Inference from Small Samples 17

17.A Appendix: Descendants of the Standard Normal Distribution

We now introduce three new families of distributions that are used in inference procedures in settings with normally distributed trials. These include not only the procedures introduced in this chapter, but also ones for inference about differences in means (Chapter 18) and about regression parameters (Chapter 20). All three families of distributions can be defined as the distributions of functions of independent standard normal random variables, and their definitions suggest some of the contexts in which they will be useful for inference.

The three new families of distributions are the χ^2 (or chi-squared), t, and F distributions.[1] Members of the χ^2 and t families of distributions are distinguished by a parameter d, while F distributions are distinguished by a pair of parameters (k, d). All of these parameters are called **degrees of freedom**.

Definition.

Let $\{Z_i\}_{i=1}^{d}$ be a sequence of i.i.d. standard normal random variables, and let the random variable C be the sum of the squares of the Z_i:

$$C = \sum_{i=1}^{d} (Z_i)^2. \quad (17.A.1)$$

Then C has a **χ^2** (**or chi-squared**) **distribution with d degrees of freedom** (denoted $C \sim \chi^2(d)$).

Definition.

Let Z be a standard normal random variable, and let C be a random variable that has a $\chi^2(d)$ distribution and that is independent of Z. If we define the random variable T by

$$T = \frac{Z}{\sqrt{C/d}}, \quad (17.A.2)$$

[1] χ is the lowercase version of the Greek letter *chi*. Although it looks like an x, the sound it represents is closer to that of a Scottish ch. “Chi” is pronounced to rhyme with “buy,” but with the b replaced by a k.

then T is said to have a **t distribution with d degrees of freedom** (denoted $T \sim t(d)$).

Definition.

Let C_k and C_d be independent χ^2 random variables with k and d degrees of freedom, respectively. If we define the random variable $\mathcal{F}$ by

$$\mathcal{F} = \frac{C_k/k}{C_d/d}, \tag{17.A.3}$$

then $\mathcal{F}$ is said to have an **F distribution with k and d degrees of freedom** (denoted $\mathcal{F} \sim F(k,d)$).

While these families of distributions do not have very memorable names, they appear so frequently in statistical analyses that you quickly get used to them.[2]

Unlike standard normal and t distributions, χ^2 and F distributions are not symmetric. Therefore, while it was enough to define z-values and t-values for the right tails of their distributions, we need to define separate χ^2-values and F-values for the left and right tails. This leads to a different notation for these values. If $C \sim \chi^2(d)$, we write $\underline{c}_\alpha^d$ for the left-tail value and $\bar{c}_\alpha^d$ for the right-tail value. These are defined by

$$\mathrm{P}(C \leq \underline{c}_\alpha^d) = \alpha \text{ and } \mathrm{P}(C \geq \bar{c}_\alpha^d) = \alpha.$$

In the language of Section 6.2, $\underline{c}_\alpha^d$ is the (100α)th percentile of the $\chi^2(d)$ distribution, and $\bar{c}_\alpha^d$ is the $(100(1-\alpha))$th percentile of this distribution.

Similarly, for $\mathcal{F} \sim F(k,d)$, we define the left-tail value $\underline{F}_\alpha^{k,d}$ and the right-tail value $\bar{F}_\alpha^{k,d}$ by

$$\mathrm{P}(\mathcal{F} \leq \underline{F}_\alpha^{k,d}) = \alpha \text{ and } \mathrm{P}(\mathcal{F} \geq \bar{F}_\alpha^{k,d}) = \alpha.$$

Excel calculation: *Finding χ^2 and F probabilities and values*

Probabilities and values for χ^2 and F distributions can be computed using the appropriate worksheets in the `distributions.xlsx` workbook. They can also be obtained using built-in Excel functions:

Suppose $C \sim \chi^2(d)$.
To obtain $\mathrm{P}(C < a)$, enter "`=CHISQ.DIST(a,d,1)`".
To obtain $\mathrm{P}(C > b)$, enter "`=CHISQ.DIST.RT(b,d)`".
To obtain $\underline{c}_\alpha^d$, enter "`=CHISQ.INV(α,d)`".

[2]Where do these names come from? The name of the χ^2 distribution is explained by the fact that the Greek letter χ is an old notation for the standard normal distribution. The F distribution was named in honor of R. A. Fisher, the father of modern statistics, whom we introduced in Section 14.3.2. But the choice of the letter t for the t distribution seems to have been made by Gosset for no good reason at all. See Churchill Eisenhart, "On the Transition from 'Student's' z to 'Student's' t," *American Statistician* 33 (1979), 6–10.

