Handout 4: Estimation, Hypothesis Testing & Confidence Intervals

EC 282: Introduction to Econometrics

Spring 2026

Instructions: Run the provided R code and answer the questions. Show your work for calculations.

1 Setup: The Data

A university wants to know the average monthly rent paid by its students. Since surveying all 30,000 students is impractical, we will draw samples and use the tools from this course—estimation, hypothesis testing, and confidence intervals—to learn about the population.

Run the code below to create the (hidden) population and draw a sample.

```
library(ggplot2)
set.seed(282)

# Population of monthly rents (unknown to us in practice)
N <- 30000
population <- data.frame(
  id = 1:N,
  rent = round(rnorm(N, mean = 950, sd = 300), 2)
)
pop_mean <- mean(population$rent)  # True mu_Y (pretend we don't know this)
pop_sd   <- sd(population$rent)

# Draw a random sample of n = 50 students
n <- 50
our_sample <- sample(population$rent, n)
```

Question 1.1: Compute the sample mean $\bar{Y}$, the sample standard deviation S_Y, and the standard error $SE[\bar{Y}]$.

```
y_bar <- mean(our_sample)
s_y   <- sd(our_sample)
se    <- s_y / sqrt(n)

cat("Sample mean (Y_bar):", round(y_bar, 2), "\n")
cat("Sample SD (S_Y):", round(s_y, 2), "\n")
cat("Standard error SE[Y_bar]:", round(se, 2), "\n")
```

Question 1.2: Explain the difference between σ_Y and S_Y, and between $\sigma_{\bar{Y}}$ and $SE[\bar{Y}]$. Why do we need the sample-based versions?

2 Estimators and Their Properties

Question 2.1: What is the difference between an *estimator* and an *estimate*? In the rent example, identify each.

Question 2.2: Three properties make an estimator "good": unbiasedness, consistency, and efficiency. For each property:

(a) Define the property mathematically.

(b) Explain what it means in plain language.

Question 2.3: Run the code below to compare two estimators of μ_Y: the sample mean $\bar{Y}$ and the sample median.

```
n_sims <- 10000

means   <- replicate(n_sims, mean(sample(population$rent, 50)))
medians <- replicate(n_sims, median(sample(population$rent, 50)))

cat("--- Sample Mean ---\n")
cat("E[Y_bar] ~=", round(mean(means), 2), "\n")
cat("Var(Y_bar) ~=", round(var(means), 2), "\n")

cat("\n--- Sample Median ---\n")
cat("E[Median] ~=", round(mean(medians), 2), "\n")
cat("Var(Median) ~=", round(var(medians), 2), "\n")

cat("\nPopulation mean:", round(pop_mean, 2), "\n")
```

(a) Are both estimators unbiased?

(b) Which estimator is more efficient? What does this mean for a researcher choosing between them?

Question 2.4: Run the code below. What property of $\bar{Y}$ does this demonstrate?

```
sample_sizes <- c(10, 50, 200, 1000, 5000)

consistency <- data.frame(
  n = rep(sample_sizes, each = 500),
  y_bar = unlist(lapply(sample_sizes, function(n) {
    replicate(500, mean(sample(population$rent, n)))
  }))
)

ggplot(consistency, aes(x = y_bar)) +
  geom_histogram(bins = 40, fill = "steelblue",
                 color = "white", alpha = 0.7) +
  geom_vline(xintercept = pop_mean, color = "red",
             linetype = "dashed") +
  facet_wrap(~ paste("n =", n), scales = "free_y") +
  labs(title = "Sampling Distribution of Y-bar for Different n",
       x = "Sample Mean", y = "Count") +
  theme_minimal()
```

3 Hypothesis Testing

The university currently budgets for student housing based on an assumed average rent of $900 per month. The housing office suspects rents may have increased. We use our sample to test this.

Question 3.1: Write down the null and alternative hypotheses. Is this a one-sided or two-sided test? Explain your choice.

Question 3.2: Compute the t-statistic using the formula:

$$t\text{-stat} = \frac{\bar{Y} - \mu_{Y,0}}{SE[\bar{Y}]}$$

where $\mu_{Y,0} = 900$.

```
mu_0   <- 900
t_stat <- (y_bar - mu_0) / se
cat("t-statistic:", round(t_stat, 4), "\n")
```

Show the calculation by hand, then verify with the R output.

