Sampling Distribution of OLS and Coefficient Tests

21 Sep 2026 | Chapter 5

ECO 351 – Introduction to Econometrics

Shi Qiu – Fall 2026

Recall: The regression error

From Lecture 8

The class-size equation requires $E(u_i \mid STR_i) = 0$ for its conditional-mean interpretation.

$$Score_i = \beta_0 + \beta_1 STR_i + u_i$$

Suppose districts with greater student needs receive extra teachers. Why could this violate the assumption? Does an average residual of zero rule that out?

From one fitted line to repeated-sample uncertainty

Answer

Student need affects scores and the assignment of teachers, so its omitted contribution can be related to STR. OLS residuals sum to zero by construction; that identity does not test the unobserved population error assumption.

Suppose the model and sampling assumptions are defensible. How much would the slope change in another sample?

Is the negative slope distinguishable from sampling noise?

The California sample estimates an STR slope of -2.28 score points. A new sample would produce a different set of districts and a different fitted slope.

Our question

Could a slope this far below zero arise often if the population slope were zero?

The slope estimate changes from sample to sample

Sampling distribution

Even if the population regression line is fixed, a new random sample produces a new cloud of points and therefore a new OLS slope. Its standard error estimates the spread of those hypothetical slopes.

Large-sample normality

Under the least-squares assumptions, the standardized slope estimator is approximately normal. This is the basis for regression t tests and confidence intervals.

Why errors make the slope move across samples

Under the linear model, let $S_{XX} = \sum_i (X_i - \bar{X})^2$.

$$\widehat{\beta}_1 - \beta_1 = \frac{\sum_i (X_i - \bar{X}) u_i}{S_{XX}}$$

Numerator

Observations far from the mean of X give their errors more weight.

Denominator

More variation in X helps distinguish a slope from noise.

The SE measures how much this weighted average fluctuates across repeated samples.

Two standard deviations with different jobs

The slope SE estimates how much $\widehat{\beta}_{STR}$ would vary across repeated samples under the maintained model.

Quantity	School estimate	What varies?
SER	18.58 points	Individual score residuals
Slope SE	0.519	Estimated STR slope

$$\widehat{\beta}_{STR} = -2.280$$

Use the slope SE, not the SER, to test a claim about the population slope.

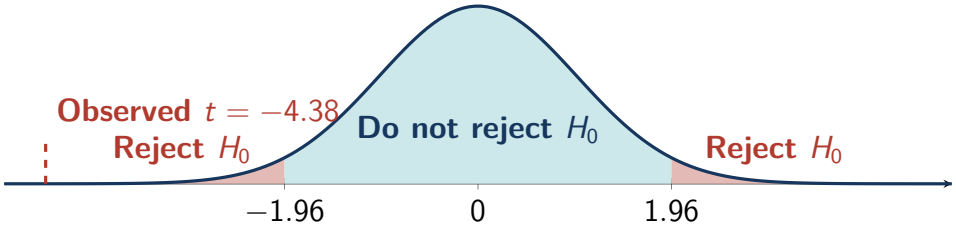
State and standardize the no-association benchmark

$$H_0 : \beta_{STR} = 0, \quad H_1 : \beta_{STR} \neq 0$$

$$t = \frac{\widehat{\beta}_{STR} - 0}{\widehat{SE}(\widehat{\beta}_{STR})} = \frac{-2.28}{0.52} \approx -4.38$$

The estimate is about 4.38 estimated standard errors below zero.

A two-sided test divides the null distribution into regions



$$|t| = 4.38 > 1.96 \quad \Longrightarrow \quad \text{reject } H_0 : \beta_{STR} = 0 \text{ at } 5\%.$$

The two red tails contain 5% of the null distribution in total. The p-value is the null probability of a statistic at least as extreme as the one observed.

Test a different proposed class-size slope

Use the same estimate -2.28 and estimated SE 0.52. Test $H_0 : \beta_{STR} = -1$ against a two-sided alternative at 5% using the normal approximation.

Solution: Test a different proposed class-size slope

$$t = \frac{-2.28 - (-1)}{0.52} = -2.46$$

Reject at 5%, because $|t| > 1.96$. The subtraction must use the stated null value, not zero by habit.

Collect the slope values consistent with the test

$$\widehat{\beta}_{STR} \pm 1.96\widehat{SE}(\widehat{\beta}_{STR})$$

$$-2.28 \pm 1.96(0.52) = [-3.30, -1.26]$$

The interval is expressed in score points per one-unit increase in STR. It excludes both 0 and -1, consistent with the two tests.

The policy comparison needs the same change of scale

The opening policy question concerns a reduction of two students per teacher, so define the fitted contrast $\Delta = -2\beta_{STR}$.

$$\widehat{\Delta} = 4.56, \qquad \widehat{SE}(\widehat{\Delta}) = 2(0.52) = 1.04$$

$$4.56 \pm 1.96(1.04) = [2.52, 6.60] \text{ score points}$$

Multiplying an interval by a negative number reverses the order of its endpoints.

Stata: test the slope and estimate the policy contrast

Predict which null test rejects and what sign the two-student contrast will have.

```
do "setup.do" schools
regress score str, vce(robust)
test str = -1
lincom -2*str
```

Read the result

test evaluates the stated restriction. lincom reports a rescaled estimate, its standard error, and its interval.

Stata output excerpt (1/2)

```
. regress score str, vce(robust)
Number of obs     =         420

------------------------------------------------------------------------------
             |               Robust
       score | Coefficient  std. err.      t    P>|t|     [95% conf. interval]
-------------+----------------------------------------------------------------
         str |  -2.279808   .5194892    -4.39   0.000    -3.300945   -1.258671
       _cons |   698.933   10.36436     67.44   0.000     678.5602    719.3057
------------------------------------------------------------------------------

. test str = -1
 ( 1)  str = -1

       F(  1,   418) =    6.07
            Prob > F =    0.0142
```

Stata output excerpt (2/2)

```
. lincom -2*str
 ( 1)  - 2*str = 0

------------------------------------------------------------------------------
       score | Coefficient  Std. err.      t    P>|t|     [95% conf. interval]
-------------+----------------------------------------------------------------
         (1) |   4.559617   1.038978     4.39   0.000     2.517343     6.60189
------------------------------------------------------------------------------
```

Compare a three-student reduction

Using -2.28 and SE 0.52, calculate the fitted effect and estimated SE for reducing STR by three students. Give a 95% interval using the normal approximation.

Solution: Compare a three-student reduction

$$\widehat{\Delta} = 6.84, \quad SE = 1.56, \quad CI \approx [3.78, 9.90]$$

The scale changes, but evidence against a zero contrast is unchanged: estimate and SE are multiplied by the same positive magnitude.

What the test does not establish

Claim

A p-value below 0.001 proves that hiring more teachers will raise scores.

What the test actually says

The estimated population regression slope is difficult to reconcile with the zero-slope null under the maintained model.

A causal claim still needs a credible argument that class size is separated from student needs and other score determinants.

Inference for a regression slope

Conclusion

The slope is precisely estimated relative to a zero benchmark. Expressing the result as a two-student contrast gives 4.56 points with an approximate interval [2.52,6.60].

Why must the standard error change when the reported policy contrast changes?