The Central Limit Theorem

If a random sample of size n is taken from *any* distribution with mean μ and variance σ^2, the sampling distribution of the mean will be *approximately* $\sim \mathrm{N}(\mu, \sigma^2/n)$, where $\sim$ means 'is distributed as'. The larger n is, the better the approximation.

The standard normal and Student's t distributions

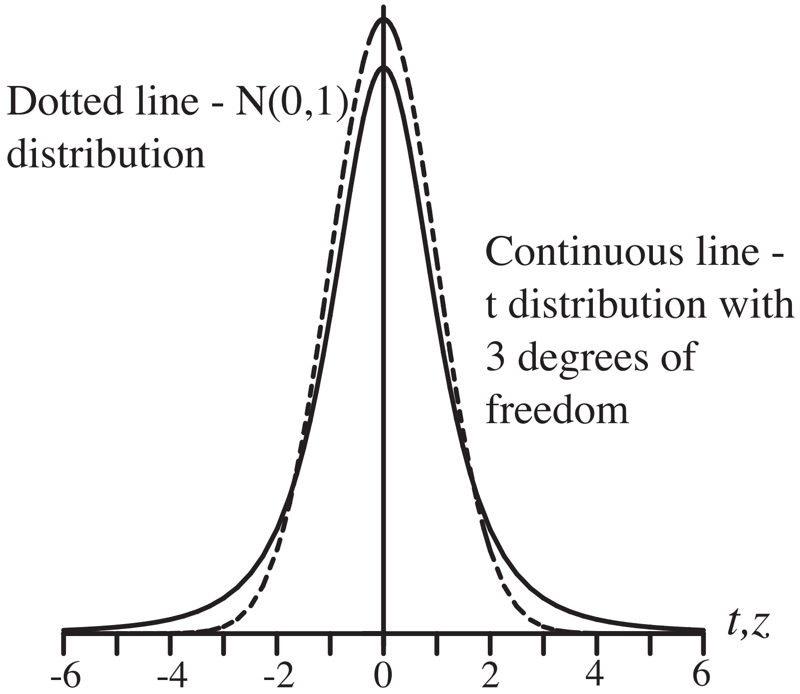

If a random variable $X \sim \mathrm{N}(\mu, \sigma^2)$, $z = (X - \mu)/\sigma \sim \mathrm{N}(0, 1)$, the *standard normal distribution*. The t distribution with $(n-1)$ degrees of freedom is used in place of z for small samples size n from normal populations when σ^2 is unknown. As n increases the distribution of t converges to $\mathrm{N}(0, 1)$. These distributions are used, e.g., for inference about means, differences between means and in regression.

Fisher's F distribution

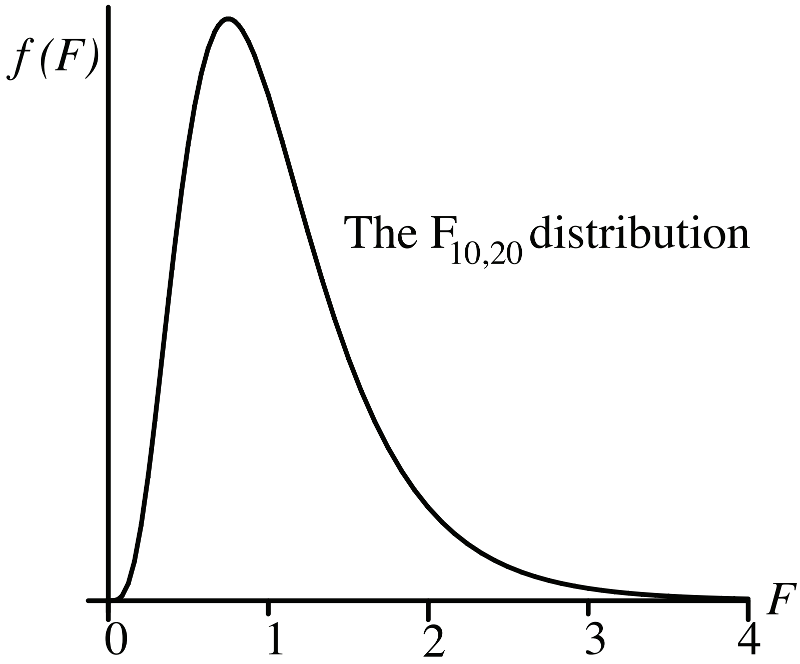

If $X_1 \sim \chi^2_{\nu_1}$ and $X_2 \sim \chi^2_{\nu_2}$ are independent random variables then

$$\frac{X_1/\nu_1}{X_2/\nu_2} \sim F_{\nu_1,\nu_2}$$

the F distribution with (ν_1, ν_2) degrees of freedom. This distribution is used, for example, for inference about the ratio of two variances, in Analysis of Variance (ANOVA) and in simple and multiple linear regression.