4.3 Some Common Distributions

Most test statistics in econometrics follow one of four well-known distributions, at least approximately. These are the standard normal distribution, the chi-squared (or χ^2) distribution, the Student's t distribution, and the F distribution. The most basic of these is the normal distribution, since the other three distributions can be derived from it. In this section, we discuss the standard, or **central**, versions of these distributions. Later, in Section 4.7, we will have occasion to introduce **noncentral** versions of all these distributions.

The Normal Distribution

The **normal distribution**, which is sometimes called the **Gaussian distribution** in honor of the celebrated German mathematician and astronomer Carl Friedrich Gauss (1777–1855), even though he did not invent it, is certainly the most famous distribution in statistics. As we saw in Section 1.2, there is a whole family of normal distributions, all based on the **standard normal distribution**, so called because it has mean 0 and variance 1. The PDF of the standard normal distribution, which is usually denoted by $\phi(\cdot)$, was defined in (1.06). No elementary closed-form expression exists for its CDF, which is usually denoted by $\Phi(\cdot)$. Although there is no closed form, it is perfectly easy to evaluate Φ numerically, and virtually every program for doing econometrics and statistics can do this. Thus it is straightforward to compute the P value for any test statistic that is distributed as standard normal. The graphs of the functions ϕ and Φ were first shown in Figure 1.1 and have just reappeared in Figure 4.2. In both tails, the PDF rapidly approaches 0. Thus, although a standard normal r.v. can, in principle, take on any value on the real line, values greater than about 4 in absolute value occur extremely rarely.

In Exercise 1.7, readers were asked to show that the full normal family can be generated by varying exactly two parameters, the mean and the variance. A random variable X that is normally distributed with mean μ and variance σ^2 can be generated by the formula

$$X = \mu + \sigma Z, \tag{4.09}$$

where Z is standard normal. The distribution of X, that is, the normal distribution with mean μ and variance σ^2, is denoted $\mathrm{N}(\mu, \sigma^2)$. Thus the standard normal distribution is the $\mathrm{N}(0, 1)$ distribution. As readers were asked to show in Exercise 1.8, the PDF of the $\mathrm{N}(\mu, \sigma^2)$ distribution, evaluated at x, is

$$\tfrac{1}{\sigma}\phi\Big(\frac{x-\mu}{\sigma}\Big) = \frac{1}{\sigma\sqrt{2\pi}}\exp\Big(-\frac{(x-\mu)^2}{2\sigma^2}\Big). \tag{4.10}$$

In expression (4.10), as in Section 1.2, we have distinguished between the random variable X and a value x that it can take on. However, for the following discussion, this distinction is more confusing than illuminating. For