To obtain $\bar{c}_\alpha^{d}$, enter "=CHISQ.INV.RT(α,d)" (or "=CHISQ.INV(1-α,d)").

Suppose $\mathcal{F} \sim F(k, d)$.
To obtain P($\mathcal{F} < a$), enter "=F.DIST(a,k,d,1)".
To obtain P($\mathcal{F} > b$), enter "=F.DIST.RT(b,k,d)".
To obtain $\underline{F}_\alpha^{k,d}$, enter "=F.INV($\alpha$,k,d)".
To obtain $\bar{F}_\alpha^{k,d}$, enter "=F.INV.RT(α,k,d)" (or "=F.INV(1-α,k,d)").

17.A.1 χ^2 distributions

The $\chi^2(d)$ distribution is defined as the distribution of the sum of the squares of d independent standard normal random variables. Because they represent sums of positive random variables, χ^2 distributions place all of their mass on positive outcomes.

Figure 17.A.1 presents the density functions of χ^2 distributions with various numbers of degrees of freedom.[3] As the number of degrees of freedom increases, we are summing the squares of more independent standard normal random variables, so the mass in the distribution shifts to the right and becomes more spread out. To be more precise, it can be shown that if C has a $\chi^2(d)$ distribution, then its mean and variance are $\mathrm{E}(C) = d$ and $\mathrm{Var}(C) = 2d$ (see Exercise 17.A.12).

Figure 17.A.1: Some χ^2 density functions.

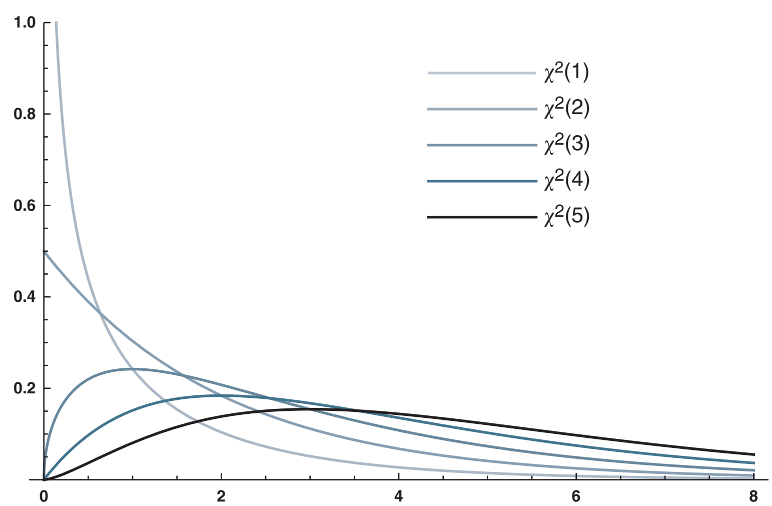

[3]The density functions for χ^2, t, and F distributions are stated in Appendix 17.A.6.

Rather than consider the sum of d squared standard normal random variables, as in (17.A.1), we can instead focus on their sample mean,

$$\frac{C}{d} = \frac{1}{d}\sum_{i=1}^{d}(Z_i)^2.$$

In the definitions of the t and F distributions, the χ^2 random variables appear in this form. The random variable C/d is said to have an **averaged χ^2 distribution with d degrees of freedom**. Since $\mathrm{E}(C) = d$ and $\mathrm{Var}(C) = 2d$, our formulas from Chapter 3 imply that

$$\mathrm{E}\left(\frac{C}{d}\right) = \frac{1}{d}\mathrm{E}(C) = 1 \quad \text{and} \quad \mathrm{Var}\left(\frac{C}{d}\right) = \frac{1}{d^2}\,\mathrm{Var}(C) = \frac{2}{d}.$$

That is, an averaged $\chi^2(d)$ random variable has a mean of 1 and a variance inversely proportional to d. These properties are illustrated in Figure 17.A.2, which presents the densities of averaged χ^2 distributions for various choices of d. While averaged chi-squared random variables are more convenient than unaveraged ones in certain respects, statistical tables and procedures are usually presented in terms of the unaveraged ones.

