LECTURE 4.2

Specification and Data Issues

ECON30130 Econometrics I

Sam Deegan
University College Dublin

Autumn 2027

Where We Are

1. Foundations

- 1.1 Data and Variation
- 1.2 Conditional Means
- 1.3 Sampling Variation

2. Simple Regression

- 2.1 Least Squares
- 2.2 Units and Fit
- 2.3 Bias and Precision

3. Multiple Regression

- 3.1 Estimation
- 3.2 Inference
- 3.3 Functional Form
- 3.4 Dummies

4. Broken Assumptions

- 4.1 Heteroskedasticity
- **4.2 Specification**

The Last Lecture of Part I
Lecture 4.1 covered heteroskedasticity. Today covers a wrong shape, which the residual plot shows, and four faults it never reveals.

Last Week's Lecture

Lecture 4.1 sorted three violations of Lecture 2.3's assumptions by whether each costs unbiasedness or precision:

1. **Heteroskedasticity.**
   The error variance changes with schooling, and the estimate stays centred on 1.30.

2. **Robust Standard Errors.**
   Squared residuals replace the unseen variances, with HC2 or HC3 in small samples.

3. **Serial Correlation.**
   Correlated errors leave the estimate centred too, and a missing variable is sought first.

Reasons 1.5227 Might Differ from the True 1.30

Our 120 give 1.5227 euro an hour a year; the true return is 1.30.

> **Question**
> What reasons has this module given you for doubting 1.5227? Answer from memory, with nothing looked up.

Lecture Plan

1. **Functional Form and Outliers.**
   A wrong shape costs unbiasedness, and the residual plot shows it.
2. **Ability, Selection and Measurement Error.**
   Ability raises the estimate above 1.30; reporting noise pulls it toward zero.
3. **Reverse Causality.**
   Regression is symmetric, causation is not, and no line of the output records which direction you meant.
4. **The Research Designs of Part II.**
   Instruments, differencing and cutoffs each answer a fault no diagnostic detects.

SECTION 1

Functional Form and Outliers

The Line Sits Too Low Mid-Range and Too High at the Ends

Returns to schooling diminish, so the curve is concave and our line is straight:

- **Mid-Range:** the curve runs above the line, so the errors there average positive.
- **Both Ends:** the curve runs below the line, so the errors there average negative.

The error therefore does not average zero at every level of schooling.

Why the Wrong Functional Form Costs Unbiasedness

An error whose average changes with schooling breaks zero conditional mean.

Assumption Four Fails and Unbiasedness Goes

Lecture 2.3 names assumption four the only one that delivers unbiasedness. With assumption four broken, the estimate is no longer centred on the true 1.30. No standard error re-centres it, because assumption four is what failed.

Heteroskedasticity leaves the estimate centred on 1.30; a wrong shape does not.

Curvature and Outliers in the Residual Plot

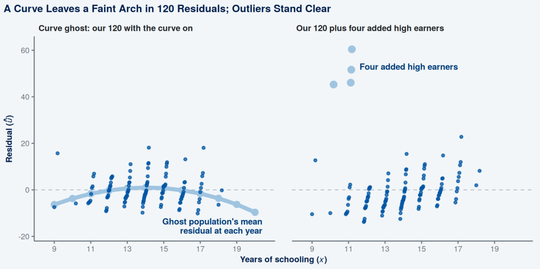

Figure 1: Left: our 120 in a counterfactual curve ghost. Right: our 120 plus four high earners.

> **An Arch Means the Wrong Form, Four Points Mean Outliers**
> In 120 people the arch is faint; the pale line, the ghost population's mean residual at each year, shows it plainly. Summary numbers hide both faults, and the residual plot shows them (Anscombe, 1973).

RESET Adds Squared and Cubed Fitted Values to the Regression

RESET (Ramsey, 1969), the named test for that arch, runs in two steps:

- **Add:** the squared and cubed fitted values join schooling in the regression.
- **Test:** an F-test asks whether those two terms together raise the fit.

> **A Significant F Suggests the Wrong Shape**
> A significant F suggests the straight line is the wrong shape (Wooldridge, 2020, p. 297).

RESET Tests the Shape of the Line and Nothing Else

RESET detects a wrong shape and cannot detect two other faults (Wooldridge, 2020, p. 298):

- **A Missing Variable:** leave out a variable whose mean is itself a straight-line function of the regressors already in the model, and RESET fails to reject, however large the sample.
- **Heteroskedasticity:** with the shape right, RESET does not detect it.

