confidence curves, confidence distributions, exponential family, meta-analysis, t-bootstrapping

Confidence distributions and confidence curves are fruitful statistical inference summaries. Suppose in general terms that data y stem from a model $f(y, \theta)$, with model parameter $\theta = (\theta_1, \ldots, \theta_p)$, and that $\mu = \mu(\theta_1, \ldots, \theta_p)$ is a parameter of particular interest. A *confidence distribution* for μ, a CD, for short, is a function $C(\mu, y)$, such that (i) it is a c.d.f. in μ, for each dataset y, and (ii) the distribution of $U = C(\mu_0, Y)$ is uniform at the true value $\mu_0 = \mu(\theta_0)$. In other words,

$$\Pr_{\theta_0}\{a \le C(\mu_0, Y) \le b\} = b - a \quad \text{for each } a, b \in [0, 1].$$

Assuming this random c.d.f. has a unique inverse, then, we have

$$\Pr_\theta\{C^{-1}(0.05, Y) \le \mu \le C^{-1}(0.95, Y)\} = 0.90, \tag{7.1}$$

and of course similarly for other choices of quantiles. This is by definition making $[C^{-1}(0.05, y_{\rm obs}), C^{-1}(0.95, y_{\rm obs})]$ a 90 percent confidence interval for the focus parameter μ. The CD concept is hence related to and an extension of the confidence intervals, see Ch. 4. The *confidence curve* is a related summary graph, most often computed from the CD via $\mathrm{cc}(\mu, y) = |1 - 2\,C(\mu, y)|$. It has the practical property that $\{\mu\colon \mathrm{cc}(\mu, y) \le 0.90\}$ give the 90 percent interval directly; similarly, all intervals at any desired confidence level can be read off from the confidence curve. We sometimes write simply $C(\mu)$ and $\mathrm{cc}(\mu)$, omitting the data argument, when clear from the context what the data are, in $C(\mu) = C(\mu, \mathrm{data})$.

Construction a CD is not always an easy or automatic task, but we develop several practical recipes, some of which are based on approximate normality, or on more general methods of likelihood theory. Just as tests have detection power, also CDs have power, and theory is developed below to find optimal CDs in classes of situations. This is partly parallelling the optimal testing methodology of Ch. 4. All in all we develop and illustrate the following recipes: (i) Via the c.d.f. of an estimator; (ii) normal approximation; (iii) based on a pivot; (iv) deviance and Wilks theorem; (v) t-bootstrapping; (vi) the optimal CD via conditional distributions, if inside the exponential family.

The CDs are post-data graphical summaries of the level of uncertainty for any focus parameter, and can be seen as frequentist parallels to the Bayesian posterior distributions; here there is no prior, however. We illustrate this 'clear data-only based posteriors without priors' aspect of the CDs through theory and applications (xx perhaps point to a few Stories xx).

Combining different information sources is a broad statistical theme, going back to the first meta-analysis concepts and methods of Karl Pearson just after 1900 (Simpson and Pearson, 1904). The more familiar meta-analysis methods aim at combining independent estimators for the same quantity, or for providing a broader population assessment of similar but not identical parameters. CDs are useful for such endeavours, and we provide methods for combining sources more general than the traditional ones.

Recipes for constructing CDs

Ex. 7.1. *The probability transform.* Some of the following facts are related to various operations for confidence distributions and confidence curves

(a). Suppose X has a continuous and increasing cumulative distribution function F, i.e. $F(x) = \Pr(X \le x)$. Show that $U = F(X)$ is uniform on the unit interval. Any continuously distributed random variable can hence be transformed to uniformity, via this *probability transform*.

(b). Show that also $U_2 = 1 - F(X)$ and $U_3 = |1 - 2\,F(X)|$ have uniform distributions.

(c). Simulate a million copies of $x_i \sim \mathrm{N}(0,1)$, and check the histogram of $\Gamma_1(x_i^2)$, where Γ_ν is the cumulative distribution function of a χ^2_ν. Comment on what you find.

Ex. 7.2. *Recipe One: via the c.d.f. of an estimator.* Suppose θ is a one-dimensional parameter, for which we need a CD, after having observed data $y_{\rm obs}$. If there is an estimator $\widehat\theta$, with a distribution depending only on this θ, there is a clear recipe.

(a). Assume therefore that $\widehat\theta$ has a continuous distribution function $K_\theta(x) = \Pr_\theta(\widehat\theta \le x)$; its distribution is here required to depend only on θ, not on other aspects of the underlying model employed. Consider Recipe One, the construction

$$C(\theta, y_{\rm obs}) = \Pr_\theta(\widehat\theta \ge \widehat\theta_{\rm obs}) = 1 - K_\theta(\widehat\theta_{\rm obs}),$$

a curve that can be computed and plotted post-data, where $\widehat\theta_{\rm obs} = \widehat\theta(y_{\rm obs})$ is the observed estimate. Show that it has the property that the random $C(\theta, Y)$ is uniformly distributed, for each fixed θ.

(b). To illustrate, go through the details for the case of using $\widehat\theta = 1/\bar Y$, with i.i.d. observations $Y_1, \ldots, Y_n$ from the exponential $\theta \exp(-\theta y)$. Show first that $2\theta Y_i \sim \chi^2_2$, and derive $K_\theta(x) = 1 - \Gamma_{2n}(2n\theta/x)$, with Γ_{2n} the c.d.f. of the χ^2_{2n}. Simulate data and plot the CD $C(\theta, y_{\rm obs}) = \Gamma_{2n}(2n\theta/\widehat\theta_{\rm obs})$. From the CD, find a 95 percent interval for θ.

(c). Assume $X_1, \ldots, X_m$ are i.i.d. Expo(θ_1) and that $Y_1, \ldots, Y_n$ are i.i.d. Expo(θ_2). Find the distribution of the estimator $\widehat{\rho} = \widehat{\theta}_1/\widehat{\theta}_2$ for the ratio $\rho = \theta_1/\theta_2$, and derive the associated CD.

(d). Generate $n = 25$ datapoints from the double exponential density $f(y, \theta) = \frac{1}{2}\exp(-|y-\theta)$, using your favourite true θ_0. Compute and display the CD for θ based on the median M_n.

(e). For a simpler and more fundamental illustration, suppose $\widehat{\theta}$ has a normal distribution centred at θ, with a known variance, say $\widehat{\theta} \sim \mathrm{N}(\theta, \kappa^2)$. Show that Recipe One gives $C(\theta) = \Phi((\theta - \widehat{\theta})/\kappa)$. Check that the famous 95 percent interval $\widehat{\theta} \pm 1.96\,\kappa$ agrees with this.

Ex. 7.3. *Confidence distribution and confidence curve for the normal standard deviation.* The confidence distribution C and the confidence curve cc are close cousins, and they do not need to be both displayed for each new statistical application. Here is a simple illustration. You observe the $n = 6$ data points 4.09, 6.37, 6.87, 7.86, 8.28, 13.13 from a normal distribution and wish to assess the underlying spread parameter, the standard deviation σ.

(a). For the empirical variance, use $\widehat{\sigma}^2 \sim \sigma^2\chi^2_m/m$, with $m = n - 1$, to build the CD

$$C(\sigma, y_{\mathrm{obs}}) = \Pr_\sigma(\widehat{\sigma} \geq \widehat{\sigma}_{\mathrm{obs}}) = 1 - \Gamma_m(m\widehat{\sigma}^2_{\mathrm{obs}}/\sigma^2).$$

Here y_{obs} represents the observed data, and $\widehat{\sigma}_{\mathrm{obs}}$ the observed point estimate. Show that $C(\sigma, Y) \sim$ unif, where Y represents a random data set $Y_1, \ldots, Y_n$, from the σ in question. In particular, the distribution of $C(\sigma, Y)$ does not depend on σ. Make a graph, also of the associated confidence curve

the confidence curve

$$\mathrm{cc}(\sigma, y_{\mathrm{obs}}) = |1 - 2\,C(\sigma, y_{\mathrm{obs}})| = |1 - 2\,\Gamma_m(m\widehat{\sigma}^2_{\mathrm{obs}}/\sigma^2)|.$$

Compute the *median confidence estimate* $\widehat{\sigma}_{0.50} = C^{-1}(0.50, y_{\mathrm{obs}})$ and the natural 90 percent confidence interval $[C^{-1}(0.05, y_{\mathrm{obs}}), C^{-1}(0.95, y_{\mathrm{obs}})]$. Find and display also the *confidence density* $c(\sigma, y_{\mathrm{obs}})$, the derivative of the CD.

(b). Compute also the *confidence density* $c(\sigma, y_{\mathrm{obs}})$ associated with the CD. Compute furthermore its mode, say σ^*, and briefly assess its properties as an estimator of σ.

(c). A Bayesian approach to the same problem, i.e. finding a posterior distribution for σ, is to start with a prior $\pi(\sigma)$ and then compute $\pi(\sigma \,|\, y_{\mathrm{obs}}) \propto \pi(\sigma)g(\widehat{\sigma}, \sigma)$, where $g(\widehat{\sigma}, \sigma)$ is the likelihood, here the density function for $\widehat{\sigma}$ as a function of σ. When does such a Bayesian approach agree with the confidence density?

(d). Suppose there are two independent normal samples, with standard deviations σ_1 and σ_2. Construct a CD for $\rho = \sigma_1/\sigma_2$. Invent a second simple small dataset, to complement the first dataset given above, and then compute and display the confidence curve $\mathrm{cc}(\rho, \mathrm{data})$.

Ex. 7.4. *A CD for mean by standard deviation.* As in Ex. 3.39, consider i.i.d. $Y_1, \ldots, Y_n$ from a normal (ξ, σ^2), with inference needed for $\kappa = \xi/\sigma$. The natural estimator is $\widehat{\kappa} = \widehat{\xi}/\widehat{\sigma}$, with the usual sample mean and sample standard deviation. As we saw in the exercise pointed to, $X = \sqrt{n}\widehat{\kappa} \sim t_m(\sqrt{n}\kappa)$, with $m = n - 1$ degrees of freedom for the noncentral t.

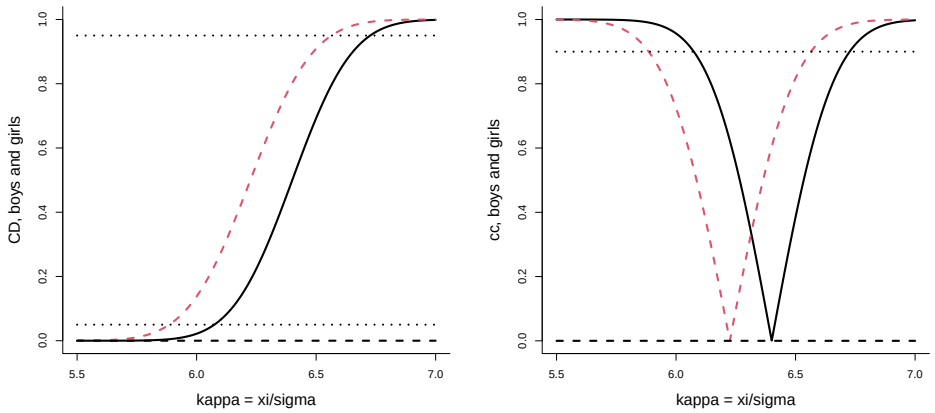


Figure 7.1: *For 548 boys and 480 girls born in Oslo, 2001–2008, CDs (left panel) and confidence curves (right panel), for the birthweight ratio $\kappa = \xi/\sigma$ (full curves for boys, slanted curves for girls). One may read off 0.05 and 0.95 confidence points from the CDs, or the 0.90 confidence interval directly from the ccs; 6.075 to 6.726 for boys, 5.886 to 6.564 for girls. Median confidence estimates are 6.400 and 6.225.*

(a). We start with a stylised version of the problem, with a single $T \sim t_m(a)$, with known m and unknown a. Show that $C(a, t_{\rm obs}) = \Pr_a(T \geq t_{\rm obs}) = 1 - G_m(t_{\rm obs}, a)$ is a CD for a, in terms of the c.d.f. for the $t_m(a)$ distribution. The associated confidence curve is then $\mathrm{cc}(a, t_{\rm obs}) = |1 - 2\,G_m(t_{\rm obs}, a)|$.

(b). Going back to the normal sample, show that the natural CD for $\kappa = \xi/\sigma$ becomes

$$C(\kappa) = \Pr_\kappa(\widehat{\kappa} \geq \widehat{\kappa}_{\rm obs}) = 1 - G_m(\sqrt{n}\widehat{\kappa}_{\rm obs}, \sqrt{n}\kappa).$$

(c). For an illustration, we use the data from Story i.5, with birthweights for 548 boys and 480 girls born in Oslo, 2001–2008. The required sufficient information $(\widehat{\xi}, \widehat{\sigma})$ is $(3.669, 0.573)$ for the boys and $(3.494, 0.561)$ for the girls. Compute and display the CDs and confidence curves, as with Figure 7.1. Find also 99 percent confidence intervals for the κ parameter in the two groups.

(d). Consider a linear regression setup, with $y_i \sim \mathrm{N}(x_i^{\rm t}\beta, \sigma^2)$ for $i = 1, \ldots, n$. For two covariate vectors x_a and x_b, examine the mean difference $\mathrm{E}\,(Y_b \,|\, x_b) - \mathrm{E}\,(Y_a \,|\, x_a) = (x_b - x_a)^{\rm t}\beta$, along with $\kappa = (x_b - x_a)^{\rm t}\beta/\sigma$, a natural quantity informing us about the essential statistical difference between the two. Construct a CD for this parameter. (xx now brief notes on solution. xx) writing $d = x_b - x_a$, note first that $Y_b - Y_a \sim \mathrm{N}(d^{\rm t}\beta, 2\sigma^2)$, so $p = \Pr(Y_b > Y_a) = \Phi(\kappa/\sqrt{2})$, with $\kappa = d^{\rm t}\beta/\sigma$. we wish a CD for this κ and hence for p. for this, $Y_b - Y_a - d^{\rm t}\widehat{\beta}$ is zero-mean normal with variance $(2 + \tau_n^2/n)\sigma^2$, where $\tau_n^2 = d^{\rm t}\Sigma_n^{-1}d$, with $\Sigma_n = n^{-1}\sum_{i=1}^n x_i x_i^{\rm t}$, as per usual linear regression theory of Ch3. via this we have a CD for $Y_b - Y_a$, for large n close to simply $Y_b - Y_a \sim \mathrm{N}(d^{\rm t}\beta, 2\sigma^2)$. also, with $\widehat{\kappa} = d^{\rm t}\widehat{\beta}/\widehat{\sigma}$, we have

$$\sqrt{n}\widehat{\kappa}/\tau_n = \frac{\sqrt{n}d^{\rm t}\widehat{\beta}/(\sigma\tau_n)}{\widehat{\sigma}/\sigma} = \frac{\mathrm{N}(\sqrt{n}\kappa/\tau_n, 1)}{(\chi^2_m/m)^{1/2}}$$

this then leads to the clear CD

$$C(\kappa) = \Pr_\kappa(\sqrt{n}\widehat{\kappa} \geq \sqrt{n}\kappa_{\rm obs}) = 1 - G_m(\sqrt{n}\widehat{\kappa}_{\rm obs}/\tau_n, \sqrt{n}\kappa/\tau_n).$$

this should be done for New Haven, regarding $p = \Pr(Y_b > Y_a)$.

Ex. 7.5. *A CD for the Mahalanobis distance between populations.* (xx nils rant; may need an easy application to illustrate. For two multinormal populations, say $\mathrm{N}_p(\xi_1, \Sigma)$ and $\mathrm{N}_p(\xi_2, \Sigma)$, the Mahalanobis distance between them is $\delta = \{(\xi_2 - \xi_1)^{\rm t}\Sigma^{-1}(\xi_2 - \xi_1)\}$. check Reiser and Faraggi, Biometrics 1997, and 'confidence intervals for the Mahalanobis distance', Reiser 2001. also point to other parameters, smooth functions of δ. so we solve many problems, precisely. xx)

(a). The natural estimator is $\widehat{\delta}^2 = (\bar{Y}_2 - \bar{Y}_1)^{\rm t}\widehat{\Sigma}^{-1}(\bar{Y}_2 - \bar{Y}_1)$, with $\widehat{\Sigma} = c_1\widehat{\Sigma}_1 + c_2\widehat{\Sigma}_2$, with $c_1 = (m-1)/(m+n-2)$ and $c_2 = (n-1)/(m+n-2)$. Then (xx via other things xx) show that

$$D^2 = \frac{mn}{m+n}\frac{m+n+p-1}{m+n-2}\widehat{\delta}^2 \sim F_{p,m+n-p-1}(\lambda),$$

with excentre parameter $\lambda = \{mn/(m+n)\}\delta^2$.

(b). (xx with the above result is should be easy to construct a clear CF for δ. xx)

(c).

Ex. 7.6. *Computing a CD with simulation and isotonic repair.* (xx to be polished. we use this is Story iii.10 and perhaps in other places, where simulations are expensive. xx) Suppose one observes $y_1, \ldots, y_n$ from the one-parameter Weibull distribution with c.d.f. $F(y, b) = 1 - \exp(-y^b)$, with sample size $n = 25$, and computes the data mean $\bar{y}_{\rm obs} = 1.313$.

(a). Though we do not actually need this in the CD computations here, find an estimate of b based on $\mathrm{E}\, Y_i = \Gamma(1 + 1/b)$; see Ex. 1.59. Show that $C(b) = \Pr_b(\bar{Y} \leq \bar{y}_{\rm obs})$ is a CD for b.

(b). The practical obstacle here is that $\bar{Y}$ does not have a simple distribution. But we're saved by simulation. Show that the simulation recipe $Y_i^* = V_i^{1/b}$ produces outcomes from the weibull $F(y, b)$, where the V_i are unit exponential. For a grid of b values, e.g. from 0.20 to 1.20, compute the simulation based $C^*(b)$, the proportion of B cases where the simulated $\bar{Y}^*$ is below $\bar{y}_{\rm obs}$. Compute also the confidence curve $\mathrm{cc}^*(b) = |1 - 2\,C^*(b)|$. For this simple example it is easy to accomplish this with a high B, say 10^5, to make $C^*(b)$ and $\mathrm{cc}^*(b)$ smooth and very close to the real $C(b)$ and $\mathrm{cc}(b)$; for this illustration, however, make the simulation size as relatively small as $B = 100$, and plot the curves, as in Figure 7.2.

(c). We learn that with a low or moderate simulation size B, the $C^*(b)$ and $\mathrm{cc}^*(b)$ will be wiggly. We can do better, using the prior knowledge that $C(b)$ is increasing. There are several repair mechanisms, which from the potentially wiggly $C^*(b)$ create a monotonically increasing curve. A simple scheme is so-called *isotonic regression*, the details of which we do need to get into here. Supposing you have first created `bval` and `Cval` in your `R` session, you may use `Cvaliso=isoreg(bval,Cval)$yf`, which repairs your $C^*(b)$ and $\mathrm{cc}^*(b)$ to ensure monotonicity. Produce versions of Figure 7.2, left and right panels.

isotonic regression

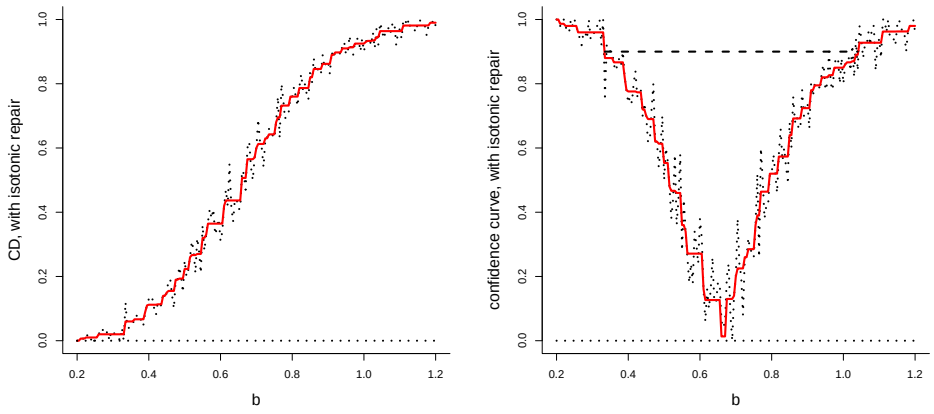


Figure 7.2: *Simulation based confidence distribution* $C^*(b)$ *and confidence curve* $\text{cc}^*(b)$ *for the Weibull parameter* b*, based on the observed sample mean* $\bar{y}_{\text{obs}} = 1.313$ *for* $n = 25$ *data points, along with isotonic repairs. The simulation size here is the low* $B = 100$*.*

(d). (xx round off. explain salient points about generalisability. we need reduction to one-parameter situation. xx)

Ex. 7.7. *An extension of Recipe One.* In Ex. 7.2 we saw that the simple construction $C(\theta, y) = \Pr_\theta(\widehat{\theta} \geq \widehat{\theta}_{\text{obs}})$ gives a CD, in the case of one-dimensional setups with a well-defined estimator $\widehat{\theta}$.

(a). When working with estimators, finetuning efforts are often exuded to trim away biases, getting the scaling right, etc. In a sense this is not needed here, when constructing the CD. Show that if $\widehat{\alpha} = g(\widehat{\theta})$, with any smooth increasing g, the recipe $C^*(\theta) = \Pr_\theta(\widehat{\alpha} \geq \widehat{\alpha}_{\text{obs}})$ gives precisely the same CD as without the g transformation.

(b). So this CD recipe relies merely on having an informative statistic, say Z, with a distribution stochastically increasing in θ; it does not really have to be an estimator for that parameter. Show that $C(\theta, y) = \Pr_\theta(Z \geq z_{\text{obs}})$ is a bona fide CD.

(c). Show also that the construction works, if there are other parameters at play too, as long as the distribution of the chosen Z only depends on θ. Go through the details for the case of the Y_i being $\text{N}(\mu, \sigma^2)$, with $Z = \sum_{i=1}^n (Y_i - \bar{Y})^2$, and also for $Z' = \sum_{i=1}^n |Y_i - M_n|$, where M_n is the empirical median. Compute, display, compare both CDs, based on Z and on Z', for the simple dataset of Ex. 7.3 (with $n = 6$). For the Z case, there is a formula, but for the Z' case you would need simulation, for a grid of σ values; see Ex. 7.6.

(d). (xx one more example, where there is a Z carrying information, but not qua estimator. xx)

Ex. 7.8. *Recipe Two: the normal approximation CD.* Applying Recipe One of Ex. 7.2 to the case of the estimator having a normal distribution leads as we saw there to a clear CD, provided the variance is known. But this is at least approximately so, for large classes of situations, as we've seen in Chs. 2 and 5.

(a). Suppose in general terms that $\widehat{\theta}$ estimates θ, and that its distribution is approximately a $\mathrm{N}(\theta, \kappa^2)$. Explain that $C(\theta) = \Phi((\theta - \widehat{\theta})/\kappa)$ then is an approximate CD for θ. More formally, if $(\widehat{\theta} - \theta_0)/\widehat{\kappa} \to_d \mathrm{N}(0,1)$, at the true parameter θ_0, show that $C(\theta, Y) = \Phi((\theta - \widehat{\theta})/\widehat{\kappa})$ has the property that it converges in distribution to the uniform, at θ_0. In typical applications of these arguments, there is a $\sqrt{n}$ scaling in terms of an underlying sample size, with $\sqrt{n}(\widehat{\theta} - \theta_0) \to_d \mathrm{N}(0, \tau^2)$, say, and $\widehat{\kappa} = \widehat{\tau}/\sqrt{n}$, with $\widehat{\tau} \to_{\mathrm{pr}} \tau$. So this is Recipe Two, the normal approximation CD, most typically of this type $\Phi(\sqrt{n}(\theta - \widehat{\theta})/\widehat{\tau})$.

(b). Simulate a moderate or small dataset from a normal distribution. Compute and display two (approximate) CDs for the mean parameter ξ, (i) using the data mean, (ii) using the data median.

(c). We have seen in Chs. 2 and 5 that approximate normality is highly common, for large classes of estimators, typically along with consistent estimators for the variances. In particular, the delta method implies approximate normality of smooth functions of background estimators (see Ex. 2.71, **??**), making in its turn approximate normality CDs easily available. For a simple illustration, suppose you throw your nearest die, which has probability p of giving a '6', until you get your first '6'. You carry out this geometric experiment $n = 10$ times, giving you the counts $Y_1, \ldots, Y_n$ equal to 1, 2, 17, 18, 20, 4, 3, 1, 15, 3. Use the normal approximation for $\bar{Y}$ to give an approximate CD for p. You may also compare this to what one achieves working with the exact distribution of $\bar{Y}$.

(d). (xx point to logistic and poisson regression, with delta method. estimate β and also $p = H(x_0^{\mathrm{t}}\beta)$. xx)

Ex. 7.9. *Recipe Three: from a pivot to a CD.* (xx check that we're not repetitive regarding pivot. xx) Suppose in general terms that μ is some parameter of interest, in a model for observations Y, and that a function $A = \mathrm{piv}(\mu, Y)$ of the parameter and the data has the property that its distribution does not depend on the model parameters (in particular, therefore, not on μ, which might itself be a function of other model parameters). We call A a pivot, in more pedantic detail a pivot for the parameter μ.

(a). With $Y_1, \ldots, Y_n$ independent from the normal (μ, σ^2), let $R_n = \sum_{i=1}^n |Y_i - \bar{Y}|$ with the sample mean $\bar{Y}$. Show that $(\bar{Y} - \mu)/R_n$ is a pivot. Invent yet another pivot involving μ, with a different denominator.

(b). With two normal samples, say $X_1, \ldots, X_m$ from $\mathrm{N}(\mu_1, \sigma_1^2)$ and $Y_1, \ldots, Y_n$ from $\mathrm{N}(\mu_2, \sigma_2^2)$, suppose $\rho = \sigma_1/\sigma_2$ is in focus. Show that $(V_1/V_2)/\rho$ is a pivot for ρ, where V_1 and V_2 are the interquartile ranges for the two datasets.

(c). Consider $Y_1, \ldots, Y_n$ from the Cauchy model with density $(1/\pi)/\{1 + (y-\theta)^2\}$. Show that $R_n - \theta$ is a pivot, where $R_n = \frac{1}{2}(Q_{n,0.10} + Q_{n,0.90})$ is the average of the 0.10 and 0.90 quantiles.

(d). Back to the generalities, consider a pivot $A = \mathrm{piv}(\mu, Y)$ for μ in some model, increasing in μ. Assume the situation is continuous, not discrete, so that the pivot's distribution function K is continuous. Show that $C(\mu, y_{\mathrm{obs}}) = 1 - K(\mathrm{piv}(\mu, y_{\mathrm{obs}}))$ is a proper CD for μ.

(e). In clean cases we may derive the precise distribution for the pivot in question, but the CD recipe given above may be used also in more complicated setups, as long as

$A = \text{piv}(\mu, Y)$ may be simulated. Make an illustration of this, with the ratio of standard deviations above. Suppose two normal datasets, both of size $n = 100$, lead to interquartile ranges $V_{1,\text{obs}} = 4.44$ and $V_{2,\text{obs}} = 3.33$. Construct and display $C(\rho)$ and $\text{cc}(\rho)$.

(f). (xx make the point that various constructions, involving large-sample approximations to the normal and to chisquares, lead to *approximate pivots*, and then again to approximate CDs and ccs. in particular, Method One, with $\Phi((\mu - \widehat{\mu})/\widehat{\kappa})$ and Method Two, with $\Gamma_1(D(\mu))$, can be seen via approximate pivots. also Method Three, construction of a t type ratio and then bootstrapping. xx)

Ex. 7.10. *CDs from the t pivot.* In Ex. 7.2 we saw that the simple construction $C(\theta, y) = \Pr_\theta(\widehat{\theta} \geq \widehat{\theta}_{\text{obs}})$ gives a CD, in the case of one-dimensional setups with a well-defined estimator $\widehat{\theta}$.

(a). For a normal sample from $\text{N}(\mu, \sigma^2)$, we see that several $\Pr_{\mu,\sigma}(Z \geq z_{\text{obs}})$ schemes work, in that the Z in question has a distribution depending on σ, but not μ. Attempt to work with $C^*(\mu, y) = \Pr_{\mu,\sigma}(\bar{Y} \geq \bar{y}_{\text{obs}})$ – and explain that it will not really work (unless σ is known).

(b). But of course there *are* natural CD constructions for μ here. What is needed is a *pivot*, say $A = \text{piv}(\mu, y)$, a function binding the focus parameter and data together in a way which makes its distribution not depend on the parameters. Study indeed

$$t_n = t_n(\mu, Y) = (\bar{Y} - \mu)/(\widehat{\sigma}/\sqrt{n}),$$

with $\widehat{\sigma}^2 = (n-1)^{-1} \sum_{i=1}^n (Y_i - \bar{Y})^2$ the classical empirical variance. Pretend that you in all your cleverness have not seen this t_n before, and are unaware of its relation to a t distribution – but show that the distribution of t_n, call it K_n, does not depend on (μ, σ).

(c). Then show that $C(\mu, y_{\text{obs}}) = K_n(t_n(\mu, y_{\text{obs}}))$ is a CD for μ. Even if you do not see the connection to the classic t of Student (1908), see Ex. 1.50, you may still carry through this, by simulating $B = 10^5$ realisations of t_n, and use

$$C(\mu, y_{\text{obs}}) = K_n^*(t_n(\mu, y_{\text{obs}})) = \frac{1}{B} \sum_{j=1}^B I\{t_{n,j} \leq t_n(\mu, y_{\text{obs}})\}.$$

Show however that by all means K_n is a t_m, with $m = n - 1$, so the canonical CD for μ is and remains $C(\mu, y_{\text{obs}}) = G_m(\sqrt{n}(\mu - \bar{y}_{\text{obs}})/\widehat{\sigma}_{\text{obs}})$, with G_m the c.d.f. for the t_m.

Ex. 7.11. *Recipe Four: confidence curves via Wilks theorems.* Consider data from a parametric model, leading to the log-likelihood function $\ell_n(\theta)$, and that there is a focus parameter $\mu = g(\theta)$. We have seen likelihood profiling and Wilks theorems in Ch. 5, and know that the deviance $D_n(\mu) = 2\{\ell_{\max} - \ell_{\text{prof}}(\mu)\}$ has the property that $D_n(\mu_0) \to_d \chi^2_1$ at the true value $\mu_0 = g(\theta_0)$; see Ex. 5.29.

(a). Recipe Four, utilising the Wilks theorems, is to form $\text{cc}(\mu, y) = \Gamma_1(D_n(\mu))$, with Γ_1 the c.d.f. for the χ^2_1. Show that $\Pr_{\theta_0}(\text{cc}(\mu_0, Y) \leq \alpha) \to \alpha$ for each α, and explain that this makes $\text{cc}(\mu, y)$ and approximate confidence curve.

(b). For an illustration, consider the model $F(y, \theta) = y^\theta$ for observations on $[0, 1]$, where θ is an unknown positive parameter. Write down the log-likelihood function and find a formula for the maximum likelihood (ML) estimator $\widehat{\theta}$. Use also theory of Ch. 5 to write down a normal approximation to the distribution of $\widehat{\theta}$.

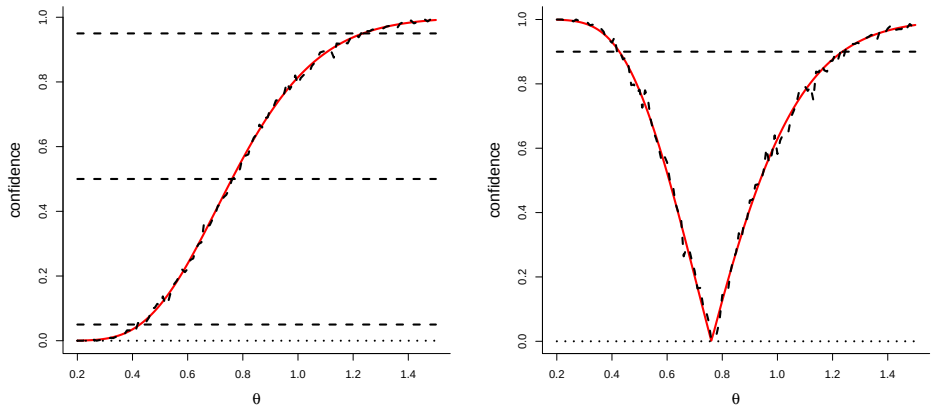


Figure 7.3: *For the simple data example of Ex. 7.11: Left panel: confidence distribution $C(\theta)$, via simulations (black and wiggly curve) and via exact calculations (red and smooth curve); right panel: the two versions of the associated confidence curve* $\mathrm{cc}(\theta)$. *From these we read off the median confidence estimate* $\widehat{\theta}_{0.50} = 0.76$, *and the 90 percent confidence interval* $[0.43, 1.24]$.

(c). Consider the data set

```
0.013 0.054 0.234 0.286 0.332 0.507 0.703 0.763 0.772 0.920
```

Estimate θ and compute the CD $C(\theta) = \mathrm{Pr}_\theta(\widehat{\theta} \geq \widehat{\theta}_{\mathrm{obs}})$, along with the confidence curve $\mathrm{cc}(\theta) = |1-2\,C(\theta)|$, (i) using simulations, (ii) using exact probability calculus. Reproduce a version of Figure 7.3.

(d). Supplement these two curves with approximations based (i) on the normal approximation for $\widehat{\theta}$ and (ii) on the chi-squared approximation for the deviance.

(e). (xx somewhere, if not here, then separately: from CD to cc, and from cc to CD, with $C(\mu) = \frac{1}{2} - \frac{1}{2}\mathrm{cc}(\mu)$ for $\mu \leq \widehat{\mu}_{0.50}$ and $\frac{1}{2} + \frac{1}{2}\mathrm{cc}(\mu)$ for $\mu \geq \widehat{\mu}_{0.50}$. particularly useful with these deviance based ccs. xx)

Ex. 7.12. *Median age for men and women in Roman Era Egypt.* In Story ii.11 we work with a rare dataset of lifetimes for 82 men and 59 women in Roman Era Egypt, a century B.C. For the present illustration of constructing confidence curves via log-likelihood profiling, take these lifetimes T_i to have arisen from Weibull distributions, with c.d.f.s $F_m(t) = 1 - \exp\{-(t/a_m)^{b_m}\}$ for men and $F_w(t) = 1 - \exp\{-(t/a_w)^{b_w}\}$ for women.

(a). For a Weibull, with parameters (a, b), show that the q level quantile becomes $\mu(q) = F^{-1}(q) = ac(q)^{1/b}$, with $c(q) = -\log(1-q)$. In particular, the median is $a(\log 2)^{1/b}$.

(b). For the men and women separately, write down the log-likelihood functions $\ell_m(a_m, b_m)$ and $\ell_w(a_w, b_w)$, and then carry out profiling to compute the deviance functions, for the medians μ_m and μ_w. Use Recipe Four, from Ex. 7.11, to construct the confidence curves, and make a version of Figure 7.4, left panel; also, read off 90 percent confidence intervals.

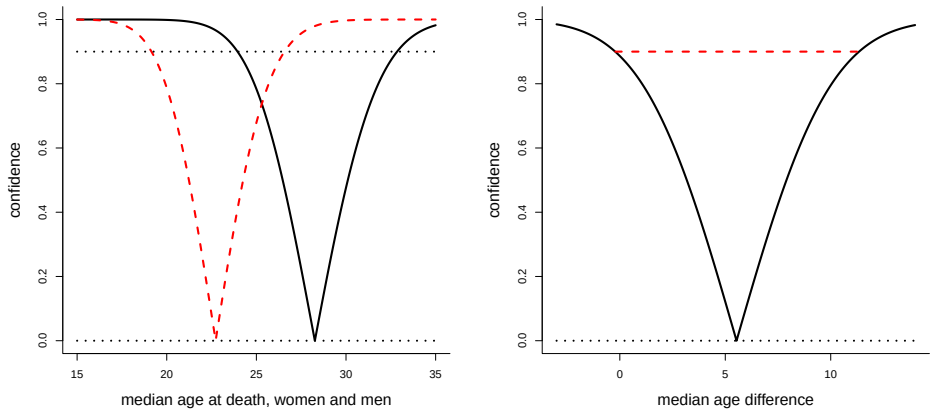


Figure 7.4: *Lifetimes in Roman Era Egypt, a century B.C.: Left panel: confidence curves for the median age μ_w and μ_m, for women and men, via log-likelihood-profiling from fitting separate Weibull distributions, with 90 percent confidence intervals* $[19.95, 26.56]$ *and* $[23.95, 32.84]$. *Right panel: confidence curve for the median age difference* $d = \mu_m - \mu_w$, *via log-likelihood profiling, and 90 percent confidence interval* $[-0.21, 11.34]$.

(c)**.** It appears indeed that men had longer lives than women, in Egypt 2100 years ago; see Story ii.11 for more discussion and details. For this exercise, carry out log-likelihood profiling, for the median difference $d = \mu_m - \mu_w$, from the four-parameter model with the two Weibulls. Compute the deviance function and then the confidence curve $\mathrm{cc}(\cdot)$, leading to Figure 7.4, right panel. Read off a 90 percent interval for d.

(d)**.** Constructing the confidence curves above involved the somewhat laborious computation of log-likelihood profiles. Compare these curves with the easier ones, computationally speaking, based on normal approximations.

Ex. 7.13. *The Wilks confidence recipe for the normal model.* Consider the familiar setup with i.i.d. data $Y_1, \ldots, Y_n$ from the $\mathrm{N}(\xi, \sigma^2)$ model. We have handled the log-likelihood profiles and accompanying Wilks theorems approximations in Ex. 5.30, and can apply results for constructing confidence curves. In the following, let as usual $\widehat{\xi}$ and $\widehat{\sigma}^2$ be the sample mean and sample variance.

(a)**.** Starting with the mean parameter ξ, explain first that $C(\xi, \mathrm{data}) = G_m(\sqrt{n}(\xi - \bar{Y})/\widehat{\sigma})$ is the natural exact CD, with G_m the c.d.f. for the t_m, with $m = n - 1$ degrees of freedom. Then use the exercise pointed to to show that the Wilks recipe leads to

$$\mathrm{cc}_w(\xi) = \Gamma_1\big(n \log(1 + n^{-1} K_n(\xi))\big),$$

with $K_n(\xi) = n(\xi - \widehat{\xi})^2/\widehat{\sigma}^2$.

(b)**.** Then considering σ, explain that the natural exact CD is $C(\sigma, \mathrm{data}) = 1 - \Gamma_m(m\widehat{\sigma}^2/\sigma^2)$, with Γ_m the c.d.f. for χ^2_m. Show that the Wilks recipe leads to

$$\mathrm{cc}_w(\sigma) = \Gamma_1(nh(\widehat{\sigma}^2/\sigma^2)),$$

with $h(x) = x - \log x - 1$.

(c). (xx then compare the exact with the Wilks. approximation is good even for moderately low n. ask perhaps for doing this with the $n = 64$ word-length averages of Story ii.2. xx)

Ex. 7.14. *Recipe Five: CDs via approximate pivots and t-bootstrapping.* Consider a parametric model $f(y, \theta)$ for data, with a model parameter of length say p. Suppose there is a focus parameter $\mu = g(\theta)$, estimated as $\widehat{\mu} = g(\widehat{\theta})$, for which we need a CD.

(a). Suppose first that $\widehat{\mu} - \mu$ has some distribution, say G_0, not depending on θ. Under this simple pivotal assumption for 'estimator minus estimand', show that

$$C(\mu) = \Pr_\theta(\widehat{\mu} \geq \widehat{\mu}_{\rm obs}) = 1 - G_0(\widehat{\mu}_{\rm obs} - \mu).$$

As indicated in e.g. Ex. 7.6, this can be computed even without knowing the form of G_0, through simulation of many realisations of $\widehat{\mu}^* - \widehat{\mu}$, where the $\widehat{\mu}^*$ is computed from a dataset drawn from the estimated distribution at $\widehat{\theta}$. Explain that this recipe works, even if G_0 is nonsymmetric and not centred well at zero. Simulating the distribution of $\widehat{\mu} - \mu$ at different points in the θ parameter space may also be helpful for checking the assumption needed for the $C(\mu)$ constructed here to be a clear CD.

(b). For a special and simpler case, if $\widehat{\mu} \sim {\rm N}(\mu, \kappa^2)$, to a good approximation, with known or well estimated κ, show that the general recipe above leads to

$$C(\mu) = \Pr_\theta(\mu + \kappa N \geq \widehat{\mu}_{\rm obs}) = \Pr(N \geq (\widehat{\mu} - \mu)/\kappa) = \mu((\mu - \widehat{\mu})/\kappa),$$

and argue that this is really a CD. Note that this requires $Z = (\widehat{\mu} - \mu)/\kappa$ having distribution equal to or close to a standard normal, regardless of where θ is in its parameter space.

(c). Often the standard deviation in the normal approximation is not that sharply estimated from the available data. Consider a Student type ratio $t = (\widehat{\mu} - \mu)/\widehat{\kappa}$, with some appropriate scale estimator $\widehat{\kappa}$. Assume first that t really is a pivot, i.e. that its distribution G is independent or nearly independent of where θ is in its parameter space. Show that

$$C(\mu) = \Pr_\theta((\widehat{\mu} - \mu)/\widehat{\kappa} \geq (\widehat{\mu}_{\rm obs} - \mu)/\widehat{\kappa}_{\rm obs}) = 1 - G((\widehat{\mu}_{\rm obs} - \mu)/\widehat{\kappa}_{\rm obs})$$

is a CD. If G is not known, or too difficult to derive, use simulations, of $t^* = (\widehat{\mu}^* - \widehat{\mu}_{\rm obs})/\widehat{\kappa}^*$, via datasets simulated at position $\widehat{\theta}_{\rm obs}$. We call this a CD computed from t-bootstrapping.

(d). The previous recipe works well if $t = (\widehat{\mu} - \mu)/\widehat{\kappa}$ is close to pivotal, i.e. its distribution is nearly constant over the parameter region. In other cases we may take the t-bootstrapping argument one step further. Write for emphasis $\theta = (\mu, \gamma)$, perhaps in a reparametrisation, where μ is in focus and γ is of length $p-1$. The t has some distribution, depending on θ, and we write $\Pr_\theta(t \leq u) = G(u, \mu, \gamma)$. Show that

$$H(\mu, \gamma) = \Pr_{\mu,\gamma}((\widehat{\mu} - \mu)/\widehat{\kappa} \geq (\widehat{\mu}_{\rm obs} - \mu)/\widehat{\kappa}_{\rm obs}) = 1 - G((\widehat{\mu}_{\rm obs} - \mu)/\widehat{\kappa}_{\rm obs}, \mu, \gamma).$$

This is not necessarily a CD, in the strict sense, as this probability may depend not only on μ but also on aspects of γ. Often the distribution of t is approximately the same, though, in a neighbourhood around the true value. Argue that this leads to

$$C^*(\mu) = 1 - \widehat{G}((\widehat{\mu}_{\rm obs} - \mu)/\widehat{\kappa}_{\rm obs}, \mu) \quad \text{where } \widehat{G}(u, \mu) = G(u, \mu, \widehat{\gamma}).$$

Such an estimated distribution can be computed via bootstrapping, i.e. simulated datasets at position $\widehat{\theta}$ in the parameter space. With B such simulated datasets, leading to simulated values $\widehat{\theta}^*, \widehat{\mu}^*, \widehat{\kappa}^*$, and hence $t^* = (\widehat{\mu}^* - \widehat{\mu}_{\rm obs})/\widehat{\kappa}^*$.

Ex. 7.15. *Recipe Six: CDs in exponential families.* We have worked with the general exponential family in previous chapters, see Ex. 1.55. In particular we learned in Ex. 4.32 that there are uniformly optimal tests, for individual parameters in such models. The same holds in the present framework of CDs. Suppose data stem from a model of the form $f(y,a,b) = \exp\{aU(y) + b^{\rm t}V(y)\}h(y)$, with U one-dimensional and V of dimension say p. The optimal recipe for a is

$$C^*(a) = \Pr_a\{U(Y) \geq u_{\rm obs} \,|\, V(Y) = v_{\rm obs}\}.$$

This construction is actually optimal, in a power risk function sense we come back to in Ex. 7.32, but we can already start working with this definition and see how it applies in various situations.

(a). Verify from arguments in Ex. 4.32 that $C^*(a)$ indeed depends only on a, not on b.

(b). For an illustration, consider the pair of exponentials of Ex. 4.27. To avoid confusion with the parametrisation, use now $X \sim {\rm Expo}(\theta)$, $Y \sim {\rm Expo}(\theta+\delta)$, with sum $Z = X+Y$. Show that the joint density is indeed of exponential form, and that the recipe leads to

$$C^*(\delta) = \Pr_\delta(Y \leq y_{\rm obs} \,|\, Z = z_{\rm obs}) = \frac{1-\exp(-\delta y_{\rm obs})}{1-\exp(-\delta z_{\rm obs})}.$$

Find the positive confidence pointmass at $\delta = 0$.

(c). Suppose there are m independent pairs of such exponentials, with $X_i \sim {\rm Expo}(\theta_i)$, $Y_i \sim {\rm Expo}(\theta_i+\delta)$, and sums $Z_i = X_i + Y_i$. We need a CD for the difference parameter δ. Show that the joint density of the $2m$ variables is on the exponential form, and that the resulting CD must be of the form

$$C^*(\delta) = \Pr_\delta(U \leq u_{\rm obs} \,|\, Z_1 = z_{1,{\rm obs}}, \ldots, z_{m,{\rm obs}}),$$

with $U = \sum_{i=1}^m Y_i$. There is no clear formula for this conditional distribution, but show that $Y_i \,|\, z_i$ has density $\delta\exp(-\delta y_i)/\{1-\exp(-\delta z_i)\}$ for $y_i \in [0,z_i]$. To show how the CD can be computed, via simulations, suppose as in Ex. 4.27 that the data are the three pairs $(0.927, 0.819), (1.479, 0.408), (3.780, 1.311)$. In that exercise we worked with the optimal test for $\delta = 0$ vs. $\delta > 0$, and needed only the null distribution of U given the three sums z_1, z_2, z_3, i.e. where $\delta = 0$. Now we need to tabulate this conditional distribution also for each $\delta > 0$, however.

(d).

Ex. 7.16. *Optimal CD for a bivariate model.* (xx spell out. we do get the natural answer $C(a) = \Pr_a(\widehat{a} \geq \widehat{a}_{\rm obs} \,|\, \widehat{b})$. then multivariate. xx)

Ex. 7.17. *Bayesian posteriors as approximate CDs.* (xx to come here: Consider a setup with data y from a model with parameter $\theta = (\theta_1, \ldots, \theta_p)$, and with $\mu = \mu(\theta_1, \ldots, \theta_p)$ a focus parameter. A CD for μ has the property $\Pr_\theta\{C^{-1}(0.05, Y) \leq \mu \leq C^{-1}(0.95, Y)\} = 0.90$, etc., as with (7.1), thus delivering confidence intervals with the right coverage. This is also akin to how Bayesian posterior distributions are used. If a Bayesian prior for θ

leads to a posterior for θ, and hence for a cumulative $B(\mu \,|\, y_{\rm obs})$, then the Bayesian can read off $[B^{-1}(0.05, y_{\rm obs}) \le \mu \le B^{-1}(0.95, y_{\rm obs})]$. A question of interest and relevance also in Bayesian contexts is whether such intervals make sense also in the frequentist sense. point to Bernshteĭn–von Mises things in Ex. 6.27. the answer is 'ok' under such conditions, but not outside. xx)

(a). For a simple start example, consider $Y_1, \ldots, Y_n$ which given θ are i.i.d. from the Pois(θ), and with a prior $\theta \sim {\rm Gam}(a, b)$; see Ex. (xx suitable exercise Ch 6 xx). Show that the posterior cumulative for θ becomes $B_n(\theta \,|\, {\rm data}) = G(\theta, a + n\bar y_{\rm obs}, b+n)$, in terms for the cumulative Gamma, and with $\bar y_{\rm obs}$ the observed data average. Let $\widehat\theta_B$ and $\widehat\tau_B$ be the posterior mean and standard deviation. Assume now that data $Y_1, Y_2, \ldots$ come from the Pois(θ_0), for a certain θ_0. Show that

$$\begin{aligned} B_n(\theta_0 \,|\, Y_1, \ldots, Y_n) &= \Pr(\theta \le \theta_0 \,|\, Y_1, \ldots, Y_n) = G(\theta_0, a + n\bar y, b + n) \\ &\doteq \Phi((\theta_0 - \widehat\theta_B)/\widehat\tau_B) \to_d \Phi({\rm N}(0,1)) \sim {\rm unif}, \end{aligned}$$

with probability 1.

(b). For a similar adventure, start with the Beta(a, b) prior for a binomial probability θ. Show that the posterior cumulative for θ becomes $B_n(\theta \,|\, {\rm data}) = {\rm Be}(\theta, a + y, b + n - y)$, in terms of the Beta cumulative. Assuming that Y_n is really from a binomial with some true θ_0, show that

$$\begin{aligned} B_n(\theta_0 \,|\, Y_n) &= \Pr(\theta \le \theta_0 \,|\, Y_n) = {\rm Be}(\theta_0, a + Y_b, b + n - Y_n) \\ &\doteq \Phi((\theta_0 - \widehat\theta_B)/\widehat\tau_B) \to_d \Phi({\rm N}(0,1)) \sim {\rm unif}, \end{aligned}$$

with probability 1.

(c). Consider then the general parametric situation, supposing $Y_1, \ldots, Y_n$ being i.i.d. from the density $f(y, \theta)$, with $\mu = \mu(\theta)$ a one-dimensional focus parameter. Let θ_0 and $\mu_0 = \mu(\theta_0)$ denote the true parameters. As with Recipe Two, see Ex. 7.8, there is convergence to the standard normal of the standardised $\sqrt n(\widehat\mu_{\rm ml} - \mu_0)/\widehat\kappa$, say, with the ML estimator for μ and $\widehat\kappa^2$ consistently estimating $\kappa^2 = c^{\rm t} J^{-1} c$ (point back to delta method things in Ch2 xx). Now use the Bernshteĭn–von Mises theorem setup of Ex. 6.27 to explain that $\sqrt n(\mu - \widehat\mu_{\rm ml}) \,|\, {\rm data}$ tends in distribution to ${\rm N}(0, \kappa^2)$. The Bayesian posterior has a c.d.f. $B_n(\mu \,|\, {\rm data})$. Show that at the true value,

$$\begin{aligned} B_n(\mu_0 \,|\, {\rm data}) &= \Pr\big(\sqrt n(\mu - \widehat\mu_{\rm ml})/\widehat\kappa \le \sqrt n(\mu_0 - \widehat\mu_{\rm ml})/\widehat\kappa \,|\, {\rm data}\big) \\ &= \mu(\sqrt n(\mu_0 - \widehat\mu_{\rm ml})/\widehat\kappa) + \varepsilon_n, \end{aligned}$$

with $\varepsilon_n \to_{\rm pr} 0$. Deduce that $B_n(\mu)$ is a CD in the large-sample sense. In this sense we may think of any sensible Bayesian posterior distribution, in regular parametric models, as Recipe Seven for creating a CD. (xx a little more here. xx)

Ex. 7.18. *Confidence distribution for the ratio of explained variation.* (xx to be polished. used in Story **??**. xx) For the classic linear regression model $Y_i = x_i^{\rm t}\beta + \varepsilon_i$, with i.i.d. ${\rm N}(0, \sigma^2)$ noise terms, theory was developed in Ex. 4.37 to estimate the fraction of the variance explained via the covariates. The statistic R^2 given there can be seen as an estimator of the explained variation ratio $\rho = \beta^{\rm t}\Sigma_n\beta/(\beta^{\rm t}\Sigma_n\beta + \sigma^2)$, with notation from that exercise. Here we construct full CDs for such parameters.

(a). (xx one covariate at a time. note that σ changes value and interpretaion, depending on whichi covariates are used in the linear regression equation. xx) $\rho = M_n\beta^2/(M_n\beta^2 + \sigma^2)$, with $M_n = (1/n)\sum_{i=1}^n (x_i - \bar{x})^2$. then $F = nM_n\widehat{\beta}^2/\widehat{\sigma}^2 \sim F(1, m, n\rho/(1-\rho))$.

(b). (xx then general. xx) Explain that

$$C(\rho) = \Pr_\rho(\widehat{\lambda} \geq \widehat{\lambda}_{\rm obs}) = \Pr_\rho(F \geq n\widehat{\lambda}_{\rm obs}/p) = 1 - F(n\widehat{\lambda}_{\rm obs}/p, p, m, n\rho/(1-\rho))$$

becomes a CD for ρ.

(c). (xx one more thing. xx)

CDs for quantiles

Ex. 7.19. *CDs for quantiles.* Let $Y_1, \ldots, Y_n$ be independent observations from a smooth density f, with c.d.f. F. How can we construct CDs for its quantiles, the $\mu_q = F^{-1}(q)$? We wish such a CD to be nonparametric, without further assumptions on the f. We go through the main ideas for the case of the median $\mu = F^{-1}(\frac{1}{2})$, before extending methods and results to a general quantile $q \in (0,1)$. (xx a bit more prose here; several methods; some better than others in terms of precision and coverage; we draw but briefly on density estimators from Ch. 13. nils needs to check Price and Bonett (2001, 2002). xx)

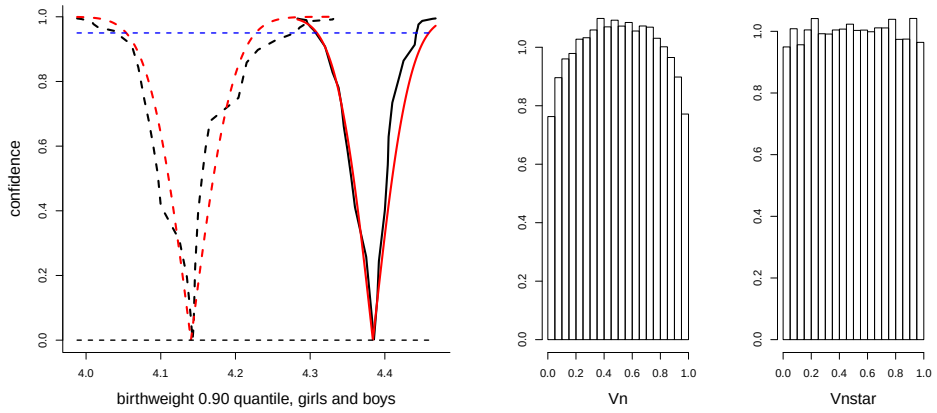


Figure 7.5: Left panel: For Ex. 7.19, confidence curves $\mathrm{cc}(\mu_{0.90})$, for the 0.90 quantiles of the birthweight distributions for girls (to the left) and boys (to the right). The black curves use the Beta method $\mathrm{cc}_n^*(\mu, \mathrm{data})$, with linear interpolation, whereas the slanted curves use the large-sample approximation. The former yields more accurate coverage than the latter. 95 percent intervals for the two 0.90 quaniles are indicated via the blue horizontal line. Right panel: For Ex. 7.20, for the case of f being standard normal, $n = 100$, $q = 0.50$: histograms of V_n and V_n^* based on sim $= 10^5$ simulations. Also in other cases the V_n^* is much closer to uniformity than is V_n.

(a). It is not difficult to construct a first-order correct CD via large-sample results reached in Chapter 3, see in particular Ex. 3.21. With M_n the sample median, we have $\sqrt{n}(M_n -$

$\mu) \to_d \mathrm{N}(0, \frac14/f(\mu)^2)$. Show that as long as $\widehat\tau$ is a consistent estimator for $f(\mu)$, then $\sqrt{n}(M_n-\mu)/(\frac12/\widehat\tau) \to_d \mathrm{N}(0,1)$, and that this leads to the approximate CD

approximate CD for quantiles

$$C_n(\mu, y) = \mu\big(\sqrt{n}(\mu - M_n)/(\tfrac12/\widehat\tau)\big).$$

One of several choices is to take $\widehat\tau = \widehat f(M_n)$, with $\widehat f(y) = n^{-1}\sum_{i=1}^n h^{-1}K(h^{-1}(y_i-y))$ a kernel density estimator, with some kernel function K and bandwidth h. The best size for this fine-tuning parameter is of the type $h = c/n^{1/5}$, as seen in Chapter 13, and a classic rule of thumb which we typically might resort to here is to take $h = 1.059\,\widehat\sigma/n^{1/2}$, with $\widehat\sigma$ the empirical standard deviation. Show that the confidence intervals from this CD take the form $M_n \pm z_0(\frac12/\widehat\tau)/\sqrt{n}$, with z_0 the relevant normal quantile, like 1.96 for intended 95 percent intervals.

(b). A different idea starts out as follows. For the ordered observations $Y_{(1)} < \cdots < Y_{(n)}$, show that

$$\Pr_f(\mu \le Y_{(i)}) = \Pr(\tfrac12 \le U_{(i)}) = 1 - \mathrm{Be}(\tfrac12, i, n-i+1) \quad \text{for } i = 1,\ldots,n.$$

Here $U_{(i)} = F(Y_{(i)})$; these form an ordered sample from the standard uniform, and we saw in Ex. 3.21 that they have Beta distributions. The $\mathrm{Be}(x,a,b)$ is the c.d.f. of a $\mathrm{Beta}(a,b)$. Define a full CD for μ, say $C_n^*(\mu, \mathrm{data})$, via linear interpolation between the $C_n^*(y_{(i)}, \mathrm{data}) = 1 - \mathrm{Be}(\frac12, i, n-i+1)$ points. This also yields a confidence curve $\mathrm{cc}_n^*(\mu, \mathrm{data}) = |1 - 2\,C_n^*(\mu, \mathrm{data})|$.

(c). Extend the two methods above, constructed there to deal with the median, to a general quantile $\mu_q = F^{-1}(q)$. For the first CD, use $\sqrt{n}(Q_{n,q} - \mu_q) \to_d \mathrm{N}(0, q(1-q)/f(\mu_q)^2)$, with $Q_{n,q} = F_n^{-1}(q)$ the empirical q quantile, and estimate $\tau_q = f(\mu_q)$ via $\widehat f(F_n^{-1}(q))$. For the second CD, show first that $\Pr_f(\mu_q \le Y_{(i)}) = 1 - \mathrm{Be}(q, i, n-i+1)$, and use linear interpolation:

$$C_n^*(\mu_q, \mathrm{data}) = \text{interpolation with } 1 - \mathrm{Be}(q, i, n-i+1) \text{ at } y_{(i)}, \tag{7.2}$$

the Beta method CD for quantiles

for $i = 1,\ldots,n$. For μ_q inside $(y_{(i)}, y_{(i+1)})$, therefore, the CD value is interpolation between $1 - \mathrm{Be}(q, i, n-i+1)$ and $1 - \mathrm{Be}(\mu_1, i+1, n-i)$. We call this *the Beta method CD* for quantiles.

(d). (xx a bit more. for birthweights oslo boys and girls, compute, display, and interpret the confidence curves for the 0.90 quantile, using both of the CD methods. Reproduce a version of Figure 7.5. point to Story i.5. xx)

Ex. 7.20. *CDs for quantiles: how well do they work?* In Ex. 7.19 we found two non-parametric CD recipes, for any quantile $\mu_q = F^{-1}(q)$. Here we investigate how well they work, in terms of actual coverage probabilities for confidence intervals. For the methods $C_n(\mu_q, \mathrm{data})$ and $C_n^*(\mu_q, \mathrm{data})$, define

$$V_n = C_n(\mu_{q,\mathrm{true}}, Y_1, \ldots, Y_n) \quad \text{and} \quad V_n^* = C_n^*(\mu_{q,\mathrm{true}}, Y_1, \ldots, Y_n),$$

with $Y_1,\ldots,Y_n$ drawn from the density in question. Accurate coverage, at all levels, means that the distribution of these two random CDs, at the true value, should be close to the uniform.

(a). To check the precision of these two CDs, carry out a simple simulation experiment. Take f equal to the standard normal, with $\mu_q = \Phi^{-1}(q)$ to be estimated with uncertainty; use $\widehat{\tau} = \widehat{f}(Q_{n,q})$ as above, with K the standard normal kernel and bandwidth $h = 1.059\,\widehat{\sigma}/n^{1/5}$ (which is optimal for the normal case), using the ordinary standard deviation from the data; and then simulate say sim $= 10^5$ values of $V_n = C_n(\mu_{q,\text{true}}, Y_1, \ldots, Y_n)$ and $V_n^* = C_n^*(\mu_{q,\text{true}}, Y_1, \ldots, Y_n)$. Check, perhaps for $n = 50, 100, 500, 1000$, how close the distributions of V_n and V_n^* are to the uniform. – For computing the C_n^*, and hence for executing that part of the simulation experiment, the `approx` algorithm of `R` is handy, carrying out linear approximation between the two values $1 - \text{Be}(\frac{1}{2}, i, n-i+1)$ and $1 - \text{Be}(\frac{1}{2}, i+1, n-i)$, for any μ inside the $[Y_{(i)}, Y_{(i+1)})$ interval.

(b). Conduct a few similar simulation experiments, to see how close V_n and $V_m n^*$ are to uniformity, with different density f, quantile μ_q, sample size n.

(c). To assess 'closeness to uniformity' more accurately, use the monitoring processes of Ex. 3.9. For each of your simulation experiments, in addition to displaying histograms of V_n and V_n^*, compute and display the functions $Z_{\text{sim}}(t) = (\text{sim})^{1/2}\{G_{\text{sim}}(t) - t\}$ and $Z_{\text{sim}}^*(t) = (\text{sim})^{1/2}\{G_{\text{sim}}^*(t) - t\}$, where G_{sim} and G_{sim}^* are the empirical distribution functions of the V_n and V_n^*. Compute also $D_{\text{sim}} = \max_t |Z_{\text{sim}}(t)|$ and $D_{\text{sim}}^* = \max_t |Z_{\text{sim}}^*(t)|$. It will transpire (i) that V_n often is not particularly close to uniformity, unless n is rather large; (ii) that V_n^* is often so close to uniformity, even for moderate n and q near 0 or 1, that we cannot see that the distribution is not uniform, even with 10^5 simulated values.

(d). xx

Ex. 7.21. *Large-sample equivalence for two CDs for quantiles.* (xx to come. details for why the two CDs are large-sample equivalent. harder to show clearly that the 2nd is better than the 1st. nils thinks that it is, though, as of 12-March-2025. xx)

CDs in some nonregular setups (change title in a while)

Ex. 7.22. *Aboriginals and invaders in Watership Down.* Suppose a population of rabbits has been living for a long time on an island, in Hardy–Weinberg equilibrium $(p_0^2, 2p_0q_0, q_0^2)$, which means that pairs of alleles aa, Aa, AA occur with these frequencies, with $q_0 = 1 - p_0$. Suppose next that there's an invading population of new rabbits, with their separate Hardy–Weinberg equilibrium $(p, 2pq, q^2)$, with $q = 1 - p$. We assume that the two populations do not mix, but live on, on the same island, and that rabbitologists don't see the difference. One is interested in learning the fraction λ of newcomers (so the fraction of aboriginals is $1 - \lambda$).

(a). Explain that when one samples n rabbits independently, and find their allele pairs aa, Aa, AA, then these numbers (X, Y, Z) have a trinomial distribution with parameters

$$\text{pr}_1 = (1-\lambda)p_0^2 + \lambda p^2, \quad \text{pr}_2 = (1-\lambda)2p_0q_0 + \lambda 2pq, \quad \text{pr}_3 = (1-\lambda)q_0^2 + \lambda q^2.$$

Note that $\text{pr}_1 + \text{pr}_2 + \text{pr}_3 = 1$.

(b). For the case of $(X, Y, Z) = (118, 438, 444)$, and assuming not only $(p_0, q_0) = (0.25, 0.75)$ known, but also $(p, q) = (0.40, 0.60)$ known, find an estimate and construct a confidence curve $\text{cc}_1(\lambda)$, as with the black smooth Figure 7.6, left panel. Assume next, with the same counts (X, Y, Z), that the home population parameters $(p_0, q_0) =$

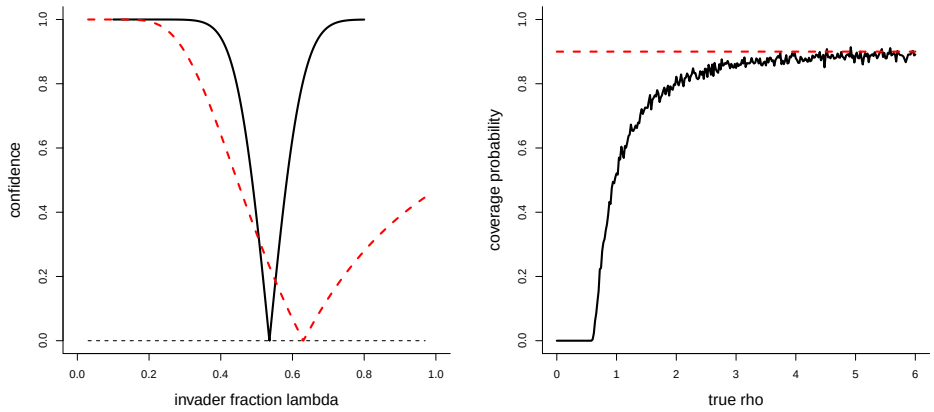


Figure 7.6: *Left panel: For Ex. 7.22, confidence curves for the unknown fraction λ of newcomers, after having counted $(X, Y, Z) = (118, 438, 444)$ of allele pairs aa, Aa, AA. The start population has HW parameters $(p_0, q_0) = (0.25, 0.75)$. (i) The black smooth $\mathrm{cc}_1(\lambda)$ is computed using the knowledge that the new population has HW parameters $(p, q) = (0.40, 0.60)$. (ii) The red slanted $\mathrm{cc}_2(\lambda)$ is computed using only knowledge about (p_0, q_0), i.e. both p, $q = 1-p$, and λ are unknown. Right panel: For Ex. 7.23, as a function of the true ρ, for dimension $p = 3$, the figure shows the actual coverage probability of the Bayesian 90 percent credibility interval, based on the posterior stemming from a flat prior for the θ.*

$(0.25, 0.75)$ are known, but that the HW parameters $(p, q) = (p, 1-q)$ for the new population are unknown. Again, estimate λ and find a confidence curves $\mathrm{cc}_2(\lambda)$, as for the red slanted curve of Figure 7.6, left panel. Comment on your findings. For your computer script, play a bit with different sample sizes, and with different degrees of difference between (p_0, q_0) and (p, q).

(c). Explain why it is not possible to estimate all (p_0, p, λ) from (X, Y, Z).

Ex. 7.23. *The length problem.* (xx might make a Satellite Collision Story, based on Cunen et al. (2020b). contrasting CD with Bayes. xx) There are several situations, of varying degrees of complexity, where the heart of the matter is, or can be transformed to, the following: with Y having the $\mathrm{N}_p(\theta, \Sigma)$ distribution, with unknown mean vector and known or partly known variance matrix, reach inference for the length $\rho = \|\theta\| = (\theta_1^2 + \cdots + \theta_p^2)^{1/2}$. See e.g. Cunen et al. (2020b) for an application involving the computation and real-time monitoring of the probability that two satellites will collide.

(a). Take first $\Sigma = I_p$, so $y \sim \mathrm{N}_p(\theta, I_p)$, which means independent $Y_i \sim \mathrm{N}(\theta_i, 1)$ for $i = 1, \ldots, p$. Show that the ML estimator of ρ is $\widehat{\rho} = \|Y\|$. Show also that $\widehat{\rho}^2 \sim \chi^2_p(\rho^2)$, the noncentral chi-squared.

(b). Deduce that $\widehat{\rho}^2$ is overshooting its target ρ^2, with mean and variance $p + \rho^2$ and $2p + 4\rho^2$. Find also an expression for $\mathrm{E}\,\widehat{\rho}$, and show that it overshoots ρ.

(c). Show that the natural CD becomes $C(\rho, y) = 1 - \Gamma_p(\widehat{\rho}^2, \rho^2)$.

(d). A typical Bayesian analysis would start with a flat prior for $\theta_1, \ldots, \theta_p$ (xx calibrate and xref Ch 5 for this detail xx). Show that $\theta \,|\, y \sim \mathrm{N}_p(y, I)$, and that this entails $\rho^2 \,|\, y \sim \chi^2_p(\widehat{\rho}^2)$.

(e). For $p = 5$ and $\widehat{\rho} = 7.77$, compute and draw both the CD and the Bayesian posterior distribution,

$$C(\rho, y) = 1 - \Gamma_p(\widehat{\rho}^2, \rho^2) \quad \text{and} \quad B(\rho, y) = \Gamma_p(\rho^2, \widehat{\rho}^2).$$

Comment on what you find.

(f). (xx simulate to illustrate that the CD by construction works, producing confidence intervals with the correct coverage; $U_C = C(\rho_0, Y) \sim$ unif when data stem from the model, at position ρ_0. show however that the Bayesian posterior distribution here risks being very far from producing intervals with the right coverage; $U_B = B(\rho_0, Y)$ is very far from being uniform. point to Figure 7.6, right panel, for the too low coverage probability of the Bayesian 90 percent credibility interval. the CD based intervals have exact coverage. link to Bernshteĭn–von Mises things in Ch. 6; here we're outside BvM terrain. more on why and how. xx)

(g). Generalise the above to the case where $Y \sim \mathrm{N}_p(\theta, \sigma^2 I_p)$.

(h). More generally, with $Y_1, \ldots, Y_n$ being i.i.d. from the $\mathrm{N}_p(\theta, \sigma^2 I_p)$, with σ known, show first that $\bar{Y} \sim \mathrm{N}_p(\theta, (\sigma^2/n) I_p)$. Then show that $\widehat{\rho} = \|\bar{y}\|$ is the ML estimator, with distribution given by $n\widehat{\rho}^2/\sigma^2 \sim \chi^2_p(n\rho^2/\sigma^2)$. On the Bayesian side, show that a flat prior for θ leads to $\theta \,|\, \text{data} \sim \mathrm{N}_p(y, (\sigma^2/n) I_p)$. Show that these statements lead to these generalisations of the above

$$\begin{aligned} C_n(\rho, y) &= 1 - \Gamma_p(n\widehat{\rho}^2/\sigma^2, n\rho^2/\sigma^2), \\ B_n(\rho \,|\, y) &= \Gamma_p(n\rho^2/\sigma^2, (n/\sigma^2)\widehat{\rho}^2). \end{aligned}$$

(i). (xx a bit more, regarding BvM, which holds for fixed p and ρ, with growing n. but misleading picture for finite n. do something to see interplay with n and p. xx)

(j). (xx also, briefly, to the case of $Y \sim \mathrm{N}_p(\theta, \sigma^2 I_p)$, with σ estimated via an independent $\widehat{\sigma}^2 \sim \sigma^2 \chi^2_m/m$. xx)

Ex. 7.24. *Ratio of normal means.* (xx edit and clean. about two exercises here on this. mention Fieller name. start with $x_0 = -a/b$, the point at which a regression type equation $a + bx = 0$. then application to bioassay or similar. xx)

(a). Consider the prototype setup for such questions, where $\widehat{a} \sim \mathrm{N}(a, 1)$ and $\widehat{b} \sim \mathrm{N}(b, 1)$ are independent. Show first that the log-likelihood is a simple $\ell(a, b) = -\frac{1}{2} Q(a, b)$, with $Q(a, b) = (a - \widehat{a})^2 + (b - \widehat{b})^2$, and find a formula for $\ell_{\mathrm{prof}}(x_0) = -\frac{1}{2} Q_{\mathrm{prof}}(x_0)$, where $Q_{\mathrm{prof}}(x_0) = \min\{Q(a, b) \colon x_0 = -a/b\}$.

(b). Show that $(\widehat{a} + \widehat{b} x_0)/(1 + x_0^2) \sim \chi^2_1$, at the true x_0, and hence that

$$\mathrm{cc}(x_0) = \Gamma_1((\widehat{a} + \widehat{b} x_0)/(1 + x_0^2))$$

is a clear confidence curve for x_0. (xx illustrate, with the mildly peculiar confidence regions. find max confidence level. more. xx)

(c). (xx with $\widehat{\sigma}$ on top. with dependence. things fine as long as $(\widehat{a}, \widehat{b})$ is binormal. xx)

(d). (xx bioassay. xx)

Ex. 7.25. *CDs and posterior distributions with boundary constraints.* Here we learn about construction of CDs when there is a boundary condition on the focus parameter. This is sometimes an easy task, involving a natural positive post-data probability on the boundary point. We also compare with Bayesian procedures. Matters may of course be extended and generalised in several directions here, but for simplicity and conciseness we study a very simple prototype situation: y is $\mathrm{N}(\theta, 1)$, and $\theta \geq 0$ a priori.

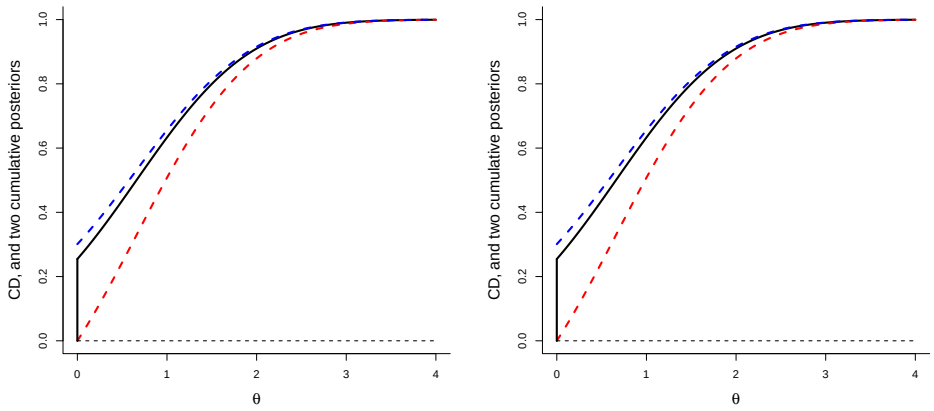


Figure 7.7: Left panel: For Ex. 7.25, with $y_{\mathrm{obs}} = 0.66$ for the $\mathrm{N}(\theta, 1)$ model, the black curve is the natural CD, with positive point mass 0.255 at zero. The red and the blue curves are Bayesian posterior distributions, for the flat prior on the halfline, and for the mixture prior with $\frac{1}{2}$ at zero and $\frac{1}{2}$ flat on the halfline, respectively. Right panel: For Ex. 7.26, (xx to come xx).

(a). Before we come to the parameter constraint, we deal with the more normal situation where there is no a priori constraint. The classical CD is then $C(\theta, y) = \Phi(\theta - y)$. Show that the Bayesian starting with a flat prior for θ finds the posterior distribution $\theta \,|\, y \sim \mathrm{N}(y, 1)$, with cumulative $B(\theta \,|\, y) = \Phi(\theta - y)$, i.e. identical to the canonical CD. – The point below will partly be that this is *not* the same for the constrained problem.

(b). For the remaining points here, assume indeed that $\theta \geq 0$ a priori. Argue that the canonical CD should be $C(\theta, y) = \Phi(\theta - y)$ for $\theta \geq 0$. Its point mass at zero is $\Phi(-y)$. Graph the CD, for the three cases y_{obs} equal to $-0.22, 0.66, 1.99$.

(c). One Bayesian approach in this situation, where $\theta \geq 0$ a priori, is to let θ be flat on $[0, \infty)$. Show that then

$$\theta \,|\, y \sim \frac{\phi(\theta - y)}{\int_0^\infty \phi(\theta - y)\,\mathrm{d}\theta} = \frac{\phi(\theta - y)}{\Phi(y)} \quad \text{for } \theta \geq 0,$$

and that the cumulative posterior distribution becomes

$$B(\theta \,|\, y) = \frac{\Phi(\theta - y) - \Phi(-y)}{1 - \Phi(-y)} = \frac{\Phi(\theta - y) - \Phi(-y)}{\Phi(y)} \quad \text{for } \theta \geq 0.$$

For the three cases of $y_{\rm obs}$ given above, graph the CD along with the Bayesian $B(\theta \,|\, y_{\rm obs})$, and comment on what you find.

(d). In general terms, for the case of $y \,|\, \theta \sim {\rm N}(\theta, 1)$, let θ have the mixture prior distribution $p_0\pi_0 + p_1\pi_1$, with the sub-priors π_0 and π_1 having their individual posteriors $\pi_0(\theta \,|\, y)$ and $\pi_1(\theta \,|\, y)$. Show that the posterior has a natural mixture form,

$$\theta \,|\, y \sim p_0^*(y)\pi_0(\theta \,|\, y) + p_1^*(y)\pi_1(\theta \,|\, y),$$

where

$$p_0(y) = \frac{p_0 f_0(y)}{p_0 f_0(y) + p_1 f_1(y)} \quad \text{and} \quad p_1(y) = \frac{p_1 f_1(y)}{p_0 f_0(y) + p_1 f_1(y)},$$

and with $f_0(y) = \int \phi(y-\theta)\pi_0(\theta)\,{\rm d}\theta$ and $f_1(y) = \int \phi(y-\theta)\pi_1(\theta)\,{\rm d}\theta$ the marginal densities following from the two priors. (This structure generalises to general mixture priors in general models, though that does not concern us just now.)

(e). For the prior $p_0\pi_0 + p_1\pi_1$, with π_0 a unit pointmass at zero and π_1 a flat prior on the halfline, show that $f_0(y) = \phi(y)$ and $f_1(y) = \Phi(y)$. With a 50-50 mixture, show hence that

$$p_0(y) = \frac{\phi(y)}{\phi(y) + \Phi(y)} \quad \text{and} \quad p_1(y) = \frac{\Phi(y)}{\phi(y) + \Phi(y)}.$$

Draw curves of these two posterior probabilities, one for the zero-point and the other for the halfline-based part, as y goes from say -5 to 5. Show that the posterior cumulative distribution becomes $B^*(\theta \,|\, y) = p_0(y) + p_1(y)B(\theta \,|\, y)$ for $\theta \geq 0$. In particular, there is a pointmass $p_0(y)$ at zero. Construct a version of Figure 7.7.

(f). Show that there is no choice of (p_0, p_1) which makes the Bayesian cumulative posterior $B^*(\theta \,|\, y)$ agree with the CD $C(\theta, y)$. Devise a method for selection (p_0, p_1) such that the distance between $B^*(\theta \,|\, y)$ and $C(\theta, y)$ is small, for a relevant range of θ and possible observed $y_{\rm obs}$.

(g). Generalise the formulae above to the case of $y_1, \ldots, y_n$ i.i.d. ${\rm N}(\theta, \sigma^2)$, with known σ.

Ex. 7.26. *CDs for regression parameters with boundary constraints.* (xx to come here: more on boundary parameters, now in simple regression models. pointer to Story iii.10. xx) (xx the point we wish to convey is that the Tore-Sims phenomenon is a general one, easier to understand and analyse in simpler models, separately. so we can have separate points for a model like $y_i = \beta_0 + \beta_1 x_{i,1} + \beta_2 x_{i,2} + \varepsilon_i$, where one has prior knowledge $\beta_2 \geq 0$. There is a clear and exact CD for β_2, of the type $C(\beta_2) = G_{\rm df}((\beta_2 - \widehat\beta_2)/\widehat\kappa_2)$ for $\beta_2 \geq 0$, with a pointmass $G_{\rm df}(-\widehat\beta_2/\widehat\kappa_2)$ at zero. the Bayesian with flat priors on $\beta_0, \beta_1, \log\sigma$ and a flat prior on $(0, \infty)$ for β_2, a la Sims, will not be able to detect that $\beta_2 = 0$; there's a clear discrepancy between the CD and the Bayesian posterior for that parameter. xx)

Ex. 7.27. *CDs in the truncated exponential model.* Here we consider a model sometimes called the truncated exponential model. We start with its simplest form, with data $Y_1, \ldots, Y_n$ i.i.d. from the density $\exp\{-(y-a)\}$ for $y \geq a$. The a is the unknown start point for the distribution.