17.A.2 The sample variance of normal trials

The χ^2 distributions describe the estimators of dispersion for i.i.d. normal trials $\{X_i\}_{i=1}^n$, $X_i \sim N(\mu, \sigma^2)$. First consider the known-mean sample variance introduced in Section 14.4:

$$V_n = \frac{1}{n}\sum_{i=1}^{n}(X_i - \mu)^2.$$

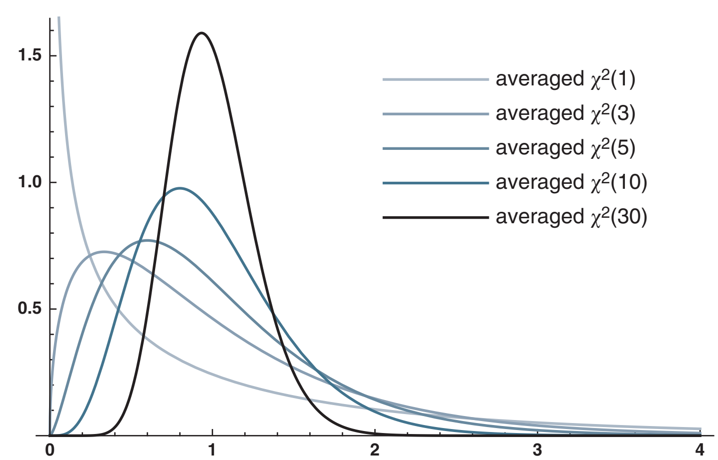

Figure 17.A.2: Some averaged χ^2 density functions.

Dividing both sides of this equation by σ^2, we obtain

$$\frac{V_n}{\sigma^2} = \frac{1}{n}\sum_{i=1}^{n}\left(\frac{X_i - \mu}{\sigma}\right)^2. \tag{17.A.4}$$

Since $X_i \sim N(\mu, \sigma^2)$, the summands in (17.A.4) are squared standard normal random variables, so V_n/σ^2 has an averaged χ^2 distribution with n degrees of freedom. This fact is usually expressed in terms of the original χ^2 distribution:

$$\text{If } \{X_i\}_{i=1}^n \text{ is i.i.d. with } X_i \sim N(\mu, \sigma^2), \text{ then } \frac{n}{\sigma^2}V_n \sim \chi^2(n). \tag{17.A.5}$$

Since the $\chi^2(n)$ distribution has mean n, (17.A.5) implies that $\mathrm{E}(V_n) = \sigma^2$: the known-mean sample variance is an unbiased estimator of the variance σ^2. We established the unbiasedness of V_n already, in Section 14.4. Fact (17.A.5) does much more than this. It fully describes the distribution of V_n, provided that the trials are normally distributed.

If the mean μ were known, this fact would allow us to use V_n to construct interval estimators and hypothesis tests about the variance σ^2. But in reality, since we almost never know μ in advance, we estimate σ^2 using the sample variance,

$$\mathcal{S}_n^2 = \frac{1}{n-1}\sum_{i=1}^{n}(X_i - \bar{X}_n)^2.$$

When the trials are normally distributed, the distribution of the sample variance can also be described using a χ^2 distribution, but this time one with $n - 1$ degrees of freedom.

Distribution of the sample variance of normal trials.

$$\textit{If } \{X_i\}_{i=1}^n \textit{ is i.i.d. with } X_i \sim N(\mu, \sigma^2), \textit{ then } \frac{n-1}{\sigma^2}\mathcal{S}_n^2 \sim \chi^2(n-1). \tag{17.A.6}$$

It follows from (17.A.6) that $\mathrm{E}(\mathcal{S}_n^2) = \sigma^2$: the sample variance $\mathcal{S}_n^2$ is an unbiased estimator of the variance σ^2. But as before, (17.A.6) fully describes the distribution of $\mathcal{S}_n^2$, provided the trials are normally distributed. Among other applications, this fact can be used to define interval estimators and hypothesis tests about an unknown variance of normally distributed trials—see Exercise 17.A.13.

■ Example *Dispersion in driving speeds.*

Driving speeds of vehicles on a stretch of I-94 near Madison are normally distributed with a standard deviation of 10 miles per hour, and hence a variance of 100 (miles per hour)^2. If we randomly sample the speeds of 13 vehicles, what is the probability that the sample variance exceeds 150?