> **The Test Names No Fault**
> The residual plot comes first, and RESET second. A test run first sends you correcting a fault nobody has named.

Which Fault Costs the Estimate and Which Costs the Interval

Most people worry more about the standard error, which has a test attached:

Heteroskedasticity Costs the Interval
The printed standard error no longer tells you how far the slope typically sits from 1.30.

A Wrong Shape Costs the Estimate
The estimate itself is no longer centred on 1.30.

King & Roberts (2015) read a large gap between classical and robust standard errors as evidence of misspecification.

What to Do About the Four Highest Earners

The right-hand panel showed the sample's four highest earners, far above the residual cloud.

> **Question**
> What is your next step with those four people? What would you have to believe about them to justify it?

Look at Who the Four Are Before Changing the Fit

Nothing on the plot separates a wrong line from a real tail of high earners:

- **A Robust Standard Error:** changes the interval and leaves the estimate where it was.
- **Dropping Them:** Lecture 2.3 showed that deleting the highest earner biases the slope.
- **Leaving Them In:** their pull on the slope grows with their distance from 13.45 years.

Belief About the Four Decides the Step

No test tells you whether the four are errors or a real tail. Find out who the four are first.

SECTION 2

Ability, Selection and Measurement Error

The Same Generator with Its Switches Turned On

Your 120 rows came from this generator with every switch off. Ability (a_i) now raises actual schooling (x_i^*) and pay:

$$x_i^* = \pi_0 + \pi_1 a_i + v_i \qquad y_i = \beta_0 + \beta_1(x_i^* - x_0) + \gamma a_i + u_i \tag{1}$$

- π_0, π_1: average schooling, and how much ability raises it
- β_0, β_1: earnings at x_0, and the true return to a year
- x_0: the reference level of schooling, thirteen years in our population
- γ: how much ability raises pay; v_i: everything else in schooling

The generator sets every value. Nobody observes a_i or x_i^*.

Reported Schooling and the Wage Floor

Least squares runs on reported schooling (x_i), actual schooling plus reporting error (m_i):

$$x_i = x_i^* + m_i \tag{2}$$

A person appears in the data only if earnings clear the wage floor (f):

$$y_i \geq f \tag{3}$$

The Regression Never Sees Actual Schooling

Earnings are built from x_i^*, and the regression is run on x_i.

The Generator as a Causal Diagram

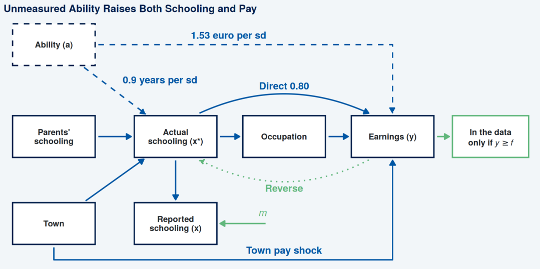

Figure 2: Ability dashed: no survey measures it. Green: switches, all off for our 120.

Two Arrows Out of Unmeasured Ability

Ability raises schooling (π_1, 0.9 years per sd) and earnings (γ, 1.53 euro per sd). The coefficient on schooling therefore includes part of ability's effect on earnings.

The Estimated Return Rises Above 1.20 as Ability’s Effects Grow

Figure 3: Sweep ghost of ability’s pay effect: our 120 solid, population pale, 1.20 dotted.

> **Ability Pushes the Estimate Above 1.20**
> Lecture 3.1 found the same upward bias in our own 120 and the same population: 1.5744 without ability, 1.3935 with it.

Dropping People by Their Schooling Costs Only Sample Size

Dropping people on their schooling would cost nothing but sample size:

- **The Reason:** at each year of schooling the average earnings are the same in that subgroup as in the whole population (Wooldridge, 2020, p. 315).
- **The Wage Floor:** cuts on earnings instead, the variable the regression explains (Heckman, 1979).

> **Exogenous and Endogenous Sample Selection**
> Wooldridge calls selection on schooling exogenous and selection on earnings endogenous. Only selection on earnings biases least squares.

Sample Selection Caused by the Wage Floor

Figure 4: The 16 euro floor applied to our 120: 102 survivors solid, the 18 dropped pale.

