

Lecture 06

Statistics II

STAT-221 • UG

Dr. Rizwan Jahangir
NUST Business School
National University of Sciences & Technology

Contents

- **1. Warm-up & Review:**
  - The Reality of Sampling.
  - The Student's t-Distribution.
- **2. One-Sample Tests (Unknown σ):**
  - The t-Statistic Formula.
  - Example 1: McFarland Insurance Claims.
  - Example 2: Digital Clock Batteries.
- **3. Break (15 Minutes)**
- **4. Introduction to Two-Sample Tests:**
  - Why Compare Two Populations?
  - Independent vs. Dependent Samples.
- **5. Comparing Two Independent Means (Unknown σ):**
  - The Pooled Variance (s_p^2).
  - Example 3: Schadek vs. Bowyer.



Warm-up & Review

The Reality of Sampling

In Lecture 4 and 5, we learned the Six-Step Procedure for testing a hypothesis about a population mean.

- We used the standard normal distribution (the z-statistic).
- **However**, using the z-statistic requires us to know the population standard deviation (σ).

The Real World: As we discussed in Chapter 9, if we do not know the population mean (μ), it is highly unlikely we would somehow know the population standard deviation (σ)!



Bringing Back the t-Distribution

When the population standard deviation (σ) is **unknown**, we must substitute it with the sample standard deviation (s). Because s is only an estimate, it introduces extra uncertainty. We account for this by using the **Student's t-Distribution**.

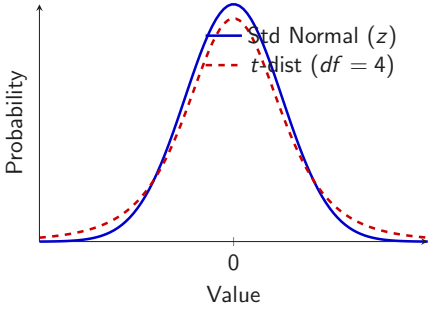
Characteristics of the t-Distribution:

- It is continuous, bell-shaped, and symmetrical about zero.
- It is **flatter** and more spread out than the z-distribution.
- As the sample size gets larger, it looks more and more like the z-distribution.



Visualizing the t-Distribution

- Notice how the t-distribution has "thicker tails" than the standard normal (z) curve.
- Because there is more area in the tails, we have to go further out from the mean to capture the same level of confidence or to reach our rejection region.





Degrees of Freedom (df)

Unlike the standard normal distribution (which is just one single curve), the t-distribution is actually a *family* of curves.

- The exact shape depends on the **Degrees of Freedom (df)**.
- For a one-sample test of a mean, the degrees of freedom are calculated as:

$$df = n - 1$$

Where n is the sample size.



Pop-Up Question 1: z vs. t

Discussion

Pop-Up Question!

Suppose we are conducting a two-tailed hypothesis test at the $\alpha = 0.05$ significance level.

1. If we know σ, what is the critical z-value?
2. If we do *not* know σ and our sample size is $n = 15$, what is the critical t-value?

Check your tables!

Answers:

- **1:** The critical z-value is ± 1.96.
- **2:** For $df = 14$ $(15 - 1)$ at $\alpha = 0.05$ (Two-Tailed), the critical t-value is ± 2.145. Notice we have to go further out (± 2.145) than the z-distribution (± 1.96) to find the boundary of our 5% risk!



One-Sample Tests (Unknown σ)

The t-Statistic Formula

When conducting a hypothesis test for a population mean where σ is unknown, we use the t-statistic in Step 3 of our procedure.

$$t = \frac{\bar{x} - \mu}{s/\sqrt{n}}$$

Where:

- $\bar{x}$ = Sample Mean.
- μ = Hypothesized Population Mean (from H_0).
- s = Sample Standard Deviation.
- n = Sample Size.



Example 1: McFarland Insurance Claims

Textbook Example: McFarland Insurance
The McFarland Insurance Company Claims Department reports the mean cost to process a claim is \$60. To evaluate recent cost-cutting measures, the supervisor selected a random sample of **26 claims** ($n = 26$) processed last month.