To compute the probability that $P(\mathcal{S}^2_{13} > 150)$, we need to rewrite the inequality in terms of a random variable C with a $\chi^2(12)$ distribution:

$$P(\mathcal{S}^2_{13} > 150) = P\left(\frac{n-1}{\sigma^2}\mathcal{S}^2_{13} > \frac{12}{100} \cdot 150\right) = P(C \geq 18).$$

Using the `distributions.xlsx` workbook, we find that this probability is .1157. ■

In Section 14.4.2, we discussed the meaning of the term "degrees of freedom," noting[4] that the distribution of the sample variance of n i.i.d. normal trials can be defined in terms of $n - 1$ i.i.d. normal random variables. This is the precisely the idea underlying fact (17.A.6) above, which shows that after a suitable rescaling, the sample variance $\mathcal{S}^2_n$ follows a χ^2 distribution with $n - 1$ degrees of freedom—that is, the distribution of the sum of $n - 1$ independent squared standard normals. Deriving this fact is beyond the scope of this book, but we will discuss some related ideas in Section 17.A.4.

17.A.3 *t* distributions

Our small-sample inference procedures in this chapter are based on *t* distributions. As we stated above, the *t* distribution with d degrees of freedom is the distribution of the random variable

$$T = \frac{Z}{\sqrt{C/d}}, \tag{17.A.2}$$

where Z is standard normal, C has a $\chi^2(d)$ distribution, and Z and C are independent. Since $E(C/d) = 1$, T is obtained by taking a standard normal random variable and dividing it by the square root of a random variable whose mean is 1. If C/d were always equal to 1, then T would simply be standard normal. In fact, *t* distributions resemble standard normal distributions, in that they are symmetric, bell-shaped curves centered at 0. But C/d has some dispersion—as we saw earlier, $\text{Var}(C/d) = \frac{2}{d}$. This dispersion leads the *t* distributions to exhibit greater variation about zero than the $N(0, 1)$ distribution, as shown in Figure 17.1.

As d grows large, the variance $\text{Var}(C/d) = \frac{2}{d}$ shrinks to zero. Thus, in the language of Section 7.3, C/d converges in probability to 1 as d approaches infinity. Thus as d grows large, the $t(d)$ distribution comes to closely resemble the standard normal distribution. We can state this point formally using the notion of convergence in distribution, which we introduced in Section 7.4 in order to state the central limit theorem.

[4]See footnote 13 and Exercise 14.M.3.

Convergence of *t* distributions to the standard normal distribution.

If $T_d \sim t(d)$, then as d approaches infinity, T_d converges in distribution to $Z \sim N(0,1)$.

17.A.4 Why does the *t*-statistic have a *t* distribution?

At last, we are prepared to explain the fact underlying the inference procedures developed in this chapter: that the t-statistic of a random sample of n normal trials, defined by

$$\frac{\bar{X}_n - \mu}{\frac{1}{\sqrt{n}}\mathcal{S}_n},$$

has a t distribution with $n-1$ degrees of freedom.

Recall that the $t(d)$ distribution is defined as the distribution of the random variable

$$\frac{Z}{\sqrt{C/d}}, \tag{17.A.7}$$

where the random variables Z and C satisfy the following three assumptions:

(i) $Z \sim N(0,1)$,
(ii) $C \sim \chi^2(d)$, and
(iii) Z and C are independent.

Let's relate these conditions to the ingredients in our t-statistic for i.i.d. normal trials. To start, recall from Section 17.1 that the z-statistic of such trials has a standard normal distribution:

$$\frac{\bar{X}_n - \mu}{\frac{\sigma}{\sqrt{n}}} \sim N(0,1). \tag{17.A.8}$$

This corresponds to item (i). Next, we observed in Section 17.A.1 that a rescaled version of the sample variance has a χ^2 distribution:

$$\frac{n-1}{\sigma^2}\mathcal{S}_n^2 \sim \chi^2(n-1). \tag{17.A.9}$$

This corresponds to item (ii). Now, we combine these two random variables the same way that Z and C are combined in equation (17.A.7) to define the t distribution, and then cancel whatever we can:

$$\frac{\frac{\bar{X}_n - \mu}{\cancel{\sigma}/\sqrt{n}}}{\sqrt{\frac{\cancel{n-1}}{\cancel{\sigma^2}}\mathcal{S}_n^2 / \cancel{(n-1)}}} = \frac{\bar{X}_n - \mu}{\frac{1}{\sqrt{n}}\mathcal{S}_n}. \tag{17.A.10}$$

The right-hand side of (17.A.10) is our t-statistic.[5]

But wait a minute… What happened to (iii)? For the argument above to establish that the t-statistic has a $t(n-1)$ distribution, we need to know that the random variables (17.A.8) and (17.A.9) are independent. This is far from obvious, but it is true.

