ISE 315: Engineering Statistics

Lecture 18: Simple Linear Regression (Part 2) — Inference and ANOVA

Instructor: Mansur M. Arief, PhD
Industrial and Systems Engineering, KFUPM

Office: 22-219 — Email: mansur.arief@kfupm.edu.sa

Based on Montgomery & Runger, Applied Statistics and Probability for Engineers, 6th Ed.

Lecture 18

Simple Linear Regression — Part 2: Inference on the Slope and Intercept, and the ANOVA Table (Chapter 6, cont'd)

Lecture 18 Outline

- Quick review: the truck fuel fit (from Lecture 17)

Lecture 18 Outline

- Quick review: the truck fuel fit (from Lecture 17)
- Is the relationship real? The t-test on the slope: Procedure 6.2

Lecture 18 Outline

- Quick review: the truck fuel fit (from Lecture 17)
- Is the relationship real? The t-test on the slope: Procedure 6.2
- Confidence intervals for β_1 and β_0 ((6.23), (6.24))

Lecture 18 Outline

- Quick review: the truck fuel fit (from Lecture 17)
- Is the relationship real? The t-test on the slope: Procedure 6.2
- Confidence intervals for β_1 and β_0 ((6.23), (6.24))
- ANOVA decomposition for regression ((6.19)) and the F-test ((6.20))

Lecture 18 Outline

- Quick review: the truck fuel fit (from Lecture 17)
- Is the relationship real? The t-test on the slope: Procedure 6.2
- Confidence intervals for β_1 and β_0 ((6.23), (6.24))
- ANOVA decomposition for regression ((6.19)) and the F-test ((6.20))
- Exam-style practice problem (AI safety red-team data)

Quick Review

What We Did in Lecture 17

Review: The Truck Fuel Fit

Load x (tonnes) vs. fuel y (L/100 km), $n = 8$ trips

Last lecture we ran Procedure 6.1 on the fleet analyst's data, start to finish:

Review: The Truck Fuel Fit

Load x (tonnes) vs. fuel y (L/100 km), $n = 8$ trips

Last lecture we ran Procedure 6.1 on the fleet analyst's data, start to finish:

Quantity	Book eq.	Value	Meaning
$\bar{x}$, $\bar{y}$	—	13, 30.65	sample means
S_{xx}	(6.7)	168	spread of x
S_{xy}	(6.8)	188.8	co-movement
S_{yy}	(6.12)	217.48	total spread of y
$\hat{\beta}_1$	(6.9)	1.124	L/100 km per extra tonne
$\hat{\beta}_0$	(6.10)	16.04	empty-truck anchor
SS_E	(6.13)	5.30	unexplained scatter
$\hat{\sigma}^2$	(6.14)	0.884	error variance ($n - 2 = 6$ df)

Review: The Truck Fuel Fit

Load x (tonnes) vs. fuel y (L/100 km), $n = 8$ trips

Last lecture we ran Procedure 6.1 on the fleet analyst's data, start to finish:

Quantity	**Book eq.**	**Value**	**Meaning**
$\bar{x}$, $\bar{y}$	—	13, 30.65	sample means
S_{xx}	(6.7)	168	spread of x
S_{xy}	(6.8)	188.8	co-movement
S_{yy}	(6.12)	217.48	total spread of y
$\hat{\beta}_1$	(6.9)	1.124	L/100 km per extra tonne
$\hat{\beta}_0$	(6.10)	16.04	empty-truck anchor
SS_E	(6.13)	5.30	unexplained scatter
$\hat{\sigma}^2$	(6.14)	0.884	error variance ($n - 2 = 6$ df)

The Natural Next Question

We have a fitted line: $\hat{y} = 16.04 + 1.124\,x$.

The Natural Next Question

We have a fitted line: $\hat{y} = 16.04 + 1.124\,x$.

But **fitting a line is always possible** — least squares produces a slope even when x and y are pure noise.

The Natural Next Question

We have a fitted line: $\hat{y} = 16.04 + 1.124\,x$.

But **fitting a line is always possible** — least squares produces a slope even when x and y are pure noise.

With only $n = 8$ trips, could a slope of 1.124 appear **by chance** even if the true β_1 were zero?

The Natural Next Question

We have a fitted line: $\hat{y} = 16.04 + 1.124\,x$.

But **fitting a line is always possible** — least squares produces a slope even when x and y are pure noise.

With only $n = 8$ trips, could a slope of 1.124 appear **by chance** even if the true β_1 were zero?

This is exactly the hypothesis testing question from Chapter 4 — new parameter, same machinery.

Same Machinery, New Parameter

Chapter 4 vs. Chapter 6

Step	**One-sample t (Ch. 4)**	**Slope test (Ch. 6)**
Parameter	μ	β_1
Hypothesized value	μ_0	$\beta_{1,0}$ (usually 0)
Estimator	$\bar{X}$	$\hat{\beta}_1$
Standard error	$s/\sqrt{n}$	$\sqrt{\hat{\sigma}^2/S_{xx}}$ (6.15)
Test statistic	T_0 (4.6)	T_0 (6.18)
Distribution under H_0	t_{n-1}	t_{n-2}

Same Machinery, New Parameter

Chapter 4 vs. Chapter 6

Step	**One-sample t (Ch. 4)**	**Slope test (Ch. 6)**
Parameter	μ	β_1
Hypothesized value	μ_0	$\beta_{1,0}$ (usually 0)
Estimator	$\bar{X}$	$\hat{\beta}_1$
Standard error	$s/\sqrt{n}$	$\sqrt{\hat{\sigma}^2/S_{xx}}$ (6.15)
Test statistic	T_0 (4.6)	T_0 (6.18)
Distribution under H_0	t_{n-1}	t_{n-2}

The structure is **identical**. Only three things change: the parameter, the standard error formula, and the degrees of freedom ($n-2$, because we estimated two coefficients).