The sample mean ($\bar{x}$) was **\$56.42** with a sample standard deviation (s) of **\$10.04**.

Task: At the **.05 significance level**, has the cost-cutting measure successfully reduced the mean cost to process a claim?



McFarland Insurance: Steps 1 & 2

Step 1: State the Hypotheses.

The supervisor wants to prove the cost is **less than** \$60. This is a **Left-Tailed Test**.

- $H_0 : \mu \geq 60$
- $H_1 : \mu < 60$

Step 2: Select the Level of Significance.

- $\alpha = 0.05$



McFarland Insurance: Steps 3 & 4

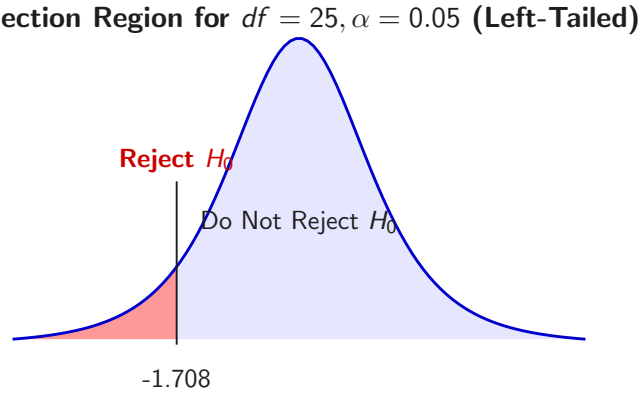
Step 3: Select the Test Statistic.
Because we only have the *sample* standard deviation ($s = 10.04$), we must use the ***t*-distribution**.

Step 4: Formulate the Decision Rule.

- Degrees of freedom: $df = n - 1 = 26 - 1 = \mathbf{25}$.
- One-Tailed test, $\alpha = 0.05$.
- Look at the t-table (Appendix B.5): The critical value is -1.708.
- **Decision Rule:** Reject H_0 if our computed $t < -1.708$.



McFarland Insurance: Visualizing the Decision Rule



McFarland Insurance: Step 5

Step 5: Make a Decision.

Compute the t-statistic using our sample data:

$$t = \frac{\bar{x} - \mu}{s/\sqrt{n}}$$

$$t = \frac{56.42 - 60.00}{10.04/\sqrt{26}}$$

$$t = \frac{-3.58}{1.969} = \textbf{-1.818}$$

Mathematical Decision: Since -1.818 is less than our critical limit of -1.708, it falls in the red rejection region. We **Reject** H_0.

McFarland Insurance: Step 6

Step 6: Interpret the Result.

What does this mean for the business?

- We rejected the null hypothesis that the mean cost is \$60 or more.
- The sample evidence is strong enough to conclude that the mean cost is now significantly *less* than \$60.
- **Conclusion:** The supervisor can confidently report to management that the new cost-cutting measures are working!



Example 2: Digital Clock Batteries

Textbook Example: Battery Life
The mean life of a battery used in a digital clock is 305 days. The battery was recently modified to last longer. A sample of **20 modified batteries** ($n = 20$) had a mean life ($\bar{x}$) of **311 days** with a standard deviation (s) of **12 days**.

Task: Did the modification significantly increase the mean life of the battery? Use the **.05 significance level**.

Group Task: Work through the Six-Step procedure with your neighbor!



Digital Clock Batteries: Solution

1 & 2. Hypotheses & Sig Level:

- $H_0 : \mu \leq 305$
- $H_1 : \mu > 305$ (Right-Tailed)
- $\alpha = 0.05$

3 & 4. Statistic & Decision Rule:

- Use t-distribution (s is used).
- $df = 20 - 1 = \mathbf{19}$.
- Critical t-value = **1.729**.
- Reject H_0 if computed $t > 1.729$.