Independence of the sample mean and sample variance of normal trials.

If $\{X_i\}_{i=1}^n$ is i.i.d. with $X_i \sim N(\mu, \sigma^2)$, then $\bar{X}_n$ and $\mathcal{S}_n^2$ are independent random variables.

With this fact in hand, the fact that functions of independent random variables are themselves independent random variables (see Appendix 4.A.1) implies that the random variables in (17.A.8) and (17.A.9) are independent. This gives us item (iii), allowing us to conclude that the t-statistic indeed has a t distribution.

The fact that $\bar{X}_n$ and $\mathcal{S}_n^2$ are independent when the trials are normally distributed has a beautiful geometric explanation. This explanation takes quite a bit of work to develop, and with regret we do not provide it here.[6]

17.A.5 *F* distributions

As we stated earlier, the F distribution with k and d degrees of freedom is the distribution of the random variable

$$\mathcal{F} = \frac{C_k/k}{C_d/d}, \tag{17.A.3}$$

where $C_k \sim \chi^2(k)$ and $C_d \sim \chi^2(d)$ are independent random variables. Put differently, an $F(k, d)$ distribution is the distribution of the ratio of an averaged $\chi^2(k)$ random variable and an averaged $\chi^2(d)$ random variable. For obvious reasons,

[5] Something subtle happens in calculation (17.A.10). If the variance σ^2 is unknown, then neither the standard normal random variable from (17.A.8) nor the χ^2 random variable from (17.A.9) is something we can observe. But when we take their ratio, the variances from (17.A.8) and (17.A.9) cancel. If we are performing a hypothesis test about the mean μ, then the null hypothesis imposes an assumption about the value of μ. Under this assumption, the t-statistic (17.A.10) is something we can compute from the results of our sample. This is important, because otherwise we would not be able to run the hypothesis test.

[6] We do not know of an explanation of this point at the level of our book. For a clear but advanced treatment, see Ronald Christensen, *Plane Answers to Complex Questions: The Theory of Linear Models*, 3rd ed., Springer, 2002.

k and d are often called the *numerator degrees of freedom* and the *denominator degrees of freedom*, respectively.

The most basic use of the F distribution is in tests of equality of variances (see Exercise 17.A.14). In this book, we use F distributions in the context of inference in regression models (Chapter 20). To conclude this chapter, we note two relations between F distributions and the families of distributions introduced above which will be useful later on.

To obtain a link with t distributions, we take the definition (17.A.2) of the $t(d)$ distribution and square both sides, obtaining

$$T^2 = \frac{Z^2}{C/d}, \tag{17.A.11}$$

where Z is standard normal, C is $\chi^2(d)$, and Z and C are independent random variables. Now, since Z is standard normal, Z^2 has a $\chi^2(1)$ distribution, and since Z and C are independent, so are Z^2 and C (see Appendix 4.A.1). Thus, comparing (17.A.11) and (17.A.3), we conclude that the square of a random variable with a $t(d)$ distribution has an $F(1, d)$ distribution.

$$\text{If } T \sim t(d), \text{ then } T^2 \sim F(1, d). \tag{17.A.12}$$

The link between F and χ^2 distributions is quite similar to the one between t distributions and the standard normal distribution (see Section 17.A.3). We illustrate the link in Figure 17.A.3, which presents the density functions of $F(5, d)$ distributions for various choices of d, along with the density function of the averaged $\chi^2(5)$ distribution (which we drew earlier in Figure 17.A.2). You can see that as d grows large, the former densities increasingly resemble the latter.