Standard Errors of the Estimators

The estimated standard error of the **slope** — equation (6.15):

$$se(\hat{\beta}_1) = \sqrt{\frac{\hat{\sigma}^2}{S_{xx}}} \tag{6.15}$$

Standard Errors of the Estimators

The estimated standard error of the **slope** — equation (6.15):

$$se(\hat{\beta}_1) = \sqrt{\frac{\hat{\sigma}^2}{S_{xx}}} \tag{6.15}$$

The estimated standard error of the **intercept** — equation (6.16):

$$se(\hat{\beta}_0) = \sqrt{\hat{\sigma}^2 \left(\frac{1}{n} + \frac{\bar{x}^2}{S_{xx}} \right)} \tag{6.16}$$

Standard Errors of the Estimators

The estimated standard error of the **slope** — equation (6.15):

$$se(\hat{\beta}_1) = \sqrt{\frac{\hat{\sigma}^2}{S_{xx}}} \tag{6.15}$$

The estimated standard error of the **intercept** — equation (6.16):

$$se(\hat{\beta}_0) = \sqrt{\hat{\sigma}^2 \left(\frac{1}{n} + \frac{\bar{x}^2}{S_{xx}} \right)} \tag{6.16}$$

Read the design lesson in (6.15): a larger S_{xx} (more spread in x) makes the slope estimate more precise. Space out your x values when you design the study.

Standard Errors of the Estimators

The estimated standard error of the **slope** — equation (6.15):

$$se(\hat{\beta}_1) = \sqrt{\frac{\hat{\sigma}^2}{S_{xx}}} \tag{6.15}$$

The estimated standard error of the **intercept** — equation (6.16):

$$se(\hat{\beta}_0) = \sqrt{\hat{\sigma}^2 \left(\frac{1}{n} + \frac{\bar{x}^2}{S_{xx}} \right)} \tag{6.16}$$

Read the design lesson in (6.15): a larger S_{xx} (more spread in x) makes the slope estimate more precise. Space out your x values when you design the study.

Textbook: §6.5 and §6.7 ((6.15) onward)

Hypothesis Test on the Slope β_1

Procedure 6.2

Hypotheses — equation (6.17):

$$H_0\colon \beta_1 = \beta_{1,0} \qquad H_1\colon \beta_1 \neq \beta_{1,0} \tag{6.17}$$

The default choice $\beta_{1,0} = 0$ asks: *is there any linear relationship at all?*

Hypothesis Test on the Slope β_1

Procedure 6.2

Hypotheses — equation (6.17):

$$H_0\colon \beta_1 = \beta_{1,0} \qquad H_1\colon \beta_1 \neq \beta_{1,0} \tag{6.17}$$

The default choice $\beta_{1,0} = 0$ asks: *is there any linear relationship at all?*
Test statistic — equation (6.18):

$$T_0 = \frac{\hat{\beta}_1 - \beta_{1,0}}{\sqrt{\hat{\sigma}^2 / S_{xx}}} \tag{6.18}$$

Under H_0, $T_0 \sim t_{n-2}$.

Hypothesis Test on the Slope β_1

Procedure 6.2

Hypotheses — equation (6.17):

$$H_0\colon \beta_1 = \beta_{1,0} \qquad H_1\colon \beta_1 \neq \beta_{1,0} \tag{6.17}$$

The default choice $\beta_{1,0} = 0$ asks: *is there any linear relationship at all?*
Test statistic — equation (6.18):

$$T_0 = \frac{\hat{\beta}_1 - \beta_{1,0}}{\sqrt{\hat{\sigma}^2 / S_{xx}}} \tag{6.18}$$

Under H_0, $T_0 \sim t_{n-2}$.

Decision rule (level α, two-sided): reject H_0 if $|T_0| > t_{\alpha/2,\, n-2}$.

What Does "$\beta_1 = 0$" Mean Visually?

Compare Figure 6.5 in the textbook

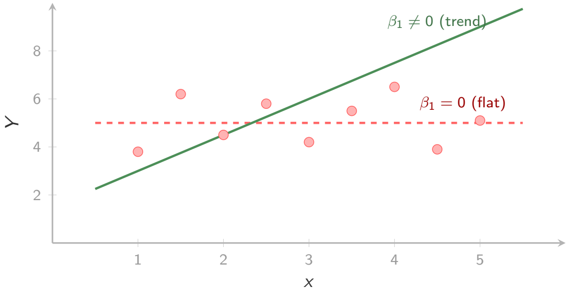

What Does "$\beta_1 = 0$" Mean Visually?

Compare Figure 6.5 in the textbook

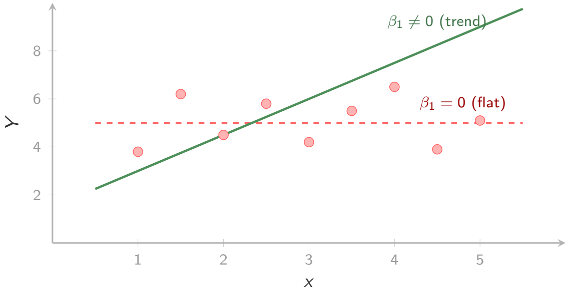

What Does "$\beta_1 = 0$" Mean Visually?

Compare Figure 6.5 in the textbook

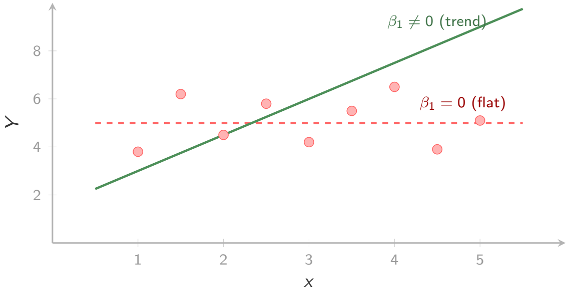

If $\beta_1 = 0$, knowing x tells you **nothing** about y: the line is flat and x has no predictive value. The test asks whether the observed slope is far enough from zero — *measured in standard errors* — to rule out chance.

Procedure 6.2: The Slope Test, Step by Step

1. State the hypotheses (6.17); the default is H_0: $\beta_1 = 0$ vs. H_1: $\beta_1 \neq 0$.

Procedure 6.2: The Slope Test, Step by Step

1. State the hypotheses (6.17); the default is H_0: $\beta_1 = 0$ vs. H_1: $\beta_1 \neq 0$.
2. Choose the significance level α.

Procedure 6.2: The Slope Test, Step by Step

1. State the hypotheses (6.17); the default is H_0: $\beta_1 = 0$ vs. H_1: $\beta_1 \neq 0$.
2. Choose the significance level α.
3. Compute $\hat{\sigma}^2 = SS_E/(n-2)$ (6.14), then $se(\hat{\beta}_1) = \sqrt{\hat{\sigma}^2/S_{xx}}$ (6.15).

