University of Oslo / Department of Economics / NCF 2009–04–24

ECON5160: Continuous time: generators, optimal control and mathematical finance

Purpose of note: give short form and essentials of the topics stated in the headline (covered in the March 27th lecture).

1 The generator

For a Markov process X, its generator – taking a function as argument and returning a function – can be defined (for those f for which the limit exists) as

$$\mathsf{A}_X f(x) = \lim_{h \searrow 0} \frac{\mathsf{E}[f(X(h)) - f(X(0))|X(0) = x]}{h}.$$

- **Essential:** Dynkin's formula

  $$\mathsf{E}[f(X(\tau))|X(0)] = f(X(0)) + \mathsf{E}\int_0^\tau \mathsf{A}_X f(X(t))\,\mathrm{d}t$$

  valid for deterministic times and also for integrable stopping times τ.

- **Essential:** If X is the solution of a stochastic differential equation $\mathrm{d}X(t) = \alpha(X(t))\,\mathrm{d}t + \sigma(X(t))\,\mathrm{d}B(t)$, then the generator will take the form $\mathsf{A}_X f(x) = \alpha(x)f'(x) + \frac{1}{2}\sigma^2(x)f''(x)$ (again, assuming f nice enough). There is also a multidimensional version: if $\boldsymbol{X}$, $\boldsymbol{\alpha}$ and $\boldsymbol{B}$ are vector-valued and $\boldsymbol{\sigma}$ is a matrix, then the form is

  $$\mathsf{A}_{\boldsymbol{X}} f = \alpha \cdot \boldsymbol{\nabla} f + \frac{1}{2}\sum_i \sum_j \sigma_i \sigma_j f''_{ij}$$

- **Nice to know:** For a Poisson process N with intensity λ, we have $\mathsf{A}_N f(x) = (f(x+1) - f(x))\lambda$. Then if X solves a SDE $\mathrm{d}X(t) = \alpha(X(t))\,\mathrm{d}t + \sigma(X(t))\,\mathrm{d}B(t) + \gamma(X(t^-))\,\mathrm{d}N(t)$, the generator (in the scalar version) will take the form

  $$\mathsf{A}_X f(x) = \alpha(x)f'(x) + \frac{1}{2}\sigma^2(x)f''(x) + (f(x + \gamma(x)) - f(x))\lambda.$$

2 Optimal control: Superoptimality and optimality

Assume that a strong Markov process X can be controlled by a parameter u, which affects the generator A_X; write A^u to signify this dependence. In this course, it means that the coefficients α and σ (and if applicable γ, but our main focus is the case without Poisson noise) are functions of $(X(t), u(t))$ where u is our control, chosen free in some given set. Assume we are to maximize

$$J^u(x) = \mathsf{E}^x[\int_0^T f_0(X(t), u(t))\,\mathrm{d}t + F(X(T))]$$

where «E^x» denotes expectation conditioned on $X(0) = x$.

It turns out that we need only consider (Markov) feedback controls $u = u(X(t))$. Let $V = V(x)$ be the value function, for the time being unknown. Consider some function v for which $\mathsf{A}^u v(x) + f_0(x, u) \le 0$ for all u. Then

$$0 \ge \mathsf{E}^x[\int_0^T (\mathsf{A}^u v(X(t)) + f_0(X(t), u(t)))\,\mathrm{d}t]$$

which by Dynkin's formula equals $\mathsf{E}^x[v(X(T)) - v(x)] + \mathsf{E}^x[\int_0^T f_0(X(t), u(t)))\,\mathrm{d}t]$, so that

$$v(x) \ge \mathsf{E}^x[v(X(T))] + \mathsf{E}^x[\int_0^T f_0(X(t), u(t)))\,\mathrm{d}t].$$

If now $v(X(T)) \ge F(X(T))$, we have that $v(x)$ exceeds J^u – for *all* u, hence $v \ge V$.
If furthermore $v(X(T)) = F(X(T))$, and there is some u^* such that $\mathsf{A}^u v(x) + f_0(x, u^*(x)) = 0$, then we get the same calculations but with *equality,* so that $J^{u^*}(x) = v(x)$, implying that $v = V$ and that u^* is an optimal control.