(a). Show that the ML estimator is equal to $U_n = \min_{i \le n} Y_i$, the smallest data point. Show that $n(U_n - a)$ has a unit exponential distribution. Build from this a natural CD for a.

(b). Construct a predictive CD for the next sample point Y_{n+1}. Illustrate by computing and displaying the confidence curve for the text sample point, after having observed the six data points `3.735, 3.338, 10.634, 3.839, 5.667, 5.808`.

(c). Then consider the more realistic two-parameter version of the model, with density

$$f(y_i, a, b) = (1/b) \exp\{-(y_i - a)/b\} \quad \text{for } y_i \ge a,$$

with a being the unknown start-point and b a scale parameter. Show that the ML estimators become $\widehat{a} = U_n$ and $\widehat{b} = (1/n) \sum_{i=1}^n (Y_i - U_n)$, again with U_n being the smallest observation.

(d). Construct accurate CDs and confidence curves for a, for b, and for the next datapoint Y_{n+1}. If some of your formulae cannot be given very explicit mathematical forms, this is ok, as long as numerical solutions can be found via numerical integration or simulation. Give approximations for these CDs for large sample sizes n.

(e). Ignoring these large-sample approximations, compute and display confidence curves for a, b, Y_{n+1} with the simple $n = 6$ dataset above.

Ex. 7.28. *CD inference for the exponential rate, with censored data.* The lifelength distribution for a certain type of technical components is considered exponential, i.e. with density $\theta \exp(-\theta t)$ for $t > 0$, on a priori grounds. To arrive at a point estimate and a confidence curve for θ, the firm producing these components sets in motion the simple experiment where n such items are set to work, under controlled natural conditions. One cannot wait until all components have died out, however, and the firm needs to report what can be said about the lifelength distribution, via θ, a certain time t_0 after project start.

(a). With data of the form observed t_i for the N of the items which have died within t_0, and the information $t_i > t_0$ for the $n - N$ which are still alive and well, show that the combined likelihood function may be expressed as

$$\theta^N \exp\Big[-\theta\Big\{\sum_{t_i \le t_0} t_i + (n - N)t_0\Big\}\Big].$$

(b). Show that the ML estimator is

$$\widehat{\theta} = N/R = N/\Big\{\sum_{t_i \le t_0} t_i + (n - N)t_0\Big\}.$$

With increasing sample size, and fixed t_0, find expressions for the probability limits of N/n and R/n, and show that $\widehat{\theta}$ is consistent.

(c). Show in fact that there is a limiting normal distribution here, with $\sqrt{n}(\widehat{\theta} - \theta) \to_d \mathrm{N}(0, \tau(t_0, \theta)^2)$, and attempt to find an explicit (though not necessarily quick and simple) formula for the limit variance.

(d). Explain why the construction $C_n(\theta) = \mathrm{Pr}_\theta(\widehat{\theta} \ge \widehat{\theta}_{\mathrm{obs}})$ yields a CD, and also how it can be computed in practice.

(e). Suppose the experiment described involves $n = 20$ such items, and that the life-lengths for the $N = 11$ of these that conk out before the deadline of $t_0 = 2.00$ years are

```
0.528 0.743 0.869 1.180 0.602 0.133 0.327 1.115 0.117 0.208 1.808
```

Compute and display perhaps as many as three (exact or approximate) confidence curves for θ, for this little experiment: the one described in (c); one based on the normal approximation to the distribution of the ML estimator; and a t-bootstrap based version. Comment on your findings.

Ex. 7.29. *Estimating n based on observing the first r.* Suppose $Y_1, \ldots, Y_n$ are i.i.d., from some known distribution with density f and cumulative F, but that one only observes the first r order statistics, $Y_{(1)} < \cdots < Y_{(r)}$. Can we estimate n? Such nonstandard problems turn up in various context, from estimating the size of a vocabulary to the number of unseen species. In this exercise we consider the special case of the unit exponential distribution, where the Y_i can be seen as waiting times, so the question may be phrased as how long time do we need to wait, until we've seen all items, when we have used a certain time to observe the first r.

(a). Let then $Y_1, \ldots, Y_n$ be i.i.d. from the unit exponential, and assume $Y_{(1)} < \cdots < Y_{(r)}$ are observed, with unknown n. Observing these r first data points is equivalent to observing the spacings $D_1 = Y_{(1)}$, $D_2 = Y_{(2)} - Y_{(1)}$, up to $D_r = Y_{(r)} - Y_{(r-1)}$. Use Ex. 1.15 to show that the joint distribution of these f first spacings may be written

$$\begin{aligned} g_r(d_1, \ldots, d_r) &= n(n-1)\cdots(n-r+1) \\ &\qquad \exp\big[-\{n(d_1 + \cdots + d_r) - d_2 - 2d_3 - \cdots (r-1)d_r\}\big], \end{aligned}$$

and deduce from this that $Y_{(r)}$ is sufficient for n.

(b). With $F(x) = 1 - \exp(-x)$, show that $F(Y_{(r)})$ has a Beta distribution with parameters $(r, n - r + 1)$.

(c). Show that the optimal CD for n, based on having observed the smallest r datapoints, is

$$C_r(n) = \Pr_n(Y_{(r)} \le Y_{(r),\mathrm{obs}}) = \mathrm{Be}(F(Y_{(r),\mathrm{obs}}), r, n - r + 1).$$

(d). (xx an example or two. suppose $Y_{(r),\mathrm{obs}} = 0.348$ with $r = 33$. estimate n. see nils com87a or thereabouts. give normal approximation. but these are not good for r/n close to zero or one. can we characterise ML estimator. xx)

Ex. 7.30. *(xx another discrete model thing. xx)* (xx to come. xx)

Risk functions for CDs

Ex. 7.31. *Risk functions for CDs.* This exercise looks into risk functions for and hence comparisons between CDs, in simple prototype situations where calculations are easier than for general cases. We start out with $Y_1, \ldots, Y_n$ being i.i.d. from the $\mathrm{N}(\theta, 1)$ model. For a CD $C_n(\theta, y)$, where y denotes the full dataset, the risk function used is

$$\mathrm{risk}_n(C_n, \theta) = \mathrm{E}_\theta \int (\theta' - \theta)^2 \, \mathrm{d}C_n(\theta', Y) = \mathrm{E}_\theta(\theta_{\mathrm{cd}} - \theta)^2,$$

where θ_{cd} is the result of a two-stage random process: data Y lead to the CD $C_n(\theta, Y)$, and then θ_{cd} is drawn from this distribution.

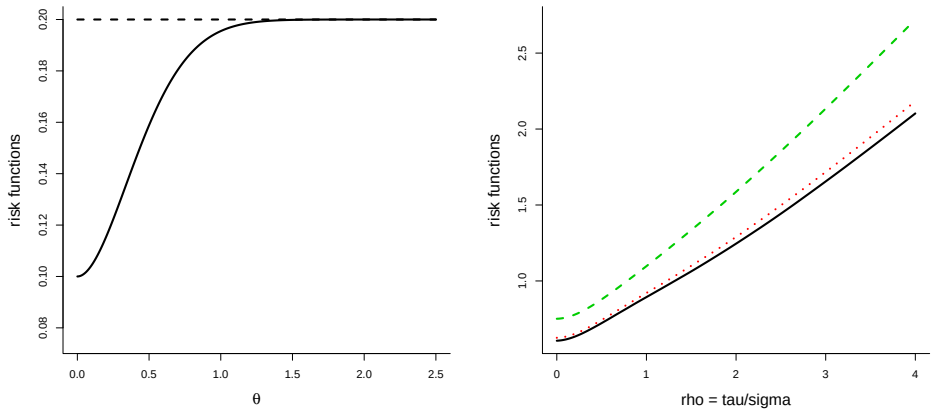


Figure 7.8: *Left panel: For Ex. 7.31, risk function for the CD, in the setup with* $Y_1, \ldots, Y_n$ *being i.i.d. from* $\mathrm{N}(\theta, 1)$*, with* $n = 10$*, but with the restriction* $\theta \geq 0$*. It starts out at* $1/n$ *and then grows to the* $2/n$ *risk of the unrestricted case as* θ *grows. Right panel: For Ex. 7.33, for the variance component model, with* $p = 4$ *and* $\sigma = 1$*, risk functions* $r(C, \tau)$ *for three CDs for* τ*. The one based on* $Z = \sum_{i=1}^{p} Y_i^2$ *is best, closely followed by the one using* $A = \sum_{i=1}^{p} |Y_i|$*, whereas the one using the range* $R = \max Y_i - \min Y_i$ *does worse.*

(a). Show that the natural CD based on the observed sample mean $\bar{y}_{\mathrm{obs}} = n^{-1} \sum_{i=1}^{n} y_i$ is $C_n(\theta, y_{\mathrm{obs}}) = \Phi(\sqrt{n}(\theta - \bar{y}_{\mathrm{obs}}))$. Prove that its risk function is $\mathrm{risk}_n(C_n, \theta) = 2/n$.

(b). More generally, assume θ^* is some unbiased estimator of θ, with finite variance τ_n^2, with the property that $\widehat{\theta}^* - \theta$ has a distribution H_n symmetric around zero. Show that the associated CD becomes $C_n^*(\theta, y_{\mathrm{obs}}) = H_n(\theta - \theta^*_{\mathrm{obs}})$, and show that its risk function becomes $2\tau_n^2$. The case of $\bar{Y}$ corresponds to $2/n$. Find the risk function for the case of the median based CD, with say $n = 10$, as for Figure 7.8, left panel.

(c). (xx fix this: Relate the above results to the optimality theorem for CDs, in certain situations, from CLP's Chapter 5. xx)

(d). Now we change gears a bit, by putting the a priori assumption $\theta \geq 0$ on the table. Show that the ML estimator becomes $\widehat{\theta} = \max(0, \bar{y})$, i.e. the sample mean truncated, if necessary, to zero. Argue that this leads to the natural CD

$$\widetilde{C}_n(\theta, y) = \Phi(\sqrt{n}(\theta - \bar{y}_{\mathrm{obs}})), \quad \text{for } \theta \geq 0,$$

in particular having a positive point-mass at zero.

(e). With θ_{cd} drawn from this CD, for given data, show that it may be expressed as $\max(0, \bar{y} + N/\sqrt{n})$, with N a standard normal. Show next that in the two-stage setup, with random data followed by θ_{cd} drawn from the $\widetilde{C}_n$ CD, we have $\theta_{\mathrm{cd}} - \theta = \max(0, \theta + (N + N')/\sqrt{n}) - \theta$, with N' another and independent standard normal. Use this to show that the $\mathrm{risk}_n(\widetilde{C}_n, \theta)$ can be expressed as

$$\begin{aligned} r_n &= \int \big[\{\max(0, \theta + (2/n)^{1/2} x)\} - \theta\big]^2 \phi(x)\, \mathrm{d}x \\ &= \theta^2 \Phi(-(\tfrac{1}{2}n)^{1/2}\theta) + (2/n)\{-(\tfrac{1}{2}n)^{1/2}\theta\phi((\tfrac{1}{2}n)^{1/2}\theta) + 1 - \Phi(-(\tfrac{1}{2}n)^{1/2}\theta)\}. \end{aligned}$$

Compute and display the risk functions for $\widetilde{C}_n$ and C_n, for say $n = 10$, constructing a version of Figure 7.8, left panel. Comment on what we learn from this.

(f). (xx not fully sure about this one. xx) There are various other estimators and CDs worth considering in this $\theta \geq 0$ setting. To simplify matters, take $n = 1$, and consider the Bayes estimator $\widehat{\theta}_B$, the conditional mean of $\theta \,|\, y$, with a flat prior on $(0, \infty)$. Show in fact that $\widehat{\theta}_B = y + \phi(y)/\Phi(y)$, and verify that this is positive even when y is negative. Work out an expression for the naturally associated CD $C_B(\theta) = \Pr_\theta(\widehat{\theta}_B \geq \widehat{\theta}_{B,\mathrm{obs}})$, and comment.

Ex. 7.32. *Optimal CDs in exponential family models.* (xx relate to Ex. 7.15. spell out NP things to demonstrate optimality. xx)

Ex. 7.33. *Risk functions for three CDs in a variance components model.* Consider the simple variance component model with independent observations $y_i \sim \mathrm{N}(0, \sigma^2 + \tau^2)$ for $i = 1, \ldots, p$, with σ known and τ the unknown parameter of interest; see Schweder and Hjort (2016, Example 4.1 and Exercise 5.8). The aim here is first to construct CDs based on (i) $Z = \sum_{i=1}^p y_i^2$, (ii) $A = \sum_{i=1}^p |y_i|$, and (iii) the range $R = \max y_i - \min y_i$; and then to compute and compare their risk functions. These are defined as

$$\mathrm{risk}(C, \tau) = \mathrm{E}_\tau \, |\tau_{\mathrm{cd}} - \tau| = \mathrm{E}_\tau \int |\tau_{\mathrm{cd}} - \tau| \, \mathrm{d}C(\tau_{\mathrm{cd}}, Y),$$

with τ_{cd} a random draw from the $C(\tau, Y)$ distribution, and with Y itself denoting a dataset drawn from the distribution indexed by τ.

(a). Show that the natural CDs, based on Z, A, R respectively, are

$$\begin{aligned} C_Z(\tau, \mathrm{data}) &= 1 - \Gamma_p(Z_{\mathrm{obs}}/(\sigma^2 + \tau^2)), \\ C_A(\tau, \mathrm{data}) &= 1 - G_p(A_{\mathrm{obs}}/(\sigma^2 + \tau^2)^{1/2}), \\ C_R(\tau, \mathrm{data}) &= 1 - H_p(R_{\mathrm{obs}}/(\sigma^2 + \tau^2)^{1/2}). \end{aligned}$$

Here Γ_p is the cumulative distribution function of $Z_0 = \sum_{i=1}^p N_i^2$, with the N_i being i.i.d. and standard normal, which means $Z_0 \sim \chi^2_p$. Similarly, G_p and H_p are the cumulative distribution functions of $A_0 = \sum_{i=1}^p |N_i|$ and of $R_0 = \max N_i - \min N_i$, respectively.

(b). Show that a random draw τ_{cd} from the first of these, i.e. C_Z, for a given dataset, can be represented as $\tau_{\mathrm{cd}} = (Z_{\mathrm{obs}}/K - \sigma^2)_+^{1/2}$, where x_+ is notation for the truncated-to-zero quantity $\max(x, 0)$, and where $K \sim \chi^2_p$. In the situation where data are random, from the model at position τ, deduce that

$$\tau_{\mathrm{cd}} - \tau = \{(\sigma^2 + \tau^2)K_0/K - \sigma^2\}_+^{1/2} - \tau = \sigma\big[\{(1 + \rho^2)K_0/K - 1\}_+^{1/2} - \rho\big],$$

where $\rho = \tau/\sigma$, and K_0, K are two independent draws from the χ^2_p. In other words, $F = K_0/K \sim F_{p,p}$, a F distribution with degrees of freedom (p, p). Use this to compute the risk function $\mathrm{risk}(C_Z, \tau)$, for $p = 4$ and $\sigma = 1$; this is the lowest of the three risk functions of Figure 7.8, right panel.

(c). Then consider the C_A option. Show that a random draw from an observed $C_A(\tau, \mathrm{data})$ can be written $\tau_{\mathrm{cd}} = \{(A_{\mathrm{obs}}/A)^2 - \sigma^2\}_+^{1/2}$. Deduce that for random data behind the CD, we have the representation

$$\tau_{\mathrm{cd}} - \tau = \{(\sigma^2 + \tau^2)(A_0/A)^2 - \sigma^2\}_+^{1/2} - \tau = \sigma\big[\{(1 + \rho^2)(A_0/A)^2 - 1\}_+^{1/2} - \rho\big],$$

with A and A_0 two independent draws from the G_p distribution. Use this to compute risk(C_A, τ). There is no simple expression for the density of A_0/A, so use simulation.

(d). Carry out similar analysis for the third CD, based on the range R. Construct a version of Figure 7.8, right panel.

(e). Use your programme to explore the three risk functions for other values of p.

Meta-analysis, combining confidence distributions

(xx need more sorting and polish. xx)

Ex. 7.34. *CD and cc for binomial probabilities.* Suppose y is observed from a binomial (n, θ). The task is to construct a CD and a cc for θ.

(a). Show that the standard normal approximations for y (xx give pointer here to large-sample chapter) lead to

$$C_a(\theta, y) = \Phi\Big(\frac{n\theta - y}{\{n\theta(1-\theta)\}^{1/2}}\Big) \quad \text{and} \quad C_b(\theta, y) = \Phi\Big(\frac{\sqrt{n}(\theta - \widehat{\theta})}{\{\widehat{\theta}(1-\widehat{\theta})\}^{1/2}}\Big),$$

with $\widehat{\theta} = y/n$ the standard estimator for θ.

(b). The recipe of Ex. 7.7 does not quite work here since y has a discrete distribution. This invites the half-correction method

$$C(\theta, y) = \Pr_\theta(y > y_{\rm obs}) + \tfrac{1}{2}\Pr_\theta(y = y_{\rm obs}).$$

For say $n = 20$ and $y = 12$, compute and display this CD, along with (i) the same CD but without the half-correction, (ii) the two simple normal approximations above. Try other combinations of (n, y), and demonstrate that they are approximately equal for moderate to large n.

(c). To investigate the basic CD property, take $n = 20$ and $\theta_{\rm true} = 0.33$. Simulate a large number of $C(\theta_0, y)$, $C_a(\theta_0, y)$, $C_b(\theta_0, y)$, to check for their approximate uniform distribution. Try other values of (n, θ_0), and summarise your findings.

Ex. 7.35. *CD and cc for comparing binomials.* In Ex. 7.34 we learn how to construct CDs for separate binomial parameters. Consider now the 2×2 table setup with two binomials, as with Ex. 4.30, say $y_0 \sim \text{binom}(m_0, p_0)$ and $y_1 \sim \text{binom}(m_1, p_1)$. How do we reach precise inference for the extent to which p_0 and p_1 differ?

(a). We first use the logistic transform $p_0 = H(\theta_0)$ and $p_1 = H(\theta_0 + \gamma)$, with $H(u) = \exp(u)/\{1 + \exp(u)\}$. Show that

$$\gamma = \log\frac{p_1/(1-p_1)}{p_0/(1-p_0)} = \log\frac{p_1}{1-p_1} - \log\frac{p_0}{1-p_0},$$

the log-odds difference. Write up the likelihood function for the observed (Y_0, Y_1) to deduce (xx via the optimal CD exercise xx) that the optimal CD for γ takes the form

$$C(\gamma) = \Pr_\gamma(Y_1 > y_{1,\rm obs} \,|\, Z = z_{\rm obs}) + \tfrac{1}{2}\Pr_\gamma(Y_1 = y_{1,\rm obs} \,|\, Z = z_{\rm obs}),$$

with $Z = Y_0 + Y_1$. The conditional distribution in question is the eccentric hypergeometric, found in Ex. 4.30. (xx do simple example here. this CD is used in both Stories i.1 and i.11. we use Ex. 4.30. xx)

Ex. 7.36. *Meta-analysis for Lidocain data.* The following data table is from Normand (1999), and pertains to prophylactic use of lidocaine after a heart attack. The aim is to evaluate mortality from prophylactic use of lidocaine in acute myocardial infarction. We view the data here as pairs of binomials, with $y_{1,i} \sim \text{binom}(m_{i,1}, p_{1,i})$ and $y_{1,0} \sim \text{binom}(m_{i,0}, p_{1,0})$.

m1	m0	y1	y0
39	43	2	1
44	44	4	4
107	110	6	4
103	100	7	5
110	106	7	3
154	146	11	4

(a). Write the probabilities in logistic fashion, i.e. $p_{i,0} = H(\theta_{i,0})$ and $p_{i,1} = H(\theta_{i,0} + \gamma_i)$, with $H(u) = \exp(u)/\{1 + \exp(u)\}$. Show that

$$\gamma_i = H^{-1}(p_{i,1}) - H^{-1}(p_{0,i}) = \log \frac{p_{i,1}}{1 - p_{i,1}} \Big/ \frac{p_{0,i}}{1 - p_{i,0}},$$

the log-odds difference. Construct and display the optimal CD for the γ_i, and also for the odds ratio $\rho_i = \exp(\gamma_i)$, for each of the six studies.

(b). Assume then that the log-odds parameter γ is the same, across studies, so that the six binomial data pairs relate to seven parameters. Find the optimal CD for this γ, and for the common odds ratio $\rho = \exp(\gamma)$. Translate the CDs to confidence curves, and display the six + one curves in a diagram. How would you conclude?

Ex. 7.37. *Comparing Poisson parameters.* (xx ranting on a bit, to be edited. point to Story i.12, application to sucide attempts rates. xx) Suppose $Y_0 \sim \text{Pois}(m_0\theta_0)$ and $Y_1 \sim \text{Pois}(m_1\theta_1)$. In what precise way is θ_1 different from θ_0? Writing $\gamma = \theta_1/\theta_0$, show that the likelihood is proportional to $\exp\{-\theta_0(m_0 + m_1\gamma)\}\theta_0^{y_0+y_1}\gamma^{y_1}$. Explain that the optimality recipe tells us inference should be made based on the distribution of $Y_1 \,|\, (Z = z)$, where $Z = Y_0 + Y_1$. Show that $Y_1 \,|\, (Z = z)$ has the binomial distribution $(z, m_1\gamma/(m_0 + m_1\gamma))$. Show how this leads to the optimal CD

$$\begin{aligned} C(\gamma) &= \text{Pr}_\gamma(Y_1 > y_{1,\text{obs}} \,|\, z) + \tfrac{1}{2}\text{Pr}_\gamma(Y_1 = y_{1,\text{obs}} \,|\, z) \\ &= 1 - B_z(y_{1,\text{obs}}, m_1\gamma/(m_0 + m_1\gamma)) + \tfrac{1}{2}b_z(y_{1,\text{obs}}, m_1\gamma/(m_0 + m_1\gamma)). \end{aligned}$$

(xx point to Story i.12, for $y_0 = 1$, $y_1 = 7$, for the patient years m_0, m_1, in Aursnes et al. (2005). compare with their Bayes gamma priors, both informative and less informative. xx)

Ex. 7.38. *Basic meta-analysis.* (xx to come, and calibrated with later stuff. xx) There is a very wide literature on combining information, with different names and labels, including meta-analysis, data fusion, etc. This exercise looks into some of the more basic versions, and where CDs will be helpful in later extensions below.

(a). Suppose $y_j \sim \text{N}(\mu, \sigma_j^2)$, for $j = 1, \ldots, k$ independent sources, with the same focus parameter μ, and with variances taken to be known or well estimated. Consider the linear combination estimator $\widehat{\mu} = \sum_{j=1}^k a_j y_j$. Show that it is unbiased, provided $\sum_{j=1}^k a_j = 1$,

and find its variance. Show that the best choice, yielding minimal variance among the unbiased ones, is $a_j \propto 1/\sigma_j^2$, leading to

$$\widehat{\mu} = \frac{\sum_{j=1}^k y_j/\sigma_j^2}{\sum_{j=1}^k 1/\sigma_j^2}.$$

Show indeed that $\widehat{\mu} \sim \mathrm{N}(\mu, \kappa^2)$, with this minimal variance being $\kappa^2 = (\sum_{j=1}^k 1/\sigma_j^2)^{-1}$. Comment on what this leads to for the case where the σ_j are equal.

(b). In various settings there is a need to generalise the setting above to one where $y_j \,|\, \mu_j \sim \mathrm{N}(\mu_j, \sigma_j^2)$, with these individual mean parameters not being equal, but having their own distribution, say $\mu_j \sim \mathrm{N}(\mu_0, \tau^2)$. The task is then to reach inference for both the overall mean μ_0 and the spread τ among the μ_j. Show that $y_j \sim \mathrm{N}(\mu_0, \sigma_j^2 + \tau^2)$, and that the log-likelihood function becomes

$$\ell(\mu_0, \tau) = -\tfrac{1}{2} \sum_{j=1}^k \Big\{ \log(\sigma_j^2 + \tau^2) + \frac{(y_j - \mu_0)^2}{\sigma_j^2 + \tau^2} \Big\}.$$

(c). Considering the spread parameter τ first, show that the profiled log-likelihood can be written

$$\ell_{\mathrm{prof}}(\tau) = -\tfrac{1}{2} \sum_{j=1}^k \Big[\log(\sigma_j^2 + \tau^2) + \frac{\{y_j - \widehat{\mu}_0(\tau)\}^2}{\sigma_j^2 + \tau^2} \Big], \quad \text{for } \tau \geq 0,$$

in which

$$\widehat{\mu}_0(\tau) = \frac{\sum_{j=1}^k y_j/(\sigma_j^2 + \tau^2)}{\sum_{j=1}^k 1/(\sigma_j^2 + \tau^2)}$$

is the best linear combination estimator for μ_0, for the fixed τ under inspection.

(d). (xx fix this, needs to be clearer. xx) Consider this profiled log-likelihood as a function of $\gamma = \tau^2$, rather than of τ, and show that its derivative at zero is

$$D = \tfrac{1}{2} \sum_{j=1}^k \frac{1}{\sigma_j^2} \Big\{ \frac{(y_j - \widetilde{\mu}_0)^2}{\sigma_j^2} - 1 \Big\}.$$

Here $\widetilde{\mu}_0 = \widehat{\mu}(0)$. A small D means that the y_j data have a low spread, and vice versa. Show that if $D \leq 0$, then $\widehat{\tau}_{\mathrm{ml}} = 0$, and if that $D > 0$, then $\widehat{\tau}_{\mathrm{ml}}$ is positive.

(e). (xx fix this, needs to be clearer. xx) Show next that

$$Q(\tau) = \sum_{j=1}^k \frac{\{y_j - \widehat{\mu}(\tau)\}^2}{\sigma_j^2 + \tau^2} \sim \chi^2_{k-1}.$$

(f). (xx work with $\ell_{\mathrm{prof}}(\tau)$. partly from CLP. derivative at zero. xx) By maximising over μ_0, for each given τ, show that

$$\ell_{\mathrm{prof}}(\tau) = -\tfrac{1}{2} \sum_{j=1}^k \Big[\log(\sigma_j^2 + \tau^2) + \frac{\{y_j - \widehat{\mu}(\tau)\}^2}{\sigma_j^2 + \tau^2} \Big].$$

Ex. 7.39. *Combining CDs for the same parameter.* (xx a few exercises here. first for the same parameter, basic. then to CLP Ch 13 settings, then CDs to likelihood; then II-CC-FF. xx) Suppose that $C_1(\mu), \ldots, C_k(\mu)$ are independent CDs for the same parameter μ, perhaps based on different sets of data. How can these be properly combined?

(a). Show that $N_j = \phi^{-1}(C_j(\mu_{\rm true}))$ is standard normal, at the true position in the parameter space underlying the C_j. With $w_1, \ldots, w_k$ numbers such that $\sum_{j=1}^k w_j^2 = 1$, show that

$$\bar C(\mu) = \Phi\Big(\sum_{j=1}^k w_j \Phi^{-1}(C_j(\mu))\Big)$$

is a proper combination CD for μ.

(b). (xx point back to Ex. 7.38. xx)

(c).

Ex. 7.40. *The problem of the Nile.* (xx starting a rant and will see how it goes. xx) Let $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_n$ be independent and exponential, the X_i with parameter θ, the Y_i with parameter $1/\theta$.

(a). Show that the log-likelihoods, for the X_i and for the Y_i parts, become

$$\ell_1(\theta) = n(\log\theta - \theta\bar X), \quad \ell_2(\theta) = n(-\log\theta - \bar Y/\theta),$$

with separate ML estimators $\widehat\theta_1 = 1/\bar X$ and $\widehat\theta_2 = \bar Y$. Show further that we may represent these as $\widehat\theta_1 \sim \theta A$, $\widehat\theta_2 \sim \theta B$, where $A = 2n/K$ and $B = L/(2n)$, with K and L being χ^2_{2n}.

(b). Show that the canonical CDs for θ, based on the two parts, are

$$C_1(\theta) = \Gamma_{2n}(2n\theta/\widehat\theta_1), \quad C_2(\theta) = 1 - \Gamma_{2n}(2n\widehat\theta_2/\theta).$$

The combination recipe hence leads to

$$\bar C(\theta) = \Phi\big((1/\sqrt2)\{\Phi^{-1}(C_1(\theta)) + \Phi^{-1}(C_2(\theta))\}\big).$$

(c). The full log-likelihood becomes $\ell(\theta) = -n(\theta\bar X + \bar Y/\theta)$. Show that this is maximised for

$$\widehat\theta_{\rm ml} = (\bar Y/\bar X)^{1/2} = (\widehat\theta_1\widehat\theta_2)^{1/2} \sim \theta F^{1/2},$$

where $F = A/B = L/K \sim F_{2n,2n}$. Show that the associated with this ML estimator is $C_{\rm ml}(\theta) = 1 - F_{2n,2n}(\widehat\theta^2_{\rm ml}/\theta^2)$.

(d). Run some simple experiments, where you for a few values of n simulate data and then plot the four CDs, and the four confidence curves (the two separate ones, the combination one, and the ML based one).

(e). (xx a bit more. ML etc. work, even though the situation is slightly irregular. $(\bar X, \bar Y)$ is sufficient, but not complete. we need two informants for the single parameter. xx)

Ex. 7.41. *From CD to likelihood.* (xx to come. with illustrations. normal conversion. $\ell(\mu) = -\frac12\Gamma_1^{-1}({\rm cc}_j(\mu))$. xx)

Ex. 7.42. *II-CC-FF: Independent Inspection, Confidence Conversion, Focused Fusion.* (xx to come. using Cunen and Hjort (2022). point to Bayesian updating being part of this, but allows user keeping only prior for the focus parameter. aim to demonstrate iiccff in Story i.12. xx)

Ex. 7.43. *Private attributes.* (xx to be checked with care. xx) The probability ψ of cheating at exams might be hard to estimate, but a bit of randomisation might grant anonymity and yield valid estimates. Suppose there are three cards, two with the statement 'I did cheat', and 'I did not cheat' on the third. Students are asked to draw one of the three cards randomly and answer either *true* or *false* to the drawn statement, without revealing it.

(a). Show that the probability of *true* is $(1 + \psi)/3$. Assume a binomial model for the number of students answering *true*, and devise a CD for ψ.

(b). Assume 1000 students go through the simple post-exam exercise above (anonymously). Find and display CDs for ψ for the cases of respectively 300, 350, 400 out of the 1000 answered *true*.

Notes and pointers

[xx A few remarks. xx]

we point to Schweder and Hjort (2016, Example 3.11), Fisher (1930), Xie and Singh (2013), Hjort and Schweder (2018), Cunen and Hjort (2022), Singh et al. (2005), ...

(xx repair. we point to Story iii.10. xx) There's a notable discrepancy between the frequentist Schweder-Hjort CD and the Bayesian posterior distribution associated with a flat prior on the $[0, \infty)$ interval, in cases where the y_{obs} is close to, or perhaps even to the left of, the boundary point.

I.8

Loss, risk, performance, optimality

Statistics is a mathematical formalisation of how to make good decisions under uncertainty. One source of uncertainty is that the true state of nature, say θ, is not known when we are to make our decisions, and since the utility or loss of a decision depends on θ, we need to be clear about how bad it is when our decision is off. This is the role of loss functions, of which the squared error is the most well known example. Decisions are based on data, and it is not anodyne how we use the data: Some procedures for going from data to a decision are better than others. Therefore, it makes sense to see how a certain procedure performs on average. This is the role played by risk functions. Risk functions average out the data, but they still depend on θ, so the risk function of various decision procedures are often difficult to order. Some might be preferable for certain values of θ, while others might be better for other values of θ. This chapter introduces criteria that let us, nevertheless, say something about how good a decision procedure is. A decision procedure is said to be *admissible* if there is no other that does better, in terms of the risk function, whatever the truth or future may be. One says that a decision procedure is *minimax* if it is the the best one in the most unfortunate situation. In proving that certain decision procedures are admissible or minimax, Bayesian thinking is an essential tool. This includes the concepte of Bayes risk, where not only the data is average out, but also the various possible states of nature are averaged out.

Key words: admissibility, Rao–Blackwell, loss functions, minimaxity, multiparameter estimation, UMVU estimators, risk functions

A statistical decision, as most decisions in life when you think about it, is a function of what we observe to the space of all possible decisions we can make in a given setting: I look out the window and see grey clouds, and choose to take my umbrella with me when I go out. I see that you are smiling and I think that you are happy. Formally, an action is a function $\widehat{a}\colon \mathcal{Y} \to \mathcal{A}$, where $\mathcal{Y}$ is the space in which the data take its values, the *sample space*; while $\mathcal{A}$ is the *action space*, that is the collection of all possible actions we might take. How wise our choice of $\widehat{a} \in \mathcal{A}$ is or turns out to be depends on the true *state of nature* θ. This θ lives in a *parameter space* Θ, and is an unknown *parameter* governing the probability distribution P_θ from which the data $Y \in \mathcal{Y}$ are generated. A *loss function* $L(\theta, a)$ measures 'how much' we lose by choosing action a when the true state of nature is θ. We assume that

loss function

$$L\colon \Theta \times \mathcal{A} \to [0, \infty),$$

so that the best possible loss is zero (in general, loss functions do not need to be non-negative, see e.g. Schervish (1995, Chapter 3.1) for a more general introduction). To be concrete, consider the point estimation problem with data $Y_1, \ldots, Y_n$ independent $\mathrm{N}(\theta, 1)$, where θ is an unknown parameter to be estimated under the loss function $L(\theta, a) = (a - \theta)^2$. Here, the amount lost is the squared distance between a and θ. If we wish to test $H_0 \colon \theta \leq 0$ versus $H_A \colon \theta > 0$, the action space is $\mathcal{A} = \{\text{keep } H_0, \text{ reject } H_0\}$, and a natural loss function can be described by

$$L(\theta, a) = \begin{array}{r|cc} & \theta \leq 0 & \theta > 0 \\ \hline \text{keep } H_0 & 0 & 1 \\ \text{reject } H_0 & 1 & 0 \end{array}$$

With this loss function, we lose the same amount when rejecting a true null hypothesis as when failing to reject a null hypothesis that is false. In statistical jargon that you probably already know, failing to reject a null hypothesis is called a Type I error, while when we fail to reject a null hypothesis that is false, we commit a Type II error.

The general setup so far is essentially as for what we built with the Bayesian framework of Ch. 6, see Ex. 6.2, where one in addition to the elements above pushes in one more ingredient, that of a prior distribution for the θ. We do comment more below, and in various exercises, on Bayesian components and connections, but the main spirit of the present chapter is the classical or frequentist one, with no priors. Here different decision procedures are compared by the loss they incur for each value of θ, that is, by their *risk function*

$$R(\theta, \widehat{a}) = \mathrm{E}_\theta \, L(\theta, \widehat{a}(Y)) = \int_{\mathcal{Y}} L(\theta, \widehat{a}(y)) \, \mathrm{d}P_\theta(y).$$

Notice that for each decision procedure $\widehat{a}$, the risk function $R(\theta, \widehat{a})$ is a function of θ, so that for two estimators $\widehat{a}_1$ and $\widehat{a}_2$ their respective risk functions may cross, that is $R(\theta, \widehat{a}_1) < R(\theta, \widehat{a}_2)$ for some θ, while $R(\theta, \widehat{a}_1) > R(\theta, \widehat{a}_2)$ for other values of θ. Thus, comparison of risk functions only provides a partial ordering of decision procedures, and as such does not point clearly at a best decision procedure. What is clear, however, is that if $R(\theta, \widehat{a}_1) \leq R(\theta, \widehat{a}_2)$ for all θ, with strict inequality for at least one value of θ, then $\widehat{a}_2$ should be discarded from the competition: We say that $\widehat{a}_1$ *dominates* $\widehat{a}_2$, and $\widehat{a}_2$ is said to be *inadmissible*. A decision procedure that is not dominated by any other decision procedure is *admissible*. In and of itself admissibility does not tell us much about an estimator. Consider for example the estimator $\widehat{a}'(Y) = \theta'$ that returns the value θ' whatever the data. Clearly, no estimator can perform better than $\widehat{a}'$ at position θ', but that does not, for obvious reasons, make it an estimator we would like to use.

admissibility

One (not very principled) fix to the problem of comparing estimators is to limit the search of a best estimator to the class of estimators that are unbiased: An estimator $\widehat{\theta}$ is unbiased for the parameter θ if $\mathrm{E}_\theta \, \widehat{\theta} = \theta$ for all θ. Another principle by which to compare decision procedures is the the *minimax principle*. According to this principle, the estimator with the best performance in the worst possible scenario ought to be chosen. We say that a decision rule $a^\star$ is *minimax* if it minimises the maximum risk, that is if there are no competitors $\widetilde{a}$ with $\sup_{\theta \in \Theta} R(\theta, \widetilde{a})$ smaller than $\sup_{\theta \in \Theta} R(\theta, a^*)$.

unbiased estimator

minimax

(xx calibrate text here in view of Ch6. xx) If you are a Bayesian and venture into the Ch. 6 terrain of constructing prior distributions $\pi(\theta)$ over the parameter space Θ, the problem of risk functions only being partially ordered can be circumvented. What

you are interested in then is the *Bayes risk* of the decision procedures you are comparing, that is

Bayes risk

$$\mathrm{BR}(\widehat{a}, \pi) = \mathrm{E}_\pi \, R(\theta, \widehat{a}) = \int_\Theta R(\theta, \widehat{a}) \, \pi(\theta) \, \mathrm{d}\theta.$$

For a given $\widehat{a}$ and prior π, the Bayes risk is a number, and under certain conditions that we will explore below, there will be a unique decision procedure $\widehat{a}$ that, for a given prior π, minimises $\mathrm{BR}(\widehat{a}, \pi)$. This decision procedure is the *Bayes solution*. As we shall see, Bayes solutions are, despite the name, of high importance also for the frequentists.

Bayes solution

UMVU estimator, Rao–Blackwell, and Lehmann–Scheffé

Ex. 8.1. *Coin tossing.* To get a feeling for some of the basic challenges and concepts concerning the comparison of various estimator, we start out with the emblematic problem of estimating the probability of heads in $n = 10$ independent tosses of a coin. Let $Y_1, \ldots, Y_n$ be independent Bernoulli random variables with expectation θ.

(a). Sketch the risk function of the maximum likelihood (ML) estimator $\bar{Y}_n = n^{-1} \sum_{i=1}^n Y_i$ under squared error loss $L(\theta, \delta) = (\delta - \theta)^2$. Recall that $n = 10$.

(b). Suppose we have some intuition about where on the unit interval the probability θ might be located, close to a value $0 < \theta_0 < 1$, say. One way in which such a prior hunch might be employed is by taking as our estimate a convex combination of the ML estimator and θ_0, that is

$$\delta_a(Y_1, \ldots, Y_n) = a\bar{Y}_n + (1 - a)\theta_0,$$

for some $0 \leq a \leq 1$. For $a = 1/2$ and $\theta_0 = 1/2$, sketch the risk function of this estimator. Suppose that your task, as was Pierre-Simon Laplace's in 1781 or so, is to estimate the probability of giving birth to a boy. Which of the two above estimators do you prefer, the ML estimator or δ_a with $a = 1/2$ and $\theta_0 = 1/2$?

(c). Based on the risk functions you sketched in (a) and (b), we see that the two estimators are difficult to compare. The ML estimator has lower risk than δ_a for certain values of θ, while δ_a performs better for other values of θ. The risk functions cross and none of the two is uniformly better than the other. An easy fix to this problem of comparison is to limit our search for an estimator to the class of estimators that are unbiased for what we are estimating. Look back at Ex. 5.15 and explain why, when the search is restricted to the class of unbiased estimators, the ML estimator is the clear winner.

(d). The risk of the estimators from (a) and (b) vary widely with what the true θ is. Choosing a best estimator, when the yardsstick is the squared error loss function, seems therefore to require some prior hunch about where θ really is. A criterion for risk function comparison that does not require such a prior hunch, is minimaxity: An estimator is minimax if it minimises the maximum risk. Estimators whose risk functions are constant are, as we will soon see, good candidates for being minimax. Consider the estimator from Ex. (b) with $\theta_0 = 1/2$. Find a function $a = a(n)$ such that the risk function $R(\theta, \delta_{a(n)})$ is constant. For $n = 10$, sketch the risk function of your estimator (that is, draw a line). Suppose you have absolutely no idea whatsoever about where in the unit interval θ may be located. Which of your three estimators of θ do you prefer? In Ex. 8.9 we learn that $\delta_{a(n)}$ is indeed minimax.

Ex. 8.2. *Uniformly minimum variance unbiased estimators.* As the sketch in Ex. 8.1 illustrates, comparing risk functions is not always straightforward, and a somewhat *ad hoc* way of making the problem of finding a best estimator tractable is by limiting the search for a best estimator to the class of unbiased estimators. What is variably called a best unbiased estimator, the uniformly minimum variance unbiased estimator, the UMVU estimator, is defined asfollows: An estimator $\delta^\star(Y)$ is the UMVU estimator for $g(\theta)$ if it is unbiased for $g(\theta)$, and for any other estimator $\delta(Y)$ that is unbiased for $g(\theta)$, it holds that $\mathrm{Var}_\theta\, \delta^\star(Y) \leq \mathrm{Var}_\theta\, \delta(Y)$ for all θ. When feasible (see Ex. 5.14 and 5.15–5.16), the easiest way of establishing that an unbiased estimator is an UMVU estimator is to verify that it achieves the Cramér–Rao lower bound.

UMVU estimator

(a). Let us look at a few examples of the Cramér–Rao approach: (i) Suppose $Y \sim \mathrm{N}(\theta, 1)$, and show that Y is UMVU for θ. (ii) Suppose Y has an exponential distribution with mean θ. Show that Y is UMVU for θ. (iii) Let $Y_i = \beta_0 + \beta_1 x_i + \sigma\varepsilon_i$ for $i = 1, \ldots, n$, where the covariates $x_1, \ldots, x_n$ are fixed numbers, and the $\varepsilon_1, \ldots, \varepsilon_n$ are independent standard normal random variables. Show that the least squares estimator for (β_0, β_1) is the UMVU estimator.

(b). Let $Y_1, \ldots, Y_n$ be i.i.d. $\mathrm{N}(\mu, \sigma^2)$. It is immediate from (a) that $\bar{Y}_n = n^{-1}\sum_{i=1}^n Y_i$ is UMVU for μ. In Ex. 8.3 we will see that there is no unbiased estimator of σ^2 attaining the Cramér–Rao lower bound.

Ex. 8.3. *Cramér–Rao and Cauchy–Schwarz.* The proof of the Cramér–Rao inequality that we met in Ex. 5.15–5.16 is a clever application of the Cauchy–Schwarz inequality.

(a). Let X and Y be two square integrable random variables with expectation zero. Show that $|\mathrm{E}\, XY| = \{\mathrm{Var}\,(X)\mathrm{Var}\,(Y)\}^{1/2}$ if and only if X and Y are linearly related, $Y = a + bX$, for example.

(b). Explain why the Cramér–Rao inequality is an equality if and only if the estimator $\delta(y)$ and the score function are linearly related, that is

$$u(y, \theta) = \partial \log f(y, \theta)/\partial\theta = a(\theta) + b(\theta)\delta(Y), \quad \text{for all } \theta,$$

for some functions $a(\theta)$ and $b(\theta)$. Solve the differential equation above for $f(y, \theta)$, and state what this entails for estimators and distributions when it comes to possible attainment of the Cramér–Rao lower bound.

(c). In Ex. 8.2(b) we saw that with $Y_1, \ldots, Y_n$ i.i.d. $\mathrm{N}(\mu, \sigma^2)$, the unbiased estimator $\widehat{\sigma}_n^2 = \sum_{i=1}^n (Y_i - \bar{Y})^2/(n-1)$ for σ^2 does not attain the Cramér–Rao lower bound. Use the result from (b) to argue that the Cramér–Rao lower bound may only be attained when μ is known.

Ex. 8.4. *Sufficiency and Rao–Blackwell.* Suppose that Y has a distribution from a family $\{P_\theta : \theta \in \Theta\}$ of distributions. Recall from Ex. **??** that $T = T(Y)$ is a sufficient statistic for this family distributions if the conditional distribution of Y given T does not depend on θ. The Rao–Blackwell theorem says that any estimator can be improved upon by conditioning on a sufficient statistic. This is an important result, as it tells us that in our search for a best estimator, we need only consider those estimators that are functions of a sufficient statistic.

(a). Let $\delta'(Y)$ be an unbiased estimator for $g(\theta)$, and suppose that $T(Y)$ is sufficient for θ. Consider the estimator given by $\delta(Y) = \mathrm{E}_\theta\{\delta'(Y) \,|\, T\}$. Then $\delta(Y)$ is better than $\delta'(Y)$. The proof proceeds in three steps. First, explain why $\delta(Y)$ is an estimator; second, show that $\delta(Y)$ is unbiased, and third, show that $\mathrm{Var}_\theta\, \delta(Y) \leq \mathrm{Var}_\theta\, \delta'(Y)$ for all θ. You have now proven the Rao–Blackwell theorem for unbiased estimators. In Ex. 8.8 we look at this theorem in a more general decision theoretic framework.

Rao–Blackwell theorem

(b). Suppose that $Y_1, \ldots, Y_n$ are i.i.d. uniforms on $[0, \theta]$, with $\theta > 0$ an unknown parameter. Show that the estimators

$$\delta_1 = \frac{2}{n}\sum_{i=1}^{n} Y_i \quad \text{and} \quad \delta_2 = \frac{n+1}{n}\max_{i\leq n} Y_i$$

are both unbiased for θ. There are (at least) two ways of showing that δ_2 is a better estimator than δ_1. Try them both. First, compute the variances of both estimators. Second, appeal to the Rao–Blackwell theorem, that is, more concretely, use the results from Ex. 3.17 to establish that $\mathrm{E}\,(\delta_1 \mid \max_{i\leq n} Y_i) = \delta_2$ almost surely.

(c). Let $Y_1, \ldots, Y_n$ be i.i.d. Pois(λ). We seek to estimate $\theta = \Pr(Y_1 = 0) = \exp(-\lambda)$. Find the ML estimator for θ, say $\widehat{\theta}$, and show that $\mathrm{E}\,\widehat{\theta} = \theta\{1 + O(1/n)\}$, meaning the the ML estimator for θ is biased. Next, find the Cramér–Rao lower bound for unbiased estimators of θ. Look back at Ex. 8.3 and consider whether this lower bound can be attained. Another estimation strategy is to estimate $\theta = \Pr(Y_1 = 0)$ by the share of zero counts, that is $\widetilde{\theta} = n^{-1}\sum_{i=1}^{n} I\{Y_i = 0\}$. This estimator is clearly unbiased for θ, why is it not best unbiased? Finally, with the aid of the sufficient statistic $T(Y) = \sum_{i=1}^{n} Y_i$ we Rao–Blackwellise the estimator $\widetilde{\theta}$. Derive an expression for this Rao–Blackwellised estimator, $\widetilde{\theta}_{\mathrm{rb}}$, say. Does $\widetilde{\theta}_{\mathrm{rb}}$ attain the Cramér–Rao lower bound?

Ex. 8.5. *Best unbiased, completeness, and Lehmann–Scheffé.* From the Rao–Blackwell theorem we know that any candidate for being an UMVU estimator must be a function of a sufficient statistic. This limits our search. In this exercise we establish that in our search for the UMVU estimator, we are only looking for one estimator. Thereafter, it is shown that an estimator is best unbiased if and only if it is uncorrelated with all unbiased estimators of zero. This yields a characterisation of the best unbiased estimators, albeit one of limited utility as it is, without further conditions, hard to describe all unbiased estimators of zero. Finally, completeness – a condition on the distribution of the data – is introduced, ensuring that the only unbiased estimator of zero is zero itself.

(a). Suppose that δ is UMVU for $g(\theta)$, and that so is δ'. Let $\delta'' = \delta/2 + \delta'/2$, and use the Cauchy–Schwarz inequality to show that $\mathrm{Var}_\theta\, \delta'' \leq \mathrm{Var}_\theta\, \delta$. But since δ is an UMVU estimator and δ'' is unbiased (check it), it must be the case that $\mathrm{Var}_\theta\, \delta'' = \mathrm{Var}_\theta\, \delta$. Look back at the results in Ex. 8.3, and use this to establish that if δ and δ' are both UMVU, then $\delta = \delta'$ almost surely, for all θ.

UMVU estimator is unique

(b). [xx rewrite xx] Suppose that δ is an unbiased estimator for $g(\theta)$ and a function of a sufficient statistic. How may we improve on δ? Well, the family of estimators $\delta_a = \delta + a\varepsilon$ as a ranges of the real numbers, and ε is some mean zero random variable, constitute a class of unbiased estimators. Show that if $\mathrm{cov}_\theta(\delta, \varepsilon) \neq 0$ for some θ, then a may be chosen so that $\mathrm{Var}_\theta(\delta_a) < \mathrm{Var}_\theta(\delta)$ for some value(s) of θ, which entails that δ is not best unbiased. Prove the converse, namely that if δ is unbiased and $\mathrm{cov}_\theta(\delta, \varepsilon)$ for all θ and all mean zero random variables ε, then δ is UMVU.

characterisation of the UMVU estimator

(c). It is in general no easy task to show that an unbiased estimator, or more generally, a statistic $T = T(Y)$ say, is uncorrelated with all unbiased estimators of zero. Since the correlation between any random variable and zero is zero, the task would be much easier if we knew of the distribution of T that the only unbiased estimator of zero, is zero itself. That is, if for any measurable function h, $\mathrm{E}_\theta\, h(T) = 0$ implies $\Pr_\theta\{h(T) = 0\} = 1$, for all θ. We recall from Ex. 4.23 that a family of distributions with this property is called *complete*. Alternatively, we just say that the statistic $T(Y)$ is complete. Suppose that T is sufficient and complete for θ. Let $\delta = \delta(T)$ be unbiased for $g(\theta)$. Prove the Lehmann–Scheffé theorem, that is, show that the estimator δ is the unique UMVU estimator for $g(\theta)$.

Lehmann–Scheffé theorem

(d). Look back at Ex. 8.4(b). Show that the estimator δ_2 is the UMVU estimator.

(e). The completeness requirement in the Lehmann–Scheffé theorem was motivated by the characterisation in (b), saying that an estimator is UMVU if and only if it is uncorrelated with all unbiased estimators of zero. A perhaps more illuminating motivation comes from the fact, proven in Ex. 8.4(a), that a best estimator must be based on a sufficient statistic. Intuitively, by getting rid of information irrelevant to the estimation problem at hand, we reduce the variance of our estimation procedure. Taking this intuition to its logical conclusion, we deduce that a best estimator must be based on a statistic achieving the maximum amount of data compression, while still retaining all the information in the data about the parameter we seek to estimate. In other words, a best estimator must be based on a minimal sufficient statistic. Recall from Ex. 4.21 that a statistic S is minimal sufficient if for any sufficient statistic T there exists a measurable function g so that $S = g(T)$. Show that if δ is an unbiased estimator, we form the estimators $\delta' = \mathrm{E}\,(\delta \,|\, T)$ and $\delta'' = \mathrm{E}\,(\delta \,|\, S)$, where T is sufficient and S is minimal sufficient, then $\mathrm{Var}(\delta'') \leq \mathrm{Var}(\delta')$. Now, suppose that T is sufficient *and* complete. Use the Lehmann–Scheffé theorem and (a) to conclude that δ' and δ'' must be almost surely equal. In view of this equality, it may not come as a surprise that if T is sufficient and complete, then T is minimal sufficient. A fact we will prove in (g).

(f). Here is a toy example illustrating some of the points made in (e). Let X_1 and X_2 be independent Bernoulli(θ) random variables and consider the estimator $\widehat{\theta} = (X_1 + X_2)/2$ and the estimator $\delta = \delta(X_1, X_2)$ given by

$$\delta(x_1, x_2) = \begin{cases} 1, & (x_1, x_2) = (1, 1), \\ 2/3, & (x_1, x_2) = (1, 0), \\ 1/3, & (x_1, x_2) = (0, 1), \\ 0, & (x_1, x_2) = (0, 0). \end{cases}$$

Explain why both $\widehat{\theta}$ and δ are sufficient for θ. Show that δ is unbiased for θ, and show that the variance of δ exceeds the variance of $\widehat{\theta}$ for all values of $\theta \in (0, 1)$.

(g). The results quoted at the end of (e) is Bahadur's theorem: If T is sufficient and complete, then T is minimal sufficient. To prove this, let W be another sufficient statistic, and assume, with out loss of generality, that T and W are real valued. We must show that there is a function g such that $T = g(W)$. If $T = \mathrm{E}_\theta\,(T \,|\, W)$, then we have found a function g, it is $g(w) = \mathrm{E}_\theta(T \,|\, W = w)$, and g does not depend on θ since W is sufficient. Let us therefore prove that T equals $\mathrm{E}_\theta\,(T \,|\, W)$ almost surely for all θ. To this end, assume that T has finite variance, and define $g(W) = \mathrm{E}_\theta\,(T \,|\, W)$ and

Bahadur's theorem

$h(T) = \mathrm{E}_\theta\{g(W) \,|\, T\}$. Now, use the tower property of conditional expectation a couple of times and that T is complete to show that $T = h(T)$ almost surely, for all θ. Next, combine the above with the variance decomposition formula to obtain

$$\mathrm{Var}_\theta\, g(W) = \mathrm{E}_\theta\, \mathrm{Var}_\theta(g(W) \,|\, T) + \mathrm{E}_\theta\, \mathrm{Var}_\theta(T \,|\, W) + \mathrm{Var}_\theta\, g(W),$$

from which we conclude that $T = \mathrm{E}_\theta\,(T \,|\, W)$ almost surely for all θ. To get rid of the finite variance assumption on T, replace T by $f(T) = 1/\{1 + \exp(-T)\}$ (which clearly has finite variance) throughout the proof, to conclude that $T = f^{-1}(g(W))$.

Ex. 8.6. *A uniform mean.* Let $Y_1, \ldots, Y_n$ be i.i.d. unif(a, b). We are to estimate the mean $\mu = (a + b)/2$.

(a). Show that $(\min_{i\leq n} Y_i, \max_{i\leq n} Y_i)$ is sufficient and complete.

(b). Propose an unbiased estimator for μ, and find its variance.

Ex. 8.7. *A weird unbiased estimator.* Limiting our search for a best estimator to the class of unbiased estimators lacks the decision theoretic foundation that the principle of minimising expected loss enjoys. More on this in the Notes and Pointers section. Sometimes, the search for unbiasedness might lead us astray. Let Y be a random variable with density

$$f(y, \theta) = \frac{\theta^y \exp(-\theta)}{y!\{1 - \exp(-\theta)\}}, \quad \text{for } y = 1, 2, \ldots,$$

with $\theta > 0$. This is a Poisson distribution truncated at zero, and the probability of being truncated is $\exp(-\theta)$.

(a). Show that $\delta_u(Y) = (-1)^{Y+1}$ is the unique unbiased estimator for $\exp(-\theta)$.

(b). Find an expression for the risk function of $\delta_1(Y)$.

(c). Propose an estimator with uniformly smaller risk than $\delta_1(Y)$.

Ex. 8.8. *More Rao–Blackwellisation.* Recall Jensen's inequality from Ex. 1.30, that $g(\mathrm{E}\, X) \leq \mathrm{E}\, g(x)$ for convex functions g. Jensen's inequality also holds for conditional expectations, see Ex. A.28(h). In this exercise we will look at loss functions $L(\theta, \delta)$ that are convex in δ for all θ. Think of your favourite loss function, and you will realise that this is a quite natural requirement.

(a). Let $\delta = \delta(X)$ be an estimator of θ. Suppose that $T = T(X)$ is sufficient for θ and define $\delta^\star(T) = \mathrm{E}\{\delta(X) \,|\, T\}$. Explain why $\delta^\star(T)$ is an estimator.

(b). Suppose $L(\theta, \delta)$ is convex in δ for all θ. Show that $R(\delta^\star, \theta) \leq R(\delta, \theta)$ for all θ.

(c). When will the inequality in (b) be strict for all θ?

(d). Suppose that $\mathrm{E}_\theta\, \delta^\star(T) = h(\theta)$, and that T is complete. Show that, in the class of estimators $\{\delta \colon \mathrm{E}_\theta\, \delta = h(\theta)\}$, the estimator $\delta^\star$ is the unique estimator minimising the risk. [xx check this for the general case here presented xx].

(e). [xx Let $L(\theta, \delta) = (\delta - \theta)^2$ and specialise to UMVU estimator xx]

(f). Suppose $L(\theta, \delta)$ is convex in δ for all θ, and that δ_π is the unique Bayes solution under the prior π. Show that δ_π must be a function of a sufficient statistic.

Minimaxity

Ex. 8.9. *Tools for minimaxity.* In Ex. 8.1 we compared three different estimators, which is fine, but what we ultimately want to say something about is the performance of an estimator compared to *all* other estimators. To do so, we need some more tools. The setup is as in the general introduction to the chapter, with data Y from a model with parameters θ, and a loss function $L(\theta, a)$ for actions a.

(a). Consider any prior π and any decision function $\widehat{a} = \widehat{a}(Y)$, with consequent risk function $R(\theta, \widehat{a})$ and Bayes risk $\mathrm{BR}(\widehat{a}, \pi)$. Writing $R_{\max}(\widehat{a})$ for the maximum risk $\sup_{\theta\in\Theta} R(\theta, \widehat{a})$, show that $\mathrm{BR}(\widehat{a}, \pi) \leq R_{\max}(\widehat{a})$, and that $\mathrm{MBR}(\pi) \leq R_{\max}(\widehat{a})$.

(b). Let δ_π be a Bayes solution with respect to the prior distribution π, and suppose that

$$\mathrm{BR}(\delta_\pi, \pi) = \sup_\theta R(\theta, \delta_\pi), \tag{8.1}$$

i.e. $\mathrm{MBR}(\pi) = R_{\max}(\delta_\pi)$. Show that δ_π is minimax.

(c). Show that if δ_π satisfies (8.1) and is the unique Bayes solution with respect to π, then δ_π is the unique minimax procedure.

(d). Show that if a Bayes solution has constant risk, then it is minimax.

(e). Show that if an estimator has constant risk and is admissible, it is minimax.

(f). Show that if an estimator is unique minimax, it is admissible.

Ex. 8.10. *Minimax estimators for binomial and multinomial parameters.* In Ex. 6.8 and 6.12 we worked through certain Bayesian priors and their associated posteriors and posterior means, along with results for their risk functions. Those efforts now lead quickly to the additional quality marks for the derived estimators: they are minimax. The points below merely spell this out.

(a). First, consider $Y \sim \mathrm{binom}(n, p)$, with p to be estimated with quadratic loss. Show that the Bayes estimator with the $\mathrm{Beta}(\frac{1}{2}\sqrt{n}, \frac{1}{2}\sqrt{n})$ prior is $p^* = (\frac{1}{2} + Y/\sqrt{n})/(\sqrt{n}+1)$, that its risk function is the constant $\frac{1}{4}/(\sqrt{n}+1)^2$, and that it is minimax. Show also that it is fact admissible. For which interval of p is the minimax risk smaller than $R(p, Y/n)$ for the usual default estimator?

(b). Then consider a multinomial $(Y_1, \ldots, Y_k)$, with parameters $(n, p_1, \ldots, p_k)$, with sums $\sum_{j=1}^k Y_j = n$ and $\sum_{j=1}^k p_j = 1$, i.e. there are $k-1$ free parameters for the k categories. With joint estimation of the probabilities, with loss function $\sum_{j=1}^k (a_j - p_j)^2$, start with the Dirichlet prior $\sqrt{n}(1/k, \ldots, 1/k)$, show that the posterior mean estimator takes the form

$$p_j^* = \frac{1/k}{\sqrt{n}+1} + \frac{\sqrt{n}}{\sqrt{n}+1}\frac{Y_j}{n} \quad \text{for } j = 1, \ldots, k.$$

and verify that its risk function is contant $(1-1/k)/(\sqrt{n}+1)^2$. Explain that it is minimax and admissible. In which region of $(p_1, \ldots, p_k)$ is the minimax estimator better than the ML estimator, i.e. $(Y_1/n, \ldots, Y_k/n)$?

(c). Explain the implications of this minimax find when it comes to constructing histograms for continuous data, see Ex. 1.6 for pertinent details, say with equal binwidths.

Ex. 8.11. *Minimaxity and sequences of priors.* In Ex. 8.9(c) we assumed that the equality in (8.1) is attained. A prior distribution that succeeds in attaining this equality is, for natural reasons, called a *least favourable* prior distribution. If no such prior distribution exists, we cannot use the conclusion of that exercise to prove minimaxity. Consider independent $X_1, \ldots, X_n \,|\, \theta$ from $\mathrm{N}(\theta, 1)$. It seems reasonable that a least favourable prior for θ should spread its mass evenly out on the real line, that is

$$\int_a^{a+c} \pi(\theta)\,\mathrm{d}\theta = \int_b^{b+c} \pi(\theta)\,\mathrm{d}\theta, \qquad \text{for all } a, b \in \mathbb{R} \text{ and } c > 0.$$

This distribution is Lebesgue measure on $\mathbb{R}$, and is not a proper probability distribution. This hints at the result above not being applicable. To fix this, the idea is to approximate an improper distributions with proper ones. In the case of the normals, one may try $\theta \sim \pi_k(\theta)$, where $\pi_k(\theta)$ is the density of a uniform distribution over $[-k, k]$, then let k grow.

(a). Suppose that δ is an estimator and $(\pi_k)_{k\geq 1}$ a sequence of prior distributions such that

$$\sup_\theta R(\theta, \delta) = \lim_{k\to\infty} \mathrm{BR}(\delta_{\pi_k}, \pi_k),$$

with δ_{π_k} being the Bayes solution for π_k. Show that δ is minimax.

(b). Let $X_1, \ldots, X_n$ be independent $\mathrm{N}(\xi, \sigma^2)$. We are to estimate ξ under the squared error loss $L(\xi, \widehat{\xi}) = (\widehat{\xi} - \xi)^2$. You may consider the sequence of priors $\mu \sim \mathrm{N}(0, \tau_k)$ for $k = 1, 2, \ldots$ to show that the estimator $\widehat{\xi} = \bar{X}_n$ is minimax.

(c). Let $X_1, \ldots, X_n$ be independent $\mathrm{Pois}(\theta)$. We want to estimate θ under the weighted loss function $L(\theta, \widehat{\theta}) = \theta^{-1}(\widehat{\theta} - \theta)^2$. Use Gamma priors to show that $\widehat{\theta} = \bar{X}_n$ is minimax.

Admissibility and Bayes

Ex. 8.12. *Some Bayes and some admissibility.* If we have at hand an estimator δ, the most convenient way of showing that δ is admissible is to show that it is Bayes. In fact, it is almost true that an estimator is admissible if and only if it is Bayes. We'll get to the cases where this implication fails, but as a rule of thumb it is pretty safe.

(a). Suppose that $X \sim f_\theta$, where $\theta \in \Theta = \{\theta_1, \ldots, \theta_k\}$ for some finite $k \geq 2$. Consider the estimator δ_π that is Bayes for the prior $\pi = \{\pi_1, \ldots, \pi_k\}$, where π_j is the prior mass given to θ_j. Show that if $\pi_j > 0$ for $j = 1, \ldots, k$, then δ_π is admissible.

(b). Why does the conclusion of Ex. 8.12(a) fail if $\pi_j = 0$ for one or more j?

(c). Show that if a Bayes solution is *unique*, then it is admissible. Or, equivalently, if every Bayes rule with respect to a prior π has the same risk function, then they are all admissible.

(d). To clarify what the uniqueness of the Bayes solution refers to, let's look at an example where there are several Bayes solutions. Let $X \sim \mathrm{unif}(0, \theta)$ and suppose that we want to estimate θ under the squared error loss function $L(\theta, \delta) = (\delta - \theta)^2$. Suppose θ is given the

prior distribution that is uniform on $(0, c)$. Find the posterior distribution $\theta \,|\, (X = x)$ and derive at least two different Bayes solutions (there are uncountably many). [xx comment on this exercise, more relevant in BNP, point to Nils' 1976-proof and the mistake made by Lehmann. Could also mention Lindley and his so-called Cromwell's rule xx].

(e). Recall that a parameter value θ is in the *support* of π (a probability density in our notation) if it is contained in the set $\{\theta \in \Theta \colon \pi(\theta) > 0\}$. Let $\{f_\theta \colon \theta \in \Theta\}$ be a model, and suppose that (i) the support of the prior π is Θ; and that (ii) the risk function $R(\theta, \delta)$ is continuous in θ for all estimators δ. Show that if δ_π is Bayes with respect to π and have finite Bayes risk, then δ_π is admissible.

Ex. 8.13. *Generalised Bayes.* (xx need a bit more nils-emil Ch6-Ch8 calibration; some of this might be in Ch6. xx) Suppose that in some experiment involving data from a normal distribution with expectation θ, you have no idea whatsoever about where on the real line θ might be located. A natural 'prior' is therefore $\pi(\theta) \propto 1$ (or just take it equal to one) that spreads the 'probability' mass uniformly over the real line. Now, $\pi(\theta) \propto 1$ corresponds to Lebesgue measure on the real line, and is not a probability measure because its integral is infinite. The fact that $\pi(\theta)$ does not integrate to one does not, however, stop us from using it to derive estimators using the Bayes theorem, or a stretched version thereof. Priors that are not probability distributions are called improper priors. improper priors

(a). Suppose $Y_1, \ldots, Y_n$ are i.i.d. $\mathrm{N}(\theta, \sigma^2)$, with known σ but unknown mean. Suppose θ is given the improper prior $\pi(\theta) = 1$. With $\bar{Y}$ as usual the sample mean, show that $\pi(\theta \,|\, y_1, \ldots, y_n)$ is the normal density $\mathrm{N}(\bar{y}, \sigma^2/n)$, noting that this is a real density, as soon as $n \geq 1$ datum is observed, even though the prior is improper. A generalised Bayes estimator is the estimator δ minimising the posterior expected loss $\mathrm{E}\{L(\theta, \delta) \,|\, \text{data}\}$. With squared error loss, show that the generalised Bayes estimator for θ is $\bar{y}$.

(b). The estimator $\widehat{\theta} = \bar{Y}$ you found in (a) is generalised Bayes; why is this not enough to conclude that it is admissible?

(c). (xx decide on balance Ch6 Ch8 here. xx) Let us generalise the above to normal data from $\mathrm{N}(\theta, \sigma^2)$ with both parameters unknown. It is mathematically convenient to turn to the parametrisation $\lambda = 1/\sigma^2$. With (θ, λ) having the improper prior $\pi(\theta, \lambda) = 1/\lambda$, and writing $Q_0 = \sum_{i=1}^n (y_i - \bar{y})^2$, show that the posterior becomes

$$\begin{aligned} \pi_n(\theta, \lambda) &\propto \lambda^{n/2-1} \exp\{-\tfrac{1}{2}\lambda\{Q_0 + n(\theta - \bar{y})^2\} \\ &\propto \lambda^{1/2} \exp\{-\tfrac{1}{2}\lambda(\theta - \bar{y})^2\} \lambda^{(n-1)/2-1} \exp(-\tfrac{1}{2}\lambda Q_0). \end{aligned}$$

Translate this to saying (i) $(1/\sigma^2) \,|\, \text{data} \sim \mathrm{Gam}(\frac{1}{2}(n-1), \frac{1}{2}Q_0)$; (ii) $\theta \,|\, (\text{data}, \sigma) \sim \mathrm{N}(\bar{y}, \sigma^2/n)$. Show that the posterior mean of $1/\sigma^2$ is $1/\widehat{\sigma}^2$ with the familiar empirical variance $Q_0/(n-1)$.

Ex. 8.14. *The normal sample mean is admissible, in dimension 1.* We tend to general sample size n below, but start wout considering the prototype setup with a single Y from the normal $(\theta, 1)$. The natural estimatoris $\bar{Y}$ itself. Here we show that it is admissible, under squared error loss, but the arguments used do not work for dimension $k \geq 2$.

(a). We have seen in Ex. 8.11 that $\bar{Y}$ is minimax, with constant risk function 1. Explain via Ex. 6.10 that it is not a pure Bayes estimator.

(b). We employ the prior $\pi_\tau \sim \mathrm{N}(0, \tau^2)$ for θ, with increasing τ, resempling the flat prior under which we have seen that $\bar{Y}$ is generalised Bayes. Show that the Bayes estimator is $\widehat{\theta}_\tau = \rho Y$, with $\rho = \tau^2/(1+\tau^2)$, with minimum Bayes risk equal to this ρ. You may redo that argument leading from this to minimaxity of of Y.

(c). Assume now that Y is not admissible. Explain that there must be a competitor $\widetilde{\theta}$ with risk at most 1 and with $R(\theta, \widetilde{\theta}) \leq 1 - \varepsilon$, for some small ε, for some interval $[a, b]$. Deduce from this that

$$\frac{\tau^2}{\tau^2+1} \leq \int R(\theta, \widetilde{\theta}) \pi_\tau(\theta)\, \mathrm{d}\theta \leq 1 - \varepsilon \int_{[a,b]} \phi(\theta/\tau)(1/\tau)\, \mathrm{d}\theta,$$

and that

$$\varepsilon(1/\tau) \int_{[a,b]} \phi(\theta/\tau)\, \mathrm{d}\theta \leq 1/(\tau^2+1) \quad \text{for all } \tau.$$

Show that this leads to a contradiction: Y is admissible.

(d). Generalise the above to any sample size; with $Y_1, \ldots, Y_n$ being i.i.d. $\mathrm{N}(\theta, \sigma^2)$, with given σ, then the sample mean $\bar{Y}$ is minimax and amissible.

(e). Consider the k-dimensional situation, with independent components $Y_i \sim \mathrm{N}(\theta_i, 1)$ for $i = 1, \ldots, k$, and with the vector of means to be estimated with loss $\sum_{i=1}^k (a_i - \theta_i)^2$. Using a sequence of independent priors $\mathrm{N}(0, \tau^2)$, for $i = 1, \ldots, k$, show that $Y = (Y_1, \ldots, Y_k)$ is minimax, with constant rate k. Then try the argument above, and explain that it does not work, for $k \geq 2$. We lean indeed in Ex. 8.20 that Y is inadmissible, for $k \geq 3$; there are competitors which uniformly dominate Y.

Ex. 8.15. *Blyth's method.* We saw in Ex. (b) that admissibility may be established via contradiction. Here we formalise that arguemt, sometimes called Blyth's method. When it works, it works, often in a case by case modus; we saw for the case of a multidimensional normal mean that the Blyth method only worked for the one-dimensional case. In general, say that δ is a *limiting Bayes* estimator if there is a sequence $\{\pi_k\}_k$ of possibly improper priors such that the corresponding Bayes estimators δ_{π_k} converge almost surely to δ.

(a). Let $\delta^\star$ be an estimator. Suppose $\Theta \subset \mathbb{R}^p$ is open, and that $R(\delta, \theta)$ is continuous in θ for all estimators δ. Let $(\pi_k)_{k \geq 1}$ be a sequence of (possibly improper) prior distributions such that $\mathrm{BR}(\delta^\star, \pi_k) < \infty$ for all k, and for any open set $\Theta_0 \subset \Theta$,

$$\frac{\mathrm{BR}(\delta^\star, \pi_k) - \mathrm{BR}(\delta_{\pi_k}, \pi_k)}{\int_{\Theta_0} \pi_k(\theta)\, \mathrm{d}\theta} \to 0, \quad \text{as } k \to \infty;$$

Prove that $\delta^\star$ is admissible.

(b). Check arguments from Ex. again, to see these general arguments in action, leading to admissibility for normal means in dimension 1 but not in higher dimension.

(c). Consder $Y \sim \mathrm{Pois}(\theta)$ with loss function $(a - \theta)^2/\theta$. Show that the estimator Y has constant risk. Use Gamma priors $(1, b)$, with growing b, to show that Y is minimax and admissible. Attempt to see how far the argument goes for the case of $k \geq 2$ Poisson variables, with loss function $\sum_{i=1}^k (a_i - \theta_i)^2/\theta_i$; see Ex. 8.22.

Ex. 8.16. *Bernoulli mean with weighted risk.* Let $X_1, \ldots, X_n$ be i.i.d. Bernoulli(θ). We wish to estimate θ, and we are particularly interested in precise estimates of very small and very large values of θ. Therefore, we'll work with the loss function

$$L(\theta, \delta) = (\delta - \theta)^2 / \{\theta(1-\theta\}.$$

(a). Compute the risk function of the ML estimator. What's noticeable about this risk function?

(b). We now take a Bayesian point of view and give θ a Beta$(a\theta', a(1-\theta'))$ prior distribution. Compute the expectation and variance of this prior.

(c). With the prior introduced in (b), find the posterior distribution $\pi(\theta \,|\, x_1, \ldots, x_n)$. Find also the Bayes solution δ_π, i.e., the minimiser of the Bayes risk $\mathrm{BR}(\delta, \theta) = \int R(\delta, \theta)\pi(\theta)\,\mathrm{d}\theta$. [xx introduce Bayes risk earlier xx].

(d). Tweak the parameters of the Beta prior distribution, so that the Bayes solution you found above equals the ML estimator from (a). What desirable properties does the ML estimator possess?

Ex. 8.17. *Estimating the standard deviation.* Suppose $X_1, \ldots, X_n$ are i.i.d. from $\mathrm{N}(0, \sigma^2)$. We are to estimate σ under the loss function $L(\sigma, \delta) = (\delta - \sigma)^2/\sigma$.

(a). Find the ML estimator, say $\widehat{\sigma}_{\mathrm{ml}}$, and show that its risk function is

$$R(\sigma, \widehat{\sigma}) = \sigma\{(n-1)/n + b_n^2 + (b_n - 1)^2\}, \quad \text{where} \quad b_n = \sqrt{\frac{2}{n}}\frac{\Gamma(n/2)}{\Gamma((n-1)/2)}.$$

You may now use Stirling's formula $\Gamma(z) = (2\pi/z)^{1/2}(z/e)^z$ to show that $b_n \to 1$, and $R(\sigma, \widehat{\sigma}) \to 2\sigma$ as $n \to \infty$, as we already knew from ML-theory (see Ex. xx in Chapter 5).

(b). Consider the prior distribution $\sigma \sim \pi(\sigma)$, whose density is

$$\pi(\sigma) \propto (1/\sigma)^{a+1} \exp(-b/\sigma^2).$$

with $a > 1$ and $b > 0$. Find the prior expectation of σ. Find also the prior expectation of $1/\sigma$.

(c). Find the posterior distribution $\sigma \,|\, x_1, \ldots, x_n$, and derive the Bayes solution under the loss function above.

(d). Show that the ML estimator is inadmissible by exhibiting an estimator, say $\delta^\star$, with uniformly smaller risk. *Hint*: Consider $\delta_\alpha = \alpha\widehat{\sigma}_{\mathrm{ml}}$.

(e). Is $\delta^\star$ admissible? *Hint*: Use Blyth's method.

Ex. 8.18. *Bayes and minimax normal estimation with the linex loss.* We worked out some basic properties of the linex loss function $\exp\{c(t-\theta)\} - 1 - c(t-\theta)$ in Ex. 6.5. Here we use the Bayesian machinery to find a minimax estimator for the normal mean.

(a). Consider the simple prototype setup where a single X has the $\mathrm{N}(\theta, 1)$ distribution. Show that the estimator $X + d$ has risk function

$$r_c(\theta) = \mathrm{E}_\theta\left[\exp\{c(X+d-\theta)\} - 1 - c(X+d-\theta)\right] = \exp(cd + \tfrac{1}{2}c^2) - 1 - cd,$$

constant in θ, and that the best estimator of this sort is $\theta^* = X - \frac{1}{2}c$. Show also that the risk achieved, by this estimator, is $\frac{1}{2}c^2$.

(b). Now consider Bayes estimation, with the prior $\theta \sim \mathrm{N}(0, \tau^2)$. Show via Ex. 6.16 that $(\theta \,|\, x) \sim \mathrm{N}(\rho x, \rho)$, with $\rho = \tau^2/(\tau^2+1)$. Show, perhaps via expressing $\theta \,|\, x$ as $\rho x + \rho^{1/2} N$ with N a standard normal, that the posterior expected loss is

$$\mathrm{E}\,\{L_c(\theta, t) \,|\, x\} = \exp\{c(t-\rho x) + \tfrac{1}{2}\rho c^2\} - 1 - c(t-\rho x).$$

Deduce that the Bayes estimator is $\widehat{\theta}_B = \rho x - \frac{1}{2}\rho c$, and that the posterior expected loss is $\mathrm{E}\,\{L_c(\theta, \widehat{\theta}_B) \,|\, x\} = \frac{1}{2}\rho c^2$, independent of x.

(c). Show that $\theta^* = X - \frac{1}{2}c$ is minimax. Show also, via Blyth's method, that it is in fact admissible.

(d). Generalise the above to the case of a full sample $X_1, \ldots, X_n$ from $\mathrm{N}(\theta, 1)$. Find the Bayes estimator and associated minimum Bayes risk, for the $\mathrm{N}(0, \tau^2)$ prior, and prove that $\theta^* = \bar{X} - \frac{1}{2}c/n$ is minimax. What is its minimax risk?

(e). Find the distribution of $Z_n = \sqrt{n}(\theta^* - \theta)$, and comment on its limit, (i) when the loss-skewness parameter c is fixed, (ii) when $c = \sqrt{n}$.

(f). (xx perhaps another example with linex loss. and something where we see certain arguments lead to a choice of c. xx)

Ex. 8.19. *More on the linex loss.* For the linex loss, studied initially in Ex. 6.5 and then in Ex. 8.18 for the normal case, we now find out more.

(a). Suppose $Y_1, \ldots, Y_n$ are i.i.d. from the $\mathrm{Pois}(\theta)$, with prior $\theta \sim \mathrm{Gam}(a, b)$, as with Ex. 6.1. Show that the Bayes estimator with the linex loss is $\widehat{\theta}_B = (1/c)(a + n\bar{y}) \log\{1 + c/(b+n)\}$. Verify that when $c \to 0$, we retrieve the posterior means of Ex. 6.1.

(b). (xx one more case with a clear formula. perhaps σ in normal. xx)

(c). As we know from Ex. 6.27, an approximation to the posterior distribution is $\theta \,|\, \mathrm{data} \sim \mathrm{N}(\widehat{\theta}_{\mathrm{ml}}, \widehat{\sigma}^2/n)$, in terms of the ML estimate and estimated inverse Fisher information. Deduce that $M_n(-c) \doteq \exp(-c\widehat{\theta}_{\mathrm{ml}} + \frac{1}{2}c^2\widehat{\sigma}^2/n)$, in the notation above, and that this leads to the approximation $\widehat{\theta}_B = \widehat{\theta}_{\mathrm{ml}} - \frac{1}{2}c\widehat{\sigma}^2/n$ for the Bayes estimator under linex loss. Show also that the posterior expected loss is approximately $\frac{1}{2}c^2\widehat{\sigma}^2/n$.

(d). (xx an example where we can check the approximation with the exact Bayes estimator, e.g. with Poisson and gamma. xx)

Multiparameter estimation and the Stein phenomenon

Ex. 8.20. *The James–Stein estimator.* When estimating the price of apples in Oslo, the height of women in Bergen, and the unemployment rate in Trondheim, it is sometimes advantageous to use information about apples is Oslo and women in Bergen to say something about the unemployment rate in Trondheim. The point is that when estimating an ensemble of unrelated quantities, we can sometimes do better by borrowing information across unrelated things. This phenomenon is known as Stein's paradox or the Stein effect; (i) it might be clever to combine these independent estimators, gaining precision for each case; and (ii) even in cases where not much is gained, it will never be worse than not-combining-at-all. In the present exercise we'll look at Stein's 1956–1961 result, a result that initiated a whole field of statistical research known by the labels

of shrinkage, ensemble, small area estimation. So let $Y_i \sim \mathrm{N}(\theta_i, 1)$ be independent for $i = 1, \ldots, k$. We are to estimate $\theta_1, \ldots, \theta_k$ under the combined loss function

$$L(\theta, a) = \sum_{i=1}^{k} (a_i - \theta_i)^2.$$

We actually studied several aspects of this situation in Ex. 6.28, from Bayesian and empirical Bayesian perspectives, and succeeded in finding good ensemble procedures. Presently the focus is more on risk functions and uniform dominance.