Procedure 6.2: The Slope Test, Step by Step

1. State the hypotheses (6.17); the default is H_0: $\beta_1 = 0$ vs. H_1: $\beta_1 \neq 0$.
2. Choose the significance level α.
3. Compute $\hat{\sigma}^2 = SS_E/(n-2)$ (6.14), then $se(\hat{\beta}_1) = \sqrt{\hat{\sigma}^2/S_{xx}}$ (6.15).
4. Compute the test statistic T_0 (6.18).

Procedure 6.2: The Slope Test, Step by Step

1. State the hypotheses (6.17); the default is H_0: $\beta_1 = 0$ vs. H_1: $\beta_1 \neq 0$.
2. Choose the significance level α.
3. Compute $\hat{\sigma}^2 = SS_E/(n-2)$ (6.14), then $se(\hat{\beta}_1) = \sqrt{\hat{\sigma}^2/S_{xx}}$ (6.15).
4. Compute the test statistic T_0 (6.18).
5. Find the critical value $t_{\alpha/2,\, n-2}$ from Table B.3.

Procedure 6.2: The Slope Test, Step by Step

1. State the hypotheses (6.17); the default is H_0: $\beta_1 = 0$ vs. H_1: $\beta_1 \neq 0$.
2. Choose the significance level α.
3. Compute $\hat{\sigma}^2 = SS_E/(n-2)$ (6.14), then $se(\hat{\beta}_1) = \sqrt{\hat{\sigma}^2/S_{xx}}$ (6.15).
4. Compute the test statistic T_0 (6.18).
5. Find the critical value $t_{\alpha/2,\, n-2}$ from Table B.3.
6. Reject H_0 if $|T_0| > t_{\alpha/2,\, n-2}$; otherwise fail to reject.

Procedure 6.2: The Slope Test, Step by Step

1. State the hypotheses (6.17); the default is H_0: $\beta_1 = 0$ vs. H_1: $\beta_1 \neq 0$.
2. Choose the significance level α.
3. Compute $\hat{\sigma}^2 = SS_E/(n-2)$ (6.14), then $se(\hat{\beta}_1) = \sqrt{\hat{\sigma}^2/S_{xx}}$ (6.15).
4. Compute the test statistic T_0 (6.18).
5. Find the critical value $t_{\alpha/2,\, n-2}$ from Table B.3.
6. Reject H_0 if $|T_0| > t_{\alpha/2,\, n-2}$; otherwise fail to reject.
7. State the conclusion **in context**, with units.

Procedure 6.2: The Slope Test, Step by Step

1. State the hypotheses (6.17); the default is H_0: $\beta_1 = 0$ vs. H_1: $\beta_1 \neq 0$.
2. Choose the significance level α.
3. Compute $\hat{\sigma}^2 = SS_E/(n-2)$ (6.14), then $se(\hat{\beta}_1) = \sqrt{\hat{\sigma}^2/S_{xx}}$ (6.15).
4. Compute the test statistic T_0 (6.18).
5. Find the critical value $t_{\alpha/2,\, n-2}$ from Table B.3.
6. Reject H_0 if $|T_0| > t_{\alpha/2,\, n-2}$; otherwise fail to reject.
7. State the conclusion **in context**, with units.

Seven steps — the same skeleton as every test since Chapter 4. Let's run it on the fuel data.

Worked Example: Is Load Really Driving Fuel Use?

Steps 1–4 of Procedure 6.2

Given (from Lecture 17): $\hat{\beta}_1 = 1.124$, $\hat{\sigma}^2 = 0.884$, $S_{xx} = 168$, $n = 8$.

Worked Example: Is Load Really Driving Fuel Use?

Steps 1–4 of Procedure 6.2

Given (from Lecture 17): $\hat{\beta}_1 = 1.124$, $\hat{\sigma}^2 = 0.884$, $S_{xx} = 168$, $n = 8$.

Step 1–2: H_0: $\beta_1 = 0$ vs. H_1: $\beta_1 \neq 0$, at $\alpha = 0.05$.

Worked Example: Is Load Really Driving Fuel Use?

Steps 1–4 of Procedure 6.2

Given (from Lecture 17): $\hat{\beta}_1 = 1.124$, $\hat{\sigma}^2 = 0.884$, $S_{xx} = 168$, $n = 8$.

Step 1–2: H_0: $\beta_1 = 0$ vs. H_1: $\beta_1 \neq 0$, at $\alpha = 0.05$.

Step 3 — standard error (6.15):

$$se(\hat{\beta}_1) = \sqrt{\frac{\hat{\sigma}^2}{S_{xx}}} = \sqrt{\frac{0.884}{168}} = \sqrt{0.005262} = 0.0725$$

Worked Example: Is Load Really Driving Fuel Use?

Steps 1–4 of Procedure 6.2

Given (from Lecture 17): $\hat{\beta}_1 = 1.124$, $\hat{\sigma}^2 = 0.884$, $S_{xx} = 168$, $n = 8$.

Step 1–2: H_0: $\beta_1 = 0$ vs. H_1: $\beta_1 \neq 0$, at $\alpha = 0.05$.

Step 3 — standard error (6.15):

$$se(\hat{\beta}_1) = \sqrt{\frac{\hat{\sigma}^2}{S_{xx}}} = \sqrt{\frac{0.884}{168}} = \sqrt{0.005262} = 0.0725$$

Step 4 — test statistic (6.18):

$$T_0 = \frac{\hat{\beta}_1 - 0}{se(\hat{\beta}_1)} = \frac{1.124}{0.0725} = 15.50$$

Worked Example: Is Load Really Driving Fuel Use?

Steps 1–4 of Procedure 6.2

Given (from Lecture 17): $\hat{\beta}_1 = 1.124$, $\hat{\sigma}^2 = 0.884$, $S_{xx} = 168$, $n = 8$.

Step 1–2: H_0: $\beta_1 = 0$ vs. H_1: $\beta_1 \neq 0$, at $\alpha = 0.05$.