Most essential: The *Hamilton-Jacobi-Bellman (HJB) equation* $\sup_u\{\mathbf{A}^u v + f_0\} = 0$ for optimality (necessary; sufficient conditions together with the boundary condition).

Extensions:

- Above, we calculated as if T were a non-random time. However, it works just as well for certain stopping times; if τ is the first exit time from a set S and we have $v(x) = F(x)$ on the boundary of[1] S (note: only «$\ge$» is required for superoptimality), then everything works just fine if Dynkin's formula holds – and it turns out that even if $\mathsf{E}^x[\tau] = \infty$, everything works well if (I) a boundedness, «(P)» or «(N)» condition in Seierstad pp. 10–11 holds, and (II) for the event $\tau = \infty$, we interpret the bequest value F as zero. Then infinite horizon is covered, and also «possibly infinite» horizon (where both $\tau < \infty$ and $\tau = \infty$ have positive probability).
- Explicit t-dependence: Consider this a problem in $Y = (t, X)$ instead of in X.

[1] «boundary» if X is continuous, or if we know that it will not «jump out of» S; otherwise, we need to replace «on the boundary of S» by «outside S».

- **Essential:** The case where the *only* time-dependence is through discounting: $f_0(t,x,u) = e^{-\delta t} g_0(x,u)$. Then the HJB equation is modified into

$$-\delta v + \sup_u \{\mathbf{A}^u v + f_0\} = 0.$$

- **Nice to know:** In the general case, you must consider as criterion $J^u(t_0, x) = \int_{t_0}^T$ [etc. where the integral starts from t_0]. Then the value function will depend on t, which is in the very last step put equal to 0.

3 Mathematical finance and arbitrage-free pricing

A single period preliminary: Consider a market consisting of a *safe* investment opportunity with price 1 both today and tomorrow, and a *risky* investment opportunity which today has price $s > 0$, but tomorrow goes either «up» to su or «down» to $sd < su$ (both with positive probability, but no other states tomorrow are possible). Assuming that there is no arbitrage opportunities, i.e. no riskless free lunch, we must have $u > 1 > d$.
Q: What is the price of a security that tomorrow pays 1 if «up» and 0 if «down»?
A: Start with endowment a, and invest ξ in the risky opportunity, and the rest $(a - \xi)$ in the safe. Tomorrow you will have $a - \xi + \xi u$ or $a - \xi + \xi d$. Now try to set a and ξ so that in the «up» state, you have 1 and in the «down» state you have 0; then a will be the only *arbitrage-free* price for the $1_{\text{«up»}}$ security, that is: since you can generate a perfect copy of the security by starting with a, then any other price will either give away a riskless free lunch or require someone to give you a riskless free lunch.
Solving $a - \xi(1-d) = 0$ and $a - \xi(1-u) = 1$, we get $\xi = (u-d)^{-1}$ and $a = (1-d)/(u-d)$. Notice that $0 < a < 1$ since $u - d > 1 - d > 0$. Notice also that $s = asu + (1-a)sd$ – so that today's price is a weighted average (with a and $(1-a)$ as weights) between tomorrow's two possible prices. It is of course nothing deep in the statement that an investment is priced somewhere in between the worst-case and the best-case – the striking fact is that the same weight a can be used for all such securities, including the underlying investment opportunity itself. Also, since a and $(1-a)$ are positive and sum to 1, they can be interpreted as *probabilities*. They have nothing[2] to do with the actual probability of the «up» event, but pricing is done by this fictitious probability measure Q defined by $\mathsf{Pr}_Q[\text{«up»}] = a$, and any random variable F is priced by $\mathsf{E}_Q[F]$ (discounted by risk free interest, here $= 0$).
Then, you may ask, what about the preferences towards risk? Don't they have any impact whatsoever? Yes they do. It is the market agents who, through their preferences towards risk and their judgement of the riskiness, decide upon a price of 1 for a security which pays u in case «up» and d in case «down». Q – and hence derivatives prices – will be determined through the pricing of this «underlying» investment opportunity[3].

