Instrumental Variables and Two-Stage Least Squares

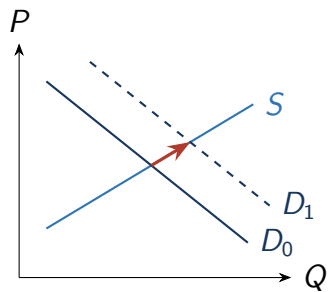
02 Nov 2026 | Chapter 12

ECO 351 – Introduction to Econometrics

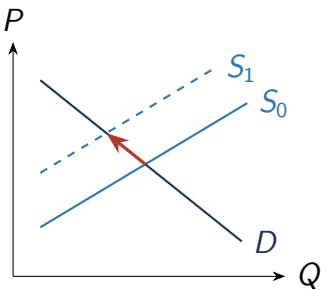
Shi Qiu – Fall 2026

Which price change helps us learn about demand?

A. Consumers want more cigarettes



B. A tax raises sellers' costs



In both cases price rises. Which movement follows the *same demand curve*? What must remain unchanged for that interpretation?

The target is a demand elasticity

For 48 US states in 1995, let Q_i be packs sold per person and P_i the real retail price, including tax.

$$\ln Q_i = \beta_0 + \beta_1 \ln P_i + u_i.$$

- β_1: the price elasticity in the structural demand equation.
- u_i: other demand determinants, such as smoking preferences.

A demand shock can raise both sales and equilibrium price:

$$\text{Cov}(\ln P_i, u_i) \neq 0.$$

We need price variation that moves consumers along demand, while leaving those other demand determinants unchanged.

A tax is useful only if it supplies the right price variation

Write $X = \ln P$, $Y = \ln Q$, and $Z =$ the real general-sales-tax component per pack.

$$Z\text{: tax} \longrightarrow X\text{: log price} \longrightarrow Y\text{: log sales}$$

Relevance	Tax predicts price: $\text{Cov}(Z, X) \neq 0$.
Exogeneity	Tax is unrelated to omitted demand determinants: $\text{Cov}(Z, u) = 0$.

Exclusion: in this demand model, tax must affect sales through price, without an additional demand channel.

Discuss: two states, two reasons for raising a tax

Hypothetical policy stories. Both tax changes raise retail prices.

State A	Raises a general sales tax to finance road maintenance. No smoking policy changes are reported.
State B	Raises the tax during an anti-smoking campaign and introduces new smoke-free rules.

1. For State B, describe one route from the policy change to sales that does *not* pass through price.
2. Is State A automatically a valid comparison? Name one fact you would investigate before accepting it.

Give a specific economic mechanism, not just a claim of statistical significance.

Debrief: name the extra route, then separate the two questions

For State B, the proposed comparison can include

policy package → campaign / restrictions ⟶ sales

This route can change demand even at the same price. A strong tax–price relationship does not remove it.

- **Economic argument:** why should the tax leave omitted demand determinants unchanged?
- **Empirical question:** how much does the tax actually move price?

We now examine the second question, while keeping the first as a maintained assumption.

First stage: tax and price

$$X_i = \pi_0 + \pi_1 Z_i + v_i.$$

π_1 measures the tax-price relationship. The data give:

```
. regress logp salestax, vce(robust)
Linear regression                               Number of obs     =         48
------------------------------------------------------------------------------
             |               Robust
        logp | Coefficient  std. err.      t    P>|t|     [95% conf. interval]
-------------+----------------------------------------------------------------
    salestax |   .0307289   .0048354     6.35   0.000     .0209956    .0404621
       _cons |   4.616546   .0289177   159.64   0.000     4.558338    4.674755
------------------------------------------------------------------------------
```

Which row provides evidence of relevance?