> **The Wage Floor Selects on Earnings Before the Regression Runs**
> Least squares describes only the people who cleared the floor. At low schooling only large positive errors clear it, so the survivors' line is flatter: 1.07 against 1.57. In the whole population the floor takes 1.60 to 0.99.

Reporting Noise Adds to the Variance and Not to the Covariance

The slope is a covariance over a variance, 965.3 over 633.9 in our 120:

- **The Covariance on Top:** reporting noise is unrelated to earnings, so the covariance does not change.
- **The Variance Underneath:** the same noise, unrelated to actual schooling, adds its own variance.

> **Reporting Noise Pulls the Estimate Toward Zero**
> The same numerator over a bigger denominator gives a smaller slope. Noise in schooling therefore attenuates the estimate toward zero.

The Regression Recovers a Share of the True Return

Least squares recovers the true return (β_1) times a share below one:

$$\beta_1 \times \underbrace{\frac{\sigma^2_{x^*}}{\sigma^2_{x^*} + \sigma^2_m}}_{\text{the reliability ratio}} \tag{4}$$

- $\sigma^2_{x^*}$: the variance of actual schooling
- σ^2_m: the variance of the reporting error

Neither variance comes from our 120; reported schooling's sample variance is 5.33.

Why the Reliability Ratio Shrinks the Estimate

With no reporting error the ratio is one and the target is 1.30. Reporting error as variable as actual schooling halves the ratio:

$$\beta_1 \times \frac{\sigma_{x^*}^2}{\sigma_{x^*}^2 + \sigma_{x^*}^2} = \beta_1 \times \frac{1}{2} \qquad (5)$$

More Data Cannot Cure Attenuation

More data only narrows the spread around 0.65. The printed standard error gives no sign of the attenuation.

Worked example: our population

How the Estimate Falls as Reporting Error Grows

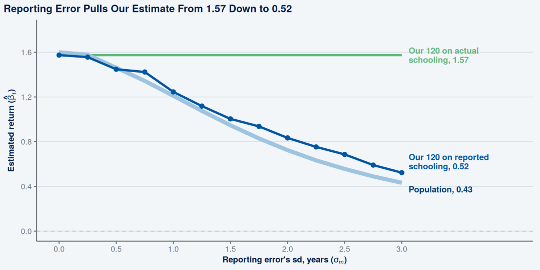

Figure 5: Reporting-error sweep ghost: our 120 solid, population pale, actual schooling green.

> **The Estimate Falls Toward Zero**
> More reporting error raises only the ratio's denominator, so the estimate falls toward zero: at an sd of three years, from 1.57 to 0.52 in our 120 and from 1.60 to 0.43 in the population.

Three Cases Where Attenuation Does Not Apply

Attenuation assumes classical measurement error, argued over since Griliches (1977). Three cases fall outside that result:

- **Noise in Earnings:** the slope is unbiased, and only the standard error grows.
- **Noise in a Control:** the bias can run either way, on any coefficient.
- **Non-Classical Error:** people who round their schooling up are not a random selection (Bound et al., 2001).

> **Differencing Twins Attenuates More, Not Less**
> Differencing two siblings removes everything the family shares, and most of the real schooling difference. The reporting error survives, so the differenced estimate is more attenuated (Ashenfelter & Krueger, 1994).

Self-Reported Schooling Against the School Register

One survey takes schooling from people's own reports, the other from the school register.

> **Question**
> Which survey reports the smaller return? Would that gap repeat in a new sample?

The Self-Reported Survey Gives the Smaller Return

Self-reported schooling carries reporting error, so its return sits closer to zero than the register's.

> **More Data Keeps the Gap**
> The shortfall is in the target, so new samples scatter around it again. A larger sample narrows both spreads, and the gap between the targets stays.

SECTION 3

Reverse Causality

The Same Data Fitted in Both Directions

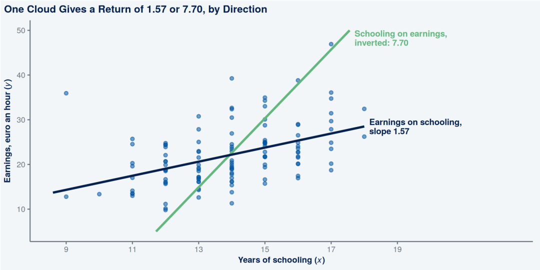

Figure 6: Our 120: earnings on schooling in navy; schooling on earnings, inverted, in green.

