Econometrics II

Experiments and Quasi-Experiments

Lasha Chochua

2026

Introduction

Why Study Experiments in an Econometrics Course?

- In medicine and psychology, causal effects are routinely estimated by **randomized controlled experiments** (RCTs): some subjects randomly get the treatment, others a placebo.

- Three reasons economists study RCTs:
  - **Benchmark.** An ideal RCT is the conceptual standard against which we judge causal estimates from observational data.
  - **Influence and limits.** Real experiments, when run, are influential – so we must understand their threats to validity, not just their strengths.
  - **As-if randomness.** External events sometimes mimic random assignment, producing a **quasi-experiment** (natural experiment) we can analyze with the same toolkit.

Program Evaluation and the Plan

- **Program evaluation:** estimating the effect of a program, policy, or intervention ("treatment"). Examples:
  - Effect on earnings of a job-training program?
  - Effect on employment of a higher minimum wage?
  - Effect on college attendance of low-cost student loans?
- The statistical tools are familiar – OLS, panel data, IV. What is new is the **type of data** and the **special opportunities and challenges** of experiments and quasi-experiments.
- Roadmap: potential outcomes $\rightarrow$ threats to RCTs $\rightarrow$ Project STAR $\rightarrow$ quasi-experiments $\rightarrow$ their threats $\rightarrow$ heterogeneous effects and LATE.

Potential Outcomes and Causal Effects

The Thought Experiment

- Should you take a drug, enroll in a training program, or do the optional problem set?
- Imagine **two parallel worlds** for individual i:
  - one where i **receives** the treatment, with outcome $Y_i(1)$;
  - one where i **does not**, with outcome $Y_i(0)$.

Definition 1: Potential outcomes and the individual causal effect

A **potential outcome** is the outcome for an individual under a potential treatment. For a binary treatment $X_i \in \{0, 1\}$, the two potential outcomes are $Y_i(1)$ and $Y_i(0)$. The **causal effect for individual** i is

$$\beta_{1i} = Y_i(1) - Y_i(0).$$

The Fundamental Problem of Causal Inference (1/2)

- The causal effect $\beta_{1i} = Y_i(1) - Y_i(0)$ generally **varies across individuals** (e.g. a drug's effect depends on age, smoking, health).

The fundamental problem of causal inference

For any individual we observe **only one** potential outcome – $Y_i(1)$ if treated, $Y_i(0)$ if not – never both. So β_{1i} can **never** be measured for a single individual.

The Fundamental Problem of Causal Inference (2/2)

- The way out: we often do not need each β_{1i}, only its population **mean**.

Definition 2: Average treatment effect

The **average causal effect** (or **average treatment effect**, ATE) is the population mean of the individual causal effects:

$$\text{ATE} = E[Y_i(1) - Y_i(0)] = E[\beta_{1i}].$$

The Switching Equation: What We Actually Observe

- We never see $Y_i(1)$ and $Y_i(0)$ together. We see the **realized** outcome Y_i, linked to the potentials by the **switching equation**:

$$Y_i = Y_i(1)\,X_i + Y_i(0)\,(1 - X_i). \tag{1}$$

- Reading (1):
  - if $X_i = 1$: $Y_i = Y_i(1)$ (treated outcome observed, control outcome **missing**);
  - if $X_i = 0$: $Y_i = Y_i(0)$ (control outcome observed, treated outcome **missing**).

- The missing potential outcome is the **counterfactual**. All of causal inference is a **missing-data problem**: how do we stand in for the counterfactual we cannot see?

The Naive Comparison and Its Bias (1/2)

- A tempting estimator: compare mean outcomes of the treated and the untreated.
- Using (1), this **observed** difference equals

$$E[Y_i \mid X_i = 1] - E[Y_i \mid X_i = 0] = E[Y_i(1) \mid X_i = 1] - E[Y_i(0) \mid X_i = 0].$$

The Naive Comparison and Its Bias (2/2)

- Add and subtract the counterfactual $E[Y_i(0) \mid X_i = 1]$ (mean control outcome **for the treated**):

$$\underbrace{E[Y_i \mid X_i{=}1] - E[Y_i \mid X_i{=}0]}_{\text{naive difference}} = \underbrace{E[Y_i(1){-}Y_i(0) \mid X_i{=}1]}_{\text{ATT}} + \underbrace{E[Y_i(0) \mid X_i{=}1] - E[Y_i(0) \mid X_i{=}0]}_{\text{selection bias}}$$

Note

The naive difference = effect on the treated (ATT) + **selection bias**. Selection bias is the difference in *untreated* potential outcomes between groups – nonzero whenever who gets treated is related to how they would have fared anyway.