Step 3 — standard error (6.15):

$$se(\hat{\beta}_1) = \sqrt{\frac{\hat{\sigma}^2}{S_{xx}}} = \sqrt{\frac{0.884}{168}} = \sqrt{0.005262} = 0.0725$$

Step 4 — test statistic (6.18):

$$T_0 = \frac{\hat{\beta}_1 - 0}{se(\hat{\beta}_1)} = \frac{1.124}{0.0725} = 15.50$$

The estimated slope sits **15.5 standard errors** away from zero

Worked Example: Decision and Conclusion

Steps 5–7 of Procedure 6.2

Step 5 — critical value: with $n - 2 = 6$ df, Table B.3 gives

$$t_{0.025,\,6} = 2.447$$

Worked Example: Decision and Conclusion

Steps 5–7 of Procedure 6.2

Step 5 — critical value: with $n-2=6$ df, Table B.3 gives

$$t_{0.025,\,6} = 2.447$$

Step 6 — decide:

$$|T_0| = 15.50 > 2.447 \quad \Longrightarrow \quad \textbf{reject } H_0$$

Worked Example: Decision and Conclusion

Steps 5–7 of Procedure 6.2

Step 5 — critical value: with $n - 2 = 6$ df, Table B.3 gives

$$t_{0.025,\,6} = 2.447$$

Step 6 — decide:

$$|T_0| = 15.50 > 2.447 \quad \Longrightarrow \quad \textbf{reject } H_0$$

P-value check: even $t_{0.0005,\,6} = 5.959$ is far below 15.50, so the P-value is **far below 0.001**. The data would be extremely unlikely if load had no effect.

Worked Example: Decision and Conclusion

Steps 5–7 of Procedure 6.2

Step 5 — critical value: with $n - 2 = 6$ df, Table B.3 gives

$$t_{0.025,\,6} = 2.447$$

Step 6 — decide:

$$|T_0| = 15.50 > 2.447 \quad \Longrightarrow \quad \textbf{reject } H_0$$

P-value check: even $t_{0.0005,\,6} = 5.959$ is far below 15.50, so the P-value is **far below 0.001**. The data would be extremely unlikely if load had no effect.

Step 7 — in context: there is very strong evidence that payload has a real linear effect on fuel consumption. The relationship the analyst saw in the scatter plot is not an artifact of an eight-trip sample.

Confidence Interval for the Slope

Using (6.23)

A $100(1-\alpha)\%$ CI for β_1 — equation (6.23):

$$\hat{\beta}_1 \pm t_{\alpha/2,\,n-2}\sqrt{\frac{\hat{\sigma}^2}{S_{xx}}} \tag{6.23}$$

Same pattern as always: estimate $\pm$ (critical value) $\times$ (standard error).

Confidence Interval for the Slope

Using (6.23)

A $100(1-\alpha)\%$ CI for β_1 — equation (6.23):

$$\hat{\beta}_1 \pm t_{\alpha/2,\, n-2} \sqrt{\frac{\hat{\sigma}^2}{S_{xx}}} \tag{6.23}$$

Same pattern as always: estimate $\pm$ (critical value) $\times$ (standard error).
Fuel data, 95%:

$$1.124 \pm 2.447 \times 0.0725 = 1.124 \pm 0.177 \quad \Longrightarrow \quad \boxed{(0.947,\ 1.301)}$$

Confidence Interval for the Slope

Using (6.23)

A $100(1-\alpha)\%$ CI for β_1 — equation (6.23):

$$\hat{\beta}_1 \pm t_{\alpha/2,\, n-2} \sqrt{\frac{\hat{\sigma}^2}{S_{xx}}} \tag{6.23}$$

Same pattern as always: estimate $\pm$ (critical value) $\times$ (standard error).
Fuel data, 95%:

$$1.124 \pm 2.447 \times 0.0725 = 1.124 \pm 0.177 \quad \Longrightarrow \quad \boxed{(0.947,\ 1.301)}$$

Interpretation: we are 95% confident that each extra tonne of payload adds between 0.95 and 1.30 L/100 km. At SAR 1.15 per litre of diesel, that pins down the cost per tonne on a 400-km run to within about SAR 1.6.

Confidence Interval for the Slope

Using (6.23)

A $100(1-\alpha)\%$ CI for β_1 — equation (6.23):

$$\hat{\beta}_1 \pm t_{\alpha/2,\,n-2}\sqrt{\frac{\hat{\sigma}^2}{S_{xx}}} \tag{6.23}$$

Same pattern as always: estimate $\pm$ (critical value) $\times$ (standard error).
Fuel data, 95%:

$$1.124 \pm 2.447 \times 0.0725 = 1.124 \pm 0.177 \quad \Longrightarrow \quad \boxed{(0.947,\ 1.301)}$$

Interpretation: we are 95% confident that each extra tonne of payload adds between 0.95 and 1.30 L/100 km. At SAR 1.15 per litre of diesel, that pins down the cost per tonne on a 400-km run to within about SAR 1.6.

Quick Check: Back to the Startup Pricing Experiment

From Lecture 17's practice problem

Recall: discount x (%) vs. signup conversion y (%), $n = 5$, $\hat{y} = 2.55 + 0.13\,x$, $S_{xx} = 250$, $S_{xy} = 32.5$, $S_{yy} = 4.30$. **Is the slope significant at $\alpha = 0.05$?**

Quick Check: Back to the Startup Pricing Experiment

From Lecture 17's practice problem

Recall: discount x (%) vs. signup conversion y (%), $n = 5$, $\hat{y} = 2.55 + 0.13\,x$, $S_{xx} = 250$, $S_{xy} = 32.5$, $S_{yy} = 4.30$. **Is the slope significant at $\alpha = 0.05$?**

Error variance first: $SS_E = S_{yy} - \hat{\beta}_1 S_{xy} = 4.30 - 0.13(32.5) = 0.075$ (6.13), so $\hat{\sigma}^2 = 0.075/3 = 0.025$ (6.14).

Quick Check: Back to the Startup Pricing Experiment

From Lecture 17's practice problem

Recall: discount x (%) vs. signup conversion y (%), $n = 5$, $\hat{y} = 2.55 + 0.13\,x$, $S_{xx} = 250$, $S_{xy} = 32.5$, $S_{yy} = 4.30$. **Is the slope significant at $\alpha = 0.05$?**

Error variance first: $SS_E = S_{yy} - \hat{\beta}_1 S_{xy} = 4.30 - 0.13(32.5) = 0.075$ (6.13), so $\hat{\sigma}^2 = 0.075/3 = 0.025$ (6.14).

Standard error (6.15): $se(\hat{\beta}_1) = \sqrt{0.025/250} = \sqrt{0.0001} = 0.01$.