The argument that establishes this connection is similar to the corresponding one from Section 17.A.3. In definition (17.A.3), the $F(k, d)$ random variable $\mathcal{F}$ has an averaged $\chi^2(k)$ random variable as its numerator and an averaged $\chi^2(d)$ random variable as its denominator. If we let d grow large, then we know from Section 17.A.1 that the variance of the denominator of $\mathcal{F}$ goes to zero; in the

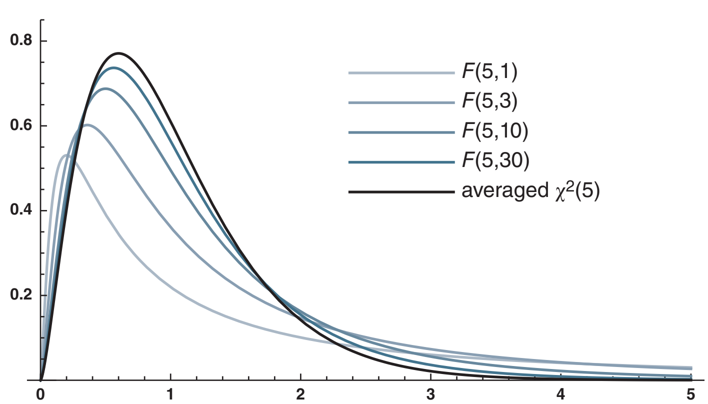

Figure 17.A.3: Some $F(5, d)$ density functions and the averaged $\chi^2(5)$ density function.

language of Section 7.3, this denominator converges in probability to 1. This suggests that as d grows large, the $F(k, d)$ distribution should approach an averaged $\chi^2(k)$ distribution, just as we see in Figure 17.A.3. We state this point formally using the notion of convergence in distribution from Section 7.4.

Convergence of F distributions to averaged χ^2 distributions.
If $\mathcal{F}_d \sim F(k, d)$, then as d approaches infinity, $\mathcal{F}_d$ converges in distribution to C/k, where $C \sim \chi^2(k)$.

SPEEDS OF CONVERGENCE.

The claim above tells us that when the number of degrees of freedom d is large enough, values from the $F(k, d)$ distribution are well approximated by values from the averaged $\chi^2(k)$ distribution. When this approximation is close, we can substitute the latter for the former in our calculations without much effect on the results. How large must d be for this to be true?

In Table 17.A.1, we present right-tail values from $F(1, d)$ for various choices of the number of degrees of freedom d and tail probabilities. As d grows large, these values converge to right-tail values from the $\chi^2(1)$ distribution (= the averaged $\chi^2(1)$ distribution). We report these χ^2-values in the last row of the table.

As with the convergence of t-values to z-values (Table 17.1), the convergence of F-values to χ^2 values becomes slower as the tail probability α becomes closer to zero. But compared to what we see in Table 17.1, the convergence in Table 17.A.1 is considerably slower. To get an $F(1, d)$ distribution that is about as close to the $\chi^2(1)$ distribution as the $t(100)$ distribution is to the standard normal distribution, we need to choose $d = 500$. Thus in statistical applications, replacing F-values with χ^2 values is only innocuous at fairly large sample sizes. We will return to this point when in the context of inference in regression models in Section 20.6.

Table 17.A.1: $F(1, d)$ values and $\chi^2(1)$ values for various numbers of degrees of freedom and tail probabilities

	$\alpha = .10$	$\alpha = .05$	$\alpha = .025$	$\alpha = .01$	$\alpha = .005$
$F^{1,1}_{1-\alpha}$	39.863	161.448	647.789	4052.181	16210.723
$F^{1,2}_{1-\alpha}$	8.526	18.513	38.506	98.503	198.501
$F^{1,5}_{1-\alpha}$	4.060	6.608	10.007	16.258	22.785
$F^{1,10}_{1-\alpha}$	3.285	4.965	6.937	10.044	12.826
$F^{1,20}_{1-\alpha}$	2.975	4.351	5.871	8.096	9.944
$F^{1,30}_{1-\alpha}$	2.881	4.171	5.568	7.562	9.180
$F^{1,50}_{1-\alpha}$	2.809	4.034	5.340	7.171	8.626
$F^{1,100}_{1-\alpha}$	2.756	3.936	5.179	6.895	8.241
$F^{1,200}_{1-\alpha}$	2.731	3.888	5.100	6.763	8.057
$F^{1,500}_{1-\alpha}$	2.716	3.860	5.054	6.686	7.950
$c^{1}_{1-\alpha}$	2.706	3.841	5.024	6.635	7.879

17.A.6 The density functions

To close this appendix, we state formulas for the density functions of the χ^2, t, and F distributions. To do so, we need to introduce an important mathematical function called the **gamma function**:

$$\Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt.$$

Although this function looks nasty, we can get a handle on it by focusing on the values it takes at integers: it can be shown that $\Gamma(n) = (n-1)!$ for each positive integer n. Thus, the gamma function provides a way of extending the factorial function to allow non-integer arguments. Below, the gamma function is only used as a normalizing term, making sure that the total areas under the density functions equal 1. Notice in particular that the arguments of the gamma functions never include x, the argument of the density function.