Randomization Eliminates Selection Bias (1/2)

- Random assignment of X_i makes treatment status **independent** of personal attributes, and in particular of the potential outcomes:

$$\big[\, Y_i(1),\, Y_i(0) \,\big] \;\perp\; X_i.$$

- Two consequences:
  - **Selection bias** $= 0$**:** $E[Y_i(0) \mid X_i{=}1] = E[Y_i(0) \mid X_i{=}0]$, since groups are comparable in their potentials.
  - **ATT** $=$ **ATE:** the treated are a random sample, so their average effect equals the population average effect.

Randomization Eliminates Selection Bias (2/2)

Why the experiment works

Random *selection* from the population + random *assignment* to treatment $\Rightarrow$ the expected difference in observed outcomes between the treatment and control groups equals the average causal effect in the population.

Identification Under Random Assignment

- Formally, with $[Y_i(1), Y_i(0)] \perp X_i$:

$$\begin{aligned} E[Y_i \mid X_i = 1] - E[Y_i \mid X_i = 0] &= E[Y_i(1) \mid X_i = 1] - E[Y_i(0) \mid X_i = 0] \\ &= E[Y_i(1)] - E[Y_i(0)] \\ &= E\big[Y_i(1) - Y_i(0)\big] = \mathsf{ATE}. \end{aligned} \tag{2}$$

- Step 1: switching equation (1).
- Step 2: independence – conditioning on X_i drops out.
- Step 3: linearity of expectations.

- The observable conditional means $E[Y_i \mid X_i{=}1]$ and $E[Y_i \mid X_i{=}0]$ are exactly what the treatment and control sample averages estimate. Random assignment is what turns a **correlation** into a **causal effect**.

From Potential Outcomes to the Regression Model

- The framework maps **directly** into the regression notation used all term. Define

$$\beta_0 = E[Y_i(0)], \qquad u_i = Y_i(0) - E[Y_i(0)], \qquad \beta_{1i} = Y_i(1) - Y_i(0).$$

- Starting from the switching equation (1):

$$\begin{aligned} Y_i &= Y_i(1)\, X_i + Y_i(0)\,(1 - X_i) \\ &= Y_i(0) + [Y_i(1) - Y_i(0)]X_i \\ &= E[Y_i(0)] + [Y_i(1) - Y_i(0)]X_i + \{Y_i(0) - E[Y_i(0)]\} \\ &= \beta_0 + \beta_{1i}\, X_i + u_i. \end{aligned} \tag{3}$$

- This is the **random coefficients model** – the regression is not an assumption, it is just **renamed potential outcomes**.

Two Special Cases of the Regression Model

Constant treatment effect

If $\beta_{1i} = \beta_1$ for all i, then (3) collapses to the **single-regressor model**

$$Y_i = \beta_0 + \beta_1 X_i + u_i, \qquad i = 1, \dots, n.$$

- By construction $E[u_i] = 0$. If X_i is randomly assigned, then $X_i \perp [Y_i(1), Y_i(0)]$, hence $X_i \perp (\beta_{1i}, u_i)$, so $E[u_i \mid X_i] = 0$.
- The first OLS assumption holds **by design** $\Rightarrow \hat{\beta}_1$ is unbiased and consistent for the causal effect.
- Measurement error in Y_i simply gets absorbed into u_i in the last line of (3) – it does not bias $\hat{\beta}_1$ under random assignment.

Conditioning on Covariates: Unconfoundedness (1/2)

- In many designs X_i is random only **conditional** on covariates W_i. The independence we need then weakens to

$$\big[\, Y_i(1),\, Y_i(0) \,\big] \;\perp\; X_i \;\big|\; W_i, \tag{4}$$

called **unconfoundedness** (or conditional independence / selection on observables).