(a). The standard approach is to use Y_i as an estimator of θ_i. This Y_i is the ML estimator, it is admissible, it is the UMVU estimator, etc. For this standard estimator $Y = (Y_1, \ldots, Y_k)$, show that its risk function is the constant k.

(b). For a single $Y \sim \mathrm{N}(\theta, 1)$, show using integration by parts that under very mild conditions on the function $b(y)$, one has $\mathrm{E}_\theta\,(Y - \theta)b(Y) = \mathrm{E}_\theta\, b'(Y)$m where b' is the derivative of b.

(c). Let now $b(y) = (b_1(y), \ldots, b_k(y))$. Generalise what you found in (b) to

$$\mathrm{E}_\theta\,(Y_i - \theta_i)b_i(Y) = \mathrm{E}_\theta\, b_{i,i}(Y), \quad \text{where } b_{i,i}(y) = \partial b_i(y)/\partial y_i.$$

(d). What you found in (b) and (c) is known as Stein's lemma. We shall use this to construct an estimator that uniformly dominates Y. Consider a general competitor to Y of the form $\delta(Y) = (\delta_1(Y), \ldots, \delta_k(Y))$, with

$$\delta_i(Y) = Y_i - b_i(Y) \quad \text{for } i = 1, \ldots, k.$$

Show that the difference in risk between Y and estimators of his form can be expressed as $R(\delta, \theta) - R(Y, \theta) = \mathrm{E}_\theta\, D(Y)$, where

$$D(y) = \sum_{i=1}^{k} \{b_i(y)^2 - 2b_{i,i}(y)\}.$$

The clever thing about $R(\delta, \theta) = k + \mathrm{E}_\theta\, D(Y)$ is that $D(y)$ does not depend on the unknown $\theta_1, \ldots, \theta_k$. We can therefore try to find a data dependent function $b(y)$ such that $D(y) < 0$ for all y, and in such a case identify an estimator that uniformly dominates the standard estimator. This turns out to be possible for $k \geq 3$, but not for $k \leq 2$.

(e). Try $b_i(y) = ay_i/\|y\|^2$, with $\|y\|^2$ being the squared Euclidian norm $\sum_{i=1}^{p} y_i^2$, leading to the estimator $\delta(y) = y - b(y) = (1 - a/\|y\|^2)y$. With this choice of $b(y)$, show that

$$D(y) = \{a^2 - 2a(k-2)\}/\|y\|^2.$$

Show that this is negative for a range of a values, provided $k \geq 3$. Demonstrate that the optimal a is $a = k - 2$, corresponding to the estimator

$$\delta_{\mathrm{JS}}(Y) = \{1 - (k-2)/\|Y\|^2\}\, Y. \tag{8.2}$$

This estimator is known as the James–Stein estimator. Show that its risk can be expressed as

$$R(\delta_{\mathrm{JS}}, \theta) = k - (k-2)^2 \mathrm{E}_\theta\, 1/\|Y\|^2.$$

In particular, explain that the classic estimator $(Y_1, \ldots, Y_k)$ for $(\theta_1, \ldots, \theta_k)$ is inadmissible, for any $k \geq 3$. Show that the greatest reduction in risk from using δ_{JS} instead of Y takes place when $\theta_1 = \cdots = \theta_k = 0$, and compute the risk $R(\delta_{\mathrm{JS}}, 0)$ in this point.

(f). (xx more; link to Ex. 6.28. xx) We'll now make a connection to empirical Bayes procedures. Start with a prior that takes $\theta_1, \dots, \theta_k$ independent from $\mathrm{N}(0, \tau^2)$. Show that the Bayes solution is $\delta^B = (\delta_1^B, \dots, \delta_k^B)$, with

$$\delta_i^B(Y) = \rho\, Y_i \text{ for } i = 1, \dots, k, \quad \text{where} \quad \rho = \tau^2/(\tau^2 + 1).$$

The empirical Bayes approach consists in estimating hyperparameters from data. Hyperparameters are those parameters set by the statistician in a pure Bayesian approach. Show that the marginal distribution of $y_1, \dots, y_k$ is a product of $\mathrm{N}(0, 1 + \tau^2)$ distributions. Find the ML estimator of α. Use the ML estimator to find an unbiased estimator, say $\tilde\alpha$, of α. The empirical Bayes estimator is then $\delta_{\mathrm{EB}}(Y) = \tilde\alpha Y$. What's noticeable about this estimator?

Ex. 8.21. *Resolving the paradox.* [xx make an exercise based on insights from Stigler (1990), perhaps?]

Ex. 8.22. *Poisson means and inadmissibility of ML estimator.* (xx nils emil check all details. xx) To show that an estimator is inadmissible it suffices to showcase one estimator that dominates it. Let $Y_1, \dots, Y_k$ be independent Poisson with means $\theta_1, \dots, \theta_k$. We are to estimate the $\theta = (\theta_1, \dots, \theta_k)$ with the weighted loss function $L(\theta, \delta) = \sum_{i=1}^k (\delta_i - \theta_i)^2/\theta_i$, where $\delta = (\delta_1, \dots, \delta_k)$. We demonstrate below that the usual estimator, with Y_i for θ_i, is minimax, but inadmissible; there are competitors doing uniformly better.

(a). Show first that the ML estimator in this problem is $\delta_{\mathrm{ml}} = (Y_1, \dots, Y_k)$, i.e. using Y_i for θ_i. Show that its risk function is the constant k. Via gamma independent priors $\theta_i \sim \mathrm{Gam}(1, b)$, find the Bayes estimator and the Bayes risk, and demonstrate that the ML estimator is minimax.

(b). Let $Z = \sum_{i=1}^k Y_i$ be the sum of these independent Poisson observations, write $\gamma = \sum_{i=1}^k \theta_i$ for the sum of the p Poisson means, and define $\pi_i = \theta_i/\gamma$ for $i = 1, \dots, k$. Show that $(Y_1, \dots, Y_k)$ given $Z = z$ is a multinomial $(z, \pi_1, \dots, \pi_k)$; this is valid for any $z \geq 1$. This establishes that $\mathrm{E}(Y_i \,|\, Z) = Z\pi_i$ and $\mathrm{Var}(Y_i \,|\, Z) = Z\pi_i(1 - \pi_i)$, for $i = 1, \dots, k$, with these formulae holding trivially also for $Z = 0$.

(c). Prove the following little lemma. If $X \sim \mathrm{Pois}(\theta)$ and g is a function such that $g(0) = 0$, then $\mathrm{E}_\theta\, g(X)/\theta = \mathrm{E}_\theta\, g(X + 1)/(X + 1)$.

(d). Consider the estimator $\delta^\star$ whose components are

$$\delta_j^\star = \{1 - \phi(Z)\}Y_j \quad \text{for } j = 1, \dots, k.$$

The game to be played now (as with the James–Stein estimator worked with in Ex. 8.20) is to find an expression for the risk difference $R(\theta, \delta^\star) - R(\theta, \delta_{\mathrm{ml}})$ that is independent of the unknown parameters. Using the results from (b) and (c) show indeed that the risk difference $R(\theta, \delta^\star) - R(\theta, \delta_{\mathrm{ml}})$ can be expressed as $\mathrm{E}_\gamma\, D(Z)$, with

$$D(z) = \{\phi(Z + 1)^2 - 2\phi(Z + 1)\}\{(Z + 1) + (k - 1)\} + 2\phi(Z)Z.$$

(e). Suppose that the function ϕ is such that $\phi(z)z$ is increasing. Under this assumption, find a function ϕ that ensures that $D(z) < 0$ for all $z \in \{0, 1, 2, \dots\}$. Conclude that the ML estimator is inadmissible. (xx check with $\phi(z) = c/(p - 1 + z)$, e.g. $c = 1$. we require $k \geq 2$. give risk expression. smallest for $\theta = 0$. x)

(f). We have shown that δ_{CZ} uniformly [xx nytt begrep xx] dominates the ML estimator, however, we do not yet now whether or not there exists and estimator that dominates δ_{CZ}. Show that δ_{CZ} is admissible.

Ex. 8.23. *Multiparameter inference, III.* (xx to be calibrated with clevenson-zidek Ch8. we do Bayes, empirical Bayes, full Bayes. xx) let $Y_i \sim \mathrm{Pois}(\theta_i)$ for $i = 1, \ldots, k$, aiming to estimate the ensemble of means, with loss $L = \sum_{i=1}^k (\widehat{\theta}_i - \theta_i)^2/\theta_i$. standard is Y_i, constant risk k. we do $\theta_i \sim \mathrm{Gam}(a, b)$, $\widehat{\theta}_i(a, b) = (y_i + a - 1)/(b + 1)$, i think. empirical Bayes yields almost Clevenson–Zidenk. then throw in full Bayes, with prior for (a, b) too.

(a). The default estimator would be $\widehat{\theta}_j = Y_j$, for $j = 1, \ldots, k$. Show that its risk function is constant, equal to k.

(b). Consider then the prior which takes $\theta_1, \ldots, \theta_k$ independent from a Gamma (a, b). Show that $\theta_j \,|\, \mathrm{data}$ is a $\mathrm{Gam}(a_j + y_j, b+1)$, find the Bayes estimator, and its risk function. (xx jotting down details here. xx)

$$\widehat{\theta}_j = \widehat{\theta}_j(a, b) = 1/(\mathrm{E}\,(1/\theta_j) \,|\, y_j) = \frac{a + y_j - 1}{b + 1}.$$

empirical Bayes: $\mathrm{E}\,(Y_j \,|\, \theta_j) = \theta_j$, $\mathrm{Var}\,(Y_j \,|\, \theta_j)) = \theta_j$, leading to $\mathrm{E}\, Y_j = a/b$, $\mathrm{Var}\, Y_j = a/b + a/b^2$; can then use moment matching to find $\widetilde{a}, \widetilde{b}$. this leads to empirical Bayes estimator $\theta_j^* = \widehat{\theta}_j(\widetilde{a}, \widetilde{b}) = (\widetilde{a} + y_j - 1)/(\widetilde{b} + 1)$. also:

$$\bar{f}(y_j, a, b) = \int_0^\infty \exp(-\theta)\theta^{y_j}/y_j! \frac{b^a}{\Gamma(a)} \theta^{a-1} \exp(-b\theta)\, \mathrm{d}\theta = \frac{b^a}{\Gamma(a)} \frac{\Gamma(a + y_j)}{(b + 1)^{a+y_j}} \frac{1}{y_j!}$$

Full Bayes, a prior for (a, b):

$$\pi_0(a, b) \prod_{j=1}^k g(\theta_j, a, b) f(y_j, \theta_j) = \pi_0(a, b) \prod_{j=1}^k g(\theta_j, a + y_j, b + 1) \bar{f}(y_j, a, b)$$

which leads to $\pi(a, b \,|\, \mathrm{data}) \propto \pi_0(a, b) \prod_{j=1}^k \bar{f}(y_j, a, b)$. then mcmc for sampling from $(a, b, \theta_1, \ldots, \theta_k)$ given data.

Ex. 8.24. *Multiparameter inference, IV.* Suppose one observes independent $Y_i \sim \mathrm{unif}[0, \theta_i]$ for $i = 1, \ldots, k$, aiming to estimate all parameters with sum of squares loss function $L(\theta, a) = \sum_{i=1}^k (a_i - \theta_i)^2$. Below use $\bar{\theta} = k^{-1} \sum_{i=1}^k \theta_i$, $z = \sum_{i=1}^k (\theta_i - \bar{\theta})^2 = s - k^2$, with $s = \sum_{i=1}^k \theta_i^2$.

(a). A direct start procedure is to use $\widetilde{\theta}_i = 2Y_i$ for the different components. Show that is is unbiased, with combined risk function $R(\theta, \widetilde{\theta}) = s/3 = (1/3)(z + k\bar{\theta}^2)$.

(b). Then consider independent priors $\theta_i \sim \mathrm{Gam}(2, b)$ for the components. Shos that the posterior distributions can be characterised as $\theta_i \,|\, \mathrm{data} \sim y_i + V_i$, where $V_i \sim \mathrm{Expo}(b)$, leading also to the Bayes estimator $\widehat{\theta}_i(b) = y_i + 1/b$. Find also the Bayes risk.

(c). Since specifying b might be hard, consider empirical Bayes estimators of the type $\theta_i^* = \widehat{\theta}_i(\widehat{b}) = y_i + \widehat{1}/b$, using data to estimate $1/b$. Explain that the mean of $2Y_i$ is $2/b$, in the marginal distribution for the data, and argue that setting $\bar{Y} = 1/\widehat{b}$ is a good scheme for estimating b. Explain that this leads to the empirical Bayes estimator $\theta_i^* = Y_i + \bar{Y}$.

(d). We then study the performance of the $\theta_i^* = Y_i + \bar{Y}$ estimator, in the spirit of the James–Stein and multi-Poisson methods worked with in Ex. 8.20 and 8.22. Show that the risk function becomes

$$\sum_{i=1}^{n} \mathrm{E}_\theta\,(Y_i + \bar{Y} - \theta_i)^2 = \sum_{i=1}^{k} [(\tfrac{1}{2}\theta_i + \tfrac{1}{2}\bar{\theta} - \theta_i)^2 + \theta_i^2/12 + (1/k)^2 s/12 + 2(\theta_i^2/12)/k]$$

and that this after some algebra may be expressed as $(1/4)z + (1/2)(1 + 3/k)(z + k\bar{\theta}^2)$. Show that this empirical Bayes method performs better than the direct one in the part of the parameter space where $w(\theta) = (1/k)z(\theta)/\bar{\theta}^2 \le 1 - 1/k$, the variance to the mean squared. Show that when the θ_i actually stem from such a $\mathrm{Gam}(2, b)$ prior, the $w(\theta)$ will be around $\frac{1}{2}$. Discuss this aspect of the methods, that the empirical Bayes scheme is expected to win over the direct $2Y_i$ method, but on this occasion not in all of the parameter space.

(e). Above we used the empicial Bayes scheme to motivate the estimator $\theta_i^* = Y_i + \bar{Y}$. Build also a full Bayesian prior, taking a background prior $\pi_0(b)$ and then $\theta_1, \ldots, \theta_k \sim \mathrm{Gam}(2, b)$. Find the posterior distribution of b and set up a simulation scheme to estimate the θ_i parameters. Compare its performance to the $Y_i + \bar{Y}$ method. (xx round off suitably. xx)

Hypothesis testing

Ex. 8.25. *Testing a simple hypothesis.* (xx make connection to Bayes Ch6 and Neyman-Pearson Ch4. xx) Let Y have density $f(y, \theta)$ and consider the simple hypothesis $H_0\colon \theta = \theta_0$ versus the simple alternative $\theta = \theta_1$. The statistical tests ϕ, with $\phi(x) = 1$ meaning reject H_0, and $\phi(x) = 0$ keep H_0, are to be evaluated under the loss function

$$L(\theta_0, \phi) = \begin{cases} 0, & \text{if } \phi(x) = 0, \\ K_1, & \text{if } \phi(x) = 1, \end{cases} \qquad L(\theta_1, \phi) = \begin{cases} K_2, & \text{if } \phi(x) = 0, \\ 0, & \text{if } \phi(x) = 1. \end{cases}$$

(a). Let $0 < \pi_0 < 1$ be your prior probability of H_0 being true. Derive an expression for the posterior expected loss, and show that the Bayes solution ϕ_π is of the likelihood ratio type

$$\phi_\pi(x) = \begin{cases} 1, & \text{if } f(y, \theta_1) > k_\pi f(y, \theta_0), \\ 0, & \text{if } f(y, \theta_1) < k_\pi f(y, \theta_0). \end{cases}$$

Find k_π and relate this quantity to the level of a test.

(b). Let now Y have the $\mathrm{N}(\theta, 1)$ distribution, for given θ. We want to test $H_0\colon \theta = 0$ versus $\theta_1 = \frac{1}{2}$ using the Bayes solution when the prior is $\pi_0 = \frac{1}{2}$. Find K_1 and K_2 such that $\mathrm{E}_{\theta_0}\,\phi_\pi(Y) = 0.05$.

(c). Show that any Bayesian test, with a prior giving weight to both the null- and the alternative hypothesis, is the most powerful test of its size. Hence the Bayesian method has somehow rediscovered the Neyman–Pearson Lemma of Ch. 4. *Hint*: Use what you know about Bayes solutions and admissibility.

Estimation in a few stylised prototype problems

Ex. 8.26. *Estimating the squared bias.* Assume $Y \sim \mathrm{N}(b,1)$, and that $\phi = b^2$ is to be estimated. There are links from the present prototype exercise about squared biases to the FIC theory of model selection in Ch. 11.

(a). Go through the relevant motivation for the following three estimators for the squared bias. (i) Since Y estimates b, use $\widehat{\phi}_1 = Y^2$. (ii) Explain that Y^2 has expected value b^2+1. This invites $\widehat{\phi}_2 = Y^2 - 1$. Show also that this is the only unbiased estimator for b^2. (iii) The trouble with the unbiased estimator for a nonnegative quantity is that it may be negative; what is the probability that this happens? Motivate the truncated estimator $\widehat{\phi}_3 = \max(Y^2 - 1, 0)$.

(b). For these three estimators, find expressions for their risk functions $r_j(b) = \mathrm{E}_b\,(\widehat{\phi}_j - b^2)^2$, and plot them in a diagram. Comment on what you find.

(c). Compute also risk functions numerically under absolute loss, i.e. using $\mathrm{E}_b\,|\widehat{\phi}_j - b^2|$, where the better behaviour of $\widehat{\phi}_3$ shows up more clearly.

Ex. 8.27. *Estimating a normalised distance, from a narrow to a wide model.* Consider a situation where $Y \sim \mathrm{N}(\theta,1)$ and where it is thoughtt not unreasonable that θ is not far from zero. One may then consider estimating θ via a smooth mixture of the simplest estimator, 0, and the natural wider model estimator, Y. There are links from this stylised exercise to FIC and model averaging themes of Ch. 11, just as for Ex. 8.26.

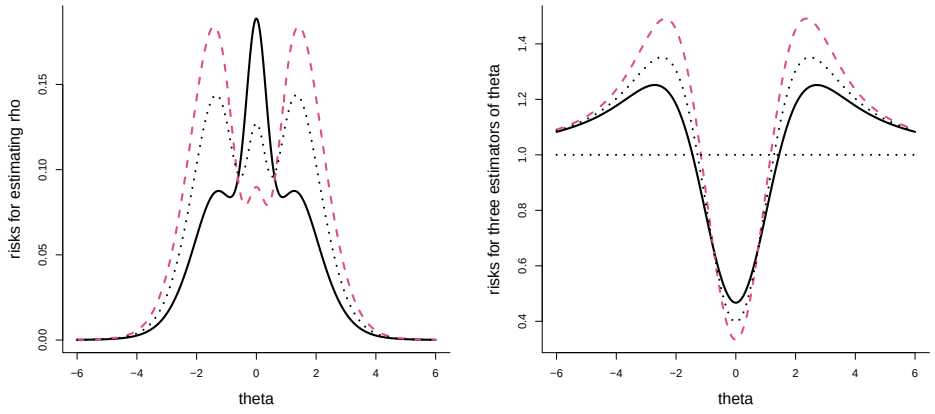


Figure 8.1: *Left panel: risk functions for the three estimators of* $\rho = \theta^2/(\theta^2+1)$ *(the first full, the second slanted, the third dotted). Right panel: risk functions for the three consequent estimators* $\widehat{\theta}_j = \widehat{\rho}_j Y$ *of* θ *(the first, the second, the rhird in increasing order of their maxima).*

(a). Explain that the estimator $\widehat{\theta}(a) = aY$ has mean squared error $(1-a)^2\theta^2 + a^2$. Show that this is minimised for $a_0 = \rho = \theta^2/(\theta^2+1)$, with minimum mse equal to this same ρ.

(b). There is hence a need to estimate ρ well. Since Y estimates θ, a direct estimator is $\widehat{\rho}_1 = Y^2/(Y^2+1)$. Since Y^2 has mean $1+\theta^2$, however, argue that $\widehat{\rho}_2 = \max(Y^2-1,0)/Y^2$

might also work well. For good measure we include one more, between the two first, namelig $\widehat{\rho}_3 = \max(Y^2 - \frac12, 0)/(Y^2 + c)$, with $c = \frac12$. Compute and display the risk functions for these three estimators, under squared error loss, as with Figure 8.1, left panel.

(c). Yet another estimator to be contemplated is the Bayesian solution, starting with a flat prior for θ, often used in situations with an unknown normal mean. Show that $\theta \,|\, y \sim \mathrm{N}(y, 1)$, and explain that the resulting estimator for ρ, under squared error, is

$$\widehat{\rho}_4 = \mathrm{E}\,(\rho \,|\, Y) = \int_0^\infty \frac{x}{x+1} g_1(x, Y^2)\,\mathrm{d}x,$$

the mean of $K/(K+1)$, where $K \sim \chi^2_1(Y^2)$.

(d). Coming back to the intended use of estimating ρ, consider the two estimators of θ in the mean estimation problem,

$$\widehat{\theta}_1 = \frac{Y^2}{Y^2+1} Y, \quad \widehat{\theta}_2 = \frac{\max(Y^2-1, 0)}{Y^2} Y, \quad \widehat{\theta}_3 = \frac{\max(Y^2 - \frac12, 0)}{Y^2 + \frac12} Y.$$

The second estimator may be written $(1 - 1/Y^2)Y$ for $|Y| \geq 1$ and zero for $|Y| \leq 1$; similarly the third estimator is $\{(Y^2 - \frac12)/(Y^2 + \frac12)\}Y$ for $Y^2 \geq \frac12$ and zero for $Y^2 \leq \frac12$. Compute and display the risk functions of these three estimators, under squared error loss, as with Figure 8.1, right panel.

(e). We learn in these two situations, estimating $\rho = \theta^2/(\theta^2+1)$ and then θ via $\widehat{\rho} Y$, that there are mo clear winners. Explain that the minimax principle may be applied, leading to $\widehat{\rho}_1 Y$ being the best for the second competition. Explain further that to find a clear ranking, one solution is to apply a Bayesian prior, say a suitable $\mathrm{N}(0, \kappa^2)$ for θ, here seen as a sensible weight function and not as a prior knowledge translated in to a prior, per se. With the standard normal as such a weight function, seeing values of θ outside $(-3, 3)$ as relatively unimportant, for example, compute the Bayes risks and rank the estimators. Find also the genuine Bayes solutions, for this prior, when estimating ρ and θ.

Ex. 8.28. *Estimating a bounded mean.* With $Y \sim \mathrm{N}(\theta, 1)$, the task is to estimate the mean under quadratic loss, when it is a priori known that $|\theta| \leq 1$.

(a). Show that the ML estimator is Y truncated to $[-1, 1]$. Compute and find a workable formula for its risk function

$$R(\theta) = \int_{-1}^1 (y-\theta)^2 \phi(y-\theta)\,\mathrm{d}y + \int_{-\infty}^{-1} (-1-\theta)^2 \phi(y-\theta)\,\mathrm{d}y + \int_1^\infty (1-\theta)^2 \phi(y-\theta)\,\mathrm{d}y.$$

(b). Then examine the Bayesian prior with θ taking values ± 1 with equal probabilities. What is the marginal density $\bar{f}(y)$ for Y? Show that $\Pr(\theta = 1 \,|\, y)$ is $p(y) = \exp(2y)/\{1 + \exp(2y)\}$, and that this leads to the Bayes estimator

$$\widehat{\theta}_B = \mathrm{E}\,(\theta \,|\, y) = \frac{\exp(2y) - 1}{\exp(2y) + 1} = \frac{\sinh(Y)}{\cosh(Y)}.$$

Plot first the Bayes and ML estimators, as functions of y, as in Figure 8.2, left panel; then compute and display both risk functions, as in the right panel.

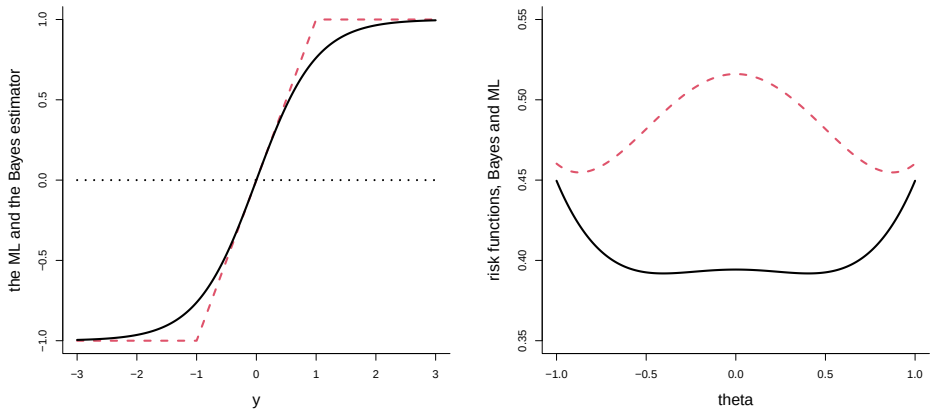


Figure 8.2: *Left panel: the Bayes (full curve) and the ML estimators (slanted), as functions of y. Right panel: their risk functions (with full curve for Bayes and slanted for the ML), with minimum Bayes risk 0.4496 for the two-point Bayes prior equal to the risk value at these two points.*

(c)**.** Find the Bayes risk, perhaps in two ways: (i) via the risk function $R_B(\theta)$, averaged over the prior; (ii) via MBR $= \mathrm{E}\,\mathrm{Var}\,(\theta \,|\, Y) = 4\,\mathrm{E}\,p(Y)\{1-p(Y)\}$. Show that $\widehat{\theta}_B$ is a minimax estimator.

(d)**.** (xx generalisations. could bother to do this. The computations above may be mimicked for the case where $\theta \in [-m, m]$ a priori, but we learn in the process that the minimax argument successfully used for $m = 1$ does not work, in cases where $m > m_0$, where $m_0 \doteq 1.06$ (and indeed the Bayes estimator is then not minimax). Start out taking the θ prior to be symmetrically concentrated in $\{-m, m\}$. Show that the posterior probability of $\theta = m$ is equal to $p(y) = \exp(2my)/\{1+\exp(2my)\}$, and that the posterior mean is $\mathrm{E}\,(\theta \,|\, y) = m\tanh(my)$. Show that the minimum Bayes risk is equal to $\mathrm{E}\,4m^2\,p(Y)\{1-m(Y)\}$ and compute this, for values of m equal to say $0.90, 1.00, 1.10$. The point is now that this minimum Bayes risk is equal to the maximum risk for $m = 0.90, 1.00$, but not for $m = 1.10$. For values $m > m_0$, the risk function starts having three modes, at $-m, 0, m$, with the one at zero becoming bigger than at the endpoints.

Ex. 8.29. *Estimating a binormal mean with nonnegativity constraints.* (xx to be done and polished. xx) Let (X, Y) be binormal with mean (a, b), unit variances, and known correlation ρ. The task is to estimate (a, b), with the natural quadratic loss function

$$\begin{pmatrix} \widehat{a}-a \\ \widehat{b}-b \end{pmatrix}^{\mathrm{t}} \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}^{-1} \begin{pmatrix} \widehat{a}-a \\ \widehat{b}-b \end{pmatrix} = \frac{1}{1-\rho^2}\{(\widehat{a}-a)^2 + (\widehat{b}-b)^2 - 2\rho(\widehat{a}-a)(\widehat{b}-b)\},$$

when it is known a priori that $a \geq 0$, $b \geq 0$. We omit the $1/(1-\rho^2)$ constant in the loss and risk function questions below.

(a)**.** A simple direct method is to truncate X and Y to zero, if negative, i.e. using $\widetilde{a} = \max(X, 0)$ and $\widetilde{b} = \max(Y, 0)$. Find formulae for $\mathrm{E}_a\,\widetilde{a}$, $\mathrm{Var}_a\,\widetilde{a}$, with of course similar results for $\widetilde{b}$. Find the risk function

$$R_0(a, b) = \mathrm{E}_{a,b}\,\{(\widetilde{a}-a)^2 + (\widetilde{b}-b)^2 - 2\rho(\widetilde{a}-a)(\widetilde{b}-b)\}.$$

(b). Then consider the ML estimators $(\widehat{a}, \widehat{b})$, maximising the log-likelihood under the given constraints. Show that this, for a given observed (x, y), is equivalent to minimising

$$Q(a, b) = (a - x)^2 + (b - y)^2 - 2\rho(a - x)(b - y) \quad \text{with } a \geq 0,\ b \geq 0.$$

(c). Though $(\widehat{a}, \widehat{b})$ may be found via numerical routines, e.g. the 'nlm in box' algorithm `nlminb` in `R`, it is helpful to find explicit formulae for the ML. Going through the four situations, for the four quadrants into which (x, y) might fall, show this: (i) For $x > 0, y > 0$, we have $(\widehat{a}, \widehat{b}) = (x, y)$; (ii) for $x > 0, y < 0$, we have $\widehat{b} = 0$ and $\widehat{a} = \max(x - \rho y, 0)$; (iii) for $x < 0, y > 0$, we have $\widehat{a} = 0$ and $\widehat{b} = \max(y - \rho x, 0)$; (iv) for $x < 0, y < 0$, we have $\widehat{a} = \max(x - \rho y, 0)$ and $\widehat{b} = \max(y - \rho x, 0)$.

(d). There is now a well-defined risk function for the ML method,

$$R(a, b) = \mathrm{E}_{a,b}\,\{(\widehat{a} - a)^2 + (\widehat{b} - b)^2 - 2\rho(\widehat{a} - a)(\widehat{b} - b)\},$$

though with no easy closed-form formula. Show that for $a \geq 2.50, b \geq 2.50$, say, the chance is so high for (X, Y) to land in the positive quadrant that both R_0 and R come close to $2(1 - \rho^2)$.

(e). Carry out some numerical computation and comparison of the two risk functions, via simulation of binormal (X, Y), for given value of ρ, at grid positions (a, b). (xx more. ML is often better, sometimes significantly so, but not uniformly. application to regression situation with constraints on coefficients. xx)

Estimation of parametric densities

Ex. 8.30. *Unbiased estimation of a parametric density.* Suppose $Y_1, \ldots, Y_n$ are i.i.d. from a parametric density $f(y, \theta)$, like the normal or the Gamma or the Beta. How can we construct an unbiased estimator of the density function itself? For this exercise, assume there is a sufficient and complete statistic, say $T = (T_1, \ldots, T_p)$.

(a). Let $\widehat{\theta}$ be the ML estimator. Explain why it needs to be a function of $T_1, \ldots, T_p$, and that the ML estimator of the density itself is $f(y, \widehat{\theta})$.

(b). For a given window A (make it an interval, if the distribution is one-dimensional), a very simple estimator for $p(\theta) = \Pr(Y \in A) = \int_A f(y, \theta)\,\mathrm{d}y$ is $\widehat{p} = I(Y_1 \in A)$, using a single data point. Show that it is unbiased. This also invites the somewhat more sensible estimator $\bar{p} = n^{-1} \sum_{i=1}^n I(Y_i \in A)$, the binomial proportion of data points inside the A window. Show that it is unbiased and find a formula for its variance.

(c). Typically this estimator can be beaten, however. Consider indeed

$$p^* = \mathrm{E}\,(\widehat{p}\,|\,T) = \Pr(Y_1 \in A\,|\,T).$$

Explain why this is actually an estimator, i.e. that it does not depend on the parameter θ, that it is unbiased, and indeed the UMVU estimator for $p(\theta)$. Show furthermore that the construction $\mathrm{E}\,(\bar{p}\,|\,T)$ leads to the very same p^*. In particular, there is a well-defined UMVU estimator for the c.d.f. $F(y, \theta)$, with this recipe being valid for any parametric distribution (again, provided there is a sufficient and complete statistic).

(d). Assume the situation is regular enough for the density $f_n(y\,|\,T)$ of a Y_i given T to exist. Explain why it does not depend on the parameter, and that

$$p^* = \int_A f_n(y\,|\,T)\,\mathrm{d}y, \quad \text{for all windows } A.$$

Show that $f_n(y\,|\,T)$ is unbiased, and the UMVU for $f(y,\theta)$.

(e). Let $Y_1,\ldots,Y_n$ be i.i.d. from the exponential density $\theta\exp(-\theta y)$. Show that $T = \sum_{i=1}^n Y_i$ is complete and sufficient, and that $(Y_1/T,\ldots,Y_n/T)$ has a Dirichlet $(1,\ldots,1)$ distribution, with each Y_i/T having a Beta distribution. Show from this that the UMVU estimator for the density takes the form

$$f_n(y\,|\,T) = \mathrm{be}(y/T,1,n-1)/T = (n-1)(1-y/T)^{n-2} \quad \text{for } y\in[0,T].$$

Why does the fact that the estimator is zero, outside the $[0,T]$ window, not contradict its unbiasedness? Show also that $f_n(y\,|\,T)$ comes close to the ML estimator $f(y,\widehat{\theta})$, for growing n. Give an example where the ML estimator is unbiased, but where $f(y,\widehat{\theta})$ is not unbiased for $f(y,\theta)$.

(f). Above we managed to find a formula $f_n(y\,|\,T)$ for the unbiased estimator of the exponential density, via the fact that Y_1/T has a Beta distribution $(1,n-1)$. Suppose, for the sake of the argument, that you did not quite have the time to do the last part of the mathematics, deriving that formula. Explain and show via an illustration in your computer that the following would work. For the given dataset (perhaps with a small n, to see more interesting shapes), compute T, and draw a million realisations of Y_1' given that values of T. Then form a very fine histogram of these Y_1'. That fine histogram is a numerical approximation to the $f_n(y\,|\,T)$ you need. This argument is valid for any such parametric density estimation situation; as long as you can simulate Y_1' given the sufficient and complete statistic, you have your unbiased density estimator.

(g). (xx may make a separate short exercise for the discrete case, umvu estimation of $f(y,\theta)$ in two cases. xx) (i) Suppose $Y_1,\ldots,Y_n$ are i.i.d. Pois(θ). With $T=\sum_{i=1}^n Y_i$, show that $Y_1\,|\,T \sim \mathrm{binom}(T,1/n)$, leading to the UMVU estimator $f_n(y\,|\,T) = \binom{T}{y}(1/n)^y(1-1/n)^{T-y}$. Show that this is close to $\exp(-\bar{Y})(\bar{Y})^y/y!$ for n growing. (ii) Consider $Y_1,\ldots,Y_a$ being geometric with $(1-p)^y p = q^y p$ for $y=0,1,\ldots$. Show that $Z=\sum_{i=1}^a Y_i$ is sufficient and complete. Use formulae from Ex. 1.28 to work out that the UMVU estimator of the point probability $q^y p$ is

$$f_a(y\,|\,Z) = \binom{Z-y+a-2}{a-2} \Big/ \binom{Z+a-1}{a-1} = \frac{(Z-y+a-2)!}{(Z-y)!}\,\frac{Z!\,(a-1)!}{(Z+a-1)!}$$

for $y=0,1,\ldots,Z$. (xx note: not hypergeometric?; close to $(1-\widehat{p})^y\widehat{p}$, with $\widehat{p}=1/(1+\bar{y})$. xx)

Ex. 8.31. *UMVU estimation of the uniform density and cumulative.* In Ex. 8.30 we learned how to construct UMVU estimators for the cumulative and densities in parametric families, as long as these is a sufficient and complete statistic. This works also for the uniform case, but with certain extra componenents and nuances. The uniform density on $[0,\theta]$ is $1/\theta$ for $y\le\theta$ and zero for $y>\theta$, and there is a subtle difference between estimating this density and estimating the parameter $1/\theta$.

(a). With $Y_1, \ldots, Y_n$ i.i.d. from this uniform on $[0, \theta]$, show that $M_n = \max_{i \le n} Y_i$ is sufficient and complete. Find the UMVU estimator for θ and for $1/\theta$.

(b). Show that the distribution of Y_1 given M_n is a mixture, equal to M_n with probability $1/n$ and uniform on $[0, M_n]$ with probability $1 - 1/n$. Explain that the c.d.f. is $F_n(y \,|\, M_n) = (1 - 1/n)(y/M_n)$ for $y < M_n$ and then 1 for $y \ge M_n$. Why is this the UMVU for the parametric c.d.f. $F(y, \theta)$?

(c). Then set up a clear formula for the UMVU estimator of the density, and verify that it is indeed unbiased.

(d). Suppose i.i.d. pairs $(X_1, Y_1), \ldots, (X_n, Y_n)$ are observed from the distribution uniform on the interior of the circle $\{(x, y) \colon x^2 + y^2 \le \theta^2\}$. Writing $R_i = \|(X_i, Y_i)\|$, show that the ML estimator is $\widehat{\theta} = M_n = \max_{i \le n} R_i$, find its distribution, and show that it is sufficient and complete. Then find the distribution of $(X_1, Y_1) \,|\, M_n$, and from this the UMVU estimators of the density and cumulative of R_i.

Ex. 8.32. *UMVU estimation of normal densities.* Suppose $Y_1, \ldots, Y_n$ are i.i.d. $\mathrm{N}(\xi, \sigma^2)$. How can we estimate the underlying density $f(y) = \sigma^{-1}\phi(\sigma^{-1}(y - \xi))$ unbiasedly? For this we use the general results of 8.30. (xx the following to be polished. xx)

(a). For a simpler separate question first, suppose $Y_1, \ldots, Y_n$ are i.i.d. $\mathrm{N}(\xi, 1)$. Here the sample average $\bar{Y}$ is sufficient and and complete, so the UMVU estimator is $f_n(y)$, the density of $Y_1 \,|\, \bar{Y}$. Writing $Y_1 = \bar{Y} + (Y_1 - \bar{Y})$, show that $Y_i - \bar{Y}$ is independent of $\bar{Y}$. Explain that $f_n(y) = g_n(y - \bar{y})$, with g_n the density of $Y_1 - \bar{Y}$. Show next that g_n is zero-mean normal with variance $1 - 1/n$, leading to

$$f_n(y) = \Big(\frac{n}{n-1}\Big)^{1/2} \phi\Big(\Big(\frac{n}{n-1}\Big)^{1/2}(y - \bar{y})\Big).$$

It is of course close to the simpler ML estimator based density $\phi(y - \bar{y})$.

(b). Then consider a dataset Y_i from $\mathrm{N}(\xi, \sigma^2)$, with sufficient and complete estimators $\bar{Y}$ and $\widehat{\sigma}^2 = \sum_{i=1}^n (Y_i - \bar{Y})^2/(n-1)$. Write $Y_i = \bar{Y} + \widehat{\sigma} W_i$, with $W_i = (Y_i - \bar{Y})/\widehat{\sigma}$. Show via the Basu lemma of Ex. 4.25 that W_i is independent of $(\bar{Y}, \widehat{\sigma})$. Show from this that the UMVU estimator must be $f_n(y) = \widehat{\sigma}^{-1} h_n(\widehat{\sigma}^{-1}(y - \bar{y}))$, with $h_n(z)$ the density of W_1. We find a formula for this density below. Explain nevertheless how that the numerical approximation trick via simulations, from Ex. 8.30(f), can be made to work here, without having the actual density formula.

(c). As an introductory exercise, before returning to the residuals, suppose $U_1, \ldots, U_m$ are i.i.d. $\mathrm{N}(0, 1)$, and consider $V_i = U_i/(U_1^2 + \cdots + U_m^2)^{1/2}$. Working first with their squares, show that V_i^2 can be written $\chi_1^2/(\chi_1^2 + \chi_{m-1}^2)$, which is a Beta; more broadly, show that $(V_1^2, \ldots, V_m^2) \sim \mathrm{Dir}(\frac{1}{2}, \ldots, \frac{1}{2})$. Show next that the density of V_1 becomes

$$g_m(v) = \frac{\Gamma(\frac{1}{2}m)}{\Gamma(\frac{1}{2})\Gamma(\frac{1}{2}(m-1))}(1 - v^2)^{(m-1)/2-1} \quad \text{on } (-1, 1).$$

Also relate what you find here to the following: with $U_1, \ldots, U_m$ i.i.d. $\mathrm{N}(0, 1)$, the distribution of $(U_1, \ldots, U_m)$ given its length $\|U\| = r$ is a uniform on the sphere of radius r.

(d). Write $G = \sum_{i=1}^n (Y_i - \bar{Y})^2$ and $V_i = (Y_i - \bar{Y})/G^{1/2}$. From $\sum_{i=1}^n V_i^2 = 1$ note that $\sum_{i=1}^n W_i^2 \le m$, with $m = n-1$ the degrees of freedom for $G \sim \chi^2_m$. Show that the distribution of V_i does not depend on the parameters, so for simplicity set (ξ, σ) to $(0,1)$. With row vectors $a_1 = (1/\sqrt{n}, \ldots, 1/\sqrt{n})$ and $a_2 = \{n/(n-1)\}^{1/2}(1 - 1/n, -1/n, \ldots, -1/n)$, show that these are orthogonal and of length 1. Hence there is an orthonormal matrix A, with a_1, a_2 the two first rows. Transforming to $Z = AX$, with $Z_1, \ldots, Z_n$ i.i.d. N$(0,1)$ by Ex. 1.49, show that $Z_1 = \sqrt{n}\bar{Y}$ and $Z_2 = \{n/(n-1)\}^{1/2}(Y_1 - \bar{Y})$, and that this leads to

$$V_1 = \{(n-1)/n\}^{1/2} Z_2/(Z_2^2 + \cdots + Y_n^2)^{1/2}.$$

Derive a formula for V_1 based on this, using results from the previous point. Conclude that the UMVU estimator for the normal density is

$$f_n(y) = \frac{1}{\widehat{\sigma}}\Big(\frac{n}{n-1}\Big)^{1/2} \frac{\Gamma(\frac{1}{2}(n-1))}{\Gamma(\frac{1}{2})\Gamma(\frac{1}{2}(n-2))}\Big\{1 - \frac{n}{n-1}\Big(\frac{y-\bar{y}^2}{\widehat{\sigma}}\Big)\Big\}^{(n-2)/2-1}$$

(e). (xx can attempt multinormal density, would be nice. xx) $Y_i = \bar{Y} + \widehat{\Sigma}^{1/2} W_i$, with $W_i = \widehat{\Sigma}^{-1/2}(Y_i - \bar{Y})$. independent of the sufficient statistics, has a density $h_n(w)$. which implies

$$f_n(y) = |\widehat{\Sigma}|^{-1/2} h_n(\widehat{\Sigma}^{-1/2}(y - \bar{Y})).$$

would like to see this more clearly: let $U_1, \ldots, U_m$ be i.i.d. N$_p(0, I)$, with $A_m = \sum_{i=1}^m U_i U_i^{\mathrm{t}}$. Then we should understand the distribution of $V_i = U_i^{\mathrm{t}} A_m^{-1} U_i$, with $\sum_{i=1}^m V_i = p$. the distribution of $A_m^{-1} U_1 U_1^{\mathrm{t}}$ is a multivariate Beta type, related to Roy's largest root test, with 1 and $m-1$ degrees of freedom, but careful with notation and terms. also keen on the root of V_i, i.e. $W_i = A_m^{-1/2} U_i$, with $W_i^{\mathrm{t}} W_i = V_i$. it should have a density with a formula.

Notes and pointers

[xx some notes and pointers here xx]

Stein: See Stein (1956); James and Stein (1961) for the original articles, and, for example Efron and Morris (1977) and Stigler (1990) for lucid presentations. Clevenson and Zidek (1975), Stoltenberg and Hjort (2020), more.

(xx Dan Roy Annals. Lehmann, Blyth. Bhattacharyya bounds, generalising Cramer-Rao. BNP chapter for minimax for $\theta = \mathrm{E}_F Y_i$, and admissibility. minimax for F. xx)

(xx there is a literature on estimating bounded means. bickel, casella, more. xx)

I.9

Brownian motion and empirical processes

[xx Here's something on Brownian motion, Brownian bridges, partial-sum processes, empirical processes, with applications. We also include some on Poisson and Gamma type processes. The C and D spaces. Applications to construction of good overall tests, for constancy of distributions against changepoint scenarios, log-likelihood bridges, etc. xx]

Time has come for us to consider *stochastic processes*, as opposed to 'only' random variables and vectors. One important result, motivating the process view, tended to and applied several times in this chapter, is as follows. For $Y_1, Y_2, \ldots$ being i.i.d., with mean ξ and variance σ^2, form the partial sum process $X_n(t) = (1/\sqrt{n}) \sum_{i \le [nt]} (Y_i - \xi)$, for $t \in [0,1]$. The process starts at zero and has the value $\sqrt{n}(\bar{Y}_n - \xi)$ at $t = 1$, which we know tends to $\mathrm{N}(0, \sigma^2)$ by the CLT of Ch. 2. Then *Donsker's theorem* says that the full empirical process X_n tends in distribution to scaled Brownian motion $X = \sigma W$, in a well-defined space of functions $[0,1] \to \mathbb{R}$. The $\{W(t) \colon t \in [0,1]\}$ process in question has the basic property that increments over disjoint intervals are independent and normal, with $W(t) - W(s) \sim \mathrm{N}(0, t-s)$.

Establishing such a $X_n \to_d \sigma W$ result, and many others, takes some work, from clear definitions and initial preparations to lemmas, criteria for securing other conditions, and theorems. Processes considered live in function spaces, as opposed to Euclidean vector spaces. Below we study $C[0,1]$, the space of all continuous functions on $[0,1]$, and then the bigger space $D[0,1]$ of all $x \colon [0,1] \to \mathbb{R}$ which are continuous to the right and with left-hand limits, i.e. $x(t-) = \lim_{\varepsilon \to 0} x(t-\varepsilon)$. The natural metric for the $C[0,1]$ space is the uniform one, $d(x,y) = \max_t |x(t) - y(t)|$, and the more complicated Skorokhod metric $d^*(x,y)$ for $D[0,1]$ is worked with in (xx pointer xx). Having equipped the function spaces with metrics, the general convergence in distribution definition is now that

$$X_n \to_d X \quad \text{if } \mathrm{E}\, g(X_n) \to \mathrm{E}\, g(x) \text{ for all continuous bounded } g. \tag{9.1}$$

This is one of several equivalent criteria for convergence in distribution in finite-dimensional spaces, as we have seen in the Portmanteau theorem, see Ex. 2.18. In exercises (xx pointer xx) we work out different more practical criteria for process convergence.

Process results of the $X_n \to_d X$ type are arguably deeper than convergence results in Euclidean spaces, since they imply convergene and hence useful approximations also for all smooth enough functionals. If one needs the limit distribution for say the range $R_n = \max_t X_n(t) - \min_t X_n(t)$, this might be easier to work with via convergence for the full process, as opposed to attempting to work directly with R_n.

(xx Applications abound. xx)

(xx then more intro text. function spaces $C[0,1]$ and $D[0,1]$. tightness and finite-dimensional convergence secures $X_n \to_d X$. many applications. Poisson process. Gamma process. xx)

(xx nils will push this, with some real reorganisation and start prose. i consider moving basic compactness characterisations and tightness conditions closer to the front. 90 percent of our readers will read and understand and use and construct and prove and use $X_n \to_d X$ things, via these lemmas and techniques for tightness, but not care about the details of the proofs behind these. also put in qq plots, and the Durbin things with $\sqrt{n}\{F_n(t) - F(t,\widehat{\theta})\}$. xx)

(xx put in an exercise for independent increment processes that convergence of Laplace is enough for finite-dimensional convergence, $L_n(u,t) = \mathrm{E}\exp\{-uZ_n(t)\} \to L(u,t)$. this would then simplify a few details for Gamma, Beta, Poisson. xx)

[xx There will be a one-page intro section to the chapter, telling the readers about the themes to be worked with. we decide a bit later whether we also will include some 3-4 exercises about the Poisson process, seen as a limit, etc., and e.g. converge to Gamma and extended Gamma processes, in this chapter. xx]

The space C, Brownian motion, tightness, partial sums, Donsker

the C[0,1] space of functions

Ex. 9.1. *The function space* $C[0,1]$. Many of the more important stochastic processes live in the space $C[0,1]$, the set of all continuous functions $x\colon [0,1] \to R$. There are also various situations where a random process $X_n = \{X_n(t)\colon t \in [0,1]\}$ may have perhaps small discontinuities, but where there is a clear limit process $X = \{X(t)\colon t \in [0,1]\}$ lies in $C[0,1]$ with probability one. We need a machinery for studying and handling the elements of $C[0,1]$.

(a). Show that each $x \in C[0,1]$ is bounded. Define the uniform metric as the distance $d(x,y) = \max_t |x(t) - y(t)|$. Show that $x_n \to x$ if and only if $x_n(t)$ tends uniformly to $x(t)$. Demonstrate also the triangle inequality, $d(x,z) \le d(x,y) + d(y,z)$.

(b). Go through the mechanics of the following simple example, where the point is that x_n tends to some x_0 pointwise, but not uniformly. Let x_0 be the zero function, and let $x_n(t)$ be zero on $[2/n,1]$, and with a 'tent', from 0 to 1 on $[0,1/n]$ and then from 1 to 0 on $[1/n,2/n]$. Find $d(x_n,x_0)$. Make some useful variations.

(c). Plot your favourite function du jour $x_0\colon [0,1] \to R$, along with its 0.03-neighbourhood, $\{x\colon d(x_0,x) \le 0.03\}$.

(d). One sometimes works with other intervals than $[0,1]$, and the basic definitions and properties go through also for say $C[a,b]$, with $0 < a < b$. Show that the function $x_n(t) = n(t^{1/n} - 1)$ converges to the function $x(t) = \log t$ in the space $C[0.01,1]$. Show also that the function $x_n(t) = t^n$ converges to the zero function in the space $C[0,0.99]$.

(e).

Ex. 9.2. *Continuous functions on* $C[0,1]$. As mentioned in the introduction, $X_n \to_d X$ for processes X_n and X in $C[0,1]$ is formally defined via (9.1). We shall indeed derive more practical criteria in exercises to come, but we may already put this somewhat abstract definition to use.

(a). Show that if $X_n \to_d X$, then $h(X_n) \to_d h(X)$, for each continuous function h on $C[0,1]$. Show also that it is sufficient that h is continuous on a set A where $\Pr(X \in A) = 1$. [xx cross-reference to something in Ch. 2. xx]

(b). For the functions $h_1(x) = \max_t x(t)$, $h_2(x) = \min_t x(t)$, $h_3(x) = \max_t |x(t)|$, $h_4(x) = \int_0^1 x(t)\,\mathrm{d}t$, show that they are continuous.

(c). To anticipate a typical use of such results, suppose $X_n \to_d W$, with W being the Brownian motion process introduced in Ex. 9.3. Show that the range $R_n = \max_t X_n(t) - \min_t X(t)$ tends in distribution to the range $R = \max_t W(t) - \min_t W(t)$ of the Brownian motion. It is also to be noted that if somewhat different X_n processes have the same W limit, then also the different R_n will have the same R limit. Similar remarks apply to $\max_t |X_n(t)| \to_d \max_t |W(t)|$.

(d). Show that the projection function $\pi_t(x) = x(t)$ is continuous, for each t. More generally, show that the projection function which takes x to its values $x(t_1), \ldots, x(t_k)$, for any finite set $t_1 < \cdots < t_k$, is continuous. Thus finite-dimensional convergence,

$$(X_n(t_1), \ldots, X_n(t_k)) \to_d (X(t_1), \ldots, X(t_k)) \text{ for each finite set } t_1 < \cdots < t_k,$$

is necessary for process convergence (but not sufficient, as we shall see below).

(e). Consider the function $h\colon C[0,1] \to C[0,1]$ which transforms $x(t)$ to $x(t) - tx(1)$. Show that this h is continuous, and generalise.

(f). Explain that the integral operation $h(x) = \int_0^1 x(t)\,\mathrm{d}t$ is continuous. More generally, with g any continuous function on the real line, show that $h(x) = \int_0^1 g(x(t))\,\mathrm{d}t$ is continuous.

(g). Let $h(x) = \int_0^1 I(x(t) > 0)\,\mathrm{d}t$ be the total time $x(t)$ has spent above zero. Show that h is 'almost continuous' [xx this to be made precise xx].

Ex. 9.3. *Brownian motion.* Here we introduce the Brownian motion process, also called the Wiener process. They are zero-mean Gaussian processes with independent increments over disjoint intervals. We shall see in Ex. 9.7 that this process can be reached as a limit in distribution of random functions in $C[0,1]$, and by completeness is then also inside, i.e. its paths, though highly wiggly, are continuous.

the Brownian motion process

(a). Consider then $W = \{W(t)\colon t \geq 0\}$, a random process evolving over time t, with the following properties: (i) $W(0) = 0$; (ii) increments $W(t) - W(s)$ have normal $\mathrm{N}(0, t-s)$ distributions, i.e. with variance equal to the length of the interval in question; (iii) increments over disjoint intervals are independent. Show that $\mathrm{cov}\{W(s), W(t)\} = \min(s,t)$, and for the random triple $(W(s), W(t), W(u))$, with $s < t < u$, find the 3×3 covariance matrix.

(b). It is incidentally fruitful and useful to think of the W process via cumulative small movements, as in $W(t)$ being the sum of many tiny and independent $\mathrm{d}W(s) = W(s + \mathrm{d}s) - W(s) \sim \mathrm{N}(0, \mathrm{d}s)$. Express $W(1)$ as the sum of 1000 tiny such $W((i+1)/n) - W(i/n)$, and show that you get the right variance.

(c). Show that with $t_1 < \cdots < t_k$, then by necessity the vector $(W(t_1), \ldots, W(t_k))$ must be multinormal, and give its $k \times k$ covariance matrix.

(d). Define $I = \int_0^1 W(t)\,\mathrm{d}t$, a random integral. Find its distribution, perhaps via Riemann approximations.

(e). Suppose a fellow reader of this book tries to define a process $W^* = \{W^*(t)\colon t \geq 0\}$ by putting up (i), (ii), (iii) above, but now with $\mathrm{Var}\,\{W^*(t)-W^*(s)\} = (t-s)^{1/2}$, rather than $t-s$. Show that this would backfire solidly. – So it's not enough to put up 'something' for the variance or covariance function, as it may lead astray; one needs to check for logical coherency, and each covariance matrix for finite-dimensional distributions need to be nonnegative definite. Show that with $\mathrm{Var}\,\{W^*(t) - W^*(s)\} = (t-s)^\alpha$, only $\alpha = 1$ secures such coherence.

(f). Note that *the existence* of the Brownian motion process is not entirely obvious, and there is no easy way to put down a full joint probability density for the full thing. The pure existence of a stochastic process is hence a more delicate and complex matter than when working in one- and finite-dimensional cases. Brownian motion indeed exists, however, which may be proved in many ways, including via the Donsker theorem below. Argue in fact that if there is a well-defined sequence $X_n = \{X_n(t)\colon t \geq 0\}$ of elements in $C[0,1]$, which converges to a limit, then that limit X must be inside $C[0,1]$. In Ex. 9.7 below we establish Brownian motion in this fashion, verifying also that this limit behaves according to (i), (ii), (iii) given above.

Ex. 9.4. *Modulus of continuity, compactness: the Arzelà–Ascoli theorem.* In a fair range of important instances, our random functions $X_n = \{X_n(t)\colon t \in [0,1]\}$ and their limits live in $C[0,1]$. To establish clear results we need the machinery of tightness, as per Ex. 2.24 and related Ch. 2 exercises. A key component is the Arzelà–Ascoli theorem, which we prove here, characterising compactness. The somewhat abstract definition is that a subset K is compact if each covering $K \subset \cup_e U_e$ of open sets has a finite subcover, i.e. $K \subset \cup_{j=1}^m U_j$ with these U_j belonging to the set of U_e. We start defining the modulus of continuity as

$$w(x,\delta) = \sup_{|t-s|<\delta} |x(t) - x(s)|,$$

for any function x and positive δ. For any susbset K, let K^* be its closure, i.e. the smallest closed set containing K. The statement is that a subset K has a compact closure, with the uniform metric, if and only if (i) $\{x(0)\colon x \in K\}$ is bounded; (ii) $\sup_{x\in K} w(x,\delta) \to 0$ as $\delta \to 0$. (xx point perhaps to (Billingsley, 1968, p. 220). xx)

(a). For an example of non-compactness, let x_n be the tent function which grows linearly from 0 to 1 on $[0,1/n]$, then back from 1 to 0 linearly on $[1/n,2/n]$, and which is zero on $[2/n,1]$. Show that $\omega(x_n,1/n) = 1$, and that the set $K = \{x_1,x_2,\ldots\}$ cannot be compact.

(b). Explain that x having small $w(x,\delta)$ means a bound of how much it can oscillate on small intervals. Show that for any x, we have $w(x,\delta) \to 0$ as $x \to 0$. Demonstrate also that $|w(x,\delta) - w(y,\delta)| \leq 2d(x,y)$, and use this to show that $w(x,\delta)$ is continuous in x, for each δ.

(c). Show that if K^* is compact, then $\{x(t_0)\colon x \in K\}$ is bounded, for each t_0. Show also that $A_K(\delta) = \sup_{x\in K} w(x,\delta)$ must go to zero with $\delta \to 0$.

(d). (xx then the converse, that (i) and (ii) imply compactness for K^*. technical but doable; worse for $D[0,1]$ to come below. xx)

Ex. 9.5. *Tightness and convergence.* For $X_n \to_d X$ in $C[0,1]$, finite-dimensional convergence is necessary, but not sufficient. An example illustrating this is the case looked at in Ex. 9.1, where $x_n \to x_0$ pointwise, but not in the $d(x_0, x_n)$ metric.

(a). With the uniform $d(x,y)$ metric, show that $C[0,1]$ is separable: there is a countable subset S such that inside every open ball $B(x,\varepsilon) = \{y\colon d(x,y) < \varepsilon\}$, there are some functions from S. Show also that $C[0,1]$ is complete: if x_n are all in $C[0,1]$, and have a limit x, then x is also inside $C[0,1]$; this is the statement that if x_n are all continuous, and converge uniformly to x, then x is also continous.

(b). We say that the sequence $\{X_n : n \geq 1\}$ sequence is *tight* provided there for each $\varepsilon > 0$ is a compact subset K such that $\Pr(X_n \in K) \geq 1 - \varepsilon$ for all large n. Suppose now there is a random element X in $C[0,1]$ such that (i) there is finite-dimensional convergence, for X_n to X, and (ii) the X_n sequence is tight. Show that $X_n \to_d X$. (xx pointer to Ex. 2.24 and more. xx)

(c). Use the Arzelà–Ascoli theorem of Ex. 9.4 to show the X_n sequence is tight provided (i) the sequence $X_n(0)$ in $\mathbb{R}$ is tight (bounded in probability); (ii) for each positive ε, $\limsup \Pr(w(X_n,\delta) \geq \varepsilon) \to 0$ as $\delta \to 0$.

(d). So we need to control $\sup_{|t-s|<\delta} |X(t) - X_(s)|$, for various processes, and attempt to do this via subintervals. Show that with $0 = t_0 < t_1 < \cdots < t_{k-1} < t_k = 1$, with $t_i - t_{i-1} \geq \delta$, then

$$\Pr(w(X,\delta) \geq 3\varepsilon) \leq \sum_{i=1}^{k} \Pr\Big(\sup_{t_{i-1} \leq t \leq t_i} |X(t) - X(t_{i-1})| \geq \varepsilon \Big).$$

In particular, for $\delta = 1/k$,

$$\Pr(w(X,1/k) \geq 3\varepsilon) \leq \sum_{i=1}^{k} \Pr\Big(\sup_{(i-1)/k \leq t \leq i/k} |X(t) - X((i-1)/k)| \geq \varepsilon \Big).$$

Clearly, bounding such probabilities, for a sequence of X_n, depends in various ways on the probabilisitic particularities of the sequence.

Ex. 9.6. *Tail inequalities for maxima of partial sums.* We learn from Ex. 9.5 that criteria for tightness must involve bounds for tail probabilities for random variables of the type $\sup_{t_1 \leq t \leq t_2} |X_n(t) - X_n(t_1)|$, and these again relate to partial sums of smaller increments $X_n(s+\varepsilon) - X_n(s)$. To set the stage for several useful inequalities of this type, start out with i.i.d. variables U_i, with mean zero and variance one, and their random walk partial sums $S_j = U_1 + \cdots + U_j$. Define $A_m = \max_{j \leq m} |S_j|/\sqrt{m}$, and put $p_m(a) = \Pr(A_m \geq a)$.

(a). Show that $p_m(a) \leq 1/a^2$; this is a special case of the so-called Kolmogorov inequality. Note that this also generalises the Chebyshov inequality. (xx may give a hint. may use martingale type arguments. generalisation to different variances: $p_m(a) \leq (1/a^2)(1/m)\sum_{j=1}^{m} \sigma_j^2$. xx)

(b). The following leads to what in various cases is a stronger result. First verify that the event $\{A_m \geq a\}$ can be decomposed into $\cup_{i=1}^{m} E_i$, with $E_i = \{\max_{j<i} |S_j|/\sqrt{m} < a \leq |S_i|/\sqrt{m}\}$. For $c \in (0,a)$, show then that

$$p_m(a) \leq \Pr(|S_m|/\sqrt{m} \geq a - c) + \sum_{i=1}^{m-1} \Pr(E_i, |S_m|/\sqrt{m} < a - c),$$

with the sum further bounded by

$$\sum_{i=1}^{m-1} \Pr(E_i, |S_m - S_i|/\sqrt{m} \geq c) \leq \Pr(|S_m|/\sqrt{m} \geq a - c) + \sum_{i=1}^{m-1} \Pr(E_i) \frac{m-i}{mc^2}$$
$$\leq \Pr(|S_m|/\sqrt{m} \geq a - c) + (1/c^2) p_m(a).$$

Show that this leads to $(1 - 1/c^2)p_m(a) \leq \Pr(|S_m|/\sqrt{m} \geq a - c)$. For $c = \sqrt{2}$, you find $p_m(a) \leq 2\Pr(|S_m|/\sqrt{m} \geq a - \sqrt{2})$, which we shall use below.

(c). The probability of $|S_m|/\sqrt{m} \geq a - \sqrt{2}$ can be bounded separately, either by Markov type inequalities, if e.g. the third moment is finite, or via the CLT fact that $\Pr(|S_m|/\sqrt{m} \geq a - \sqrt{2}) \to \Pr(\mathrm{N}(0,1) \geq a - \sqrt{2})$. This is less than e.g. $3/(a - \sqrt{2})^4$ for any $a > \sqrt{2}$. Show hence that $p_m(a) \leq 2 \cdot 3.01/(a - \sqrt{2})^4$ for all big enough m, even when we do not assume more than the finite variance.

(d). (xx we include useful variations, with different variances σ_j^2 and phrase results in terms of $s_m^2 = \sigma_1^2 + \cdots + \sigma_m^2$. the arguments are essentially the same as for the i.i.d. case. xx)

Ex. 9.7. *Partial-sum processes and the Donsker theorem.* Let $U_1, U_2, \ldots$ be i.i.d., with mean zero and variance one (and we do not need to say anything more regarding their distribution). Then we know from the CLT that $n^{-1/2}(U_1 + \cdots + U_n) \to_d \mathrm{N}(0,1)$. But we also see that e.g. $n^{-1/2}(U_1 + \cdots + U_{[n/2]})$ must have a normal limit, and similarly with other partial sums. This exercise goes through the basics for the full partial-sum process, and leads to the fundamental and very useful Donsker theorem. Define indeed the random process $X_n = \{X_n(t) \colon t \geq 0\}$ via

$$X_n(t) = n^{-1/2} \sum_{i \leq [nt]} U_i = (U_1 + \cdots + U_{[nt]})/\sqrt{n} \quad \text{at values } t = i/n,$$

with linear interpolation for $t \in [i/n, (i+1)/n]$ to keep X_n inside $C[0,1]$. We shall indeed prove *Donsker's theorem*, that $X_n \to_d W$, with W the Brownian motion process introduced in Ex. 9.3. The linear interpolation is essentially merely cosmetic, to keep X_n inside $C[0,1]$. With $X_n^*(t)$ more naturally defined as $n^{-1/2} \sum_{i \leq [nt]} U_i$ for all t, without linear interpolation but with jumps of size $O_{\mathrm{pr}}(1/\sqrt{n})$, we shall see in Ex. 9.14 that we still have $X_n^* \to_d W$, but then in the bigger natural space $D[0,1]$.

the Donsker theorem

(a). Use one or two of your favourite distributions for U_i, with mean zero and variance one, to generate some sample paths of X_n. We've done this for Figure 9.1, with $n = 50$ and $n = 10^4$, actually using $U_i = V_i - 1$, with the V_i being unit exponential; the intended point is to use something with a skewness, though that aspect disappears in the limit.

(b). Fix t positive. Show that $X_n(t)$ has mean zero and variance $[nt]/n$ converging to t. Show also that $X_n(t) \to_d \mathrm{N}(0,t)$. In view of the Brownian motion process introduced in Ex. 9.3, we may write $X_n(t) \to_d W(t)$.

(c). Show that if $s < t < u < v$, then $(X_n(t) - X_n(s), X_n(v) - X_n(u))$ tends to $(W(t) - W(s), W(v) - W(u))$, two independent normal pieces with variances $t - s$ and $v - u$. More generally, show that there is finite-dimensional convergence in distribution; for $t_1 < \cdots < t_k$, we have $(X_n(t_1), \ldots, X_n(t_k)) \to_d (W(t_1), \ldots, W(t_k))$.

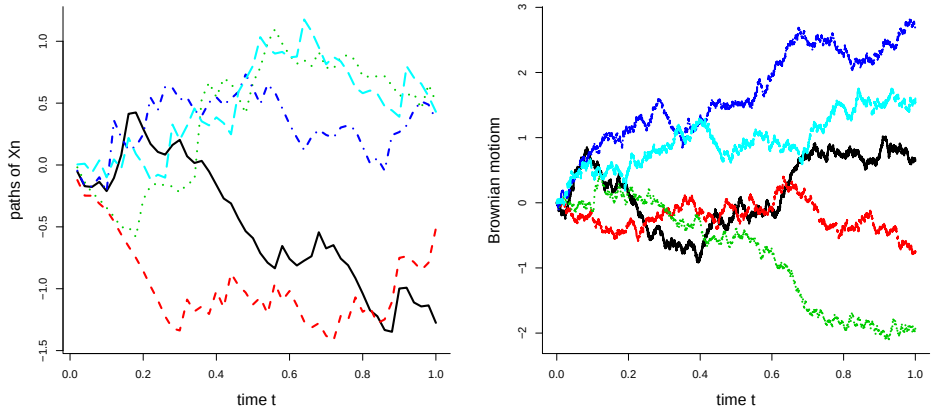


Figure 9.1: *Simulated partial-sum processes X_n, as defined in Ex. 9.7, with $n = 50$ (left panel) and $n = 10^4$ (right panel), with i.i.d. components $U_i = V_i - 1$ and $V_i \sim \text{Expo}(1)$. For large n the eye can barely see the difference between these paths and real Brownian motion.*

(d). It remains to demonstrate that the X_n sequence is tight, where we shall use the characterisations and oscillation bounds from Ex. 9.5, along with probability tail inequalities from Ex. 9.6. Show that $\Pr(w(X_n, 1/k) \geq 3\varepsilon) \leq k \Pr(B_k \geq \varepsilon)$, where $B_k = \sup_{t \leq 1/k} |X_n(t)|$. We have $X_n(t) = S_j/\sqrt{n}$ for these $t = j/n \in [0, 1/k]$, up to $t_0 = j_0/n = 1/k$, so indexes at work are those with $j/n \leq 1/k$, or $j \leq n/k$. Show hence that $B_k = \max_{j \leq n/k} |S_j|/\sqrt{n} = A_{[n/k]}/k^{1/2}$, in terms of $A_{[n/k]} = \max_{j \leq n/k} |S_j|/([n/k])^{1/2}$, worked with in Ex. 9.6; the integer value modification $[n/k]$ is there merely to avoid saying that n/k needs to be an integer. With notation from that exercise, deduce that

$$k \Pr(B_k \geq \varepsilon) = k p_{[n/k]}(k^{1/2}\varepsilon) \leq 2k \Pr(|S_{[n/k]}|/([n/k])^{1/2} \geq k^{1/2}\varepsilon - \sqrt{2}),$$

Show that we for any givem $\eta > 0$ have $k \Pr(B_k \geq \varepsilon) < \eta$ for all large enough k, and that this secures tightness, hence concluding the proof for the Donsker theorem.

(e). Explain that the $X_n \to_d W$ result also implies the existence of the Brownian motion, as a random element inside $C[0, 1]$, i.e. with continuous paths.

Ex. 9.8. *Applying the Donsker theorem.* In Ex. 9.7 we have established that the natural normalised partial-sum process X_n tends to Brownian motion W, and we may read off corollaries via the continuous functionals pointed to in Ex. 9.2, and yet others.

(a). A very simple continuous function is $h(x) = x(1)$, reading off the value of $x(t)$ at the endpoint $t = 1$. It is continuous. Hence $X_n(1) \to_d X(1)$. Go back to Donsker and deduce the CLT. In this light, the Donsker theorem is really a much bigger sibling of the CLT; the Donsker has a thousand corollaries, some simple, many complex, and one of these is the CLT.

(b). Consider $h(x) = \max_{0 \leq t \leq 1} x(t)$. Deduce that with $M_n = \max_{i \leq n} S_i$, where $S_i = U_1 + \cdots + U_i$, we have $M_n/\sqrt{n} = \max_{0 \leq t \leq 1} X_n(t) \to_d M = \max_{0 \leq t \leq 1} W(t)$. Note that this limit ensues regardless of the distribution of the i.i.d. sequence U_i, as long as these

have mean zero and variance one. The limit is the same, with U_i being standard normal, or symmetric ± 1 variables, of $U_i = V_i - 1$ with unit exponential V_i, etc.

(c). Often enough the limit $h(W)$ of $h(X_n)$ may be complicated, not admitting a clear explicit representation, in which case one may resort to simulations. In this particular case, with $M_m/\sqrt{n} \to_d M = \max_{0\le t\le 1} W(t)$, an explicit solution can be found, however. Choose $U_i = \pm 1$ with probabilities $\frac12, \frac12$, so that the partial sums $S_i = U_1 + \cdots + U_i$ form a random walk, with each of the 2^n possible paths for S_n having the same probability $1/2^n$. Argue first that

$$\Pr(M_n \ge a) = \Pr(M_n \ge a, S_n > a) + \Pr(M_n \ge a, S_n < a) + \Pr(M_n \ge a, S_n = a),$$

and then show that the two first terms must be equal, whereas the third is small. Deduce from this that $\Pr(M_n/\sqrt{n} \ge u) = 2\Pr(S_n > \sqrt{n}u) + \varepsilon_n$, with $\varepsilon_n \to 0$, with resulting limit $\Pr(M \ge u) = 2\Pr(\mathrm{N}(0,1) \ge u) = 2\{1 - \Phi(u)\}$ for $u \ge 0$. Verify that this also means that $M \sim |N(0,1)|$.

(d). Alice and Bob play a long game. Alice's winning in game i is U_i, with a distribution having mean zer0 and variance one; if it's positive, good for Alice, it it's negative, good for Bob. After i games, Alice's account has $U_1 + \cdots + U_i$, and Bob's account has minus this sum. They play a few thousand times. – How much of the time has Alice been in the lead? This is food for the Donsker theorem, since this concerns

$$T_n = (1/n)\sum_{i=1}^n I\{U_1 + \cdots + U_i > 0\} = m\{t \in [0,1] \colon X_n(t) > 0\},$$

the Lebesgue measure of how much time the process has been above zero. This is a continuous functional, so $T_n \to_d T = m\{t \in [0,1] \colon W(t) > 0\}$. The problem hence has a clear-cut general solution: Alice has been in the lead, over Bob, a portion T_n of the time, and T_n tends to T in distribution. So it's 'only' a matter of finding the distribution of T. This is non-trivial, but the solution is the Beta$(\frac12, \frac12)$, intriguingly, with U-shaped density $(1/\pi)/\{t(1-t)\}^{1/2}$ on $(0,1)$. Thus there's a high chance that *one of the two* has been leading for a very long time; the least likely outcome is that they've been leading about half the time.

Ex. 9.9. *The Brownian bridge.* We say that a process $X = \{X(t) \colon t \in [0,1]\}$ is normal, or Gaußian, if all its finite-dimensional distributions are normal. This is the same as saying that all linear combinations are normal. Thus the Brownian motion process is normal, for example.

(a). To define and describe a normal process X, show that it is sufficient to give (i) the mean function $\xi(t) = \mathrm{E}\, X(t)$ and (ii) the covariance function $K(s,t) = \mathrm{cov}\{X(s), X(t)\}$. So nothing more is required than these two functions.

(b). Describe the Brownian motion process via its mean and covariance functions.

(c). Then consider the process $W^0(t) = W(t) - tW(1)$ for $t \in [0,1]$, with W being the Brownian motion. Show that W^0 is normal, with zero mean, and covariance function $\min(s,t) - st$, i.e. $s(1-t)$ for $s \le t$. In particular, its variance is $t(1-t)$. The W^0 is called the Brownian bridge.

the Brownian bridge

(d). For two disjoint time windows, say $[s_1, s_2]$ and $[t_1, t_2]$, with $s_1 < s_2 < t_1 < t_2$, find the covariance and correlation between the two bridge increments $W^0(s_2) - W^0(s_1)$ and $W^0(t_2) - W^0(t_1)$.

(e). Show that W^0 also can be characterised as W conditional on $W(1) = 0$.

Ex. 9.10. *Some functionals of Brownian motion and Brownian bridge.* We have constructed and properly defined the Brownian motion W and Brownian bridge W^0, random continuous functions on $[0, 1]$. Various functionals, say $h_1(W)$ and $h_2(W^0)$, continuous functions $C[0, 1] \to \mathbb{R}$, are of sufficient interest to warrant examination of their distributions. For a given such $h(W)$, say, we may either attempt to work out its distribution directly, of find it via the Donsker theorem and continuity, since $h(X_n) \to_d h(W)$ when $X_n \to_d W$; we have seen some such examples in Ex. 9.8. It is also typically possible and not complicated to simulate say 10^5 paths of W (or very close approximations thereof), and read off the distribution of $h(W)$ in that way.

(a). Consider $I = h(W) = \int_0^1 W(t)\,\mathrm{d}t$. Show that $I_n \to_d I$, where $I_n = n^{-1}\sum_{i=1}^n W(i/n)$. Explain that I_n must be a normal $(0, \sigma_n^2)$; find the variance in question, and its limit. Conclude that $I \sim \mathrm{N}(0, 1/3)$. Show similarly that $\int_0^1 W^0(t)\,\mathrm{d}t$ has the $\mathrm{N}(0, 1/12)$ distribution.

(b). Find mean and variance for the random integrals $A_1 = \int_0^1 W(t)^2\,\mathrm{d}t$, $A_2 = \int_0^1 W^0(t)^2\,\mathrm{d}t$, $A_3 = \int_0^1 W^0(t)^2/\{t(1-t)\}\,\mathrm{d}t$, $A_4 = \int_0^1 \exp(W(t))\,\mathrm{d}t$.

(c). With e.g. `sim = 100`, `N = 1000`, explain that the following simple `R` code produces `sim` simulated paths of Brownian motion $W(t)$ and Brownian bridge $W^0(t)$, at the fine grid of all points j/N:

```
tval = seq(1/N,1,by=1/N)
Wsim = 0*tval %*% t(1:sim)
W0sim = 0*tval %*% t(1:sim)
for (ss in 1:sim)
{
Wsim[ ,ss] = cumsum(rnorm(N))/sqrt(N)
W0sim[ ,ss] = Wsim[ ,ss] - tval*Wsim[N-1,ss]
}
matplot(tval,W0sim,type="l")
```

This may now be followed by computing and plotting the required $h(W)$ or $h(W^0)$. With `Wmax = 0*(1:sim)`, `W0absmax = 0*(1:sim)`, and similar, verify that

```
for (ss in 1:sim)
{ Wmax[ss] = max(Wsim[, ss])
W0absmax[ss] = max(abs(W0sim[ ,ss])) }
```

produces the intended simulated realisations of these functions. Carry out this for a suitably high number of `sim`, perhaps also with a high `N`, to simulate the distribution of $W^0_{\max} = \max_{0\le t\le 1} |W^0(t)|$. Record the quantiles at levels 0.50, 0.90, 0.95, 0.975, 0.99; your answers should be close to 0.828, 1.224, 1.358, 1.480, 1.627.

(d). We have $W^0(s) \sim \mathrm{N}(0, s(1-s))$, so the variable $W^0(s)/\{s(1-s)\}^{1/2}$ is standard normal. Now consider

$$D = \max_{\varepsilon \le s \le 1-\varepsilon} |W^0(s)|/\{s(1-s)\}^{1/2},$$

for a fixed $\varepsilon \in (0, \frac12)$. Various naturally built test statistics have limit distributions related to this D, so it might be necessary to work out say the 0.90, 0.95, 0.99 quantiles of its distribution. There are no explicit formulae for this, so the apparently simplest method is to simulate a large number of W^0 paths (or close approximations to such) and then read off A values. Explain how this can be done. Find the 0.90, 0.95, 0.99 quantiles for intervals $[0.025, 0.975]$, $[0.05, 0.95]$, $[0.10, 0.90]$, displaying also the A distributions via fine histograms. (xx Answers, needed below, are xx.)

```
               0.50   0.90   0.95   0.99
eps = 0.025:  2.137  2.927  3.184  3.692
eps = 0.05:   2.051  2.858  3.121  3.649
eps = 0.10:   1.931  2.768  3.041  3.577
```

(e). (xx one more example, showing simulations in action. can do $A = \{t \in [0,1] \colon W(t) > 0\}$, which has a Beta$(\frac12, \frac12)$ distribution, the random amount of time Alice leads over Bob. xx)

Ex. 9.11. *Simulating a Gaussian process.* If $Z = \{Z(t) \colon t \in [0,1]\}$ is a process for which each finite-dimensional $(Z(t_1), \ldots, Z(t_k))^{\mathrm{t}}$ is multinormal, i.e. for any k and any $t_1 < \cdots < t_k$, we say that Z is a *Gaussian process*, or a normal process. Instances include of course the Brownian motion and the Brownian bridge, worked with above. There are also schemes for simulating these, as in Ex. 9.10, exploting the structure of the process. We shall also need general recipes for this.

(a). Explain first that any process Z in $C[0,1]$ is characterised by its finite-dimensional distributions. Then argue that a Gaussian process Z is characterised by its mean function $\xi(t) = \mathrm{E}\, Z(t)$ and covariance function $K(s,t) = \mathrm{cov}\{Z(s), Z(t)\}$.

(b). Which functions $K(s,t)$ are bona fide covariance functions? First, it needs to be symmetric, with $K(s,t) = K(t,s)$. We say that such a K is *nonnegative definite* provided each matrix finite-dimensional matrix $\Sigma = (K(t_i,t_j))_{i,j=1,\ldots,k}$ is nonnegative definite, i.e. $\sum_{i,j} a_i a_j K(t_i,t_j) \geq 0$ for all $a = (a_1, \ldots, a_k)$. There is then a Gaussian process with this covariance function. (xx sometning with continuity. xx)

(c). How can we simulate many paths, from a given Gaussian process Z? Take for simplicity Z to have mean function zero. If Z is known to be the limit of some Z_n that can be simulated, we may use a high n and then simulate. A version of this scheme is to start with a grid of $t_i = i/n$, for such a high n. Consider Z_n, the process with values $Z_n(t_i) = Z(t_i)$, completed by using linear interpolation. Explain that $Z_n \to_d Z$. Hence it is sufficient to find a scheme to simulate such Z_n.

