at the 5% significance level (Exercise 3.17). Third, the gap is large if it is measured instead in percentage terms: According to the estimates in Table 3.1, in 2008 women earned 16% less per hour than men did ($4.11/$24.98), slightly more than the gap of 14% seen in 1992 ($3.22/$23.27). Fourth, the gender gap is smaller for young college graduates (the group analyzed in Table 3.1) than it is for all college graduates (analyzed in Table 2.4): As reported in Table 2.4, the mean earnings for all college-educated women working full-time in 2008 was $23.93, while for men this mean was $30.97, which corresponds to a gender gap of 23% [= (30.97 − 23.93)/30.97] among all full-time college-educated workers.

This empirical analysis documents that the "gender gap" in hourly earnings is large and has been fairly stable (or perhaps increased slightly) over the recent past. The analysis does not, however, tell us *why* this gap exists. Does it arise from gender discrimination in the labor market? Does it reflect differences in skills, experience, or education between men and women? Does it reflect differences in choice of jobs? Or is there some other cause? We return to these questions once we have in hand the tools of multiple regression analysis, the topic of Part II.

[1]Because of inflation, a dollar in 1992 was worth more than a dollar in 2008, in the sense that a dollar in 1992 could buy more goods and services than a dollar in 2008 could. Thus earnings in 1992 cannot be directly compared to earnings in 2008 without adjusting for inflation. One way to make this adjustment is to use the CPI, a measure of the price of a "market basket" of consumer goods and services constructed by the Bureau of Labor Statistics. Over the 16 years from 1992 to 2008, the price of the CPI market basket rose by 53.4%; in other words, the CPI basket of goods and services that cost $100 in 1992 cost $153.40 in 2008. To make earnings in 1992 and 2008 comparable in Table 3.1, 1992 earnings are inflated by the amount of overall CPI price inflation, that is, by multiplying 1992 earnings by 1.534 to put them into "2008 dollars."

3.6 Using the *t*-Statistic When the Sample Size Is Small

In Sections 3.2 through 3.5, the *t*-statistic is used in conjunction with critical values from the standard normal distribution for hypothesis testing and for the construction of confidence intervals. The use of the standard normal distribution is justified by the central limit theorem, which applies when the sample size is large. When the sample size is small, the standard normal distribution can provide a poor approximation to the distribution of the *t*-statistic. If, however, the population distribution is itself normally distributed, then the exact distribution (that is, the finite-sample distribution; see Section 2.6) of the *t*-statistic testing the mean of a single population is the Student t distribution with $n - 1$ degrees of freedom, and critical values can be taken from the Student t distribution.

The *t*-Statistic and the Student t Distribution

The t-statistic testing the mean. Consider the *t*-statistic used to test the hypothesis that the mean of Y is $\mu_{Y,0}$, using data $Y_1, \ldots, Y_n$. The formula for this statistic