Question 3.3: Compute the p-value for:

(a) A two-sided test ($H_A : \mu_Y \neq 900$)

(b) A one-sided test ($H_A : \mu_Y > 900$)

```
# Two-sided p-value
p_two <- 2 * pnorm(-abs(t_stat))
cat("Two-sided p-value:", round(p_two, 4), "\n")

# One-sided p-value (right tail)
p_one <- 1 - pnorm(t_stat)
cat("One-sided p-value:", round(p_one, 4), "\n")
```

Hint: For the two-sided test, p-value $= 2 \times \Phi(-|t\text{-stat}|)$. For the one-sided test ($H_A : \mu_Y > \mu_{Y,0}$), p-value $= 1 - \Phi(t\text{-stat})$.

Question 3.4: At the $\alpha = 0.05$ significance level, what is your decision? What about at $\alpha = 0.01$? Explain.

Question 3.5: Run the code below to visualize where our test statistic falls under the null distribution:

```
x_seq <- seq(-4, 4, length.out = 300)
df <- data.frame(x = x_seq, y = dnorm(x_seq))

ggplot(df, aes(x = x, y = y)) +
  geom_line(linewidth = 1) +
  geom_area(data = subset(df, x <= -1.96),
            fill = "red", alpha = 0.3) +
  geom_area(data = subset(df, x >= 1.96),
            fill = "red", alpha = 0.3) +
  geom_vline(xintercept = t_stat, color = "blue",
             linetype = "dashed", linewidth = 1) +
  annotate("text", x = t_stat + 0.3, y = 0.35,
           label = paste("t =", round(t_stat, 2)),
           color = "blue", hjust = 0) +
  labs(title = "Standard Normal Distribution Under H0",
       subtitle = "Red = rejection region (alpha = 0.05, two-sided)",
       x = "Z", y = "Density") +
  theme_minimal()
```

Where does our t-statistic fall relative to the rejection region? Is this consistent with your decision?

4 Type I and Type II Errors

Question 4.1: Complete the table below with descriptions of each outcome:

	Fail to Reject H_0	**Reject H_0**
H_0 True		
H_0 False		

Question 4.2: In our rent example, describe in plain language what a Type I error and a Type II error would mean for the university housing office.

Question 4.3: Run the code below to simulate the Type I error rate. We draw 10,000 samples from a population where $\mu_Y = 900$ (i.e., H_0 is true) and count how often we incorrectly reject.

```
pop_null <- rnorm(N, mean = 900, sd = 300)

rejections <- replicate(10000, {
  samp <- sample(pop_null, 50)
  t <- (mean(samp) - 900) / (sd(samp) / sqrt(50))
  abs(t) > 1.96
})

cat("Type I error rate:", mean(rejections), "\n")
cat("Nominal alpha:", 0.05, "\n")
```

Is the simulated Type I error rate close to $\alpha = 0.05$? Explain why this makes sense.

5 Confidence Intervals

Question 5.1: Using your sample, construct a 95% confidence interval for μ_Y:

$$\bar{Y} \pm 1.96 \times SE[\bar{Y}]$$

```
ci_lower <- y_bar - 1.96 * se
ci_upper <- y_bar + 1.96 * se
cat("95% CI: [", round(ci_lower, 2), ",", round(ci_upper, 2), "]\n")
```

Show the calculation by hand, then verify with R.

Question 5.2: Construct a 90% and a 99% confidence interval. How do they compare to the 95% CI?

```
cat("90% CI: [", round(y_bar - 1.65 * se, 2), ",",
    round(y_bar + 1.65 * se, 2), "]\n")

cat("95% CI: [", round(y_bar - 1.96 * se, 2), ",",
    round(y_bar + 1.96 * se, 2), "]\n")

cat("99% CI: [", round(y_bar - 2.576 * se, 2), ",",
    round(y_bar + 2.576 * se, 2), "]\n")
```

What is the trade-off between confidence level and the width of the interval?

Question 5.3: Interpret the 95% confidence interval in plain language. Which of the following is the correct interpretation?

(a) There is a 95% probability that μ_Y lies in our interval.

(b) If we drew many samples and computed a 95% CI from each, about 95% of those intervals would contain the true μ_Y.

(c) 95% of students pay rents within this interval.