(d). For a given and perhaps high n, and writing $t_i = i/n$, form the $n \times n$ covariance matrix Σ with elements $K(t_i,t_j)$. A recipe for simiulating from such a $\mathrm{N}_n(0,\Sigma)$ has actually been given in Ex. 1.46, working in principle in any dimension, as for the intended high n here. In essence, use $\Sigma^{1/2}N$, with $N = (N_1, \ldots, N_n)^{\mathrm{t}}$ an i.i.d. string of standard normals.

this can be accomplished even with 1000×1000 matrices

(e). To illustrate this, consider the Brownian bridge again, with covariance function $\min(s,t) - st$. Use e.g. $n = 100$ in your computer, form the $n \times n$ matrix, carry out the spectral decomposition as above; then simulate and display say 1000 paths.

Ex. 9.12. *Extending and scaling the time axis.* Above we have worked with Brownian motion as a process W defined as a process $W(t)$ for $t \in [0,1]$. It may however easily be defined also for longer stretches of time, directly, or via scaling.

(a). A full Brownian motion process $W = \{W(t) : t \geq 0\}$, on any time interval, is defined via the basic properties laid out in Ex. 9.3: it starts at zero, and its increments are independent, with $W(t_2) - W(t_1) \sim \mathrm{N}(0, t_2 - t_1)$. If W is such a process, and $c > 0$, show that the process $W^*(t) = W(ct)/c^{1/2}$, defined for $t \in [0, 1]$, is a standard Brownian motion in $C[0, 1]$.

(b). Let $M_t = \max_{0 \leq s \leq t} W(s)$ be the maximum value obtained during the time interval $[0, t]$. Show that M_t must have the same distribution as $t^{1/2} M_1$. In particular, $\mathrm{E}\, M_t = \gamma_0 t^{1/2}$ and $\mathrm{Var}\, M_t = \gamma_1 t$, where γ_0 and γ_1 are the mean and variance of M_1. Here we actually know the distribution; check with Ex. 9.8 to see that $M_1 \sim |\mathrm{N}(0, 1)|$, and deduce that $\max_{a \leq t \leq b} W(t) \sim (b - a)^{1/2} |\mathrm{N}(0, 1)|$.

(c). With W a Brownian motion process, define $W^*(t) = tW(1/t)$ for $t > 0$. Show that its variance is t, which also allows the definition $W^*(0) = 0$. Show that W^* is another Brownian motion. It somehow has the ability to time-transform itself.

(d). (xx might not need this, unless used for something more. xx) Consider the Donsker theorem setup of Ex. 9.7, with an i.i.d. sequence $U_1, U_2, \ldots$ used to form the partial-sum process $X_n(t) = n^{-1/2} \sum_{i \leq [nt]} U_i$. The standard formulation, which we have seen and used, is that $X_n \to_d W$ in $D[0, 1]$. Explain how this may be extended, from details of definitions to details of proof, to reach a wider theorem: $X_n = \{X_n(t) \colon t \in [0, c]\}$ tends in distribution to $W = \{W(t) \colon t \in [0, c]\}$, in each of the function spaces $D[0, c]$.

The space D, independent increments, Gamma and Poisson processes

the D[0,1] space of functions

Ex. 9.13. *The function space* $D[0, 1]$. A long list of crucial processes have discontinuities, for which we need a bigger space to work with than the $C[0, 1]$. Let indeed $D[0, 1]$ consist of all functions $x \colon [0, 1] \to R$ which are right-continuous with left-hand limits.

(a) (xx examples. the metric. xx).

(b). (xx separable and complete. xx)

(c). (xx compact sets. tightness. the $w'(x, \delta)$ and $w''(x, \delta)$, following Billingsley (1968, Ch. 4). xx)

Ex. 9.14. *Tightness in* $D[0, 1]$. (xx Tightness in $D[0, 1]$. establishing criteria, in several steps and details and tail inequalities; will then be used extensively. xx)

(a). (xx finish and polish this. xx) Consider a sequence X_n of processes in $D[0, 1]$. Assume there are nondecreasing functions G_n convgering to a continuous increasing function G on $[0, 1]$, such that, for some $\gamma > 0$ and $\alpha > 1$,

$$\mathrm{E}\, |X_n(t) - X_n(s)|^\gamma \, |X_n(u) - X_n(t)|^\gamma \leq \{G_n(u) - G_n(s)\}^\alpha$$

holds for all $s < t < u$. Then X_n is tight.

(b). We reached Donsker's theorem in Ex. 9.7, for the process $X_n^*(t) = n^{-1/2} \sum_{i \leq [nt]} U_i$, but there cosmetically repaired via linear interpolation in each $[(i-1)/n, i/n]$ cell to secure $X_n^* \in C[0, 1]$. Show now, without this linear interpolation, that $X_n^* \to_d W$, in $D[0, 1]$; each X_n^* has jumps, of size $O_{\mathrm{pr}}(1/\sqrt{n})$, and the limit W is in the subspace $C[0, 1]$. Show in fact, working with $X_n(t) - X_n(s) = n^{-1/2} \sum_{[ns] < i \leq [nt]} U_i$, that

$$\begin{aligned} \mathrm{E}\, |X_n(t) - X_n(s)|^2 \, |X_n(u) - X_n(t)|^2 &\leq \{G_n(t) - G_n(s)\}\{G_n(u) - G_n(t)\} \\ &\leq \{G_n(u) - G_n(s)\}^2, \end{aligned}$$

where $G_n(t) = [nt]/n$, converging to $G(t) = t$.

Ex. 9.15. *Independent increment processes.* Consider processes X_n in $D[0,1]$ with independent increments; $X_n(u) - X_n(t)$ is independent of $X_n(t) - X_n(s)$ are independent, for any $s < t < u$, and similarly for any collection of disjoint intervals.

(a). Demonstrating tightness for independent increment processes becomes easier than for the general cases. Suppose there are nondecreasing functions G_n, tending to a continuous increasing G on $[0,1]$, such that for some $\gamma < 0$

$$\mathrm{E}\,|X_n(t) - X_n(s)|^\gamma \le \{G_n(t) - G_n(s)\}^\alpha \quad \text{for all } s < t,$$

this time for some $\alpha > \frac{1}{2}$. Show that the X_n sequence is tight.

(b). Consider once more the Donsker theorem setup of Ex. 9.7, with $X_n(t) = n^{-1/2} \sum_{i \le [nt]} U_i$, now defined without linear interpolation. Demonstrate tightness of X_n, again, verifying the simpler criterion $\mathrm{E}\,|X_n(t) - X_n(s)|^2 = ([nt] - [ns])/n$.

(c).

Ex. 9.16. *Gamma processes.* Construct a random process as follows. With $V_1, \ldots, V_m$ i.i.d. $\mathrm{Gam}(a/m, 1)$, let $Z_m(t) = \sum_{i \le [mt]} V_i$.

(a). Show that $Z_m(i/m) \sim \mathrm{Gam}(ai/m, 1)$, and that $Z_m(t) - Z_m(s) \to_d \mathrm{Gam}(a(t - s), 1)$ for each interval $[s,t]$. Generalise this to a statement with all finite-dimensional distributions. Then demonstate that the Z_m sequence is tight. Argue thus that $Z_m \to_d Z$, in $D[0,1]$, with a well-defined Gamma process limit, having independent Gamma distributed increments.

(b). The paths of such Gamma processes are increasing but not continuous; one may indeed show that there are infinitely many random jumps on each interval. To work with the size of the biggest jump, consider $X_\varepsilon \sim \mathrm{Gam}(\varepsilon, 1)$, with ε small. Writing Γ_ε for its c.d.f., i.e. $\Gamma_\varepsilon(v) = G(v, \varepsilon, 1)$ in terms of the Gamma c.d.f. with parameters $(\varepsilon, 1)$, show that

$$\Gamma_\varepsilon(v) = 1 - \int_v^\infty \frac{1}{\Gamma(\varepsilon)} x^{\varepsilon - 1} \exp(-x)\,\mathrm{d}x = 1 - \varepsilon \int_v^\infty \frac{1}{\Gamma(1+\varepsilon)} x^\varepsilon x^{-1} \exp(-x)\,\mathrm{d}x,$$

and that this may be written $1 - \varepsilon E_1(v)\{1 + O(\varepsilon)\}$, in terms of the exponential integral function $E_1(v) = \int_v^\infty x^{-1} \exp(-x)\,\mathrm{d}x$.

(c). Let now $J_m(t)$ be the biggest jump for Z_m, in the construction above, over the window $[0,t]$. Show that $\Pr(J_m(t) \le v) = \prod_{i/m \le t} \Gamma_{a/m}(v)$, and that this converges to $\exp\{-aE_1(v)\}$ for $v > 0$. Show that the median $v_{0.50}$ in this biggest jump distribution is the solution to $E_1(v_{0.50}) = (\log 2)/a$.

(d). We now generalise the above to Gamma processes with non-constant rate. First, if $V \sim \mathrm{Gam}(A, b)$, show that its Laplace transform becomes

$$\mathrm{E}\,\exp(-uV) = b^A/(b+u)^A = \exp\{-A \log(1 + u/b)\}.$$

Consider $A(t) = \int_0^t a(s)\,\mathrm{d}s$, for some nonnegative function $a(s)$, and define $Z_m(t) = \sum_{i \le [mt]} V_i$, with independent small Gamma contributions $V_i \sim \mathrm{Gam}(a(i/m)(1/m), b)$.

Show first that $\mathrm{E}\, Z_m(t) \to A(t)/b$ and $\mathrm{Var}\, Z_m(t) \to A(t)/b^2$, and then that the Laplace transform converges,

$$\mathrm{E}\, \exp\{-uZ_m(t)\} = \prod_{i/m \le t} \exp\{-a(i/m)(1/m)\log(1+u/b)\} \to \exp\{-A(t)\log(1+u/b)\}.$$

Show as for the simpler case of $\alpha(s) = 1$ and $b = 1$ above that there is finite-dimensional convergence for Z_m, that is tight, implying convergence $Z_m \to_d Z$ to a Gamma process with independent increments $Z(t) - Z(s) \sim \mathrm{Gam}(A(t) - A(s), b)$.

(e). For this more general Gamma process, work with the biggest jump $J_m(t)$ of Z_m, over the window $[0, t]$, and find its limit distribution, which is then associated with the biggest jump of the $\mathrm{Gam}(A(\cdot), b)$ process.

Ex. 9.17. *Extended Gamma processes.* The construction used in Ex. 9.16 led to well-defined Gamma processes, of the type $Z(t) \sim \mathrm{Gam}(A(t), b)$, with $A(t) = \int_0^t a(s)\,\mathrm{d}s$. This may be extended further to the case of non-constant shape parameter b. (xx could be pointer to BNP. xx)

(a). Start with independent $V_i \sim \mathrm{Gam}(a(i/m)(1/m), b(i/m)$, for continuous positive functions $a(s), b(s)$, and define $Z_m(t) = \sum_{i/m \le t} V_i$ for $t \ge 0$. Show that the mean and variance converge,

$$\begin{aligned} \mathrm{E}\, Z_m(t) &= \sum_{i/m \le t} \frac{a(i/m)(1/m)}{b(i/m)} \to \int_0^t \frac{a(s)}{b(s)}\,\mathrm{d}s, \\ \mathrm{Var}\, Z_m(t) &= \sum_{i/m \le t} \frac{a(i/m)(1/m)}{b(i/m)^2} \to \int_0^t \frac{a(s)}{b(s)^2}\,\mathrm{d}s. \end{aligned}$$

(b). Then show that there is convergence of the Laplace transform,

$$\mathrm{E}\, \exp\{-uZ_m(t)\} = \prod_{i/m \le t} \mathrm{E}\, \exp(-uV_i) \to \exp\Big\{-\int_0^t a(s)\log(1+u/b(s))\,\mathrm{d}s\Big\}.$$

Use this to argue that there is finite-dimensional convergence in distribution for the Z_m process.

(c). Demonstrate tightness of Z_m via Ex. 9.15. This leads to the clear existence of such an extended Gamma process, with mean and variance found above. Note however that increments $Z(t) - Z(s)$ are not Gamma distributed; show that this would require b to be constant over that interval.

Ex. 9.18. *Beta processes.* (xx to be done. perhaps pushed to the BNP chapter. $A_m(t) = \sum_{i/m \le t} B_{m,i}$, where $B_{m,i} \sim \mathrm{Beta}(ca_0(i/m)(1/m), c\{1 - a_0(i/m)(1/m)\})$. Show $\mathrm{E}\, A_m(t) \to A_0(t)$ and $\mathrm{Var}\, A_m(t) \to A_0(t)/(c+1)$. then convergence $A_m \to_d A$. pointers to BNP. xx)

Ex. 9.19. *Poisson processes.* Poisson processes are fundamental in probability theory and statistics, on their own, also via extensions and further constructions. Here we establish existence and basic properties, via limits of simpler processes and the convergence machinery for $D[0, 1]$.

(a). For each given m, let $Y_{m,i}$ be i.i.d. Pois(λ/m), and consider $Z_m(t) = \sum_{i/m \le t} Y_{m,i}$ for $t \ge 0$, a well-defined random process in $D[0,1]$. Show that Z_m has independent increments, that $Z_m(t) \sim \text{Pois}(\lambda[mt]/m)$, and that $Z_m(t) - Z_m(s) \to_d \text{Pois}(\lambda(t-s))$ for each cell $[s,t]$.

(b). From the simple expression $\text{E}\{Z_m(t) - Z_m(s)\} = \lambda([mt] - [ms])/m$, verify that Z_m is tight, via criteria in Ex. 9.15. Deduce that there is convergence to a well-defined process $Z = \{Z(t)\colon t \ge 0\}$, with independent increments $Z(t) - Z(s) \sim \text{Pois}(\lambda(t-s))$.

(c). With an infinite sequence $Y_{m,i} \sim \text{Pois}(\lambda/m)$, we may define $Z_m(t) = \sum_{i/m \le t} Y_{m.i}$ for all $t \ge 0$, i.e. beyond the $[0,1]$ interval. Argue that the $Z_m \to_d Z$ convergence above holds for each $D[0,\tau]$, for any finite τ, and hence that $Z = \{Z(t)\colon t \ge 0\}$ may be seen as a well-defined Poisson process for the full half-line. For Z such a Poisson process, let T_1 be the time for the first event. Show that $T_1 \sim \text{Expo}(\lambda)$. More generally, let $T_1 < T_2 < \cdots$ be the random time-points at which the Poisson events occur. Representing these via waiting times $W_i = T_i - T_{i-1}$, show that $W_1, W_2, \ldots$ are i.i.d. Expo(λ).

(d). Such processes may also be interpreted, constructed, and analysed via transition probabilities. Conditioning on what has happened up to time t, then considering what will happen over $[t, t+h]$, show that

$$\begin{aligned} \Pr(Z(t+h) = z \,|\, Z(t) = z\} &= 1 - \lambda h + o(h), \\ \Pr(Z(t+h) = z+1 \,|\, Z(t) = z\} &= \lambda h + o(h), \\ \Pr(Z(t+h) \ge z+2 \,|\, Z(t) = z\} &= o(h). \end{aligned}$$

(e). (xx formulate clear questions giving equivalence here; start from transition things, or start from waiting times, will lead to Poisson process. xx)

(f). For t inside a wider time window $[0,T]$, show that $Z(t) \,|\, \{Z(T) = n\}$ is a binomial $(n, t/T)$. More generally, if $0 = t_0 < t_1 < \cdots < t_k = T$, and N_j the number of events inside $(t_{j-1}, t_j]$, show that if $Z(t) = n$ is given, then $(N_1, \ldots, N_k)$ is multinomial $(n, p_1, \ldots, p_k)$, with $p_j = (t_j - t_{j-1})/T$.

(g). Given $Z(T) = n$, as in the previous points, with Poisson events occurring at time points $T_1 < \ldots < T_n$, show that the distribution of $(T_1, \ldots, T_n)$ is identical to that of $(U_{(1)}, \ldots, U_{(n)})$, the order statistics for a uniformly distributed sample on $[0,T]$. (xx pointer to Story v.8. amend text in that story to point here. xx)

Ex. 9.20. *Nonhomogeneous Poisson processes and Bernoulli proccesses.* We know that many small probabilities give a Poisson, see e.g. Ex. 2.8. Here we exhibit process generalisations of such findings, leading also to nonhomogeneous Poisson processes.

(a). For $Y \sim \text{binom}(n,p)$, show that its Laplace transformation is $L_n(u) = \text{E}\exp(-uY) = \{\exp(-u)p + 1 - p\}^n$. Use this to show that with n increasing and p decreasing in such a way that $np \to \lambda$, then $L_n(u)$ converges to $\exp[-\lambda\{1 - \exp(-u)\}]$, which is the Laplace transform of a Pois(λ).

(b). Now study a Bernoulli event process, of the form $Z_m(t) = \sum_{i/m \le t} B_{m,i}$, with independent Bernoulli components $B_{m,i} \sim \text{binom}(1, \lambda(i/m)(1/m))$. Here $\lambda(s)$ is a function, perhaps evolving over time, with cumulative intensity function $\Lambda(t) = \int_0^t \lambda(s)\,\text{d}s$. Show that Z_m has independent increments, and that Laplace transform converges,

$$\text{E}\exp\{-uZ_m(t)\} = \prod_{i/m \le t} [1 - \lambda(i/m)\{1 - \exp(-u)\}] \to \exp[-\Lambda(t)\{1 - \exp(-u)\}].$$

(c). Demonstrate also that the Z_m sequence is tight. So there is convergence of the fine-grid cumulative Bernoulli process, with small probabilities, to a full process with $Z(t) - Z(s) \sim \text{Pois}(\int_s^t \lambda(u)\,\mathrm{d}u)$.

(d).

Ex. 9.21. *Mapping a process to another process.* We've seen in Ex. 9.8 that if $X_n \to_d X$, in function spaces $C[0,1]$ or $D[0,1]$, and if g is continuous mapping to the real line, then $g(X_n) \to_d g(X)$. We are also taking an interest in mappings from a function to another function.

(a). Assume a mapping $g\colon C[0,1] \to C[0,1]$ is continuous. Show that if $X_n \to_d X$ in $C[0,1]$, then indeed the new random process $g(X_n)$ tends to the random limit process $g(X)$.

(b). Assume $X_n \to_d W$ in $C[0,1]$, and consider the new random process $Y_n(t) = \exp\{X_n(t)\}$. Show that $Y_n \to_d Y$, with $Y(t) = \exp\{W(t)\}$. Note that each $Y(t)$ has a log-normal distribution. Find the mean and variance of $Y(t)$, and also the covariance betwen $Y(t_1)$ and $Y(t_2)$.

(c). Assume $X_n \to_d W$, and consider $Z_n(t) = X_n(t) - tX_n(1)$ for $t \in [0,1]$. Show that the Z_n process converges in distribution to a well-defined Z zero-mean normal process, and find its covariance function.

(d). xx

Ex. 9.22. *A Gamma process of a Poisson process.* Let $X = \{X(t)\colon t \geq 0\}$ be a Poisson process with rate λ, and let $Y = \{Y(t)\colon t \geq 0\}$ be a Gamma process with $Y(t) \sim \text{Gam}(at, b)$. From these define the combined process $Z(t) = Y(X(t))$. Find its mean and variance. [xx a bit more. show that it can be reached as a limit of $Y_n(X_n(t))$, with small things placed at i/m for $i = 0, 1, 2, \ldots$. also its distribution. do some simulations. xx]

The empirical distribution process and the Brownian bridge

the e.c.d.f. process

Ex. 9.23. *The uniform empirical distribution process.* (xx some cleaning needed. xx) We have encountered the empirical distribution function in Ex. 3.9, including its use to form monitoring processes. Here we go further. Consider $U_1, U_2, \ldots$ being i.i.d. from the uniform distribution. We study the empirical distribution process $G_n(t) = n^{-1} \sum_{i=1}^n I(U_i \leq t)$ for $t \in [0,1]$. We have seen in the exercise pointed to that the process $Z_n(t) = \sqrt{n}\{G_n(t) - t\}$ has the property that $Z_n(t) \to_d \text{N}(0, t(1-t))$, for each t. Here we establish that there is a well-defined process limit, actually $Z_n \to_d W^0$, the Brownian bridge. Note indeed that Z_n starts and ends at zero.

(a). From the binomial CLT, show that $Z_n(t) \to_d \text{N}(0, t(1-t))$. Show also that for $s < t$, $(Z_n(s), Z_n(t)) \to_d (A, B)$, say, a binormal zero-mean with variances $s(1-s)$ and $t(1-t)$, and covariance $s(1-t)$.

(b). As an interlude, turning out to be relevant in a minute, study a zero-mean binormal (X, Y), with variances σ_1^2 and σ_2^2 and correlation ρ. Show that $\text{E}\, X^2Y^2 = \sigma_1^2\sigma_2^2(1+2\rho^2) \leq 3\sigma_1^2\sigma^2$. If now $(\widehat{p}, \widehat{q}) = (A/n, B/n)$, in a setup where (A, B) is trinomial (n, p, q), consider $K_n(p,q) = \text{E}\,|\sqrt{n}(\widehat{p} - p)\sqrt{n}(\widehat{q} - q)|^2$. Show that $K_n(p,q) \to K(p,q) = \text{E}\, X^2Y^2$, where

(X, Y) is binormal, with $\sigma^2 = p(1-p)$, $\sigma_2^2 = q(1-q)$, and covariance $-pq$. Deduce that $K(p,q) \le 3pq(1-p)(1-q) \le 3pq$, which is what we need below.

(c). Now back to G_n and the empirical process Z_n. There is actually full convergence in distribution here, $Z_n \to_d W^0$ in $D[0,1]$, to the Brownian bridge. Show this – which means demonstrating (i) finite-dimensional convergence and (ii) tightness. For the latter, use the above trinomial work to show that

$$\mathrm{E}\,|Z_n(t) - Z_n(s)|^2\,|Z_n(u) - Z_n(t)|^2 \to K(t-s, u-t) \le 3(t-s)(u-t) \le 3(u-s)^2.$$

and appeal to Ex. 9.14.

(d) (xx perhaps a separate exercise: constructing and deriving the Pearson chi-squared test from the basic bridge limit. in that case link to Story vii.1. xx).

Ex. 9.24. *The empirical distribution process.* In generalisation of what we worked with in Ex. 9.23, consider i.i.d. variables $X_1, \ldots, X_n$ from some continuous and strictly monotone c.d.f. F, with empirical c.d.f. $F_n(t) = n^{-1}\sum_{i=1}^n I(X_i \le t)$. We shall establish a limit process for $Z_n^*(t) = \sqrt{n}\{F_n(t) - F(t)\}$.

(a). Write $X_i = F^{-1}(U_i)$, with the U_i being i.i.d. and uniform. Show that $F_n(t) = G_n(F(t))$, in terms of the empirical c.d.f. for the U_i. From this deduce that $Z_n^*(t) = \sqrt{n}\{G_n(F(t)) - F(t)\} = Z_n(F(t))$, with Z_n the uniform empirical c.d.f. process of Ex. 9.23.

(b). Show from this that $Z_n^*(t) \to_d W^0(F(t))$, both pointwise and as a process; the limit is time-transformed Brownian motion.

(c). (xx edit and fix the following. we use relevant detail for Story vi.1. $D_n = \sqrt{n}\max|i/n - F(x_{(i)})|$ has right limit for D_n. xx) So how much can $F_n(t)$ deviate from its mean function $F(t)$? One answer is to apply the max functional. Start with

$$D_n = \sqrt{n}\max_t |F_n(t) - F(t)| = \max_u |G_n(u) - u| = \max_{i \le n}\{|G_n(i/n-) - i/n|, |G_n(i/n) - i/n|\},$$

the maximum distance from $F_n(t)$ to $F(t)$, also equal to the maximum distance from $G_n(t)$ to t. (xx a bit more here. xx) Then deduce

$$D_n = \max_{0 \le t \le 1} |Z_n(t)| \to_d D = \max_{0 \le t \le 1} |W^0(t)|.$$

The distribution of this limit D is sometimes called the Kolmogorov–Smirnov distribution, and due to its importance has been tabulated since pre-computer days. It can easily be simulated, without knowing any formulae for its c.d.f. or density, but we do return to such (xx in some exercise xx). We record here that 0.8276, 1.2238, 1.3580, 1.6271 are its 0.50, 0.90, 0.95, 0.99 quantiles. So $\max_t |G_n(t) - t| \le 1.358/\sqrt{n}$, with probability close to 0.95. If you have data and compute D_n, with D_n bigger than 1.358, you might be skeptical about the assumption that your data are uniform.

(d). Pretend that we have lost the Kolmogorov–Smirnov tables, and that we don't have the time to derive its (rather complicated) distribution. Simulate 10^5 paths in your computer, and read off the 0.95 point, which should then be close to the 1.358 number.

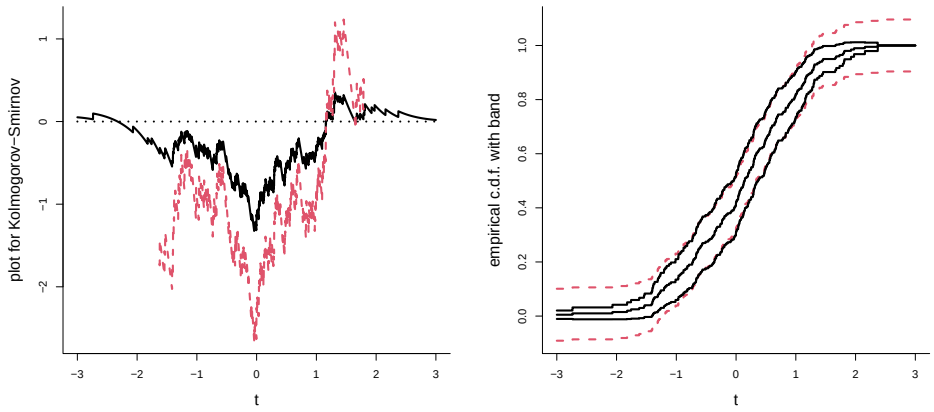


Figure 9.2: *Left panel: plots of the Kolmogorov–Smirnov related processes* $Z_n = \sqrt{n}(F_n - F)$ *and its weighted version* Z_n^**, as in Ex. 9.25 and 9.26, to check if 200 data points from the standard normal look nonnormal or not. Here the* $\max_t |Z_n(t)|$ *and* $\max |Z_n^*(t)|$ *values are 1.324 and 2.652, both lower than the respective uppper 0.05 points, 1.358 and 3.121, of their limit distributions under normality. Right panel: simultaneous 95 percent confidence bands for* F*, one with constant width valid for all* t*, the other of width proportional to the standard deviation, valid in the central 90 percent data range.*

Ex. 9.25. *The Kolmogorov–Smirnov test and band.* Results and insights from Ex. 9.24 can be used for purposes of testing and confidence bands. Suppose you have i.i.d. data $Y_1, \ldots, Y_n$ and need to test the hypothesis H_0 that their distribution F is equal to some given F_0, like the standard normal. Consider the empirical c.d.f. $F_n(t) = (1/n)\sum_{i=1}^n I(y_i \leq t)$. Under F_0, we have established above that $Z_n = \sqrt{n}(F_n - F_0) \to_d Z = W^0(F_0(\cdot))$.

the Kolomogorov–Smirnov test

(a). Consider the test for $F = F_0$ consisting in rejecting if $F_0(t)$ for some part of the t domain is outside the Kolmogorov–Smirnov band. Show that this test has level 0.05 for large n. Show also that it is equivalent to rejecting if $D_n = \sqrt{n}\max_t |F_n(t) - F_0(t)| > 1.358$.

(b). We have learned above that $\sqrt{n}\max_t |F_n(t) - F(t)|$ tends in distribution to $D = \max_t |W^0(F(t))| = \max_s |W^0(s)|$. Explain that this implies

$$\Pr\{F(t) \in F_n(t) \pm 1.358/\sqrt{n}, \text{ for all } t\} \to 0.95.$$

This is the Kolmogorov–Smirnov simultaneous confidence band.

(c). Simulate $n = 100$ points from a standard normal, and construct a version of Figure 9.2, with the empirical process $Z_n = \sqrt{n}(F_n - F_0)$, left panel, and the confidence band, right panel.

(d). (xx this is too remarky; perhaps push to Notes. xx) Note in particular that since D_n has a limit distribution, we must have $D_n/\sqrt{n} \to_{\rm pr} 0$; the maximal difference, between the e.c.d.f. F_n and the true F, goes to zero. The stronger version of this statement, that this happens with probability 1, is the Glivenko–Cantelli theorem of Ex. 2.61. It tells us that with lots of data, we can uncover any tiny little aspect of their underlying

distribution, we can be as sophisticated as we might wish to. The result above is however more informative, since we learn how quickly it goes to zero, namely with speed $1/\sqrt{n}$. This is incidentally impossible for density estimators, where the best one can hope for is that the maximal distance $\max_t |f_n(t) - f(t)|$ goes to zero with speed $1/n^{2/5}$.

Ex. 9.26. *Weighting the Kolmogorov–Smirnov test and band.* (xx to be done. also since we care about $W^0(s)/\{s(1-s)\}^{1/2}$ for certain monitoring bridges to come. calibrate with later exercises. point to early exercise with quantile values. xx) Consider as in Ex. 9.25 i.i.d. data $Y_1, \ldots, Y_n$ from some F, with F_n the empirical c.d.f.

(a). With $v(s)$ a nonnegative bounded function on the unit interval, with finitely many discontinuities, show that $Z_n^*(t) = \sqrt{n}v(F_n(t))\{F_n(t)-F(t)\}$ tends to $Z^*(t) = v(F(t))W^0(F(t))$. From this explain that

$$D_n^* = \sqrt{n}\max_t v(F_n(t))|F_n(t) - F(t)| \quad \text{has } D_n^* \to_d D^* = \max v(s)|W^0(s)|,$$

with the same limit obtained if $v(F_n(t))$ is replaced by $v(F(t))$.

(b). This invites using weighted versions of the Kolmogorov–Smirnov test and band, with $v(t)$ chosen to reflect wished-for features, like giving more importance to some data regions than others. Consider in particular $v(s) = 1/\{s(1-s)\}^{1/2}$ on say $[\varepsilon, 1-\varepsilon]$, with ε small and positive, and $v(s) = 0$ outside that interval. Show that

$$D_n^* = \sqrt{n}\max_{\varepsilon \le F(t) \le 1-\varepsilon} \frac{|F_n(t) - F(t)|}{[F(t)\{1-F(t)\}]^{1/2}} \quad \text{has } D_n^* \to_d D^* = \max_{\varepsilon \le s \le 1-\varepsilon} \frac{|W^0(s)|}{\{s(1-s)\}^{1/2}}.$$

the absolute maximum of the naturally normalised Brownian bridge, to make the variance constant and equal to one. This is a process and a distribution we have met in Ex. 9.10, along with a few quantiles based on simulations, for e.g. the $[0.05, 0.95]$ interval. Explain how this may be used to test if the data come from some given $F = F_0$, say the standard normal.

(c). Explain that

$$\Pr_F\{F(t) \in F_n(t) \pm 3.121\,[F_n(t)\{1-F_n(t)\}]^{1/2}/\sqrt{n}, \text{ for all } t \in A_n\} \to 0.95,$$

where A_n contains those t for which $\varepsilon \le F_n(t) \le 1-\varepsilon$. We have hence constructed a simultaneous confidence band, with approximate 95 percent level, proportional to the standard deviation of $F_n(t)$, as opposed to the fixed bandwidth construction of Ex. 9.25.

(d). (xx redo the Figure 9.2 things to include both versions. xx)

Ex. 9.27. *Local power of the Kolmogorov–Smirnov test.* For an i.i.d. data sample $y_1, \ldots, y_n$, testing a distributional assumption, of the form $F = F_0$, may be carried out using the method developed in Ex. 9.25. In particular, with $D_n = \sqrt{n}\max_t |F_n(t)-F_0(t)|$, we have $D_n \to_d D = \max_s |W^0(s)|$ under the null hypothesis.

(a). Consider a so-called local alternative, where the real c.d.f. takes the form $G_n(t) = F_0(t) + \delta(t)/\sqrt{n}$, for some discrepancy function $\delta(t)$. Show that

$$Z_n(t) = \sqrt{n}\{F_n(t) - F_0(t)\} \to_d W^0(F(t)) + \delta(t).$$

Let $d_0 = 1.358$ be the upper 0.05 point of the distribution for D. For the local power function, show that

$$\pi_n = \Pr_{G_n}(D_n \ge d_0) \to \pi(\delta) = \Pr\{\max_t |W^0(F_0(t)) + \delta(t)| \ge d_0\}.$$

(b). As an example, suppose one tests $F = F_0 = \Phi$, the standard $\mathrm{N}(0,1)$, where an alternative is a certain contamination with data from $\mathrm{N}(a,1)$. We formalise this as data coming from

$$G_n(t) = (1 - c/\sqrt{n})F_0(t) + (c/\sqrt{n})F_0(t-a) = F_0(t) + (c/\sqrt{n})\{F_0(t-a) - F_0(t)\}.$$

Show that the local power in question has a clear limit,

$$\pi_n(a,c) \to \pi(a,c) = \Pr\{\max_s |W^0(s) + c\{\Phi(\Phi^{-1}(s) - a) - s\}| \geq d_0\}.$$

(c). (xx need to do this carefully, and then phrasing the questions better. xx) Use simulation to compute approximate power curves for the Kolmogorov–Smirnov test, for say $n = 200$, (i) as a function of contamination level $c/\sqrt{n}$, for fixed $a = 1$, and (ii) as a function of contamination distance a, for fixed contamination level $c = 0.05\,n$.

(d). (xx generalise. invent another alternative. xx)

Monitoring processes for models with estimated parameters

Ex. 9.28. *Monitoring processes and Kolmogorov–Smirnov with estimated parameters.* (xx need cleaning and polish. xx) Suppose again that $Y_1, \ldots, Y_n$ are i.i.d. from some c.d.f. F, which potentially belongs to some given parametric class, with density $f(y,\theta)$ and c.d.f. $F(\cdot,\theta)$. How can we test that modelling assumption? Assuming the model holds, with θ_0 the true parameter, let $\widehat{\theta}$ be the ML estimator, for which we know (xx from Ch5 xx) that $\widehat{\theta} - \theta_0 = J^{-1}U_n + o_{\mathrm{pr}}(n^{-1/2})$, with J the Fisher information matrix at θ_0 and $U_n = n^{-1}\sum_{i=1}^n u(Y_i,\theta_0)$, with $u(y,\theta)$ the score function.

(a). Show that

$$F(t,\widehat{\theta}) \doteq F(t,\theta_0) + H(t)^{\mathrm{t}}(\widehat{\theta} - \theta_0) \doteq F(t,\theta_0) + H(t,\theta_0)^{\mathrm{t}}J^{-1}U_n,$$

in which

$$H(t,\theta_0) = \partial F(t,\theta_0)/\partial\theta = \int_{-\infty}^t f(y,\theta_0)u(y,\theta_0)\,\mathrm{d}y.$$

The natural monitoring process to examine is $Z_n(t) = \sqrt{n}\{F_n(t) - F(t,\widehat{\theta})\}$. With $Z_{n,0}(t)$ the simpler $\sqrt{n}\{F_n(t) - F(t,\theta_0)\}$, explain that

$$Z_n(t) = Z_{n,0}(t) - H(t,\theta_0)^{\mathrm{t}}J^{-1}\sqrt{n}U_n + o_{\mathrm{pr}}(1).$$

Supplement the $Z_{n,0}$ results of Ex. 9.24 with the broader statement $(Z_{n,0}(t), U_n) \to_d (Z_0(t), U)$, in terms of joint convergence, where $Z_0(t) = W^0(F(t,\theta_0))$ is time-transformed Brownian bridge, $U \sim \mathrm{N}_p(0,J)$, and

$$\mathrm{cov}(Z_0(t), U) = \mathrm{cov}\{I(Y_i \leq t), u(Y,\theta_0)\} = H(t,\theta_0).$$

Show from this that there is indeed process convergence

$$Z_n(t) \to_d Z(t) = Z_0(t) - H(t,\theta_0)^{\mathrm{t}}J^{-1}U.$$

Show furthermore that the covariance function for the limit process is

$$\mathrm{cov}\{Z(t_1), Z(t_2)\} = F(t_1,\theta_0)\{1 - F(t_2,\theta_0)\} - H(t_1,\theta_0)^{\mathrm{t}}J^{-1}H(t_2,\theta_0)$$

for $t_1 \leq t_2$.

(b). (xx to be done and cleaned, and used later. should point to a lemma about $A_n(B_n(v)) \to_d A(b(v))$ when $A_n \to_d A$ and $B_n(v) \to_{\rm pr} b(v)$. xx) We now carry out a convenient time transformation, bringing processes to $[0,1]$. Transforming by $F^{-1}(v,\widehat{\theta})$, show that

$$A_n(v) = Z_n(F^{-1}(v,\widehat{\theta})) = \sqrt{n}\{G_n(v) - v\},$$

in terms of the empirical c.d.f. G_n of the $\widehat{v}_i = F(y_i,\widehat{\theta})$. Show that

$$A_n(v) \to_d A(v) = Z(F^{-1}(v,\theta_0)) = W^0(v) - H_0(v)^{\rm t} J^{-1} U,$$

writing $H_0(v) = H(F^{-1}(v,\theta_0))$, and that this transformed limit process has covariance function

$$\mathrm{cov}\{A(v_1), A(v_2)\} = v_1(1-v_2) - H_0(v_1)^{\rm t} J^{-1} H_0(v_2) \quad \text{for } v_1 \le v_2.$$

(xx also: $\max_t |Z_n(t)| = \max_v |A_n(v)|$. so can carry out computation and monitoring on the scale of $v \in [0,1]$. xx)

(c). (xx repair in view of the above. xx) Explain that this leads to a Kolmogorov–Smirnov test for cases with estimated parameters. Specifically, with $D_n = \sqrt{n}\max_t |F_n(t) - F(t,\widehat{\theta})|$, show that $D_n \to_d D = \max_t |Z(t)|$. It is also convenient to transform these processes to the $[0,1]$ interval, using $A(v) = Z(F^{-1}(v,\theta_0))$. Writing F_0 and F_0^{-1} for $F(t,\theta_0)$ and its inverse, show that this process has covariance function

$$\mathrm{cov}\{A(v_1), A(v_2)\} = v_1(1-v_2) - H(F^{-1}(v_1))^{\rm t} J^{-1} H(F^{-1}(v_2)) \quad \text{for } v_1 \le v_2.$$

(d). For an illustration, consider testing that data come from the $\mathrm{N}(\xi,1)$, for some ξ. In the notation above, show that $u(y,\xi) = y - \xi$, that $J = 1$, and that

$$Z_n(t) = \sqrt{n}\{F_n(t) - \Phi(t - \widehat{\xi})\} \to_d Z(t) = Z_0(t) + \phi(t-\xi_0)U,$$

in which U is standard normal and $\mathrm{cov}\{Z_0(t), U\} = -\phi(t-\xi_0)$.

Ex. 9.29. *Monitoring processes and tests for normality.* We now apply the general methods and results of Ex. 9.28 to the important case of testing whether data follow a normal distribution, i.e. with both parameters unknown.

(a). For the two-parameter $\mathrm{N}(\xi,\sigma^2)$ model, show that $Z_n(t) = \sqrt{n}\{F_n(t) - \Phi((y - \widehat{\xi})/\widehat{\sigma})\} \to_d Z(t)$, with the limit process being

$$Z(t) = W^0(\Phi(s)) + \phi(s)U_1 + \phi(s)sU_2, \quad \text{writing } s = (t-\xi_0)/\sigma_0,$$

in which U_1, U_2 are independent normals with variances 1 and 2, and explain how these are connected to $W^0(\Phi(s))$. Deduce that

$$D_n = \sqrt{n}\max_t |F_n(t) - \Phi((t-\widehat{\xi})/\widehat{\sigma})| \to D = \max_s |Z^*(s)|,$$

where $Z^*(s)$ has a distribution independent of the parameters. Show indeed that with this notation, for $s_1 \le s_2$,

$$\mathrm{cov}\{Z^*(s_1), Z^*(s_2)\} = \Phi(s_1)\{1-\Phi(s_2)\} - \phi(s_1)\phi(s_2) - \tfrac{1}{2}\phi(s_1)\phi(s_2)s_1 s_2.$$

(b). We need upper quantiles of the D distribution, for testing normality with this type of Kolmogorov–Smirnov with estimated parameters. Explain that there are two ways of simulating close approximations to this D; (i) to simulate from the actual D_n, for a high n, and (ii) to ignore the Z_n processes but simulate via the limit process. For the latter, show that the limit process $A(v) = Z^*(\Phi^{-1}(v))$ on $[0,1]$ has covariance structure

$$\operatorname{cov}\{A(v_1), A(v_2)\} = v_1(1-v_2) - \phi(\Phi^{-1}(v_1))\phi(\Phi^{-1}(v_2))\{1 + \tfrac{1}{2}\Phi^{-1}(v_1)\Phi^{-1}(v_2)\}.$$

To simulate from the A we may use the general recipe of Ex. 9.11, for a fine grid of v in $[0,1]$. Carry out versions of both schemes, to arrive at numbers close to the following. Here are quantiles at levels $0.50, 0.90, 0.95, 0.99$ found from from having simulated $2 \cdot 10^4$ times, with a fine grid $0, 0.001, \ldots, 0.999, 1$ for the Gaussian limit process.

```
method (i) from Z_n        method (ii) from limit process with grid
0.50  0.90  0.95  0.99     0.50  0.90  0.95  0.99
0.822 1.202 1.360 1.613    0.805 1.205 1.341 1.614
0.632 0.854 0.949 1.107    0.638 0.872 0.949 1.101
0.601 0.822 0.885 1.044    0.600 0.812 0.888 1.039
```

(xx a few more lines to explain features. first line is the traditional KS. with more parameters, D becomes smaller. xx)

(c). (xx to be done and polished. xx)

(d). The basic process result, that Z_n tends to a well-defined zero-mean Gaussian process Z, was used above to identify the limit distribution of the Kolmogorov–Smirnov type test statistic $\max_t |Z_n(t)|$. It may however also be used to form other tests. To illustrate this, consider (xx do this, with simulations, for absolute integral or CvM xx).

Ex. 9.30. *Monitoring processes and tests for exponentiality.* For another application of the general methodology of Ex. 9.28, consider the exponential model with c.d.f. $F(t,\theta) = F_0(\theta t)$, writing $F_0(t) = 1 - \exp(-t)$ for $t \geq 0$.

(a). With i.i.d. data $Y_1, \ldots, Y_n$, show that the ML estimator is $\widehat{\theta} = 1/\bar{Y}$ and that the Fisher information is $1/\theta^2$. Show that the results above lead to

$$Z_n(t) = \sqrt{n}\{F_n(t) - F_0(\widehat{\theta}t)\} \to_d Z_0(t) + \exp(-\theta_0 t)\theta_0 t\, U_0,$$

where U_0 is standard normal, with $\operatorname{cov}\{Z_0(t), \theta_0 U_0\} = -\exp(-\theta_0 t)t$. Write down the variance of the limit process. It is convenient here to transform this further, to the monitoring process

$$Z_n^*(t) = Z_n(t/\widehat{\theta}) = \sqrt{n}\{G_n(t) - F_0(t)\}, \quad \text{with } G_n(t) = F_n(t/\widehat{\theta}) = n^{-1}\sum_{i=1}^{n} I(\widehat{\theta}Y_i \leq t),$$

which has a limit process free of the θ_0 parameter. Show indeed that $Z_n^*(t) \to_d Z^*(t) = Z(t/\theta_0)$, with

$$\operatorname{cov}\{Z^*(t_1), Z^*(t_2)\} = F_0(t_1)\{1 - F_0(t_2)\} - t_1 t_2 \exp(-t_1 - t_2) \quad \text{for } t_1 \leq t_2.$$

With $D_n^* = \sqrt{n}\max_t |G_n(t) - F_0(t)|$, explain that $D_n^* \to_d D^* = \max_t |Z^*(t)|$.

(b). To carry out such a test in practice, one needs quantiles from the limit distribution D. As with the normality tests of Ex. 9.29 one may use method (i), choosing a high n and simulating the Z_n process, of method (ii), ignoring sampling exponentials and instead simulate from an approximation to the Gaussian limit process, for this purpose again using the general recipe of Ex. 9.11. Carry out such simulations, to find 0.50, 0.90, 0.95, 0.99 quantiles of the D^* distribution, as those given below. It might be simpler, though, to simulate actual paths of Z_n, for the given situation, with its n, at the ML estimated position $\widehat{\theta}$, and then read off realisations of D_n. You may check Story i.5 to see such normal monitoring processes Z_n, checking for normality of birthweight distributions.

```
method (i) from Z_n        method (ii) from limit process with grid
0.50  0.90  0.95  0.99     0.50  0.90  0.95  0.99
0.822 1.202 1.359 1.612    0.811 1.201 1.336 1.610
0.695 0.980 1.075 1.296    0.686 0.977 1.075 1.289
```

Ex. 9.31. *Monitoring processes for linear regression residuals.* (xx ranting first, to be polished. we use the smoothing trick to allow Taylor arguments. xx) Consider the linear regression model $Y_i = x_i^{\rm t}\beta + \sigma\varepsilon_i$, with x_i of length p, where the ε_i are independent standard normals; see Ex. 3.34–3.35. We assume here that the first column of the $n \times p$ covariate matrix X is $(1, \ldots, 1)^{\rm t}$, i.e. there is an intercept in the model. After having found least squares estimators $\widehat{\beta}$ and the consequent $\widehat{\sigma}^2 = Q_0/m$, with $Q_0 = \sum_{i=1}^n (Y_i - x_i^{\rm t}\widehat{\beta})^2$ and $m = n - p$, we may form the estimated residuals $\widehat{\varepsilon}_i = (Y_i - x_i^{\rm t}\widehat{\beta})/\widehat{\sigma}$. Under model conditions, these have the t_m distribution, but with negative correlations, so the Brownian bridge result of Ex. 9.24 does not hold for their empirical c.d.f. G_n. Below we put in the extra work to understand the behaviour of G_n, which then may be used to check model conditions.

(a). We start by taking σ known and equal to 1, so $\widehat{\varepsilon}_i = Y_i - x_i^{\rm t}\widehat{\beta}$, seen as estimated versions of the $\varepsilon_i = Y_i - x_i^{\rm t}\beta_0$, with β_0 the true parameter. With $\Sigma_n = n^{-1}\sum_{i=1}^n x_i x_i^{\rm t}$, write the hat matrix of Ex. 3.35 as $H = X(X^{\rm t}X)^{-1}X^{\rm t} = n^{-1}M$, with $M = X\Sigma_n^{-1}X^{\rm t}$. Explain that $\widehat{\varepsilon} = Y - X\widehat{\beta} = (I - H)Y \sim {\rm N}_n(0, I - n^{-1}M)$; in particular, $\widehat{\varepsilon}_i \sim {\rm N}(0, 1 - n^{-1}m_{i,i})$. Show that $n^{-2}\sum_{i,j} m_{i,j} = 1$, the average value of the M matrix. For one more useful fact here, let as usual $\bar{x} = n^{-1}\sum_{i=1}^n x_i$, the first component of which being 1. Show that $\bar{x}^{\rm t}\Sigma_n^{-1}\bar{x} = 1$. Below we shall assume that the $m_{i,j}$ are essentially stable in size, in the sense that $\max_{i,j} m_{i,j}/n \to 0$.

(b). There are indeed connections between

$$V_n(t) = \sqrt{n}\{F_n(t) - \Phi(t)\} \quad \text{and} \quad W_n(t) = \sqrt{n}\{G_n(t) - \Phi(t)\},$$

where F_n and G_n are the empirical c.d.f.s for the ε_i and $\widehat{\varepsilon}_i$, respectively. Explain that we already know that $V_n(t) \to_d V(t) = W^0(\Phi(t))$, with the Brownian bridge. Intuition says G_n should be only a delta method step away from F_n, but analysis is not easy in that the $I(\widehat{\varepsilon}_i \le t)$ functions do not have simple derivatives. For this technical reason we introduce a smoothing operation. Let $k(u)$ be a kernel density function, having smooth derivatives, symmetric and supported on $[-1, 1]$, with c.d.f. K. Explain that $K_h(u) = K(u/h)$ tends to $K_0(u)$, which is 0 for $u < 0$, $\frac{1}{2}$ for $u = 0$, and 1 for $u > 0$, as h shrinks to zero. Explain that in $\sum_{i=1}^n \{K_h(t - \varepsilon_i) - K_0(t - \varepsilon_i)\}$, most of terms are zero, and that only those for which $|t - \varepsilon_i| \le h$ make a difference. Explain further that the number of such terms is a binomial $(n, p_{n,h}(t))$, with $p_{n,h}(t) \doteq 2\phi(t)h$ for small h. Deduce from this that with

$$F_{n,h}(t) = n^{-1}\sum_{i=1}^n K_h(t - \varepsilon_i) \quad \text{and} \quad G_{n,h}(t) = n^{-1}\sum_{i=1}^n K_h(t - \widehat{\varepsilon}_i)$$

the kernel smoothed versions of F_n and G_n, we have

$$\sqrt{n}\{F_{n,h}(t) - F_n(t)\} \to_{\rm pr} 0 \quad \text{and} \quad \sqrt{n}\{G_{n,h}(t) - G_n(t)\} \to_{\rm pr} 0,$$

as long as $\sqrt{n}h \to 0$, as e.g. with $h \propto 1/n^c$ with $c \in (\frac12, 1)$. Show from this again that under the $\sqrt{n}h \to 0$ condition,

$$V_{n,h}(t) = \sqrt{n}\{F_{n,h}(t) - \Phi(t)\} \to V(t).$$

(c)**.** For these smoothed versions we may use Taylor approximations. With $k_h(u) = h^{-1}k(h^{-1}u)$ the density of K_h, show that

$$G_{n,h}(t) = F_{n,h}(t) - A_{n,h}(t)^{\rm t}(\widehat\beta - \beta_0) + o_{\rm pr}(n^{-1/2}),$$

with $A_{n.h}(t) = n^{-1}\sum_{i=1}^n k_h(\varepsilon_i - t)x_i$. (xx careful with details here. xx) With $\bar x \to x_0$, show that $A_{n,h}(t) \to_{\rm pr} \phi(t)x_0$. Explain that $x_0^{\rm t}\Sigma^{-1}x_0 = 1$, and that $\sqrt{n}(\widehat\beta - \beta_0) \to_d J^{-1}U$, with $U \sim {\rm N}_p(0, \Sigma)$. show

$$W_{n,h}(t) = \sqrt{n}\{G_{n,h}(t) - \Phi(t)\} \to_d W(t) = V(t) - \phi(t)x_0^{\rm t}\Sigma^{-1}U = V(t) + \phi(t)U^*,$$

where U^* is standard normal, with ${\rm cov}\{V(t), U\} = -\phi(t)$.

(d)**.** (xx then $\widehat\sigma$ too. G_n the empirical c.d.f. for $\widehat\varepsilon_i = (Y_i - x_i^{\rm t}\widehat\beta)/\widehat\sigma$. same limit. then KS and other tests. so can use $D_n = \sqrt{n}\max_t |G_n(t) - \Phi(t)| \to_d D = \max_t |V(t) + \phi(t)U^*|$. can simulate from limit distribution, or perhaps more simply simulate D_n for the given n. other few details. xx)

Ex. 9.32. *Monitoring processes for general regression models.* (xx from rant to polish. xx) Consider a general regression setting for pairs (x_i, y_i), with continuous outcomes Y_i following distributions with densities $f(y, x_i, \theta)$ and c.d.f.s $F(y, x_i, \theta)$. We do assume the model is correct, with θ_0 denoting the true parameter. Let $\widehat\theta$ be the ML estimator. For moderate to large n, $\widehat v_i = F(y_i, x_i, \widehat\theta)$ will be close to the real $v_i = F(y_i, x_i, \theta_0)$, and these again will be i.i.d. uniform under model conditions. Below let F_n and G_n the empirical c.d.f.s for these v_i and $\widehat v_i$, respectively. The aim is to understand the behaviour of

$$V_n(t) = \sqrt{n}\{F_n(v) - v\} \quad \text{and} \quad W_n(t) = \sqrt{n}\{G_n(v) - v\} \text{ for } v \in [0,1],$$

and their relation; limit process results for W_n may then be used for testing model conditions.

Below we shall partly simplify notation, using $f(y,x)$ and $F(y,x)$ for $f(y,x,\theta_0)$ and $F(y,x,\theta_0)$, i.e. the density and c.d.f. at the true parameter, and similarly $H(y,x)$ for $\partial F(y,x,\theta_0)/\partial\theta$; in particular, $H(y_i,x_i) = \partial v_i/\partial\theta$. There will also be use for the (xx Ch5 xx) result that $\widehat\theta - \theta_0 = J_n^{-1}U_n = o_{\rm pr}(n^{-1/2})$, with $U_n = n^{-1}\sum_{i=1}^n u(Y_i, x_i, \theta_0)$ the average score function at the true parameter and J_n the Fisher information matrix $-n^{-1}\sum_{i=1}^n {\rm E}\,\partial^2 \log f(Y_i, x_i, \theta_0)/\partial\theta\,\partial\theta^{\rm t}$. Furthermore, as for Ex. 9.31, it is technically helpful with smoothed versions, say $F_{n,h}(v) = n^{-1}\sum_{i=1}^n K_h(v - v_i)$ and $G_{n,h}(v) = n^{-1}\sum_{i=1}^n K_h(v - \widehat v_i)$. We need h to shrink to zero with $\sqrt{n}h \to 0$. We also assume below that the x_i are regular enough to ensure that averages $\bar g_n = n^{-1}\sum_{i=1}^n g(x_i)$ tend to well-defined limits $g_0 = \int g(x)\,{\rm d}Q(x)$, means with respect to the the long-term covariate distribution.

(a). Explain that we know from Ex. 9.24 that V_n tends to $V = W^0$, the Brownian bridge. Show that $V_{n,h}(t) = \sqrt{n}\{F_{n,h}(v) - v\}$ tends to the same limit V. Show further that

$$G_{n,h}(v) = F_{n,h}(v) - A_{n,h}(v)^{\mathrm{t}}(\widehat{\theta} - \theta_0) + o_{\mathrm{pr}}(n^{-1/2}),$$

with $A_{n,h}(v) = n^{-1}\sum_{i=1}^n k_h(v_i - v)H(y_i, x_i)$. To understand the $A_{n,h}(v)$ random function, show that $k_h(v_i - v)H(y_i, x_i)$ has expected value

$$\int k_h(w - v)H(F^{-1}(w, x_i), x_i)\,\mathrm{d}w = \int k(u)H(F^{-1}(v + hu, x_i), x_i)\,\mathrm{d}u,$$

and that this for small h leads to

$$\mathrm{E}\,A_{n,h}(v) = A_n(v) + O(h^2), \quad \text{with } A_n(v) = n^{-1}\sum_{i=1}^n H(F^{-1}(v, x_i), x_i).$$

By ergodic assumptions, the $A_n(v)$ tends in probability to a well-defined limit function $A(v)$.

(b). Establish that

$$W_{n,h}(v) = \sqrt{n}\{G_{n,h}(t) - v\} = V_{n,h}(v) - A_{n,h}(v)\sqrt{n}J_n^{-1}U_n + o_{\mathrm{pr}}(1)$$

with $A_{n,h}(v) \to_{\mathrm{pr}} A(v)$. Furthermore,

$$\mathrm{cov}\{I(v_i \le v), u(Y_i, x_i, \theta_0)\} = \int_{-\infty}^{F^{-1}(v,x_i)} f(y, x_i)u(y, x_i)\,\mathrm{d}y = H(F^{-1}(v, x_i), x_i).$$

From this establish that $\mathrm{cov}\{F_n(v), U_n\} = n^{-1}A_n(v)$. With $\sqrt{n}U_n \to_d U \sim \mathrm{N}_p(0, J)$, conclude that

$$W_n(v) \to_d V(v) - A(v)^{\mathrm{t}}J^{-1}U, \quad \text{where } \mathrm{cov}\{V(v), U\} = A(v).$$

For this limit process, establish that

$$\mathrm{cov}\{W(v_1), W(v_2)\} = v_1(1 - v_2) - A(v_1)^{\mathrm{t}}J^{-1}A(v_2) \quad \text{for } v_1 \le v_2.$$

(c). (xx do this. xx) show that W can be represented as the Brownian bridge V conditional on $U = 0$.

Ex. 9.33. *Monitoring processes for linear regression.* Here we apply the general results developed in Ex. 9.32 to the classic linear regression model, with y given x having c.d.f. $F(y, x, \beta, \sigma) = \Phi((y - x^{\mathrm{t}}\beta)/\sigma)$, with ML estimators $\widehat{\beta}$ and $\widehat{\sigma}$ as per Ex. 3.34–3.35. We compute the estimated residuals $\widehat{r}_i = (y_i - x_i^{\mathrm{t}}\widehat{\beta})/\widehat{\sigma}$ and from these $\widehat{v}_i = \Phi(\widehat{r}_i)$. For their empirical c.d.f. $G_n(v) = n^{-1}\sum_{i=1}^n I(\widehat{v}_i \le v)$ we shall identify the limit for the model monitoring process, $W_n(v) = \sqrt{n}\{G_n(v) - v\}$, under model conditions. (xx note. assume x_i has a 1, intercept, $\bar{x}^{\mathrm{t}}\Sigma_n^{-1}\bar{x} = 1$. can also use $F_{n-p}(\widehat{r}_i)$, with the c.d.f. for the t_{n-p}, with the same limit. xx)

(a). Show that $F^{-1}(v, x_i) = x_i^{\mathrm{t}}\beta + \sigma\Phi^{-1}(v)$, and that $H(y, x)$ has components $-(1/\sigma)x\phi(\varepsilon)$ and $-(1/\sigma)\varepsilon\phi(\varepsilon)$, in terms of $\varepsilon = (y - x^{\mathrm{t}}\beta)/\sigma$. With notation as in the general exercise pointed to, show that

$$H(F^{-1}(v, x), x) = -(1/\sigma)\begin{pmatrix} \phi(\Phi^{-1}(v))x \\ \phi(\Phi^{-1}(v)\Phi^{-1}(v)) \end{pmatrix},$$

and that this leads to

$$A_n(v) = n^{-1} \sum_{i=1}^{n} H(F^{-1}(v, x_i), x_i) = -(1/\sigma)\phi(\Phi^{-1}(v)) \begin{pmatrix} \bar{x} \\ \Phi^{-1}(v) \end{pmatrix}.$$

Use this to establish that $W_n \to_d W$, a zero-mean Gaussian process with covariance function

$$v_1(1 - v_2) - \phi(\Phi^{-1}(v_1))\phi(\Phi^{-1}(v_2))\{1 + \tfrac{1}{2}\Phi^{-1}(v_1)\Phi^{-1}(v_2)\} \quad \text{for } v_1 \le v_2.$$

(xx should now simulate these from general recipe, yielding upper quantiles for checking regression assumptions. xx)

(b). (xx another illustration: exponential regression, to be polished. perhaps we have a story with such data. xx) Consider the setup with $Y_i \sim \text{Expo}(\theta_i)$, with $\theta_i = \exp(x_i^{\mathrm{t}}\beta)$. With $F_0(u) = 1 - \exp(-u)$ and $f_0(u) = \exp(-u)$ the c.d.f. and density for the unit exponential, show that $F(y, x_i) = F_0(\theta_i y)$, $f(y, x_i) = f_0(\theta_i y)\theta_i$, $v_i = F_0(\theta_i y_i)$, $H(y, x_i) = f_0(\theta_i y)\theta_i y x_i$, $F^{-1}(v, x_i) = \theta_i^{-1} F_0^{-1}(v)$, $H(F^{-1}(v, x_i), x_i) = f_0(F_0^{-1}(v))F_0^{-1}(v)x_i$. this yields $A_n(v) = F_0^{-1}(v)(1 - v)\bar{x}_n$. With $\widehat{\beta}$ the ML estimator, leading to $\widehat{\theta}_i = \exp(x_i^{\mathrm{t}}\widehat{\beta})$, let G_n the empirical c.d.f. for $\widehat{v}_i = F_0(\widehat{\theta}_i Y_i)$. With $J_n = n^{-1}\sum_{i=1}^n x_i x_i^{\mathrm{t}}$, show that

$$W_n(v) = \sqrt{n}\{G_n(v) - v\} = \sqrt{n}\{F_n(v) - v\} - F_0^{-1}(v)(1 - v)\bar{x}_n^{\mathrm{t}} J_n^{-1}\sqrt{n}U_n + o_{\mathrm{pr}}(1),$$

with $\text{cov}\{F_n(v), U_n\} = n^{-1}A_n(v)$.

Ex. 9.34. *Forming new motions from old ones.* [xx a couple of things, to be polished and edited. show that $(1/\sqrt{2})W_1 + (1/\sqrt{2})W_2$ is a new Brownian motion. more generally a linear combination. xx]

(a).

Ex. 9.35. *Forming new bridges from old ones.* [xx a couple of things, to be polished and edited. more or less as for previous exercise, but now for bridges. xx]

(a).

Ex. 9.36. *Chi-squared processes and F processes.* Just as various distributions may be constructed and used based on the normal distribution, like the log-normal, the chi-squared, the t, the F, we may also use Brownian motions to construct further processes of interest.

(a). With W a Brownian motion, consider $Z(t) = W(t)^2/t$. Show that $Z(t)$ has a χ_1^2 distribution, for each t.

(b). More generally, let $W_1, \ldots, W_p$ be independent Brownian motions. Show that $Z_p(t) = \sum_{j=1}^p W_j(t)^2/t$ has a χ_p^2 distribution, for each t. Simulate and display some paths, for say $p = 6$.

(c). With integers p, q, construct a process $F = \{F(t) : t > 0\}$ such that for each t, $F(t)$ has the F distribution with parameters p, q.

(d) (xx briefly on the process $V(t) = \exp\{W(t)\}$. its multiplicative nature. xx).

Delta method for processes, order statistics processes

Ex. 9.37. *The delta method for processes.* We have developed and used versions of the delta method multiple times in earlier chapters; in its simplest form, if if $\sqrt{n}(X_n - a) \to_d Z$, then $\sqrt{n}\{g(X_n) - g(a)\} \to_d g'(a)Z$. Here are versions suited for empirical processes.

(a). Suppose $Z_n = \sqrt{n}(A_n - a) \to_d Z$, in $C[0,1]$, with $A_n(t)$ and $a(t)$ being a process and a function. Show that $\max_t |A_n(t) - a(t)| \to_{\rm pr} 0$. Assume g is a function with a smooth and bounded derivative on $[0,1]$. Show that the process Z_n^*, defined by $Z_n^*(t) = \sqrt{n}\{g(A_n(t)) - g(a(t))\}$, converges to Z^*, defined by $Z^*(t) = g'(a(t))Z(t)$. The same arguments should show that there is process convergence $Z_n^* \to_d Z^*$ in the function space $D[\varepsilon, 1-\varepsilon]$ in cases where g has a smooth and bounded derivative only on the subinterval $[\varepsilon, 1-\varepsilon]$.

(b). (xx one or two applications, before the order statistics thing. xx) let me try:

$$\sqrt{n}\{A(F_n(t)) - A(F(t))\} \to_d A'(F(t))W^0(F(t)).$$

with $A(u) = 2\arcsin(u^{1/2})$: variance is 1. covariance more complicated, but fully doable, for testing and bands.

Ex. 9.38. *The order statistics process.* (xx to be done, with a few applications. linear combinations of order statistics, or sample quantiles. briefly: $\sqrt{n}(\widehat{\mu}_p - \mu_p) \to_d c(p)W^0(p)$, with $c(p) = (F^{-1})'(p) = 1/f(F^{-1}(p))$. from this follow results for $\widehat{\theta} = \int_0^1 w_n(p)\widehat{\mu}_p \, {\rm d}p$, for $(1/m)\sum_{j=1}^m w(j/m)\widehat{\mu}_{j/m}$, etc. xx) Suppose $V_1, V_2, \cdots$ are i.i.d. with mean 1 and variance 1, and construct from these the bridge process, starting and ending at zero,

$$Z_n(t) = \sqrt{n}\Big(\sum_{i=1}^{[nt]} V_i / \sum_{i=1}^n V_i - t\Big) \quad \text{for } t \in [0,1].$$

(a). Studying first $A_n(t) = n^{-1/2}\sum_{i\le[nt]}(V_i - 1)$, show via Donsker's theorem that $B_n \to_d W^0$, the Brownian bridge, where $B_n(t) = A_n(t) - tA_n(1)$. (xx check if this is in place already. xx) Show that

$$Z_n(t) = \sqrt{n}\Big\{n^{-1}\sum_{i=1}^{[nt]}(V_i - 1) - tn^{-1}\sum_{i=1}^n (V_i - 1) + ([nt]/n - t)\Big\}/n^{-1}\sum_{i=1}^n V_i,$$

and use this to demonstrate that $Z_n \to_d W^0$.

(b). Consider $U_1, \ldots, U_n$ being i.i.d. from the uniform distribution, with order statistics $U_{(1)} < \cdots < U_{(n)}$. Use the representation of the $U_{([nq])}$ process from Ex. 3.17, in concert with results above, to demonstrate that the uniform order statistics process $Q_n(q) = \sqrt{n}(U_{([nq])} - q)$ converges to the Brownian bridge.

(c). With observations $Y_1, \ldots, Y_n$ from a smooth increasing c.d.f. F, with density f, consider the quantile process $R_n(q) = \sqrt{n}\{Y_{([nq])} - F^{-1}(q)\}$. Use Ex. 9.37 to show that $R_n \to_d R$, on each $D[\varepsilon, 1-\varepsilon]$, where $R(q) = c(q)W^0(q)$, and $c(q) = (F^{-1})'(q) = 1/f(F^{-1}(q))$.

(d). (xx clean this. xx) Show that with i growing large with n such that $q_i = i/(n+1)$ is bounded away from 0 and 1, the event

$$Y_{(i)} \in F^{-1}(q_i) \pm 1.96\, c(q_i)\{q_i(1-q_i)\}^{1/2}/\sqrt{n}$$

will have probability tending to 0.95. Argue in particular from this that a plot of $(F^{-1}(q_i), Y_{(i)})$ should be approximately linear. This, along with variations, is called a q-q or quantile-quantile plot.

(e)**.** In the situation just described, show that with a positive weight function $w(t)$, we have

$$H_n = \sqrt{n} \max_{\varepsilon \le q \le 1-\varepsilon} w(q)|Y_{([nq])} - F^{-1}(q)| \to_d H = \max_{\varepsilon \le t \le 1-\varepsilon} w(q)c(q)|W^0(q)|.$$

(f)**.** (xx an application or illustration. both the QQ plot thing with pointwise band, and functionals of the type $V_n = \int_0^1 w(q)\,\mathrm{d}R_n(q) \to_d V = \int_0^1 w(q)\,\mathrm{d}R(q)$, almost linear combinations of order statistics. can check for normality via QQ plots. see nils com5* for birthweights for nonsmoking and smoking mothers. xx)

Ex. 9.39. *The top ten percent do fifty percent of the total job.* Statements are sometimes made along the lines of 'the top ten percent in a workforce carry out fifty percent of the total job' (with variations). There are settings where this can be seen to hold. Suppose efforts of i.i.d. workers can be seen as $X_1, \ldots, X_n$ from some distribution on the halfline, with density f and c.d.f. F. Then consider the ratio $R_n = \sum_{i \ge [0.90\,n]} X_{(i)} / \sum_{i=1}^n X_i$, in terms of the order statistics $X_{(i)}$.

(a)**.** First assume that the X_i are from the uniform distribution on the unit interval. Show that $R_n \to_{\mathrm{pr}} 0.19$.

(b)**.** For the general case, show that $n^{-1} \sum_{i \ge [0.90\,n]} X_{(i)} \to_{\mathrm{pr}} \int_{0.90}^1 F^{-1}(u)\,\mathrm{d}u$, and that

$$R_n \to_{\mathrm{pr}} \rho(0.90) = A(0.90)/A(0), \quad \text{with } A(q) = \int_q^1 F^{-1}(u)\,\mathrm{d}u = \int_{F^{-1}(q)}^\infty x f(x)\,\mathrm{d}x.$$

(c)**.** Examine the Beta$(a, 1)$ distribution, with $F(x) = x^a$ and $F^{-1}(u) = u^{1/a}$. Find an expression for $\rho(0.90)$, and show that for large n, the top 10 percent in a sample amount to 50 percent of the total, for $a_0 = 0.179$. Carry out some simulations to illustrate this phenomenon.

(d)**.** On any given day the efforts of the n i.i.d. workers might not correspond precisely to the 10 percent do 50 percent formula. Find the limit distribution of $\sqrt{n}\{R_n - \rho(0.90)\}$.

Ex. 9.40. *Correlated Brownian motions.* Let X and Y be independent Brownian motions.

(a)**.** For a $\rho \in (-1, 1)$, define $Z(t) = \rho X(t) + (1 - \rho^2)^{1/2} Y(t)$. Show that Z is another Brownian motion, with constant correlation ρ between $X(t)$ and $Z(t)$.

(b)**.** For $\rho = 0.90$, display simulated paths of $(X, Z) = \{(X(t), Z(t)) \colon t \in [0, 1]\}$. Then do this for $\rho = -0.90$, and comment on what you see.

(c)**.** Make similar constructions also for Brownian bridges.

Monitoring bridges, likelihood bridges, running t tests

Ex. 9.41. *Bridges from partial-sum processes.* [xx a couple of things, to be polished and edited. xx]

(a). Let $U_1, U_2, \ldots$ be i.i.d. with mean zero and variance one. Consider the process $V_n(t) = n^{-1} \sum_{i \le [nt]} (U_i - \bar{U}_n)$, with the sample mean $\bar{U}_n = n^{-1} \sum_{i=1}^n U_i$, and note that it starts and ends at zero. Show indeed that $V_n \to_d W^0$. You may use (xx Ex. 9.2 check this xx).

(b). Now suppose $Y_1, \ldots, Y_n$ are i.i.d. with mean θ and standard deviation σ. Show that

$$H_n(t) = n^{-1/2} \sum_{i \le [nt]} (Y_i - \bar{Y}_n) \to_d \sigma W^0(t). \tag{9.2}$$

Show hence that both $H_n(t)/\sigma$ and $H_n(t)/\widehat{\sigma}$ have the same Brownian bridge limit process, where $\widehat{\sigma}$ is any consistent estimator for σ. Note that the θ parameter cancels out in the H_n construction.

(c). Consider the null hypothesis H_0 of constancy, that the distribution of the Y_i has not changed over time. Explain why the test which rejects the null if $\max_{0 \le t \le 1} |H_n(t)|/\widehat{\sigma} > 1.358$ has test level 0.05.

(d). Using $H_n/\widehat{\sigma} \to_d W^0$, valid under H_0, construct a couple of other 0.05 level test statistics.

Ex. 9.42. *Detecting non-constant behaviour.* [xx a couple of things, to be polished and edited. xx] Suppose $Y_1, \ldots, Y_n$ are independent, but with means $\theta_1, \ldots, \theta_n$ which do not necessarily stay constant over time. The monitoring process $H_n(t) = n^{-1/2} \sum_{i \le [nt]} (Y_i - \bar{Y})$ studied in Ex. 9.41 can be used to test the constancy hypothesis H_0 that $\theta_1 = \cdots = \theta_n$. To learn about its detection abililities outside H_0, suppose now that $\theta_i = \theta_0 + \delta m(i/n)/\sqrt{n}$, for some underlying function $m(s)$ on $[0,1]$.

(a). Show that the mean of $n^{-1/2} \sum_{i \le [nt]} Y_i$ tends to $\theta_0 + \delta \int_0^t m(s)\,\mathrm{d}s$, and that

$$\mathrm{E}\, H_n(t) \to \delta \int_0^t \{m(s) - \bar{m}\}\,\mathrm{d}s, \quad \text{writing } \bar{m} = \int_0^1 m(s)\,\mathrm{d}s.$$

(b). Now consider the particular case of there being a changepoint in the mean structure, with say $\theta_i = a$ for $i \le \tau$ and $\theta_i = b$ for $i \ge \tau + 1$, which we translate to $m(s)$ having values a and b on $[0, t_0]$ and $(t_0, 1]$, with $t_0 = \tau/n$. Find and display the limiting mean function for $H_n(t)$. Show in particular that it has a triangular shape, with the triangle pointing upwards if $a > b$ and downwards if $a < b$; also, the point where the mean function attains its maximum (for $a > b$) or minimum (for $a < b$) is precisely the changepoint position t_0.

(c). Consider the case where $Y_i = \theta_0 + \delta m(i/n)/\sqrt{n} + U_i$, where the U_i are i.i.d. with mean zero and standard deviation σ. Show that

$$H_n(t) = n^{-1/2} \sum_{i \le [nt]} (Y_i - \bar{Y}) \to_d \delta \int_0^t (m - \bar{m})\,\mathrm{d}s + \sigma W^0(t).$$

Ex. 9.43. *More bridges.* (xx nils ranting at the moment. wish to have things for $\widehat{\theta}_j - \widehat{\theta}_n$, to monitor if parameter has changed. some is in Hjort and Koning (2002), but a bit more direct here. xx)

(a). Consider again the setup of Ex. 9.41, for i.i.d. observations $Y_1, \ldots, Y_n$ with variance σ^2, We saw there that $H_n(t) = n^{-1/2} \sum_{i \le [nt]} (Y_i - \bar{Y}_n)$ tends to $\sigma W^0(t)$. Writing $\bar{Y}_j$ for

the mean of the first j data points, show that the value of H_n at position $t = j/n$ can be expressed as

$$H_n(t) = n^{-1/2} j(\bar{Y}_j - \bar{Y}_n) = n^{-1/2} j(1 - j/n)(\bar{Y}_j - \widetilde{Y}_j) = \sqrt{n} t(1-t)(\bar{Y}_j - \widetilde{Y}_j),$$

in which $\widetilde{Y}_j$ is the mean of the $n - j$ observations falling to the right of j. Show that the variance of $H_n(t)$ is $\sigma^2 t(1-t)$.

(b). These findings, where several natural constructions for checking constancy of means lead to Brownian bridges, can now be lifted to other settings, and specifically to classes of estimators. Suppose $\widehat{\theta}_j$ is an estimator of some θ, after j data points. Assume these have the property that $\widehat{\theta}_j - \theta_0 = (1/j)\sum_{i \le j} U_i + \varepsilon_j$, for i.i.d. U_i with zero mean and some variance τ^2, with $j^{1/2}\varepsilon_j \to_{\rm pr} 0$ as j increases. Construct the process H_n with value $H_{n,j} = n^{-1/2} j(\widehat{\theta}_j - \widehat{\theta}_n)$ at $t = j/n$, and then constant over cells $[j/n, (j+1)/n]$. Show that $H_n \to_d \tau W^0$, over each $D[\delta, 1]$, for small δ. (xx extra argument for what might take place in the beginning. xx)

(c). Explain that for any estimator $\widehat{\theta}_n$ having an influence function with finite variance τ^2, see Ex. 5.57, the monitoring process M_n with values $(1/\sqrt{n}) j(\widehat{\theta}_j - \widehat{\theta}_n)/\widehat{\tau}$ at $t = j/n$, will converge to the Brownian bridge. Here $\widehat{\tau}$ is some consistent estimator of τ. Departure from Brownian bridge behaviour of this bridge plot thus signals that the model has not remained constant.

(d). (xx find a dataset where the standard deviation σ has changed over time. plot $n^{-1/2} j(\widehat{\sigma}_j - \widehat{\sigma}_n)$. xx)

Ex. 9.44. *A running t test process to check left versus right.* Consider an i.i.d. sequence $Y_1, \ldots, Y_n$ with mean μ and standard deviation σ. For any $\tau \in \{1, \ldots, n\}$, divide data into a left part, $Y_1, \ldots, Y_\tau$, and a right part, $Y_{\tau+1}, \ldots, Y_n$, with estimates $\widehat{\mu}_L = (1/\tau)\sum_{i \le \tau} Y_i$ and $\widehat{\mu}_R = (1/(n-\tau))\sum_{i > \tau} Y_i$. The usual way of checking that the left and right parts actually have the same mean is by comparing $\widehat{\mu}_R - \widehat{\mu}_L$ to an estimate of its standard deviation, as with the classical t test, say

$$t_n(\tau) = \frac{\widehat{\mu}_L - \widehat{\mu}_R}{\{\widehat{\sigma}_L^2/\tau + \widehat{\sigma}_R^2/(n-\tau)\}^{1/2}}.$$

One may use separate standard deviation estimates for the left and right parts, or a common one. The idea now is to consider the full running t test process. Below we work under the null hypothesis assumption that the full sequence is i.i.d., but the point will later be to find changepoints, where the mean changes from one value to another one.

(a). With one test at the time, so to speak, show that $t_n(\tau) \to_d {\rm N}(0,1)$, as long as τ and $n - \tau$ both increase.

(b). Writing $s = \tau/n$, form the process

$$T_n(s) = \frac{\widehat{\mu}_L - \widehat{\mu}_R}{\{\widehat{\sigma}_L^2/\tau + \widehat{\sigma}_R^2/(n-\tau)\}^{1/2}} = \frac{\bar{Y}_{1,[ns]} - \bar{Y}_{[ns]+1,n}}{\{\widehat{\sigma}^2_{1,[ns]}/[ns] + \widehat{\sigma}^2_{[ns]+1,n}/(n-[ns])\}^{1/2}},$$

where subscripts L and R as above refer to sample parts $Y_1, \ldots, Y_{[ns]}$ and $Y_{[ns]+1}, \ldots, Y_n$. Show that $T_n(s) \to_d {\rm N}(0,1)$ for each $s \in (0,1)$.

(c). From the Donsker theorem, we know that the process $A_n(s) = n^{-1/2}\sum_{i\le[ns]}(Y_i-\mu)$ tends to $\sigma W(s)$. With $s=\tau/n$, show that

$$\sqrt{n}(\widehat{\mu}_L-\mu) = A_n(s)/s \quad \text{and} \quad \sqrt{n}(\widehat{\mu}_R-\mu) = \{A_n(1)-A_n(s)\}/(1-s).$$

(d). Show from this that

$$\sqrt{n}(\widehat{\mu}_L-\widehat{\mu}_R) = \frac{A_n(s)}{s} - \frac{A_n(1)-A_n(s)}{1-s} = \frac{A_n(s)-sA_n(1)}{s(1-s)}.$$

(e). Let $v_n(s)$ be a suitable weight function, deterministic or depending on the data, in such a way that there is a limit in probability function $v_n(s)\to_{\rm pr} v(s)$. Show that the following weighted difference test process converges in distribution:

$$Z_n(s) = \sqrt{n}v_n(s)(\widehat{\mu}_L-\widehat{\mu}_R) \to_d Z(s) = v(s)\sigma\frac{W(s)-sW(1)}{s(1-s)} = \sigma\frac{v(s)W^0(s)}{s(1-s)},$$

with the Brownian bridge $W^0(s)=W(s)-sW(1)$. Show also the there is full process convergence $Z_n\to_d Z$ in each $D[\varepsilon,1-\varepsilon]$, for $\varepsilon>0$, provided $\max_{\varepsilon\le s\le 1-\varepsilon}|v_n(s)-v(s)|\to_{\rm pr}0$, i.e. we need uniform convergence in probability of the weight function.

(f). For the special case of the running t test process $T_n(s)$, write out the weight function $v_n(s)$ indirectly used there, and show that $v_n(s)\to_p\{s(1-s)\}^{1/2}/\sigma$. Hence

$$T_n(s) = \frac{\widehat{\mu}_L-\widehat{\mu}_R}{\{\widehat{\sigma}_L^2/\tau+\widehat{\sigma}_R^2/(n-\tau)\}^{1/2}} \to_d T(s) = \frac{W^0(s)}{\sqrt{s(1-s)}},$$

the naturally normalised Brownian bridge, to make the variance constant and equal to one. This is a process we met in Ex. 9.10.