Conditioning on Covariates: Unconfoundedness (2/2)

- Suppose $Y_i(0) = \beta_0 + \gamma W_i + u_i$ and a constant effect $\beta_{1i} = \beta_1$. The same algebra as (3) gives

$$Y_i = \beta_0 + \beta_1 X_i + \gamma W_i + u_i.$$

- Unconfoundedness then delivers **conditional mean independence**:

$$E[u_i \mid X_i, W_i] = E[Y_i(0) - \beta_0 - \gamma W_i \mid X_i, W_i] = E[Y_i(0) - \beta_0 - \gamma W_i \mid W_i] = E[u_i \mid W_i].$$

Note

Under CMI, $\widehat{\beta}_1$ is unbiased for the causal effect, but $\widehat{\gamma}$ on the control W_i generally is **not** causal: controls earn their keep by closing back-doors, not by being interpreted.

Estimating Causal Effects from Experiments

The Differences Estimator

Definition 3: Differences estimator

The **differences estimator** is the difference in sample averages for the treatment and control groups – equivalently, the OLS estimator of β_1 in

$$Y_i = \beta_0 + \beta_1 X_i + u_i, \qquad i = 1, \ldots, n.$$

- If X_i is randomly assigned, $E[u_i \mid X_i] = 0$, so $\hat{\beta}_1$ is **unbiased and consistent** for the causal effect.
- $\hat{\beta}_1 = \bar{Y}_{\text{treat}} - \bar{Y}_{\text{control}}$: the regression and the mean-comparison give the **same** number.

Differences Estimator with Additional Regressors

Definition 4: Differences estimator with additional regressors

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 W_{1i} + \cdots + \beta_{1+r} W_{ri} + u_i.$$

$\hat{\beta}_1$ is the OLS estimator of the causal effect with controls $W_{1i}, \ldots, W_{ri}$.

- **Why add controls W?**
  - **Efficiency:** if W explains variation in Y, the regression standard error falls and typically so does $\mathrm{SE}(\hat{\beta}_1)$ – random assignment guarantees unbiasedness, controls sharpen precision.
  - **Validity under conditional randomization:** they restore unbiasedness when assignment depends on W (next slides).

- **Requirement:** W_i must satisfy **conditional mean independence**, $E[u_i \mid X_i, W_i] = E[u_i \mid W_i]$. This holds when W_i are **pretreatment** characteristics (e.g. sex). **Never include post-treatment outcomes as controls.**

Effects That Depend on Observables

- Causal effects may differ across observable groups. Estimate this with **interactions**.
- Let W_{1i} be a binary indicator (e.g. 1 for women). Adding $W_{1i} \times X_i$:

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 W_{1i} + \beta_3 (W_{1i} \times X_i) + u_i.$$

- Effect of treatment for men ($W_{1i} = 0$): β_1. For women ($W_{1i} = 1$): $\beta_1 + \beta_3$.

Note

This is the **observable** heterogeneity. We will discuss the harder case where effects vary with **unobservables**.

Randomization Based on Covariates

- **Randomization based on covariates:** assignment probability depends on observable W.
- Example: estimate the effect of mandatory vs. optional homework. Assignment is random, but economics majors ($W_i = 1$) are sent to the treatment group ($X_i = 1$) with higher probability.
- If majors do better anyway, the **simple** differences estimator suffers **omitted variable bias**: being treated is correlated with being a major.

Fix

Because X_i is random **given** W_i, including W_i restores $E[u_i \mid X_i, W_i] = E[u_i \mid W_i]$, so $\widehat{\beta}_1$ is unbiased. If the effect differs by major, add $X_i \times W_i$ as well.

Threats to Validity of Experiments

Internal Validity – Threats (1)

- Recall: a study is **internally valid** if its causal inferences are valid for the population studied.

- **Failure to randomize.** If assignment depends on subject characteristics, X_i correlates with $u_i \Rightarrow$ bias.
  - **Test:** regress X_i on pretreatment $W_1, \ldots, W_r$ and F-test that all coefficients are 0. (X_i binary $\Rightarrow$ linear probability model with robust SEs, or use probit/logit.)

- **Partial compliance.** Subjects do not follow the protocol (assigned to treatment but skip it, or vice versa). Choice enters X_i, so X_i correlates with u_i.
  - **Fix:** use the random **assignment** Z_i as an **instrument** for treatment **received** X_i. Z_i is relevant (assignment shifts uptake) and exogenous (random) $\Rightarrow$ valid IV.

Internal Validity – Threats (2)

- **Attrition.** Subjects drop out after assignment. If drop-out is **related to treatment** (e.g. the ablest trainees leave town for jobs), the remaining sample is nonrandom $\Rightarrow$ **selection bias**.