Reduced form: tax and sales

$$Y_i = \rho_0 + \rho_1 Z_i + e_i.$$

ρ_1 measures the tax-sales relationship:

```
. regress logq salestax, vce(robust)
Linear regression                               Number of obs     =         48
------------------------------------------------------------------------------
             |               Robust
        logq | Coefficient  std. err.      t    P>|t|     [95% conf. interval]
-------------+----------------------------------------------------------------
    salestax |  -.0332974   .0102541    -3.25   0.002    -.0539378    -.012657
       _cons |   4.717449   .0650586    72.51   0.000     4.586493    4.848405
------------------------------------------------------------------------------
```

Under the instrument restrictions, $\rho_1 = \beta_1 \pi_1$.

Why divide the sales response by the price response?

For the *same* tax change, the two regressions describe

$$\Delta\widehat{X} = \widehat{\pi}_1 \Delta Z, \qquad \Delta\widehat{Y} = \widehat{\rho}_1 \Delta Z.$$

If tax affects sales only through price, the population slopes satisfy

$$\rho_1 = \beta_1 \pi_1.$$

The single-instrument IV estimate is therefore

$$\widehat{\beta}_{IV} = \frac{\widehat{\rho}_1}{\widehat{\pi}_1} = \frac{-0.0332974}{0.0307289} \approx -1.084.$$

The tax units cancel. Under the instrument assumptions, a 1% tax-induced price increase implies about 1.084% lower sales in this model.

Try the same idea: a tutoring invitation

Constructed example. A university randomly invites students to extra tutoring. Attendance is voluntary.

	Invited	Not invited
Mean tutoring hours, X	12	8
Mean final score, Y	74	72

Let $Z = 1$ denote an invitation. Assume a linear model with a common effect of an extra tutoring hour.

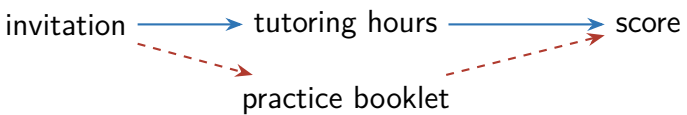
1. Use the invitation-related differences to estimate points gained per extra tutoring hour. Why is the score gap alone not that effect?
2. Now suppose only invited students also receive a practice booklet. Does random assignment of the invitation settle the IV argument?

Debrief: the invitation changes hours; hours change scores

$$\widehat{\pi}_1 = 12 - 8 = 4 \text{ hours}, \qquad \widehat{\rho}_1 = 74 - 72 = 2 \text{ points}.$$

$$\widehat{\beta}_{IV} = \frac{2 \text{ points}}{4 \text{ hours}} = 0.5 \text{ points per hour}.$$

The **2-point gap** is the invitation-related difference. The **0.5-point ratio** expresses that difference per induced tutoring hour.



The booklet can violate exclusion. Random assignment alone does not close that additional route.

Two-stage least squares implements this comparison

1. Predict log price using the tax.

$$X_i = \pi_0 + \pi_1 Z_i + v_i \quad \Longrightarrow \quad \widehat{X}_i = \widehat{\pi}_0 + \widehat{\pi}_1 Z_i.$$

$\widehat{X}_i$ retains the price variation associated with the instrument.

2. Relate log sales to that predicted component.

$$Y_i = b_0 + b_1 \widehat{X}_i + \text{second-stage residual}.$$

With one instrument and one endogenous regressor, b_1 equals the IV ratio we just calculated.

The two regressions explain the coefficient calculation. Use `ivregress 2sls` for its standard errors; an ordinary second regression does not generally give valid IV standard errors.

Stata: the IV estimate and uncertainty

```
. ivregress 2sls logq (logp = salestax), vce(robust)
Instrumental-variables 2SLS regression            Number of obs   =         48
------------------------------------------------------------------------------
             |               Robust
        logq | Coefficient  std. err.      z    P>|z|     [95% conf. interval]
-------------+----------------------------------------------------------------
        logp |  -1.083587   .3122036    -3.47   0.001    -1.695495    -.471679
       _cons |   9.719877   1.496143     6.50   0.000     6.78749    12.65226
------------------------------------------------------------------------------
Endogenous: logp
Exogenous:  salestax
```

`logq` is the outcome. Inside parentheses, `logp` is endogenous and `salestax` is the excluded instrument.