5. Make Decision:

$$t = \frac{311 - 305}{12/\sqrt{20}} = \frac{6}{2.683} = \mathbf{2.236}$$

6. Interpretation:

Because $2.236 > 1.729$, we **Reject** H_0. The data proves the modification successfully increased the battery life beyond 305 days!



15 Minute Break

Introduction to Two-Sample Tests

Why Compare Two Populations?

Until now, we have tested a single sample against a known or claimed population target (e.g., "Is the mean salary equal to \$45,000?").

In Chapter 11, we shift our focus to comparing **two separate populations** to see if they are different from each other.

Real-World Questions:

- Is there a difference in the mean value of residential real estate sold by male agents vs. female agents?
- Do young workers have more absences than older workers?
- Does a new machine assemble parts faster than an old machine?



Independent vs. Dependent Samples

Before applying a formula, you must ask: Are the two samples connected?

Independent Samples:

- The samples are drawn from two completely unrelated groups.
- Example: Comparing the salaries of nurses vs. teachers. The nurses have nothing to do with the teachers.
- *This is our focus for today.*

Dependent Samples:

- The samples are paired or matched. Often, it's the exact same person measured "before and after" an intervention.
- Example: Weighing people before starting a diet, and weighing those *exact same people* after.
- *We will cover this in Lecture 7.*



Pop-Up Question 2: Independent or Dependent?

Discussion

Pop-Up Question!

Determine if the following scenarios require an Independent or Dependent two-sample test:

1. You compare the fuel efficiency (MPG) of 15 Honda Civics and 18 Toyota Corollas.

2. You measure the typing speed of 10 secretaries, send them to a typing workshop, and measure the typing speed of those same 10 secretaries again.

Answers:

- **1: Independent**. The Hondas and Toyotas are completely separate entities.
- **2: Dependent**. The data points are paired to specific individuals (before and after the workshop).



Comparing Two Independent Means

Testing Two Independent Means (Unknown σ)

We want to test if the mean of Population 1 (μ_1) is equal to the mean of Population 2 (μ_2).

- Hypotheses generally look like: $H_0 : \mu_1 = \mu_2$ vs. $H_1 : \mu_1 \neq \mu_2$

Just like one-sample tests, we rarely know the true population standard deviations (σ_1 and σ_2). We must rely on sample standard deviations (s_1 and s_2).

To use the t-distribution for two samples, we must make an assumption: **The two populations have equal (but unknown) variances.**



The Pooled Variance (s_p^2)

Because we assume the two populations have the same variance, we combine (or "pool") our two sample variances (s_1^2 and s_2^2) into a single, better estimate of that shared variance.

Pooled Variance Formula:

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

Notice how the formula acts as a weighted average. If sample 1 is much larger than sample 2 ($n_1 > n_2$), the pooled variance will lean closer to s_1^2.



The Two-Sample t-Statistic

Once we have our pooled variance, we plug it into our t-statistic formula to compare our sample means ($\bar{x}_1$ and $\bar{x}_2$):

Two-Sample t-Statistic:

$$t = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{s_p^2 \left(\frac{1}{n_1} + \frac{1}{n_2} \right)}}$$

Degrees of Freedom:

- For two independent samples, the degrees of freedom are $n_1 + n_2 - 2$.



Example 3: Schadek vs. Bowyer

Textbook Example: Comparing Home Appraisals

We want to investigate if there is a difference in the mean appraised values of homes done by two different appraisers: Schadek and Bowyer.

A sample of 10 homes appraised by **Schadek** showed a mean of **226.8** ($000s) with a sample variance (s_1^2) of **208.844**.

A sample of 10 homes appraised by **Bowyer** showed a mean of **222.2** ($000s) with a sample variance (s_2^2) of **204.178**.

Task: Use the **.05 significance level** to test if there is a difference in the mean appraisals.



Schadek vs. Bowyer: Steps 1 & 2

Step 1: State the Hypotheses.
We are checking if there is *any* difference (not specifically greater or less than). This is a **Two-Tailed Test**.

- $H_0 : \mu_{\text{Schadek}} = \mu_{\text{Bowyer}}$
- $H_1 : \mu_{\text{Schadek}} \neq \mu_{\text{Bowyer}}$

Step 2: Select the Level of Significance.