> **One Cloud, Two Least-Squares Lines**
> Least squares does not know which variable you meant to explain. One fit minimises vertical distances, the other horizontal ones.

The Algebra Behind the Two Fitted Slopes

The fits share a numerator, not a denominator, so the slopes are not reciprocals:

$$\hat{\beta}_{y|x} = \frac{s_{xy}}{s_x^2} \qquad \hat{\beta}_{x|y} = \frac{s_{xy}}{s_y^2} \tag{6}$$

- $\hat{\beta}_{y|x}$: the slope from regressing earnings on schooling
- $\hat{\beta}_{x|y}$: the slope from regressing schooling on earnings

> **The Product Is the R-Squared, Not One**
> Our 120 give 1.5227 and 0.2319. Inverting 0.2319 gives 4.31 euro an hour a year, not 1.5227. The two slopes multiply to 0.3531, the R-squared.

Simultaneity: When Causation Runs in Both Directions

Sometimes each variable causes the other at the same time.

> **Each Observed Price and Quantity Is One Crossing**
> Price and quantity are settled together, where demand and supply cross (Haavelmo, 1943). Each observed pair is one crossing of the two curves.

A regression of quantity on price therefore traces out neither curve (Working, 1927).

Earnings Expectations Can Cause Schooling

Two causal arrows can join schooling and earnings at once:

- **Schooling to Earnings:** a year of school raises pay, the return we want.
- **Earnings to Schooling:** expecting higher pay keeps a person in school longer.

> **The Fitted Slope Carries Both Directions**
> A regression of earnings on schooling mixes the two arrows into one slope.

The Regression Output Is Identical Either Way

No line of the regression table separates the two fits:

- **R-Squared and t:** 0.3531 and the same t-statistic, whichever variable sits on the left.
- **The Residual Plot:** clean in both directions.
- **A Larger Sample:** narrows both intervals and still names no direction.

> **Direction Comes from Outside the Data**
> Timing, institutions and what cannot have caused what settle the direction. Part II adds designs that fix the direction before the estimate is computed.

Police Numbers and Recorded Crime

Areas with more police officers per head have more recorded crime.

> **Question**
> Which way does the causation run? Could a regression on those two columns settle the direction?

Police and Recorded Crime Can Each Cause the Other

Three stories fit the same positive correlation:

- **Crime to Police:** areas with more crime hire more officers.
- **Police to Recorded Crime:** more officers record more of the crime that happens.
- **A Confounder:** a city's size and poverty raise both columns.

> **No Regression on Two Columns Settles the Direction**
> The two columns give one R-squared and one t-statistic, whichever story is true.

Ability, Mismeasurement, Selection and Reverse Causality Leave No Pattern

The residual plot separates two kinds of fault:

- **Visible:** heteroskedasticity, serial correlation and the wrong shape each leave a pattern.
- **Invisible:** ability, mismeasurement, selection and reverse causality leave none.

> **No Diagnostic Catches Any of These Four Faults**
> These four faults are in how the data arose, not in the residuals. Part II answers those four faults by designing the comparison instead.

SECTION 4

The Research Designs of Part II

Instrumental Variables Answer an Unmeasured Confounder

Ability is unmeasured, so an instrument uses schooling differences unrelated to ability:

- **The Example:** compulsory-schooling laws keep some people in school longer (Angrist & Krueger, 1991).
- **The Assertion:** those laws reach wages only through schooling, the exclusion restriction.
- **The Dispute:** economists argue over which instruments satisfy it (Card, 1999).

> **An Instrument Reaches Earnings Only Through Schooling**
> An instrument changes schooling and reaches earnings no other way. The estimator uses only the part of schooling that the instrument changes.

Difference-in-Differences and Fixed Effects: Constant Confounders

Differencing removes any unmeasured cause that stays constant across the comparison:

- **A State Border:** Card & Krueger (1994) compared fast-food jobs on both sides after one state raised its minimum wage.
- **Twins:** Ashenfelter & Krueger (1994) compared siblings with the same parents and different schooling.

> **A Confounder That Changes over Time Survives**
> What is differenced away never has to be measured. A confounder that changes over time survives the differencing.

Regression Discontinuity Uses a Rule That Assigns Treatment

A published threshold decides who is treated.

> **The Jump at the Cutoff Is the Effect**
> People just above and just below the threshold differ only in treatment. The survey to start from is Lee & Lemieux (2010).