Question 5.4: Run the code below to visualize the coverage property. We draw 100 samples, compute a 95% CI from each, and check which ones contain μ_Y:

```r
ci_data <- do.call(rbind, lapply(1:100, function(i) {
  samp <- sample(population$rent, 50)
  m <- mean(samp)
  s <- sd(samp) / sqrt(50)
  data.frame(
    sample = i, y_bar = m,
    lower = m - 1.96 * s, upper = m + 1.96 * s,
    covers = (m - 1.96 * s <= pop_mean) &
             (pop_mean <= m + 1.96 * s)
  )
}))

ggplot(ci_data, aes(x = sample, y = y_bar, color = covers)) +
  geom_point(size = 1.5) +
  geom_errorbar(aes(ymin = lower, ymax = upper), width = 0.3) +
  geom_hline(yintercept = pop_mean, linetype = "dashed",
             color = "black") +
  scale_color_manual(values = c("TRUE" = "steelblue",
                                "FALSE" = "red")) +
  labs(title = "95% Confidence Intervals from 100 Samples",
       subtitle = "Red intervals miss the true mean",
       x = "Sample Number", y = "Rent ($)",
       color = "Contains mu") +
  theme_minimal() +
  theme(legend.position = "bottom")

cat("Coverage rate:", mean(ci_data$covers), "\n")
```

How many of the 100 intervals contain the true μ_Y? Is this close to what you expected?

6 The Connection Between CIs and Hypothesis Tests

Question 6.1: There is a deep connection between confidence intervals and hypothesis testing. Using your 95% CI from Question 5.1, can you determine whether you would reject $H_0 : \mu_Y = 900$ at the 5% significance level *without computing a t-statistic*? Explain.

Question 6.2: What about $H_0 : \mu_Y = 950$? Would you reject at $\alpha = 0.05$? Check using both the CI and the t-test:

```r
# CI check
cat("95% CI: [", round(ci_lower, 2), ",", round(ci_upper, 2), "]\n")
cat("Is 950 inside the CI?",
    ci_lower <= 950 & 950 <= ci_upper, "\n")
```

```

# t-test check
t_950 <- (y_bar - 950) / se
p_950 <- 2 * pnorm(-abs(t_950))
cat("t-stat (H0: mu = 950):", round(t_950, 4), "\n")
cat("p-value:", round(p_950, 4), "\n")
```

State the general rule connecting the 95% CI and the two-sided test at $\alpha = 0.05$.

7 Comparing Two Groups

The housing office also wants to know whether students living off-campus pay different rents than students living in university-adjacent housing. Run the code below to set up the two groups.

```
# Two populations
off_campus <- rnorm(15000, mean = 1020, sd = 280)
adjacent   <- rnorm(15000, mean = 950, sd = 250)

# Draw samples
n1 <- 60; n2 <- 55
samp1 <- sample(off_campus, n1)
samp2 <- sample(adjacent, n2)

y_bar1 <- mean(samp1);  s1 <- sd(samp1)
y_bar2 <- mean(samp2);  s2 <- sd(samp2)
diff   <- y_bar1 - y_bar2
se_diff <- sqrt(s1^2/n1 + s2^2/n2)

cat("Off-campus:  Y_bar1 =", round(y_bar1,2),
    " S1 =", round(s1,2), " n1 =", n1, "\n")
cat("Adjacent:    Y_bar2 =", round(y_bar2,2),
    " S2 =", round(s2,2), " n2 =", n2, "\n")
cat("Difference:  ", round(diff,2), "\n")
cat("SE(diff):    ", round(se_diff,2), "\n")
```

Question 7.1: Write down the null and alternative hypotheses for testing whether there is a difference in average rent between the two groups.

Question 7.2: Compute the t-statistic and p-value:

$$t\text{-stat} = \frac{(\bar{Y}_1 - \bar{Y}_2) - 0}{SE[\bar{Y}_1 - \bar{Y}_2]} \qquad \text{where } SE[\bar{Y}_1 - \bar{Y}_2] = \sqrt{\frac{S_1^2}{n_1} + \frac{S_2^2}{n_2}}$$

```
t_diff <- diff / se_diff
p_diff <- 2 * pnorm(-abs(t_diff))
cat("t-statistic:", round(t_diff, 4), "\n")
cat("p-value:", round(p_diff, 4), "\n")
```

Question 7.3: Construct a 95% confidence interval for the difference in means $\mu_1 - \mu_2$.

```
ci_diff_l <- diff - 1.96 * se_diff
ci_diff_u <- diff + 1.96 * se_diff
cat("95% CI for difference: [", round(ci_diff_l,2), ",",
    round(ci_diff_u,2), "]\n")
```

Does the interval contain 0? Is this consistent with your hypothesis test result?

Question 7.4: Interpret your findings. What should the housing office conclude about the difference in rents between off-campus and university-adjacent housing?