- $\alpha = 0.05$



Schadek vs. Bowyer: Steps 3 & 4

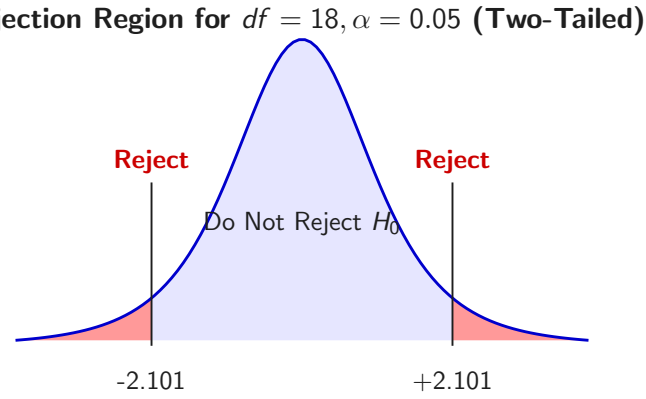
Step 3: Select the Test Statistic.
We use the Two-Sample t-statistic.

Step 4: Formulate the Decision Rule.

- Degrees of freedom: $df = n_1 + n_2 - 2 = 10 + 10 - 2 = \mathbf{18}$.
- Two-Tailed test, $\alpha = 0.05$.
- Look at the t-table (Appendix B.5): The critical values are ± 2.101.
- **Decision Rule:** Reject H_0 if $t < -2.101$ or $t > 2.101$.



Schadek vs. Bowyer: Visualizing the Decision Rule



Schadek vs. Bowyer: Calculating Pooled Variance

Step 5: Make a Decision (Part 1 - Pooled Variance).
First, we must combine the sample variances:

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

$$s_p^2 = \frac{(10 - 1)(208.844) + (10 - 1)(204.178)}{10 + 10 - 2}$$

$$s_p^2 = \frac{9(208.844) + 9(204.178)}{18}$$

$$s_p^2 = \frac{1879.596 + 1837.602}{18} = \frac{3717.198}{18} = \mathbf{206.511}$$



Schadek vs. Bowyer: Calculating t-Statistic

Step 5: Make a Decision (Part 2 - t-Statistic).
Now plug $s_p^2 = 206.511$ into the final formula:

$$t = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{s_p^2\left(\frac{1}{n_1} + \frac{1}{n_2}\right)}}$$

$$t = \frac{226.8 - 222.2}{\sqrt{206.511\left(\frac{1}{10} + \frac{1}{10}\right)}}$$

$$t = \frac{4.6}{\sqrt{206.511(0.20)}} = \frac{4.6}{\sqrt{41.302}}$$

$$t = \frac{4.6}{6.426} = \mathbf{0.716}$$



Schadek vs. Bowyer: Final Decision & Interpretation

Step 6: Interpret the Result.

What does $t = 0.716$ tell us?

- Our computed value of 0.716 falls between our critical limits of -2.101 and 2.101. It is in the "Do Not Reject" region.
- We **Do Not Reject** H_0.
- **Business Conclusion:** There is no statistically significant difference in the mean appraisals given by Schadek and Bowyer. The \$4,600 difference observed between their sample means ($\$226,800 - \$222,200$) is likely just due to sampling error.



Software Application: Excel Output

Software instantly computes pooled variance, the t-statistic, and the exact p-value! Compare our manual calculations with this Excel output for the same problem.

t-Test: Two-Sample Assuming Equal Variances		
	Schadek	Bowyer
Mean	226.800	222.200
Variance	208.844	204.178
Observations	10.000	10.000
Pooled Variance	**206.511**	
Hypothesized Mean Difference	0.000	
df	18.000	
t Stat	**0.716**	
P(T¡=t) one-tail	0.242	
t Critical one-tail	1.734	
P(T¡=t) two-tail	**0.483**	
t Critical two-tail	**2.101**	

