Multiple-Regression Inference: From Single to Joint Hypotheses

12 Oct 2026 | Chapter 7

ECO 351 – Introduction to Econometrics

Shi Qiu – Fall 2026

From the school-board result to today's question

Lecture 13 distinguished two class-size plans. For the plan holding EL and spending fixed, two fewer students per teacher gave a fitted increase of about 0.57 score point.

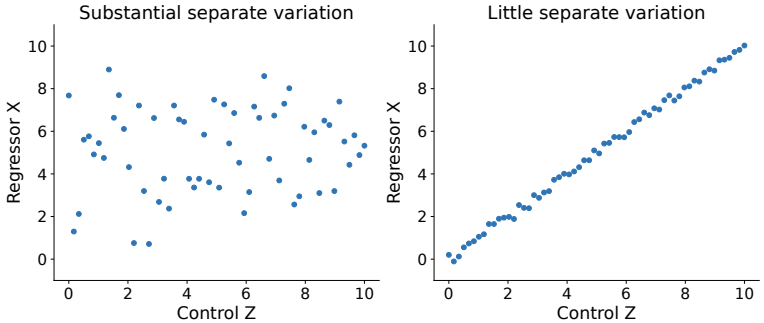
What we still need

How uncertain is that contrast? Why can related regressors make separate slopes difficult to estimate?

A broader claim

A taxpayer argues that neither class size nor spending matters after accounting for EL. That is a claim about two coefficients together.

Related regressors leave less separate variation



Read the two clouds

When Z nearly determines X, it is harder to find observations with different X but similar Z. Those comparisons are what help estimate X's separate slope.

Stata: quantify uncertainty in the fixed-budget contrast

Reestimate the full model, then ask for the two-student contrast and inspect regressor overlap.

```
do "setup.do" schools
regress score str el spend, vce(robust)
lincom -2*str
estat vif
```

Read the result

lincom reports the contrast's estimate, SE, test, and interval. VIF concerns regressor overlap, not causal credibility.

Stata output: the full model and class-size contrast

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

------------------------------------------------------------------------------
             |               Robust
       score | Coefficient  std. err.      t    P>|t|     [95% conf. interval]
-------------+----------------------------------------------------------------
         str |  -.2863992   .4820728    -0.59   0.553    -1.234002     .661203
          el |  -.6560227   .0317844   -20.64   0.000    -.7185008    -.5935446
       spend |   .0038679   .0015807     2.45   0.015     .0007607     .0069751
       _cons |   649.5779   15.45834    42.02   0.000     619.1917     679.9641
------------------------------------------------------------------------------

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

------------------------------------------------------------------------------
       score | Coefficient  Std. err.      t    P>|t|     [95% conf. interval]
-------------+----------------------------------------------------------------
         (1) |   .5727985   .9641456      0.59   0.553    -1.322406     2.468003
------------------------------------------------------------------------------
```

Variance inflation factor (VIF)

```
. estat vif
    Variable |       VIF        1/VIF
-------------+----------------------
         str |      1.68     0.594960
       spend |      1.63     0.613529
          el |      1.04     0.961510
-------------+----------------------
    Mean VIF |      1.45
```

Calculation

Regress X_j on the other regressors and an intercept; call its fit R_j^2.

$$VIF_j = \frac{1}{1 - R_j^2}$$

This is not the outcome regression's R^2.

Interpret the STR result

$R^2_{STR} \approx 0.405$: EL and spending explain about 40.5% of STR variation. Under homoskedasticity, $VIF \approx 1.68$ means 68% higher slope variance (about 30% higher SE) than with no regressor overlap, holding error variance and STR spread fixed.

VIF is an overlap diagnostic, not a bias test or an exact robust-SE multiplier.

What the output does—and does not—say

$$\widehat{\Delta} = 0.573, \qquad SE(\widehat{\Delta}) = 0.964, \qquad 95\%\ \text{CI} = [-1.322,\ 2.468]$$

The contrast

The interval includes zero and positive changes of more than two score points. Nonrejection is not proof of no class-size effect.

The precision diagnosis

STR has VIF 1.68: it is not almost determined by EL and spending. Its SE rose only from 0.433 to 0.482 when spending was added; the point estimate moved much more substantially.

Write the taxpayer's claim as two restrictions

$$Score_i = \beta_0 + \beta_1 STR_i + \beta_2 EL_i + \beta_3 Spend_i + u_i$$

$$H_0 : \beta_1 = 0 \text{ and } \beta_3 = 0$$

$$H_1 : \beta_1 \neq 0 \text{ or } \beta_3 \neq 0 \quad \text{(or both)}.$$

Count the restrictions

Two coefficients are fixed at specified values. The EL coefficient and intercept remain unrestricted. Rejecting this null does not say that both resource coefficients are nonzero.

What would the model look like if the joint null held?

Unrestricted specification

$$Score_i = \beta_0 + \beta_1 STR_i + \beta_2 EL_i + \beta_3 Spend_i + u_i$$

Restricted specification

$$Score_i = \alpha_0 + \alpha_1 EL_i + v_i$$

The restricted regression removes STR and Spend together. It still allows scores to vary with EL.

A joint test asks whether the two zero restrictions are compatible with the data, allowing for sampling uncertainty.

Stata: test the two restrictions together

The full regression is still the active estimation result. Test the STR and Spend coefficients jointly.

