Econ 103: Introduction to Econometrics

Lecture 15 — Hypothesis Testing in Multiple Regression

Ryan Longmuir

UCLA

Summer Session C, 2026

Reading: Hill, Griffiths & Lim (5th ed.), §5.4–5.5; Stock & Watson (4th ed.), §7.1, 7.3.

Where we are

We have Big Andy's coefficients *and* their standard errors:

$$\widehat{\text{SALES}} = \underset{(6.35)}{118.91} \; - \underset{(1.096)}{7.908} \text{ PRICE } + \underset{(0.683)}{1.863} \text{ ADVERT}.$$

Good news: the inference machinery is exactly Lectures 9–10 — same t-statistic, same logic. *One* change: degrees of freedom are now $N - K$.

New power: multiple regression lets us test and estimate combinations of several coefficients — the kind of question a manager actually asks ("what if I cut price *and* raise advertising?"). Today:

- t-tests and CIs on a single coefficient (df $= N - K$),
- estimating and testing linear combinations of coefficients,
- single vs. joint restrictions (a teaser for the F-test).

Today's plan

1. t-tests on a single coefficient
2. Confidence intervals in multiple regression
3. Linear combinations of coefficients
4. Single vs. joint restrictions

Part 1

t-tests on a single coefficient

The same t-statistic, new degrees of freedom

Under MR1–MR6, for each coefficient

$$t = \frac{b_k - c}{\mathrm{se}(b_k)} \sim t_{(N-K)} \quad \text{under } H_0 : \beta_k = c$$

The only difference from simple regression: degrees of freedom are $N - K$ (subtract one per estimated coefficient), not $N - 2$. For Big Andy's, $N - K = 75 - 3 = 72$.

Everything else is identical:

- rejection regions by the direction of H_1 $(>, <, \neq)$;
- the p-value rule (reject if $p \leq \alpha$);
- "reject" vs. "fail to reject" (never "accept").

Tests of significance (automatic in the output)

The default test for each coefficient is $H_0 : \beta_k = 0$ vs. $H_1 : \beta_k \neq 0$ — "does this variable matter, holding the others fixed?" With $t_c = t_{(0.975,72)} = 1.993$:

Price

$$t = \frac{-7.908}{1.096} = -7.22,$$

$|-7.22| > 1.993$ ($p \approx 0.000$) — reject. Sales depend on price.

Advertising

$$t = \frac{1.863}{0.683} = 2.73,$$

$2.73 > 1.993$ ($p = 0.008$) — reject. Sales depend on advertising.

Both regressors are statistically **significant** — the t's and p's in the regression table *are* these tests, computed automatically.

One-tail tests of economic hypotheses

The interesting questions are usually *directional* — put the conjecture in H_1.

Is demand elastic? $\beta_2 < 0$

$H_0 : \beta_2 \geq 0$ vs. $H_1 : \beta_2 < 0$.

$$t = \frac{-7.908}{1.096} = -7.22 < -1.666.$$

Reject — demand is price-elastic; a price cut *raises* revenue.

Is advertising profitable? $\beta_3 > 1$

$H_0 : \beta_3 \leq 1$ vs. $H_1 : \beta_3 > 1$.

$$t = \frac{1.863 - 1}{0.683} = 1.26 < 1.666.$$

Fail to reject — can't prove ads return more than they cost.

The advertising case is the cautionary one: $b_3 = 1.86 > 1$ *looks* profitable, but once we account for its imprecision ($\mathrm{se} = 0.68$), the data cannot rule out $\beta_3 \leq 1$. Significant $\neq$ economically decisive, and a point estimate is not proof.

Part 2

Confidence intervals in multiple regression

Interval estimates, df $= N - K$

Identical to Lecture 9, with the new degrees of freedom:

$$b_k \pm t_c \, \mathrm{se}(b_k), \qquad t_c = t_{(1-\alpha/2,\, N-K)} = t_{(0.975,72)} = 1.993.$$

Price β_2

$$-7.908 \pm 1.993(1.096) = [-10.09,\ -5.72].$$

A \$1 price cut raises mean revenue by \$5,723 to \$10,093 (a dime: \$572–\$1,009). A tight, useful interval.

Advertising β_3

$$1.863 \pm 1.993(0.683) = [0.50,\ 3.22].$$

Wide — \$1,000 of ads could return as little as \$501 (a loss) or over \$3,000. Reflects the large $\mathrm{se}(b_3)$.

