

- Outline for Today
  - Review the Central Limit Theorem
    - Think-Pair-Share
  - Chapter 8
    - Section 8-2: Confidence Interval on the Mean of a Normal Distribution, Variance Known
      - Development of confidence intervals (z-distribution)
        - Example
      - Choice of sample size
      - One-sided confidence bounds
      - Large sample CI
        - Think-Pair-Share
    - Section 8-3: Confidence Interval on the Mean of a Normal Distribution, Variance Unknown
      - t-distribution
    - Section 8-4: Confidence Interval on the Variance and Standard Deviation of a Normal Distribution
      - chi-square distribution
- Learning Objectives
  - Be able to explain to a 6th grader what a confidence interval is.
  - Describe how the central limit theorem makes it possible to calculate large-sample confidence intervals (CI) for  $\mu$  without assuming a normal population distribution with a known variance. Specifically, explain why this still works when the population is non-normal.
  - Understand the appropriate situations to use a z-test, t-test, and chi-squared test.
- Readings/Videos/Websites
  - Chapter 8: Applied Prob. Stats: Followed the book fairly precisely. Go back into the chapters for more depth.
  - Central Limit Theorem
    - Video: [Introduction to the Central Limit Theorem](#)
    - Online simulator: [https://rpubs.com/RamiroJC/CLT\\_Slides](https://rpubs.com/RamiroJC/CLT_Slides)
  - Confidence Intervals
    - Videos:
      - [Introduction to Confidence Intervals](#)
      - [Intro to Confidence Intervals for One Mean \(Sigma Known\)](#)
      - And the rest of the JB statistics confidence interval series
  - Degrees of Freedom: <https://www.khanacademy.org/math/ap-statistics/quantitative-data-ap/measuring-spread-quantitative/v/review-and-intuition-why-we-divide-by-n-1-for-the-unbiased-sample-variance>
  - R Examples
    - Interval estimation (aka confidence intervals): <http://www.r-tutor.com/elementary-statistics/interval-estimation>
    - Calculating confidence intervals: <http://www.cyclismo.org/tutorial/R/confidence.html#calculating-a-confidence-interval-from-a-normal-distribution>
    - Hypothesis testing (for later on in HW 4): <http://www.r-tutor.com/elementary-statistics/hypothesis-testing>

- Helpful R functions:
  - `qqnorm()` : produces a normal QQ plot to test normality. If points fall on a line [see: `qqline()`], the distribution is normal.
    - Package: stats (built-in)
  - `z.test()` : creates confidence intervals and tests hypotheses for both one and two sample problems for a population that is normal
    - Package: BSDA (must download)
    - Documentation: <https://www.rdocumentation.org/packages/BSDA/versions/1.01/topics/z.test>
  - `t.test()` : performs one and two sample t-tests on vectors of data
    - Package: stats (built-in)
  - `qnorm()` : use to answer: What is the z-score of the pth quantile of the normal distribution. For CIs, we are generally concerned with the (1-alpha)th quantile.
    - E.g. For an alpha of 0.05 (95% CI)
 

```
qnorm(0.975) == 1.96
qnorm(0.025) == -1.96
```
    - Package: stats (built-in)
  - `qt()` : similar to `qnorm()` above but used when the variance is unknown.
    - Package: stats (built-in)
  - `qchisq()` : similar to `qnorm()` and `qt()` above but used for calculating CIs of variance instead of mean
    - Package: stats (built-in)
  - `power.t.test()` : compute the power of the one or two sample t-test
    - Package: stats (built-in)
    - Documentation: <https://stat.ethz.ch/R-manual/R-devel/library/stats/html/power.t.test.html>