Review for Exam II
Stat 205: Statistics for the Life Sciences

Tim Hanson, Ph.D.

University of South Carolina

Logistics

* Multiple choice, 35 questions. Bring a number 2 pencil. An extra pencil wouldn't hurt, or maybe a sharpener.
* Be at least 5 minutes early; we need to go over how to fill the scantron sheet out.
* You can bring one page (both sides) of notes.
* No hats, no phones.
* Exam II covers Chapters 6, 7, and 8.

6.3: Confidence interval for μ

* Data are generated $Y_1, \ldots, Y_n \sim N(\mu, \sigma^2)$.
* Use $Y_1, \ldots, Y_n$ to come up with plausible range for μ, called a confidence interval.
* A 95% confidence interval for μ is given by

$$\bar{y} \pm t_{0.025} SE_{\bar{Y}} \text{ where } SE_{\bar{Y}} = \frac{s}{\sqrt{n}}.$$

* If $n < 30$ then data need to be $\boxed{\ ?\ }$ Can check this with a $\boxed{\ ?\ }$.
* A 99% confidence interval is $\boxed{\ ?\ }$ than a 95% confidence interval.
* True or False: a confidence interval always contains the unknown μ.
* W.S. Gossett invented the t-distribution doing quality control for the $\boxed{\ ?\ }$ brewery.

6.3: Confidence interval for μ

* t-distribution is used because we estimate σ by s in $SE_{\bar{y}}$; t has fatter tails than normal.
* Probability of confidence interval covering μ is 95% before we conduct experiment. After experiment the interval either covers μ or not, we don't know which.
* After we conduct experiment and compute $\bar{Y} \pm t_{0.025} SE_{\bar{y}}$, we call refer to "confidence" instead of "probability."
* HW: 6.3.4, 6.3.5 (use R for both) .

6.7: Confidence interval for $\mu_1 - \mu_2$

* Now have *two random samples* from *two populations*:

$$\begin{array}{ll} \text{Population 1:} & \mu_1 \text{ and } \sigma_1 \\ \text{Population 2:} & \mu_2 \text{ and } \sigma_2 \end{array}$$

* Have sample statistics:

$$\begin{array}{ll} \text{Sample 1:} & \bar{y}_1 \text{ and } s_1 \text{ and } n_1 \\ \text{Sample 1:} & \bar{y}_2 \text{ and } s_2 \text{ and } n_2 \end{array}$$

* (p. 201) Standard error of $\bar{y}_1 - \bar{y}_2$ is

$$SE_{\bar{y}_1 - \bar{y}_2} = \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}.$$

Confidence interval for $\mu_1 - \mu_2$

* 95% confidence interval for $\mu_1 - \mu_2$ is given by

$$\bar{y}_1 - \bar{y}_2 \pm t_{0.025} SE_{\bar{y}_1 - \bar{y}_2}.$$

* The degrees of freedom for the t-distribution is (p. 206)

$$df = \frac{(s_1^2/n_1 + s_2^2/n_2)^2}{(s_1^2/n_1)^2/(n_1 - 1) + (s_2^2/n_2)^2/(n_2 - 1)}.$$

R will do the work for us.

* HW: 6.7.11, 6.7.12, 6.7.13, 6.7.14 (use R for all of these).

μ_1 & μ_2 versus $\bar{y}_1$ & $\bar{y}_2$

* We do not know μ_1 or μ_2. These are unknown population means.
* We *do know* the sample means $\bar{y}_1$ and $\bar{y}_2$.
* Don't write something like $\mu_1 = 142$ miles per hour.
* Write: $\mu_1 =$ population mean tennis ball serve speed using the new composite racquet, $\mu_2 =$ population mean tennis ball serve speed using the old-type racquet.

7.2: The t-test for $H_0 : \mu_1 = \mu_2$

* Initially consider $H_0 : \mu_1 = \mu_2$ versus $H_A : \mu_1 \neq \mu_2$.
* $t_s = \frac{\bar{y}_1 - \bar{y}_2}{SE_{\bar{y}_1 - \bar{y}_2}}$ is the *test statistic*.
* The p-value is $\Pr\{|T| \geq |t_s|\}$, where T is a student t random variable with degrees of freedon df given by the Welch-Satterthwaite formula on slide 6.
* The P-value will be computed for you. Recall that the P-value is the probability of seeing two sample means $\bar{Y}_1$ and $\bar{Y}_2$ *even further apart than what we saw* given that $H_0 : \mu_1 = \mu_2$ is true.
* Reject $H_0 : \mu_1 = \mu_2$ in favor of $H_A : \mu_1 \neq \mu_2$ if P-value $< \alpha$ (otherwise accept H_0). α is called the *significance level* of the test, usually $\alpha = 0.05$.
* HW: 7.2.3(a,b), 7.2.4(a,b), 7.2.9 (use R), 7.2.10 (use R), 7.2.11 (use R), 7.2.14, 7.2.17 (use R).