The estimated price elasticity is -1.084, with 95% interval approximately $[-1.695, -0.472]$.

OLS and IV use different price comparisons

The same 48-state data give

Estimator	Price coefficient	Robust SE
OLS	-1.213	0.195
IV using sales tax	-1.084	0.312

OLS uses all price variation. IV uses the price variation associated with the proposed instrument.

A larger IV standard error does not make OLS unbiased. A difference between the point estimates does not prove that the instrument is valid.

Why add income, and where does it enter?

Higher-income states may use a different tax mix, and income can also affect cigarette demand. A tax–income association can threaten the simple IV comparison.

Let W_i be log real income per person, treated here as an exogenous included control:

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 W_i + u_i,$$
$$X_i = \pi_0 + \pi_1 Z_i + \pi_2 W_i + v_i.$$

Stage 1 holds income fixed; stage 2 includes fitted price and the same income control.

The tax now needs relevant, valid variation *conditional on income*. Adding the control addresses a proposed channel; it does not prove the remaining instrument assumptions.

Stata: IV with income held fixed

```
. ivregress 2sls logq logincome (logp = salestax), vce(robust)
Instrumental-variables 2SLS regression            Number of obs   =         48
------------------------------------------------------------------------------
             |               Robust
        logq | Coefficient  std. err.      z    P>|z|     [95% conf. interval]
-------------+----------------------------------------------------------------
        logp |  -1.143375   .3604805    -3.17   0.002    -1.849904   -.4368463
   logincome |   .2145153   .3018477     0.71   0.477    -.3770953    .8061258
       _cons |   9.430658   1.219402     7.73   0.000     7.040675    11.82064
------------------------------------------------------------------------------
Endogenous: logp
Exogenous:  logincome salestax
```

The fitted price elasticity is about -1.143, conditional on log income and the maintained instrument assumptions.

First-stage evidence conditional on income

```
. estat firststage
  First-stage regression summary statistics
  ------------------------------------------------------------------------------
               |             Adjusted      Partial       Robust
      Variable |   R-sq.       R-sq.        R-sq.       F(1,45)   Prob > F
  -------------+----------------------------------------------------------------
          logp |  0.6389      0.6228       0.5009       44.7305    0.0000
  ------------------------------------------------------------------------------
```

The excluded-instrument robust $F(1, 45)$ is about 44.73.

This is evidence about relevance. It cannot establish exclusion or exogeneity. Weak instruments can make 2SLS unstable and routine confidence intervals unreliable.

Discussion: write the conclusion the evidence supports

A policy analyst asks: *"What did this study find, and why should I trust its price comparison?"*

IV price coefficient, controlling for income	-1.143
95% confidence interval	$[-1.850, -0.437]$
Excluded-instrument first-stage $F(1, 45)$	44.73

Prepare a three-sentence briefing:

1. Interpret the price coefficient, with the held-fixed variable.
2. Explain what the first-stage result tells us.
3. State one causal assumption and a concrete way it could fail.

Do not treat a small p-value as evidence that every identifying assumption holds.

Debrief: estimate, first stage, identifying argument

One defensible briefing

Holding income fixed, the IV estimate implies about **1.14% lower cigarette sales for a 1% tax-induced price increase**, under the maintained demand model and instrument assumptions.

The first-stage statistic of **44.73** provides evidence that the excluded tax instrument predicts price after income is included.

The causal interpretation still requires tax variation unrelated to omitted demand shifts, with no additional demand channel; accompanying smoking restrictions are one potential threat.

The 95% interval for the price elasticity is about $[-1.85, -0.44]$. It describes sampling uncertainty *conditional on the model and instrument assumptions*.

IV in an applied project

Before using IV, state the endogenous regressor, the excluded instrument, and the economic argument for its validity. Then examine the first stage and report the IV estimate with uncertainty.

IV is available as an optional project method. A carefully justified regression remains preferable to an unconvincing instrument.

Next class: repeated observations offer another way to address confounding that stays constant within a place over time.