The width is the story: β_3's wide interval is the same imprecision that made the "$\beta_3 > 1$" test inconclusive — and the remedy is the Lecture-14 one (more, more independent, data).

Part 3

Linear combinations of coefficients

Questions that mix several coefficients

Real decisions move *several* regressors at once. Big Andy plans to drop price 40 cents and raise advertising $800. The change in expected sales is a **linear combination**:

$$\lambda = -0.4\,\beta_2 + 0.8\,\beta_3.$$

The estimator $\hat{\lambda} = \sum_k c_k b_k$ is BLUE, so the point estimate is immediate:

$$\hat{\lambda} = -0.4(-7.908) + 0.8(1.863) = 4.65.$$

Expected sales rise about $4,653 from the strategy.

The standard error needs the covariances

The variance of a combination of *several* estimators carries a covariance term for *every* pair (the Lecture-3 rule, generalized):

$$\operatorname{Var}\Big(\sum_k c_k b_k\Big) = \sum_k c_k^2 \operatorname{Var}(b_k) + 2\sum_{j<k} c_j c_k \operatorname{Cov}(b_j, b_k).$$

For $\lambda = -0.4b_2 + 0.8b_3$, using Big Andy's var–cov matrix:

$$\widehat{\operatorname{Var}}(\hat{\lambda}) = (-0.4)^2(1.201) + (0.8)^2(0.467) - 2(0.4)(0.8)(-0.020) = 0.504,$$

$$\operatorname{se}(\hat{\lambda}) = \sqrt{0.504} = 0.71.$$

A 90% interval ($t_c = t_{(0.95,72)} = 1.666$):

$$4.65 \pm 1.666(0.71) = [\,3.47,\ 5.84\,] \;\Rightarrow\; \$3{,}471 \text{ to } \$5{,}835.$$

Drop the covariance term and you get the standard error *wrong* — here it would understate the spread.

Testing a linear combination

To test a claim about a combination, use the same t with the combination's standard error:

$$t = \frac{\sum_k c_k b_k - c_0}{\mathrm{se}\left(\sum_k c_k b_k\right)} \sim t_{(N-K)}.$$

S&W's reparametrization trick

To test, say, $H_0 : \beta_1 = \beta_2$ (equal effects), rewrite the restriction as a *single* new coefficient $\gamma = \beta_1 - \beta_2$ and regress so that γ appears directly. Then it's an ordinary single-coefficient t-test, and $\hat{\gamma} \pm 1.96\,\mathrm{se}(\hat{\gamma})$ is a CI for the difference.

Either way — direct standard error or reparametrization — a single restriction, even one spanning several coefficients, is a t-test.

Part 4

Single vs. joint restrictions

One restriction vs. many

Everything today tests a **single restriction** — one equation about the β's (even if it mixes several, like $-0.4\beta_2 + 0.8\beta_3$). That is always a t-test.

What a t-test *cannot* do

Test several restrictions at once — a **joint** hypothesis like

$$H_0: \ \beta_2 = 0 \ \text{ and } \ \beta_3 = 0.$$

Checking each with its own t "one at a time" is unreliable (the error rates compound). This needs a new statistic — the F-**test**.

That joint machinery — overall significance, nested model comparisons, testing economic restrictions like constant returns to scale — is Lecture 17.

Recap

Single-coefficient inference

- $t = \dfrac{b_k - c}{\mathrm{se}(b_k)} \sim t_{(N-K)}$ — only df changes
- significance: PRICE ($t = -7.2$), ADVERT ($t = 2.7$) both reject
- one-tail: demand elastic; ads *not* proven profitable ($t = 1.26$)
- CI: $\beta_2 \in [-10.1, -5.7]$, $\beta_3 \in [0.50, 3.22]$ (wide)

Linear combinations

- $\hat{\lambda} = \sum c_k b_k$ (BLUE)
- Var includes all covariances
- Andy's strategy: \$4653, 90% CI $[\$3471, \$5835]$
- single restriction ⇒ t-test (even across coefficients)

Joint restrictions

- several at once ⇒ F-test (Lecture 17), not many t's

Next time (Lecture 16): interaction terms

Coefficients so far are *constant* partial effects. With interactions ($x_2 \times x_3$) the effect of one variable depends on another — and marginal effects, like Andy's optimal advertising, come from *combinations* of coefficients we now know how to handle.

Questions?