(g). Suppose $Y_i=\mu_i+\varepsilon_i$, with the ε_i being i.i.d., and that we are to test the the null hypothesis that all the means μ_i are equal, against the alternative that there is changepoint τ with one value μ_L for $i\le\tau$ and another value μ_R for $i\ge\tau+1$. A natural test statistic is

$$t_{\max} = \max_{\varepsilon\le s\le 1-\varepsilon}|T_n(s)| = \max_{c\le\tau\le d}|T_n(\tau)|,$$

with $c=[n\varepsilon]$ and $d=[n(1-\varepsilon)]$, for a suitable and perhaps small ε. Under the null hypothesis, show that $t_{\max}\to_d\max_{\varepsilon\le s\le 1-\varepsilon}|T(s)|$, in the notation above. Exemplify this by drawing $y_i\sim{\rm N}(a,1)$ for $i=1,\ldots,75$ and $y_i\sim{\rm N}(b,1)$ for $i=76,\ldots,100$, for a,b of your own choice. Draw the $t_n(\tau)$ curve, and test the null, using $[\varepsilon,1-\varepsilon]=[0.05,0.95]$. This calls for separate simulation of $\max_{0.05\le s\le 0.95}|T(s)|$; see Ex. 9.10.

(h). For some applications we encounter processes of the above type, with running tests based on $\widehat{\mu}_R-\widehat{\mu}_L$, but based on observations with different variances. This happens e.g. in Story ii.3, where we compare word lengths for a sequence of chapters of different sizes. Consider therefore Y_i with mean μ_i and variance σ^2/m_i, with left-and-right averages $\widehat{\mu}_L=\sum_{i\le\tau}m_iY_i/\sum_{i\le\tau}m_i$ and $\widehat{\mu}_R=\sum_{i\ge\tau+1}m_iY_i/\sum_{i\ge\tau+1}m_i$. As above, write $s=\tau/n$, with τ growing with n. The natural test ratio to monitor is then

$$T_n(s) = \frac{\widehat{\mu}_L-\widehat{\mu}_R}{(\widehat{\sigma}_L^2/\sum_{i\le\tau}m_i+\widehat{\sigma}_R^2/\sum_{i\ge\tau+1}m_i)^{1/2}},$$

where $\widehat{\sigma}_L^2 = \tau^{-1}\sum_{i\le\tau} m_i(Y_i-\widehat{\mu}_L)^2$ and $\widehat{\sigma}_R^2 = (n-\tau)^{-1}\sum_{i\ge\tau+1} m_i(Y_i-\widehat{\mu}_R)^2$. The relevant empirical process to examine is

$$A_n(s) = (\sum_{i=1}^n m_i)^{-1/2}\sum_{i\le ns} m_i(Y_i-\mu) \quad \text{for } s\in[0,1].$$

For technical reasons we assume the background sample sizes $m_1,\ldots,m_n$ are well-behaved, with $\bar{m}_n = n^{-1}\sum_{i=1}^n m_i$ tending to some m_0. Show via Ex. 9.46 that $A_n \to_d \sigma W$, and that arguments as above extend to the result $T_n(s) \to_d \sigma W^0(s)/\{s(1-s)\}^{1/2}$. Explain how this leads to overall running t tests type statistics.

(i). (xx pointers, more. to Cunen et al. (2020a) and to Story ii.3. xx)

Ex. 9.45. *Generalisations and uses of the running t test process.* [xx something here. generalised version, valid for any estimand. make active use of influence functions from Ex. 5.57. xx] The apparatus of Ex. 9.44 can be used to test the homogeneity hypothesis against the existence of one of more changepoints *for the mean*. The machinery extends nicely and fruitfully to other parameters, as developed here. Consider in general terms $Y_1,\ldots,Y_n$ being i.i.d. from some distribution F, with $\theta=\theta(F)$ a parameter of interest, like the mean, the standard deviation, a quantile, or a threshold probability. We also assume that there is a well-defined estimator $\widehat{\theta}_{a,b}$ for θ based on the data stretch $Y_a,\ldots,Y_b$. We work under the null hypothesis here, but the aim is to build a running t test process for monitoring the θ_i associated with Y_i. Such a process should lead to tests to pick up that the θ_i might have changed.

(a). We assume that the estimand θ and its estimator admit an influence function $\mathrm{IF}(F,y)$, with finite variance, as per Ex. 5.57. Show that

$$\widehat{\theta}_{a,b} = \theta + \frac{1}{b-a+1}\sum_{i=a}^b \mathrm{IF}(F,Y_i) + \varepsilon_{a,b},$$

with the remainder term small enough to have $(b-a+1)^{1/2}\varepsilon_{a,b} \to_{\mathrm{pr}} 0$ when the data stretch length $b-a+1$ increases. Show that $(b-a+1)^{1/2}(\widehat{\theta}_{a,b}-\theta) \to_d \mathrm{N}(0,\kappa^2)$, where $\kappa^2 = \mathrm{Var}_F\,\mathrm{IF}(F,Y)$.

(b). Now construct, for each τ, the t type test statistic

$$t_n(\tau) = \frac{\widehat{\theta}_L - \widehat{\theta}_R}{\{\widehat{\kappa}_L^2/\tau + \widehat{\kappa}_R^2/(n-\tau)\}^{1/2}},$$

with $\widehat{\theta}_L$ and $\widehat{\kappa}_L$ defined for the left data $Y_1,\ldots,Y_\tau$ and similarly $\widehat{\theta}_R$ and $\widehat{\kappa}_R$ for the right data $Y_{\tau+1},\ldots,Y_n$. [xx then the point is to find the same limit as for Ex. 9.44. xx]

(c) (xx illustration: normal data, with shift in μ and/or σ. xx).

(d) (xx illustration: following the median over time. xx).

(e).

Ex. 9.46. *A Lindeberg type extension of the Donsker theorem.* Often enough one encounters partial-sum processes where the summands are not i.i.d. Assume that $U_1,U_2,\ldots$ are independent with zero means, but perhaps with different distributions. With $\sigma_1,\sigma_2,\ldots$ their standard deviations, we suppose these are sufficiently regular to have $n^{-1}\sum_{i=1}^n \sigma_i^2 \to \kappa^2$, say. Now build the process $X_n(t) = n^{-1/2}\sum_{i\le[nt]} U_i$ for $t\ge 0$.

(a). Show that for each t, $X_n(t) \to_d \mathrm{N}(0, \kappa^2 t)$. Under the Lindeberg type condition that $n^{-1}\sum_{i=1}^n \mathrm{E}\, U_i^2 I(|U_i| \geq \sqrt{n}\varepsilon) \to 0$ for each positive ε, show that $X_n \to_d \kappa W$.

(b). (xx give an example. point to monitoring bridges for regression models. and perhaps generalise to weaker conditions. perhaps the most natural bridge uses $m_n(t) = \sum_{i\leq [nt]} \kappa_i^2 / \sum_{i=1}^n \kappa_i^2$, which tends to t, and then the bridge constriction $Z_n(t) = X_n(t) - m_n(t)X_n(1)$, tending to $Z(t) = X(t) - tX(1) = \kappa W^0(t)$. xx)

Ex. 9.47. *A Donsker type theorem for regression coefficients.* Consider a regression type sequence (x_i, y_i), where the x_i are increasing, as in a time series context, and where it is assumed that $y_i = bx_i + \varepsilon_i$, with the ε_i being i.i.d. with zero mean and variance σ^2. We may then estimate the regression coefficient based on the first j datapoints, as $\widehat{b}_j = \sum_{i\leq j} x_i y_i / M_j$, with $M_j = \sum_{i\leq j} x_i^2$, and these may be monitored over time. We assume that the x_i are well-behaved, with (i) $n^{-1}M_n$ tending to some positive m_0 for growing n and (ii) $\max_{i\leq n} x^2/M_n \to 0$.

(a). Show first that $n^{-1}M_{[nt]} \to m_0 t$. Consider then the process $A_n(t) = n^{-1/2}\sum_{i\leq [nt]} x_i\varepsilon_i$. Explain that it also be written

$$n^{-1/2}\sum_{i\leq [nt]} x_i(y_i - bx_i) = n^{-1}M_{[nt]}\sqrt{n}(\widehat{b}_{[nt]} - b).$$

Show that variance of $A_n(t)$ tends to $m_0\sigma^2 t$, and that in fact $A_n \to_d m_0^{1/2}\sigma W$, a scaled Brownian motion. You may use results from Ex. 2.41.

(b). We then generalise to the more natural setup with a full two-parameter regression line, with $y_i = a + bx_i + \varepsilon_i$, and again with the ε_i being i.i.d. with zero mean and variance σ^2. Writing $\bar{x}_j = j^{-1}\sum_{i\leq j} x_i$, we assume that $M_n = \sum_{i=1}^n (x - \bar{x}_n)^2$ tends to some m_0, and that $\max_{i\leq n}(x_i - \bar{x}_n)^2/M_n \to 0$. Show that

$$B_n(t) = n^{-1/2}\sum_{i\leq [nt]} (x_i - \bar{x}_{[nt]})\varepsilon_i \to_d m_0^{1/2}\sigma W(t),$$

again a scaled Brownian motion. With $\widehat{b}_j = \sum_{\leq j}(x_i - \bar{x}_j)y_i$ estimating b, show also that $B_n(t) = n^{-1}M_{[nt]}\sqrt{n}(\widehat{b}_{[nt]} - b)$.

(c). Explain that the methods and results of Ex. 9.44–9.45 may be extended to the present case, monitoring the regression estimators $\widehat{b}_j$ over time. For given τ, divide data into left, with $i \leq \tau$, and right, with $i \geq \tau + 1$, and ensuing $\widehat{b}_L$ and $\widehat{b}_R$, with variances $\sigma^2\kappa_L^2/\tau$ and $\sigma^2\kappa_R^2/(n-\tau)$, writing $\kappa_L^2 = \sum_{i\leq\tau}(x_i - \bar{x}_L)^2$ and $\kappa_R^2 = \sum_{i\geq\tau+1}(x_i - \bar{x}_R)^2$, in which $\bar{x}_L$ and $\bar{x}_R$ are the averages to the left and right. Then define the monitoring process

$$T_n(\tau) = \frac{\widehat{b}_R - \widehat{b}_L}{\widehat{\sigma}\{\kappa_L^2/\tau + \kappa_R^2/(n-\tau)\}^{1/2}},$$

using $\widehat{\sigma}^2 = (\tau/n)\widehat{\sigma}_L^2 + (1 - \tau/n)\widehat{\sigma}_R^2$. Explain that with $\tau/n \to s$, we have $T_n(\tau) \to_d T(s) = W^0(s)/\{s(1-s)\}^{1/2}$, and that $T_{n,\max} = \max_{c\leq\tau/n\leq d}|T_n(\tau)|$ tends in distribution to $T_{\max} = \max_{c\leq s\leq d}|T(s)|$. Quantiles for this distribution, for some intervals $[c, d]$, were found in Ex. 9.10. (xx point to Oslo temperatures story. xx)

Ex. 9.48. *A log-likelihood maxima bridge.* (xx the following needs a bit of cleaning. pointer to use of poisson monitoring bridge for Story iii.2. xx) Suppose i.i.d. data $Y_1, \ldots, Y_n$ arise in orderly sequential fashion, and consider a parametric distribution $f(y, \theta)$ for these, of dimension say p. Here we build a log-likelihood-maxima bridge, which may be used to check whether something unusual has taken place in the course of time points $1, \ldots, n$.

(a). Start out considering $\ell_j = \sum_{i \le j} \log f(Y_i, \theta_0)$, with θ_0 the true parameter, along with the attained maximum $\widehat{\ell}_j = \sum_{i \le j} \log f(Y_i, \widehat{\theta}_j)$, where $\widehat{\theta}_j$ is the ML estimator after j datapoints. Show that ℓ_j/j as well as $\widehat{\ell}_j/j$ tends to $\xi = \mathrm{E} \log f(Y_i, \theta_0)$. Plotting $\widehat{\ell}_j$ as a function of j should accordingly be approximately a straight line, with derivative ξ.

(b). Use Ex. 9.41 to show that the process $A_n(t) = n^{-1/2}\{\ell_{[nt]} - ([nt]/n)\ell_n\}$, with values $A_{n,j} = n^{-1/2}\{\ell_j - (j/n)\ell_n\}$ at j/n, tends to $\kappa W^0(t)$, with $\kappa^2 = \mathrm{Var} \log f(Y_i, \theta_0)$.

(c). From the log-likelihood maxima construct the process B_n with values $B_{n,j} = n^{-1/2}\{\widehat{\ell}_j - (j/n)\widehat{\ell}_n\}$ at j/n. The $B_n(t)$ process needs at least p data points to get started, so technically its time domain is $[p/n, 1]$, and if pedantry demands can be put to zero in the small interval before p/n. Show that the $\widehat{\ell}_j - \ell_j$ differences are bounded in probability (xx point to Wilks Ch5 and chi-squared things here xx). Use this to demonstrate that B_n and A_n are close enough for them to have the same process limit, i.e. κW^0.

(d). Show how this leads to model monitoring bridge plots

$$M_{n,j} = B_{n,j}/\widehat{\kappa} = (1/\sqrt{n})\{\widehat{\ell}_j - (j/n)\widehat{\ell}_n\}/\widehat{\kappa} \quad \text{for } j = p, \ldots, n,$$

which under nothing-has-changed model conditions will behave as a Brownian bridge. In particular, $\max_{p \le j \le n} M_{n,j} \to_d M = \max_{0 \le s \le 1} |W^0(s)|$, with 1.358 as upper 0.05 point, etc., see Ex. 9.24. Here $\widehat{\kappa}$ is an estimator for κ, computed either under model conditions, i.e. as the appropriate $\kappa(\widehat{\theta})$, or as the empirical standard deviation of the observed $\log f(y_i, \widehat{\theta})$.

(e). (xx possible bias correction: argue via Wilks theorems of Ch5 that $\widehat{\ell}_j - \ell_j = \frac{1}{2}W_j + \varepsilon_j$, say, with W_j tending to a χ^2_p. this invites the bias correction $\widehat{\ell}_j - \frac{1}{2}p$, to come closer to ℓ_j when constructing the bridge. show that this leads to the modified bridge plot

$$M'_{n,j} = B'_{n,j}/\widehat{\kappa} = (1/\sqrt{n})\{\widehat{\ell}_j - (j/n)\widehat{\ell}_n - \tfrac{1}{2}p(1 - j/n)\}.$$

verify that it has the same Brownian bridge limit; the intention is to come closer to this limit. xx)

(f). The ℓ_j and $\widehat{\ell}_j$ grow linearly over time, so it may be convenient to compute and present the plots in terms of $H_j = \widehat{\ell}_j/j$. Show that $M_{n,j} = n^{-1/2} j(H_j - H_n)/\widehat{\kappa}$.

Ex. 9.49. *Likelihood monitoring bridges for some models.* Here we use the general setup of Ex. 9.48 to set up clear likelihood monitoring bridges for a few parametric models.

(a). For the case of normal $\mathrm{N}(\xi, \sigma^2)$ data, show that $M_{n,j} = n^{-1/2} j(\log \widehat{\sigma}_n - \log \widehat{\sigma}_j)/\widehat{\kappa}$, where $\widehat{\sigma}_j^2 = (1/j)\sum_{i \le j}(y_i - \bar{y}_j)^2$ is the empirical variance after j data points. Also, show that $\kappa = 1/\sqrt{2}$, under model conditions, but we may also take $\widehat{\kappa}$ to be the empirical standard deviation for the observed $\frac{1}{2}(y_i - \widehat{\xi})^2/\widehat{\sigma}^2$.

(b). For the Pois(θ) model, show that $M_{n,j} = n^{-1/2} j\{H(\bar{y}_j) - H(\bar{y}_n)\}/\widehat{\kappa}$, with $\bar{y}_j$ the mean of the first j data points and $H(u) = u \log u - u$. Also, $\widehat{\kappa} = (\bar{y}_n)^{1/2} \log \bar{y}_n$, under model conditions, or one may use the model robust $\widehat{\kappa} = \widehat{\sigma}_n \log \bar{y}_n$, with $\widehat{\sigma}_n$ the standard deviation of the data.

(c). You observe Bernoulli 0-1 outcomes $Y_1, \ldots, Y_n$ and wonder whether the underlying probability p has remained unchanged during the process. With $K(p) = p \log p + (1-p) \log(1-p)$, show that the monitoring bridge takes the form $M_{n,j} = n^{-1/2} j\{K(\widehat{p}_j) - K(\widehat{p}_n)\}/\widehat{\kappa}$, with $\widehat{p}_j$ the estimate after j data points, and $\widehat{\kappa}^2$ the estimate of $\kappa^2 = p(1-p)[\log\{p/(1-p)\}]^2$. Set up a simple simulation experiment, to check out how this works: take $n = 200$, with p equal to 0.33 for the first 100 and 0.44 for the next hundred. See if the $M_{n,j}$ plot picks up that p has not remained constant.

(d). Consider a trinomial process, with $(Y_{i,1}, Y_{i,2}, Y_{i,3})$ being $(1,0,0), (0,1,0), (0,0,1)$ in experiment i, for cases A, B, C, say, with probabilities p_1, p_2, p_3 summing to one. Show that the likelihood monitoring plot takes the form

$$M_{n,j} = n^{-1/2} j\{K(\widehat{p}_{1,j}, \widehat{p}_{2,j}, \widehat{p}_{3,j}) - K(\widehat{p}_{1,j}, \widehat{p}_{2,j}, \widehat{p}_{3,j})\}/\widehat{\kappa},$$

with $(\widehat{p}_{1,j}, \widehat{p}_{2,j}, \widehat{p}_{3,j})$ the estimated probabilities after j experiments, and $K(p_1, p_2, p_3) = p_1 \log p_1 + p_2 \log p_2 + p_3 \log p_3$. Also, show that the required $\widehat{\kappa}$ should be the estimate of $\kappa^2 = \sum_j p_j (\log p_j)^2 - (\sum_j p_j \log p_j)^2$. Set up a simple simulation experiment to see how the plots look and to see if they can pick up a modest change of (p_1, p_2, p_3).

(e). (xx then one more. find a dataset to work through, where the pedagogical point would be to demonstrate that these monitoring plots are fine to set up and use just using the numerical tools from ML estimation, without clear formulae. something with history, or emil's basketball data over time. point to Stories iv.5 and ii.3. xx)

Ex. 9.50. *Likelihood monitoring bridges for regression models.* (xx point here is to generalise the above monitoring plots to regression setups in the right way. at the moment nils ranting away. see nils note; sum of two Brownian bridges in the limit. xx) consider a model $f(y_i \,|\, x_i, \theta_0)$, with mean ξ_i and variance κ_i^2 for $\log f(Y_i \,|\, x_i, \theta_0)$. then build the partial-sum process $A_n(t)$, taking values

$$n^{-1/2}(\ell_j - a_j) = n^{-1/2} \sum_{i \le j} \{\log f(y_i \,|\, x_i, \theta_0) - \xi_i\}$$

at positions $t = j/n$, with $a_j = \sum_{j \le i} \xi_j$. the variance is $n^{-1} \sum_{i \le [nt]} \kappa_i^2$, tending to say $\kappa^2 t$, under ergodic assumptions. by the Lindeberg extension of the Donsker theorem, see Ex. 9.46. $A_n(t) \to_d W(\kappa^2 t)$. from this the bridge construction, but careful. $B_n(t) = A_n(t) - m_n(t) A_n(1)$.

(a). (xx check nils note; sum of two Brownian bridges. xx)

(b). (xx but this is going to be super useful for time series models, where the basic construction still works, though needing more arguments for the basic Brownian bridge lemma. xx)

Ex. 9.51. *Monitoring processes for constancy or parameters over time.* [xx the basics from Hjort and Koning (2002). processes tending to Brownian bridges. application to goodness of fit, and to testing the hypothesis of θ_i having remained unchanged over time. xx]

Ex. 9.52. *An extension of the Donsker theorem with positive kurtosis.* [xx to come later. nils construction, perhaps touched in chapter 2, with

$$X_n(t) = \sum_{i \le [nt]} J_i (i/n)^\alpha U_i, \quad \text{for } t \ge 0,$$

with the U_i being i.i.d. with mean zero and variance one, and the J_i being an independent sequence of independent Bernoulli variables, with $\Pr(J_i = 1) = 1/i$. There is $X_n \to_d X$ process convergencce, with extra kurtosis. xx]

Ex. 9.53. *Criteria for existence of processes in $D[0,1]$.* Suppose we set up a system of finite-dimensional distributions for a stochastic process $X = \{X(t)\colon t \in [0,1]\}$ on the unit interval. These need to have or imply a minimum of regularity in order to ensure that yes, such an X process really exists. Let $\mu_{t_1,\ldots,t_k}$ be the probability measure for the envisaged $(X(t_1),\ldots,X(t_k))$. First, these need to be coherently defined, in the Kolmogorov sense: when integrating out the third element in $\mu_{s,t,u}$ we should get $\mu_{s,t}$, etc., and if we for say $s < t < u < v$ start out with the assigned distribution $\mu_{s,t,u,v}$, and then compute the distribution at (v,u,t,s), we ought to get the specified $\mu_{v,u,t,s}$, etc. Suppose

$$\mu_{t,t+h}\{(u_1,u_2)\colon |u_1-u_2| \ge \varepsilon\} \to 0 \quad \text{as } h \to 0, \tag{9.3}$$

and that we for all $t_1 \le t \le t_2$ have

$$\int |u-u_1|^\gamma \, |u_2-u|^\gamma \, \mathrm{d}\mu_{t_1,t,t_2}(u_1,u,u_2) \le |G(t_2)-G(t_1)|^\alpha, \tag{9.4}$$

for some $\gamma > 0$ and $\alpha > 1$, where G is a nondecreasing and continuous function on $[0,1]$. Then Theorem 15.7 in Billingsley (1968) secures the existence of an X in $D[0,1]$ having these $\mu_{t_1,\ldots,t_k}$ as its finite-dimensional distributions.

(a). Suppose that such an $X = \{X(t)\colon t \in [0,1]\}$ really exists. Show that the left hand sides of (9.3) and (9.4) then can be written in the more familiar ways $\Pr(|X(t+h) - X(t)| \ge \varepsilon)$ and $\mathrm{E}\,|X(t)-X(t_1)|^\gamma \, |X(t_2)-X(t)|^\gamma$. The point is that we in principle do not yet know that there is such an X in $D[0,1]$, hence the somewhat pedantic notation.

(b). Specify finite-dimensional distributions by saying that for each t, $\mu_t \sim \mathrm{Pois}(\lambda t)$, that for $s < t$ we have $\mu_{s,t}$ equal to the distribution of $(A, A+B)$ with $A \sim \mathrm{Pois}(\lambda s)$ and $B \sim \mathrm{Pois}(\lambda(t-s))$ independent, and so on, with independent Poisson increments for any $t_1 < \cdots < t_k$. Show that conditions (9.3) and (9.4) hold, actually with $\gamma = 1$ and $\alpha = 2$. This ensures *the existence of a Poisson process* in $D[0,1]$, with independent increments and $X(t) - X(s) \sim \mathrm{Pois}(\lambda(t-s))$ for any $s < t$.

(c). Generalise the above to demonstrate the existence in $D[0,1]$ of any time-inhomogeneous Poisson process, with mean function $\mathrm{E}\,X(t) = \Lambda(t) = \int_0^t \lambda(s)\,\mathrm{d}s$, where $\lambda(s)$ is any non-negative integrable intensity function. The standard Poisson process then corresponds to this intensity function being a constant λ.

(d). Use the criteria above to give another proof of the existence of the Brownian motion, as a process W in $D[0,1]$. Here we already know that $\Pr(W \in C[0,1]) = 1$.

(e). Show the existence of a *Gamma process* in $D[0,1]$, defined via independent increments and $Z(t)-Z(s) \sim \mathrm{Gam}(a(t-s), b)$, where a, b are given positive parameters. Show that its mean and variance are $(a/b)t$ and $(a/b^2)t$.

(f). Let $X = \{X(t)\colon t \in [0,1]\}$ be a simple stationary Markov process with values in $\{0,1\}$, starting in 0, and with transition probabilities $p_{a,b}(t) = \{X(s+t) = b \,|\, X(s) = a\}$ for $a, b \in \{0,1\}$. We have $\sum_b p_{a,b}(t) = 1$ and $p_{a,b}(s+t) = \sum_c p_{a,c}(s)p_{c,b}(t)$ for all a, b and s, t. Show that X is really a random process in $D[0,1]$. [xx perhaps a helpful hint here. xx] Generalise to any time-continuous Markov chain.

(g). xx make the point, via a simple exercise, that another way to prove existence in $D[0,1]$ of some X is to derive it as a limit, $X_n \to_d X$, with $X_n \in D[0,1]$, but this needs convergence criteria, to which we turn now. the point, from $\limsup P_n(F) \le P(F)$, for any closed F, is that if all $X_n \in F$, then certainly also $X \in F$. xx

Ex. 9.54. *Convergence in distribution, for random indexes.* (xx to be worked through and polished. now placed here in Ch9. point to where used in Ch12. Sorting out such themes is useful when studying Markov chains from different perspectives. nils is ranting on a bit so far. xx) Suppose $Z_n = \sqrt{n}(A_n - a) \to_d Z_0$, say, and that N_n is a sequence of random indexes, growing to infinity. Will it be the case that $Z_{N_n} = N_n^{1/2}(A_{N_n} - a)$ tends to the same limit?

(a). To realise that the statement above cannot be taken for granted, without further conditions, consider $Y_1, Y_2, \ldots$ being independent standard normal, for which of course $Z_n = \sqrt{n}\bar{Y}_n \to_d \mathrm{N}(0,1)$. Show that there will be infinitely many n at which Z_n is positive, and let $N_1 < N_2 < \cdots$ be these time points. Show that the N_n sequence increases without bounds. Clearly Z_{N_n} does not have the normal limit; the CLT does not work for this increasing subsequence.

(b). (xx here we are tung paa labben inntil videre. check Billingsley (1968, Section 17). xx) Suppose $Z_n = \sqrt{n}(A_n - a) \to_d Z_0$, and that a sequence of random indexes N_n satisfies $N_n/n \to_{\mathrm{pr}} c$, with $c \in (0,1)$. Define the random functions $Z_n(t) = \sqrt{n}(A_{[nt]} - a)$ and $V_n(t) = N_n t/n$, in the function space $D[\varepsilon, 1]$, with some fixed $\varepsilon \in (0, c)$. Show that $V_n \to_d V$, where V is the deterministic function $V(t) = ct$, and that $Z_n \to_d Z$, where $Z(t) = Z_0/t^{1/2}$.

(c). Show that $(Z_n, V_n) \to_d (Z, V)$, in the product function space $D[\varepsilon, 1]^2$, and that this limit is in the subspace $C[\varepsilon, 1]^2$. (xx i hope this is correct. xx) Deduce from this that $Z_n \circ V_n \to_d Z \circ V$, i.e. that there is process convergence $Z_n(V_n(t)) \to_d Z(V(t))$. In particular, show from this that

$$Z_n(V_n(1)) = \sqrt{n}(A_{N_n} - a) \to_d Z(V(1)) = Z_0/c^{1/2},$$

and that $N_n^{1/2}(A_{N_n} - a) \to_d Z_0$ follows from this. Show in particular that $A_{N_n} \to_{\mathrm{pr}} a$.

(d). (xx give one or two easy applications. xx)

More material and leftovers (as of 12-March-2025)

Ex. 9.55. *Testing exponentiality with Kolmogorov–Smirnov.* (xx to be polished, perhaps to be seen as a consequence of Ex. 9.28. xx) Consider nonnegative i.i.d. data $Y_1, \ldots, Y_n$. It is clear how one can use the Kolmogorov–Smirnov test of Ex. 9.25 to test whether the underlying distribution is some given $\mathrm{Expo}(\theta_0)$, say. Testing exponentiality in general involves one more level of uncertainty, however, related to estimating the parameter too.

(a). For $Y_i \sim \text{Expo}(\theta)$, write down the log-likelihood function, and establish that the ML estimator $\widehat{\theta} = 1/\bar{Y}$. With $\theta Y_i = V_i$, in terms of V_i being i.i.d. unit exponential, show that $\widehat{\theta} Y_i = V_i/\bar{V}$. These have the same distribution, but with dependence. Show that actually $\widehat{\theta} Y_i \sim (1/n)\text{Beta}(1, n-1)$, with c.d.f. $1-(1-t/n)^n$, going to $G_0(t) = 1-\exp(-t)$, the c.d.f. for a unit exponential.

(b). Consider the empirical c.d.f. for the $\widehat{\theta} Y_i$, i.e. $H_n(t) = n^{-1}\sum_{i=1}^n I(\widehat{\theta} Y_i \le t)$. Show that $H_n(t) \to_{\text{pr}} G_0(t)$. With $Y_{(1)} < \cdots < Y_{(n)}$ the order statistic, show that $H_n(\widehat{\theta} Y_{(i)}) = i/n$, leading to the natural test $K_n = \sqrt{n}\max_{i\le n} |i/n - G_0(\widehat{\theta} y_{(i)})|$. Explain that the null distribution of this K_n, i.e. under exponentiality, is well-defined, and does not depend on the θ. We may in particular simulate this distribution, for any given n, to test if a given sample is exponential or not; for an application, see Story iii.5.

(c). Is is fruitful to learn how to find the limit distribution of such a K_n. Let θ_0 be true value for the exponential parameter, and consider the empirical c.d.f. G_n for the Y_i. Explain that we have

$$\begin{pmatrix} A_n(t) \\ B_n \end{pmatrix} = \begin{pmatrix} \sqrt{n}\{G_n(t) - G_0(\theta_0 t)\} \\ \sqrt{n}(\bar{Y} - 1/\theta_0) \end{pmatrix} \to_d \begin{pmatrix} A(t) \\ B \end{pmatrix},$$

a jointly normal limit, with $A(t) = W^0(G(\theta_0 t))$ and

$$\text{cov}(\{A(t), B\} = \text{cov}\{I(Y_i \le t), Y_i\} = (1/\theta_0)\{\int_0^{\theta_0 t} u g_0(u)\,\text{d}u - G_0(u)\}.$$

Also, $\text{Var}\, B = \theta_0^2$. Deduce from this that

$$\begin{aligned} Z_n(t) &= \sqrt{n}\{G_n(t) - G_0(\widehat{\theta} t)\} \\ &= A_n(t) - g_0(\theta_0 t) t\sqrt{n}(\widehat{\theta} - \theta_0) + o_{\text{pr}}(1) \to_d A(t) + g_0(\theta_0 t)t\theta_0^2 B. \end{aligned}$$

(d). Show from this that $K_n \to_d \max_{t\ge 0} |Z(t)|$.

(e). (xx bit more. from $G_n(t) - G_0(\widehat{\theta} t)$ to $H_n(t) - G_0(t)$, same maxima. general case pointed to. using ML representation from Ch. 5 to reach

$$\begin{aligned} Z_n(t) &= \sqrt{n}\{F_n(t) - F(t, \widehat{\theta})\} \\ &= \sqrt{n}\{F_n(t) - F(t, \theta_0)\} - F^*(t, \theta_0)^{\text{t}} J^{-1} n^{-1/2} \sum_{i=1}^n u(Y_i, \theta_0) + o_{\text{pr}}(1) \\ &\to_d W^0(F(t, \theta_0)) - F^*(t, \theta_0)^{\text{t}} J^{-1} U, \end{aligned}$$

with $F^*(t, \theta) = \partial F(t, \theta)/\partial\theta$, and where we can write down covariance structure. this leads to $K_n = \sqrt{n}\max_{i\le n} |i/n - F(y_{(i)}, \widehat{\theta})| \to_d \max_t |Z(t)|$. some durbin some hjort. xx)

Ex. 9.56. *The maximum and the maximum absolute value of Brownian motion.* [xx some but not all details; good to have results written down. xx] Let W be a standard Brownian motion on $[0, 1]$. We shall consider the random variables

$$H = \max_{0\le t\le 1} W(t) \quad \text{and} \quad M = \max_{0\le t\le 1} |W(t)|.$$

(a) (xx some details, via random walk, and limit. xx). Show that $\Pr(H \le x) = 2\Phi(x) - 1$ for $x \ge 0$, which also means that $H \sim |N(0, 1)|$.

(b). Explain why this also gives results for $H_0 = \min_{0\le t\le 1} W(t)$.

(c). The distribution of M takes other forces. The result is

$$\Pr(M \le x) = \sum_{j=-\infty}^{\infty} (-1)^j \Pr((2j-1)x < N < (2j+1)x),$$

with N a standard normal. Show that this may also be written

$$\Pr(M \le x) = 1 - \frac{4}{\pi}\sum_{j=1}^{\infty} \frac{(-1)^j}{2j+1} \exp\{-\pi^2(2j+1)^2/(8x^2)\}.$$

[xx check more here, and see Billingsley (1968, Section 11). the point would also be to find expression for mean and variance. xx]

(d). (xx explore and polish. find the mean of M. i think it is $\mathrm{E}\, M = 2G$, with G the Catalan constant. at any rate we can simulate. xx)

Ex. 9.57. *The maximum and the maximum absolute value of a Brownian bridge.* (xx testing a bit: xref to Chapters 3, 4, 2, 5, 7, 8, 13. xx)

(a). (xx don't know yet how we most readily can ask for this, but would be nice to have on board, without having to point to a reference. Billingsley (1968, Section 13) isn't entirely good here. xx) We saw in Ex. 9.24 that $D_n \to_d D = \max_t |W^0(t)|$. Show know that

$$\Pr(D \le x) = 1 - 2\sum_{j=1}^{\infty} (-1)^{j+1} \exp(-2j^2x^2) \quad \text{for } x \ge 0.$$

(xx a bit more. point to Kolmogorov (1933b). compute it, and find 0.50, 0.90, 0.95, 0.99 quantiles. answers: 0.8276, 1.2238, 1.3580, 1.6271. see com23a of nilswork21. xx)

(b). Show that $\mathrm{E}\, D = \frac12\sqrt{2\pi}\log 2 = 0.8687$, and find also its variance. (xx is it here Catalan's constant comes in? brief lines for solution:

$$\begin{aligned}\mathrm{E}\, D &= \int_0^\infty \{1 - F(x)\}\,\mathrm{d}x = 2\sum_{j=1}^{\infty}(-1)^{j+1}\int_0^\infty \exp(-2j^2x^2)\,\mathrm{d}x \\ &= \tfrac12\sqrt{2\pi}\sum_{j=1}^{\infty}(-1)^{j+1}(1/j),\end{aligned}$$

$$\begin{aligned}\mathrm{E}\, D^2 &= \int_0^\infty \{1 - F(\sqrt{x})\}\,\mathrm{d}x = 2\sum_{j=1}^{\infty}(-1)^{j+1}\int_0^\infty \exp(-2j^2x)\,\mathrm{d}x \\ &= \sum_{j=1}^{\infty}(-1)^{j+1}(1/j^2) = \pi^2/12.\end{aligned}$$

with a bit of care. xx)

(c).

Ex. 9.58. *Arcsine laws for Brownian motion.* (xx will see during the opploepsside whether we include this or not. xx) For Brownian motion W on the unit interval, consider the maximum position $T = \operatorname{argmax} W$, the last crossing point $U = \max\{t \in [0,1]\colon W(t) = 0\}$, and the time above zero $V = \int_0^1 I(W(t) > 0)\,\mathrm{d}t$. We shall see that these three have the same arcsine distribution, with c.d.f. $\Pr(Z \le z) = (2/\pi)\arcsin(z^{1/2})$ and U-shaped density $(1/\pi)z^{-1/2}(1-z)^{-1/2}$, which is the Beta$(\frac12,\frac12)$.

(a). (xx an attempt. xx) for fixed t, show via previous results that $\max_{[0,t]} W(s) \sim |W(t)| = t^{1/2}N_1$ and $\max_{[t,1]} W(s) \sim |W(1-t)| = (1-t)^{1/2}N_2$, where N_1 and N_2 are independent standard normals. With $C = N_2/N_1$ being Cauchy, deduce that

$$\begin{aligned}\Pr(T \le t) &= \Pr(\max_{[0,t]} W(s) > \max_{[t,1]} W(s))\\ &= \Pr(t^{1/2}|N_1| > (1-t)^{1/2}|N_2|)\\ &= \Pr(|C| < t^{1/2}/(1-t)^{1/2}),\end{aligned}$$

and from this derive the Beta$(\frac12,\frac12)$ density. note: can also be expressed via $F_{1,1}$.

(b).

Notes and pointers

A few remarks. Read more about risk functions in DeGroot (1970).

Brown, Einstein, Wiener, Donsker, Billingsley. Empirical procesess with estimated parameters: Durbin 1973, Hjort 1990, Loynes 1980.

Robert Brown (1773–1858) is the British botanist, who once had a cup of tea; Norbert Wiener (1894–1964) the father of cybernetics etc.; also Albert Einstein (1879–1955) belongs on the list of famous scientists who early on worked with this process.

Tightness: what are nice sufficient conditions for X_n, if it has independent increments, beyond $X_n(t) - X_n(s)$ converging for each interval?

[xx need care with what to use where etc. cross-check with other chapters, also Ch. 15. xx]

[xx make a point of Glivenko–Cantelli, $D_n/\sqrt{n} \to_{\text{a.s.}} 0$, which we perhaps have as an exercise in Ch. 2, and that Kolmogorov's $D_n \to_d D$ is more informative. i think both results appeared in 1933. xx]

[xx cool log-likelihood monitoring bridge: from Cunen et al. (2018). pointer to Story iv.5. xx]

[xx should have gamma process and poisson process, to illustrate the use of such limit processes inside $D[0,1]$. pointer to bayesian nonparametrics too. xx]

[xx can have Nils' Beta process construction as one of the Stories. xx]

[xx may consider pushing math details for C and D spaces to an envisaged Technical Appendix, with various nitty-gritty details. here description of compacta; the $w''(x,\delta)$ modulus for D; the Skorohod topology; the Prohorov theorem about tightness. xx]

[xx at some point we need to point to compactness characterisations, in C and in D, without making too much of it. mention Hanche-Olsen and Holden (2010), which binds together the Arzelà-Ascoli theorem and the Kolmogorov compactness theorem. xx]

ToDo: nils does more cleaning. nils makes some exercises with Poisson process, with Gamma process, pointing also to chapter on Bayesian nonparametrics.

I.14

Bootstrapping

Part II

Stories

Part III

Appendix

III.A

Mini-primer on measure and integration theory

[xx Mini-primer on measures, probabilities on spaces, integration theory. Background for rest of the book. xx]

Chapter introduction

(xx mini intro to measure theory and integration, background for probability measures, distributions, densities, models, etc. we also explain in one paragraph that yes, these things matter, and without them we cannot work properly; on the other hand, for most of the work we do, also in later chapters, we do not need to think too much about it. it's also a matter of becoming *basically literate* in the probability language underlying theoretical and also applied statistics. xx)

Various aspects of probability theory, and hence statistics methodology, rest on the general theory of measure and integration. If all random variables we meet have nice distributions and densities, on regular domains, like an interval, the real line, or open subsets of Euclidean spaces, we can get pretty far without this underlying measure and integration theory. To formulate concepts in natural generality, and to develop tools and demonstrate basic properties for these, however, one needs this more general theory. In particular, the business of defining probabilities, for perhaps complicated events in not-so-standard spaces, demands theory beyond 'ordinary' integration.

Essentials of measure, integration, and probability

Ex. A.1. *Some set theory.* Let Ω be a set, and $A, B, A_1, A_2, \ldots$ subsets thereof. The *union* $A \cup B$ consists of the points $\omega \in \Omega$ that are in A or in B (or in both A and B). So $\cup_{n=1}^{\infty} A_n = \{\omega \in \Omega\colon \omega \in A_n \text{ for at least one } n\}$. The *intersection* $A \cap B$ consits of the points $\omega \in \Omega$ that are in A and in B, so $\cap_{n=1}^{\infty} A_n = \{\omega \in \Omega\colon \omega \in A_n \text{ for all } n\}$. The *complement* A^c consists of all the points $\omega \in \Omega$ that are not in A, $A^c = \{\omega \in \Omega\colon \omega \notin A\}$. The *set difference* $A \setminus B = A \cap B^c$, and the *symmetric difference* of A and B is $A \Delta B = (A \setminus B) \cup (B \setminus A)$. We say that A is a *subset* of B if all the elements of A are also elements of B, and denote this $A \subset B$. If $A \subset B$ and $B \subset A$, then $A = B$. The collection of all subsets of Ω, is called the *power set* and is denoted 2^{Ω}.

distributive laws

(a). Prove the *distributive laws*

$$B \cap (\cup_{n=1}^{\infty} A_n) = \cup_{n=1}^{\infty} (B \cap A_n), \quad \text{and} \quad B \cup (\cap_{n=1}^{\infty} A_n) = \cap_{n=1}^{\infty} (B \cup A_n),$$

and *de Morgan's laws*, de Morgan's laws

$$(\cup_{n=1}^{\infty} A_n)^c = \cap_{n=1}^{\infty} A_n^c, \quad \text{and} \quad (\cap_{n=1}^{\infty} A_n)^c = \cup_{n=1}^{\infty} A_n^c.$$

(b). The *empty set*, denoted $\emptyset$, is the set with no element, think of it as $\emptyset = \{\}$. Write down a truth table with the columns $P\colon \omega \in \emptyset$, $Q\colon \omega \in A$, and 'if P then Q', to prove the vacuous truth that the empty set is a subset of any set.

(c). To any function $f\colon \Omega \to \mathcal{X}$, there is an associated inverse image. The inverse image, denoted f^{-1}, is a set mapping $f^{-1}\colon 2^{\mathcal{X}} \to 2^{\Omega}$, defined by

$$f^{-1}(B) = \{\omega \in \Omega\colon f(\omega) \in B\}, \quad \text{for } B \in \mathcal{X}.$$

Show that that for subsets $B, B_1, B_2, \ldots$ of $\mathcal{X}$, the inverse image f^{-1} preserves the set operations complement, union, and intersection in the sense that $f^{-1}(B^c) = (f^{-1}(B))^c$, $f^{-1}(\cup_{n=1}^{\infty} B_n) = \cup_{n=1}^{\infty} f^{-1}(B_n)$, and $f^{-1}(\cap_{n=1}^{\infty} B_n) = \cap_{n=1}^{\infty} f^{-1}(B_n)$.

(d). The sets Ω or $\mathcal{X}$ will often be the real line or a subset thereof. For $a < b$, the open interval is $(a, b) = \{x \in \mathbb{R}\colon a < x < b\}$, the closed interval is $[a, b] = \{x \in \mathbb{R}\colon a \le x \le b\}$, and so on. If $a = b$, then $(a, b) = \emptyset$ and $[a, b] = a$. Show that $(a, b) = \cup_{n=1}^{\infty}(a, b - 1/n]$, $[a, b] = \cap_{n=1}^{\infty}(a - 1/n, b]$, and that $\{a\} = \cap_{n=1}^{\infty}(a - 1/n, a] = \cap_{n=1}^{\infty}(a - 1/n, a + 1/n)$.

(e). A subset B of $\mathbb{R}$ is open if for any $x \in B$ you can fit a little ε-interval $(x - \varepsilon, x + \varepsilon)$ around x. Show that any open set in $\mathbb{R}$ is a countable union of open intervals. To do so, you may consider open intervals with rational endpoints, and recall that a countable union of countable sets, is countable.

(f). The *Cartesian product* of two sets A and B is the set $A \times B = \{(a, b)\colon a \in A, b \in B\}$, of three sets A, B and C it is $A \times B \times C = \{(a, b, c)\colon a \in A, b \in B, c \in C\}$, and so on. If $A = \{1, 2, 3\}$ and $B = \{4, 5\}$ what is $A \times B$? Make small sketches of the Cartesian products $\mathbb{R}^2 = \mathbb{R} \times \mathbb{R}$, and also of $\mathbb{R} \times \mathbb{N}$, and $\mathbb{Z}^2 = \mathbb{Z} \times \mathbb{Z}$. Cartesian product

Ex. A.2. *Measurable spaces.* Underlying all of statistics and probability is a mathematical model for randomness. This model consists of three things: a set Ω, called the *sample space*, containing all possible outcomes of the random phenomenon we are interested in; a family $\mathcal{A}$ of subsets of Ω, whose members are called *events*; and a function Pr having $\mathcal{A}$ as its domain and the unit interval as its range, called a *probability measure*. We start with a *measurable space*, say $(\Omega, \mathcal{A})$, consisting of a non-empty set Ω and a collection $\mathcal{A}$ of subsets of Ω. The subsets in $\mathcal{A}$ are later to be given values, perhaps probabilities, in terms of a measure. For $\mathcal{A}$, we demand that a σ-algebra of sets

(i) if $A \in \mathcal{A}$, then $A^c = \Omega \setminus A \in \mathcal{A}$;
(ii) if $A_1, A_2, \ldots$ are in $\mathcal{A}$, then $\cup_{j=1}^{\infty} A_j \in \mathcal{A}$.

A family of subsets $\mathcal{A}$ with these properties is called a σ-algebra. Thus, a σ-algebra is family of subsets that is closed under complements and countable unions.

(a). Show that if $\mathcal{A}$ is a σ-algebra, then Ω and the emptyset $\emptyset$ are in $\mathcal{A}$.

(b). Show that the power set 2^{Ω} is a σ-algebra.

(c). At the other extreme, show that the trivial σ-algebra $\{\emptyset, \Omega\}$ is a σ-algebra.

(d). Somewhat intermediately, show that if $A \subset \mathcal{A}$ is some nonempty set, then $\{\emptyset, A, A^c, \Omega\}$ is an σ-algebra.

(e). Consider $\mathcal{A}$, all subsets of $\mathbb{R}$ which are either empty, finite, or countably infinite, or whose complements are either empty, finite, or countably infinite. Show that $\mathcal{A}$ is a σ-algebra.

(f). Show that if $\mathcal{A}$ is a σ-algebra, with $B, A_1, A_2, \ldots \in \mathcal{A}$, then also $A_1 \cap A_2$, $A_1 \cap A_2 \cap A_3$, and even $\cap_{i=1}^{\infty} A_i$, is in $\mathcal{A}$. Show that sets like $A_1 \cap (A_2 \cup A_3 \cup A_4)^c \cap A_5$, $B \cap (\cup_{n=1}^{\infty} A_n)$, and $A \setminus B$ are in $\mathcal{A}$.

(g). Show that an intersection of σ-algebras must be a σ-algebra. Hence we may start by identifying a list of basis events, finite or infinite, say $\mathcal{B}_0$, and then define $\mathcal{B} = \sigma(\mathcal{B}_0)$, as *the smallest σ-algebra* containing all sets in $\mathcal{B}_0$. The σ-algebra $\mathcal{B}$ is said to be *generated* by $\mathcal{B}_0$, and $\mathcal{B}_0$ is called a generating family. Working with basis events that generate a σ-algebra is much more convenient than trying to somehow list all types of subsets of the σ-algebra.

the Borel-σ-algebra

(h). A famous and important example of a generated σ-algebra is the *Borel-σ-algebra on the real line*, denoted $\mathcal{B}(\mathbb{R})$, defined as the σ-algebra generated by all open intervals (a, b). Show that sets $\{a\}$, $[a, b]$, $[a, b)$, $(a, b]$, $(-\infty, b)$, $(-\infty, b]$, (a, ∞), $[a, \infty)$, as well as all countable unions and intersections of these, are then also in $\mathcal{B}(\mathbb{R})$. Show also that any of the families $\{(a, \infty)\colon a \in \mathbb{R}\}$, $\{[a, \infty)\colon a \in \mathbb{R}\}$, $\{(-\infty, a)\colon a \in \mathbb{R}\}$, and $\{(-\infty, a]\colon a \in \mathbb{R}\}$ also generates $\mathcal{B}(\mathbb{R})$.

(i). Similarly to the real case, we define $\mathcal{B}(\mathbb{R}^k)$, the Borel-$\sigma$-algebra on $\mathbb{R}^k$, as the σ-algebra generated by all rectangles $(a_1, b_1) \times \cdots \times (a_k, b_k)$. Show that $\mathcal{B}^k$ then also must contain all closed rectangles $[a_1, b_1] \times \cdots \times [a_k, b_k]$, all rectangles of half-open sets $(a_1, b_1] \times \cdots \times (a_k, b_k]$, etc., and also all open sets of $\mathbb{R}^k$.

(j). Let $\mathcal{B}_0$ be the collection of all intervals of the form $(-\infty, q]$, with $q \in \mathbb{Q}$. Show that $\sigma(\mathcal{B}_0) = \mathcal{B}(\mathbb{R})$. A σ-algebra generated by a countable collection of sets is said to be separable. Let (Ω, d) be a separable metric space, i.e., Ω contains a countable and dense subset. Show that the smallest σ-algebra on Ω that contains all the open sets is a separable σ-algebra.

(k). If $(\Omega_1, \mathcal{F})$ and $(\Omega_2, \mathcal{G})$ are two measurable spaces, then we can construct the product space $(\Omega_1 \times \Omega_2, \mathcal{F} \otimes \mathcal{G})$, where the *product σ-algebra* $\mathcal{F} \otimes \mathcal{G}$ is the smallest σ-algebra on the Cartesian product $\Omega_1 \times \Omega_2$ that contains all sets of the form $F \times G$, where $F \in \mathcal{F}$ and $G \in \mathcal{G}$. Show that if $\Omega_1 = \Omega_2 = \mathbb{R}$ and $\mathcal{F} = \mathcal{G} = \mathcal{B}(\mathbb{R})$, then $\mathcal{B}(\mathbb{R}) \otimes \mathcal{B}(\mathbb{R}) = \mathcal{B}(\mathbb{R}^2)$. Generalise to dimension $k > 2$.

extended real numbers

(l). Plus and minus infinity will occur naturally as limits of sequences of real functions. Consider, for example, $f_n(x) = (\sqrt{n/2\pi}) \exp(-\frac{1}{2} n x^2)$ for $n = 1, 2, \ldots$, i.e., the density of the mean of n independent standard normals. Define the *extended real numbers* $\bar{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$. Similarly, $\bar{\mathbb{R}}_+ = \mathbb{R} \cup \infty$. If we allow for functions taking values in $\bar{\mathbb{R}}$, then $f_n(x)$ converges pointwise to $f(x)$, where $f(x) = 0$ if $x \neq 0$, and $f(x) = \infty$ if $x = 0$. Show that the Borel-σ-algebra on $\bar{\mathbb{R}}$ is $\mathcal{B}(\bar{\mathbb{R}}) = \sigma(\mathcal{B}(\mathbb{R}), -\infty, \infty)$; or, equivalently, the σ-algebra generated by $\mathcal{B}_0 = \{(a, b), [-\infty, b), (a, \infty]\colon a, b \in \mathbb{R}\}$. Show similarly, that the Borel-σ-algebra on $\bar{\mathbb{R}}_+$ is $\sigma(\mathcal{B}(\mathbb{R}_+), \infty)$.

(m). In general, the Borel-σ-algebra on a set $\mathcal{X}$ is defined as the σ-algebra generated by all the open sets in $\mathcal{X}$, where open is defined in terms of a metric or a topology on $\mathcal{X}$. Consider the space $C[0,1] = \{\text{all continuous functions } f\colon [0,1] \to \mathbb{R}\}$, equipped with $d(f,g) = \sup_{x\in[0,1]} |f(x) - g(x)|$. The space $(C[0,1], d)$ is a separable metric space (as we will see in Ex. 9.5). The open ε-ball around $g \in C[0,1]$ is $B_\varepsilon(g) = \{f \in C[0,1]\colon d(f,g) < \varepsilon\}$, and the Borel-$\sigma$-algebra on $C[0,1]$ is the smallest σ-algebra containing all the open balls. Show that this σ-algebra is generated by $\cup_{\varepsilon\in\mathbb{Q}\cap(0,\infty)}\{B_\varepsilon(g_1), B_\varepsilon(g_2), \ldots\}$, where $\{g_1, g_2, \ldots\}$ is dense in $C[0,1]$; and also by the collection of sets of the form $\{f \in C[0,1]\colon (f(x_1), \ldots, f(x_n)) \in B_1 \times \cdots B_n\}$ where $B_1, \ldots, B_n \in \mathcal{B}(\mathbb{R})$ and $0 \le x_1 < x_2 < \cdots < x_n \le 1$.

Ex. A.3. *Measurable functions.* Let $(\Omega, \mathcal{A})$ and $(\mathcal{X}, \mathcal{B})$ be two measurable spaces. A function $f\colon \Omega \to \mathcal{X}$ is *$\mathcal{A}/\mathcal{B}$-measurable* if $f^{-1}(B) \in \mathcal{A}$ for all $B \in \mathcal{B}$. When it is clear what σ-algebras are involved, we simply say that f is measurable. When $\mathcal{X}$ is $\mathbb{R}$, $\bar{\mathbb{R}}$, or $\mathbb{C}$, $\mathcal{B}$ will always be the Borel-σ-algebra.

(a). In view of the efforts in Ex. A.2, the following lemma, tells much of the story of this exercise. Let $(\Omega, \mathcal{A})$ and $(\mathcal{X}, \mathcal{B})$ be two measurable spaces, and suppose that $\mathcal{B}_0$ is a collection of subsets of $\mathcal{X}$ that generates $\mathcal{B}$. Then f is measurable if and only if $f^{-1}(B_0) \in \mathcal{A}$ for all $B_0 \in \mathcal{B}_0$. Look back at Ex. A.1(c), and prove this lemma.

(b). On a measurable space $(\Omega, \mathcal{A})$, let $f\colon \Omega \to \mathbb{R}$ be a measurable function. Show that $f^{-1}(a,b) \in \mathcal{A}$ for all open intervals (a,b), and that this might as well be taken as our definition of measurability for real valued functions. In particular, show that the sets $f^{-1}\{a\}$, $f^{-1}[a,b)$, $f^{-1}(a,b]$, $f^{-1}(-\infty,b)$ $f^{-1}(-\infty,b]$, $f^{-1}(a,\infty)$, and $f^{-1}[a,\infty)$ are all in $\mathcal{A}$. Generalise to extended real valued measurable functions $f\colon \Omega \to \bar{\mathbb{R}}$.

(c). Suppose that (Ω, d) is a a metric space equipped with its Borel-σ-algebra, and let $f\colon \Omega \to \bar{\mathbb{R}}$ be a continuous function. Show that f is measurable.

(d). We will frequently use the characteristic function, a function taking values in the complex plane. An open subset of $\mathbb{C}$ can be identified with the open subsets of $\mathbb{R}^2$. For example, and open box in $\mathbb{C}$ is $\{x + iy\colon a < x < b, c < y < d\}$. Therefore, any collection of subsets of $\mathbb{R}^2$ that generates $\mathcal{B}(\mathbb{R}^2)$ can, by making the appropriate identifications, be used to generate $\mathcal{B}(\mathbb{C})$. Let $f = g + ih$ be a complex valued function on the measurable space $(\Omega, \mathcal{A})$, for measurable functions $g, h\colon \Omega \to \mathbb{R}$. Show that f is measurable if and only if g and h are measurable. In particular, use (c) to show that function $\exp(ix) = \cos(x) + i\sin(x)$ is measurable.

(e). Let $f\colon \Omega \to \mathcal{X}$ be a measurable function between $(\Omega, \mathcal{A})$ and $(\mathcal{X}, \mathcal{B})$. The σ-algebra generated by f is the smallest σ-algebra on Ω such that f is measurable, i.e., the intersection of all σ-algebras with respect to which f is measurable. We denote it $\sigma(f)$, and clearly, $\sigma(f) \subset \mathcal{A}$. Show that $\sigma(f) = \{f^{-1}(B)\colon B \in \mathcal{B}\}$, and that if $\mathcal{X} = \bar{\mathbb{R}}$ and $\mathcal{B} = \mathcal{B}(\bar{\mathbb{R}})$, then $\sigma(f)$ may be generated by the collection of sets $\pi(f) = \{\{\omega \in \Omega\colon f(\omega) \le x\}\colon x \in \mathbb{R}\}$.

(f). First, let $f_1, \ldots, f_n$ be finitely many measurable functions. Show that $\max(f_1, \ldots, f_n)$ and $\min(f_1, \ldots, f_n)$ are measurable. Next, let $f_1, f_2, \ldots$ be a sequence of extended real valued measurable functions. Show that $\sup_{n\ge1} f_n$ and $\inf_{n\ge1} f_n$ are also measurable functions. Show that if $0 \le g_1 \le g_2 \le \cdots$ is a sequence of functions where $g_n(\omega)$ is nondecreasing (in n) for each ω, then the limit function g, with $g(\omega) = \lim_{n\to\infty} g_n(\omega)$, is also measurable. Next, show that $\limsup_{n\to\infty} f_n$ and $\liminf_{n\to\infty} f_n$ are measurable.

Finally, show that, if it exists, the limit function $f(\omega) = \lim_{n\to\infty} f_n(\omega)$ is a measurable function.

Simple functions

(g). A simple function is a function taking on only finitely many values, meaning that if g is a simple function it can be written $g = \sum_{j=1}^{k} c_j I_{A_j}$ for sets $A_1, \ldots, A_k \in \mathcal{A}$ such that $\cup_{j=1}^{k} A_j = \Omega$ and real constants $c_1, \ldots, c_k$. If the $A_1, \ldots, A_k$ are disjoint we say that g is a simple function on standard form. With f any nonnegative measurable function on $(\Omega, \mathcal{A})$, show that the sequence of simple functions

Approximation by simple functions

$$f_n(\omega) = \sum_{k=1}^{n2^n} \frac{k-1}{2^n} I_{A_{n,k}}(\omega) + n I_{A_n}(\omega),$$

with $A_{n,k} = \{\omega \in \Omega\colon (k-1)/2^n \le f(\omega) < m/2^n\}$ and $A_n = \{\omega \in \Omega\colon f(\omega) \ge n\}$, is measurable, and is such that $0 \le f_1 \le f_2 \le \cdots$ and $f(\omega) = \lim_n f_n(\omega)$ for each ω.

(h). On $(\Omega, \mathcal{A})$, let $f, g\colon \Omega \to \mathbb{R}$ be measurable functions, and let a, b be constants. Show that $af + bg$, fg, and f/g if $g \neq 0$ are measurable.

(i). Let $f\colon \Omega \to \mathcal{X}$ and $g\colon \mathcal{X} \to \bar{\mathbb{R}}$ be measurable functions, on the measurable spaces $(\Omega, \mathcal{A})$ and $(\mathcal{X}, \mathcal{C})$, respectively. Show that the composition $g(f(\cdot))\colon \Omega \to \bar{\mathbb{R}}$ is measurable.

(j). In (i), it is important that g is $\mathcal{C}$-measurable. Let $\mathcal{C}'$ be a σ-algebra on $\mathcal{X}$, containing at least one set, for example C', that is not in $\mathcal{C}$. Define $g = I_{C'}$, and let $h = g(f(\cdot))$, with f the $\mathcal{A}/\mathcal{C}$-measurable function from (i). Show that h is not measurable. [xx can make this more concrete, see, e.g. Romano and Siegel (1986, p. 36) xx]

Ex. A.4. *Measure and measure spaces.*. A *measure* μ on a measurable space $(\Omega, \mathcal{A})$ is a function $\mu\colon \mathcal{A} \to [0, \infty]$, giving values to all sets A in $\mathcal{A}$, with the following properties:

(i) $\mu(\emptyset) = 0$, for the empty set;
(ii) $\mu(\cup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty} \mu(A_i)$, for $A_1, A_2, \ldots$ disjoint sets in $\mathcal{A}$, .

The resulting triple $(\Omega, \mathcal{A}, \mu)$ is called a *measure space*. We say that μ is a *finite measure* if $\mu(\Omega)$ is finite. Prime examples of finite measures are those with $\mu(\Omega) = 1$, such measures are *probability measures*, to be returned to in the rest of this book. If Ω can be represented as a countable union $\cup_{i=1}^{\infty} A_i$ of measurable sets, with each $\mu(A_i)$ finite, we say that μ is a σ-finite measure.

σ-finite measure

(a). Show that countable additivity implies finite additivity, namely $\mu(A_1 \cup \cdots \cup A_n) = \mu(A_1) + \cdots \mu(A_n)$ for any finite sequence of disjoint measurable sets. Show also that if $A, B \in \mathcal{A}$ and $A \subset B$, then $\mu(A) \le \mu(B)$. Deduce that $\mu(\cup_{i=1}^{\infty} A_i) \le \sum_{i=1}^{\infty} \mu(A_i)$, for any sequence of sets $A_1, A_2, \ldots$ in $\mathcal{A}$.

Continuity of measure

(b). Suppose that $A_1 \subset A_2 \subset \cdots$ is a nondecreasing sequence of sets in $\mathcal{A}$. Show that $\nu(\cup_{n=1}^{\infty} A_n) = \lim_{n\to\infty} \nu(A_n)$. Suppose now that $A_1 \supset A_2 \supset \cdots$ is a decreasing sequence in $\mathcal{A}$. Show that $\nu(\cap_{n=1}^{\infty} A_n) = \lim_{n\to\infty} \nu(A_n)$, provided $\nu(A_n)$ is finite for some n.

(c). Let ν be a *finitely additive* measure. That is, a function $\nu\colon \mathcal{A} \to [0, \infty]$, such that (i) $\nu(\emptyset) = 0$, and (ii) $\nu(A \cup B) = \nu(A) + \nu(B)$ for all disjoint sets $A, B \in \mathcal{A}$. Suppose that for any nondecreasing sequence $A_1 \subset A_2 \subset \cdots$ in $\mathcal{A}$ it is the case that $\nu(\cup_{n=1}^{\infty} A_n) = \lim_{n\to\infty} \nu(A_n)$. Show that ν is a measure. Suppose that for any nonincreasing sequence $B_1 \supset B_2 \supset \cdots$ in $\mathcal{A}$ such that $\cap_n^{\infty} A_n = \emptyset$, we have $\lim_{n\to\infty} \nu(B_n) = 0$. Show that ν is a measure, provided $\nu(\Omega)$ is finite [xx check xx]. In other words, countably additivity is a continuity property in disguise.

(d). Let $(\Omega, \mathcal{A}, \mu)$ be a measure space with no atoms, that is, $\mu(\{\omega\}) = 0$ for all $\omega \in \Omega$. Show that $\mu(A) = 0$ for any countable set $A \in \mathcal{A}$.

(e). Let μ be a measure on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ such that $\mu((a,b)) = b - a$. Show that $(\mathbb{R}, \mathcal{B}(\mathbb{R}), \mu)$ has no atoms, and that $\mu([a,b]) = \mu((a,b]) = \mu([a,b)) = b - a$. This measure, called Lebesgue measure, is constructed in Ex. A.7.

(f). Consider the σ-algebra $\mathcal{A}$ of those subsets of $\mathbb{R}$ which are empty, or finite, or countably infinite, or complements of such sets. Let $\nu(A)$ be the simple measure which counts the number of elements in A. Show that it defines a measure, called the *counting measure*, and that it is not finite nor σ-finite. counting measure

(g). Consider the σ-algebra $2^{\mathbb{N}}$ of all subsets of $\mathbb{N} = \{1, 2, \ldots\}$. Let ν be the counting measure. Show that ν is σ-finite.

(h). Let $(\Omega, \mathcal{A}, \nu)$ be a measure space, $(\mathcal{X}, \mathcal{B})$ a measurable space, and $f\colon \Omega \to \mathcal{X}$ a measurable function. Define the set function $\nu f^{-1}(B) = \mu(f^{-1}(B))$ for $B \in \mathcal{B}$. Show that νf^{-1} is a measure on $\mathcal{B}$.

Ex. A.5. *π-systems, d-systems, monotone classes.* We often need to prove that a certain property of a measure, an integral, or the like, holds for all sets in a σ-algebra. Without further tools this is a tall order, simply because it can be exceedingly hard to give 'closed form' characterisations of all the elements of a σ-algebra. Think, for example, of the Borel-σ-algebra on the real line, how are you to describe all its elements without recourse to the family of sets that generates it? This is the motivation for the definitions and results we introduce in this exercise. We start with the definitions we need: Let $\mathcal{C}$ be a collection of subsets of a set Ω, then $\mathcal{C}$ is

- an algebra if it is closed under complements and finite unions;
- a π-system if $A \cap B \in \mathcal{C}$ for every $A, B \in \mathcal{C}$;
- a d-system if (i) $\Omega \in \mathcal{C}$; (ii) $A \subset B$ then $B \setminus A \in \mathcal{C}$ for $A, B \in \mathcal{C}$; and (iii) if $A_1 \subset A_2 \subset \cdots$ are in $\mathcal{C}$ then $\cup_{n=1}^{\infty} A_n \in \mathcal{C}$.
- a monotone class if (i) $A_1 \subset A_2 \subset \cdots$ are in $\mathcal{C}$ then $\cup_{n=1}^{\infty} A_n \in \mathcal{C}$, and (ii) if $A_1 \supset A_2 \supset \cdots$ are in $\mathcal{C}$ then $\cap_{n=1}^{\infty} A_n \in \mathcal{C}$;

(a). Show that the collection of sets $\pi(f)$ defined in A.3(e) is a π-system.

(b). Show that (i) an algebra is a π-system; (ii) a d-system is a monotone class; (iii) a σ-algebra is a π-system, a d-system, and a monotone class; (iv) a collection of sets that is both an algebra and a monotone class is a σ-algebra; (v) a collection of sets that is both a π-system and a d-system is a σ-algebra.

(c). If Π is a π-system, and $\mathcal{D}$ is a d-system that contains Π, then $\sigma(\Pi) \subset \mathcal{D}$. In particular $\sigma(\Pi) = d(\Pi)$, where $d(\Pi)$ is the smallest d-system containing Π. In (d) we will get a glimpse of the power of this lemma. Here, we prove it, step by step. Step 1: Explain why it is enough to prove that $d(\Pi)$ is a π-system. Step 2: To prove that $d(\Pi)$ is a π-system, form the set $d(\Pi)_A = \{B \in d(\Pi)\colon A \cap B \in d(\Pi)\}$, and show that if $A \in d(\Pi)$, then $d(\Pi)_A$ is a d-system. Step 3: Show that if $A \in \Pi$, then $d(\Pi) \subset d(\Pi)_A$. Step 4: Show that if $A \in d(\Pi)$, then we still have that $d(\Pi) \subset d(\Pi)_A$. Step 4: Conclude from the above that $d(\Pi)$ is a π-system. Dynkin's lemma

(d). Let ν and μ be measures on $(\Omega, \mathcal{A})$ such that $\nu(\Omega) = \mu(\Omega)$ is finite, and suppose ν and μ agree on a π-system that generates $\mathcal{A}$, i.e., $\mathcal{A} = \sigma(\Pi)$. Consider the collection of sets given by

$$\mathcal{D} = \{A \in \mathcal{A} \colon \nu(A) = \mu(A)\}.$$

Since $\Pi \subset \sigma(\Pi) = \mathcal{A}$ and $\nu(A) = \mu(A)$ for all $A \in \Pi$, it is clearly the case that $\Pi \subset \mathcal{D}$. Show that $\mathcal{D}$ is a d-system and conclude from (c) that $\nu = \mu$.

Monotone class theorem

(e). We now turn to a theorem that is quite similar to that proved in (c), the difference being that our basis events form an algebra, and an algebra is a π-system, but the reverse need not hold. Here is the theorem: Let Ω be a set, and $\mathcal{A}_0$ an algebra of subsets of Ω. Suppose that $\mathcal{M}$ is a monotone class of subsets of Ω and that $\mathcal{A}_0 \subset \mathcal{M}$. Then $\sigma(\mathcal{A}_0) \subset \mathcal{M}$. To prove this, we proceed in steps. Step 1: Let $m(\mathcal{A}_0)$ be the smallest montone class containing $\mathcal{A}_0$, and argue that it suffices to prove that $m(\mathcal{A}_0)$ is a σ-algebra. Step 2: Show that $m(\mathcal{A}_0)$ is closed under countable unions. Step 3: Form the set $\mathcal{M}_c = \{B \in m(\mathcal{A}_0) \colon B^c \in m(\mathcal{A}_0)\}$, argue why it is now sufficient to prove that $\mathcal{M}_c$ is a monotone class, and prove that $\mathcal{M}_c$ is indeed a monotone class.

(f). If $\mathcal{A}_0$ is an algebra, a natural question is how much bigger than $\mathcal{A}_0$ the σ-algebra it generates is. A partial answer to this question can be given by an application of the monotone class theorem from which we conclude that if $\mathcal{A}_0$ is an algebra, then every set in the σ-algebra generated by $\mathcal{A}_0$ can be approximated arbitrarily well by sets from $\mathcal{A}_0$. To make this precise, let $(\Omega, \mathcal{A}, \nu)$ be a finite measure space. Let $\mathcal{A}_0$ be an algebra contained in $\mathcal{A}$, and form the family of subsets

$$\mathcal{M} = \{A \in \mathcal{A} \colon \text{for any } \varepsilon > 0 \text{ there is an } A_0 \in \mathcal{A}_0 \text{ such that } \nu(A \Delta A_0) < \varepsilon\}.$$

Here, $A \Delta B = (A \setminus B) \cup (B \setminus A)$ is the symmetric difference of the sets A and B. Show that $\nu(\cup_{j=1}^{\infty} A_j \Delta \cup_{j=1}^{\infty} B_j) = \nu(\cap_{j=1}^{\infty} A_j^c \Delta \cap_{j=1}^{\infty} B_j^c) \leq \sum_{j=1}^{\infty} \nu(A_j \Delta B_j)$, which you may use to show that $\mathcal{M}$ is a monotone class, and appeal the monotone class theorem to conclude that given $A \in \sigma(\mathcal{A}_0)$ and $\varepsilon > 0$, we can find $A_0 \in \mathcal{A}_0$ so that $\nu(A \Delta A_0) < \varepsilon$.

(g). Let ν and μ be two finite measures on $(\Omega, \mathcal{A})$ such that $\nu = \mu$ on an algebra $\mathcal{A}_0$ contained in $\mathcal{A}$, and consider the collection of sets $\mathcal{M} = \{A \in \mathcal{A} \colon \mu(A) = \nu(A)\}$. Use a monotone class argument to show that $\nu = \mu$ on the σ-algebra generated by $\mathcal{A}_0$.

(h). The measures in (g) need not be finite. Suppose that $\Omega = \cup_{n=1}^{\infty} A_n$ for sets $A_1, A_2, \ldots$ in the algebra $\mathcal{A}_0$ on which μ and ν agree, and that $\mu(A_n)$ and $\nu(A_n)$ are finite for $n = 1, 2, \ldots$, then $\mu = \nu$ on $\sigma(\mathcal{A}_0)$. To prove this, fix C_n, show that collection of sets $\mathcal{M}_n = \{A \in \mathcal{A} \colon \mu(A \cap C_n) = \nu(A \cap C_n)\}$ forms a monotone class, and take it from there.

Ex. A.6. *Lifting good candidates to bona fide measures.* An important and useful general result from measure theory is Carathéodory's Extension Theorem. The point of this theorem is to lift a set function defined on simpler subsets of a set than those contained in a σ-algebra, to bona fide measures defined on the full σ-algebra. We start out with a set function $\nu_0 \colon \mathcal{A}_0 \to [0, \infty]$ working on subsets A_0 of an algebra $\mathcal{A}_0$ of subsets of a set Ω. Suppose that ν_0 has the properties (i) $\nu_0(\emptyset) = 0$; and (ii) if $A_1, A_2, \ldots$ is a disjoint sequence in $\mathcal{A}_0$ whose union happens to be in $\mathcal{A}_0$, then $\mu_0(\cup_{n=1}^{\infty} A_n) = \sum_{n=1}^{\infty} \mu_0(A_n)$. Under these conditions, Carathéodory's Extension Theorem says that ν_0 on $\mathcal{A}_0$ can be lifted to a full measure ν on $\sigma(\mathcal{A}_0)$, with $\nu(A) = \nu_0(A)$ for all $A \in \mathcal{A}_0$. Also, if

Carathéodory's Extension Theorem

$\Omega = \cup_{n=1}^{n} A_n$ for sets $A_1, A_2, \ldots$ in $\mathcal{A}_0$ and $\mu_0(A_n) < \infty$ for all n, then this extension is unique.

Notice that if $\mathcal{A}_0$ had been a σ-algebra, and not merely an algebra, then ν_0 would have been a measure. This is what distinguishes the present theorem from the results of Ex. A.5(d) and Ex. A.5(g), where ν and μ were assumed to be defined on σ-algebras, that is, they were assumed to be measures. In the extension theorem, by contrast, ν_0 has measure-like properties, but is only defined on an algebra.

Let us also mention the variations of Carathéodory's Extension Theorem where the function μ_0, instead of being defined on an algebra, as here, is defined on a semialgebra or on a semi-ring. A *semialgebra*, say $\mathcal{S}$, is a collection of subsets of a set Ω such that (i) if $A, B \in \mathcal{S}$, then $A \cap B \in \mathcal{S}$; and (ii) if $A \in \mathcal{S}$, then $A^c = \cup_{j=1}^{n} B_j$ for $B_1, \ldots, B_n \in \mathcal{S}$. A *semi-ring*, say $\mathcal{R}$, is a collection of subsets of a set Ω such that (i) $\emptyset \in \mathcal{R}$; (ii) if $A, B \in \mathcal{R}$, then $A \cap B \in \mathcal{R}$; and (iii) if $A, B \in \mathcal{R}$, then $A \setminus B = \cup_{j=1}^{n} C_j$ for $C_1, \ldots, C_n \in \mathcal{R}$. The advantage with these two variations of Carathéodory's Extension Theorem is that natural basis events often constitute a semialgebra or a semi-ring, but not an algebra. For example, $\mathcal{S} = \{\text{all intervals on } \mathbb{R}\}$ is a semialgebra, but not an algebra; and $\mathcal{R} = \{(a, b] \colon a, b \in \mathbb{R}\}$ is a semi-ring, but not an algebra. An algebra, on the other hand, is both a semialgebra and a semi-ring. The reader might verify these claims. Assuming that the basis events constitute a semialgebra or a semi-ring, rather than an algebra, therefore amounts to imposing weaker conditions than we do here, and, as a consequence, the proof is more involved than with the algebra version of the extension theorem. We now prove the extension theorem, as stated above, through a string of exercises.

(a). Let $\mathcal{A}_0$ be an algebra of subsets of a set Ω, and let $\lambda \colon \mathcal{A}_0 \to [0, \infty]$ be such that $\lambda(\emptyset) = 0$. Consider the collection of elements of $\mathcal{A}_0$ that splits every element of $\mathcal{A}_0$ 'as it should', namely

$$\mathcal{A}_0^\lambda = \{B \in \mathcal{A}_0 \colon \lambda(B \cap A) + \lambda(B^c \cap A) = \lambda(A) \text{ for all } A \in \mathcal{A}_0\}.$$

Show that $\mathcal{A}_0^\lambda$ is an algebra, and that λ is finitely additive on $\mathcal{A}_0^\lambda$.

(b). If $\mathcal{A}$ is a σ-algebra of subsets of Ω, the set function $\lambda \colon \mathcal{A} \to [0, \infty]$ is called an outer measure if

(i) $\lambda(\emptyset) = 0$;
(ii) $A, B \in \mathcal{A}$ with $A \subset B$, then $\lambda(A) \subset \lambda(B)$; outer-measure
(iii) if $A_1 \subset A_2 \subset \cdots$ are sets in $\mathcal{A}$, then $\lambda(\cup_{j=1}^{\infty} A_j) \leq \sum_{j=1}^{\infty} \lambda(A_j)$.

Carathéodory's lemma (not the theorem yet) says that if λ is an outer measure on the measurable space $(\Omega, \mathcal{A})$, then the collection of elements of the σ-algebra $\mathcal{A}$ that splits every element of $\mathcal{A}$ 'as it should', that is Carathéodory's lemma

$$\mathcal{A}^\lambda = \{B \in \mathcal{A} \colon \lambda(B \cap A) + \lambda(B^c \cap A) = \lambda(A) \text{ for all } A \in \mathcal{A}\},$$

form a σ-algebra, and λ is countably additive on $\mathcal{A}^\lambda$, meaning that $(\Omega, \mathcal{A}^\lambda, \lambda)$ is a measure space. Let $B_1, B_2, \ldots$ be a sequence of sets in $\mathcal{A}^\lambda$, and set $B = \cup_{j=1}^{\infty} B_j$. To show that $\lambda(A) = \lambda(A \cap B) + \lambda(A \cap B^c)$ for any $A \in \mathcal{A}$, show that $\lambda(A) \leq \lambda(A \cap B) + \lambda(A \cap B^c)$ and $\lambda(A) \geq \lambda(A \cap B) + \lambda(A \cap B^c)$. For the latter inequality, consider sets $C_n = \sum_{j=1}^{n} B_j$, and use that from (a), λ is finitely additive.

(c). We are now ready for the extension theorem, as stated in the introduction to this exercise, and prove this theorem in three steps. Step 1: For $G \in 2^\Omega$, define

$$\lambda(G) = \inf\{\sum_{j\geq 1} \mu_0(F_j)\colon F_1, F_2, \ldots \in \mathcal{A}_0 \text{ such that } G \subset \cup_{j\geq 1} F_j\},$$

and prove that λ is an outer measure. By (b), λ is a measure on $(\Omega, \mathcal{A}^\lambda)$, with $\mathcal{A}^\lambda$ defined as in (b) (so $\mathcal{A} = 2^\Omega$). This observation leads to the two final steps. Step 2: Show that $\mathcal{A}_0 \subset \mathcal{A}^\lambda$. Step 3: Show that $\lambda = \mu_0$ on $\mathcal{A}^\lambda$. We can then define μ to be the restriction of λ to $\mathcal{A} = \sigma(\mathcal{A}_0) \subset \mathcal{A}^\lambda$, with the inclusion coming from the already proven fact that $\mathcal{A}^\lambda$ is a σ-algebra.

(d). Point to Ex. A.5(h) to argue that if $\Omega = \cup_{n=1}^\infty A_n$ for sets $A_1, A_2, \ldots \in \mathcal{A}_0$ with $\mu(A_n) < \infty$ for $n = 1, 2, \ldots$, then the extension is unique.

Ex. A.7. *Lebesgue measure.* In this exercise we apply Carathéodory's extension theorem to the construction of Lebesgue measure on $\mathbb{R}$. The basic property of Lebesgue measure, say λ, is that for any interval

$$\lambda((a,b)) = \lambda([a,b]) = \lambda((a,b]) = \lambda([a,b)) = b - a,$$

the length of the interval in question. We want to have λ as a measure on the Borel-σ-algebra $\mathcal{B}(\mathbb{R})$, but it is not at all obvious that such a measure exists. Recall from Ex. A.2 that the only way we mange to describe the elements of $\mathcal{B}(\mathbb{R})$ are via the various generating classes (e.g. all open intervals), so how are we then to describe $\lambda(B)$ for some arbitray set B in $\mathcal{B}(\mathbb{R})$? The solution is to define λ for a simpler class of subsets of $\mathbb{R}$, and then use Carathéodory's theorem to extend the domain of λ to all of $\mathcal{B}(\mathbb{R})$.

(a). Let $\mathcal{S}$ be the collection of all half-open intervals $(a,b]$ on $\mathbb{R}$ (if $a \geq b$, then $(a,b] = \emptyset$), as well as all infinite intervals of the form $(-\infty, b]$ and (a, ∞). Let $\mathcal{A}_0$ be a collection of sets consisting of all finite disjoint unions of elements of $\mathcal{S}$. Show that $\mathcal{A}_0$ is an algebra, and that $\mathcal{B}(\mathbb{R}) = \sigma(\mathcal{A}_0)$.