Thistlethwaite & Campbell (1960) invented the design to evaluate scholarships awarded on a test score. Angrist & Lavy (1999) found it in a rule splitting classes above forty pupils.

Each Design Answers Only About the People It Compares

All three designs find people who are alike except for the treatment:

- **The Comparison:** the two groups differ in treatment and in nothing else that moves earnings.
- **The Cost:** the discontinuity answers for people at the cutoff, and Ashenfelter's twins answer for siblings.
- **The Estimator:** least squares still fits the line. Each design changes which people go into it.

> **Credibility Moves from the Model to the Data**
> A design is judged on how the data arose, not on the error term's assumptions. Leamer (1983) made that argument. Angrist & Pischke (2010) turned that argument into a research programme, and Card (2022) gave it a Nobel lecture.

Why These Designs Are Left Until Part II

Angrist & Pischke (2017) argue that assumptions and corrections leave students unable to read applied work. Most first courses still teach them (Harter & Asarta, 2022).

> **Judging an Instrument Needs a Standard Error First**
> Part I showed you what a failing regression looks like before any design arrives.

Three Replies to Angrist and Pischke

Three economists answer Angrist and Pischke:

- **Sims:** a design reaches only the questions an experiment can reach (Sims, 2010).
- **Keane:** a design rests on assumptions too, and leaves them unstated (Keane, 2010).
- **Deaton:** an estimate you can defend in one setting may not hold in another population (Deaton, 2010).

A Grant Paid Below a Points Cut-Off

The state pays a maintenance grant to every student below a published cut-off.

> **Question**
> Which design would show whether the grant raises graduation rates? What must be true of students either side of the cut-off?

A Published Cut-Off Makes the Grant a Regression Discontinuity

The design needs students just either side of the cut-off to match:

- **The Assumption:** students just above and just below are alike except for the grant.
- **Choosing a Side:** appealing a points total, re-sitting, or picking a course with another cut-off.

> **Nobody May Choose Their Side of the Line**
> A student who re-sits to cross the cut-off has chosen a side. The two sides then differ in more than the grant.

Returning to the Question from Lecture 1.1

The module opened on a question you could not answer in September:

> **Question**
> How much more does a person earn per hour, on average, for each extra year they stayed in school?

What Part I Settled About 1.5227

A causal reading of 1.5227 needs three things, which Lecture 1.2 named:

- **Mechanism:** no lecture in Part I supplied a mechanism, and no design supplies one either.
- **Rival Stories:** Lecture 3.1 controlled for ability, which only a simulated population records.
- **Chance:** Lecture 1.3 measured chance, and the standard error on 1.5227 is 0.1897.

> **Economics Supplies the Mechanism**
> A design settles who is compared, not why schooling changes their earnings. Why a year of school pays stays yours to argue.

What Would Have to Be True?

Today the standing question is put to somebody else's published estimate:

> **Question**
> What would have to be true for this number to mean what we just said it means? And what is the plausible story where it isn't?

Three Questions to Put to the Published Estimate

Answer three questions about the paper's estimate:

1. What is the number, in its units, and what does its sentence claim?
2. Which of this term's faults is the most plausible story where it isn't?
3. What design would answer it, and what would that design cost you?

Turning the Switches on and Watching Every Diagnostic Pass

The app turns ability, reporting error and the wage floor on one at a time.

> **Question**
> Dial the reporting error up. Does the estimate gain sampling noise, attenuate toward zero, or both?

Reporting Error Attenuates the Estimate and Narrows Its Spread

Two runs of the app answer the question:

- **Reporting Error Up:** the estimate falls toward zero, and its spread across samples narrows.
- **Ability On:** every diagnostic from today and Lecture 4.1 passes while the estimate sits above 1.30.

> **A Clean Residual Plot Clears No Confounder**
> The residual plot from a confounded regression looks like one from a clean regression.

Summary: Specification and Data Issues

The residual plot shows a wrong shape and hides four other faults:

1. **Functional Form and Outliers.**
   A wrong shape costs unbiasedness, and RESET tests for nothing else.
2. **Ability, Selection and Measurement Error.**
   Each biases the estimate, and none shows in the residuals.
3. **Reverse Causality.**
   No line of the regression output records which direction you meant.
4. **The Research Designs of Part II.**
   Each design exists because a fault cannot be diagnosed, only designed around.