Density functions of the standard normal distribution and its descendants.

The density function of the standard normal distribution is

$$f(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}.$$

The density function of the $\chi^2(d)$ distribution is

$$f(x) = \frac{x^{(d/2)-1} e^{-x/2}}{\Gamma(\frac{d}{2})\, 2^{d/2}} \quad \textit{if } x > 0 \textit{ (with } f(x) = 0 \textit{ otherwise).}$$

The density function of the $t(d)$ distribution is

$$f(x) = \frac{\Gamma(\frac{d+1}{2})}{\Gamma(\frac{d}{2})\sqrt{d\pi}} \left(1 + \frac{x^2}{d}\right)^{-(d+1)/2}.$$

The density function of the $F(k, d)$ distribution is

$$f(x) = \frac{\Gamma(\frac{k+d}{2})\, x^{\frac{k}{2}-1}\, k^{\frac{k}{2}} d^{\frac{d}{2}}}{\Gamma(\frac{k}{2})\, \Gamma(\frac{d}{2})\, (kx+d)^{\frac{k+d}{2}}} \quad \textit{if } x > 0 \textit{ (with } f(x) = 0 \textit{ otherwise).}$$

If you use a computer to graph these density functions for various choices of d and k, the results should look like Figures 17.1, 17.A.1, and 17.A.3.

KEY TERMS AND CONCEPTS

degrees of freedom (p. 1)
χ^2 (or chi-squared) distribution (p. 1)
t distribution (p. 2)
F distribution (p. 2)
averaged χ^2 distribution (p. 4)
sample variance of normal trials (p. 4)
t-statistic (p. 7)
gamma function (p. 11)

Exercises

Exercise 17.A.1. Use `distributions.xlsx` or built-in Excel functions to answer these questions:

a. If $C \sim \chi^2(1)$, what is $P(C > 2)$?
b. If $C \sim \chi^2(4)$, what is $P(C < 3)$?
c. If $C \sim \chi^2(20)$, what is $P(15 < C < 25)$?

Exercise 17.A.2. Use `distributions.xlsx` or built-in Excel functions to find these χ^2-values:

a. $\underline{c}^{10}_{.05}$.
b. $\bar{c}^{10}_{.02}$.
c. $\underline{c}^{30}_{.01}$.
d. $\bar{c}^{30}_{.01}$.

Exercise 17.A.3. Use `distributions.xlsx` or built-in Excel functions to answer these questions:

a. If $\mathcal{F} \sim F(1, 11)$, what is $P(\mathcal{F} < 1.00)$?
b. If $\mathcal{F} \sim F(5, 5)$, what is $P(\mathcal{F} > 2.20)$?
c. If $\mathcal{F} \sim F(10, 30)$, what is $P(3.00 < \mathcal{F} < 4.00)$?

Exercise 17.A.4. Use `distributions.xlsx` or built-in Excel functions to find these F-values:

a. $\underline{F}^{1,10}_{.05}$.
b. $\bar{F}^{1,10}_{.02}$.
c. $\underline{F}^{1,60}_{.05}$.
d. $\bar{F}^{1,60}_{.05}$.

Exercise 17.A.5. The army is testing a new self-targeting weapon. To score the weapon's performance, it conducts 10 trials, each of which involves a target 300 yards away; the score is obtained by squaring the distance of the impact point from the target in each trial and summing the results. (A low score is a good score.)

Suppose that in each trial, the distance of the impact point from the target is normally distributed with a mean of 0 feet and a standard deviation of 1.5 feet. What is the probability that the total score is less than 30?

Exercise 17.A.6. A manufacturer of SD memory cards is evaluating the variation in the cards' capacity. Suppose that the capacities of their 256-MB cards are normally distributed with a mean of 242.2 MB and a variance of .81 MB^2. Suppose that the manufacturer measures the capacities of 11 randomly chosen cards.

a. Describe the distribution of the sample variance.
b. Compute the probability that the sample variance is less than 1.00 MB^2.

Exercise 17.A.7. The time it takes an Internet payment processor to process credit card transactions is normally distributed with a mean of 2.70 seconds and a standard deviation of .85 seconds, and the times required for distinct transactions are independent of one another. What is the probability that the sample standard deviation from 55 transactions will exceed 1.00 seconds?

