Confidence Intervals

- A confidence interval provides a simple summary of how precisely a parameter, denoted θ, is estimated.
- In many situations, a $(1-\alpha)100\%$ confidence interval is of the form

$$(\hat{\theta} - s_{\hat{\theta}} t_{\alpha/2}, \quad \hat{\theta} + s_{\hat{\theta}} t_{\alpha/2})$$

  where
  - $\hat{\theta}$ is an estimate of θ
  - $s_{\hat{\theta}}$ is its standard error
  - $t_{\alpha/2}$ is the upper $\alpha/2$th quantile from a distribution like the normal or t
- $s_{\hat{\theta}}$ is usually inversely proportional to the square root of the sample size, so the interval is narrower for larger samples
- $t_{\alpha/2}$ is larger for smaller α or larger confidence level, so a 99% confidence interval is wider than a 95% confidence interval
- the interval is constructed so that in advance there is probability $1-\alpha$ that it includes the true value of the parameter
- once we get the data and evaluate the interval endpoints we don't know whether or not the interval contains the true parameter
  - but we are confident that it does

- the figure below shows 95% confidence intervals for the mean constructed using 25 different random samples

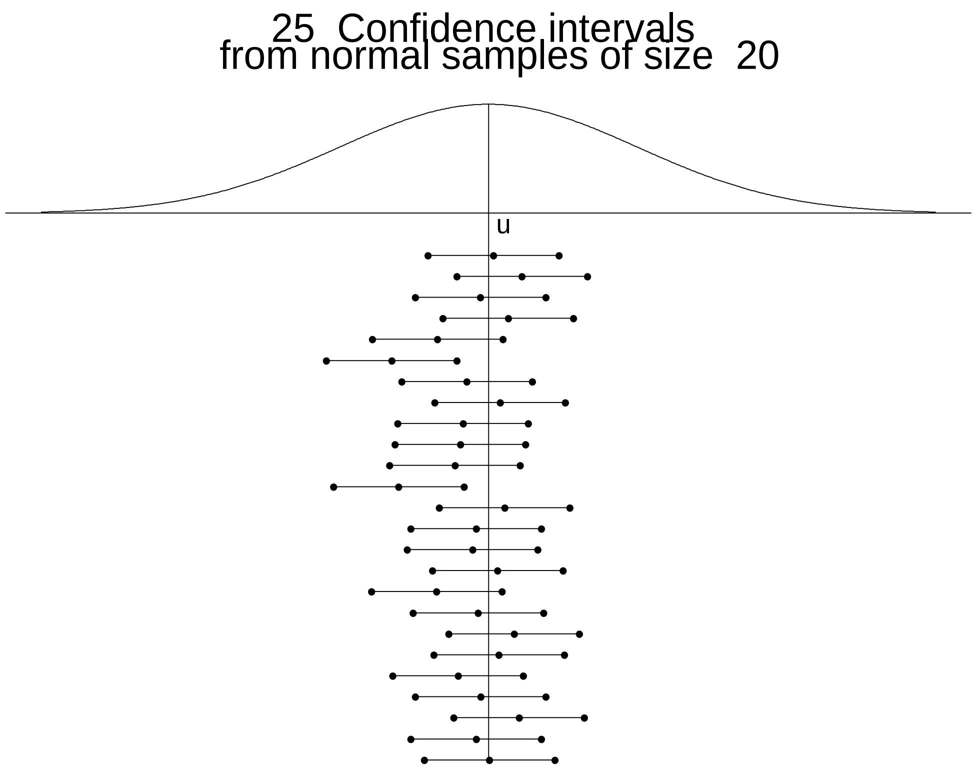

- most of these intervals do contain the true mean but two do not
- there is an important connection between confidence intervals and hypothesis

tests

- a $(1-\alpha)100\%$ confidence interval contains all values θ_0 which are not rejected in a test of $H_0 : \theta = \theta_0$ versus $H_a : \theta \neq \theta_0$ at level α