Quick Check: Back to the Startup Pricing Experiment

From Lecture 17's practice problem

Recall: discount x (%) vs. signup conversion y (%), $n = 5$, $\hat{y} = 2.55 + 0.13\,x$, $S_{xx} = 250$, $S_{xy} = 32.5$, $S_{yy} = 4.30$. **Is the slope significant at $\alpha = 0.05$?**

Error variance first: $SS_E = S_{yy} - \hat{\beta}_1 S_{xy} = 4.30 - 0.13(32.5) = 0.075$ (6.13), so $\hat{\sigma}^2 = 0.075/3 = 0.025$ (6.14).

Standard error (6.15): $se(\hat{\beta}_1) = \sqrt{0.025/250} = \sqrt{0.0001} = 0.01$.

Test (6.18): $T_0 = 0.13/0.01 = 13.0 > t_{0.025,\,3} = 3.182 \implies$ **reject** H_0.

Quick Check: Back to the Startup Pricing Experiment

From Lecture 17's practice problem

Recall: discount x (%) vs. signup conversion y (%), $n = 5$, $\hat{y} = 2.55 + 0.13\,x$, $S_{xx} = 250$, $S_{xy} = 32.5$, $S_{yy} = 4.30$. **Is the slope significant at $\alpha = 0.05$?**

Error variance first: $SS_E = S_{yy} - \hat{\beta}_1 S_{xy} = 4.30 - 0.13(32.5) = 0.075$ (6.13), so $\hat{\sigma}^2 = 0.075/3 = 0.025$ (6.14).

Standard error (6.15): $se(\hat{\beta}_1) = \sqrt{0.025/250} = \sqrt{0.0001} = 0.01$.

Test (6.18): $T_0 = 0.13/0.01 = 13.0 > t_{0.025,\,3} = 3.182 \implies$ **reject** H_0.

95% CI (6.23): $0.13 \pm 3.182(0.01) = 0.13 \pm 0.032 \Rightarrow (0.098,\ 0.162)$.
The founders can defend the discount effect: between 0.10 and 0.16 conversion points per discount point.

Inference on the Intercept β_0

(6.16), (6.25), (6.24)

Test statistic for H_0: $\beta_0 = \beta_{0,0}$ — equation (6.25):

$$T_0 = \frac{\hat{\beta}_0 - \beta_{0,0}}{\sqrt{\hat{\sigma}^2 \left(\frac{1}{n} + \frac{\bar{x}^2}{S_{xx}} \right)}} \tag{6.25}$$

Under H_0, $T_0 \sim t_{n-2}$, with $se(\hat{\beta}_0)$ from (6.16) in the denominator.

Inference on the Intercept β_0

(6.16), (6.25), (6.24)

Test statistic for H_0: $\beta_0 = \beta_{0,0}$ — equation (6.25):

$$T_0 = \frac{\hat{\beta}_0 - \beta_{0,0}}{\sqrt{\hat{\sigma}^2 \left(\frac{1}{n} + \frac{\bar{x}^2}{S_{xx}} \right)}} \tag{6.25}$$

Under H_0, $T_0 \sim t_{n-2}$, with $se(\hat{\beta}_0)$ from (6.16) in the denominator.
Confidence interval — equation (6.24):

$$\hat{\beta}_0 \pm t_{\alpha/2,\, n-2} \sqrt{\hat{\sigma}^2 \left(\frac{1}{n} + \frac{\bar{x}^2}{S_{xx}} \right)} \tag{6.24}$$

Inference on the Intercept β_0

(6.16), (6.25), (6.24)

Test statistic for H_0: $\beta_0 = \beta_{0,0}$ — equation (6.25):

$$T_0 = \frac{\hat{\beta}_0 - \beta_{0,0}}{\sqrt{\hat{\sigma}^2 \left(\frac{1}{n} + \frac{\bar{x}^2}{S_{xx}} \right)}} \tag{6.25}$$

Under H_0, $T_0 \sim t_{n-2}$, with $se(\hat{\beta}_0)$ from (6.16) in the denominator.
Confidence interval — equation (6.24):

$$\hat{\beta}_0 \pm t_{\alpha/2,\, n-2} \sqrt{\hat{\sigma}^2 \left(\frac{1}{n} + \frac{\bar{x}^2}{S_{xx}} \right)} \tag{6.24}$$

Compare with the slope: the intercept SE carries an extra $\bar{x}^2/S_{xx}$ term — the

Worked Example: 95% CI for the Intercept

Standard error (6.16), with $\bar{x} = 13$, $S_{xx} = 168$, $n = 8$:

$$se(\hat{\beta}_0) = \sqrt{0.884\left(\frac{1}{8} + \frac{13^2}{168}\right)} = \sqrt{0.884\,(0.125 + 1.006)} = \sqrt{0.9998} = 1.000$$

Worked Example: 95% CI for the Intercept

Standard error (6.16), with $\bar{x} = 13$, $S_{xx} = 168$, $n = 8$:

$$se(\hat{\beta}_0) = \sqrt{0.884\left(\frac{1}{8} + \frac{13^2}{168}\right)} = \sqrt{0.884\,(0.125 + 1.006)} = \sqrt{0.9998} = 1.000$$

95% CI (6.24):

$$16.04 \pm 2.447 \times 1.000 = 16.04 \pm 2.45 \quad \Longrightarrow \quad \boxed{(13.59,\ 18.49)}$$

Worked Example: 95% CI for the Intercept

Standard error (6.16), with $\bar{x} = 13$, $S_{xx} = 168$, $n = 8$:

$$se(\hat{\beta}_0) = \sqrt{0.884\left(\frac{1}{8} + \frac{13^2}{168}\right)} = \sqrt{0.884\,(0.125 + 1.006)} = \sqrt{0.9998} = 1.000$$

95% CI (6.24):

$$16.04 \pm 2.447 \times 1.000 = 16.04 \pm 2.45 \quad \Longrightarrow \quad \boxed{(13.59,\ 18.49)}$$

Careful with the interpretation. This brackets the model's anchor at $x = 0$ — an *empty* truck. But our loads ran from 6 to 20 tonnes, so $x = 0$ is **outside the data range**.