(b). We define $\lambda\colon \mathcal{A}_0 \to [0,\infty]$ to be such that $\lambda((a,b]) = b - a$ when a, b are finite, and to be ∞ when at least one of them are not. Thus, λ gives what we think of as the length of an interval. For disjoint intervals $I_1, \ldots, I_k \in \mathcal{S}$, define

$$\lambda(\cup_{j=1}^n I_j) = \lambda(I_1) + \cdots + \lambda(I_n),$$

so, in particular, if $I_j = (a_j, b_j]$ with $a_j, b_j \in \mathbb{R}$ for $j = 1, \ldots, n$, then $\lambda(\cup_{j=1}^n (a_j, b_j]) = \sum_{j=1}^n (b_j - a_j)$. Show that this is unambiguous, meaning that if $\cup_{j=1}^n I_j = \cup_{i=1}^m I_i'$, then $\lambda(\cup_{j=1}^n I_j) = \lambda(\cup_{i=1}^m I_i')$.

(c). Show that $\lambda(\emptyset) = 0$, and that λ is finitely additive on $\mathcal{A}_0$. In order to lift λ from the algebra $\mathcal{A}_0$ on which it is currently defined to the σ-algebra $\mathcal{B}(\mathbb{R})$ that $\mathcal{A}_0$ generates, we need to show that λ is *countably* additive on $\mathcal{A}_0$. This is the hard part. We start by showing that if $\cup_{j=1}^\infty (a_j, b_j] = (a,b]$ for disjoint intervals $(a_1, b_1], (a_2, b_2], \ldots$, then $\lambda(a,b] = \sum_{j=1}^n \lambda(a,b]$. To prove this equality, prove the inequality both ways. First, show that if $\cup_{j=1}^n (a_j, b_j] \subset (a,b]$, for disjoint $(a_1, b_1], \ldots, (a_n, b_n]$, then $\sum_{j=1}^n \lambda(a_j, b_j] \leq \lambda(a,b]$. Second, if $(a_1, b_1], (a_2, b_2], \ldots$ are disjoint, and $\cup_{j=1}^\infty (a_j, b_j] \subset (a,b]$, then $\sum_{j=1}^\infty \lambda(a_j, b_j] \leq \lambda(a,b]$. Now we get to the reverse inequality. Third, show that if $(a,b] \subset \cup_{j=1}^n (a_j, b_j]$,

then $\lambda(a,b] \le \sum_{j=1}^n (a_j,b_j]$. Fourth, if $(a,b] \subset \cup_{j=1}^\infty (a_j,b_j]$, then $\lambda(a,b] \le \sum_{j=1}^\infty (a_j,b_j]$. To prove this fourth claim, note that for any $\varepsilon > 0$ smaller than $b-a$,

$$[a+\varepsilon,b] \subset (a,b] \subset \cup_{j=1}^\infty (a_j,b_j] \subset \cup_{j=1}^\infty (a_j,b_j+\varepsilon/2^j).$$

Thus, the closed and bounded set $[a+\varepsilon,b]$ is covered by the open sets $(a_1,b_1+\varepsilon/2), (a_2,b_2+\varepsilon/4), (a_3,b_3+\varepsilon/8), \ldots$, so by the Heine–Borel theorem it must then have a finite subcover, i.e., $[a+\varepsilon,b]$ is compact. Take it from here.

(d). Show that λ extends to a measure on Borel σ-algebra $\mathcal{B}(\mathbb{R})$, and that this extension, which is Lebesgue measure on the real line, is unique (see Ex. A.6(d)).

(e). Now that we have Lebesgue measure on the real line, we can construct Lebesgue measure on any subinterval of the real line, for example $([0,1],\mathcal{B}[0,1])$ or $([0,\infty),\mathcal{B}[0,\infty)$. Construct Lebesgue measure on these two measurable spaces.

(f). Establish similarly Lebesgue measure on $(\mathbb{R}^2,\mathcal{B}^2)$, i.e. on the plane, with its Borel sets, starting from the area of rectangles $\lambda((a_1,b_1)\times(a_2,b_2)) = (b_1-a_1)(b_2-a_2)$. Via the Carathéodory lifting, this gives rise to a well-defined way of measuring the area of any Borel subset A on the plane.

(g). Once the fundamental Lebesgue measure has been properly put on the map, it will be easy to define classes of others, via *cumulative distribution functions* and *densities*; see Ex. A.15 and A.18 below. It is nevertheless useful to go through direct arguments, resembling those for the Lebesgue measure itself, for a few concrete instances. Do this for the measures μ and ν on the positive halfline, starting with respectively $\mu(a,b) = \log(b/a)$ and $\nu(a,b) = b^2 - a^2$, for intervals (a,b).

Ex. A.8. *Almost surely and infinitely often.* Let $(\Omega,\mathcal{A},\mu)$ be a measure space. A property is said to hold μ almost surely, or μ-a.s., or simply a.s. if there is no confusion about the underlying measure, if it holds for for all ω outside of a set of μ-measure zero.

(a). Let $(\mathbb{R},\mathcal{B}(\mathbb{R}),\lambda)$ be a measure space with λ Lebesgue measure. Look back at Ex. A.4(e) to show that the indicator function $I_{\mathbb{R}\setminus\mathbb{Q}} = 1$ almost surely.

(b). From Ex. A.2(l), recall the sequence $f_n(x) = (\sqrt{2\pi/n})^{-1}\exp(-\frac{1}{2}nx^2)$ for $n = 1,2,\ldots$, defined on the measure space $(\mathbb{R},\mathcal{B}(\mathbb{R}),\lambda)$. Show that $f_n \to 0$ almost surely.

(c). Suppose that the sequence of measurable functions $f_1, f_2, \ldots$ converges to f almost surely. Show that f is measurable. Compare this to what you showed in Ex. A.3(f).

(d). Let $(\Omega,\mathcal{A},\mu)$ be a measure space and $A_1, A_2, \ldots$ a sequence of sets in $\mathcal{A}$. Show that $\limsup_{n\to\infty} A_n = \emptyset$ if and only if $\lim_{n\to\infty} \mu(\cup_{m=n}^\infty A_m) = 0$.

(e). Let $A_1, A_2, \ldots$ be a sequence of sets, and $I_{A_1}, I_{A_2}, \ldots$ the corresponding sequence of indicators functions. Show that $\liminf_{n\to\infty} I_{A_n}(\omega) = 1$ if and only if $\omega \in \cup_{n\ge1}\cap_{m\ge n} A_m$. Show also that $\limsup_{n\to\infty} I_{A_n}(\omega) = 1$ if and only if $\omega \in \cap_{n\ge1}\cup_{m\ge n} A_m$. These two equivalences motivate the definitions

$$\liminf_{n\to\infty} A_n = \cup_{n\ge1}\cap_{m\ge n} A_m, \quad \text{and} \quad \limsup_{n\to\infty} A_n = \cap_{n\ge1}\cup_{m\ge n} A_m.$$

Show that these sets are measurable provided the sets $A_1, A_2, \ldots$ are. In probability and statistics one encounters the notion of something occurring *infinitely often*, or i.o. This notion is defined by

$$\begin{aligned} A_{\text{i.o.}} &= \limsup_{n\to\infty} A_n = \cap_{n\geq 1} \cup_{m\geq n} A_m \\ &= \{\omega \in \Omega\colon \text{for every } n \text{ there is an } m = m(\omega) \geq n \text{ such that } \omega \in A_m\} \\ &= \{\omega \in \Omega\colon \omega \in A_n \text{ for infinitely many } n\}. \end{aligned}$$

Borel–Cantelli lemma

(f). Let $(\Omega, \mathcal{A}, \mu)$ be a measure space and assume that $A_1, A_2, \ldots \in \mathcal{A}$ are such that $\sum_{n=1}^{\infty} \mu(A_n) < \infty$. Show that $\mu(A_{\text{i.o.}}) = 0$. In Ex. **??** we will study several illustration of the use of this lemma, and also see that it has a partial converse.

Fatou's lemma for sets

(g). Let $(\Omega, \mathcal{A}, \mu)$ be a measure space, and let $A_1, A_2, \ldots \in \mathcal{A}$. Show that

$$\mu(\liminf_{n\to\infty} A_n) \leq \liminf_{n\to\infty} \mu(A_n).$$

This inequality, known as Fatou's lemma, holds not only for sequences of indicator functions, but, as we will see in Ex. A.11(c), for any sequence of nonnegative measurable functions.

Ex. A.9. *Convergence in measure/probability.* Let $f, f_1, f_2, \ldots$ be measurable functions on a measurable space $(\Omega, \mathcal{A}, \mu)$. The sequence $f_1, f_2, \ldots$ converges to f *in measure*, if for any $\varepsilon > 0$,

$$\lim_{n\to\infty} \mu(\{\omega \in \Omega\colon |f_n(\omega) - f(\omega)| \geq \varepsilon\}) = 0.$$

convergence in probability

If $\mu(\Omega) = 1$, so that μ is a probability measure, then convergence in measure is called convergence in probability.

(a). If $f_n \to f$ almost surely, then $f_n \to f$ in measure. To show this, consider the sets $A_n = \{\omega \in \Omega\colon |f_n(\omega) - f(\omega)| \geq \varepsilon\}$ for some $\varepsilon > 0$, show that $\cap_{n\geq 1} \cup_{m\geq n} A_n \subset A$, where A is the set where $f_n(\omega)$ does not converge to $f(\omega)$, and take it from there.

(b). Show that if $f_n \to f$ in measure, then there is a subsequence f_{n_j} such that $f_{n_j} \to f$ almost surely. To construct such a subsequence, you may use the Borel–Cantelli lemma from Ex. A.8(f).

(c). Let $([0,1], \mathcal{B}, \lambda)$ be the unit interval with the Borel-σ-algebra and the Lebesgue measure. Divide the unit interval in $2, 3, \ldots$ pieces: $A_1 = [0, 1/2]$, $A_2 = (1/2, 1]$, $A_3 = [0, 1/3]$, $A_4 = (1/3, 2/3]$, $A_5 = (2/3, 1]$, and so on. Define the function $f_n(x) = I(x \in A_n)$. Show that $f_n \to 0$ in measure, but not almost surely.

Ex. A.10. *The Lebesgue integral.* After having defined measures and measurable functions, the next goal is to form a well-defined *integral*, say $\int f \, d\nu = \int f(\omega) \, d\mu(\omega)$, with $(\Omega, \mathcal{A}, \mu)$ a measure space, and with $f\colon \Omega \to \bar{\mathbb{R}}$ a measurable function.

(a). We start with f a nonnegative and simple function on standard form, that is, f is on the form $f = a_1 I_{A_1} + \cdots + a_k I_{A_k}$ for some $n \geq 1$, with disjoint sets $A_1, \ldots, A_k \in \mathcal{A}$ such that $\cup_{j=1}^{k} A_k = \Omega$, and nonnegative real constants $a_1, \ldots, a_n$. We then define the integral of f as

$$\int f \, d\mu = a_1 \mu(A_1) + \cdots + a_k \mu(A_k). \tag{A.1}$$

If any of these sets have infinite measure and coefficient zero, we follow the measure theoretical convention that $0 \cdot \infty = 0$. Show that (A.1) is unambiguous, giving the same value for different representations of the same simple function, i.e., show that if $g = b_1 I_{B_1} + \cdots + b_n I_{B_n}$ is another nonnegative simple function on standard form such that $g(\omega) = f(\omega)$ for all $\omega \in \Omega$, then $\int g \, d\mu = \int f \, d\mu$.

(b). Show that $\int f \, d\mu$, as defined in (A.1), is linear and nondecreasing. That is, for nonnegative and simple functions on standard form f and g, and nonnegative real constants a, b, we have $\int (af + bg) \, d\mu = a \int f \, d\mu + b \int g \, d\mu$; and if $f \leq g$, then $\int f \, d\mu \leq \int g \, d\mu$.

(c). Next, we extend the integral to any nonnegative measurable function $f \colon \Omega \to \bar{\mathbb{R}}_+$. To do so, pick a sequence of nonnegative simple functions $0 \leq f_1 < f_2 < \cdots$ such that $f_n \to f$, the existence of which is guaranteed by Ex. A.3(g), and define

$$\int f \, d\mu = \lim_{n \to \infty} \int f_n \, d\mu.$$

We need to prove that $\int f \, d\mu$ is independent of the approximating sequence f_n. To do this, let $0 \leq g_1 < g_2 < \cdots$ be some other sequence of nonnegative simple functions on standard form converging to f, and show, first, that $\lim_{n\to\infty} \int g_n \, d\mu \leq \int f \, d\mu$. Second, consider the sets $A_n = \{\omega \in \Omega \colon g_n(\omega) \geq f_k(\omega)\}$, argue that $\liminf_{n\to\infty} I_{A_n} f_k = f_k$, and use Fatou's lemma (Ex. A.8(g)) to show the reverse inequality, $\lim_{n\to\infty} \int g_n \, d\mu \geq \int f \, d\mu$.

(d). Let's look at some of the properties of the integral defined in (d). Let $f, g \colon \Omega \to \bar{\mathbb{R}}_+$ be measurable, and show that the integral is (i) *linear* $\int (af + bg) \, d\mu = a \int f \, d\mu + b \int g \, d\mu$; and (ii) *nondecreasing*, if $f \leq g$ almost surely, then $\int f \, d\mu \leq \int g \, d\mu$. Show also that (iii) $\int f \, d\mu = 0$ if and only if $f = 0$ almost surely; (iv) if $f > 0$ almost surely, then $\int f \, d\mu > 0$; (v) if $f = g$ almost surely, then $\int f \, d\mu = \int g \, d\mu$; (vi) if $\int_A f \, d\mu = \int_A g \, d\mu$ for all $A \in \mathcal{A}$ and μ is σ-finite, then $f = g$ almost surely; and (vii) if $\int f \, d\mu < \infty$, then $f < \infty$ almost surely.

(e). Now, we extend the integral to measurable functions $f \colon \Omega \to \bar{\mathbb{R}}$ taking on positive and negative values. This requires some more care, for if a nonnegative function f equals ∞ on a set of positive measure, then we can set $\int f \, d\mu = \infty$, but this does not work for functions taking on really large positive and really small negative values, as we may end up in $\infty - \infty$ type trouble. Therefore, the integral of functions taking values in $\bar{\mathbb{R}}$ is only defined for those functions that are *integrable*. For a fully general measurable f, we may represent it as $f = f_+ - f_-$, where $f_+ = \max(f, 0)$ and $f_- = \max(-f, 0)$ are two nonnegative measurable functions (see Ex. A.3(f)). We say that f is integrable if $\int |f| \, d\mu$ is finite. For integrable functions f, we then define

$$\int f \, d\mu = \int f_+ \, d\mu - \int f_- \, d\mu.$$

Show that f is integrable if and only if $\int f_+ \, d\mu$ and $\int f_- \, d\mu$ are finite. Show also that if μ is a finite measure, and f is bounded, then f is integrable.

(f). If A is measurable and f is a measurable function, show that fI_A is measurable too. We can hence define $\int_A f \, d\mu = \int I_A f \, d\mu$, where, in the case that f takes both positive and negative values, we require that $I_A f$ is integrable.

(g). Assume that $f, g \colon \Omega \to \bar{\mathbb{R}}$ are two integrable, measurable functions, and let $a, b \in \mathbb{R}$ be constants. Show that $af + bf$ is measurable and integrable, and that

(i) $\int (af + bg)\,d\mu = a\int f\,d\mu + b\int g\,d\mu$;
(ii) if $f \le g$, then $\int f\,d\mu \le \int g\,d\mu$;
(iii) $|\int f\,d\mu| \le \int |f|\,d\mu$;

(h). Some more properties of the integral defined in (e). Let $f, g\colon \Omega \to \bar{\mathbb{R}}$ be measurable functions, and $A, B \in \mathcal{A}$. Show that (i) if $A \subset B$ are measurable sets, then $\int_A f\,d\mu \le \int_B f\,d\mu$; (ii) if $\mu(A) = 0$, then $\int_A f\,d\mu = 0$; (iii) if f is integrable and $f = 0$ almost surely, then $\int f\,d\mu = 0$; (iv) if $f = g$ almost surely, then $\int f\,d\mu = \int g\,d\mu$; and (v) if f and g are integrable and $\int_A f\,d\mu = \int_A g\,d\mu$ for all $A \in \mathcal{A}$, then $f = g$ almost surely.

(i). Let $f = g + ih$ be a measurable complex valued function defined on the measurable space $(\Omega, \mathcal{A}, \mu)$, where $g, h\colon \Omega \to \mathbb{R}$ are measurable functions (see Ex. A.3(d)). Since $|f| \le |g| + |h|$, we see that f is integrable, i.e., $\int |f|\,d\mu < \infty$, if g and h are. For an integrable $f\colon \Omega \to \mathbb{C}$ we define $\int f\,d\mu = \int g\,d\mu + i\int h\,d\mu$. Show that linearity of the integral, (g)(i), extends to the integral of complex valued functions, where now $a, b \in \mathbb{C}$; and that the inequality $|\int f\,d\mu| \le \int |f|\,d\mu$ also holds for complex valued functions.

Ex. A.11. *Convergence theorems.* One of the main objectives of the integration theory developed in the preceding exercises is to find general criteria for when $\lim_{n\to\infty} \int f_n\,d\mu = \int \lim_{n\to\infty} f_n\,d\mu$. The theorems that give various sets of conditions for when we can pass the limit under the integral sign, are the convergence theorems of measure theory. Remember that the measure μ can be any measure on any measurable space, so the theorems that follow are very general, they will, for example, apply to sums $\sum_{j=1}^{\infty} f_n(j)$, as well as to Riemann integrals $\int f_n(x)\,dx$.

Bounded convergence

(a). Suppose that $f_1, f_2, \dots$, is a sequence of measurable functions such that $|f_n(\omega)| \le M$ for all ω and n, that μ is a finite measure, and that $f_n \to f$ in measure. Show that $\lim_{n\to\infty} \int f_n\,d\mu = \int f\,d\mu$.

Monotone convergence

(b). Suppose that $f_1, f_2, \dots$ is an nondecreasing sequence of nonnegative measurable functions such that that $f_n \to f$ almost surely. Show that $\lim_{n\to\infty} \int f_n\,d\mu = \int f\,d\mu$. Show that it is enough to assume that $f_n \to f$ in measure.

Fatou's lemma

(c). Suppose $f_1, f_2, \dots$ is a sequence of nonnegative measurable functions. Show that

$$\int \liminf_{n\to\infty} f_n\,d\mu \le \liminf_{n\to\infty} \int f_n\,d\mu.$$

Assume in addition that $f_n \to f$ in measure. Show that $\int f\,d\mu \le \liminf_{n\to\infty} \int f_n\,d\mu$.

Dominated convergence

(d). Let $f_1, f_2, \dots$ be a sequence of measurable functions such that $f_n \to f$ almost surely. Suppose there is an integrable function g so that $|f_n| \le g$ almost surely for each n. Show that the limit function f is integrable, and $\lim_{n\to\infty} \int |f_n - f|\,d\mu = 0$ and that $\lim_{n\to\infty} \int f_n\,d\mu = \int f\,d\mu$.

(e). In (d), weaken the the assumption $f_n \to f$ almost surely, to $f_n \to f$ in measure, and show that the same conclusion holds.

(f). As a corollary to the Dominated convergence theorem, suppose that the dominating function g, instead of being merely integrable, is so that g^p is integrable for some $p \ge 1$. Show that then f^p is integrable, and $\lim_{n\to\infty} \int |f_n - f|^p\,d\mu = 0$.

(g). We extend the Dominated convergence theorem to complex valued functions. That is, suppose that the $f_1, f_2, \dots$ in (d) is a sequence of measurable complex valued functions (see Ex. A.10(i)), satisfying the conditions in (d). Show that the same conclusion holds.

(h). Suppose that $f_1, f_2, \dots$, is a sequence of measurable complex valued functions such that $|f_n(\omega)| \le M$ for all ω and n, that μ is a finite measure, and that $f_n \to f$ in measure. Show that $\lim_{n\to\infty} \int f_n \, \mathrm{d}\mu = \int f \, \mathrm{d}\mu$.

Ex. A.12. *More properties of the integral/Applications of the convergence theorems.* In this exercise we apply the convergence theorems of Ex. A.11 to work out a few more properties of the Lebesgue integral and touch on a proof strategy that appears again and again when working with the Lebesgue integral, so often that the strategy sometimes goes by the name of a bootstrapping argument (not to be confused with bootstrapping in statistics). All the functions in this exercise are defined on a measure space $(\Omega, \mathcal{A}, \mu)$.

(a). First, let $A_1 \subset A_2 \subset \cdots$ be sets in $\mathcal{A}$, and let $f\colon \Omega \to \bar{\mathbb{R}}$ be integrable over $\cup_{j=1}^{\infty} A_j$. Show that $\lim_{n\to\infty} \int_{A_n} f \, \mathrm{d}\mu = \int_A f \, \mathrm{d}\mu$ where $A = \cup_{j=1}^{\infty} A_j$. Second, let $B_1 \supset B_2 \supset \cdots$ be sets in $\mathcal{A}$, and assume that $f\colon \Omega \to \bar{\mathbb{R}}$ is integrable over B_1. Show that $\lim_{n\to\infty} \int_{B_n} f \, \mathrm{d}\mu = \int_B f \, \mathrm{d}\mu$, where $B = \cap_{n=1}^{\infty} B_n$.

(b). Let $f_1, f_2, \dots$ be a sequence of extended real valued measurable functions. They are nonnegative only if explicitly mentioned. (i) If $f_n \ge 0$ for each n, show that $\int \sum_{j=1}^{\infty} f_j \, \mathrm{d}\mu = \sum_{j=1}^{\infty} \int f_j \, \mathrm{d}\mu$. Here, both sides are either finite or infinite. (ii) If $\sum_{j=1}^{\infty} f_j < \infty$ almost surely, and each partial sum is bounded by the same integrable function g, then f_n and each partial sum is integrable, and $\int \sum_{j=1}^{\infty} f_j \, \mathrm{d}\mu = \int \sum_{j=1}^{\infty} f_j \, \mathrm{d}\mu$. (iii) If $\sum_{j=1}^{\infty} \int |f_j| \, \mathrm{d}\mu < \infty$, then $\sum_{j=1}^{\infty} |f_j| < \infty$ almost surely, this sum is integrable, and $\int \sum_{j=1}^{\infty} f_j \, \mathrm{d}\mu = \int \sum_{j=1}^{\infty} f_j \, \mathrm{d}\mu$. See Ex. A.36(c) for important applications of these results.

(c). Let $f_1, f_2, \dots$ be a sequence of measurable complex valued functions. Show that if $\sum_{n=1}^{\infty} \int |f_n| \, \mathrm{d}\mu < \infty$, then $\int \sum_{n=1}^{\infty} f_n \, \mathrm{d}\mu = \sum_{n=1}^{\infty} \int f_n \, \mathrm{d}\mu$.

(d). Deduce from (b) that if f is a nonnegative measurable function, then $\nu(A) = \int_A f \, \mathrm{d}\mu$ for $A \in \mathcal{A}$ defines a measure on $(\Omega, \mathcal{A})$. One says that f is the density of ν with respect to μ. Show that if $\mu(A) = 0$ then $\nu(A) = 0$. In general, a measure ν is said to be *absolutely continuous* with respect to μ if for all $A \in \mathcal{A}$, $\mu(A) = 0$ implies $\nu(A) = 0$; we also say that ν is *dominated* by μ, and denote this $\nu \ll \mu$.

absolute continuity $\nu \ll \mu$

(e). Let $\nu(A) = \int_A f \, \mathrm{d}\mu$ be the measure introduced in (d). We want to show that that

$$\int g \, \mathrm{d}\nu = \int gf \, \mathrm{d}\mu, \tag{A.2}$$

for any measurable function g. To do so, we emply a *bootstrapping argument*. It goes like this: First, prove (A.2) for indicator functions $g = I_A$. Second, extend (A.2) to simple functions $g = \sum_{j=1}^{k} a_j I_{A_j}$. Third, use Ex. A.2(g) and the monotone convergence theorem, and deduce that (A.2) holds for measurable functions $g\colon \Omega \to \bar{\mathbb{R}}_+$. Finally, using linearity again, show that (A.2) holds for all measurable functions $g\colon \Omega \to \bar{\mathbb{R}}$, provided g is integrable with respect to ν.

bootstrapping argument

(f). Let $(\Omega, \mathcal{A}, \mu)$ be a measure space, $(\mathcal{X}, \mathcal{B})$ a measurable space, and $f\colon \Omega \to \mathcal{X}$ and $g\colon \mathcal{X} \to \mathbb{R}$ be measurable functions. Show that

change of variable formula

$$\int_{f^{-1}(B)} g(f(\omega))\mathrm{d}\mu(\omega) = \int_B g(x)\,\mathrm{d}\mu f^{-1}(x),$$

for every $B \in \mathcal{B}$, provided either integral exists. You may once again use a bootstrapping argument to show this. For the measurability of $g(f(\cdot))$, see Ex. A.3(i).

(g). [xx make it neater. See Kolmogorov (1933) p. 44 English version xx] Let $(\mathcal{Y}, \mathcal{A})$ and $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ be measurable spaces. Suppose that $g\colon \mathcal{Y} \times \mathbb{R} \to \bar{\mathbb{R}}$ is such that $y \to g(y, \theta)$ is measurable and integrable for each $\theta \in (a, b) \subset \mathbb{R}$, and that $\partial g(y, \theta)/\partial\theta$ exists for all $\theta \in (a, b)$. Show that $\partial g(y, \theta)/\partial\theta$ is measurable. Suppose further that there is an integrable function $h(y)$ such that $|\partial g(y, \theta)/\partial\theta| \leq h(y)$ for all $\theta \in (a, b)$ and all y, combine the fundamental theorem of calculus and the dominated converge theorem to show that

Derivative under the integral sign

$$\frac{\mathrm{d}}{\mathrm{d}\theta} \int g(y, \theta)\,\mathrm{d}\mu(y) = \int \frac{\partial}{\partial\theta}\, g(y, \theta)\,\mathrm{d}\mu(y), \quad \text{for } \theta \in (a, b),$$

that is, we can pass the derivative under the integral sign.

(h). [xx Bartlett identity example here, perhaps xx]

(i). [xx move this to Ex. A.22 xx]Suppose $f\colon [a, b] \to \mathbb{R}$ is a Riemann integrable function. Show that f is measurable, and that its classical Riemann-definition integral $\int_a^b f(x)\,\mathrm{d}x$ coincides with the more general integral we've worked with in this exercise, $\int_{[a,b]} f(x)\,\mathrm{d}\lambda(x) = \int_{[a,b]} f\,\mathrm{d}\lambda$, with λ the Lebesgue measure defined on the Lebesgue subsets of $[a, b]$.

Riemann and Lebesgue

Ex. A.13. *Product measure and iterated integrals.* Let $(\mathcal{X}, \mathcal{A}, \mu)$ and $(\mathcal{Y}, \mathcal{B}, \nu)$ be two σ-finite measure spaces. The measurable space $(\mathcal{X} \times \mathcal{Y}, \mathcal{A} \otimes \mathcal{B})$ consists of the Cartesian product $\mathcal{X} \times \mathcal{Y}$ (see Ex. A.1(f)) and the σ-algebra $\mathcal{A} \otimes \mathcal{B}$ generated by the measurable rectangles $A \times B$, for $A \in \mathcal{A}$ and $B \in \mathcal{B}$. In this exercise we first construct a measure $\mu \times \nu$ on $\mathcal{A} \otimes \mathcal{B}$, such that $(\mu \times \nu)(A \times B) = \mu(A)\nu(B)$ for all measurable rectangles $A \times B$. This measure is called the product measure. Second, we establish conditions under which we can compute double integrals by iterated integration,

product measure

$$\int f\,\mathrm{d}(\mu \times \nu) = \int\int f(x, y)\,\mathrm{d}\nu(y)\,\mathrm{d}\mu(x) = \int\int f(x, y)\,\mathrm{d}\mu(x)\,\mathrm{d}\nu(y). \tag{A.3}$$

for measurable functions $f\colon \mathcal{X} \times \mathcal{Y} \to \bar{\mathbb{R}}$.

(a). Show that $\Pi = \{A \times B\colon A \in \mathcal{A}, B \in \mathcal{B}\}$ is a π-system generating $\mathcal{A} \otimes \mathcal{B}$.

(b). For (A.3) to make sense, we need a technical lemma. Assume that $f\colon \mathcal{X} \times \mathcal{Y} \to [0, \infty)$ is measurable, and that μ is σ-finite. Then (i) $x \mapsto f(x, y)$ is $\mathcal{A}$-measurable for each $y \in \mathcal{Y}$; and (ii) $\int_{\mathcal{X}} f(x, y)\,\mathrm{d}\mu(x)$ is $\mathcal{B}$-measurable. To prove (i) and (ii), assume that μ is finite, check (i) and (ii) for $f = I_C$ for $C \in \Pi$, then use Dynkin's lemma (Ex. A.5(c)) to show that (i) and (ii) holds for for $f = I_C$ for $C \in \mathcal{A} \otimes \mathcal{B}$. Third, use a bootstrapping argument, and, finally, extend what you have to μ being σ-finite.

(c). Due to our efforts in (b), it makes sense to define

$$\lambda(C) = \int\int I_C(x, y)\,\mathrm{d}\nu(y)\,\mathrm{d}\mu(x), \quad \text{for } C \in \mathcal{A} \otimes \mathcal{B}.$$

Show that λ is a measure, and argue that it is the unique measure such that $\lambda(A \times B) = \mu(A)\nu(B)$ on the π-system from (a). Conclude from this that (A.3) holds for indicator functions (so we could have defined λ with the order of integration reversed), and, indeed λ is the product measure $\mu \times \nu$.

(d). Combine what you found in (c) with the monotone convergence theorem (a boostrapping argument) to show that (A.3) holds for all measurable functions $f\colon \mathcal{X}\times\mathcal{Y}\to[0,\infty]$. Tonelli's theorem

(e). Suppose that $f\colon \mathcal{X}\times\mathcal{Y}\to\bar{\mathbb{R}}$ is integrable with respect to $\mu\times\nu$. Show that (A.3) holds also in this case. Notice that the sets $\{x\in\mathcal{X}\colon \int_{\mathcal{Y}}|f(x,y)|\,\mathrm{d}\nu(y)=\infty\}$ and $\{y\in\mathcal{Y}\colon \int_{\mathcal{Y}}|f(x,y)|\,\mathrm{d}\mu(x)=\infty\}$ have μ- and ν-measure zero, respectively (see Ex. A.10(d)), so you can modify f on these sets to avoid $\infty-\infty$ type trouble. Fubini's theorem

Ex. A.14. *Probability spaces.* Mathematically speaking, we are free to define the basics of probabilities, along with axioms these should satisfy, without yet tying these to the so-called real world. So let us define a *probability space* as a triple $(\Omega,\mathcal{A},P)$, where Ω is a set; $\mathcal{A}$ a σ-algebra of subsets of Ω; and $P\colon\mathcal{A}\to[0,1]$ a *probability measure*, defined simply to be a measure, in the sense of Ex. A.4, with full measure $P(\Omega)=1$. a probability space; a probability measure

We may envisage P as a probability machine, assessing to each A a probability $P(A)$. Such a probability measure on $(\Omega,\mathcal{A})$ has axiomatic properties following those of more general measures, given in Ex. A.2, and for convenience stated again here, for the present case of $P(\Omega)=1$. We demand axioms for a probability space

(i) that $P(A)\geq 0$ for all $A\in\mathcal{A}$;
(ii) that $\Pr(\cup_{i=1}^{\infty}A_i)=\sum_{i=1}^{\infty}\Pr(A_i)$ for are disjoint sets $A_1,A_2,\ldots$ in $\mathcal{A}$;
(iii) and that $P(\Omega)=1$.

The subsets A can be given several names, including *events*; the conceptual idea is that we do not yet know whether a certain A occurs or not, but we can give it a probability.

(a). Show that $\Pr(\emptyset)=0$; and that and that if $A\subset B$, then $\Pr(A)\leq\Pr(B)$.

(b). For all events A and B show that $\Pr(A\setminus B)=\Pr(A)-\Pr(A\cap B)$; that $\Pr(A)=1-\Pr(A^c)$; and that $\Pr(A\cup B)=\Pr(A)+\Pr(B)-\Pr(A\cap B)$. Generalise this latter formula to the union of three events, and then try to generalise it to the union of four or more events.

(c). If A and B have probabilities 0.95, or more, show that $\Pr(A\cap B)\geq 0.90$. Generalise. This simple lower-bounding of certain types of probabilities is sometimes called the Bonferroni method, or Bonferroni correction.

(d). From Ex. A.4(b) we know that if $A_1\subset A_2\subset\cdots$, then $\Pr(\cup_{n=1}^{\infty}A_n)=\lim_{n\to\infty}\Pr(A_n)$; and, secondly, if $A_1\supset A_2\supset\cdots$, then $\Pr(\cap_{n=1}^{\infty}A_n)=\lim_{n\to\infty}\Pr(A_n)$. Show that either of these two statements could replace (ii) in the axiom list above.

(e). Above we have been careful to define probability measures P for large collections of events, namely σ-algebras, but also avoiding defining $P(A)$ for *every* subset A. Attempting to do that, in various natural spaces, will lead to difficulties and incoherencies, related to the existence of non-measurable sets. These issues are not present when the full space Ω is finite, however, as one can simply allow *every subset* to be included, in the set of subsets for which a probability is attached. Show indeed that if $\Omega=\{\omega_1,\ldots,\omega_m\}$, with perhaps a large m, and these singletons are attached probabilities $p_1,\ldots,p_m$ (non-negative, with sum 1), then

$$\Pr(A)=\sum_{j\colon\omega_j\in A}p_j,\quad\text{for any subset } A,$$

defines a probability measure on $(\Omega, \mathcal{A})$, where, in this case, $\mathcal{A} = 2^\Omega$ the set of all 2^m subsets (which, from Ex. A.2(a), we know is a σ-algebra). Generalise to the case of countably big spaces, say $\Omega = \{\omega_1, \omega_2, \ldots\}$, with pointmasses $p_1, p_2, \ldots$ summing to 1. In these cases the collection $\mathcal{A}$ of *all* subsets is the natural set of events.

Ex. A.15. *Distribution functions.* Consider the case where the probability space is $(\mathbb{R}, \mathcal{B}, P)$, with $\mathcal{B} = \mathcal{B}(\mathbb{R})$ the Borel σ-algebra on the real line, and P is some probability measure on this measurable space. For such a P, define the *cumulative distribution function* (c.d.f. for short) as

the c.d.f., the cumulative distribution function

$$F(x) = P(-\infty, x], \quad \text{for } x \in \mathbb{R},$$

where we use the simpler notation $P(-\infty, x]$ instead of $P((-\infty, x])$.

(a). Show that F is (i) nondecreasing, i.e., $F(x) \le F(y)$ for all $x \le y$; (ii) right continuous with left hand limits, i.e., $F(x+\varepsilon) \to P(-\infty, x]$ and $F(x-\varepsilon) \to P(-\infty, x)$ as $\varepsilon \to 0$ from above; (iii) and that $F(x) \to 1$ and $F(x) \to 0$ as $x \to \infty$ and $x \to -\infty$, respectively. Any function F satisfying these three criteria is c.d.f. on the real line.

(b). Show that

$$P(a + 1/n, b - 1/n] = F(b - 1/n) - F(a + 1/n) \to F(b-) - F(a) = P[a, b),$$

for all intervals (a, b). Here $F(b-)$ is notation for the limit of $F(b-\varepsilon)$ as ε converges to zero from above, and is also the same as $P(-\infty, b)$.

(c). Show that $\Pr(\{t\})$, the probability assigned to the fixed point t, is $F(t) - F(t-)$. This probability is often zero, as is the case for all t if F is continuous. Show that the set D_F of discontinuities for F is at most countably infinite.

(d). Suppose P_1 and P_2 are two probability measures on $(\mathbb{R}, \mathcal{B})$, with the same c.d.f., i.e. $F_1 = F_2$. An important fact (to say the least) is that if $F_1 = F_2$, then indeed $P_1(A) = P_2(A)$ for *all* $A \in \mathcal{B}(\mathbb{R})$. Show this, from Carathéodory's Extension Theorem of Ex. A.6, or, alternatively from one of the theorems of Ex. A.5. Very conveniently, this allows one to define a full probability measure P by giving only its c.d.f., or its values for all intervals. For example, saying that $P(a, b) = \int_a^b (2\pi)^{-1/2} \exp(-\frac{1}{2}x^2)\, \mathrm{d}x$, for all intervals (a, b), is a sufficient description of the standard normal distribution; we don't need to give a more laborious recipe for how to compute $P(A)$ for more complicated events A.

(e). Suppose P is a probability measure on $(\mathbb{R}^2, \mathcal{B})$, where $\mathcal{B} = \mathcal{B}(\mathbb{R}^2)$ is the Borel-σ-algebra in the plane (see Ex. A.2(i)). Define the cumulative distribution function corresponding to P by

$$F(x_1, x_2) = P((-\infty, x_1] \times (-\infty, x_2]), \quad \text{for } x_1, x_2, \in \mathbb{R}.$$

Show that for any rectangle,

$$P((a_1, b_1] \times (a_2, b_2]) = F(a_1, a_2) - F(a_1, b_2) - F(a_2, b_1) + F(a_2, b_2).$$

Use again Ex. A.5, or indeed Carathéodory Extension Theorem of Ex. A.6, to prove that if two probability measures are equal for all rectangles, then they are identical, i.e., giving the same probability to any Borel set. Thus a probability measure P on $(\mathbb{R}^2, \mathcal{B})$ is fully determined by giving its $F(x_1, x_2)$ function.

(f). Attempt to generalise (e) to dimension k, i.e., to $(\mathbb{R}^k, \mathcal{B}(\mathbb{R}^k))$; in particular, the probability attached to a rectangle $(a_1, b_1] \times (a_k, b_k]$ can be expressed as a sum of values of F computed at the 2^k vertices of the rectangle, with ± 1 signs, as seen above for $k = 2$.

(g). [xx might move this to Ex. A.22 xx] Let P be a probability measure such that the distribution function $F(x) = P(-\infty, x]$ (see Ex. A.15(d)) is montonically increasing. Suppose that g is a real valued function such that the Riemann–Stieltjes integral $\int_a^b g \, \mathrm{d}F$ exists. Show that $\int_{[a,b]} g \, \mathrm{d}P = \int_a^b g \, \mathrm{d}F$.

Ex. A.16. *Random variables.* Speaking mathematically, a *random variable* is a measurable function on a probability space. Measurable functions were defined in Ex. A.3, but let us repeat some of the details. With $(\Omega, \mathcal{A}, \Pr)$ a 'background' probability space, we construct random variables as measurable functions $X \colon \Omega \to \mathcal{X}$, where $(\mathcal{X}, \mathcal{B})$ is the measurable space where $X(\omega)$ lands. Measurability means that the inverse images $X^{-1}(B) = \{\omega \in \Omega \colon X(\omega) \in B\}$ are in $\mathcal{A}$ for any $B \in \mathcal{B}$. As is common, we often write $\{X \in B\}$ instead of the more cumbersome $\{\omega \in \Omega \colon X(\omega) \in B\}$, and $\Pr(X \in B)$ instead of $\Pr(\{\omega \in \Omega \colon X(\omega) \in B\})$.

(a). The probability distribution, distribution, or law, of a random variable X, say P, is defined by

probability distribution

$$P(B) = \Pr(X \in B) = \Pr(X^{-1}(B)), \quad \text{for } B \in \mathcal{B}.$$

With pedantic care, we define P via $(\Pr X^{-1})(B) = \Pr(X^{-1}(B))$. Even though this is just repeating Ex. A.4(h) with the extra requirement that $P(\mathcal{X}) = 1$, show that $P = \Pr X^{-1}$ indeed is a probability measure on $(\mathcal{X}, \mathcal{B})$.

(b). Often what matters is the distribution of X, rather than particularities of the background space. Indeed there may be different spaces $(\Omega_j, \mathcal{A}_j, \Pr_j)$ and random variables $X_j \colon \Omega_j \to \mathcal{X}$ inducing precisely the same distribution, i.e., the different $P_j = \Pr_j X_j^{-1}$ might be identical. For a given P on $(\mathcal{X}, \mathcal{B})$, show that the identity map $x \mapsto x$ is one such construction, leading to a random variable X with distribution P. In the case of $X_j \colon \Omega_j \to \mathbb{R}$, we have seen in Ex. A.15 that what matters is the c.d.f. $\Pr_j(X_j \leq t) = P_j(-\infty, t] = F_j(t)$; as long as these are equal, the distributions $P_j = \Pr_j X_j^{-1}$ are identical. Give three separate such constructions of the standard normal distribution.

(c). If $X \colon \Omega \to \bar{\mathbb{R}}$ is a random variable, defined on a background probability space $(\Omega, \mathcal{A}, \Pr)$, its *mean*, or *expected value*, is defined as

the expectation

$$\mathrm{E}\, X = \int X \, \mathrm{d}\Pr = \int X(\omega) \, \mathrm{d}\Pr(\omega),$$

as long as this integral is finite, i.e., X is integrable. Since the expectation is a Lebesgue integral (as defined in Ex. A.10), the convergence theorems of Ex. A.11, as well as the properties of the Lebesgue integral derived in Ex. A.10(g)–(h) and Ex. A.12 apply. In particular, with $g \colon \mathbb{R} \to \bar{\mathbb{R}}$ any measurable function, deduce from Ex. A.12(f) that

$$\mathrm{E}\, g(X) = \int g(X(\omega)) \, \mathrm{d}\Pr(\omega) = \int g(x) \, \mathrm{d}P(x), \quad \text{with } P = \Pr X^{-1},$$

provided g is P-integrable. In particular, only the distribution P of X matters, not the details associated with the background probability space.

(d). With the mean of a real random variable well defined, we may of course go on to other and higher moments. For a random variable $X\colon \Omega \to \bar{\mathbb{R}}$, as above, with $\xi = \mathrm{E}\, X$, show that

the variance

$$\mathrm{E}\,(X-\xi)^2 = \int (X-\xi)^2 \,\mathrm{dPr} = \int_{-\infty}^{\infty} (x-\xi)^2 \,\mathrm{d}P(x) = \int_0^{\infty} y \,\mathrm{d}Q(y),$$

with Q the distribution of $Y = (X-\xi)^2$; so there's no ambiguity. This quantity is of course *the variance of* X, denoted $\mathrm{Var}\, X = \mathrm{E}\,(X-\xi)^2$. The square root of $\mathrm{Var}\, X$ is called the *standard deviation* of X.

(e). Consider integrable random variables $X, Y\colon \Omega \to \bar{\mathbb{R}}$ defined on the same background probability space $(\Omega, \mathcal{A}, \mathrm{Pr})$. Show that $\mathrm{E}\,(aX+bY) = a\,\mathrm{E}\, X + b\,\mathrm{E}\, Y$ (see Ex. A.10(g)), and generalise . In particular, for integrable random variables $X_1, \ldots, X_n$, we have $\mathrm{E}\,(X_1 + \cdots + X_n) = \mathrm{E}\, X_1 + \cdots + \mathrm{E}\, X_n$, regardless of any dependencies between these variables.

Ex. A.17. *Types of distribution functions.* Let F be the c.d.f. of the distribution P on $(\mathbb{R}, \mathcal{B})$, i.e., $F(x) = P(-\infty, x]$ for all $x \in \mathbb{R}$. If F is piecewise constant, only changing its values at the (countably many) points $x_1, x_2, \ldots$, then P is a *discrete* probability distribution. If F can be expressed as $F(x) = \int_{-\infty}^{x} f(y)\,\mathrm{d}y$ for some nonnegative function f, then F is absolutely continuous, and the distribution P is absolutely continuous with respect to Lebesgue measure, with density f (see Ex. A.12(d) with more in Ex. A.18). As we will see, these two types do not exhaust all possibilities.

discrete distributions

absolutely continuous distributions

(a). Let F be a piecewise constant c.d.f. only changing its values at the points in $\mathcal{X} = \{x_1, x_2, \ldots\}$. Show that there is a nonnegative function f such that

$$F(x) = \sum_{y \in \mathcal{X}, y \le x} f(y) = \int_{A_x} f(y)\,\mathrm{d}\mu(y), \quad \text{for all } x \in \mathbb{R}.$$

where $A_x = (-\infty, x]$ and μ is counting measure on $(\mathbb{R}, \mathcal{B})$. More generally, show that $P(A) = \int_A f(y)\,\mathrm{d}\mu(y)$ for $A \in \mathcal{B}$ is a measure, and that $F(x) = P(-\infty, x]$. Perhaps a more natural definition, is to take μ as counting measure on $(\mathcal{X}, 2^{\mathcal{X}})$, so that μ is σ-finite, and $F(x) = \int_{A_x} f(y)\,\mathrm{d}\mu(y)$ where now $A_x = \{y \in \mathcal{X}\colon y \le x\}$. – The function f is the density of P with respect to counting measure, and is often called a probability mass function (p.m.f.).

p.m.f.

(b). When F is absolutely continuous, its density is simply the derivative of the cumulative distribution function, $f(x) = F'(x)$, and from the fundamental theorem of calculus,

$$F(b) - F(a) = \int_a^b f(x)\,\mathrm{d}x, \quad \text{for all } [a, b].$$

Show that $P(A) = \int_A f(x)\,\mathrm{d}x$, for all $A \in \mathcal{B}$, i.e. not merely for intervals. Giving $f(x)$, instead of the cumulative $F(x)$, is the most convenient (and traditional) way in which to define a probability distribution.

(c). Let F_{ac} and F_{d} be an absolutely continuous and a piecewise constant c.d.f., respectively. Define $F = \alpha F_{\mathrm{ac}} + (1-\alpha) F_{\mathrm{d}}$ for a $0 < \alpha < 1$, and show that F is a c.d.f. Find a measure ν and a density f such that the distribution P associated with F can be written $P(A) = \int_A f\,\mathrm{d}\nu$ for all $A \in \mathcal{B}$. Notice that the measure ν is σ-finite.

mixture distribution

(d). [xx something on the Cantor function and singular measures. See Example 31.1 p. 407 in Billingsley, og Problem 31.2. xx]

(e). [xx Kan ha noe knyttet til Problem 31.22 p. 419 og Prob 32.13 p. 427 i Billingsley. Se også Kolmogorov (1933) p. 51 nederst, English version xx]

Ex. A.18. *Probability densities.* We have seen in Ex. A.15 that probability measures on the real line are fully characterised by the cumulative distribution functions. Very often it is more practical and satisfying to define a probability distribution via its density, provided it exists. As we saw in Ex. A.17 these may be defined not only in familiar situations with continuous distributions, but with discrete data, and with measures having both continuous and discrete components. Now, if μ is a measure on $(\mathcal{X}, \mathcal{B})$ and f is a nonnegative function such that $\int f \, \mathrm{d}\mu = 1$, then it is immediate from Ex. A.12(d) that $P(A) = \int_A f \, \mathrm{d}\mu$, $A \in \mathcal{B}$ is a probability measure on $(\mathcal{X}, \mathcal{B})$. The Radon–Nikodym theorem goes in the opposite direction.

probability density function

(a). Let P be a probability measure and μ a σ-finite measure, and suppose that P is dominated by μ, i.e., $P \ll \mu$ as defined in Ex. A.12(d). Under these conditions, the Radon–Nikodym theorem says that there is a nonnegative function f, such that

$$P(A) = \int_A f(x) \, \mathrm{d}\mu(x) \quad \text{for all } A \in \mathcal{B}. \tag{A.4}$$

the Radon–Nikodym theorem

The function f, called the *density* of P with respect to μ, is often denoted $\mathrm{d}P/\mathrm{d}\mu$, to remind us that this is a density of P with respect to μ. Look back at Ex. A.10(d) and explain why the density $f = \mathrm{d}\nu/\mathrm{d}\mu$ is only unique μ-almost surely. Explain why Ex. A.17(b), where F has a derivative and is the integral of this derivative, matches this more general setup, where μ is the Lebesgue measure, with $\mu(a, b) = b - a$ for all intervals.

(b). The strength of the general $f = \mathrm{d}P/\mathrm{d}\mu$ machinery above is that it can be fruitfully used for large classes of other probability measures too, not only for those which are dominated by the Lebesgue measure. The dominating measure is often chosen by mathematical convenience, to match the situation at hand. For the Poisson and other distributions, with random variables landing in $\mathcal{X} = \{0, 1, 2, \ldots\}$, consider, for any subset of $\mathcal{X}$, $\mu(A)$ equal to the number of numbers $j \in A$, that is, μ is the counting measure on the integers, which, from Ex. A.2(g), we know is a σ-finite measure. Show that with P having a Poisson distribution P, with mean θ, that there is a density $f = \mathrm{d}P/\mathrm{d}\mu$, given by $f(x) = \exp(-\theta)\theta^x/x!$ for $x = 0, 1, 2, \ldots$, in the sense given above.

(c). Consider a probability measure P on $[0, 1]$ with probabilities 0.1 and 0.1 at positions 0 and 1, and which has $P(a, b) = 0.8\,(b - a)$ for (a, b) inside $(0, 1)$. Thus P is not continuous, and not discrete, but a mixture. Show that P is dominated by the measure μ, which has pointmasses 1 and 1 at the points 0 and 1, and is uniform inside $(0, 1)$. Find the probability density $f(x) = \mathrm{d}P(x)/\mathrm{d}\mu$.

(d). [xx but we have yet to introduce rv's xx] Suppose P is dominated by a σ-finite μ, with $f(x) = \mathrm{d}P/\mathrm{d}\mu$ the probability density. With X having distribution P, and $g(x)$ being a function for which the mean is finite (with respect to P), we can now take the change of variable formula from Ex. A.16**??** one step further, so to speak. With the help of Ex. A.12(e), show that

$$\mathrm{E}\, g(X) = \int g(x) \, \mathrm{d}P(x) = \int g(x) \, \frac{\mathrm{d}P(x)}{\mathrm{d}\mu} \, \mathrm{d}\mu(x) = \int g(x) f(x) \, \mathrm{d}\mu(x).$$

In particular, if μ is Lebesgue measure on the real line, we get from Ex. A.12(g) that $\mathrm{E}\, g(X) = \int g(x)f(x)\,\mathrm{d}x$, and we are back to the expectation formula from introductory statistics courses.

(e). (xx round this off. drive home that this makes it possible and convenient to derive results in a general manner, point to Cramér–Rao, which we need to redo, as of 12-March-2025, and also that we can handle any type of mixed distributions, not merely the classic ones, the continuous and the discrete. ask for the mean and variance of the 0.1, 0.1, 0.8 distribution above. xx)

Ex. A.19. *The Hahn and the Lebesgue decompositions.*

Ex. A.20. *The Radon–Nikodym theorem.* [xx proof here xx] Let $(\mathcal{X}, \mathcal{A})$ be a measure space. In Ex. A.12(d) we saw that if μ is a measure, and f a nonnegative function, then $\nu(A) = \int_A f\,\mathrm{d}\mu$ for $A \in \mathcal{A}$ is a measure on $\mathcal{A}$, and that $\mu(A) = 0$ implies $\nu(A) = 0$, i.e., $\nu \ll \mu$. The Radon–Nikodym theorem provides a converese to this result: If ν and μ are σ-finite measures such that $\nu \ll \mu$, then there is a nonnegative function f such that

$$\nu(A) = \int_A f(x)\,\mathrm{d}\mu(x) \quad \text{for all } A \in \mathcal{A}.$$

The function f, called the *density* of ν with respect to μ, is often denoted $\mathrm{d}\nu/\mathrm{d}\mu$, to remind us that this is a density of ν with respect to μ. If $\nu = P$ is a probability measure, then $f = \mathrm{d}P/\mathrm{d}\nu$ is a probability density function. – There are several proofs of the Radon–Nikodym theorem. The one we work through in the following few exercises relies on the notion of signed measures and the Hahn decomposition. A signed measure is a set function $s\colon \mathcal{A} \to \bar{\mathbb{R}}$ such that $s(\emptyset) = 0$ and $s(\cup_j A_j) = \sum_j s(A_j)$ for any sequence of disjoint sets in $\mathcal{A}$. The Hahn decomposition is a theorem stating that for any signed measure s, there are disjoint sets A^+ and A^- in $\mathcal{A}$ such that $A^+ \cup A^- = \mathcal{X}$ and $s(A) \geq 0$ for $A \subset A^+$ and $s(A) \leq 0$ for $A \subset A^-$.

(a).

Ex. A.21. *Radon–Nikodym derivatives* Let P, Q, and μ be σ-finite measures on the measure space $(\Omega, \mathcal{A})$.

(a). Show that if $Q \ll P$ and $P \ll \mu$, then $Q \ll \mu$, and that we have the following relation of Radon–Nikodym derivatives,

$$\frac{\mathrm{d}Q}{\mathrm{d}\mu} = \frac{\mathrm{d}Q}{\mathrm{d}P}\frac{\mathrm{d}P}{\mathrm{d}\mu}, \quad \mu\text{-almost surely.}$$

The second part here is, to a certain extent, Ex. A.12(e) in new dressing, and you might use that exercise to prove it.

(b). Show that if $Q \ll P$, then the density $\mathrm{d}Q/\mathrm{d}P$ is positive Q-almost surely. Next, show that if $Q \ll P$ and $P \ll Q$, then

$$\frac{\mathrm{d}P}{\mathrm{d}Q} = I(\mathrm{d}Q/\mathrm{d}P > 0)\Big(\frac{\mathrm{d}Q}{\mathrm{d}P}\Big)^{-1}, \quad \text{almost surely,}$$

with respect to both Q and P. You may again use Ex. A.12(e).

(c). Suppose that $Q \ll \mu$ and $P \ll \mu$. Let $N = \{\mathrm{d}P/\mathrm{d}\mu = 0\}$. Show that there is a measurable function $f \geq 0$ such that

$$Q(A) = \int_A f \,\mathrm{d}P + Q(A \cap N).$$

The pair (f, N) is called the Lebesgue decomposition of Q with respect to P. If P and Q are two σ-finite measures, in particular probability measures, then $Q + P$ is σ-finite and $Q \ll Q + P$ and $P \ll Q + P$, and a Lebesgue decomposition of Q with respect to P exists. Such constructions play an important role in parts of Ch. 2. Lebesgue decomposition

(d). Suppose that $Q \ll \mu$ and $P \ll \mu$. Show that $Q \ll P$ if and only if the set $\{\mathrm{d}Q/\mathrm{d}\mu > 0\} \cap \{\mathrm{d}P/\mathrm{d}\mu = 0\}$ has measure zero under μ. Show also that if this is the case, i.e., if $\mu(\{\mathrm{d}Q/\mathrm{d}\mu > 0\} \cap \{\mathrm{d}P/\mathrm{d}\mu = 0\}) = 0$, then

$$\frac{\mathrm{d}Q}{\mathrm{d}P} = \frac{\mathrm{d}Q/\mathrm{d}\mu}{\mathrm{d}P/\mathrm{d}\mu} I(\mathrm{d}P/\mathrm{d}\mu > 0), \quad P\text{-almost surely.}$$

(e).

(f). Let P_ξ be the $\mathrm{N}(\xi, 1)$ distributions, and $P_{\xi+n^{-1/2}h}$ be the $\mathrm{N}(\xi + n^{-1/2}h, 1)$ distribution, e.g., P_ξ has density $f(x;\xi) = (1/\sqrt{2\pi})\exp(-\frac{1}{2}(x-\xi)^2)$ with respect to Lebesgue measure on the real line. Show that $P_{\xi+n^{-1/2}h} \ll P_\xi$; that

$$\frac{\mathrm{d}P_{\xi+n^{-1/2}h}}{\mathrm{d}P_\xi}(x) = \exp\Big(\frac{h}{\sqrt{n}}(x-\xi) - \frac{1}{2}\frac{h^2}{n}\Big),$$

and that $\mathrm{E}_\xi\,(\mathrm{d}P_{\xi+n^{-1/2}h}/\mathrm{d}P_\xi)(X) = 1$, i.e., when the expectation is taken under P_ξ.

(g). Let $(\mathcal{X}, \mathcal{B}_1, \mu)$ and $(\mathcal{Y}, \mathcal{B}_2, \nu)$ be two σ-finite measurable spaces. Suppose that $X\colon \Omega \to \mathcal{X}$ and $Y\colon \Omega \to \mathcal{Y}$ are independent random variables with distributions P_1 and P_2, and that $P_1 \ll \mu$ and $P_2 \ll \nu$. Denote the densities f_1 and f_2. Show that the distribution of (X, Y) has density $f_1 f_2$ with respect to the product measure $\mu \times \nu$ on $(\mathcal{X} \times \mathcal{Y}, \mathcal{B}_1 \otimes \mathcal{B}_2)$ (see Ex. A.24(f)). Generalise to higher dimensions. This result is fundamental to the likelihood theory we cover in Chapter 5.

Ex. A.22. *Lebesgue and Riemann–Stieltjes.* [xx fix this one. can then remove Ex. A.15(g) and Ex. A.12(g) xx] In this exercise we make the connection between the Lebesgue integral introduced above and the Riemann integral we know from calculus. Let $\mathcal{P}$ denote the collection of all partitions $a = x_0 \leq x_1 \leq \cdots \leq x_{n-1} \leq x_n = b$ of the interval $[a, b]$, and let F be a monotonically increasing function on $[a, b]$. For a given partition $P = \{x_0, x_1, \ldots, x_n\} \in \mathcal{P}$, and a real valued function g that is bounded on $[a, b]$, define $M_i = \sup_{x\in[x_{i-1},x_i]} g(x)$, $m_i = \inf_{x\in[x_{i-1},x_i]} g(x)$, and set $U(P, g, F) = \sum_{i=1}^n M_i\{F(x_i) - F(x_{i-1})\}$ and $L(P, g, F) = \sum_{i=1}^n m_i\{F(x_i) - F(x_{i-1})\}$. If $\sup_{P\in\mathcal{P}} L(P, g, F) = \inf_{P\in\mathcal{P}} U(P, g, F)$, their common value, the Riemann–Stieltjes integral, is denoted $\int_a^b g \,\mathrm{d}F$ or $\int_a^b g(x)\,\mathrm{d}F(x)$. If $F(x) = x$, this is the Riemann integral, and we write $\int_a^b g\,\mathrm{d}x$ or $\int_a^b g(x)\,\mathrm{d}x$.

(a). Show that if $g\colon [a, b] \to \mathbb{R}$ is a bounded and Riemann integrable function on $[a, b]$, then g is Lebesgue measurable, and

$$\int_{[a,b]} g \,\mathrm{d}\lambda = \int_a^b g \,\mathrm{d}x,$$

where λ is Lebesgue measure on the Lebesgue subsets of the real line [xx need pointer here xx]. Thus, the Lebesgue integral and the Riemann integral coincide.

(b). Let P be a probability measure such that the distribution function $F(x) = P(-\infty, x]$ (see Ex. A.15(d)) is montonically increasing. Suppose that g is a real valued function such that the Riemann–Stieltjes integral $\int_a^b g\,\mathrm{d}F$ exists. Show that $\int_{[a,b]} g\,\mathrm{d}P = \int_a^b g\,\mathrm{d}F$.

If g and F are so that (**??**) exists, we say that g is Riemann integrable with respect to F. In what follows we assume that $g\colon [a,b] \to \mathbb{R}$ is bounded and that $F\colon [a,b] \to \mathbb{R}$ is monotonically increasing.

(c). Let g be as described above, i.e., there are numbers l and u so that $l \le g \le u$. Show that $l(b-a) \le L(P,g,F) \le U(P,g,F) \le u(b-a)$ for every partition P. Show also that if P and P' are partitions and $P \subset P'$, then $U(P,g,F) \le U(P',g,F)$ and $L(P',g,F) \le L(P,g,F)$.

(d). Show that g is Riemann integrable with respect to F if and only if we for every $\varepsilon > 0$ can find a partition P so that $U(P,g,F) - L(P,g,F) < \varepsilon$. Show also that if this inequality holds for some partition P, then it holds for all partitions containing P.

(e). Given $\varepsilon > 0$, suppose that $P = \{x_0, x_1, \dots, x_n\}$ is a partition so that the inequality in (d) holds. For $i = 1, \dots, n$ let $y_i \in [x_i, x_{i-1}]$ be arbitrary points. Show that $|\sum_{i=1}^n f(y_i)\Delta F_i - \int_a^b f\,\mathrm{d}F| < \varepsilon$.

(f). Show that g is Riemann integrable with respect to F if (i) g is continuous on $[a,b]$; (ii) g is monotonic and F is continuous on $[a,b]$; (iii) g is bounded on $[a,b]$, has only finitely many points of discontinuity, and F is continuous at these discontinuity points.

(g). Let $(\mathbb{R}, \mathcal{B}, P)$ be a probability space, and define the distribution function $F(x) = P(-\infty, x]$ (see Ex. A.15(d)). Suppose that the real valued function g is bounded and Riemann integrable with respect to F on $[a,b]$. Show that

$$\int_{[a,b]} g\,\mathrm{d}P = \int_a^b g\,\mathrm{d}F,$$

where the integral on the left is the Lebesgue integral, while the integral on the right is the Riemann–Stieltjes integral.

Ex. A.23. *Convolutions.* Let X and Y be independent real random variables with distributions P_X and P_Y, and cumulative distribution functions $F_X(x) = P_X(-\infty, x]$ and $F_Y(y) = P_Y(-\infty, y]$.

(a). For $B \in \mathcal{B}(\mathbb{R})$ write $B - x = \{b - x\colon b \in B\} = \{y \in \mathbb{R}\colon x + y \in B\}$, and use Ex. A.13(d) to show that

$$\Pr(X + Y \in B) = \int_{\mathbb{R}} P_Y(B - x)\,\mathrm{d}P_X(x) = \int_{\mathbb{R}} P_X(B - y)\,\mathrm{d}P_Y(y).$$

This defines the measure $P_Y * P_X$ on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$, called the *convolution* of P_X and P_Y, i.e., $(P_X * P_Y)(B) = \int_{\mathbb{R}} P_Y(B-x)\,\mathrm{d}P_X(x)$, for $B \in \mathcal{B}(\mathbb{R})$. Note also that the convolution is commutative, i.e., $P_Y * P_X = P_X * P_Y$.

(b). Show that $Z = X + Y$ has cumulative distribution function

$$H(z) = \int F_Y(z - x)\,\mathrm{d}F_X(x) = \int F_X(z - y)\,\mathrm{d}F_Y(y).$$

Show also that if P_X and P_Y have densities, say f_X and f_Y, with respect to Lebesgue meaures on the real line, then the distribution of Z, namely $P_Y * P_X$, has density

$$h(z) = \int_{\mathbb{R}} f_Y(z-x) f_X(x)\,\mathrm{d}x = \int_{\mathbb{R}} f_X(z-y) f_Y(y)\,\mathrm{d}y,$$

also, as we see, with respect to Lebesgue measure. The density h is often denoted $(f_X * f_Y)(z)$, which is called the convolution of f_X and f_Y.

(c). (xx just a bit more, with X discrete and Y having a density. an illustration. also pointers to mgf things and CLT etc. xx)

(d). If $f, g\colon \mathbb{R} \to \mathbb{R}$ are two functions such that $\int g(z-x)f(x)\,\mathrm{d}x$ is finite for all z, we may also define the convolution $(g*f)(z) = \int g(z-x)f(x)\,\mathrm{d}x = \int f(z-y)g(y)\,\mathrm{d}y$ for $z \in \mathbb{R}$. Show that if g is bounded and continuously differentiable with a bounded derivative g', and $\int f(x)\,\mathrm{d}x$ is finite, then $(f*g)(z)$ is also bounded and continuously differentiable, with derivative

$$(g*f)'(z) = (g'*f)(z) = \int g'(z-x)f(x)\,\mathrm{d}x.$$

You may look at Ex. A.12(g) to find weaker conditions under which the equality above holds. Notice also that $f(x)\,\mathrm{d}x$ may be replaced by any finite measure, i.e., the existence of a density is not used.

(e). As a corollary to (d), suppose that g is k times continuously differentiable function that vanishes outside of a compact set (i.e., g is nonzero only on a closed and bounded set), and that f is as before. Show that $(g*f)$ is also k times continuously differentiable, with derivatives $(g*f)^{(j)} = (g^{(j)} * f)$ for $j \le k$. Moreover, show that if f also vanishes outside of a compact set (which need not be the same as for g), then $g*f$ vanishes outside of a compact set.

(f). Let g be as in (e), and suppose that f is the density of the uniform distribution on $[a,b]$, i.e., $f(x) = 1/(b-a)$ for $x \in [a,b]$, and $f(x) = 0$ elsewhere. Show that $(g*f)$ has one more continuous derivative than g, and that these derivatives take the form

$$(g*f)^{(j+1)}(z) = \frac{g^{(j)}(z-a) - g^{(j)}(z-b)}{b-a}, \quad \text{for } j = 1, \dots, k.$$

(g). We'll use our findings above to prove the existence of a infinitely smooth density function with support $[-1,1]$. Let $U_1, U_2, \dots$ be i.i.d. uniforms on $[-1,1]$. Show that

$$Y_3 = \frac{U_1}{2} + \frac{U_2}{4} + \frac{U_3}{8},$$

has a density, say f_3, that is one time continuously differentiable. Proceed by induction to show that the density of $Y_m = \sum_{n=1}^m U_n/2^n$ is $m-2$ times continuously differentiable. Finally, argue that the density of

$$Y = \sum_{n=1}^{\infty} U_n/2^n = \sum_{n=1}^{m} U_n/2^n + \sum_{n \ge n+1} U_n/2^n,$$

is infinitely smooth, i.e., has infinitely many continuous derivatives.

Ex. A.24. *Independence.* Here we define and work through basic properties of *independence*, for events and for random variables.

(a). For a probability space $(\Omega, \mathcal{A}, \Pr)$, we start out saying that two events A and B are independent if $\Pr(A \cap B) = \Pr(A)\Pr(B)$. Show that then also A and B^c are independent, A^c and B are independent, and A^c and B^c are independent. Show that all events are independent of the emptyset and of the full set Ω.

(b). Try to exhibit an example, with a finite Ω, of events A, B, C such that A and B are independent, A and C are independent, B and C are independent, but where $\Pr(A \cap B \cap C) \neq \Pr(A)\Pr(B)\Pr(C)$. Hence care is needed when defining independence for more than two events. We say that $A_1, \ldots, A_n$ are independent if

$$\Pr(A_{i_1} \cap \cdots A_{i_k}) = \Pr(A_{i_1}) \cdots \Pr(A_{i_k}), \quad \text{for any } \{i_1, \ldots, i_k\} \subset \{1, \ldots, n\}.$$

Show that this is equivalent to requiring that $\Pr(B_1 \cap \cdots \cap B_n) = \Pr(B_1) \cdots \Pr(B_n)$ for the 2^n such equations obtained by setting $B_j = A_j$ or $B_j = A_j^c$.

(c). The definition in (b) extends to countably infinite many events. The events $A_1, A_2, \ldots$ are independent if for any finite number of distinct indices $i_1, \ldots, i_n \in \{1, 2, \ldots\}$, the events $A_{i_1}, \ldots, A_{i_n}$ are independent, in the sense defined in (b). We say that the sub-σ-algebras $\mathcal{G}_1, \mathcal{G}_2, \ldots$ of $\mathcal{A}$ are independent if for arbitrary representatives $A_1 \in \mathcal{G}_1$, $A_2 \in \mathcal{G}_2$, and so on, the events $A_1, A_2, \ldots$ are independent. Suppose that Π_1 and Π_2 are two π-systems such that $\Pr(A_1 \cap A_2) = \Pr(A_1)\Pr(A_2)$ whenever $A_1 \in \Pi_1$ and $A_2 \in \Pi_2$. Use Dynkin's lemma (see Ex. A.5(c)) to show that $\sigma(\Pi_1)$ and $\sigma(\Pi_2)$ are independent. Generalise to π-systems $\Pi_1, \Pi_2, \ldots$ and σ-algebras $\sigma(\Pi_1), \sigma(\Pi_2), \ldots$.

independent random variables

(d). Consider random variables $X_1, X_2, \ldots$ defined on the probability space $(\Omega, \mathcal{A}, \Pr)$. We say that $X_1, X_2, \ldots$ are independent if the σ-algebras they generate, $\sigma(X_1), \sigma(X_1), \ldots$, are independent in the sense defined in (c). Suppose that X and Y are two random variables defined $(\Omega, \mathcal{A}, \Pr)$, and that

$$\Pr(X \leq x, Y \leq y) = \Pr(X \leq x)\Pr(Y \leq y), \quad \text{for all } x, y \in \mathbb{R},$$

i.e., their joint c.d.f. equals the product of the marginal c.d.f.'s. Use the result from (c) to show that X and Y are independent.

(e). Consider the probability space $(\{1, 2, 3, 4, 5, 6\}, 2^\Omega, \Pr)$, where Pr is the uniform probability measure on this space. Define the random variables $X = I_{\{1,3,5\}}$ and $Y = I_{\{5,6\}}$. Write down the full σ-algebras $\sigma(X)$ and $\sigma(Y)$, and show that, indeed, $\Pr(A \cap B) = \Pr(A)\Pr(B)$ for all $A \in \sigma(X)$ and $B \in \sigma(Y)$.

(f). Let $(\mathcal{X}, \mathcal{A})$ and $(\mathcal{Y}, \mathcal{B})$ be measurable spaces, and let $X \colon \Omega \to \mathcal{X}$ and $Y \colon \Omega \to \mathcal{Y}$ be random variables on the same underlying probability space, with distributions P_1 and P_2, respectively. From Ex. A.13 we know that the measure $P_1 \times P_2$ on, defined by

$$(P_1 \times P_2)(C) = \int_{\mathcal{X}} \int_{\mathcal{Y}} I_C(x, y) \, \mathrm{d}P_2(y) \, \mathrm{d}P_1(x) = \int_{\mathcal{Y}} \int_{\mathcal{X}} I_C(x, y) \, \mathrm{d}P_1(x) \, \mathrm{d}P_2(y),$$

for $C \in \mathcal{A} \otimes \mathcal{B}$, is the unique probability measure on $(\mathcal{X} \times \mathcal{Y}, \mathcal{A} \otimes \mathcal{B})$ such that $(P_1 \times P_2)(A \times B) = P_1(A)P_2(B)$ for all measurable rectangles $A \times B$. Show that X and Y are independent if and only if (X, Y) has distribution $P_1 \times P_2$.

(g). So $(P_1 \times P_2)(C) = \Pr((X,Y) \in C)$ is now properly defined for much more complicated sets than the direct product sets $A \times B$. Let X and Y be independent, both with uniform distributions on $[-1,1]$, with subintervals of equal length having the same probability. Find the probability that (X,Y) lands inside the unit circle.

(h). Show that if X and Y are independent random variables, and g and h are measurable functions, then $g(X)$ and $h(Y)$ are also independent. Use a bootstrapping argument (see Ex. A.12(e)) to show that

$$\mathrm{E}\, g(X)h(Y) = \mathrm{E}\, g(X)\, \mathrm{E}\, h(Y),$$

provided $g(X)$ and $h(Y)$ are integrable. The *covariance* of two random variables W and Z with means $\xi_W = \mathrm{E}\, W$ and $\xi_Z = \mathrm{E}\, Z$, respectively, is defined by covariance

$$\mathrm{cov}(W,Z) = \mathrm{E}\,(W - \xi_W)(Z - \xi_Z).$$

Show that $\mathrm{cov}(W,Z) = \mathrm{E}\,(WZ) - \xi_W \xi_Z$, and conclude that if W and Z are independent, then $\mathrm{cov}(W,Z) = 0$.

(i). We must of course extend (f) and (h) to the case of more than two independent random variables. With $X_1, \ldots, X_n$ defined on the same underlying probability space $(\Omega, \mathcal{A}, \Pr)$, their distributions are $P_1 = \Pr X_1^{-1}, \ldots, P_n = \Pr X_n^{-1}$. Suppose that these random variables are independent, i.e., $\sigma(X_1), \ldots, \sigma(X_n)$ are independent σ-algebras (see definition in (d)), and show that this is equivalent to

$$\Pr(X_1 \le x_1, \ldots, X_n \le x_n) = P_1(-\infty, x_1] \cdots P_n(-\infty, x_n], \quad \text{for all } x_1, \ldots, x_n \in \mathbb{R},$$

that is, again, the joint c.d.f. equals the product of the marginal c.d.f.'s. Generalise the construction in (f): For $j = 1, \ldots, n$, suppose that $X_j \colon \Omega \to \mathcal{X}_j$ where $(\mathcal{X}_j, \mathcal{B}_j)$ are measurable spaces, and show that this gives rise to a well-defined product probability measure $Q = P_1 \times \cdots \times P_n$ on the σ-algebra $\mathcal{B}_1 \otimes \cdots \otimes \mathcal{B}_n$, generated by all rectangles $B_1 \times \cdots \times B_n$, where $B_j \in \mathcal{B}_j$ for $j = 1, \ldots, n$. Generalise also the if and only if claim from (f) to this higher dimensional setting.

(j). Show that if $X_1, \ldots, X_n$ are independent, and $g_1, \ldots, g_k$ are measurable functions, then also $g_1(X_1), \ldots, g_k(X_k)$ are independent. Show also that

$$\mathrm{E}\, g_1(X_1) \cdots g_k(X_k) = \mathrm{E}\, g_1(X_1) \cdots \mathrm{E}\, g_k(X_k),$$

when these means exist.

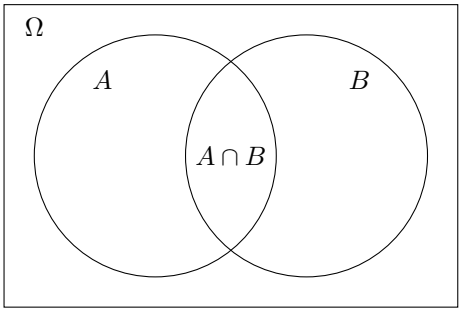
(k). Why product spaces? Because they are an efficient way of constructing probability spaces on which an arbitrary number of independent random variables with arbitrary distributions live. To see what is meant by this, consider the probability space from (e). Is it possible to construct two independent Bernoulli random variables with the same success probability on this space? Is it possible to construct more than two independent Bernoulli random variable on this space?

(l). Suppose that $A_1, A_2, \ldots, B_1, B_2, \ldots$ are independent events. Show that the σ-algebras $\sigma(A_1, A_2, \ldots)$ and $\sigma(B_1, B_2, \ldots)$ are independent. Let $A_1, A_2, \ldots$ be independent events, and consider the tail-σ-algebra $\mathcal{A} = \cap_{n=1}^{\infty} \sigma(A_n, A_{n+1}, \ldots)$. Prove Kolmogorov's zero-one law, namely that if $A \in \mathcal{A}$, then $\Pr(A) = 0$ or $\Pr(A) = 1$. [xx some hints? xx] Kolmogorov's zero-one law

Ex. A.25. *Basic conditional probability.* Let $(\Omega, \mathcal{A}, \Pr)$ be a probability space. The Venn-diagram below provides the intuition behind the definition of the conditional probability of an event A, given that the event B has occurred. If all we know is that B occurred, then the probability that A also occured is the the size of the area of B that intersects A, relative to the total size of B. Here, size is measured by, well, a probability measure. Thus, the conditional probability of A given B is defined by

$$\Pr(A \,|\, B) = \frac{\Pr(A \cap B)}{\Pr(B)}, \quad \text{provided } \Pr(B) > 0. \tag{A.5}$$

The notation $\Pr(A \,|\, B)$ might be unfortunate, because it may seem that the events A and B are on an equal footing. They are not. The event we condition on, namely B, is fixed, while the event we are computing the conditional probability of, that is A, can change. In fact, $A \mapsto P(A \,|\, B)$ is a probability measure, while $B \mapsto P(A \,|\, B)$ is not.



(a). Show that $A \mapsto \Pr(A \,|\, B)$ is a probability measure $(\Omega, \mathcal{A})$; and that $B \mapsto \Pr(A \,|\, B)$, i.e., the likelihood function, is *not* a probability measure.

(b). Suppose that B has positive probability. Show that A and B are independent if and only if $\Pr(A \,|\, B) = \Pr(A)$.

(c). Let $A_1, \ldots, A_n$ be disjoint events such that $\cup_{j=1}^n A_j = \Omega$. From the definition of conditional probability, deduce the law of total probability, that for any event B,

law of total probability

$$\Pr(B) = \Pr(B \,|\, A_1)\Pr(A_1) + \cdots + \Pr(B \,|\, A_n)\Pr(A_n).$$

Deduce also Bayes' theorem,

Bayes theorem

$$\Pr(A_j \mid B) = \frac{\Pr(B \,|\, A_j)\Pr(A_j)}{\Pr(B \,|\, A_1)\Pr(A_1) + \cdots + \Pr(B \,|\, A_n)\Pr(A_n)},$$

provided $\Pr(B) > 0$.

Ex. A.26. *Basic conditional expectation.* Let X be an integrable real random variable on a probability space $(\Omega, \mathcal{A}, \Pr)$. The probability of X landing in $B \in \mathcal{B}(\mathbb{R})$ given the event $A \in \mathcal{A}$, is, using the definition in (A.5), $\Pr(X \in B \,|\, A) = \Pr(\{X \in B\} \cap A)/\Pr(A)$, provided $\Pr(A) > 0$. We can then define the conditional expectation of X given A by

$$\mathrm{E}\,(X \,|\, A) = \frac{\mathrm{E}\, I_A X}{\Pr(A)} = \int_A X \frac{\mathrm{d}\Pr(\omega)}{\Pr(A)} = \int X \mathrm{d}\Pr(\omega \,|\, A). \tag{A.6}$$

The point being that conditioning on an event A with positive probability is just a matter of using the definition in (A.5) in the obvious manner.

(a). Suppose that we seek a function $Z\colon \Omega \to \mathbb{R}$ such that $Z(\omega) = \mathrm{E}\,(X \,|\, A)$ if $\omega \in A$, and $Z(\omega) = \mathrm{E}\,(X \,|\, A^c)$ if $\omega \in A^c$. That is, Z is the function

$$Z = \mathrm{E}\,(X \,|\, A) I_A + \mathrm{E}\,(X \,|\, A^c) I_A^c.$$

While $\mathrm{E}\,(X \,|\, A)$ and $\mathrm{E}\,(X \,|\, A^c)$ are constants, the indicator functions I_A and I_{A^c} are random variables, and so Z is also a random variable. We define this random variable Z to be the conditional expectation of X given I_A, denoted $\mathrm{E}\,(X \,|\, I_A)$. Thus, a conditional expectation given a random variable, is itself random variable. Two properties of $\mathrm{E}\,(X \,|\, I_A)$ are worth noticing because they are the defining properties of conditional expectation generally: First, show that $\mathrm{E}\,(X \,|\, I_A)$ is measurable with respect to the σ-algebra generated by I_A, namely $\sigma(I_A) = \{A, A^c, \emptyset, \Omega\}$. Second, show that $\mathrm{E}\, I_G \mathrm{E}\,(X \,|\, I_A) = \mathrm{E}\, I_G X$ for all $G \in \sigma(I_A)$.

(b). Let X be an integrable random variable, and Y a discrete random variable with range $\{y_1, y_2, \ldots\}$, and define $A_j = \{Y = y_j\} = \{\omega \in \Omega\colon Y(\omega) = y_j\}$ for $j = 1, 2, \ldots$. As a mild extension of (a), we can define the conditional expectation of X given Y as the random variable $\mathrm{E}\,(X \,|\, Y) = \sum_{j=1}^{\infty} \mathrm{E}\,(X \,|\, A_j) I_{A_j}$. Show that $\mathrm{E}\,(X \,|\, Y)$ is measurable with respect to $\sigma(Y)$, and that

$$\mathrm{E}\, I_G \mathrm{E}\,(X \,|\, Y) = \mathrm{E}\, I_G\, X, \quad \text{for all } G \in \sigma(Y). \tag{A.7}$$

Ex. A.27. *Conditional expectation.* If the random variable we want to condition on, say Y is such that $\Pr(Y = y) = 0$ for all y, then the definition of $\mathrm{E}\,(X \,|\, Y)$ given in Ex. A.26(b) does not make sense, as it would involve division by zero. The solution to this problem is to take (A.7) as the definition of conditional expectation. Here is how. On the probability space $(\Omega, \mathcal{A}, \Pr)$ let X be an integrable random variable, and let $\mathcal{G}$ be a sub-σ-algebra of $\mathcal{A}$. Then any $\mathcal{G}$-measurable random variable Z such that

conditional expectation

$$\mathrm{E}\, I_G Z = \int_G Z \,\mathrm{d}\Pr = \int_G X \,\mathrm{d}\Pr = \mathrm{E}\, I_G X, \quad \text{for all } G \in \mathcal{G}, \tag{A.8}$$

is called the conditional expectation of X given $\mathcal{G}$, denoted $\mathrm{E}(X \,|\, \mathcal{G})$. Per Ex. A.10(h), this definition entails that the conditional expectation is only defined up to null sets of $\Pr$, and we call any Z satisfying (A.8) a *version* of the conditional expectation.