Exercise 17.A.8. Let $\{X_i\}_{i=1}^{6}$ be a sequence of i.i.d. draws from a normal distribution with mean 5 and variance 100.

a. Compute $\text{P}(\bar{X}_6 > 7)$.
b. Compute $\text{P}(\mathcal{S}_6^2 > 120)$.
c. Compute $\text{P}(\bar{X}_6 > 7 \text{ and } \mathcal{S}_6^2 > 120)$. Why is this calculation straightforward given your answers to (a) and (b)?

Exercise 17.A.9. Let $T = (\bar{X}_{10} - \mu)/(\frac{1}{\sqrt{10}}\mathcal{S}_{10})$ be the t-statistic of 10 i.i.d. normal trials. What is the probability that T^2 is greater than 4? Answer this question using two different methods.

Exercise 17.A.10. Let $\{X_i\}_{i=1}^{n}$ and $\{Y_i\}_{i=1}^{m}$ be independent sequences of i.i.d. normal trials. (For instance, each sequence could represent a random sample from a distinct normally distributed population.) Suppose that the trials are drawn from normal distributions with the same variance: $\text{Var}(X_i) = \text{Var}(Y_i) = \sigma^2$. Show that the ratio of the sequences' sample variances, $\mathcal{S}_X^2/\mathcal{S}_Y^2$, has an $F(n-1, m-1)$ distribution.

Exercise 17.A.11. A precision tool manufacturer would like to know whether the variations in the lengths of knockout pins produced at two of its factories are equal. The distributions of the lengths of knockout pins at each factory are known to be normally distributed. Suppose that the variances of these distributions are equal. If the manufacturer randomly samples 15 pins from factory X and 20 pins from factory Y, what is the probability that the sample variance ratio $\mathcal{S}_X^2/\mathcal{S}_Y^2$ exceeds 2? (Hint: Use the result of Exercise 17.A.10.)

Exercise 17.A.12. Let C have a χ^2 distribution with d degrees of freedom.

a. Show that $\mathrm{E}(C) = d$. (Hint: Use the definition of the $\chi^2(d)$ distribution in terms of squared standard normal random variables and formula (4.15).)
b. It can be shown that if Z is a standard normal random variable, then $\mathrm{E}(Z^4) = 3$. Use this fact to show that $\mathrm{Var}(C) = 2d$.

Exercise 17.A.13. Let $\{X_i\}_{i=1}^n$ be i.i.d. with $X_i \sim N(\mu, \sigma^2)$.

a. Show that

$$\mathrm{P}\left(\sigma^2 \in \left[\frac{(n-1)\mathcal{S}_n^2}{\bar{c}_{\alpha/2}^{\,n-1}}, \frac{(n-1)\mathcal{S}_n^2}{\underline{c}_{\alpha/2}^{\,n-1}}\right]\right) = 1 - \alpha.$$

The random interval in brackets is thus the $(1-\alpha)$ interval estimator for σ^2 for normally distributed trials.

b. Suppose we want to test the null hypothesis $H_0 : \sigma^2 = \sigma_0^2$ against the alternative $H_1 : \sigma^2 > \sigma_0^2$ at significance level α. Show that this can be accomplished by rejecting the null hypothesis when the realized sample variance s^2 satisfies

$$s^2 > \frac{\sigma_0^2\, \bar{c}_{\alpha}^{\,n-1}}{n-1}.$$

Exercise 17.A.14. Let $\{X_i\}_{i=1}^n$ be i.i.d. $N(\mu_X, \sigma_X^2)$, let $\{Y_i\}_{i=1}^m$ be i.i.d. $N(\mu_Y, \sigma_Y^2)$, and suppose that the two sequences are independent. Exercise 17.A.10 showed that if $\sigma_X^2 = \sigma_Y^2$, then $\mathcal{S}_X^2/\mathcal{S}_Y^2$ has an $F(n-1, m-1)$ distribution.

Suppose we want to test the null hypothesis $H_0 : \sigma_X^2 = \sigma_Y^2$ against the alternative hypothesis $H_1 : \sigma_X^2 > \sigma_Y^2$ at significance level α. Show that this can be accomplished by rejecting the null hypothesis when the realized sample variance ratio s_X^2/s_Y^2 satisfies

$$\frac{s_X^2}{s_Y^2} > \bar{F}_{\alpha}^{\,n-1,m-1}.$$