Worked Example: 95% CI for the Intercept

Standard error (6.16), with $\bar{x} = 13$, $S_{xx} = 168$, $n = 8$:

$$se(\hat{\beta}_0) = \sqrt{0.884\left(\frac{1}{8} + \frac{13^2}{168}\right)} = \sqrt{0.884\,(0.125 + 1.006)} = \sqrt{0.9998} = 1.000$$

95% CI (6.24):

$$16.04 \pm 2.447 \times 1.000 = 16.04 \pm 2.45 \quad \Longrightarrow \quad \boxed{(13.59,\ 18.49)}$$

Careful with the interpretation. This brackets the model's anchor at $x = 0$ — an *empty* truck. But our loads ran from 6 to 20 tonnes, so $x = 0$ is **outside the data range**.

Rule: test or interpret β_0 only when $x = 0$ is inside (or near) the observed range, or when physics says the line must pass through a known point. Otherwise treat

ANOVA for Regression

Decomposing the Total Variability

The ANOVA Identity

Total = Regression + Error, (6.19)

The total variability in y splits into two pieces — equation (6.19):

$$\underbrace{\sum_{i=1}^{n}(y_i - \bar{y})^2}_{SS_T} = \underbrace{\sum_{i=1}^{n}(\hat{y}_i - \bar{y})^2}_{SS_R} + \underbrace{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}_{SS_E} \tag{6.19}$$

The ANOVA Identity

Total = Regression + Error, (6.19)

The total variability in y splits into two pieces — equation (6.19):

$$\underbrace{\sum_{i=1}^{n}(y_i - \bar{y})^2}_{SS_T} = \underbrace{\sum_{i=1}^{n}(\hat{y}_i - \bar{y})^2}_{SS_R} + \underbrace{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}_{SS_E} \qquad (6.19)$$

$SS_T = S_{yy}$	(6.12)	total variability in y
$SS_R = \hat{\beta}_1 \, S_{xy}$		variability **explained** by the line
$SS_E = S_{yy} - \hat{\beta}_1 \, S_{xy}$	(6.13)	variability **left over**

The ANOVA Identity

Total = Regression + Error, (6.19)

The total variability in y splits into two pieces — equation (6.19):

$$\underbrace{\sum_{i=1}^{n}(y_i - \bar{y})^2}_{SS_T} = \underbrace{\sum_{i=1}^{n}(\hat{y}_i - \bar{y})^2}_{SS_R} + \underbrace{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}_{SS_E} \tag{6.19}$$

$SS_T = S_{yy}$	(6.12)	total variability in y
$SS_R = \hat{\beta}_1\, S_{xy}$		variability **explained** by the line
$SS_E = S_{yy} - \hat{\beta}_1\, S_{xy}$	(6.13)	variability **left over**

Key idea: if the model is useful, SS_R is large relative to SS_E. If there is no relationship, almost everything stays in SS_E.

Textbook: §6.8 ((6.19) onward)

The ANOVA Table and the F-Test

Table 6.4 and (6.20)

The layout the textbook calls Table 6.4:

Source	**df**	**SS**	**MS**	F_0
Regression	1	SS_R	$MS_R = SS_R/1$	MS_R/MS_E
Error	$n-2$	SS_E	$MS_E = SS_E/(n-2)$	
Total	$n-1$	SS_T		

The ANOVA Table and the F-Test

Table 6.4 and (6.20)

The layout the textbook calls Table 6.4:

Source	**df**	**SS**	**MS**	F_0
Regression	1	SS_R	$MS_R = SS_R/1$	MS_R/MS_E
Error	$n-2$	SS_E	$MS_E = SS_E/(n-2)$	
Total	$n-1$	SS_T		

Test statistic — equation (6.20):

$$F_0 = \frac{MS_R}{MS_E} = \frac{SS_R/1}{SS_E/(n-2)} \tag{6.20}$$

Reject H_0: $\beta_1 = 0$ if $F_0 > f_{\alpha,1,n-2}$ (from Table B.5)

The ANOVA Table and the F-Test

Table 6.4 and (6.20)

The layout the textbook calls Table 6.4:

Source	**df**	**SS**	**MS**	F_0
Regression	1	SS_R	$MS_R = SS_R/1$	MS_R/MS_E
Error	$n-2$	SS_E	$MS_E = SS_E/(n-2)$	
Total	$n-1$	SS_T		

Test statistic — equation (6.20):

$$F_0 = \frac{MS_R}{MS_E} = \frac{SS_R/1}{SS_E/(n-2)} \tag{6.20}$$

Reject $H_0: \beta_1 = 0$ if $F_0 > f_{\alpha,1,n-2}$ (from Table B.5)

Worked Example: Building the ANOVA Table

Fuel data, one sum of squares at a time

$$\textbf{Total:} \quad SS_T = S_{yy} = 217.48 \tag{6.12}$$

Worked Example: Building the ANOVA Table

Fuel data, one sum of squares at a time

$$\textbf{Total:} \quad SS_T = S_{yy} = 217.48 \tag{6.12}$$

Regression (explained by the line):

$$SS_R = \hat{\beta}_1 \, S_{xy} = 1.124 \times 188.8 = 212.18$$

Worked Example: Building the ANOVA Table

Fuel data, one sum of squares at a time

Total: $SS_T = S_{yy} = 217.48$ (6.12)

Regression (explained by the line):

$$SS_R = \hat{\beta}_1 \, S_{xy} = 1.124 \times 188.8 = 212.18$$

Error (left over) — must match Lecture 17's SS_E:

$$SS_E = SS_T - SS_R = 217.48 - 212.18 = 5.30 \quad \checkmark$$

Worked Example: Building the ANOVA Table

Fuel data, one sum of squares at a time

Total: $SS_T = S_{yy} = 217.48$ (6.12)

Regression (explained by the line):

$$SS_R = \hat{\beta}_1 \, S_{xy} = 1.124 \times 188.8 = 212.18$$

Error (left over) — must match Lecture 17's SS_E:

$$SS_E = SS_T - SS_R = 217.48 - 212.18 = 5.30 \quad \checkmark$$

Mean squares:

$$MS_R = \frac{SS_R}{1} = 212.18, \qquad MS_E = \frac{SS_E}{n-2} = \frac{5.30}{6} = 0.884$$

Worked Example: The Completed ANOVA Table

Source	**df**	**SS**	**MS**	F_0
Regression	1	212.18	212.18	240.0
Error	6	5.30	0.884	
Total	7	217.48		

Worked Example: The Completed ANOVA Table

Source	**df**	**SS**	**MS**	F_0
Regression	1	212.18	212.18	240.0
Error	6	5.30	0.884	
Total	7	217.48		

Test at $\alpha = 0.05$ (6.20): from Table B.5, $f_{0.05,\,1,\,6} = 5.99$.
Since $F_0 = 240.0 > 5.99$, we **reject** H_0: $\beta_1 = 0$ — same verdict as the t-test.