(a). Prove that conditional expectation exists. More to the point, appeal to the Radon–Nikodym theorem, and show that if X is an integrable random variable on $(\Omega, \mathcal{A}, \Pr)$, and $\mathcal{G} \subset \mathcal{A}$, then there exists a $\mathcal{G}$-measurable random variable Z satisfying (A.8).

(b). Show that if X is $\mathcal{G}$ measurable, then $\mathrm{E}\,(X \,|\, \mathcal{G}) = X$, $\Pr$-almost surely.

(c). Show that $\mathrm{E}\, X$ equals the conditional expectation of X given the trivial σ-algebra, that is, $\mathrm{E}\, X = \mathrm{E}\,(X \,|\, \{\emptyset, \Omega\})$.

(d). Show that $\mathrm{E}\, X = \mathrm{E}\,\mathrm{E}\,(X \,|\, \mathcal{G})$. When $\mathcal{G} = \sigma(Y)$ for a random vector Y, we typically write $\mathrm{E}(X \,|\, Y)$ instead of the more cumbersome $\mathrm{E}\,(X \,|\, \mathcal{G})$ or $\mathrm{E}\,\{X \,|\, \sigma(Y)\}$, and we get the familiar formula

$$\mathrm{E}\, X = \mathrm{E}\,\mathrm{E}\,(X \,|\, Y). \tag{A.9}$$

(e). Prove the so-called tower property of conditional expectation: If $\mathcal{G} \subset \mathcal{B}$ are sub-σ-algebras of $\mathcal{A}$, then

tower property of conditional expectation

$$\mathrm{E}\,(X \,|\, \mathcal{B}) = \mathrm{E}(\,\mathrm{E}\,(X \,|\, \mathcal{B}) \,|\, \mathcal{G}), \quad \Pr\text{–a.s.}.$$

Deduce that (A.9) is the special case when $\mathcal{B}$ is the trivial σ-algebra.

Doob–Dynkin lemma

(f). Let X be a random variable and $Y = (Y_1, \dots, Y_k)$ a random vector. Combine Ex. A.3(f)&(g) to show that if X is measurable with respect to $\sigma(Y)$, then there exists a measurable function $g\colon \mathbb{R}^n \to \mathbb{R}$ such that $X = g(Y)$.

(g). The upshot of (f) is that since conditional expectation $\mathrm{E}(X \mid Y)$ is $\sigma(Y)$ measurable, there exists a measurable function g such that $g(Y) = \mathrm{E}(X \mid Y)$, Pr-almost surely. So when we write $\mathrm{E}(X \mid Y = y)$, we mean $g(y)$. With P_Y the distribution of Y, show that

$$\int_{Y^{-1}(B)} X \,\mathrm{dPr} = \int_B \mathrm{E}(X \mid Y = y) \,\mathrm{d}P_Y(y),$$

for every $B \in \mathcal{B}(\mathbb{R}^k)$. We can take this as the definition of conditional expectation instead of (A.8): Any function g defined on the range of Y is a version of the conditional expectation of X given Y if

$$\int_{Y^{-1}(B)} X \,\mathrm{dPr} = \int_B g(y) \,\mathrm{d}P_Y(y), \quad \text{for every } B \in \mathcal{B}(\mathbb{R}^k). \tag{A.10}$$

(h). Show that if (X, Y) has joint density $f_{X,Y}$ with respect to Lebesgue measure on $(\mathbb{R}^2, \mathcal{B}(\mathbb{R}^2))$, then the function

$$g(y) = \frac{\int x f_{X,Y}(x, y) \,\mathrm{d}x}{\int f_{X,Y}(x, y) \,\mathrm{d}x},$$

is a version of the conditional expectation of X given Y.

Ex. A.28. *Properties of conditional expectation.* Even though a conditional expectation is a random variable and an expectation is a constant, many of the same properties and theorems apply. Let $(\Omega, \mathcal{A}, \mathrm{Pr})$ be a probability space, $\mathcal{G}$ a sub-σ-algebra of $\mathcal{A}$, and let X and Y be integrable random variables.

take out what is known

(a). Use bootstrapping to show that if XY is integrable, and Y is $\mathcal{G}$-measurable, we can 'take out what is known' from the conditional expectation, that is

$$\mathrm{E}(XY \mid \mathcal{G}) = Y\mathrm{E}(X \mid \mathcal{G}), \quad \text{a.s.}.$$

(b). If $X = a$ a.s., show that $\mathrm{E}(X \mid \mathcal{G}) = a$. For constants a and b, show that

$$\mathrm{E}(aX + bY \mid \mathcal{G}) = a\,\mathrm{E}(X \mid \mathcal{G}) + b\,\mathrm{E}(Y \mid \mathcal{G}), \quad \text{a.s.}.$$

(c). Show that if $X \le Y$ almost surely, then $\mathrm{E}(X \mid \mathcal{G}) \le \mathrm{E}(Y \mid \mathcal{G})$ almost surely.

(d). Show that Fatou's lemma holds for conditional expectation, i.e., if $X_1, X_2, \dots$ are nonnegative random variables, then

$$\mathrm{E}(\liminf_{n\to\infty} X_n \mid \mathcal{G}) \le \liminf_{n\to\infty} \mathrm{E}(X_n \mid \mathcal{G}), \quad \text{a.s.}.$$

monotone convergence

(e). Show that if X_n is a monotone increasing sequence of integrable random variables such that $X_n \to X$ a.s., then $\mathrm{E}(X_n \mid \mathcal{G}) \to \mathrm{E}(X \mid \mathcal{G})$ a.s..

dominated convergence

(f). Show that if X_n is a sequence of integrable random variables such that $X_n \to X$ a.s., and $|X_n| \le Y$ a.s. for an integrable random variable Y, then $\mathrm{E}(X_n \mid \mathcal{G}) \to \mathrm{E}(X \mid \mathcal{G})$ a.s..

(g). Show that in (e) and (f), if the sequence X_n only converges in probability to X (instead of almost surely), then we have $\mathrm{E}(X_n \,|\, \mathcal{G}) \to_p \mathrm{E}(X \,|\, \mathcal{G})$.

(h). Let g be a convex function such that $g(X)$ is integrable.. Show that

$$\mathrm{E}\{g(X) \,|\, \mathcal{G}\} \geq g(\mathrm{E}\{X \,|\, \mathcal{G}\}), \quad \text{a.s..}$$

Applying this to the convex function $g(x) = |x|^p$ for some $1 \leq p < \infty$, we have that $|\mathrm{E}(X \mid \mathcal{G})| \leq \mathrm{E}(|X| \,|\, \mathcal{G})$, a.s.. Show also that $(\mathrm{E}\,|\mathrm{E}(X \mid \mathcal{G})|^p)^{1/p} \leq (\mathrm{E}(|X|^p \,|\, \mathcal{G}))^{1/p}$.

(i). Suppose that X is square integrable. The conditional variance of X given $\mathcal{G}$, denoted $\mathrm{Var}(X \,|\, \mathcal{G})$, or $\mathrm{Var}(X \,|\, Y)$ when $\mathcal{G} = \sigma(Y)$, is any version of the random variable $\mathrm{E}\{(X - \mathrm{E}(X \,|\, \mathcal{G}))^2 \,|\, \mathcal{G}\}$. Show that if Y and Z are $\mathcal{G}$-measurable random variables, such that XY is square integrable, then $\mathrm{Var}(YX + Z \,|\, \mathcal{G}) = Y^2 \mathrm{Var}(X \,|\, \mathcal{G})$ a.s..

variance decomposition

(j). Suppose that X is square integrable. Show the following very useful variance decomposition formula $\mathrm{Var}(X) = \mathrm{E}\,\mathrm{Var}(X \,|\, \mathcal{G}) + \mathrm{Var}(\mathrm{E}(X \,|\, \mathcal{G}))$ a.s..

Ex. A.29. *Conditional probability and conditional distributions.* Let $(\Omega, \mathcal{F}, \Pr)$ be a probability space. The conditional probability of the event $A \in \mathcal{F}$ given a sub-σ-algebra $\mathcal{G} \subset \mathcal{F}$ is defined as $\Pr(A \,|\, \mathcal{G}) := \mathrm{E}(I_A \,|\, \mathcal{G})$ for some version of $\mathrm{E}(I_A \,|\, \mathcal{G})$. When $\mathcal{G}$ is generated by a random vector $Y = (Y_1, \ldots, Y_k)$, a perhaps more natural definition of the conditional probability of the event A given $\mathcal{G} := \sigma(Y)$ is to call any function $y \mapsto \Pr(A \,|\, Y = y)$ satisfying

conditional probability

$$\int_B \Pr(A \,|\, Y = y)\, \mathrm{d}\Pr Y^{-1}(y) = \Pr(A \cap Y^{-1}(B)), \quad \text{for all } B \in \mathcal{B}(\mathbb{R}^k),$$

a version of the conditional probability of A given Y. Let X be a random element (variable, vector, etc.) taking values in $(\mathcal{X}, \mathcal{B})$, and let $\mathcal{G} \subset \mathcal{F}$ be a sub-σ-algebra. A conditional distribution of X given $\mathcal{G}$ is a function $P_X(B \,|\, \mathcal{G})(\omega)$ of B and ω such that

conditional distribution

(i) for each $B \in \mathcal{B}$, $P_X(B \,|\, \mathcal{G}) = \Pr(X \in B \,|\, \mathcal{G})$ Pr-a.s.;
(ii) for each $\omega \in \Omega$, $B \mapsto P_X(B \,|\, \mathcal{G})(\omega)$ is a probability measure on $(\mathcal{X}, \mathcal{B})$.

If $(\mathcal{X}, \mathcal{B}) = (\Omega, \mathcal{F})$ and $X(\omega) = \omega$ is the identity map, a conditional distribution gives a regular conditional probability. The existence of conditional distributions is not obvious. Fortunately, regular conditional distributions do exist in all the cases we care about in this book, e.g., $\mathcal{X} = \mathbb{R}^k$, $\mathcal{X} = C[0,1]$ with the supnorm, etc. [xx perhaps make an exercise on this xx].

(a). Show that (i) $\Pr(A \,|\, \mathcal{G}) \geq 0$ a.s., for each $A \in \mathcal{F}$; (ii) $\Pr(\Omega \,|\, \mathcal{G}) = 1$ a.s.; (iii) if $A_1, A_2, \ldots$ is a disjoint sequence, then $\Pr(\cup_{j=1}^{\infty} A_j \,|\, \mathcal{G}) = \sum_{j=1}^{\infty} \Pr(A_j \,|\, \mathcal{G})$ a.s..

(b). If $\mathcal{G} = \sigma(Y)$, show that (i) and (ii) in the definition of conditonal distribution may be restated as follows: A conditional distribution of X given Y is a function $P_X(B \,|\, Y = y)$ of B and y such that

(i) for each $B \in \mathcal{B}$, $P_X(B \,|\, Y = y) = \Pr(X \in B \,|\, Y = y)$, P_Y-a.s.;
(ii) for each $y \in \mathbb{R}^k$, $B \mapsto P_X(B \,|\, Y = y)$ is a probability measure on $(\mathcal{X}, \mathcal{B})$.

Show that if the famility $\{P_X(\cdot \,|\, Y = y) \colon y \in \mathbb{R}^k\}$ is dominated by a σ-finite measure μ (i.e., $P_X(\cdot \,|\, Y = y) \ll \mu$ for all y), then there is a nonnegative function $(x, y) \mapsto f(x \,|\, y)$, a conditonal density, such that for each $y \in \mathbb{R}^k$,

$$P_X(A \,|\, Y = y) = \int_A f(x \,|\, y)\, \mathrm{d}\mu(x), \quad \text{for all } A \in \mathcal{F}.$$

(c). Let $P_X(B \mid \mathcal{G})$ be a conditional distribution of $X\colon \Omega \to \mathcal{X}$ given $\mathcal{G}$, and suppose that $g\colon \mathcal{X} \to \mathbb{R}$ is a measurable function such that $g(X)$ is integrable. Show that

$$\mathrm{E}\,(g(X) \mid \mathcal{G})(\omega) = \int g(x)\, P_X(\mathrm{d}x \mid \mathcal{G})(\omega),$$

for Pr-almost all $\omega \in \Omega$. This is the conditional expectation analogue of the second expression for $\mathrm{E}\, g(X)$ in Ex. A.16(c). Show also that if $\mathcal{G} = \sigma(Y)$ for some random vector $Y = (Y_1, \ldots, Y_k)$, then

$$\mathrm{E}\,(g(X) \mid Y = y) = \int g(x)\, P_X(\mathrm{d}x \mid Y = y),$$

for P_Y almost all $y \in \mathbb{R}^k$.

(d). Let X and Y be independent random variables on the probabilty space $(\Omega, \mathcal{A}, \Pr)$. Let g be a real-valued function such that $g(X, Y)$ is integrable with respect to Pr. Show that

$$\mathrm{E}\,(g(X,Y) \mid \sigma(Y))(\omega) = \int g(x, Y(\omega))\, \mathrm{d}P_X(x),$$

for Pr-almost all $\omega \in \Omega$, where $P_X = \Pr X^{-1}$ is the distribution of X.

(e). Suppose that $(\mathcal{X}, \mathcal{A}, \mu)$ and $(\mathcal{Y}, \mathcal{B}, \nu)$ are σ-finite measure spaces, and let $X\colon \Omega \to \mathcal{X}$ and $Y\colon \Omega \to \mathcal{Y}$ be random variables. Let P_Y be the distribution of Y on $(\mathcal{Y}, \mathcal{B})$, and let $P_{X,Y}$ be their joint distribution on $(\mathcal{X} \times \mathcal{Y}, \mathcal{A} \otimes \mathcal{B})$. Suppose that $P_{X,Y} \ll \mu \times \nu$ with density $f_{X,Y}$. Show that $P_Y \ll \nu$ with density

$$f_Y(y) = \int_{\mathcal{X}} f_{X,Y}(x, y)\, \mathrm{d}\mu(x),$$

and that the conditional distribution of X given $Y = y$ has densities

$$f_{X \mid Y}(x \mid y) = \frac{f_{X,Y}(x, y)}{f_Y(y)},$$

with respect to μ. Convince yourself that the two expressions above also holds with X and Y switching roles. In particular, show that the densities above can be expressed as

$$f_{X \mid Y}(x \mid y) = \frac{f_{Y \mid X}(y \mid x) f_X(x)}{f_Y(y)}.$$

This is Bayes' theorem for densities. In the case that X is a (random) parameter, and Y are the data, $f_{X \mid Y}(x \mid y)$ is called the posterior density of the parameter given the data.

(f). Show that if $g\colon \mathbb{R} \to \mathbb{R}$ is a measurable function such that $g(X)$ is integrable, then

$$\mathrm{E}\,(g(X) \mid \sigma(Y))(\omega) = \int g(x) f_{X \mid Y}(x \mid Y(\omega))\, \mathrm{d}\mu(x),$$

for Pr-almost all $\omega \in \Omega$. Note that this is the conditional expectation analogue of the expression for $\mathrm{E}\, g(X)$ in Ex. A.18(d).

Ex. A.30. *Conditional independence.* [xx some intro text. kanskje ta dette inn i oppgavene over xx]

(a).

(b). Let $\mathcal{F}$, $\mathcal{G}$, and $\mathcal{H}$ be σ-algebras. Show that $\mathcal{F} \perp\!\!\!\perp \mathcal{H} \,|\, \mathcal{G}$ if and only if

$$\Pr(H \,|\, \mathcal{F} \vee \mathcal{G}) = \Pr(H \,|\, \mathcal{G}), \text{ a.s., for all } H \in \mathcal{H},$$

where $\mathcal{F} \vee \mathcal{G}$ is the smallest σ-algebra that contains $\mathcal{F}$ and $\mathcal{G}$.

Ex. A.31. *[xx rydd opp here xx].*

(a). The results in (e) are in the background of many of the conditional probability calculations to be carried out in this book, but they cannot always be used. In Ex. **????**, we had $X \sim \mathrm{N}(0,1)$ and the event $A = \{X \geq c\}$. Consider the random variable $Y = I(X \geq c)$. The distributions of X and Y have, of course, densities with respect to Lebesgue and counting measure, say λ and μ, respectively. Show, however, that the distribution of (X,Y) is *not* dominated by the product measure $\lambda \times \nu$ on $(\mathbb{R} \times \{0,1\}, \mathcal{B}(\mathbb{R}) \otimes 2^{\{0,1\}})$. Define the set function $\rho(C) = \lambda(\{x \in \mathbb{R} \colon (x, I\{x \geq c\}) \in C\})$ on $\mathcal{B}(\mathbb{R}) \otimes 2^{\{0,1\}}$. Show that ρ is a measure, and that $P_{X,Y} \ll \rho$. Find an expression for the density. Introduce the set function $\lambda_{X\,|\,Y}(B \,|\, y) = \lambda(B_y)$, where $B_y = \{x \in \mathbb{R} \colon I(x \geq c) = y\}$ for $B \in \mathcal{B}(\mathbb{R})$. Show that $\lambda_{X\,|\,Y}(B \,|\, y)$ is a measure on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ for $y = 0, 1$; that the conditional distribution of X given $Y = y$ is dominated $\lambda_{X\,|\,Y}(B \,|\, y)$; and find an expression for the density.

(b).

(c).

(d). Suppose that $(\mathcal{X}, \mathcal{A}, \mu)$ be a σ-finite measure space, and $(\Theta, \mathcal{B})$ a measurable space. Let $X \colon \Omega \to \mathcal{X}$ and $\theta \colon \Omega \to \Theta$ be random variables on the same underlying probability space $(\Omega, \mathcal{F}, \Pr)$. Suppose that the conditional distribution of X given θ is P_θ, and that $P_\theta \ll \mu$ for all $\theta \in \Theta$, with densities $f_\theta(x)$. Let Π be the distribution of θ. We think of θ as a random parameter, and, as is the convention, we do not distinguish between the random variable θ and the values it attains. Show that (X, θ) has joint distribution

$$\Pr((X,\theta) \in C) = \int I_C(x,\theta) f_\theta(x) \,\mathrm{d}(\mu \times \Pi)(x,\theta),$$

for $C \in \mathcal{A} \otimes \mathcal{B}$, where $\mu \times \Pi$ is the product measure on $(\mathcal{X} \times \Theta, \mathcal{A} \otimes \mathcal{B})$. Show that the posterior distribution of θ given X is

$$\Pr(\theta \in B \,|\, \sigma(X)) = \int_B \frac{f_\theta(X) \,\mathrm{d}\Pi(\theta)}{\int_{\mathcal{X}} f_\theta(X) \,\mathrm{d}\Pi(\theta)},$$

almost surely. Show also that if ν is a σ-finite measure on $(\Theta, \mathcal{B})$ such that $\Pi \ll \nu$ with density π, then the posterior density of θ given X is

$$\pi(\theta \,|\, x) = \frac{f_\theta(x)\pi(\theta)}{\int f_\theta(x)\pi(\theta) \,\mathrm{d}\nu(\theta)},$$

with respect to ν. This version of Bayes' theorem, with this notation, is the workhorse formula of Chapter 6.

(e). (xx the more general cases, and examples; Bayes' theorem for densities; include also classical transformation formula, with Jacobian etc.)

(f).

(g).

Ex. A.32. *Extension of probability spaces.* An extension of a probability space $(\Omega, \mathcal{F}, \Pr)$ is a product space $(\Omega \times \mathcal{X}, \mathcal{F} \otimes \mathcal{B})$ equipped with a probability measure $\Pr'$ such that $\Pr'(A \times \mathcal{X}) = \Pr(A)$ for all $A \in \mathcal{F}$.

(a). Let Y be a random variable (or vector, or process) on $(\Omega, \mathcal{F}, \Pr)$. We can extend Y to be defined on the extension $(\Omega \times \mathcal{X}, \mathcal{F} \otimes \mathcal{B}, \Pr')$ by setting $Y'(\omega, x) = Y(\omega)$. Show that Y' has the same distribution as Y, and, in particular, that $\mathrm{E}' g(Y') = \mathrm{E}\, g(Y)$ for all measurable $g\colon \mathcal{X} \to \mathbb{R}$.

(b). Let Q be a probability kernel (or Markov kernel) between the probability space $(\Omega, \mathcal{G}, \Pr)$ and the measurable space $(\mathcal{X}, \mathcal{B})$. That is $Q\colon \Omega \times \mathcal{X} \to [0, \infty)$ is so that

$\omega \mapsto Q(\omega, B)$ is $\mathcal{G}$-measurable for each fixed $B \in \mathcal{B}$; and
$B \mapsto Q(\omega, B)$ is a probability measure on $(\mathcal{X}, \mathcal{B})$ for each $\omega \in \Omega$.

The regular conditional distributions introduced in Ex. A.29 are probability kernels. Other examples of probability kernels are the mixed normal densities we meet in Ex. 2.62, where $\sigma\colon \Omega \to (0, \infty)$ a random variable, and

$$Q(\omega, B) = \int_B \frac{1}{\sqrt{2\pi}\sigma(\omega)} \exp(-\frac{x^2}{2\sigma(\omega)^2}\, \mathrm{d}x).$$

If we were to simulate a random variable X from $Q(\omega, \cdot)$ we would perhaps run something ressembling the following little script

```
sigma <- rgamma(1,2,2) # for example
X <- rnorm(1,0,sigma)
```

and almost without thinking about it, X would have been taken as conditionally independent of all other random variables, given σ. Since our probability space may not be rich enough to support such as conditionally independent random variable (see example below), we might need enlarge the probability space. Here is how: Suppose we have a probability space $(\Omega, \mathcal{F}, \Pr)$ and $\mathcal{G} \subset \mathcal{F}$. Let Q be a probability kernel between $(\Omega, \mathcal{G}, \Pr)$ and $(\mathcal{X}, \mathcal{B})$. Consider $(\Omega \times \mathcal{X}, \mathcal{G} \otimes \mathcal{B}, \Pr')$ with

$$\Pr'(A) = \int_{\Omega \times \mathbb{R}} I_A(\omega, x) Q(\omega, \mathrm{d}x)\, \mathrm{d}\Pr(\omega), \text{ for all } A \in \mathcal{G} \otimes \mathcal{B}.$$

Show that (i) $\Pr'(G \times \mathcal{X}) = \Pr(G)$ for all $G \in \mathcal{G}$, meaning that the product space just defined is indeed an extension; (ii) that the random variable $X(\omega, x) = x$ is so that $\Pr'(X \in B \,|\, \sigma(V)) = Q(\cdot, B)$, $\Pr'$-almost surely; and (iii) that X is conditionally independent of all $\mathcal{F}$-measurable random variables, given $\mathcal{G}$.

(c). Consider the probability space $(\{H, T\}, 2^{\{H,T\}}, \Pr)$ with $\Pr$ a probability measure such that $\Pr(H) = p$. Let $\theta\colon \{H, T\} \to \{0, 1\}$ be such that $\Pr(\theta = 1) = p$, and consider the probability kernel $Q(\omega, B) = \theta(\omega) \int_B \phi(x)\, \mathrm{d}x + (1 - \theta(\omega)) \int_B \phi(x/\sqrt{2})/\sqrt{2}\, \mathrm{d}x$, where ϕ is the standard normal density. Construct an extension on which the random variable X has conditional distribution Q, given θ.

Ex. A.33. *Regularity and approximations.* Let $(\mathcal{X}, d)$ be a metric space and consider the measurable space $(\mathcal{X}, \mathcal{B}, \mu)$, where $\mathcal{B}$ is the Borel σ-algebra, and μ is a finite measure. The distance from a point x to a set A is $d(x, A) = \inf\{d(x, y)\colon y \in A\}$. Using this

distance, we can express any closed set F as a countable intersection of open sets, as follows $F = \cap_{n=1}^{\infty}\{x \in \mathcal{X}\colon d(x,F) < 1/n\}$. So by De Morgan's laws, any open set G can be expressed as a countable union of closed sets $G = \cup_{n=1}^{\infty}\{x \in \mathcal{X}\colon d(x,G^c) \geq 1/n\}$. These facts are useful in what follows.

(a). Show that $\mu(B) = \sup_{F\subset B}\mu(F)$ for any open set B, where the supremum is taken over closed sets F. Similarly, show that $\mu(C) = \inf_{G\supset C}\mu(G)$ for any closed set C, with the infinum taken over open sets G. Recall that $\{B \subset \mathcal{X}\colon B \text{ is open}\}$ and $\{C \subset \mathcal{X}\colon C \text{ is closed}\}$ are π-systems, and that they both generate the Borel σ-algebra, facts we use below.

(b). Show that $\mathcal{D}_c = \{B \in \mathcal{B}\colon \mu(B) = \sup_{F\subset B}\mu(F), F \text{ closed}\}$ and $\mathcal{D}_o = \{B \in \mathcal{B}\colon \mu(B) = \inf_{B\subset G}\mu(G), G \text{ open}\}$ are d-systems, and conclude from Dynkin's lemma (see Ex. A.5) that for any $B \in \mathcal{B}$,

$$\mu(B) = \sup_{F\subset B}\mu(F) = \inf_{G\supset B}\mu(G),$$

where the supremum and the infimum are taken over closed sets F and open sets G, respectively.

(c). Show that for any measurable function $f\colon \mathcal{X} \to \mathbb{R}$ such that $\int f^p\,\mathrm{d}\mu < \infty$, i.e., $f \in L^p(\mathcal{X},\mathcal{B},\mu)$, there is a sequence $f_1, f_2, \ldots\colon \mathcal{X} \to \mathbb{R}$ of bounded and continuous functions such that $\int |f_n - f|^p\,\mathrm{d}\mu \to 0$.

Ex. A.34. *The mean via the cumulative distribution function.* Consider a random variable X on $[0,\infty)$, with cumulative function F.

(a). Show that the mean $\mathrm{E}\,X = \int_0^\infty x\,\mathrm{d}F(x)$ also can be expressed as $\int_0^\infty (1-F)\,\mathrm{d}x$. You may use $x = \int_0^\infty I(x > y)\,\mathrm{d}y$ and then Fubini. Show furthermore that $\mathrm{E}\,X^p = \int_0^\infty \{1 - F(x^{1/p})\}\,\mathrm{d}x$.

(b). As a simple illustration, consider X with density function $f(x) = \theta\exp(-\theta x)$, where θ is a positive parameter. Find the cumulative F, and compute $\mathrm{E}\,X$ in two ways.

(c). (xx something with a discrete distribution too. find the AmStat paper we talked briefly about, for a bit more. xx)

Ex. A.35. *Moment-generating functions: examples.* [xx we have not introduces these distributions yet. xx] (xx nils lifts these to Ch1; then need post-polish here in App. xx) For a random variable Y, with distribution P, its moment generating function is

$$M(t) = \mathrm{E}\,\exp(tY) = \int \exp(ty)\,\mathrm{d}P(y),$$

defined for each t at which the expectation exists. The moment generating function is useful for finding and characterising distributions, for finding their moments, for handling the distributions of sums of variables, and in connection with distributional limits. When Y has a density $f(y)$ (with respect to Lebesgue measure), we have $M(t) = \int \exp(ty)f(y)\,\mathrm{d}y$, and if it is discrete with pointmasses $f(y)$ for sample space S, say, then $M(t) = \sum_{y\in S}\exp(ty)f(y)$. The expectation operator is more general, however, and $M(t)$ is perfectly defined also for intermediate cases where Y can have both discrete and continuous parts; see Ex. A.16.

(a). For a standard normal $Y \sim \mathrm{N}(0,1)$, show that $M(t) = \exp(\frac{1}{2}t^2)$. When $Y \sim \mathrm{N}(\mu, \sigma^2)$, derive $M(t) = \exp(\mu t + \frac{1}{2}\sigma^2 t^2)$.

(b). For $Y \sim \mathrm{Expo}(\theta)$, show that $M(t) = 1/(1 - t/\theta)$, for $t < \theta$.

(c). For $Y \sim \mathrm{Gam}(a, b)$, with density $\{b^a/\Gamma(a)\}y^{a-1}\exp(-by)$, show that $M(t) = \{b/(b-t)\}^a$, for $t < b$. In particular, $M(t) = 1/(1-t)^a$ for $\mathrm{Gam}(a, 1)$.

(d). Suppose Y is equal to zero with probability 0.90, but a standard normal with probability 0.10. Find the $M(t)$, and generalise.

(e). For the binomial (n, p), show that $M(t) = \{1 - p + p\exp(t)\}^n$.

(f). For $Y \sim \mathrm{Pois}(\theta)$, find $M(t) = \exp\{\theta(e^t - 1)\}$. Use this, with Ex. 1.27, to find $M(t)$ also for the negative binomial (a, p). (xx hm, should give the formula here. xx)

(g). Let $Y = \pm 1$ with probabilities $\frac{1}{2}, \frac{1}{2}$. Show that

$$M(t) = \cosh(t) = \tfrac{1}{2}(e^t + e^{-t}) = 1 + (1/2)t^2 + (1/4!)t^4 + (1/6!)t^6 + \cdots .$$

(h). For the uniform distribution on the unit interval, show that $M(t) = \{\exp(t) - 1\}/t$, for $t \neq 0$, and with $M(0) = 1$.

(i). Let Y have the uniform distribution on the $[-1, 1]$ interval. Show that

$$M(t) = \frac{\exp(t) - \exp(-t)}{2t} = \frac{\sinh t}{t},$$

and that this function may be written as the infinite sum $1 + (1/3!)t^2 + (1/5!)t^4 + \cdots$.

Ex. A.36. *Moment-generating functions: properties.* Among the basic properties of moment-generating functions is that it generates moments. As we will see in this exercise, the rth derivative of $M(t) = \mathrm{E}\exp(tY)$ at the point zero equals $\mathrm{E}\,Y^r$.

(a). Suppose that the moment generating function $M(t)$ of a random variable Y is finite for all $t \in (-t_0, t_0)$, for some $t_0 > 0$. For some t in this interval, the sum $M(-t) + M(t)$ is then also clearly finite. Appeal to Ex. A.12(b) to show that the finiteness of $M(-t)+M(t)$ implies that $\mathrm{E}\,|Y|^{2k}$ is finite for all $k \geq 1$. Now, use that $|x|^{2k-1} \leq 1 + |x|^{2k}$ to fill in the odd gaps. This highlights the restrictiveness of the moment-generating function: by assuming its existence in a neighbourhood of zero, we are effectively assuming that all moments exist.

(b). To see that the converse of (a) does not hold, that is, a distribution with finite moments of all orders may not have a moment-generating function that is finite in some interval around zero, consider the log-normal distribution; see Ex. 1.58.

(c). Provided $M(t)$ is finite in an interval around zero, say $(-t_0, t_0)$, $t_0 > 0$, you may again use Ex. A.12(b) to show that

$$M(t) = \sum_{k=0}^{\infty} \frac{t^k}{k!}\,\mathrm{E}\,Y^k, \quad \text{for } |t| < t_0.$$

We now see that if we can slip the derivative inside this sum, then we have the property mentioned in the introduction. To see that we can, choose points a and b so that $|t| < a < b < t_0$, and show that

$$\Big|\frac{(t+h)^k - t^k}{h}\Big| \leq 3ka^{k-1},$$

provided $|h| < a - |t|$. Next, show that there is an $M > 0$ so that $3ka^{k-1} \leq 3Mb^k$, and use this to conclude that $\sum_{k=0}^{\infty} 3ka^{k-1}\mathrm{E}\,Y^k < \infty$. Using these results, explain why $M'(t) = \sum_{k=1}^{\infty}(t^{k-1}/(k-1)!)\mathrm{E}\,Y^k$. Finally, by induction, show that

$$M^{(r)}(t) = \sum_{k=r}^{\infty} \frac{t^{k-r}}{(k-r)!}\,\mathrm{E}\,Y^k.$$

This expression shows that moment-generating functions generate moments in the sense that $M^{(r)}(0) = \mathrm{E}\,(Y^r)$.

(d). For $X \sim \mathrm{N}(0,1)$, show that $M(t)$ for $|X|$ becomes $2\exp(\frac{1}{2}t^2)\Phi(t)$, and use this to find its mean and variance.

(e). It Y has mean ξ and standard deviation σ, and moment-generating function $M(t)$, give a formula for that of $Y' = (Y-\xi)/\sigma$. Illustrate this in the case of $Y \sim \mathrm{Pois}(\theta)$, computing and drawing the moment-generating function of $(Y-\theta)/\theta^{1/2}$, alongside $\exp(\frac{1}{2}t^2)$. Comment on what you find.

(f). If Y has a distribution symmetric around zero, such that Y and $-Y$ have the same distribution, then $M(t) = M(-t)$, so it depends on t only via $|t|$.

Ex. A.37. *Uniqueness of the Laplace transform.* For random variables taking values on the positive halfline $[0,\infty)$, it is sometimes more convenient to work with the Laplace transform instead of the moment-generating function. Let X be a random variable with support $[0,\infty)$, and denote its distribution and c.d.f. by P and F, respectively. Define the Laplace transform $L(t) = \mathrm{E}\,\exp(-tX)$, for $t \geq 0$. We here follow Billingsley (1995, pp. 284–286) in proving that $L(t)$ determines the distribution of X.

(a). Show that $L(t)$ is finite for all $t \geq 0$, and deduce from Ex. A.36(c) that the rth derivative of $L(t)$ is $L^{(r)}(t) = (-1)^r\mathrm{E}\,X^r\exp(-tX)$.

(b). To prove that the Laplace transform uniquely determines the distribution of X, we make a detour via the Poisson distribution. Let Y_θ be a Poisson random variable with mean θ, i.e., Y_θ has density $f(r;\theta) = (1/r!)\theta^r\exp(-\theta)$ with respect to counting measure on $\{0,1,2,\ldots\}$. Use the Chebyshov inequality (see Ex. 2.10) to show that $Y_\theta/\theta \to_p 1$ as $\theta \to \infty$. Let $G(t;\theta) = \Pr(Y_\theta/\theta \leq t) = \sum_{r=0}^{\lfloor\theta t\rfloor} f(r;\theta)$ be the c.d.f. of Y_θ/θ, where $\lfloor u\rfloor = \max\{m \in \mathbb{Z}\colon m \leq u\}$. Deduce from the convergence in probability result that as θ tends to infinity, $G(t;\theta) \to 1$ if $t > 1$ and $G(t;\theta) \to 0$ if $t < 1$.

(c). Show that we can write,

$$\sum_{r=0}^{\lfloor ty\rfloor} \frac{(-1)^r}{r!}t^r L^{(r)}(t) = \mathrm{E}\,G(y/X;tX).$$

and show that $G(y/x;tx) \to I\{0 \leq x \leq y\}$ as $t \to \infty$, almost everywhere. Appeal to the right convergence theorem to conclude that $\mathrm{E}\,G(y/X;tX) \to F(y)$, for all continuity points of $F(y) = P(-\infty,y]$.

(d). [xx include continuity theorem for Laplace transform, i.e., that for nonnegative random variables X_n the Laplace transforms converge to a the Laplace transform of a rv X, then $X_n \to_d X$ xx]

Ex. A.38. *Moment-generating functions for sums.* (xx point here to Ex. **??**, and more. xx) If Y_1 and Y_2 are independent, with given distributions, say with densities f_1 and f_2, then their sum $Z = Y_1 + Y_2$ have of course a well-defined distribution, and its density can be expressed as

$$g_2(z) = \int f_1(z - y_2) f_2(y_2)\,\mathrm{d}y_2 = \int f_1(y_1) f_2(z - y_1)\,\mathrm{d}y_1.$$

With algebraic patience this may e.g. be used to show that if $Y_1 \sim \mathrm{N}(\mu_1, \sigma_1^2)$ and $Y_2 \sim \mathrm{N}(\mu_2, \sigma_2^2)$, then indeed $Y_1 + Y_2$ is normal too, with parameters $\mu_1 + \mu_2$ and $\sigma_1^2 + \sigma_2^2$; see Ex. 1.2. Such convolutions quickly become convoluted in more general setups, however, and finding the density of say $Y_1 + Y_2 + Y_3 + Y_4$ from given densities f_1, f_2, f_3, f_4 may become too complicated. Pushing the matter to the domain of moment-generating functions instead makes matters simpler.

(a). When X and Y are independent, then $M_{X+Y}(t) = M_X(t) M_Y(t)$, in the obvious notation. This generalises of course to the case of more than two independent variables.

(b). Let $Y_i \sim \mathrm{N}(\mu_i, \sigma_i^2)$, for $i = 1, \ldots, n$, with these variables being independent. Find the moment-generating function for the sum $Z = Y_1 + \cdots + Y_n$, and use the characterisation property to establish that indeed $Z \sim \mathrm{N}(\sum_{i=1}^n \mu_i, \sum_{i=1}^n \sigma_i^2)$.

(c). Let $Y_1, \ldots, Y_k$ be independent Gamma distributed variables, with parameters (a_1, b), $\ldots, (a_k, b)$; see Ex. 1.9. Show that their sum is a Gamma with parameters $(\sum_{i=1}^k a_i, b)$.

(d). Suppose $Z = Y_1 + Y_2$, with these two being independent, and suppose you know that $Y_1 \sim \mathrm{N}(0, 2)$ and $Z \sim \mathrm{N}(0, 7)$. Prove that Y_2 must be a $\mathrm{N}(0, 5)$.

(e). Similarly, suppose $Z = Y_1 + Y_2$, with these two being independent, and assume it is known that $Y_1 \sim \chi^2_{10}$ and that $Z \sim \chi^2_{24}$. Prove that $Y_2 \sim \chi^2_{14}$.

Ex. A.39. *Characteristic functions.* The characteristic function of a random variable X is defined as

$$\varphi(t) = \mathrm{E}\,\exp(itX) = \mathrm{E}\,\cos(tX) + i\,\mathrm{E}\,\sin(tX),$$

with $i = \sqrt{-1}$ the complex unit, and $t \in \mathbb{R}$. As the name suggests, the characteristic function is useful for finding and characterising distributions, for finding their moments, for handling the distributions of sums of variables, and for finding with distributional limits. What distingusihes it from the moment-generating function is that it always exists, i.e., we do not need to make the assumption of all moments being finite (see Ex. A.36(a)).

(a). Show that the characteristic function always exists, in fact $|\varphi(t)| \le 1$. Establish that $|\varphi(t + h) - \varphi(t)| \le \mathrm{E}\,|\exp(ihX) - 1|$, and use this inequlity to show that $\varphi(t)$ is uniformly continuous.

(b). Show that if $Z \sim \mathrm{N}(0, 1)$, then its characteristic function is $\varphi_Z(t) = \exp(-\frac{1}{2}t^2)$, and that if $X \sim \mathrm{N}(\mu, \sigma^2)$ its characteristic function is $\varphi_X(t) = \exp(it\mu - \frac{1}{2}t^2\sigma^2)$.

(c). Show that the Cauchy distribution with density $f(x) = (1/\pi)(1 + x^2)^{-1}$ has characteristic function $\exp(-|t|)$. Note that this function does not have a derivative at zero, corresponding to the fact that the Cauchy distribution does not have a finite mean.

(d). Suppose that X has c.d.f. F. Show that the characteristic function is real-valued if and only if the distribution is symmetric, that is. $F(x) = 1 - F(-x)$ for all x.

(e). For $m = 0, 1, 2, \ldots$ define $r_m(x) = \exp(ix) - \sum_{k=0}^{m} (ix)^k/k!$, and let $r_{-1}(x) = \exp(ix)$. Convince yourself that $r_m(x) = r_{m-1}(x) - (ix)^m/m!$, and that $r_m(x) = i\int_0^x r_{m-1}(y)\,\mathrm{d}y$ for $x > 0$ and $r_m(x) = -i\int_x^0 r_{m-1}(y)\,\mathrm{d}y$ for $x < 0$. Show that $|r_0(x)| \leq \min(2, |x|)$, and proceed by induction to show that

$$|\exp(ix) - \sum_{k=0}^{m} \frac{(ix)^k}{k!}| \leq \min\big(\frac{2|x|^m}{m!}, \frac{|x|^{m+1}}{(m+1)!}\big), \quad \text{for } m = 0, 1, 2, \ldots.$$

In particular, for $m = 0, 1, 2$,

$$\begin{aligned} |\exp(ix) - 1| &\leq \min(|x|, 2), \\ |\exp(ix) - (1 + ix)| &\leq \min(\tfrac{1}{2}|x|^2, 2|x|), \\ |\exp(ix) - (1 + ix - \tfrac{1}{2}x^2)| &\leq \min(\tfrac{1}{6}|x|^3, x^2). \end{aligned}$$

(f). Use the inequality in (e) and the dominated convergence theorem to show that if $\mathrm{E}\,X^2$ is finite, then $\varphi(t) = 1 + it\mathrm{E}\,X - \frac{1}{2}t^2\mathrm{E}\,X^2 + o(t^2)$ as $t \to 0$. In particular, if $\mathrm{E}\,X = 0$ and $\mathrm{E}\,X^2 = \sigma^2$, we have

$$\varphi(t) = 1 - \tfrac{1}{2}t^2\sigma^2 + o(t^2), \quad \text{as } t \to 0.$$

Generalise to show that if $\mathrm{E}\,|X|^m < \infty$ for some $m \geq 1$, then

$$\varphi(t) = \sum_{k=0}^{m} ((it)^k/k!)\mathrm{E}\,X^k + o(t^m), \quad \text{as } t \to 0.$$

(g). Assume that $\mathrm{E}\,|X|$ is finite. For $h > 0$ write,

$$\frac{\varphi(t+h) - \varphi(t)}{h} = \mathrm{E}\,\exp(itX)\big(\frac{\exp(ihX) - 1}{h}\big),$$

and, taking limits as $h \to 0$, combine the inequality from (e) and the dominated convergence theorem to show that $\varphi'(t) = \mathrm{E}\,\{iX\exp(itX)\}$. Proceed inductively to show that, provided $\mathrm{E}\,|X|^r$ is finite, the rth derivative of the characteristic function is

$$\varphi^{(r)}(t) = \mathrm{E}\,\{(iX)^r \exp(itX)\}.$$

This shows that the moments can be read off from the characteristic function. Use the same proof technique as in (a) to show that $\varphi^{(r)}(t)$ is uniformly continuous.

Ex. A.40. *Uniqueness of characteristic functions.* If two random variables have identical characteristic functions, then their distributions are identical too. This fact is proved via so-called 'inversions theorems', providing a mechanism for finding the distribution of a random variable from its characteristic function.

(a). One such inversion formula is as follows: If the random variable X has characteristic function $\varphi(t)$ that is integrable, i.e., $\int |\varphi(t)|\,\mathrm{d}t < \infty$, then X has density f, for which a formula is

$$f(x) = \frac{1}{2\pi}\int \exp(-itx)\varphi(t)\,\mathrm{d}t.$$

Write down what this means, in the cases of a normal and a Cauchy, and verify the implied formulae. Show that f in each such case of an integrable $\varphi(t)$ necessarily becomes continuous.

(b). As a small digression, show that the characteristic function of the uniform distribution on $[-1,1]$ is $\varphi(t) = \sin(t)/t$. Deduce that

$$\int |\frac{\sin t}{t}| \, \mathrm{d}t = \infty \quad \text{even though} \quad \int \frac{\sin t}{t} \, \mathrm{d}t = \pi.$$

(c). Point (a) above gives a formula for the density f of a random variable, in the case of it having an integrable characteristic function φ. One also needs a more general formula, for the case of random variables that do not have densities, etc. Let X be any random variable, with cumulative distribution function F and characteristic function φ (but with nothing assumed about it having a density), and add a little Gaußian noise to X,

$$U_\sigma = X + \sigma Z, \text{ with } Z \sim \mathrm{N}(0,1),$$

with $\sigma > 0$ and Z is independent of X. Then U_σ has a density, even if X does not have one. Our intention is to let $\sigma \to 0$, to come back to X. Show that U_σ has cumulative distribution function and density of the form

$$F_\sigma(u) = \int \Phi\big(\frac{u-x}{\sigma}\big) \, \mathrm{d}F(x), \quad \text{and} \quad f_\sigma(u) = \int \frac{1}{\sqrt{2\pi}\sigma} \exp\big(-\frac{(u-x)^2}{2\sigma^2}\big) \, \mathrm{d}F(x),$$

where $\Phi(z) = \int_{-\infty}^{z} (1/\sqrt{2\pi}) \exp(-\frac{1}{2}x^2) \, \mathrm{d}x$ is the standard normal distribution function.

(d). Verify that the characteristic function of U_σ is $\varphi(t)\exp(-\frac{1}{2}t^2\sigma^2)$, and that it is integrable. Thus, according to (a) we must have that the density of U_σ is

$$f_\sigma(u) = \frac{1}{2\pi} \int \exp(-itu)\varphi(t) \exp(-\tfrac{1}{2}t^2\sigma^2) \, \mathrm{d}t.$$

To see that this equality is true, start by showing that $\mathrm{E} \int \exp\{it(X-u) - \frac{1}{2}t^2\sigma^2\} \, \mathrm{d}t = 2\pi f_\sigma(u)$, and apply Fubini's theorem.

(e). Deduce from (d) that for $a < b$,

$$\Pr(U_\sigma \in (a,b]) = \frac{1}{2\pi} \int \frac{\exp(-itb) - \exp(-ita)}{-it} \varphi(t) \exp(-\tfrac{1}{2}t^2\sigma^2) \, \mathrm{d}t,$$

inversion formula

and from this, derive the general inversion formula, valid for all continuity points a and b ($a < b$) of F,

$$F(b) - F(a) = \lim_{\sigma \to 0} \frac{1}{2\pi} \int \frac{\exp(-itb) - \exp(-ita)}{-it} \varphi(t) \exp(-\tfrac{1}{2}t^2\sigma^2) \, \mathrm{d}t.$$

(f). Assume that X and Y are random variables with identical characteristic functions. Show that X and Y must be equal in distribution.

(g). Let $X_1, \ldots, X_n$ be independent $\mathrm{N}(\mu_j, \sigma_j^2)$ random variables. Show that $\sum_{j=1}^n X_i$ has a normal distribution.

(h). By the inequality in Ex. A.39(e), and using that the absolute value of the complex exponential is 1, show that for all real x and y,

$$|\exp(iy) - \exp(ix)| \le |x - y|.$$

characteristic functions with finite integral

Now, let X be a random variable with characteristic function φ, and suppose that $\int |\varphi(t)| \, \mathrm{d}t$ is finite. Show that X has density $f(x) = (2\pi)^{-1} \int \exp(-itx)\varphi(t) \, \mathrm{d}t$, meaning that the claim in (a) is true.

Ex. A.41. *Characteristic functions for vector variables.* The characteristic function of a random vector $X = (X_1, \ldots, X_k)^{\mathrm{t}}$ is

$$\varphi(t) = \mathrm{E}\, \exp(it^{\mathrm{t}} X) = \mathrm{E}\, \exp\{i(t_1 X_1 + \cdots + t_k X_k)\},$$

for $t = (t_1, \ldots, t_k)^{\mathrm{t}}$, so $\varphi(t) = \varphi(t_1, \ldots, t_k)$.

(a). Show that the properties from Ex. A.39, with the appropriate amendments, generalises to the k-dimensional case.

(b). Show that if the components of X are independent, then

$$\varphi(t_1, \ldots, t_k) = \varphi_1(t_1) \cdots \varphi_k(t_k),$$

in terms of the individual characteristic functions $\varphi_1, \ldots, \varphi_k$. Thus, the characteristic function of any subset of $(X_1, \ldots, X_k)$ can be retrieved by setting the appropriate subset of $(t_1, \ldots, t_k)$ to zero.

(c). Show for the multinormal case, where $X \sim \mathrm{N}_k(\xi, \Sigma)$, that $\varphi(t) = \exp(it^{\mathrm{t}}\xi - \frac{1}{2} t^{\mathrm{t}} \Sigma t)$.

(d). The (one-dimensional) inversion formula derived in Ex. A.40 generalises to the k-dimensional case. In analogy to that exercise, let $Z_1, \ldots, Z_k$ independent standard normals, and define $U_{\sigma,j} = X_j + \sigma Z_j$ for $j = 1, \ldots, k$. Amend the proof from said exercise to show that with F_σ the c.d.f. of $U_\sigma = (U_{\sigma,1}, \ldots, U_{\sigma,k})^{\mathrm{t}}$,

$$F_\sigma(b) - F_\sigma(a) = \frac{1}{(2\pi)^k} \int_{\mathbb{R}^k} \left(\prod_{j=1}^k \frac{\exp(-it_j b_j) - \exp(-it_j a_j)}{-it_j} \right) \exp(-\tfrac{1}{2} t^{\mathrm{t}} t \sigma^2) \varphi(t)\, \mathrm{d}t,$$

where $\varphi(t)$ is the characteristic function of X and we use the shorthand $\mathrm{d}t = \mathrm{d}t_1 \cdots \mathrm{d}t_k$.

(e). Deduce the inversion formula for random vectors: Let F be the c.d.f. of $X = (X_1, \ldots, X_k)$, and φ its c.f., and show that

$$F(b) - F(a) = \lim_{\sigma \to 0} \frac{1}{(2\pi)^k} \int_{\mathbb{R}^k} \left(\prod_{j=1}^k \frac{\exp(-it_j b_j) - \exp(-it_j a_j)}{-it_j} \right) \exp(-\tfrac{1}{2} t^{\mathrm{t}} t \sigma^2) \varphi(t)\, \mathrm{d}t,$$

for all continuity points $a = (a_1, \ldots, a_k)$ and $b = (b_1, \ldots, b_k)$ of F.

(f). Show that if $X = (X_1, \ldots, X_k)$ and $Y = (Y_1, \ldots, Y_k)$ are random vectors with identical characteristic functions, then they are also identical in distribution. uniqueness of c.f.

(g). Let $X = (X_1, \ldots, X_k)$ and $Y = (Y_1, \ldots, Y_k)$ be two random vectors. Show that if $c^{\mathrm{t}} X = c_1 X_1 + \cdots c_k X_k = c_1 Y_1 + \cdots c_k Y_k = c^{\mathrm{t}} Y$ for all vectors $c = (c_1, \ldots, c_k)^{\mathrm{t}}$, then X and Y are identical in distribution.

(h). Let $X = (X_1, \ldots, X_k)^{\mathrm{t}}$ be a random vector with integrable characteristic function, i.e., $\int_{\mathbb{R}^k} |\phi(t)|\, \mathrm{d}t < \infty$. Show that X has density

$$f(x_1, \ldots, x_k) = \frac{1}{(2\pi)^k} \int_{\mathbb{R}^k} \left(\prod_{j=1}^k \exp(-it_j x_j) \right) \varphi(t)\, \mathrm{d}t,$$

with respect to Lebesgue measure on $\mathbb{R}^k$.

Ex. A.42. *Smoothness of characteristics functions.* In Ex. A.39(g) we have seen that the more moments a random variable X has, the smoother its characteristic function $\varphi(t)$ is. In (a) , which is the Riemann–Lebesgue lemma, the smoothness of $\varphi(t)$ is connected to the behaviour of $\varphi(t)$ as $|t|$ tends to infinity.

(a). Suppose that X has density f with respect to Lebesgue measure. Show that $\varphi(t) \to 0$ as $|t| \to \infty$. Show that if f has $k \geq 1$ integrable derivatives, then $\varphi(t) = o(1/t^k)$ as $|t| \to \infty$.

(b). Suppose that the characteristic function φ of the random variable X is such that $|\varphi(t)| = 1$ for all t. Show that X must be equal to a constant, i.e., it has a degenerate distribution.

(c). Show that $\sup_{|t|>\varepsilon} |\varphi(t)| < 1$ for any $\varepsilon > 0$.

(d). Generalise the above to higher dimensions, i.e., the characteristic functions of random vectors of dimension $k \geq 2$.

Ex. A.43. *Uniqueness of moment-generating functions.* Moment-generating functions characterise distributions: If X and Y are random variables such that $\mathrm{E}\exp(tX) = \mathrm{E}\exp(tY) < \infty$ for all $t \in (-t_0, t_0)$, then X and Y have the same distribution.

(a). Suppose that $\mathrm{E}\, f(X) = \mathrm{E}\, f(Y)$ for all bounded and continuous functions f. Approximate the indicator function $I\{y \leq x\}$ by a sequence of bounded and continuous functions to show that $F(x) = G(x)$ for all x. From Ex. A.15(d), this entails that X and Y have the same distribution.

(b). Let X and Y be random variables with distributions F and G on the unit interval, with identical moment-generating functions, $\int_0^1 \exp(tx)\,\mathrm{d}F(x) = \int_0^1 \exp(tx)\,\mathrm{d}G(x)$ for all $t \in (-t_0, t_0)$, say. Show that then $\int_0^1 p(x)\,\mathrm{d}F(x) = \int_0^1 p(x)\,\mathrm{d}G(x)$ for all polynomials $p(x)$. Use the Weierstraß approximation theorem, see Ex. 2.17, to show that this equality must hold for all continuous functions f. Point to (a) and conclude.

(c). Suppose that X and Y have identical moment-generating functions that are finite on $(-t_0, t_0)$. Let φ_X and φ_Y be their characteristic functions. We follow Billingsley (1995) in showing that $\varphi_X = \varphi_Y$. Appeal to Ex. A.36(c) to argue that

$$\lim_{k\to\infty} \frac{t^{2k}\,\mathrm{E}\,|X|^{2k}}{(2k)!} = 0, \quad \text{for } t \in (-t_0, t_0).$$

Let $0 < s < \min(t_0, 1)$, and fix an $0 < r < s$. Argue that there is an k_0 so that $2kr^{2k-1} < s^{2k}$ for all $k \geq k_0$. Use this and the inequality $|x|^{2k-1} \leq 1 + |x|^{2k}$ to show that $\lim_{k\to\infty} t^{2k-1}\,\mathrm{E}\,|X|^{2k-1}/(2k-1)! = 0$ for $t \in (-t_0, t_0)$, as well. Next, use the inequality in Ex. A.39(e) to show that

$$\Bigl|\varphi_X(t+h) - \sum_{k=0}^{m-1} \frac{t^k}{k!}\varphi_X^{(k)}(t)\Bigr| \leq \frac{|h|^m}{m!}\mathrm{E}\,|X|^m, \quad \text{for } h \in (-r, r),$$

with a similar inequality holding for φ_Y for the same h, where $\varphi_X^{(k)}(t)$, $k = 0, 1, \ldots$ are the derivatives from Ex. A.39(g). Since X and Y have identitical moment sequences (why?), $\varphi_X^{(k)}(0) = \varphi_Y^{(k)}(0)$ for $k = 0, 1, 2, \ldots$, and so $\varphi_X(t) = \varphi_Y(t)$ for $t \in (-r, r)$. For an arbitrary $0 < \varepsilon < r$, consider the inequality above with $t = r - \varepsilon$ and $t = -r + \varepsilon$, and argue that $\varphi_X(t) = \varphi_Y(t)$ for $t \in (-2r, 2r)$. Use the same argument with $t = 2r - \varepsilon$ and $t = -2r + \varepsilon$, to argue that $\varphi_X(t) = \varphi_Y(t)$ for $t \in (-3r, 3r)$, and so on, forevermore.

(d) (xx perhaps other conditions ensuring that identical moment sequences ensure identical distribiutions. Then a few counterexamples! xx).

Ex. A.44. *Posterior distributions without Bayes.* In Chapter 15 we will need to find posterior distributions without densities, i.e., in situations where neither of the formulae of Ex. A.29(d) are applicable.

(a).

(b).

Ex. A.45. *Something More.* (xx not yet an exercise, but a place to jot down a few comments, also as of 12-March-2025. we need the 'double variance' formula too. and show we round off ChZero with a few things to make the readers feel 'aha, so after all of this, we can do familiar things again', with ordinary integrals and sums and means and variances. perhaps a few simple but nonstandard things too. xx)

Notes and pointers

(xx we point to some of the many books on measure theory for probability and statistics, and also to where key ideas originated. Kolmogorov (1933a,b). Billingsley (1968); Royden and Fitzpatrick (2010) also Shiryaev (1996) and Williams (1991) and Kallenberg (2002). also, briefly, to 'what is a statistical model', McCullagh (2002). Cantor set and Cantor function, F is continuous on $[0,1]$ but not at all absolutely continuous. In connection with Ex. 2.5, mention that $X_n \to_d X$ and X_n uniformly integrable, implies $\mathrm{E}\, X_n \to \mathrm{E}\, X$, and that the proof of this employs a theorem of Skorokhod, see Billingsley (1995, p. 338) xx)

Notes to Ex. A.29. A conditional distribution is said to be *regular* if $P_X(B \,|\, \mathcal{G})(\omega)$ is a distribution on $(\mathcal{X}, \mathcal{B})$ for Pr-almost all $\omega \in \Omega$. Not all conditional distributions are regular (see, e.g. Dudley (2002, Problem 6, p. 351)), but, fortunately, if the measurable space $(\mathcal{X}, \mathcal{B})$ is composed of a complete and separable metric space $(\mathcal{X}, d)$, and the Borel-σ-algebra $\mathcal{B}$ on $\mathcal{X}$ (see Ex. A.2(m)), then there exists a regular conditional distribution of X given $\mathcal{G}$ (see, e.g., Dudley (2002, Theorem 10.2.2, p. 345) or Schervish (1995, Lemma B.40, p. 621)). In this book, we deal exclusively with complete and separable metric spaces, and all conditional distributions will be assumed regular without further mention.

For Ex. A.23, Emil is indebted to the lecture notes for the course Statistics 381: Measure-Theoretic Probability 1, by Steven Lalley, at the University of Chicago. Provide a citation

Leftovers

Ex. A.46. *Characteristic functions with finite integral.* [xx some of this is worked into Ex. A.40 xx]

(a). If X has an integrable characteristic function φ, that is, if $\int |\varphi(t)|\, \mathrm{d}t < \infty$, show that X has the continuous density

$$f(x) = \frac{1}{2\pi} \int \exp(-itx)\varphi(t)\, \mathrm{d}t.$$

[xx Show that this density is smooth xx]

(b). Show that $\sin t/t$ is not integrable. Then show that

$$f(x) = \frac{1}{2\pi} \int \cos(tx) \Big(\frac{\sin t}{t} \Big)^2 \, \mathrm{d}t$$

becomes the triangular density on $[-2, 2]$, linear from 0 to $\frac{1}{2}$ on $[-2, 0]$, linear from $\frac{1}{2}$ to 0 on $[0, 2]$.

(c). (xx one more. xx)

III.B

Overview of stories and data

A generous number of real datasets are used in this book to illustrate aspects of the methodology being developed, and hence also a correlated generous list of Statistical Stories. Here we provide brief descriptions of each of these real data examples, along with key points to indicate which substantive questions they relate to. Some of these datasets are small, partly meant for simpler demonstrations of certain methods, whereas other are bigger, allowing also more ambitious modelling for reaching inference conclusions. Key words are included to indicate the data sources, the types of model we apply and for what inferential goal, along with pointers to which stories the datasets are analysed.

nils emil notes to themselves, as of 12-March-2025

Not yet well organised, as of 12-March-2025, but we'll get to things and have a clearer overview.

This part of Kioskvelter needs of course serious working-over and polish near the end of the process, and we need to achieve the right level of 'oi!, this might be interesting' and conciseness. Each entry needs (a) brief description of what the data are about, with a source or two; (b) description of which questions arise naturally; (c) briefly worded pointers to which Stories touch these data.

The 2.A list will be updated and polished during autumn 2023. We get the cross-referencing to work for `data:listofstories` etc.

(xx So, rough estimate of magical story number as of 12-March-2025: 18 + 13 + 15 + 8 + 12 + 9 + 11 = 86, but some of those listed 12-March-2025 might not come to fruition. nils pushes mothers babies I and mothers babies II from Demography Eip Medicine to Biology. we're not shying away from StoryX Part I and StoryX Part II to avoid having them too long, etc. our Final Number was meant to be 77 but could even become 88 if several of these are decently short. xx)

List of All Stories

(xx annotated as we proceed. as of 12-March-2025, meant to be helpful for the authors. it will become a better kept list later on, along with brief key words and pointers. the list of these brief dataset descriptions should follow our stories. we also need to monitor the lengths of the stories, though we do accept some shorter ones and some longer ones.

which stories have points more than from (a) to (h)? as of 12-March-2025, there are ok plans for how many stories, in the separate chapters? xx)

I Demography, Epidemiology, Medicine (approx 14)

(1) Cooling of newborns. (nils, 0.90 ferdig)

(2) Overdispersed children. (nils, 0.90 ferdig)

(3) Markovian children. (nils, 0.90 ferdig)

(4) IDU expulsion. (nils, 0.90 ferdig, maa koples til Beta process)

(5) Boys are born slightly bigger than girls. (nils, 0.90 ferdig)

(6) Mothers, babies, birthweights, factors. (nils, 0.90 ferdig)

(7) Mothers, babies, birthweights, logistic or logistic. (nils, 0.90 ferdig)

(8) Mothers, babies, birthweights, birth order. (nils, 0.90 ferdig)

(9) Time to 2nd child after stillborn (nils, 0.90 ferdig)

(10) A third child? (emil, 0.60 ferdig)

(11) PCI and confidence fusion from thirteen studies. (nils, 0.80 ferdig)

(12) Suicide attempt rates for Paroxetine vs. placebo. (nils, 0.85 ferdig)

(13) Pedro og Elvira (emil, 0.50 ferdig)

(14) Onset of menarche (nils, 0.90 ferdig)

(15) Eye retinopathy, FIC for logistic regressions. (nils, 0.85 ferdig)

(16) Cigarette consumption and lung cancer. (nils, 0.85 ferdig)

(17) Norwegian women and men of the future (nils, 0.10 ferdig)

(18) A cure model (emil, 0.20 ferdig)

II Art, History, Literature, Music (approx 11)

(1) Game of Thrones and the Wars of the Roses. (nils, 0.80 ferdig)

(2) Stride towards your bookshelves. (nils, 0.90 ferdig, kan gjoere noe mer)

(3) Tirant lo Blanc: When did Author B take over for Author A? (nils, 0.90 ferdig, maa deales med, linkes til Ch9 monitoring for change)

(4) The children of Odin. (nils, 0.90 ferdig)

(5) How many Abel first-day cover envelopes from 1929? (nils, 0.90 ferdig.)

(6) Markov and Pushkin. (nils, 0.85 ferdig, henter ting fra CLP, med modifikasjoner)

(7) And Quiet Does Not Flow the Don. (nils, 0.90 ferdig)

(8) Republic, Laws, Critias, Philebus, Politicus, Sophist, Timaeus. (nils, 0.85 ferdig, maa stelles med og forkortes og synches med Ch12)

(9) Presidents of the First Republic. (nils, 0.90 ferdig)

(10) Dangerous job assignment: Roman emperor. (nils, 0.90 ferdig)

(11) Lifelengths in Roman Era Egypt, 2100 years ago. (nils, 0.90 ferdig)

(12) Bach, Reger, organ fugues, and Wohltemperierte I und II. (nils, 0.80 ferdig, maa avrundes)

(13) How many piano tuners in Oslo? (nils, 0.75 ferdig)

III Economy, Political Science, Sociology (approx 10)

(1) Power law scaling for academics and support staff. (nils, 0.85 ferdig)

(2) Poisson overdispersion for British mining. (nils, 0.90 ferdig)

(3) Changepoints for British mining. (nils, 0.90 ferdig)

(4) War and Peace and War and Peace, I. (nils, 0.80 ferdig, linkes til boundary parameter things in Ch5, Ch11)

(5) War and Peace and War and Peace, II. (nils, 0.50 ferdig, trenger mer om changepoints, Bayes too)

(6) War and Peace and War and Peace, III. (nils, 0.60 ferdig, men tre angels stories er kanskje en for meget)

(7) Psychiatric disorders and body type. (nils, 0.90 ferdig)

(8) Galton and 111 husbands and wives. (nils, 0.40 ferdig)

(9) Terbeschikkingstelling. (nils, 0.80 ferdig)

(10) Tore Sims Sveriges Riksbank, (nils, 0.90 ferdig)

(11) Does winning make you live longer? (emil, 0.75 ferdig)

(12) Minimum Wages and Employment. (emil, 0.75 ferdig)

(13) How many were killed in Srebrenica, 1991? (nils, 0.80 ferdig)

(14) How many were killed in Guatemala, 1978–1996? (nils, 0.80 ferdig)

(15) Volatility estimation from noisy financial data. (emil, 0.02 ferdig. Har data og vet omtrent hva som skal skrives)

IV Biology, Climate, Ecology (approx 6)

(1) New Haven annual temperatures 1912–1971. (nils, 0.85 ferdig)

(2) Oslo annual temperatures 1901–20024. (nils, 0.85 ferdig)

(3) Where are the snows of yesteryear? (nils, 0.80 ferdig)

(4) Mammals and their bodies and brains. (nils, 0.80 ferdig, need a a little link til Jamtveit story)

(5) Kola temperatures and the Hjort liver index time series 1859-2020. (nils, 0.75 ferdig)

(6) How many Clethrionomys glareoli? (nils, 0.80 ferdig, venter paa Ch7 polishing)

(7) Birds on islands outside Ecuador. (nils, 0.90 ferdig)

(8) Birds on islands, via square-rooting to normal nonlinear regression. (nils, 0.90 ferdig; may be subsumed in previous story)

V Sports (approx 12)

(1) Bolt from heaven. (nils, 0.90 ferdig, venter paa litt Ch7 avrunding)

(2) Will there be a new world record this season? (nils, 0.90 ferdig)

(3) The random angles of golf putting. (nils, 0.90 ferdig)

(4) NBA three point shooting averages. (emil, 0.80 ferdig)

(5) Olympic Unfairness I: Inner and outer for 1000 m speedskating. (nils, 0.80 ferdig, skal ha CD(τ) med punktmasse m.m.

(6) Olympic Unfairness II: From semifinals to finals. (nils, 0.75 ferdig)

(7) Who wins? Computing probabilities as a function match time. (nils, 0.85 ferdig)

(8) The turn-around operation: from 0-2 to 3-2. (nils, 0.85 ferdig, kan kaste inn fotballdata fra Nils-Gerda til sist)

(9) The hot hand in basketball. Myth or not? (emil, 0.05 ferdig. har data og har plottet litt)

(10) When to shoot? An optimal stopping problem in basketball. (emil, 0.02 ferdig. Har et datasett med hvilket dette sporrsmaalet kan besvares)

(11) Winning streaks in chess. (nils-emil, 0.03 ferdig; need book-keeping scripts from Alamaty 2022)

(12) BMI for Olympic speedskaters. (nils, 0.10 ferdig, tas fra II-CC-FF pluss litt mer)

VI Simulated stories (approx 7)

(1) Checking out the CLT. (nils, 0.85 ferdig)

(2) An infinite weighted sum of Bernoullis. (nils, 0.85 ferdig)

(3) Finding magic squares by MCMC. (nils, 0.85 ferdig)

(4) Reconstructing exponential decay beer foam. (nils, 0.85 ferdig)

(5) From falling ill to having recovered. (nils, 0.65 ferdig, vi maa ha noe simulation)

(6) Causal inference and potential outcomes. (emil, 0.33 ferdig, haaper noe kan skrives rundt dette)

(7) Identifying your counterfactual cousin. (nils, 0.80 ferdig)

(8) Recostructing a blurry image. (nils, 0.01 ferdig)

(9) Law, evidence, and Bayes (emil, 0.01 ferdig. Ideen er der. Noe basert paa Calina sin master. Bayesianske nett og noen simuleringer)

VII Miscelannea (approx 6)

(1) The Pearson goodness-of-fit statistic. (nils, 0.85 ferdig, vil ha link til Brownian bridge, og til ML)

(2) Decimals of pi. (nils, 0.75 ferdig, flikker inn dette med aa jukse for hver 1000ende desimal, for aa se om pearson oppdager det, local power)

(3) Random integers via prime numbers. (nils, 0.75 ferdig, gjoer noe mer)

(4) Time-to-failure for machine components. (nils, 0.90 ferdig, taken from Ch6)

(5) Speed of light in 1882, with BHHJ estimation. (nils, 0.85 ferdig)

(6) Stout's physician and the last n. (nils, 0.90 ferdig)

(7) Boys, girls, and mathematics scores (nils, 0.90 ferdig)

(8) Lean body mass, percent body fat, and correlations. (nils, 0.80 ferdig)

(9) Eddington versus Fisher, 1920. (nils, 0.85 ferdig)

(10) Elo rating and how to classify data as f_0 or f_1. (nils, 0.90 ferdig)

(11) Elo rating and logistic versus probit regression. (nils, 0.90 ferdig)

Description of datasets for stories

(xx needs careful work and checking, with crossref. also regarding the order in which we present these datasets. as of 12-March-2025, nils attempts to put these in the order of appearance through the chapters of stories, but this will be modified later. nils tentatively gives numbers to the stories too, perhaps to be take away later. xx)

I (i) The cooling of newborns

[xx Laptook study, point to two smaller papers by Walløe, Thoresen, Hjort, and the Hjort blogpost. Laptook (2017), Walløe et al. (2019a,b). basically: inference for p_1/p_2. return to this in the Bayes chapter, with informative priors on p_1, p_2. basically: $y_0 \sim \text{binom}(m_0, p_0)$ for the noncooled group, 22 of 79 cases; and $y_1 \sim \text{binom}(m_0, p_1)$ for the cooled groups, 19 of 68 cases. xx]

I (ii) Suicide attempt rates, drug vs. placebo

(xx describe data, two Poisson, in to round, from Aursnes et al. (2005, 2006). suicide rates for users of a certain antidepressant, vs. a placebo group. for Story i.12. xx)

I (iii) and (iv) Overdispersed and Markovian children

[xx data from Geißler (1889), on the number of girls in 8-children-families. we demonstrate extra-binomial variability, and assess this degree of overdispersion. see Hjort (2018a). in later chapter also the Markovian model. xx]

y	N	E1	pear1	E2	pear2
0	264	192.325	5.168	255.621	0.524
1	1655	1445.384	5.514	1657.032	-0.050
2	4948	4752.364	2.838	4909.686	0.547
3	8498	8928.902	-4.560	8683.213	-1.988
4	10263	10484.952	-2.168	10024.863	2.378
5	7603	7879.792	-3.118	7735.975	-1.512
6	3951	3701.205	4.106	3896.509	0.873
7	1152	993.421	5.031	1171.238	-0.562
8	161	116.655	4.106	160.865	0.011

(xx In this version of our manuscript we also give the girl-boy data for 195 families worked with in i.3, from Klotz (1972, 1973). In the Real Book we might be content with having only summary statistics. i.e. give our readers $M = 85$, $n - M = 110$, $(N_{0,0}, N_{0,1}, N_{1,0}, N_{1,1}) = (345, 298, 287, 333)$, not bother readers with the full dataset. here 1 is girl and 0 is boy. nei, forresten, nils ombestemmer seg, det tar for meget plass i pdf-en og det hele, so the 195 families can be found here, by nils and emil, but they've been commented away. xx)

I (v) IUD expulsion

(xx describe data, 100 women using IUD, (i, t_i, δ_i), from Peterson (1975). frailty model. check Aalen1978 sjs. used in Story i.4. xx)

I (vi) The 73 French presidents of France 1792–1795

(xx drama following the French Revolution. 73 presidents, some with short lives after elections. two covariates. data from Céline Cunen. Story: ii.9. xx)

Data, for 73 presidents: `id, birth, death, presistart, presiend, v, giro, vip`. Here v is indicator for having experienced a violent death of not (so $\delta = 1 - v$ is indicator for non-censoring, in survival analysis language); giro is indicator for belonging to the Gironde party or not; and vip is a proxy for fame, counted here as the number of languages for which there is wikipedia pages about the president in question (recorded, by Cunen, in October 2017).

I (vii) Boys and girls born in Oslo in 2001-2008

In the larger context of identifying and exploring factors that may influence the chances of babies being born too large, Voldner et al. (2008) examined a certain cohort of 1028 mothers and children (548 boys, 480 girls), all of whom born at Rikshospitalet, Oslo, in the 2001–2008 period. We have had access to the birthweight data in question (via N. Voldner and K.F. Frøslie, personal communication), and use these to illustrate non-parameetric confidence curves for quantiles.

(xx here: too much to show all the 548 + 480 data points, but we point to files, and the figures. xx)

(xx nota bene: will use these data also for other chapters. xx)

I (viii) Lifelengths in Roman era Egypt

In Spiegelberg (1901) the age at death has been recorded for 141 Egyptian mummies, 82 men and 59 women, dating from the Roman period of ancient Egypt from around year 100 B.C. These life-lengths vary from 1 to 96 years, and Pearson (1902) argued that these can be considered a random sample from one of the better-living classes in that society, at

a time when a fairly stable and civil government was in existence. Interestingly, Pearson wrote "in dealing with [these data] I have not ventured to separate the men from the women mortality, the numbers are far too insignificant". That did not stop Claeskens and Hjort (2008b, p. 33–35) from establishing partly different parametric hazard rate models for the two sexes, and where incidentally Gompertz type models were found to be better than e.g. Weibull and gamma models.

(xx then pointers to where we do what. Story i.5 for ratio of quantiles. xx)

```
 1.5  1.8  2.0  2.0  3.0  3.0  3.0  4.0  4.0  4.0  4.0  5.0  5.0  5.0  6.0
10.0 11.0 14.0 14.0 16.0 17.0 17.0 19.0 20.0 20.0 21.0 22.0 22.0 23.0 24.0
24.0 25.0 25.0 25.0 25.0 25.0 26.0 26.0 26.0 26.0 27.0 29.0 30.0 32.0 33.0
33.0 36.0 36.0 37.0 40.0 40.0 40.0 46.0 48.0 48.0 50.0 50.0 50.0 50.0 50.0
52.0 52.0 52.0 55.0 55.0 59.0 60.0 60.0 60.0 60.0 62.0 63.0 65.0 65.0 68.0
68.0 70.5 72.0 72.0 72.0 84.0 90.0  1.5  3.0  3.0  4.0  4.0  4.0  6.0  6.0

 9.0 10.0 11.0 14.0 16.0 17.0 17.0 17.0 18.0 18.0 19.0 19.0 20.0 20.0 20.0
21.0 21.0 21.0 21.0 21.0 21.0 22.0 22.0 23.0 23.0 25.0 25.0 25.0 25.0 26.0
27.0 28.0 29.0 30.0 30.0 33.0 35.0 35.0 35.0 36.0 36.0 40.0 40.0 40.0 50.0
52.0 54.0 55.0 60.0 70.0 96.0
```

I (ix) PCI and confidence fusion from meta-analyses

(xx from Schömig et al. (2008). thirteen two-by-two tables. rare events. xx)

I (xi) Terbeschikkingstelling

(xx from Hjort and Koning (2002), Dutch ombudsman. poisson with changing rates. xx)

I (xii) Brazilian kids

Data from Borgan et al. (2007), come in the three files `prevdat.txt`, `atprevrisk.txt`, and `covariates.txt`. Across these three files, the children are identified by their row number. The first file, `prevdat.txt`, contain sequences of zero and ones, a one indicating the the child on the given row was observed with diarrhoea on the day in question, and zero otherwise. The data file `atprevrisk.txt` contains indicators for noncensoring, 1 indicates that the child was observed on the day in question, 0 indicating that no observation was made (i.e., censoring). The third file contain the covariates:

`sexcat`: Gender of child (0-female , 1-male);

`age.beg`: Age of child in the begining of the study (month);

`dens2`: Number of people/beroom (0-one or two people, 1- >= three people);

`pavcat`: Quality of the street (0- good , 1- bad);

`resagcat`: Type of drinking water-reserve (0-good , 1- bad);

`tipagcat`: Quality of drinking water (0-good, 1-bad);

`valacat`: Presence of holes with dirty water-30m (0-no, 1-yes);

`corrcat`: Presence of small rivers with dirty water-30m (0-no , 1-yes);

`chovecat`: Condition of accomodation during rain (0-good , 1-bad);

`age.mcat`: Age of the mother 0- >= 25 years old, 1- < 25 years old;

`socec2`: Social economic class (0-medium or high, 1- poor);

`no.age5`: Number of other children less and equal 5 (0=none, 1- more than 1).

II (i) The Game of Thrones and the Wars of Roses

[xx via Cunen (2015) blog post and story (bringing her international fame). full dataset: $(t_i, x_{i,1}, x_{i,2}, \delta_i, o_i)$, with t_i the age at death if $\delta_i = 1$ and the last known age for the person if $\delta_i = 0$; $x_{i,1}$ indicator for nobility; $x_{i,2}$ concerns gender, and is 1 for women and 0 for men; indicator for nobility; and o_i is 1 for GoT and 0 for WoR. also with violent-death-or-not, from two worlds. see Story ii.1]. xx]

II (iii) Stride towards your bookshelves

[xx student books oblig story for chapter1: wordlengths. metaanalysis modelling, chapter2, simple analysis in chapter1, and we return to the dataset later, in chapter6. xx]

	xNor	xEng	sdNor	sdEng		xNor	xEng	sdNor	sdEng
1	4.45	2.88	4.53	2.77	33	4.76	3.10	4.10	2.02
2	4.82	2.88	4.87	2.94	34	4.58	2.79	4.03	1.81
3	4.39	2.24	3.94	2.04	35	6.32	4.02	4.49	2.41
4	4.44	2.65	4.02	1.90	36	3.94	1.96	5.30	2.78
5	4.79	3.11	4.56	2.24	37	3.54	1.77	4.73	2.86
6	4.91	3.12	4.05	1.86	38	4.41	2.15	4.16	2.20
7	4.61	2.42	4.64	3.01	39	4.57	2.98	4.28	1.95
8	3.91	1.94	4.48	2.18	40	4.34	2.24	4.31	2.17
9	4.42	2.52	3.94	2.09	41	5.12	3.16	4.33	2.26
10	4.86	2.61	5.21	2.86	42	4.47	2.69	3.81	2.00
11	6.18	4.62	5.43	2.89	43	4.29	2.33	4.08	1.95
12	5.04	3.47	4.97	2.86	44	3.87	2.06	4.72	2.26
13	3.99	2.71	4.33	2.52	45	3.96	2.19	4.20	2.49
14	4.35	2.64	4.82	2.36	46	4.43	2.40	4.74	2.53
15	4.56	2.93	5.23	3.35	47	5.01	2.88	4.47	2.35
16	5.27	3.10	5.16	3.20	48	4.10	1.98	3.69	1.69
17	4.88	3.07	4.17	2.25	49	4.98	2.99	4.15	2.06
18	4.33	2.42	4.58	2.60	50	4.17	2.21	4.15	2.43
19	3.90	1.80	4.80	2.61	51	5.26	3.80	4.92	2.32
20	4.19	2.04	4.37	2.35	52	4.13	2.25	3.74	1.92
21	4.41	2.43	4.20	2.11	53	4.65	3.07	4.57	2.55
22	4.10	2.06	4.35	1.73	54	5.68	3.92	4.79	2.38
23	3.84	1.80	4.92	2.55	55	4.60	2.44	4.61	2.42
24	3.82	1.83	3.88	1.77	56	4.18	2.10	4.88	2.78
25	4.61	3.18	3.96	2.41	57	4.55	3.18	4.32	1.87
26	4.44	2.60	4.24	2.03	58	4.40	2.74	4.39	2.15
27	4.26	2.44	3.99	2.31	59	4.81	2.82	4.76	2.42
28	5.11	3.10	4.45	2.55	60	5.11	3.26	4.68	2.40
29	4.38	2.74	4.33	2.44	61	4.29	2.48	4.38	2.57
30	4.18	2.65	4.97	2.40	62	4.27	2.35	4.07	2.16
31	4.41	2.42	3.56	1.70	63	4.84	2.73	4.21	2.29
32	5.52	3.14	4.83	2.67	64	4.17	2.30	4.07	1.81

As part of the obligatory exercises work for an introductory course on statistical methodology at the Department of Mathematics, University of Oslo, we made our students carry out the following task. Each student was told to stride towards her or his bookshelves, to pick one book in Norwegian and one in English, then record the lengths of the first 100 words on page 51. The books could be novels, collections of short stories, poetry, or prose in general, but not technical material (as with matematics or statistics); and the students were instructed to use page 52 if page 51 didn't have enough words. Do Fløgstad, Kjærstad, Solstad tend to use words with more or less the same lengths as do Miller, Lessing, Munro? What about Askeladden and Winnie Pooh?