Before Part II

1. **Play:**
   The app with every switch on at once. Check whether any diagnostic you know catches any of the four hidden faults.
2. **Apply:**
   Problem set 4, covering 4.1 and 4.2, ending on the standing question.
3. **Review:**
   Wooldridge (2020), sections 9-1, 9-4, 9-5 and 16-1. Then weigh Angrist & Pischke (2017) against Sims (2010), in two hundred words.

Next Lecture

Part II adds designs that give a reason to believe the comparison:

- **Instrumental Variables:** the estimator arrives, and with it the exclusion restriction an instrument has to satisfy.
- **Difference-in-Differences and Panel Data:** you state parallel trends, and watch that assumption fail.
- **Regression Discontinuity:** bandwidths arrive, with the assumption that nobody chose their side.

Each design rests on an assertion about the data that no test checks.

Thank you

sam.deegan@ucdconnect.ie
sam-deegan.com

BACKUP

Appendix

Notation Used in This Lecture

Symbol	What it is	First met
i	the index running over the people in the sample	Lecture 1.1
x_i	reported schooling of person i	Lecture 1.1
y_i	hourly earnings of person i, in euro	Lecture 1.1
s_x^2	the sample variance of schooling	Lecture 1.1
s_y^2	the sample variance of earnings	Lecture 1.1
s_{xy}	the sample covariance of schooling and earnings	Lecture 1.1
u_i	person i's error	Lecture 2.1
x_0	the reference level of schooling, thirteen years	Lecture 2.1
β_0	earnings at the reference level of schooling	Lecture 2.1
β_1	the true return to a year of schooling, 1.30	Lecture 2.1

Notation Used in This Lecture

Symbol	What it is	First met
a_i	ability	The Same Generator With Its Switches Turned On
x_i^*	actual schooling	The Same Generator With Its Switches Turned On
π_0	average schooling	The Same Generator With Its Switches Turned On
π_1	how much ability raises schooling	The Same Generator With Its Switches Turned On
γ	how much ability raises pay	The Same Generator With Its Switches Turned On
v_i	everything else that raises or lowers schooling	The Same Generator With Its Switches Turned On

Symbol	What it is	First met
m_i	the reporting error	Reported Schooling and the Wage Floor
f	the wage floor	Reported Schooling and the Wage Floor
$\sigma_{x^*}^2$	population variance of actual schooling	The Regression Recovers a Share of the True Return
σ_m^2	population variance of the reporting error	The Regression Recovers a Share of the True Return

Symbol	What it is	First met
$\hat{\beta}_{y\|x}$	the slope from regressing earnings on schooling	The Algebra Behind the Two Fitted Slopes
$\hat{\beta}_{x\|y}$	the slope from regressing schooling on earnings	The Algebra Behind the Two Fitted Slopes

Worked Example: Halving the Reliability Ratio

Reporting error as variable as actual schooling ($\sigma_m^2 = \sigma_{x^*}^2$) halves the reliability ratio:

$$\beta_1 \times \frac{\sigma_{x^*}^2}{\sigma_{x^*}^2 + \sigma_{x^*}^2} = \beta_1 \times \frac{1}{2} \tag{7}$$

Our population's true return, 1.30, replaces β_1:

$$1.30 \times \frac{1}{2} = 0.65 \tag{8}$$

Back to the slide

Derivation: Attenuation from Classical Measurement Error

Write the slope as a covariance over a variance. Substitute reported schooling for actual schooling, then find the term the reporting error changes.

The Noise Enters the Denominator, Not the Numerator

Substitute $x = x^* + m$ into the numerator. Independence leaves the covariance with earnings unchanged. In the denominator the two variances add.

[Full derivation. Set as homework; worked solution published with the set.]

Notes and Tools

- **The Switches:** `_apps/regression/R/model.R`, block B_04, and what each one does.
- **The App:** the measurement-error sweep, and the script behind it.
- **The Residual Plot:** `plot(model, which = 1)` draws residuals against fitted values.
- **The Lecture in Two Lines:** `lm(y ~ x)` beside `lm(x ~ y)` on the same data.
- **This Lecture's Lab:** turn the switches on yourself and try to catch each fault with a diagnostic.