Worked Example: The Completed ANOVA Table

Source	**df**	**SS**	**MS**	F_0
Regression	1	212.18	212.18	240.0
Error	6	5.30	0.884	
Total	7	217.48		

Test at $\alpha = 0.05$ (6.20): from Table B.5, $f_{0.05,\,1,\,6} = 5.99$.
Since $F_0 = 240.0 > 5.99$, we **reject** H_0: $\beta_1 = 0$ — same verdict as the t-test.

Equivalence check: $T_0^2 = (15.50)^2 = 240.2 \approx 240.0 = F_0$ ✓ (the small gap is rounding in $se(\hat{\beta}_1)$).

Practice Problem [5 pts]

Exam-Style: AI Safety Red-Team Data

Problem statement. An AI safety team red-teams each release of its perception model. For $n = 6$ releases they record $x =$ structured red-team hours and $y =$ critical defects discovered before deployment:

Practice Problem [5 pts]

Exam-Style: AI Safety Red-Team Data

Problem statement. An AI safety team red-teams each release of its perception model. For $n = 6$ releases they record x = structured red-team hours and y = critical defects discovered before deployment:

x (hours)	4	6	8	10	12	14
y (defects)	8.4	10.2	14.4	16.6	20.8	22.6

Practice Problem [5 pts]

Exam-Style: AI Safety Red-Team Data

Problem statement. An AI safety team red-teams each release of its perception model. For $n = 6$ releases they record $x =$ structured red-team hours and $y =$ critical defects discovered before deployment:

x (hours)	4	6	8	10	12	14
y (defects)	8.4	10.2	14.4	16.6	20.8	22.6

A junior analyst already fit the line and reports: $\hat{y} = 2.0 + 1.5\,x,$

$$S_{xx} = 70, \qquad S_{xy} = 105, \qquad S_{yy} = 159.42, \qquad \bar{x} = 9$$

Practice Problem [5 pts]

Exam-Style: AI Safety Red-Team Data

Problem statement. An AI safety team red-teams each release of its perception model. For $n = 6$ releases they record $x =$ structured red-team hours and $y =$ critical defects discovered before deployment:

x (hours)	4	6	8	10	12	14
y (defects)	8.4	10.2	14.4	16.6	20.8	22.6

A junior analyst already fit the line and reports: $\hat{y} = 2.0 + 1.5\,x,$

$$S_{xx} = 70, \qquad S_{xy} = 105, \qquad S_{yy} = 159.42, \qquad \bar{x} = 9$$

(a) [1 pt] Compute SS_E and $\hat{\sigma}^2$.

(b) [2 pts] Test $H_0 : \beta_1 = 0$ at $\alpha = 0.05$ using Procedure 6.2.

Practice Problem — Solution

Parts (a) and (b)

(a) Error sum of squares (6.13) and variance (6.14):

$$SS_E = S_{yy} - \hat{\beta}_1 S_{xy} = 159.42 - 1.5(105) = 159.42 - 157.5 = 1.92$$

$$\hat{\sigma}^2 = \frac{SS_E}{n-2} = \frac{1.92}{4} = 0.48$$

Practice Problem — Solution

Parts (a) and (b)

(a) Error sum of squares (6.13) and variance (6.14):

$$SS_E = S_{yy} - \hat{\beta}_1 S_{xy} = 159.42 - 1.5(105) = 159.42 - 157.5 = 1.92$$

$$\hat{\sigma}^2 = \frac{SS_E}{n-2} = \frac{1.92}{4} = 0.48$$

(b) Standard error (6.15):

$$se(\hat{\beta}_1) = \sqrt{\frac{0.48}{70}} = \sqrt{0.006857} = 0.0828$$

Practice Problem — Solution

Parts (a) and (b)

(a) Error sum of squares (6.13) and variance (6.14):

$$SS_E = S_{yy} - \hat{\beta}_1 S_{xy} = 159.42 - 1.5(105) = 159.42 - 157.5 = 1.92$$

$$\hat{\sigma}^2 = \frac{SS_E}{n-2} = \frac{1.92}{4} = 0.48$$

(b) Standard error (6.15):

$$se(\hat{\beta}_1) = \sqrt{\frac{0.48}{70}} = \sqrt{0.006857} = 0.0828$$

Test statistic (6.18): $T_0 = \dfrac{1.5}{0.0828} = 18.12$

Practice Problem — Solution

Parts (a) and (b)

(a) Error sum of squares (6.13) and variance (6.14):

$$SS_E = S_{yy} - \hat{\beta}_1 S_{xy} = 159.42 - 1.5(105) = 159.42 - 157.5 = 1.92$$

$$\hat{\sigma}^2 = \frac{SS_E}{n-2} = \frac{1.92}{4} = 0.48$$

(b) Standard error (6.15):

$$se(\hat{\beta}_1) = \sqrt{\frac{0.48}{70}} = \sqrt{0.006857} = 0.0828$$

Test statistic (6.18): $\quad T_0 = \dfrac{1.5}{0.0828} = 18.12$

Critical value $t_{0.025,4} = 2.776$. Since $18.12 > 2.776$, **reject** H_0: red-team hours

Practice Problem — Solution (Continued)

Parts (c) and (d)

(c) 95% CI for β_1 (6.23):

$$1.5 \pm 2.776 \times 0.0828 = 1.5 \pm 0.230 \quad \Longrightarrow \quad \boxed{(1.270,\ 1.730)}$$

Each extra red-team hour surfaces between 1.3 and 1.7 additional critical defects — evidence the team can use to argue for a bigger testing budget.

Practice Problem — Solution (Continued)

Parts (c) and (d)

(c) 95% CI for β_1 (6.23):

$$1.5 \pm 2.776 \times 0.0828 = 1.5 \pm 0.230 \quad \Longrightarrow \quad \boxed{(1.270,\ 1.730)}$$

Each extra red-team hour surfaces between 1.3 and 1.7 additional critical defects — evidence the team can use to argue for a bigger testing budget.