The students were asked to summarise information and to compare their own two datasets in terms of means and standard deviations. This was expected to involve tests for equality of means and of variances, confidence intervals for differences, perhaps comments on skewnesses, etc. But the experiment also gave us an interesting combined data set, where we recorded the empirical mean and standard deviation for each dataset, for the two languages, for each student. In other words, we have summary statistics data $(x_i, \widehat{\kappa}_{1,i}, y_i, \widehat{\kappa}_{2,i})$ for $i = 1, \ldots, n$, for the $n = 64$ students, with

$$\begin{aligned} x_i &= \text{average word-length for 100 Norwegian words for student } i, \\ y_i &= \text{average word-length for 100 English words for student } i, \end{aligned} \tag{B.1}$$

along with

$$\begin{aligned} \widehat{\kappa}_{1,i} &= \text{standard deviation for the 100 Norwegian word-lengths for student } i, \\ \widehat{\kappa}_{2,i} &= \text{standard deviation for the 100 English word-lengths for student } i, \end{aligned} \tag{B.2}$$

Figure [xx nemlig xx] displays the word-length averages (x_i, y_i), with grand averages $\bar{x} = 4.549$ for Norwegian and $\bar{y} = 4.436$ for English words (so these are both averages over 6400 words).

There are at least four notable aspects of these data, each rather non-trivial when it comes to their precise assessments and analyses. We point to and briefly explain these aspects or features now, and come back to precise inference methods and results in later sections.

The first and second aspects we wish to call attention to and analyse is that the overall means are reasonably similar, for Norwegian and English, but that the overall variability measures, related to the sizes of the $(\widehat{\kappa}_{1,i}, \widehat{\kappa}_{2,i})$ of (B.2), are significantly different.

The third feature is that there is positive correlation between word-lengths; if x_i is small, for student i, then y_i also tends to be small, etc. Indeed we find $\rho = \text{corr}(x, y)$ estimated in the usual fashion as 0.283. This is a deflated correlation, however, in that there is an undiluted correlation, say $\rho_0 = \text{corr}(x_0, y_0)$, in which we would be more interested, where $(x_{0,i}, y_{0,i})$ are the real word-length averages for student i, across all sentences in all books on her or his bookshelves. This ρ_0, if we may estimate it, despite the (x_0, y_0) not being visible, gives a better measure of the extent to which long-worded books in Norwegian tend to be coupled with long-worded books in English (and similarly, for short-worded texts), for a given student (and her or his bookshelves). To arrive at estimates and inference for the ρ_0 we need to accept that the directly observed (x, y) are proxies for (x_0, y_0), not the real thing, and we need to model them. We shall learn that the real ρ_0 is notably bigger than ρ.

Finally the fourth aspect worth serious examination is the level of differences in average word-lengths, across students. There is sufficient variation from student to student, as we shall see, that neither the (x_i, y_i) of (B.1) nor the underlying and not visible $(x_{i,0}, y_{i,0})$ alluded to above can be seen as samples from the same homogeneous normal distribution. The statistical task becomes how to model and assess the level of variability, among the $x_{i,0}$ and the $y_{i,0}$.

II (iv) Platon: from Republic to Laws

(xx to come. $2^5 = 32$ different clausulae, the five last syllables in Platon's sentences, corpus for corpus. here '0' and '1' indicate 'short' and 'long' (or 'light' and 'stressed'),

for these last five syllables. data from Cox and Brandwood (1959). this is about ordering the last five works, Crit, Phil, Pol, Soph, Tim, on the timeline from corpus A, Republic ($n_A = 3778$ sentences), to corpus B, Laws ($n_B = 3783$ sentences). Sample sizes for Crit, Phil, Pol, Soph, Tim are smaller, 150, 958, 770, 919, 762 sentences. Story ii.8 in Ch 4. xx)

```
0 0 0 0 0  1.1  2.4  3.3  2.5  1.7  2.8  2.4

1 0 0 0 0  1.6  3.8  2.0  2.8  2.5  3.6  3.9
0 1 0 0 0  1.7  1.9  2.0  2.1  3.1  3.4  6.0
0 0 1 0 0  1.9  2.6  1.3  2.6  2.6  2.6  1.8
0 0 0 1 0  2.1  3.0  6.7  4.0  3.3  2.4  3.4
0 0 0 0 1  2.0  3.8  4.0  4.8  2.9  2.5  3.5

1 1 0 0 0  2.1  2.7  3.3  4.3  3.3  3.3  3.4
1 0 1 0 0  2.2  1.8  2.0  1.5  2.3  4.0  3.4
1 0 0 1 0  2.8  0.6  1.3  0.7  0.4  2.1  1.7
1 0 0 0 1  4.6  8.8  6.0  6.5  4.0  2.3  3.3
0 1 1 0 0  3.3  3.4  2.7  6.7  5.3  3.3  3.4
0 1 0 1 0  2.6  1.0  2.7  0.6  0.9  1.6  2.2
0 1 0 0 1  4.6  1.1  2.0  0.7  1.0  3.0  2.7
0 0 1 1 0  2.6  1.5  2.7  3.1  3.1  3.0  3.0
0 0 1 0 1  4.4  3.0  3.3  1.9  3.0  3.0  2.2
0 0 0 1 1  2.5  5.7  6.7  5.4  4.4  5.1  3.9

1 1 1 0 0  2.9  4.2  2.7  5.5  6.9  5.2  3.0
1 1 0 1 0  3.0  1.4  2.0  0.7  2.7  2.6  3.3
1 1 0 0 1  3.4  1.0  0.7  0.4  0.7  2.3  3.3
1 0 1 1 0  2.0  2.3  2.0  1.2  3.4  3.7  3.3
1 0 1 0 1  6.4  2.4  1.3  2.8  1.8  2.1  3.0
1 0 0 1 1  4.2  0.6  4.7  0.7  0.8  3.0  2.8
0 0 1 1 1  2.8  2.9  1.3  2.6  4.6  3.4  3.0
0 1 0 1 1  4.2  1.2  2.7  1.3  1.0  1.3  3.3
0 1 1 0 1  4.8  8.2  5.3  5.3  4.5  4.6  3.0
0 1 1 1 0  2.4  1.9  3.3  3.3  2.5  2.5  2.2

0 1 1 1 1  3.5  4.1  2.0  3.3  3.8  2.9  2.4
1 0 1 1 1  4.0  3.7  4.7  3.3  4.9  3.5  3.0
1 1 0 1 1  4.1  2.1  6.0  2.3  2.1  4.1  6.4
1 1 1 0 1  4.1  8.8  2.0  9.0  6.8  4.7  3.8
1 1 1 1 0  2.0  3.0  3.3  2.9  2.9  2.6  2.2

1 1 1 1 1  4.2  5.2  4.0  4.9  7.3  3.4  1.8
```

II (v) Tirant lo Blanc: the world's first novel

(xx the 425 chapters. sentence lengths as well as relative frequencies of words of lengths 1, 2, 3, 4, 5, 6, 7, 8, 9, 10. too much to give all data here, btu we give a figure and point to a data file. xx) [xx describe data. used in Cunen et al. (2018). word length frequencies, $(\widehat{p}_1, \ldots, \widehat{p}_{10})$, for 420 chapters. when did Author B take over for Author A? xx]

II (vi) Sons of Odin

Odin had six sons (though sources are not entirely clear on the matter): Thor, Balder, Vitharr, Váli (cf. the Eddic poems and the Snorri Edda), Heimdallr, Bragi (cf. Snorri's

kennings). In Story [xx nemlig xx] we construct a confidence distribution for the number of all of Odin's children. In some of the Snorri kennings there are also references to Týr and Höd as sons of Odin (and yet other names are mentioned in the somewhat apocryphical Skáldskaparmál). [xx pointer to brief stories about this, one CD, one Bayes, which prior fits the CD answer. xx]

II (vii) Markov and Pushkin

(xx describe data, versions of which used also in Hjort and Varin (2008); Schweder and Hjort (2016). vowels and consonants in Pushkin's classic 1833 epic poem Евгений Онегин. point to Markov (1906, 1913). analysed in Story ii.6. So *Heedless of the proud world's enjoyment* becomes

```
0 1 1 0 0 1 0 0  1 0  0 0 1  0 0 1 1 0  0 1 0 0 0  0  1 0 0 1 1 0 1 0 0
```

i t.d. xx)

II (viii) And Quiet Does Not Flow the Don

[xx describe data, Hjort (2007), mention Kjetsaa et al., Kjetsaa et al. (1984); Lessing (1997). Is it Sholokhov, or is it Kruikov? Grandest scandal in all of Nobel Prize history? analysed in Story ii.7. data are sentence lengths, sorted into window 1-5, 6-10, 11-15, and so on, for three text corpora: Sh (known to be Sholokhov), Kr (known to be Kriukov), TD (the apple of discord, Tikhij Don).

from	to	Sh	Kr	TD
1	5	799	714	684
6	10	1408	1046	1212
11	15	875	787	826
16	20	492	528	480
21	25	285	317	244
26	30	144	165	121
31	35	78	78	75
36	40	37	44	48
41	45	32	28	31
46	50	13	11	16
51	55	8	8	12
56	60	8	5	3
61	65	4	5	8

II (ix) Bach and the others

A fugue, whether for a piano, an organ, or an ensemble of instruments, starts with the principal fugue theme itself, before it is imitated and varied, perhaps in complex ways, in other voices; typical Bach fugues have from three to five voices. Rydén (2020) has studied fugue subjects in organ compositions of Bach and other composers, and has defined certain features, say $x_1, \ldots, x_6$, which can be worked out for each given organ fugue. In brief, x_1 is the length; x_2 the range; x_3 the number of unique notes; x_4 the start interval; x_5 the number of unique intervals betweem successive notes; and x_6 the max interval. These are described in a bit more detail in Story ii.12, where we investigate differences between the fugues of J.S. Bach (1685–1750) and Max Reger (1873–1916), using data from J. Rydén (personal communciation).

One of the present authors has also played through all 24 preludes and fugues from Bach's Das Wohltemperierte Klavier, part I (from the Anhalt-Köthen period, c. 1722), carefully recording $x_1, \ldots, x_6$ for each fugue, and similarly for all the 24 preludes and fugues from part II (from Bach's rather different Leipzig years, c. 1742). As part of Story ii.12 we investigate whether there are any notable differences between these Bach collections, or between the Clavier fugues and those for organ. (xx decide later whether we go for two consecutive stories. xx)

WTK I, c. 1722							WTK II, c. 1743						
		x1	x2	x3	x4	x5	x6			x1	x2	x3	x4	x5	x6
1	C	14	9	6	2	4	7	1	C	29	9	6	2	6	9
2	c	20	11	7	1	7	9	2	c	9	7	5	3	5	7
3	Ciss	20	15	11	2	10	12	3	Ciss	12	7	5	4	5	7
4	ciss	5	5	4	1	3	5	4	ciss	22	12	8	1	3	7
5	D	13	9	6	2	4	9	5	D	9	10	6	0	6	7
6	d	20	9	7	2	4	9	6	d	24	12	9	2	3	5
7	Diss	24	14	10	3	7	10	7	Diss	20	9	6	7	6	7
8	diss	13	8	6	7	5	7	8	diss	13	8	6	0	3	5
9	E	21	12	7	2	4	5	9	E	6	5	4	2	3	3
10	e	26	13	12	3	7	9	10	e	42	13	9	2	8	10
11	F	21	10	7	2	3	8	11	F	21	13	8	1	6	7
12	f	11	12	9	1	5	7	12	f	28	11	9	7	5	8
13	Fiss	16	10	7	5	4	5	13	Fiss	24	14	8	1	7	12
14	fiss	18	7	7	2	3	3	14	fiss	15	12	7	4	7	8
15	G	31	14	8	2	5	10	15	G	42	17	11	3	7	7
16	g	12	8	7	1	4	8	16	g	24	8	7	4	5	5
17	Giss	7	9	5	7	5	9	17	Giss	24	14	7	3	7	12
18	giss	12	12	9	1	5	9	18	giss	27	13	8	2	7	12
19	A	15	13	10	5	6	8	19	A	19	9	6	2	5	5
20	a	31	13	8	1	6	9	20	a	12	13	10	4	7	11
21	B	38	15	8	2	8	9	21	B	24	14	9	2	5	5
22	b	10	13	6	6	4	13	22	b	27	9	7	2	5	6
23	H	14	10	7	1	3	7	23	H	12	12	7	4	6	9
24	h	21	15	12	3	6	9	24	h	27	13	8	4	5	12

III (i) Power law scaling for academics and support staff

(xx from Jamtveit et al. (2009), Jamtveit et al. (2018). xx)

III (ii) Statistical Sightings of Better Angels

[xx Correlates of War data. several stories. Story I, iii.4, for Ch. 3: about the $w_i = x_i - x_{i-1}$. Story II, iii.5, and also Story III, iii.6, for Chs. 2: estimates and intervals for $\rho_q = \mu_{L,q}/\mu_{R,q}$, ratio of quantiles, using direct large-sample theory. here we take Korea 1950 as change-point, with $n_L = 60$ wars to the left and $n_R = 35$ wars to the right. story III, for Ch. 7: the ρ_q again, but now via II-CC-FF, more precise nonparametric confidence. story IV, finding the change-point, in Ch. 9. Pinker (2011), Cunen et al. (2020a), Gleditsch (2020) volume on the work of Lewis Fry Richardson. we examine the between-war times $w_i = x_i - x_{i-1}$, fit both the simple exponential, as per Poisson process assumption, and a two-parameter model extension. later briefly in Bayes chapter. xx]

(xx here is the full data matrix, for wars $i = 1, \ldots, 95$, with onset time x_i, the between-onset times $w_i = x_i - x_{i-1}$, the battle deaths count z_i, and its logarithm $\log z_i$. from com26 of nilswork21. xx)

i	x	w	z	log z	i	x	w	z	log z
1	1823.269	--	1000	6.908	49	1932.458	0.489	92661	11.437
2	1828.322	5.053	130000	11.775	50	1934.222	1.764	2100	7.650
3	1846.319	17.997	19283	9.867	51	1935.758	1.536	20000	9.903
4	1848.247	1.928	7527	8.926	52	1937.519	1.761	1000000	13.816
5	1848.278	0.031	6000	8.700	53	1938.581	1.061	1726	7.454
6	1849.333	1.056	2600	7.863	54	1939.364	0.783	28000	10.240
7	1851.553	2.219	1300	7.170	55	1939.917	0.553	151798	11.930
8	1853.814	2.261	264200	12.484	56	1940.919	1.003	1400	7.244
9	1856.819	3.006	2000	7.601	57	1941.492	0.572	16634907	16.627
10	1859.331	2.511	22500	10.021	58	1947.822	6.331	3500	8.161
11	1859.811	0.481	10000	9.210	59	1948.375	0.553	8000	8.987
12	1860.697	0.886	1000	6.908	60	1950.825	2.450	910084	13.721
13	1860.792	0.094	1000	6.908	61	1954.675	3.850	2370	7.771
14	1862.294	1.503	20000	9.903	62	1956.836	2.161	3221	8.077
15	1863.894	1.600	1000	6.908	63	1956.844	0.008	2426	7.794
16	1864.086	0.192	4481	8.408	64	1958.111	1.267	1122	7.023
17	1865.181	1.094	310000	12.644	65	1958.647	0.536	1800	7.496
18	1865.736	0.556	1000	6.908	66	1962.806	4.158	1853	7.525
19	1866.472	0.736	44100	10.694	67	1965.103	2.297	1021442	13.837
20	1870.553	4.081	204313	12.227	68	1965.597	0.494	7061	8.862
21	1876.242	5.689	4000	8.294	69	1967.431	1.833	19600	9.883
22	1877.317	1.075	285000	12.560	70	1968.036	0.606	13866	9.537
23	1879.122	1.806	13868	9.537	71	1969.183	1.147	5368	8.588
24	1882.531	3.408	10079	9.218	72	1969.539	0.356	1900	7.550
25	1884.458	1.928	12100	9.401	73	1970.231	0.692	6525	8.783
26	1885.244	0.786	1000	6.908	74	1971.925	1.694	11223	9.326
27	1894.569	9.325	15000	9.616	75	1973.767	1.842	14439	9.578
28	1897.125	2.556	2000	7.601	76	1974.556	0.789	1500	7.313
29	1898.311	1.186	3685	8.212	77	1975.814	1.258	2700	7.901
30	1900.464	2.153	3003	8.007	78	1977.564	1.750	10500	9.259
31	1900.547	0.083	4000	8.294	79	1977.733	0.169	8000	8.987
32	1904.106	3.558	151831	11.931	80	1978.828	1.094	3000	8.006
33	1906.408	2.303	1000	6.908	81	1979.131	0.303	21000	9.952
34	1907.136	0.728	1000	6.908	82	1980.728	1.597	1250000	14.039
35	1909.519	2.383	10000	9.210	83	1982.236	1.508	1001	6.909
36	1911.747	2.228	20000	9.903	84	1982.308	0.072	1655	7.412
37	1912.797	1.050	82000	11.314	85	1986.875	4.567	8000	8.987
38	1913.542	0.744	60500	11.010	86	1987.014	0.139	4000	8.294
39	1914.581	1.039	8578031	15.965	87	1990.589	3.575	41466	10.633
40	1918.894	4.314	11750	9.372	88	1992.269	1.681	5240	8.564
41	1918.936	0.042	13246	9.491	89	1993.100	0.831	14000	9.547
42	1919.122	0.186	100000	11.513	90	1995.025	1.925	1500	7.313
43	1919.294	0.172	11000	9.306	91	1998.350	3.325	120000	11.695
44	1919.347	0.053	50000	10.820	92	1999.233	0.883	5002	8.518
45	1919.836	0.489	40000	10.597	93	1999.356	0.122	1172	7.066
46	1920.542	0.706	1000	6.908	94	2001.875	2.519	4002	8.295
47	1929.631	9.089	3200	8.071	95	2003.219	1.344	7173	8.878
48	1931.969	2.339	60000	11.002					

III (v) The assortative mating according to temper

In his classic *Natural Inheritance*, Galton (1889) gives a fascinating and entertaining analysis of 'temper', coupled with the general theme of inheritance, and courageously classifies husbands and wives as 'bad-tempered' or 'good-tempered'. We use these data in Story iii.8 to not merely test for independence of these character traits, but to provide confidence distributions and curves for relevant parameters. Are bad-tempered men better at finding good-tempered women than the good-tempered men are?

```
              wife:
              good  bad
husband: good  24    27
         bad   34    26
```

Galton, 1887: "We can hardly, too, help speculating uneasily upon the terms that our own relatives would select as most appropriate to our particular selves."

III (vi) Sims and Sveriges Riksbank

(xx perhaps, the Tore story on Sims. xx)

III (vii) Does winning make you live longer?

(xx emil story with RDD. xx)

III (viii) Minimum wages and employment

(xx emil story. xx)

III (ix) How many were killed in Guatemala?

(xx Lum et al. (2013) and Urdal. xx)

III (x) Volatility and estimation from noisy data

(xx emil story. xx)

IV (i) New Haven annual temperatures 1912–1971

(xx annual temperatures at New Haven, Connecticut, for the years 1912 to 1971, as follows. xx)

```
 9.94 11.28  9.67 10.61  9.67  8.83  9.89 10.50  9.61 11.06 10.44  9.78
 9.61 10.33  9.11 10.39 10.50 10.33 10.83 11.56 11.00 10.61  9.89 10.11
10.22 10.89 11.00 10.50  9.33 10.94 10.56 10.33 10.94 10.83 11.17 10.72
10.56 12.22 10.78 11.50 11.72 12.56 11.11 11.11 10.50 11.44 10.11 11.44
10.89 11.06 10.28 10.50 10.94 10.78 10.94 10.44 11.06 11.00 11.06 11.67
```

IV (ii) Skiing days at Bjørnholt 1897–2012

The table gives the number of skiing days at the location Bjørnholt in Nordmarka, a tramride and a skiing hour away downtown Oslo, from 1896 to 2022. A skiing day is defined as there being at least 25 cm snow on the ground. Note the 'hole' in the time series, with no counting of skiing days for the time window 1938 to 1954. Data from `rimfrost.no`. See Heger (2011), Cunen et al. (2018).

```
1896   25 1916  155 1936  151 1972  120 1992   39  2012   95
1897  147 1917  139 1937  108 1973  114 1993   70  2013  108
1898  169 1918  115 1954   37 1974  146 1994  119  2014   81
1899  135 1919  175 1955  135 1975  129 1995  123  2015   97
1900  135 1920  141 1956  111 1976  150 1996  116  2016   84
1901  137 1921  110 1957   82 1977  165 1997   69  2017   55
1902  131 1922  129 1958  125 1978  152 1998   39  2018  115
1903  147 1923  164 1959  143 1979  182 1999  121  2019  125
1904  185 1924  135 1960  141 1980  156 2000   57  2020   34
1905  127 1925   91 1961  141 1981  157 2001  150  2021  110
```

1906	116	1926	130	1962	152	1982	148	2002	129	2022	96
1907	165	1927	163	1963	150	1983	128	2003	113		
1908	170	1928	127	1964	92	1984	121	2004	108		
1909	176	1929	75	1965	150	1985	138	2005	37		
1910	182	1930	122	1966	162	1986	139	2006	120		
1911	149	1931	140	1967	142	1987	127	2007	58		
1912	174	1932	65	1968	125	1988	141	2008	131		
1913	125	1933	55	1969	153	1989	5	2009	101		
1914	128	1934	110	1970	191	1990	27	2010	144		
1915	190	1935	117	1971	178	1991	97	2011	117		

IV (iii) British coal-mining disasters, 1851 to 1962

[xx for chapter1: when data length is stretched, at which time point is it too much for the homogeneous Poisson assumption? in later chapters too. story us treated in Cunen et al. (2018), but then with a change-point perspective. xx]

1851	4	1871	5	1891	2	1911	0	1931	3	1951	1
1852	5	1872	3	1892	1	1912	1	1932	3	1952	0
1853	4	1873	1	1893	1	1913	1	1933	1	1953	0
1854	1	1874	4	1894	1	1914	1	1934	1	1954	0
1855	0	1875	4	1895	1	1915	0	1935	2	1955	0
1856	4	1876	1	1896	3	1916	1	1936	1	1956	0
1857	3	1877	5	1897	0	1917	0	1937	1	1957	1
1858	4	1878	5	1898	0	1918	1	1938	1	1958	0
1859	0	1879	3	1899	1	1919	0	1939	1	1959	0
1860	6	1880	4	1900	0	1920	0	1940	2	1960	1
1861	3	1881	2	1901	1	1921	0	1941	4	1961	0
1862	3	1882	5	1902	1	1922	2	1942	2	1962	1
1863	4	1883	2	1903	0	1923	1	1943	0		
1864	0	1884	2	1904	0	1924	0	1944	0		
1865	2	1885	3	1905	3	1925	0	1945	0		
1866	6	1886	4	1906	1	1926	0	1946	1		
1867	3	1887	2	1907	0	1927	1	1947	4		
1868	3	1888	1	1908	3	1928	1	1948	0		
1869	5	1889	3	1909	2	1929	0	1949	0		
1870	4	1890	2	1910	2	1930	2	1950	0		

IV (iv) Mammals and their bodies and brains

[xx 56 mammals, table gives average weight of body and average weight of brain, in kg. *you* are no. 25. how special are you? (log-body, log-brain) follow a binormal, more or less. for Story iv.4. xx]

1	0.480	0.0155	20	10.000	0.1150	39	60.000	0.0810
2	0.019	0.0003	21	3.300	0.0256	40	3.600	0.0210
3	600.000	0.4230	22	0.200	0.0050	41	0.320	0.0019
4	14.000	0.0700	23	85.000	0.3250	42	0.743	0.0200
5	14.800	0.0982	24	2.625	0.0123	43	0.075	0.0012
6	33.500	0.1150	25	62.000	1.3200	44	0.148	0.0012
7	0.728	0.0055	26	6654.000	5.7120	45	0.122	0.0030
8	0.420	0.0064	27	6.800	0.1790	46	0.920	0.0057
9	0.060	0.0010	28	0.120	0.0010	47	0.101	0.0040
10	1.000	0.0066	29	0.022	0.0004	48	0.048	0.0003
11	0.005	0.0001	30	0.010	0.0002	49	86.250	0.1800
12	3.500	0.0108	31	1.400	0.0125	50	4.500	0.0250

13	2.950	0.0123	32	2.500	0.0121	51	207.501	0.1690
14	1.700	0.0063	33	55.500	0.1750	52	0.900	0.0026
15	2547.000	4.6030	34	52.200	0.4400	53	0.104	0.0025
16	0.023	0.0003	35	100.000	0.1570	54	2.000	0.0175
17	521.000	0.6550	36	25.235	0.1800	55	3.380	0.0445
18	187.000	0.4190	37	0.550	0.0024	56	4.230	0.0504
19	0.770	0.0035	38	1.620	0.0114			

IV (v) Kola temperatures and The Hjort liver index time series 1859–2020

[xx The Hjort liver index. comes in two forms, the hli_{bulk} and $\text{hli}_{\text{perfish}}$. from Hermansen et al. (2016). one story on the five-parameter model with gamma marginals, Story iv.5, with ML. xx]

IV (vi) How many Clethrionomys glareoli?

(xx fill in basic reference. xx)

Bolt from heaven

Bolt (2013) is an interesting book, but the author restrains most of his attention to himself and his own achievements, so to find relevant data for our measure of surprise analyses for track and field events in Story v.1 we have needed to track down and organise our own files. Our analyses in particular utilise all sub-10.00 100 metre races from 2000 to 2008, and on the technical side involve Bartletting the deviance function derived from the extreme value distribution.

(xx Bolt did 9.72 in June 2008. how surprising ought we to have been? data: during the eight seasons 2000 to 2007, there were 20, 27, 29, 14, 38, 21, 29, 17 races clocked at 10.00-or-better, a total of 195 such. Story v.1. xx)

9.86	9.87	9.91	9.93	9.94	9.95	9.96	9.97	9.97	9.97	9.98	9.98
9.98	9.99	9.99	10.00	10.00	10.00	10.00	10.00	9.82	9.84	9.85	9.88
9.88	9.90	9.90	9.90	9.91	9.92	9.94	9.95	9.96	9.96	9.96	9.96
9.97	9.97	9.97	9.98	9.98	9.98	9.99	9.99	9.99	10.00	10.00	9.78
9.87	9.89	9.91	9.93	9.94	9.94	9.94	9.94	9.94	9.95	9.95	9.96
9.97	9.97	9.97	9.98	9.98	9.98	9.98	9.98	9.98	9.99	9.99	9.99
9.99	9.99	10.00	10.00	9.93	9.94	9.95	9.97	9.97	9.97	9.97	9.98
9.99	9.99	10.00	10.00	10.00	10.00	9.85	9.86	9.87	9.87	9.88	9.89
9.90	9.90	9.91	9.91	9.91	9.92	9.93	9.93	9.93	9.93	9.93	9.93
9.94	9.94	9.95	9.95	9.96	9.96	9.97	9.97	9.97	9.97	9.98	9.98
9.99	9.99	9.99	9.99	10.00	10.00	10.00	10.00	9.77	9.84	9.85	9.88
9.89	9.94	9.96	9.96	9.98	9.99	9.99	9.99	9.99	9.99	9.99	9.99
10.00	10.00	10.00	10.00	10.00	9.77	9.84	9.85	9.85	9.85	9.86	9.86
9.86	9.87	9.88	9.88	9.88	9.89	9.91	9.91	9.92	9.92	9.95	9.95
9.96	9.96	9.97	9.97	9.98	9.99	9.99	9.99	9.99	10.00	9.74	9.78
9.83	9.84	9.85	9.90	9.91	9.93	9.94	9.95	9.96	9.97	9.97	9.98
9.98	9.99	10.00									

Height, weight, BMI for Olympic speedskaters

(xx speedskaters taking part in Olympics, 1274 men, from 1952 to 2018, 907 women, from 1960 to 2018. we have height, weight, and hence BMI data for these, from files gathered by NLH and fellow speedskating history enthusiasts Jeroen Heijmans and Arild Gjerde. see Story v.11, with overall confidence curves for the 0.2, 0.5, 0.8 quantiles of the

BMI distributions; there is also a more complex story for the II-CC-FF story of medians, partly from Cunen and Hjort (2022). xx)

(xx the 1274 + 907 BMI scoes are available at the kioskvelter website. mean and standard deviation are 23.706, 1.445 for the males, and 22.050, 1.412 for the females. The data ranges are from 18.365 to 28.965 for the men and 17.647 to 26.675 for the women. Figure v.17 shows how clearly the bulk of women's BMI distribution is situated to the left of the men's. xx)

Inner and outer

[xx describe data, inner and outer starts for World Championships 1000 m speedskating. manually collected and curated by NLH over the years, via ISU protocols. analysed in Story v.6. xx]

From semifinals to finals

(xx describe data. type A A B B B A, skiers coming from Semifinal A or Semifinal B, their ranks in the final, for lots of events, Olympics 2022, 2018, 2014, 2010, World Championships 2021, 2019, 2017, 2015, 2013, 2011, and various FIS World Cup events. analysed in Story v.7. xx)

Golf putting

(xx fill in. Gelman and Nolan (2002), Schweder and Hjort (2016, Ch. 14). xx)

Who wins: Real Time Real Excitement Plots

(xx fill in. data I, that particular Nor-Den match; data II, some 75 matches and their scores, for correlated Poissons. xx)

Abelian envelopes

How many Abel commemorative envelopes were issued, in 1929? Prior to 2011, five such envelopes were known in the international community of stamp collectors, with so-called R numbers 280, 304, 308, 310, 328. We derive in Story ii.5 a confidence distribution for the number N of envelopes originally issued. (xx on the Bayesian version in Ch. 6. xx) Story ii.5 concerns reaching Bayesian inference for this unknown number of envelopes, based in these five R numbers. Story ii.5. are also about how these inerence summaries, the posterior for the Bayesian and the CD for the frequentist, can be updated, in the light of three more 1902 envelopes discovered in 2012, carrying R numbers 314, 334, 389. The data, along with the facsimile shown in Figure ii.9, are from colleagues Y. Reichelt and N.V. Johansen of the authors. (xx nils moves the facisimile to the story itself. xx)

Mixed effects for doughnuts

The table below provides data for eight different fats used for six consecutive working days in connection with doughnut mixing; specifically, $y_{i,j}$ given there displays the grams of fat absorbed for fat $i = 1, \ldots, r = 8$ across days $j = 1, \ldots, s = 6$. The data are from Scheffé (1959, p. 137); cf. also McCloskey (1943). (xx then point to where we have a little story about this, with mixed effects, relative influence of days, etc. xx)

	Mon	Tue	Wed	Thu	Fri	Sat	means
1	164	177	168	156	172	195	172.00
2	172	197	167	161	180	190	177.83

```
3      177    184    187    169    179    197     182.17
4      178    196    177    181    184    191     184.50
5      163    177    144    165    166    178     165.50
6      163    193    176    172    176    178     176.33
7      150    179    146    141    169    183     161.33
8      164    169    155    149    170    167     162.33
means 166.38 184.00 165.00 161.75 174.50 184.88  172.75
```

State-wise cigarette consumption and cancers

For each of 44 US states (actually, 43 states and the District of Columbia), the table below, dating from 1960, gives `cig`, number of cigarettes smoked (hundreds per capita); `blad`, deaths per 100k population from bladder cancer; `lung`, deaths per 100k population from lung cancer; `kid`, deaths per 100k population from kidney cancer; `leuk`, deaths per 100k population from leukemia. (xx used in Ch. 6. can also be used for regression stories in other chapters. the same x acts on four regressions. outliers? what happens with the verbatim, first line? xx)

```
    cig  blad  lung  kid  leuk
AL 18.20 2.90 17.05 1.59 6.15
AZ 25.82 3.52 19.80 2.75 6.61
AR 18.24 2.99 15.98 2.02 6.94
CA 28.60 4.46 22.07 2.66 7.06
CT 31.10 5.11 22.83 3.35 7.20
DE 33.60 4.78 24.55 3.36 6.45
DC 40.46 5.60 27.27 3.13 7.08
FL 28.27 4.46 23.57 2.41 6.07
ID 20.10 3.08 13.58 2.46 6.62
IL 27.91 4.75 22.80 2.95 7.27
IN 26.18 4.09 20.30 2.81 7.00
IO 22.12 4.23 16.59 2.90 7.69
KS 21.84 2.91 16.84 2.88 7.42
KY 23.44 2.86 17.71 2.13 6.41
LA 21.58 4.65 25.45 2.30 6.71
ME 28.92 4.79 20.94 3.22 6.24
MD 25.91 5.21 26.48 2.85 6.81
MA 26.92 4.69 22.04 3.03 6.89
MI 24.96 5.27 22.72 2.97 6.91
MN 22.06 3.72 14.20 3.54 8.28
MS 16.08 3.06 15.60 1.77 6.08
MO 27.56 4.04 20.98 2.55 6.82
MT 23.75 3.95 19.50 3.43 6.90
NB 23.32 3.72 16.70 2.92 7.80
NE 42.40 6.54 23.03 2.85 6.67
NJ 28.64 5.98 25.95 3.12 7.12
NM 21.16 2.90 14.59 2.52 5.95
NY 29.14 5.30 25.02 3.10 7.23
ND 19.96 2.89 12.12 3.62 6.99
OH 26.38 4.47 21.89 2.95 7.38
OK 23.44 2.93 19.45 2.45 7.46
PE 23.78 4.89 12.11 2.75 6.83
RI 29.18 4.99 23.68 2.84 6.35
SC 18.06 3.25 17.45 2.05 5.82
SD 20.94 3.64 14.11 3.11 8.15
TE 20.08 2.94 17.60 2.18 6.59
TX 22.57 3.21 20.74 2.69 7.02
UT 14.00 3.31 12.01 2.20 6.71
VT 25.89 4.63 21.22 3.17 6.56
```

```
WA 21.17 4.04 20.34 2.78 7.48
WI 21.25 5.14 20.55 2.34 6.73
WV 22.86 4.78 15.53 3.28 7.38
WY 28.04 3.20 15.92 2.66 5.78
AK 30.34 3.46 25.88 4.32 4.90
```

Mothers and babies

(xx describe the data, from Hosmer and Lemeshow (1999), 189 newborns, their weights, along with information about their mothers, including weight before pregnancy, age, an indicator for smoking, and also information about three ethnic groups. used for nonparametric confidence intervals for quantiles, and nils plans to have a fic. also something for comparing distribution of weight at birth for smoking and non-smoking mothers. xx)

Longevity returns to political office

[xx describe the data set of Barfort et al. (2020), that is analysed in Story iii.11. xx]

Onset of menarche

(xx clean and polish text, with pointer to exercise. xx) Data pertain to 3918 Warszawa girls and onset of menarche. The data have three columns, giving respectively (i) age x (or rather the midpoint in the appropriate age interval); (ii) the number m of girls in this age window; (iii) the number y among the of girls in this age window who have had their first menstruation. Thus y_j in age window j is seen as a realisation of a binom(m_j, p_j) distribution, with $p_j = p(x_j)$ the probability of onset having taken place at age x_j or earlier.

x	m	y	x	m	y	x	m	y
9.21	376	0	12.33	93	29	14.58	120	113
10.21	200	0	12.58	100	39	14.83	102	95
10.58	93	0	12.83	108	51	15.08	122	117
10.83	120	2	13.08	99	47	15.33	111	107
11.08	90	2	13.33	106	67	15.58	94	92
11.33	88	5	13.58	105	81	15.83	114	112
11.58	105	10	13.83	117	88	17.58	1049	1049
11.83	111	17	14.08	98	79			
12.08	100	16	14.33	97	90			

Decimals of π

For our Story vii.2 we have first gathered the first 6,537,216 digits of π and then via some coding found the first 223,157 values of V_{10}, the lengths of cycles required to have seen all decimals $0, 1, \ldots, 9$. These are gathered in our file `V10counts`.

Australian rowers

(xx to be edited and clarified. for Story vii.9. data from `www.statsci.org/data/oz/ais.txt`. Telford, R. D. and Cunningham, R. B. (1991) Sex, sport, and body-size dependency of hematology in highly trained athletes. Medicine and Science in Sports and Exercise, 23(7):788-794. gender is 11 for girls and 12 for boys. x is lean body mass, y is percent body fat. do we know or merely guess that no 16 and no 30 are coxswains? xx)

	gender	x	y	bmi
1	11	66.24	17.71	25.436
2	11	57.92	18.77	22.630
3	11	56.52	19.83	21.856
4	11	54.78	25.16	22.270
5	11	56.31	18.04	21.275
6	11	62.96	21.79	23.470
7	11	56.68	22.25	23.190
8	11	62.39	16.25	23.174
9	11	63.05	16.38	24.536
10	11	56.05	19.35	22.955
11	11	53.65	19.20	19.763
12	11	65.45	17.89	23.363
13	11	64.62	12.20	22.666
14	11	60.05	23.70	24.236
15	11	56.48	24.69	24.212
16	11	41.54	16.58	20.464
17	11	52.78	21.47	20.810
18	11	52.72	20.12	20.168
19	11	61.29	17.51	23.060
20	11	59.59	23.70	24.402
21	11	61.70	22.39	23.974
22	11	62.46	20.43	22.617
23	12	78.00	9.00	23.566
24	12	75.00	12.61	25.839
25	12	78.00	9.03	24.057
26	12	87.00	6.96	23.850
27	12	78.00	10.05	25.090
28	12	79.00	9.56	23.844
29	12	79.00	9.36	25.314
30	12	48.00	10.81	19.690
31	12	82.00	8.61	26.069
32	12	82.00	9.53	25.503
33	12	82.00	7.42	23.688
34	12	83.00	9.79	26.795
35	12	88.00	8.97	25.615
36	12	83.00	7.49	25.055
37	12	78.00	11.95	24.928

Other Stuff, and references to be used

(xx we mention and point to things here, not yet on board as of 12-March-2025, but probably to be factored in. we also point to references that we sooner or later will have on board. xx)

Aalen (1992), Aalen and Gjessing (2004), Cunen et al. (2020c), Frigessi and Hjort (2002), Gjessing et al. (2003), Hjort (1990b), Hjort (1990a), Hjort and Stoltenberg (2021), ... when we mention Breiman (2001), we also point to the Special Issue of *Observational Studies* (vol. 7, 2021), e.g. Bickel's comments, in `muse.jhu.edu/issue/45147`. also Gelman et al. (2022).

References

Aalen, O. O. (1992). Modelling heterogeneity in survival analysis by the compound Poisson distribution. *Annals of Applied Probability*, 2:951–972.

Aalen, O. O., Borgan, Ø., and Gjessing, H. K. (2008). *Survival and Event History Analysis. A process point of view*. Springer, New York.

Aalen, O. O. and Gjessing, H. K. (2004). Survival models based on the Ornstein–Uhlenbeck process. *Lifetime Data Analysis*, 10:407–423.

Aït-Sahalia, Y. and Jacod, J. (2014). *High-Frequency Financial Econometrics*. Princeton University Press, Princeton.

Aldous, D. J. and Eagleson, G. K. (1978). On mixing and stability of limit theorems. *The Annals of Probability*, pages 325–331.

Andersen, P. K., Borgan, Ø., Gill, R. D., and Keiding, N. (1993). *Statistical Models Based on Counting Processes*. Springer-Verlag, Berlin.

Angrist, J. D. and Pischke, J.-S. (2010). The credibility revolution in empirical economics: How better research design is taking the con out of econometrics. *Journal of Economic Perspectives*, 24:3–30.

Ashworth, S., Berry, C. R., and de Mesquita, E. B. (2021). *Theory and Credibility: Integrating Theoretical and Empirical Social Science*. Princeton University Press, Princeton.

Aursnes, I., Tvete, I. F., Gåsemyr, J., and Natvig, B. (2005). Suicide attempts in clinical trials with paroxetine randomised against placebo. *BMC Medicine*, xx:1–5.

Aursnes, I., Tvete, I. F., Gåsemyr, J., and Natvig, B. (2006). Even more suicide attempts in clinical trials with paroxetine randomised against placebo. *BMC Psychiatry*, xx:1–3.

Ball, P. (1999). *Making the Case: Investigating Large Scale Human Rights Violations Using Information Systems and Data Analysis*. American Academy for the Advancement of Science, Washington.

Barber, R. F., Candes, E. J., Ramdas, A., and Tibshirani, R. J. (2022). Conformal prediction beyond exchangeability. *arXiv preprint arXiv:2202.13415*.

Barfort, S., Klemmensen, R., and Larsen, E. G. (2020). Longevity returns to political office. *Political Science Research and Methods*, 9:658–664.

Bartolucci, F. and Lupparelli, M. (2008). Focused Information Criterion for capture-recapture models for closed populations. *Scandinavian Journal of Statistis*, 9:658–664.

Basu, A., Harris, I. R., Hjort, N. L., and Jones, M. C. (1998). Robust and efficient estimation by minimising a densithy power divergence. *Biometrika*, 85:549–559.

Basu, A., Shioya, H., and Park, C. (2011). *Statistical Inference: The Minimum Distance Approach*. Chapman & Hall/CRC Press, London.

Basu, D. (1955). On statistics independent of a complete sufficient statistic. *Sankhya*, 15:377–180.

Billingsley, P. (1961). *Statistical Inference for Markov Processes*. Chicago University Press, Chicago.

Billingsley, P. (1968). *Convergence of Probability Measures*. Wiley, New York.

Billingsley, P. (1995). *Probability and Measure. Third Edition*. Wiley, New York.

Blower, J. G., Cook, L. M., and Bishop, J. A. (1981). *Estimating the Size of Animal Populations*. Allen & Unwin, Kondon.

Boitsov, V. D., Karsakov, A. L., and Trofimov, A. G. (2012). Atlantic water temperature and climate in the barents sea, 2000–2009. *ICES Journal of Marine Science*, 69:833–840.

Bolt, U. (2013). *Faster Than Lightning: My Autobiography*. HarperSport, London.

Borgan, Ø., Fiaccone, R. L., Henderson, R., and Barreto, M. L. (2007). Dynamic analysis of recurrent event data with missing observations, with application to infant diarrhoea in brazil. *Scandinavian Journal of Statistics*, 34:53–69.

Borgan, Ø. and Keilman, N. (2019). Do Japanese and Italian women live longer than women in Scandinavia? *European Journal of Population*, 35:87–99.

Bowman, A. (1984). An alternative method of cross-validation for the smoothing of density estimates. *Biometrika*, 71:353–360.

Breiman, L. (2001). Statistical modeling: The two cultures [with comments and a rejoinder by the author]. *Statistical Science*, 16:199–231.

Brunborg, H., Lyngstad, T. H., and Urdal, H. (2003). Accounting for genocide: How many were killed in Srebrenica? *European Journal of Population*, 19:229–248.

Candès, E. J., Lei, L., and Ren, Z. (2021). Conformalized survival analysis. *arXiv preprint arXiv:2103.09763*.

Card, D. and Krueger, A. B. (1994). Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *The American Economic Review*, 84:772–793.

Casella, G. and George, E. I. (1992). Explaining the Gibbs Sampler. *American Statistician*, 46:167–174.

Chernozhukov, V., Wüthrich, K., and Zhu, Y. (2021). Distributional conformal prediction. *Proceedings of the National Academy of Sciences*, 118:e2107794118.

Claeskens, G., Cunen, C., and Hjort, N. L. (2019). Model selection via Focused Information Criteria for complex data in ecology and evolution. *Frontiers in Ecology and Evolution*, 7:415–428.

Claeskens, G. and Hjort, N. L. (2003). The focused information criterion [with discussion and a rejoinder]. *Journal of the American Statistical Association*, 98:900–916.

Claeskens, G. and Hjort, N. L. (2008a). Minimizing average risk in regression. *Econometric Theory*, 24:493–527.

Claeskens, G. and Hjort, N. L. (2008b). *Model Selection and Model Averaging*. Cambridge University Press, Cambridge.

Clauset, A. (2018). Trends and fluctuations in the severity of interstate wars. *Science Advances*, 4:1–9.

Clauset, A. (2020). On the frequency and severity of interstate wars. In Gleditsch, N. P., editor, *Lewis Fry Richardson: His Intellectual Legacy and Influence in the Social Sciences*, pages 113–128. Springer, Berlin.

Clevenson, M. L. and Zidek, J. V. (1975). Simultaneous estimation of the means of independent Poisson laws. *Journal of the American Statistical Association*, 70:698–705.

Cox, D. R. (1958). Some problems with statistical inference. *The Annals of Mathematical Statistics*, 29:357–372.

Cox, D. R. (1972). Regression models and life-tables [with discussion]. *Journal of the Royal Statistical Society: Series B*, 34:187–202.

Cox, D. R. and Brandwood, L. (1959). On a discriminatory problem connected with the worlks of Plato. *Journal of the Royal Statistical Society Series B*, 21:195–200.

Cox, D. R. and Miller, H. D. (1965). *The Theory of Stochastic Processes*. Chapman & Hall, London.

Cramér, H. (1946). *Mathematical Methods of Statistics*. Princeton University Press, Princeton.

Cramér, H. (1976). Half a century with probability theory: some personal reflections. *Annals of Probability Theory*, 4:509–546.

Cunen, C. (2015). Mortality and Nobility in the Wars of the Roses and Game of Thrones. *FocuStat Blog, University of Oslo*, iv.

Cunen, C., Hermansen, G. H., and Hjort, N. L. (2018). Confidence distributions for change points and regime shifts. *Journal of Statistical Planning and Inference*, 195:14–34.

Cunen, C. and Hjort, N. L. (2015). Optimal inference via confidence distributions for two-by-two tables modelled as Poisson pairs: fixed and random effects. In Nair, V., editor, *Proceedings of the 60th World Statistics Congress, ISI Rio*, pages xx–xx. Springer, Rio.

Cunen, C. and Hjort, N. L. (2022). Combining information from diverse sources: the II-CC-FF paradigm. *Scandinavian Journal of Statistics*, 49:625–656.

Cunen, C. and Hjort, N. L. (2024). Survival and event history models and methods via Gamma processes. Technical report, University of Oslo. Technical report.

Cunen, C., Hjort, N. L., and Nygård, H. M. (2020a). Statistical sightings of better angels. *Journal of Peace Research*, 57:221–234.

Cunen, C., Hjort, N. L., and Schweder, T. (2020b). Confidence in confidence distributions! *Proceedings of the Royal Society, A*, 476:1–5.

Cunen, C., Walløe, L., and Hjort, N. L. (2020c). Focused model selection for linear mixed models, with an application to whale ecology. *Annals of Applied Statistics*, 14:872–904.

Dæhlen, I., Hjort, N. L., and Hobæk Haff, I. (2024). Accurate bias estimation with applications to focused model selection [to appear]. *Scandinavian Journal of Statistics*, 51:724–759.

Dagsvik, J. K., Fortuna, M., and Moen, S. H. (2020). How does temperature vary over time?: Evidence on the stationary and fractal nature of temperature fluctuations. *Journal of the Royal Statistical Society A*, pages 883–908.

De Blasi, P. and Hjort, N. L. (2007). Bayesian survival analysis in proportional hazard models with logistic relative risk. *Scandinavian Journal of Statistics*, 34:229–257.

DeGroot, M. H. (1970). *Optimal Statistical Decisions*. John Wiley & Sons, Hoboken, N.J.

Dudley, R. M. (2002). *Real Analysis and Probability*. Cambridge University Press, Cambridge.

Eddington, A. S. (1914). *Stellar Movements and the Structure of the Universe*. Macmillan, London.

Efron, B. (2023). *Exponential Families in Theory and Practice*. Cambridge University Press, Cambridge.

Efron, B. and Morris, C. (1977). Stein's paradox in statistics. *Scientific American*, 236:119–127.

Einmahl, J. H. J. and Smeets, S. G. W. R. (2011). Ultimate 100 m world records through extreme-value theory. *Statistica Neerlandica*, 65:32–42.

Embrechts, P., Klüppelberg, C., and Mikosch, T. (1997). *Modelling Extremal Events for Insurance and Finance*. Springer, London.

Fagerland, M., Lydersen, S., and Laake, P. (2017). *Statistical Analaysis of Contingency Tables*. Chapman and Hall/CRC, New YOrk.

Ferguson, T. S. (1996). *A Course in Large Sample Theory*. Chapman & Hall, London.

Fisher, R. A. (1920). A mathematical examination of the methods of determining the accuracy of an observation by the mean error, and by the mean square error. *Monthly Notices of the Royal Astronomical Society*, 80:758–770.

Fisher, R. A. (1930). Inverse probability. *Proceedings of the Cambridge Philosophical Society*, 26:528–535.

Franklin, B. (1793). *The Autobiography of Benjamin Franklin*. Dover, New York. Reprinted from Dover, New York, 1996.

Friesinger, A. (2004). *Mein Leben, mein Sport, meine besten Fitness-Tipps*. Goldmann, Berlin.

Frigessi, A. and Hjort, N. L. (2002). Statistical methods for discontiniuous phenomena. *Journal of Nonparametric Statistics*, 14:1–5.

Galton, F. (1889). *Natural Inheritance*. Macmillan, London.

Geißler, A. (1889). Beiträge zur Frage des Geschlechts verhältnisses der Geborenen. *Zeitschrift des königlichen sächsischen statistischen Bureaus*, 35:1–24.

Gelman, A., Hill, J., and Vehtari, A. (2022). *Regression and Other Stories*. Cambridge University Press, Cambridge.

Gelman, A. and Nolan, D. (2002). A probability model for golf putting. *Teaching Statistics*, 24:93–95.

Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741.

Ghosh, M. (2002). Basu's theorem with applications: a personalistic review. *Sankhya*, 35:721–741. Special issue in memory of D. Basu.

Gilovich, T., Vallone, R., and Tversky, A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, 17:295–314.

Gjessing, H. K., Aalen, O. O., and Hjort, N. L. (2003). Frailty models based on Lévy processes. *Advances in Applied Probability*, 35:532–550.

Glad, I. K., Hjort, N. L., and Ushakov, N. G. (2003). Correction of density estimators that are not densities. *Scandinavian Journal of Statistics*, 30:415–427.

Gleditsch, N. P. (2020). *Lewis Fry Richardson: His Intellectual Legacy and Influence in the Social Sciences (edited book)*. Springer, Berlin.

Goudie, I. B. J. and Goudie, M. (2007). Who captures the marks for the Petersen estimator? *Journal of the Royal Statistical Society, Series A*, 170:825–839.

Gran, J. M. and Stensrud, M. J. (2022). Hva er forventet levealder? *Tidsskrift for Den norske legeforening*, page 245.

Grønneberg, S. and Hjort, N. L. (2012). On the errors committed by sequences of estimator functionals. *Mathematical Methods of Statistics*, 20:327–346.

Haavelmo, T. (1943). The statistical implications of a system of simultaneous equations. *Econometrics*, 11:1–12.

Hall, P. (1927). The distribution of means for samples of size N drawn from a population in which the variate takes values between 0 and 1, all such values being equally probable. *Biometrika*, 19:240–245.

Hall, P. G. (1983). Large-sample optimality of least squares cross-validation in density estimation. *Annals of Statistics*, 11:1156–1174.

Hanche-Olsen, H. and Holden, H. (2010). The Kolmogorov-Riesz compactness theorem. *Expositiones Mathematicase*, 28:385–394.

Hastie, T., Tibshirani, R., and Friedman, J. H. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Second Edition*. Springer, New York.

Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 9:97–109.

Haug, K. K. (2019). Focused model selection for Markov chain models, with an application to armed conflict data. Technical report, University of Oslo. Master Thesis.

Heger, A. (2011). Jeg og jordkloden. *Dagsavisen*, Dec. 16.

Hermansen, G. H., Hjort, N. L., and Kjesbu, O. S. (2016). Modern statistical methods applied on extensive historic data: Hjort liver quality time series 1859-2012 and associated influential factors. *Canadian Journal of Fisheries and Aquatic Sciences*, 73:273–295.

Hjort, J. (1914). *Fluctuations in the Great Fisheries of Northern Europe, Viewed in the Light of Biological Research*. Conseil Permanent International Pour l'Exploration de la Mer, Copenhagen.

Hjort, N. L. (1986a). Bayes estimators and asymptotic efficiency in parametric counting process models. *Scandinaavian Journal of Statistics*, 13:63–85.

Hjort, N. L. (1986b). *Notes on the Theory of Statistical Symbol Recognition*. Norwegian Computing Centre, Oslo.

Hjort, N. L. (1988a). The eccentric part of the non-central chi square. *The American Statistician*, 42:130–132.

Hjort, N. L. (1988b). On large-sample multiple comparison methods. *Scandinaavian Journal of Statistics*, 15:259–271.

Hjort, N. L. (1990a). Goodness of fit tests for life history data based on cumulative hazard rates. *Annals of Statistics*, 18:1221–1258.

Hjort, N. L. (1990b). Nonparametric Bayes estimators based on Beta processes in models for life history data. *Annals of Statistics*, 18:1259–1294.

Hjort, N. L. (1992). On inference in parametric survival data models. *International Statisical Review*, xx:355–387.

Hjort, N. L. (1994). The exact amount of t-ness that the normal model can tolerate. *Journal of the American Statistical Association*, 89:665–675.

Hjort, N. L. (2007). And quiet does not flow the Don: Statistical analysis of a quarrel between Nobel laureates. In Østreng, W., editor, *Concilience*, pages 134–140. Centre for Advanced Research, Oslo.

Hjort, N. L. (2008). Discussion of P.L. Davies' article 'Approximating data'. *Journal of the Korean Statistical Society*, 37:221–225.

Hjort, N. L. (2014). Discussion of Efron's article 'Estimation and accuracy after model selection'. *Journal of the American Statistical Association*, 110:1017–1020.

Hjort, N. L. (2017a). Cooling of Newborns and the Difference Between 0.244 and 0.278. *FocuStat Blog, University of Oslo*, xv.

Hjort, N. L. (2017b). The Semifinals Factor for Skiing Fast in the Finals. *FocuStat Blog, University of Oslo*, xv.

Hjort, N. L. (2018a). Overdispersed Children. *FocuStat Blog, University of Oslo*, xxi.

Hjort, N. L. (2018b). Towards a More Peaceful World [insert '!' or '?' here]. *FocuStat Blog, University of Oslo*, xvii.

Hjort, N. L. (2019a). The Magic Square of 33. *FocuStat Blog, University of Oslo*, xxi.

Hjort, N. L. (2019b). Sudoku Solving by Probability Models and Markov Chains. *FocuStat Blog, University of Oslo*, xxi.

Hjort, N. L. and Claeskens, G. (2003a). Frequentist model averaging [with discussion and a rejoinder]. *Journal of the American Statistical Association*, 98:879–899.

Hjort, N. L. and Claeskens, G. (2003b). Rejoinder to the discussion of the Hjort and Claeskens and Claeskens and Hjort papers. *Journal of the American Statistical Association*, 98:917–925.

Hjort, N. L. and Fenstad, G. (1992). On the last time and the number of times an estimator is more than ε from its target value. *The Annals of Statistics*, 20:469–489.

Hjort, N. L. and Glad, I. K. (1995). Nonparametric density estimation with a parametric start. *The Annals of Statistics*, 23:882–904.

Hjort, N. L. and Jones, M. C. (1996). Locally parametric nonparametric density estimation. *The Annals of Statistics*, 24:1619–1647.

Hjort, N. L. and Koning, A. J. (2002). Tests for constancy of model parameters over time. *Journal of Nonparametric Statistics*, 14:113–132.

Hjort, N. L. and Lumley, T. (1993). Normalised local hazard plots. Technical report, Department of Statistics, University of Oxford, Oxford.

Hjort, N. L., McKeague, I. W., and Van Keilegom, I. (2009). Extending the scope of empirical likelihood. *Annals of Statistics*, 37:1079–1111.

Hjort, N. L., McKeague, I. W., and Van Keilegom, I. (2018). Hybrid combinations of parametric and empirical likelihoods. *Statistica Sinica*, 27:2389–2407.

Hjort, N. L. and Petrone, S. (2007). Nonparametric quantile inference using Dirichlet processes. In Nair, V., editor, *Advances in Statistical Modeling and Inference: Essays in Honor of Kjell Doksum*, pages 463–492. World Scientific, New Jersey.

Hjort, N. L. and Pollard, D. B. (1993). Asymptotics for minimisers of convex processes. Technical report, Department of Mathematics, University of Oslo.

Hjort, N. L. and Schweder, T. (2018). Confidence distributions and related themes: introduction to the special issue. *Journal of Statistical Planning and Inference*, 195:1–13.

Hjort, N. L. and Stoltenberg, E. A. (2021). The partly parametric and partly nonparametric additive risk model. *Lifetime Data Analysis*, 27:1–31.

Hjort, N. L. and Varin, C. (2008). ML, PL, QL in Markov chain models. *Scandinavian Journal of Statistics*, 35:64–82.

Hjort, N. L. and Walker, S. G. (2009). Quantile pyramids for Bayesian nonparametrics. *Annals of Statistics*, 37:105–131.

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association*, 81:945–960.

Holum, D. (1984). *The Complete Handbook of Speed Skating*. High Peaks Cyclery, Lake Pacid.

Hosmer, D. W. and Lemeshow, S. (1999). *Applied Logistic Regression*. Wiley, New York.

Hveberg, K. (2019). *Lene din ensomhet langsomt mot min*. Aschehoug, Oslo.

Imbens, G. W. and Rubin, D. B. (2015). *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press, Cambridge.

Inlow, M. (2010). A moment generating function proof of the Lindeberg–Lévy central limit theorem. *American Statistician*, 64:228–230.

Irwin, J. O. (1927). On the frequency distribution of the means of samples from a population having any law of frequency with finite moments, with special reference to Pearson's type II. *Biometrika*, 19:225–239.

Jacod, J. and Mémin, J. (1981). Sur un type de convergence intermédiaire entre la convergence en loi et la convergence en probabilité. In *Séminaire de Probabilités (Strasbourg), tome 15*, pages 529–546. Springer.

Jacod, J. and Protter, P. (2004). *Probability Essentials. Second Edition*. Springer, Berlin.

Jacod, J. and Shiryaev, A. (2013). *Limit Theorems for Stochastic Processes. Second Edition*. Springer, Berlin.

James, G., Witten, D., Hastie, T., and Tibshirani, R. (2021). *An Introduction to Statistical Learning with Applications in R. Second Edition*. Springer, New York.

James, W. and Stein, C. (1961). Estimation with quadratic loss. *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, pages 361–379.

Jamtveit, B., Jacobsen, A. U., and Wyller, T. B. (2018). Utvikling i andel administrativt personale i norske helseforetak. *Samfunnsøkonomen*, 6:17–21.

Jamtveit, B., Jettestuen, E., and Mathiesen, J. (2009). Scaling properties of European research units. *Proceedings of the National Academy of Sciences*, 106:13160–13163.

Jansen, D. (1994). *Full Circle*. Villard Books, New York.

Jones, M. C. (1991). The roles of ISE and MISE in density estimation. *Statistics and Probability Letters*, 12:51–56.

Jones, M. C., Hjort, N. L., Harris, I. R., and Basu, A. (2001). A comparison of related density-based minimum divergence estimators. *Biometrika*, 88:865–873.

Jullum, M. and Hjort, N. L. (2017). Parametric or nonparametric: The FIC approach. *Statistica Sinica*, 27:951–981.

Jullum, M. and Hjort, N. L. (2019). What price semiparametric Cox regression? *Lifetime Data Analysis*, 25:406–438.

Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York.

Kahneman, D., Sibony, O., and Sunstein, C. R. (2020). *Noise: A Flaw in Human Judgment*. William Collins, London.

Kallenberg, O. (2002). *Foundations of Modern Probability. Second Edition*. Springer, Berlin.

Kjesbu, O. S., Opdal, A. F., Korsbrekke, K., Devine, J. A., and Skjæraasen, J. E. (2014). Making use of Johan Hjort's 'unknown' legacy: reconstruction of a 150-year coastal time-series on northeast Arctic cod (Gadus morhua) liver data reveals long-term trends in energy allocation patterns. *ICES Journal of Marine Science*, 71:2053–2063.

Kjetsaa, G., Gustavson, S., Beckman, B., and Gil, S. (1984). *The Authorship of The Quiet Don [also published in Russian]*. Solum/Humanities Press, Oslo.

Klein, R., Knudtson, M. D., Lee, K. E., Gangnon, R., and Klein, B. E. (2008). The Wisconsin epidemiologic study of diabetic retinopathy: XXII the twenty-five-year progression of retinopathy in persons with type 1 diabetes. *Ophtalmology*, 115:1859–1868.

Klotz, J. (1972). Markov chain clustering of births by year. *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability Theory*, 4:173–185.

Klotz, J. (1973). Statistical inference in Bernoulli trials with dependence. *Annals of Statistics*, 1:373–379.

Koehler, J. J. and Conley, C. A. (2003). The "hot hand" myth in professional basketball. *Journal of Sport and Exercise Psychology*, 25:253–259.

Kolmogorov, A. N. (1933a). *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Julius Springer, Berlin.

Kolmogorov, A. N. (1933b). Sulla determinazione empirica di una legge di distribuzione. *Giorn Ist Ital Attuar*, 4:83–91.

Kosorok, M. R. (2008). *Introduction to empirical processes and semiparametric inference*. Springer, New York.

Kusolitsch, N. (2010). Why the theorem of Scheffé should be rather called a theorem of Riesz. *Periodica Mathematica Hungarica*, 61:225–229.

Laptook, A. e. a. (2017). Effect of therapeutic hypothermia initiated after 6 hours of age on death and disability among newborns with hypoxic-ischemic encephalopathy: A randomized clinical trial. *Journal of the American Medical Association*, 318:1550–1560.

Larkey, P. D., Smith, R. A., and Kadane, J. B. (1989). It's okay to believe in the "hot hand". *Chance*, 2:22–30.

Le May Doan, C. (2002). *Going For Gold*. McClelland & Stewart Publisher, Toronto.

LeCam, L. (1986). The Central Limit Theorem around 1935. *Statistical Science*, 1:78–91.

Lehmann, E. L. (1950). *Notes on the Theory of Estimation*. Berkeley University Press, Berkeley. Notes recorded by Colin Blyth.

Lehmann, E. L. (1975). *Nonparametrics: Statistical Methods Based on Ranks*. Holden-Day, San Francisco.

Lei, J., G'Sell, M., Rinaldo, A., Tibshirani, R. J., and Wasserman, L. (2018). Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113:1094–1111.

Leike, A. (2001). Demonstration of the exponential decay law using beer froth. *European Journal of Physics*, 23:1–21.

Lessing, D. (1997). *Walking in the Shade: Volume Two of My Autobiography, 1949 to 1962*. xx, xx.

Lindeberg, J. W. (1922). Eine neue Herleitung des Exponentialgesetzes in der Wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift*, 15:211–225.

Lindqvist, B. H. (1978). A note on Bernoulli trials with dependence. *Scandinavian Journal of Statistics*, 5:205–208.

Linero, A. and Rosalsky, A. (2013). On the Toeplitz lemma, convergence in probability, and mean convergence. *Stochastic Analysis and Applications*, 31:684–694.

Loader, C. (1996). Local likelihood density estimation. *Annals of Statistics*, 67:1602–1618.

Lum, K., Price, M. E., and Banks, D. (2013). Applications of multiple systems estimation in human rights research. *American Statistician*, 24:191–200.

Markov, A. A. (1906). Распространение закона больших чисел на величины, зависящие друг от друга [Extending the law of large numbers for variables that are dependent of each other]. Известия Физико-математического общества при Казанском университете (2-я серия), 15:124–156.

Markov, A. A. (1913). Пример статистического исследования над текстом "Евгения Онегина", иллюстрирующий связь испытаний в цепь [Example of a statistical investigation illustrating the transitions in the chain for the 'Evgenii Onegin' text]. Известия Академии Наук, Санкт-Петербург (6-я серия), 7:153–162.

Marron, S. and Wand, M. P. (1992). Exact mean integrated squared error. *Annals of Statistics*, 20:712–736.

McCloskey, R. (1943). *Homer Price*. Scholastic Inc., New York.

McCullagh, P. (2002). What is a statistical model? [with discussion]. *Annals of Statistics*, 30:1225–1310.

Miller, J. B. and Sanjurjo, A. (2018). Surprised by the hot hand fallacy? A truth in the law of small numbers. *Econometrica*, 86:2019–2047.

Miller, J. B. and Sanjurjo, A. (2021). Is it a fallacy to believe in the hot hand in the NBA three-point contest? *European Economic Review*, 138:103771.

Mykland, P. A. and Zhang, L. (2012). The econometrics of high frequency data. In Kessler, M., Lindner, A., and Sørensen, M., editors, *Statistical Methods for Stochastic Differential Equations*, pages 109–190. CRC Press.

Mykland, P. A., Zhang, L., and Chen, D. (2019). The algebra of two scales estimation, and the S-TSRV: High frequency estimation that is robust to sampling times. *Journal of Econometrics*, 208:101–119.

Neyman, J. and Pearson, E. (1933). On the problem of the most efficient statistical hypotheses. *Philosophical Transactions of the Royal Society of London*, 68:289–337.

Normand, S.-L. T. (1999). Tutorial in biostatistics: Meta-analysis: formulating, evaluating, combining, and reporting. *Statistics in Medicine*, 18:321–359.

Pearson, K. (1900). On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *Philosophical Magazine Series*, 5(302):157–175.

Pearson, K. (1902). On the change in expectation of life in man during a period of circa 2000 years. *Biometrika*, 1:261–264.

Petersen, C. G. J. (1896). The yearly immigration of young plaice into the Limfjord from the German Sea. *Report of the Danish Biological Station*, 6:5–84.

Peterson, A. V. (1975). Nonparametric estimation in the competing risks problem. Technical report, Department of Statistics, Stanford University.

Pinker, S. (2011). *The Better Angels of Our Nature: Why Violence Has Declined*. Viking Books, Toronto.

Pollard, D. (1984). *Convergence of Stochastic Processes*. Springer, New York.

Price, R. M. and Bonett, D. G. (2001). Estimating the variance of the sample median. *Journal of Statistical Computation and Simulation*, 68:xx–xx.

Price, R. M. and Bonett, D. G. (2002). Distribution-free confidence intervals for difference and ratio of medians. *Journal of Statistical Computation and Simulation*, 72:xx–xx.

Rao, C. R. (1945). Information and the accuracy attainable in the estimation of statistical parameters. *Bulletins of the Calcutta Mathematical Society*, pages 81–91.

Reeves, R. V. (2022a). *Of Boys and Men: Why the Modern Male is Struggling, Why it Matters, and What to Do About It*. Brookings Institution Press, Washington, D.C.

Reeves, R. V. (2022b). Redshirt the boys. *The Atlantic*, October.

Romano, J. P. and Siegel, A. F. (1986). *Counterexamples in Probability and Statistics*. Wadsworth & Brooks/Cole, Belmont.

Rosenbaum, P. R. and Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70:41–55.

Royden, H. L. and Fitzpatrick, P. M. (2010). *Real Analysis [4th ed.]*. Pearson Education Asia, Beijin.

Rudemo, M. (1982). Empirical choice of histograms and kernel density estimation. *Scandinavian Journal of Statistics*, 9:65–78.

Rydén, J. (2020). On features of fugue subjects: A comparison of J.S. Bach and later composers. *Journal of Mathematics and Music*, pages 1–20.

Saleh, J. H. (2019). Statistical reliability analysis for a most dangerous occupation: Roman emperor. *Palgrave Communication*, 5:1–7.

Sanathanan, L. (1972). Estimating the size of a multinomial population. *Annals of Mathematical Statistics*, 43:142–1542.

Scheffé, H. (1947). A useful convergence theorem for probability distributions. *Annals of Mathematical Statistics*, 18:434–438.

Scheffé, H. (1959). *The Analysis of Variance*. Wiley, New York.

Schervish, M. J. (1995). *Theory of Statistics*. Springer, New York.

Schömig, A., Mehili, J., de Waha, A., Seyfarth, M., Pahce, J., and Kastrati, A. (2008). A meta-analysis of 17 randomized trials of a percutaneous coronary intervention-based strategy in patients with stable coronary artery disease. *Journal of the American College of Cardiology*, 52:894–904.

Schweder, T. (1980). Scandinavian statistics, some early lines of development. *Scandinavian Journal of Statistics*, 7:113–129.

Schweder, T. (1999). Early statistics in the Nordic countries – when did the Scandinavians slip behind the British? *Bulletin of the International Statistical Institute*, 58:1–4.

Schweder, T. and Hjort, N. L. (2016). *Confidence, Likelihood, Probability: Statistical Inference with Confidence Distributions*. Cambridge University Press, Cambridge.

Scott, D. W. (1992). *Multivariate Density Estimation: Theory, Practice, and Visualization*. Wiley, New York.

Serfling, R. J. (1980). *Approximation Theorems of Mathematical Statistics*. Wiley, London.

Shafer, G. and Vovk, V. (2008). A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9.

Shao, J. (1991). Second-order differentiability and jackknife. *Statistica Sinica*, 1:185–202.

Shiryaev, A. N. (1996). *Probability. Second edition*. Springer, Berlin.

Shumway, R. H. and Stoffer, D. S. (2016). *Time Seires Analysis and Its Applications [4th ed.]*. Springer, Heidelberg.

Silver, N. (2012). *The Signal and the Noise: Why so Many Predictions Fail, but Some Don't*. Penguin.

Silverman, B. W. (1986). *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, London.

Simpson, R. J. S. and Pearson, K. (1904). Report on certain enteric fever inoculation statistics. *British Medical Journal*, 3:1243–1246.

Sims, C. A. (2012a). Appendix: inference for the Haavelmo model. Technical report, Puplic Policy & Finance, Princeton University, Princeton, NJ.

Sims, C. A. (2012b). Statistical modeling of monetary policy and its effects [Sveriges Riksbank Prize in Memory of Alfred Nobel lecture]. *American Economic Review*, xx:1–22.

Singh, K., Xie, M., and Strawderman, W. E. (2005). Combining information from independent sources through confidence distributions. *Annals of Statistics*, 33:159–183.

Slud, E. (1989). Clipped Gaussian processes are never M-step Markov. *Journal of Multivariate Analysis*, 29:1–14.

Smith, T. D. (1994). *Scaling Fisheries: The Science of Measuring the Effects of Fishing 1855–1955*. Cambridge University Press, Cambridge.

Spiegelberg, W. (1901). *Aegyptische und Griechische Eigennamen aus Mumientiketten der Römischen Kaiserzeit*. Greek Inscriptions, Cairo.

Stein, C. (1956). Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, pages 197–206.

Stigler, S. M. (1973). Studies in the history of probability and statistics. xxxii: Laplace, Fisher and the discovery of the concept of sufficiency. *Biometrika*, 60:439–445.

Stigler, S. M. (1977). Do robust estimators work with real data? *Annals of Statistics*, 27:1055–1098.

Stigler, S. M. (1983). Who discovered Bayes's Theorem? *American Statistician*, 37:290–296.

Stigler, S. M. (1990). The 1988 Neyman memorial lecture: a Galtonian perspective on shrinkage estimators. *Statistical Science*, 5:147–155.

Stigler, S. M. (2006). How Ronald Fisher became a mathematical statistician. *Mathematics and Social Sciences*, 44:23–30.

Stoltenberg, E. A. (2019). An MGF proof of the Lindeberg theorem. Technical report, Department of Mathematics, University of Oslo.

Stoltenberg, E. A. and Hjort, N. L. (2020). Multivariate estimation of Poisson parameters. *Journal of Multivariate Analysis*, 175:xx–xx.

Stoltenberg, E. A. and Hjort, N. L. (2021). Models and inference for on-off data via clipped Ornsten–Uhlenbeck processes. *Scandinavian Journal of Statistics*, 48:908–929.

Stout, W. F. (1974). *Almost Sure Convergence*. Academic Press, New York.

Student (1908). The probable error of a mean. *Biometrika*, 6:1–25.

Swensen, A. R. (1983). A note on convergence of distributions of conditional moments. *Scandinavian Journal of Statistics*, 10:41–44.

Tibshirani, R. J., Foygel Barber, R., Candes, E., and Ramdas, A. (2019). Conformal prediction under covariate shift. *Advances in Neural Information Processing Systems*, 32.

Tversky, A. and Gilovich, T. (1989). The cold facts about the "hot hand" in basketball. *Chance*, 2:16–21.

van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge University Press, Cambridge.

Varian, H. R. (1975). Distributive justice, welfare economics, and the theory of fairness. *Philosophy and Public Affairs*, 4:223–247.

Voldner, B., Frøslie, K. F., Haakstad, L., Hoff, C., and Godang, K. (2008). Modifiable determinants of fetal macrosomia: role of lifestyle-related factors. *Acta Obstetricia et Gynecologica Scandinavica*, 87:423–429.

von Bahr, B. (1965). On the convergence of moments in the central limit theorem. *Annals of Mathematical Statistics*, xx:808–818.

von Bortkiewicz, L. (1898). *Das Gesetz der kleinen Zahlen*. B.G. Teubner, Berlin.

Vovk, V., Gammerman, A., and Shafer, G. (2005). *Algorithmic learning in a random world*. Springer Science & Business Media, Berlin/Heidelberg.

Walløe, L., Hjort, N. L., and Thoresen, M. (2019a). Major concerns about late hypothermia study. *Acta Paediatrica*, 108:588–589.

Walløe, L., Hjort, N. L., and Thoresen, M. (2019b). Why results from Bayesian statistical analyses of clinical trials with a strong prior and small sample sizes may be misleading: The case of the NICHD Neonatal Research Network Late Hypothermia Trial. *Acta Paediatrica*, 108:1190–1191.

Wand, M. P. and Jones, M. C. (1995). *Kernel Smoothing*. Chapman and Hall, London.

Wardrop, R. L. (1995). Simpson's paradox and the hot hand in basketball. *The American Statistician*, 49:24–28.

Williams, D. (1991). *Probability with Martingales*. Cambridge University Press, Cambridge.

Wilmoth, J. R., Andreev, K., Jdanov, D., Glei, D., Riffe, T., Boe, C., Bubenheim, M., Philipov, D., Shkolnikov, V., Vachon, P., C, W., and M, B. (2021). Methods protocol for the Human Mortality Database. University of California, Berkeley, US, and Max Planck Institute for Demographic Research, Rostock, Germany. `https://www.mortality.org/` [Version 6. Last revised January 26, 2021].

Wissner-Gross, Z. (2020). Can you feed the hot hand? https://fivethirtyeight.com/features/can-you-feed-the-hot-hand/. Accessed: December 12, 2020.

Xie, M. and Singh, K. (2013). Confidence distribution, the frequentist distribution estimator of a parameter: a review [with discussion and a rejoinder]. *International Statistical Review*, 81:3–39.

Zabriskie, B. N., Corcoran, C., and Senchaudhuri, P. (2021). A comparison of confidence distribution approaches for rare event meta-analysis. *Statistics in Medicine*, 40:5276–5297.

Zellner, A. (1986). Bayesian estimation and prediction using asymmetric loss functions. *Journal of the American Statistical Association*, 81:446–451.

Zhang, L., Mykland, P. A., and Aït-Sahalia, Y. (2005). A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association*, 100:1394–1411.

Name index

DeGroot, Morris H, 240, 357

Schervish, Mark J., 296

Subject index

π-system, 646
σ-finite measure, 645
d-system, 646
Monotone convergence theorem, 653
σ-algebra, 642

absolute continuity of measures, 654
absolutely continuous distribution, 659
Admissibility, 296
algebra of sets, 646
ancillary statistic, 164

Bahadur's theorem, 300
Bayes risk, 297
Bayes solution, 297
Bayes theorem, 667
Bayes' theorem, 671
Blyth's method, 305
bootstrapping argument, 654
Borel σ-algebra on of $\mathbb{R}^k$, 643
Borel–Cantelli lemma, 53
bounded in probability, 65

Carathéodory's Extension Theorem, 647
Carathéodory's lemma, 648
Cartesian product, 642
change of variable formula, 654
Characteristic functions with finite integral, 679
complement, 641
Completeness, 300
conditional distribution, 670
conditional expectation, 668, 669
conditional probability, 670
conditional sufficiency, 163
consistency of an estimator, 52
convergence in measure, 651
convergence in probability, 651
convolutions, 663
counting measure, 646
covariance, 666
Cramér–Wold theorem, 76
cumulative distribution function, 657

de Morgan's laws, 642
density, 654, 661
Derivative under the integral sign, 655
differentiability in quadratic mean, 214
discrete probability distribution, 659
distribution, 658
distributive laws, 641
Dominated convergence theorem, 653
Doob–Dynkin lemma, 669
Dynkin's lemma, 646

empty set, 642
expectation, 658
extended real numbers, 643
extension of probability spaces, 673

Factorisation theorem, 161
Fatou's lemma, 653
finitely additive measure, 645
first Bartlett identity, 193
floor function, 59
Fubini's theorem, 656

Hahn decomposition, 661

improper priors, 304
independence of σ-algebras, 665
independence, countably many events, 665
independence, finitely many events, 665
independent random variables, 665
infinitely often, 651
integrable functions, 652

intersection, 641
inverse image, 642
inversion formula, 679

Jensen's inequality for conditional expectation, 670

Kolmogorov's zero-one law, 666
kurtosis, 57

Laplace transform, 676
law, 658
law of large numbers, weak, 59
law of total probability, 667
Lebesgue decomposition, 662
Lehmann–Scheffé theorem, 300
Levy's continuity theorem, 77
Lindeberg CLT, 70
loss function, 295

Markov's inequality, 53
measurable space, 642, 645
measure, 645
method of moments, 130
Mills ratio, 54
minimal sufficient statistic, 163
Minimax, 296
mixture distribution, 659
Monotone class, 646
Monotone class theorem, 647
multivariate CLT, 77

natural parameter region, 36

outer-measure, 648

posterior density, 671
power set, 641, 642
probability density function, 660
probability distribution, 658
probability kernel, 673
probability mass function, 659
probability measure, 645
probability transform, 64
product measure, 655
product σ-algebra, 643
Prokohorov's theorem, 77

Radon–Nikodym theorem, 660
Rao–Blackwell theorem, 299
regular conditional distributions, 682
regular conditional probability, 670
Riemann–Lebesgue lemma, 681
risk, 240, 357

score function, 193
second Bartlett identity, 193
semi-ring, 648
semialgebra, 648
separable σ-algebra, 643
sigma-algebra, σ-algebra, 642
signed measure, 661
simple function on standard form, 645
Simple functions, 645
stable convergence, 88
standard deviation, 659
strong consistency of an estimator, 52
subsequence lemma, 67
subset, 641
sufficient statistic, 160

tail-σ-algebra, 666
the diagonal method, 66
tight random vectors, 76
Tightness, 65
Toeplitz lemma, 58
Tonelli's theorem, 656
tower property, 668
triangular array, 74
trivial σ-algebra, 642
Type I and Type II errors, 296

unbiased estimator, 296
uniformly minimum variance unbiased estimator, 298
union, 641
uniqueness of c.f., 680

variance, 659
variance decomposition formula, 670
version, 668

weak convergence, 61