[2] OK, in fact quite a bit: they agree on what future outcomes are possible and impossible, so in terms of probability measures they agree on the «null sets». But once they are agreed upon, then «nothing».

[3] In our cases, Q will be unique since the markets are *complete* – which is to say that you can generate any random variable by trading. In incomplete markets, there will be a range of possible Q's compatible with the price 1. The market then determines both the «1» and the Q.

The fundamental theorem, the Black–Scholes formula and arbitrage-free pricing in continuous time: The above result – that in such[4] frictionless[5] markets, there exists a probability measure Q such that all state-contingent claims F paid out at time T are priced by $\mathsf{E}_Q[F]$ discounted by risk-free interest – is called *the fundamental theorem of arbitrage-free pricing* or sometimes even *the fundamental theorem of asset pricing*. We did only treat complete markets, so Q is unique (and given in a previous problem[6]).

The *Black–Scholes market* has a riskless asset with continously compounded interest at constant rate r, and a risky asset (only one in this course) whose price behaves as a *geometric Brownian motion* (gBm) $S(t) = \exp((\mu - \frac{1}{2}\sigma^2)t + \sigma B(t))$ – this S solves the stochastic differential equation $\mathrm{d}S = \mu S \,\mathrm{d}t + \sigma S \,\mathrm{d}B(t)$.
The Black–Scholes *formula*[7] gives the theoretical price of a so-called European-type call option, which is the right (but not the duty) at a fixed future time T to buy the asset – worth $S(T)$ at a given price («strike») K. The price (consult the book for the formula itself) depends on price, strike, time and σ (but – fortunately! – not μ). Rather than estimating σ from prices directly, practitioneers will frequently infer σ from liquidly traded option prices and the Black–Scholes formula, and then use this «implied volatility» to price more exotic derivatives where there are no observable market prices.

More generally, it turns out that we might price derivatives on a gBm by a quite simple principle: the new measure Q works by replacing μ by the risk-free rate r. Outline:

Essential: Theoretical arbitrage-free price at time 0 of a claim F (a «financial derivative» on the «underlying» gBm S), to be paid out at time T:

- Replace μ by r, i.e. the process $S(t)$ by $Z(t) = S(0)\exp((r - \frac{1}{2}\sigma^2)t + \sigma B(t))$ – note that $Z(0) = S(0) =$ today's price of the underlying.
- If the claim has form $F = G(S(T))$, the price at time 0 is $e^{-rT}\mathsf{E}[G(Z(T))]$.
- More generally, assume the claim has form $F = G(S(\tau_1), S(\tau_2), \dots)$ where the τ_i are stopping times for S (i.e. «first time S does this or that») so that each τ_i is $\in [0, T]$. Then the price at time 0 is $e^{-rT}\mathsf{E}[G(Z(\tilde{\tau}_1), Z(\tilde{\tau}_2), \dots)]$ where $\tilde{\tau}_i$ are the similar stopping times for Z (i.e. «first time Z does this or that»).

[4] In the most general case, we need to assume that the market has no «free lunch with vanishing risk». An FLVR is as close you get to a free lunch: think of it as «a lunch available at any positive price».
[5] also, permitting borrowing at risk-free rate as well as short sale (i.e. selling something you don't own).
[6] http://www.uio.no/studier/emner/sv/oekonomi/ECON5160/v09/5160_20090322_problems.pdf problem 1.4 with no jump terms.
[7] We did not derive it, but it might be done through replication as in the single period example.