(d) $SS_R = \hat{\beta}_1 S_{xy} = 1.5 \times 105 = 157.5$, then (6.20):

Source	**df**	**SS**	**MS**	F_0
Regression	1	157.5	157.5	328.1
Error	4	1.92	0.48	
Total	5	159.42		

Practice Problem — Solution (Continued)

Parts (c) and (d)

(c) 95% CI for β_1 (6.23):

$$1.5 \pm 2.776 \times 0.0828 = 1.5 \pm 0.230 \quad \Longrightarrow \quad \boxed{(1.270,\ 1.730)}$$

Each extra red-team hour surfaces between 1.3 and 1.7 additional critical defects — evidence the team can use to argue for a bigger testing budget.

(d) $SS_R = \hat{\beta}_1 S_{xy} = 1.5 \times 105 = 157.5$, then (6.20):

Source	**df**	**SS**	**MS**	F_0
Regression	1	157.5	157.5	328.1
Error	4	1.92	0.48	
Total	5	159.42		

Lecture 18 Summary

- Procedure 6.2 tests whether the linear relationship is **real**: $T_0 = \hat{\beta}_1 / se(\hat{\beta}_1)$ (6.18) with $n - 2$ df, using $se(\hat{\beta}_1) = \sqrt{\hat{\sigma}^2 / S_{xx}}$ (6.15).

Lecture 18 Summary

- Procedure 6.2 tests whether the linear relationship is **real**: $T_0 = \hat{\beta}_1 / se(\hat{\beta}_1)$ (6.18) with $n-2$ df, using $se(\hat{\beta}_1) = \sqrt{\hat{\sigma}^2 / S_{xx}}$ (6.15).
- CI for the slope: $\hat{\beta}_1 \pm t_{\alpha/2,\, n-2}\, se(\hat{\beta}_1)$ (6.23). Zero outside the CI $\Leftrightarrow$ reject H_0.

Lecture 18 Summary

- Procedure 6.2 tests whether the linear relationship is **real**: $T_0 = \hat{\beta}_1 / se(\hat{\beta}_1)$ (6.18) with $n - 2$ df, using $se(\hat{\beta}_1) = \sqrt{\hat{\sigma}^2 / S_{xx}}$ (6.15).
- CI for the slope: $\hat{\beta}_1 \pm t_{\alpha/2,\, n-2}\, se(\hat{\beta}_1)$ (6.23). Zero outside the CI $\Leftrightarrow$ reject H_0.
- Intercept inference uses $se(\hat{\beta}_0)$ (6.16) with (6.25) and (6.24) — meaningful only when $x = 0$ is near the data.

Lecture 18 Summary

- Procedure 6.2 tests whether the linear relationship is **real**: $T_0 = \hat{\beta}_1 / se(\hat{\beta}_1)$ (6.18) with $n-2$ df, using $se(\hat{\beta}_1) = \sqrt{\hat{\sigma}^2 / S_{xx}}$ (6.15).
- CI for the slope: $\hat{\beta}_1 \pm t_{\alpha/2,\, n-2}\, se(\hat{\beta}_1)$ (6.23). Zero outside the CI $\Leftrightarrow$ reject H_0.
- Intercept inference uses $se(\hat{\beta}_0)$ (6.16) with (6.25) and (6.24) — meaningful only when $x = 0$ is near the data.
- ANOVA identity (6.19): $SS_T = SS_R + SS_E$ with $SS_R = \hat{\beta}_1 S_{xy}$.

Lecture 18 Summary

- Procedure 6.2 tests whether the linear relationship is **real**: $T_0 = \hat{\beta}_1 / se(\hat{\beta}_1)$ (6.18) with $n - 2$ df, using $se(\hat{\beta}_1) = \sqrt{\hat{\sigma}^2 / S_{xx}}$ (6.15).
- CI for the slope: $\hat{\beta}_1 \pm t_{\alpha/2,\, n-2}\, se(\hat{\beta}_1)$ (6.23). Zero outside the CI $\Leftrightarrow$ reject H_0.
- Intercept inference uses $se(\hat{\beta}_0)$ (6.16) with (6.25) and (6.24) — meaningful only when $x = 0$ is near the data.
- ANOVA identity (6.19): $SS_T = SS_R + SS_E$ with $SS_R = \hat{\beta}_1 S_{xy}$.
- $F_0 = MS_R / MS_E$ (6.20) equals T_0^2: two tests, one verdict.

Lecture 18 Summary

- Procedure 6.2 tests whether the linear relationship is **real**: $T_0 = \hat{\beta}_1 / se(\hat{\beta}_1)$ (6.18) with $n - 2$ df, using $se(\hat{\beta}_1) = \sqrt{\hat{\sigma}^2 / S_{xx}}$ (6.15).
- CI for the slope: $\hat{\beta}_1 \pm t_{\alpha/2,\, n-2}\, se(\hat{\beta}_1)$ (6.23). Zero outside the CI $\Leftrightarrow$ reject H_0.
- Intercept inference uses $se(\hat{\beta}_0)$ (6.16) with (6.25) and (6.24) — meaningful only when $x = 0$ is near the data.
- ANOVA identity (6.19): $SS_T = SS_R + SS_E$ with $SS_R = \hat{\beta}_1 S_{xy}$.
- $F_0 = MS_R / MS_E$ (6.20) equals T_0^2: two tests, one verdict.
- **Next lecture:** intervals for prediction, R^2, correlation, residual analysis, and how regression gets abused.

Reminders

- **Quiz on the slope test** is due before class — it walks Procedure 6.2 on a small dataset

Reminders

- **Quiz on the slope test** is due before class — it walks Procedure 6.2 on a small dataset
- Homework 6 (regression) is due next Thursday on Blackboard

Reminders

- **Quiz on the slope test** is due before class — it walks Procedure 6.2 on a small dataset
- Homework 6 (regression) is due next Thursday on Blackboard
- Office hours: Tuesdays 9–10 AM, Room 22-219 or Zoom

Reminders

- **Quiz on the slope test** is due before class — it walks Procedure 6.2 on a small dataset
- Homework 6 (regression) is due next Thursday on Blackboard
- Office hours: Tuesdays 9–10 AM, Room 22-219 or Zoom
- Read §6.9–§6.11 and §6.14 to prepare for the next lecture