References

Angrist, J. D., & Krueger, A. B. (1991). Does compulsory school attendance affect schooling and earnings? *Quarterly Journal of Economics*, *106*(4), 979–1014. https://doi.org/10.2307/2937954

Angrist, J. D., & Lavy, V. (1999). Using Maimonides' rule to estimate the effect of class size on scholastic achievement. *Quarterly Journal of Economics*, *114*(2), 533–575. https://doi.org/10.1162/003355399556061

Angrist, J. D., & Pischke, J.-S. (2010). The credibility revolution in empirical economics: How better research design is taking the con out of econometrics. *Journal of Economic Perspectives*, *24*(2), 3–30. https://doi.org/10.1257/jep.24.2.3

Angrist, J. D., & Pischke, J.-S. (2017). Undergraduate econometrics instruction: Through our classes, darkly. *Journal of Economic Perspectives*, *31*(2), 125–144. https://doi.org/10.1257/jep.31.2.125

Anscombe, F. J. (1973). Graphs in statistical analysis. *The American Statistician*, *27*(1), 17–21. https://doi.org/10.1080/00031305.1973.10478966

Ashenfelter, O., & Krueger, A. (1994). Estimates of the economic return to schooling from a new sample of twins. *American Economic Review*, *84*(5), 1157–1173.

References

Bound, J., Brown, C., & Mathiowetz, N. (2001). Measurement error in survey data. In J. J. Heckman & E. E. Leamer (Eds.), *Handbook of econometrics* (Vol. 5, pp. 3705–3843). Elsevier. https://doi.org/10.1016/S1573-4412(01)05012-7

Card, D. (1999). The causal effect of education on earnings. In O. C. Ashenfelter & D. Card (Eds.), *Handbook of labor economics* (Vol. 3A, pp. 1801–1863). Elsevier. https://doi.org/10.1016/S1573-4463(99)03011-4

Card, D. (2022). Design-based research in empirical microeconomics. *American Economic Review*, *112*(6), 1773–1781. https://doi.org/10.1257/aer.112.6.1773

Card, D., & Krueger, A. B. (1994). Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *American Economic Review*, *84*(4), 772–793.

Deaton, A. (2010). Instruments, randomization, and learning about development. *Journal of Economic Literature*, *48*(2), 424–455. https://doi.org/10.1257/jel.48.2.424

Griliches, Z. (1977). Estimating the returns to schooling: Some econometric problems. *Econometrica*, *45*(1), 1–22. https://doi.org/10.2307/1913285

Haavelmo, T. (1943). The statistical implications of a system of simultaneous equations. *Econometrica*, *11*(1), 1–12. https://doi.org/10.2307/1905714

References

Harter, C., & Asarta, C. J. (2022). Teaching methods in undergraduate intermediate theory, statistics and econometrics, and other upper-division economics courses: Results from a sixth national quinquennial survey. *The American Economist*, *67*(1), 132–146. https://doi.org/10.1177/05694345211037904

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica*, *47*(1), 153–161. https://doi.org/10.2307/1912352

Keane, M. P. (2010). Structural vs. Atheoretic approaches to econometrics. *Journal of Econometrics*, *156*(1), 3–20. https://doi.org/10.1016/j.jeconom.2009.09.003

King, G., & Roberts, M. E. (2015). How robust standard errors expose methodological problems they do not fix, and what to do about it. *Political Analysis*, *23*(2), 159–179. https://doi.org/10.1093/pan/mpu015

Leamer, E. E. (1983). Let's take the con out of econometrics. *American Economic Review*, *73*(1), 31–43.

Lee, D. S., & Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of Economic Literature*, *48*(2), 281–355. https://doi.org/10.1257/jel.48.2.281

Ramsey, J. B. (1969). Tests for specification errors in classical linear least-squares regression analysis. *Journal of the Royal Statistical Society: Series B (Methodological)*, *31*(2), 350–371. https://doi.org/10.1111/j.2517-6161.1969.tb00796.x

References

Sims, C. A. (2010). But economics is not an experimental science. *Journal of Economic Perspectives*, *24*(2), 59–68. https://doi.org/10.1257/jep.24.2.59

Thistlethwaite, D. L., & Campbell, D. T. (1960). Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational Psychology*, *51*(6), 309–317. https://doi.org/10.1037/h0044319

Wooldridge, J. M. (2020). *Introductory econometrics: A modern approach* (7th ed.). Cengage Learning.

Working, E. J. (1927). What do statistical demand curves show? *Quarterly Journal of Economics*, *41*(2), 212–235. https://doi.org/10.2307/1883501