```
test str spend
```

Read the result

Stata lists the two restrictions before reporting one F statistic and one p-value.

Stata output: one test of the two resource restrictions

```
. test str spend
 ( 1)  str = 0
 ( 2)  spend = 0

       F(  2,   416) =    5.43
            Prob > F =    0.0047
```

The two degrees of freedom answer different questions

$$F(q, \nu), \qquad q = 2, \qquad \nu = n - k_U - 1 = 420 - 3 - 1 = 416$$

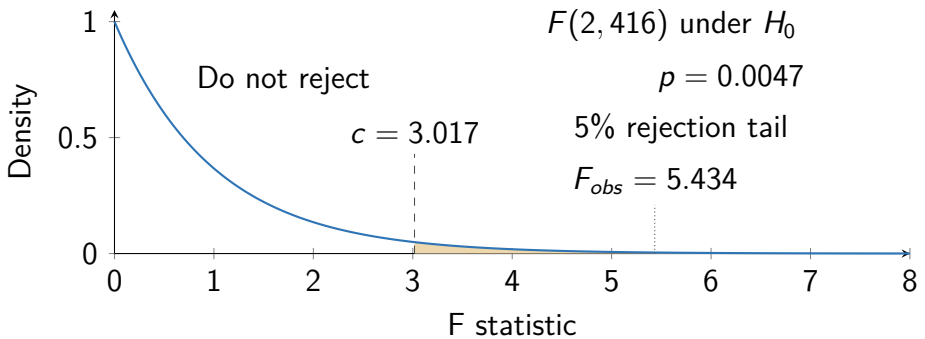
Numerator df: how many restrictions?

$q = 2$ counts $\beta_1 = 0$ and $\beta_3 = 0$. It does not count all the regressors in the model.

Denominator df: how much information remains?

$k_U = 3$ counts the unrestricted model's slopes. Estimating those slopes and the intercept uses four degrees of freedom; 416 remain for residual variation.

The F reference distribution and the rejection region



With numerator df $= 2$, the density is highest at zero and decreases. Reject in the right tail; the robust F reference is approximate.

Make the decision in the order of the question

1. State $H_0 : \beta_1 = \beta_3 = 0$; choose $\alpha = 0.05$.
2. Use $q = 2$ and $\nu = 416$ to find the cutoff: $c = 3.017$.
3. Compare: $F_{obs} = 5.434 > 3.017$. Reject H_0.
4. Equivalently, $p = 0.0047 < 0.05$. At least one resource coefficient differs from zero.

Limit of the conclusion

This does not establish that both slopes are nonzero, nor that either association is causal.

Your turn: identify the joint claim in another application

A rental model uses floor area (square metres), building age (years), and distance to a station (km):

$$Rent_i = \theta_0 + \theta_1 Area_i + \theta_2 Age_i + \theta_3 Distance_i + e_i.$$

An analyst claims: "After accounting for floor area, neither building age nor station distance is associated with rent."

1. Write the joint null and alternative. How many restrictions are imposed?
2. Which regressors remain in the restricted model?
3. With $n = 80$, identify both df. The test gives $F = 3.60$ and a 5% cutoff of 3.117. Make the decision; can you conclude that both slopes are nonzero?

Solution: rejection means at least one restriction fails

$$H_0 : \theta_2 = 0 \text{ and } \theta_3 = 0, \qquad q = 2$$

$$H_1 : \theta_2 \neq 0 \text{ or } \theta_3 \neq 0 \quad \text{(or both)}.$$

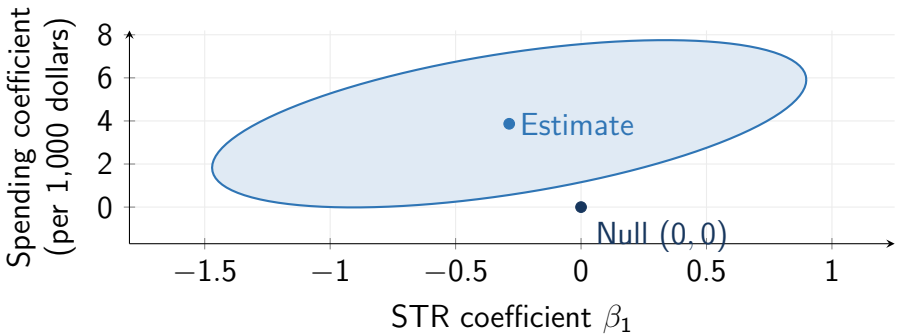
Restricted model

Regress rent on floor area and an intercept; omit age and station distance together.

Decision

$\nu = 80 - 3 - 1 = 76$. Since $3.60 > 3.117$, reject the joint null at 5%. At least one of the two restrictions is inconsistent with the data at the chosen level. The test alone does not identify which one.

A joint confidence set contains pairs, not single numbers



The shaded 95% set contains pairs not rejected by the joint test. It excludes $(0, 0)$, so the joint null is rejected.

Use the test to inform a specification, not choose it automatically

An analyst says: "STR has p = 0.553, so I will remove it and keep only the significant variables."

First ask what is being estimated

If the question is the class-size slope at fixed spending and EL, removing STR removes the coefficient of interest.

Then ask what the joint result establishes

The joint rejection says the two resource terms are not both zero. It does not prescribe which variable to drop or prove that all confounding has been removed.

Match the claim to the evidence

One specified change

The fixed-budget class-size contrast is a linear combination. Report its estimate and interval in score points.

Two restrictions together

“Neither STR nor spending matters is a joint hypothesis. Use the joint test's p-value, not the separate p-values as a substitute.

A range of pairs

A joint confidence set represents uncertainty about the two population coefficients together.