Confidence intervals for the mean (normally distributed population)

- Confidence interval for a population mean μ is an interval constructed around the sample mean $\bar{x}$ so we are reasonably sure that it contains the (unknown) population mean μ

- Any confidence interval is based on a confidence level

- The probability that the confidence interval will not contain the population mean μ is denoted by α

- The probability that the confidence interval will contain the population mean μ in repeated samples is denoted by $1 - \alpha$

- $(1 - \alpha) \times 100\%$ is called the *confidence level*. The confidence level is the success rate for the method

- A 95% confidence level is most commonly used

- In general, if the population variance σ^2 is known, for a sample size of n, the probability is $1 - \alpha$ that the (unknown) population mean μ is contained in the interval

$$\left[\bar{x} \pm z_{\alpha/2} \frac{\sigma}{\sqrt{n}} \right]$$

- See your text for extensions to other cases