7.3: Confidence interval and t test

* Pages 234–235 explains the following *important* rule:
* **Reject** $H_0 : \mu_1 = \mu_2$ **in favor of** $H_A : \mu_1 \neq \mu_2$ **at the 5% level whenever a 95% confidence interval for** $\mu_1 - \mu_2$ **does not contain zero.**
* HW: 7.3.6, 7.3.7.

7.3: Type I and Type II errors

* Type I error is rejecting $H_0 : \mu_1 = \mu_2$ when H_0 **is true**.
* Type II error is *accepting* $H_0 : \mu_1 = \mu_2$ when H_0 **is false**.
* α is the probability of making a Type I error, usually 5%. This is called the significance level of the test.
* β is the probability of a Type II error. This number depends on the *true, unknown value of* $\mu_1 - \mu_2$.
* HW: 7.3.4, 7.3.5.

7.4: Association vs. causation

* When can we ascribe causality?
* A carefully controlled experiment creates two populations that are essentially identical except for an experimental manipulation (treatment vs. control). If we're careful, we can ascribe causality.
* An observational study simply collects some data and looks for association. Here, lurking variables, or unmeasured *confounders* may be *may be* driving any association that we see.
* HW: 7.4.1.

7.9: What a P-value is and isn't...

* P-value $\boxed{\ ?\ }$ the probability that H_0 is true.
* P-value $\boxed{\ ?\ }$ the probability of seeing a test statistic as extreme or more extreme than what we saw.
* HW: 7.9.1.

7.5: One-sided alternatives

* There are two one sided alternatives, use same t_s for both.
* $H_A : \mu_1 < \mu_2$ or
* $H_A : \mu_1 > \mu_2$.
* Top one has p-value $\Pr\{T < t_s\}$, Figure 7.5.1(a).
* Bottom one has p-value $\Pr\{T > t_s\}$, Figure 7.5.1(b).
* One-sided tests give you *more power* to reject $H_0 : \mu_1 = \mu_2$.
* HW: 7.5.6, 7.5.13.

7.7: Sample size calculation

* Often we need to know how much data to collect to reject $H_0 : \mu_1 = \mu_2$.
* Need power, α, whether the alternative is two-sided or one-sided, and estimates of $\mu_1 - \mu_2$ and $\sigma = \sigma_1 = \sigma_2$.
* Here's some R code

```
> power.t.test(delta=2,sd=0.8,sig.level=0.05,power=0.9,type="two.sample",alternative="one.sided")

     Two-sample t test power calculation

              n = 3.678026
          delta = 2
             sd = 0.8
      sig.level = 0.05
          power = 0.9
    alternative = one.sided

 NOTE: n is number in *each* group
```

* HW: 7.7.1, 7.7.3(a).

When is t-test valid?

* The t test assumes both populations are normal. Can check with normal probability plot.
* If sample sizes are "large enough" $n_1 \geq 30$ and $n_2 \geq 30$, then normality isn't important.
* If sample sizes are small we can instead use either a *permutation test* (Section 7.1) or the Wilcoxin-Mann-Whitney test (7.10).

7.10: Wilcoxin-Mann-Whitney test

* The t test assumes both populations are normal. Wilcoxin-Mann-Whitney test works when they are not.
* Called "nonparametric" because it does not assume that the population densities have a specific shape like the normal distribution.
* Null is H_0 : population densities are the same.
* Alternative is something like H_A : one population tends to be larger than the other, or H_A : soil respiration is tends to be be greater in the growth area vs. the gap area.
* Uses recipe on pp. 284–285; we use R to get P-value.
* HW: 7.10.3, 7.10.4, 7.10.6(a,b,c).

8.2: Paired t-test

* Have repeated measurements *on the same subject*.
* Often "before" and "after" type experiments, e.g. heart rate measured before and after exercise.
* Look at *differences* in measurement for each subject.
* Can get confidence interval for μ_D, the mean difference among the two treatments.
* Can also test $H_0 : \mu_D$ vs. one of (a) $H_A : \mu_D \neq 0$, (b) $H_A : \mu_D < 0$, or (c) $H_A : \mu_D > 0$.
* HW: 8.2.1(b), 8.2.4, 8.2.6.

8.4: Sign test

* Paired t-test requires either a large sample size or normal data.
* With small sample size and non-normal differences, can instead do the sign test.
* (1) Take differences, (2) count how many differences are positive N_+, (3) in R type `binom.test(`N_+`,`n`)` to get P-value.
* Tests $H_0 : \eta_D$ vs. one of (a) $H_A : \eta_D \neq 0$, (b) $H_A : \eta_D < 0$, or (c) $H_A : \eta_D > 0$, where η_D is median difference across treatments.
* HW: 8.4.5, 8.4.6(a).