- **Experimental effects (Hawthorne effect).** Being in an experiment changes behavior.
  - **Fix where feasible:** a **double-blind** protocol – neither subject nor administrator knows group assignment, so both groups share the same experimental effect.
  - Often infeasible in economics (everyone knows who attends the training).

- **Small samples.** Do not bias $\widehat{\beta}_1$, but reduce precision and threaten the large-sample basis for CIs and tests.

External Validity – Threats

- **External validity:** can results generalize to other populations and settings?
- **Nonrepresentative sample.** A training program evaluated on former prison inmates may not generalize to workers with no criminal record.
- **Nonrepresentative program or policy.** A small, tightly monitored pilot may differ from a scaled-up program in quality, funding, or duration.
- **General equilibrium effects.** A small pilot holds the market environment fixed; a economy-wide program does not. A widely available training program might **displace** employer-provided training, shrinking net benefits. An internally valid pilot can still mislead about the scaled-up effect.

Project STAR

Class Size and the Project STAR Experiment (1/2)

- **Question:** does reducing class size raise test scores in early grades?
- **Project STAR** (Student–Teacher Achievement Ratio), Tennessee, late 1980s: a 4-year, $\approx \$12$ million RCT. Three arrangements, K–3:
  - regular class (22–25 students), no aide;
  - small class (13–17 students), no aide;
  - regular class with a teacher's aide.
- Students **and** teachers randomly assigned within each school. Analyzed with the differences estimator (3), optionally adding regressors.

Class Size and the Project STAR Experiment (2/2)

Note

Effect of a small class in kindergarten ≈ 0.19 SD of test scores (95% CI $[0.08, 0.30]$). Strikingly, observational studies in California and Massachusetts gave **overlapping** confidence intervals – the experiment validated the observational work.

Quasi-Experiments

What Is a Quasi-Experiment?

Definition 5: Quasi-experiment

In a **quasi-experiment** (natural experiment), randomness is introduced by variation in individual circumstances – legal quirks, location, policy timing, birth dates, weather – that makes treatment look **as if** it were randomly assigned.

- Two types:
  - **Type 1 – treatment as-if random.** Estimate the effect by **OLS** with X_i as the regressor (possibly with controls W).
  - **Type 2 – assignment as-if random.** The as-if random source only **partly** determines treatment $\Rightarrow$ estimate by **IV**, using that source as the instrument.
- Quasi-experiments are a **bridge** between observational data and true RCTs: the tools are old (OLS, panel, IV), the **interpretation and data** are new.

Three Canonical Examples

- **Card (1990) – Mariel boatlift.** The 1980 lifting of Cuban emigration restrictions sent many immigrants to Miami. Card compared the change in low-skilled wages in Miami to comparable cities $\Rightarrow$ negligible effect on wages. *(Type 1.)*

- **Urquiola (2006), Angrist–Lavy (1999) – class size.** A rule grants an extra teacher above a 30-student threshold. Crossing the threshold shifts class size $\Rightarrow$ regression discontinuity. Bolivia: a 1-SD class-size cut raises scores by up to 0.3 SD. *(Type 2, fuzzy.)*

- **McClellan–McNeil–Newhouse (1994) – cardiac catheterization.** Relative distance to a catheterization hospital influences, but does not determine, treatment $\Rightarrow$ distance is an instrument. *(Type 2.)*

The Differences-in-Differences Estimator

The Differences-in-Differences Estimator

- Even with as-if randomness, treatment and control groups may differ at baseline. Compare **changes** in Y, not levels, to net out fixed group differences.

Definition 6: Differences-in-differences estimator

$$\hat{\beta}_1^{\text{diffs-in-diffs}} = (\bar{Y}_{\text{treat,after}} - \bar{Y}_{\text{treat,before}}) - (\bar{Y}_{\text{control,after}} - \bar{Y}_{\text{control,before}}) = \Delta\bar{Y}_{\text{treat}} - \Delta\bar{Y}_{\text{control}}. \tag{5}$$

- In regression form, with ΔY_i the post- minus pre-treatment change:

$$\Delta Y_i = \beta_0 + \beta_1 X_i + u_i, \qquad \Delta Y_i = \beta_0 + \beta_1 X_i + \beta_2 W_{1i} + \cdots + \beta_{1+r} W_{ri} + u_i.$$

- If X_i is (as-if) random given W, $\hat{\beta}_1$ is unbiased and consistent. With $T > 2$ periods, use panel methods (Chapter 10).

DiD – The Picture (1/2)

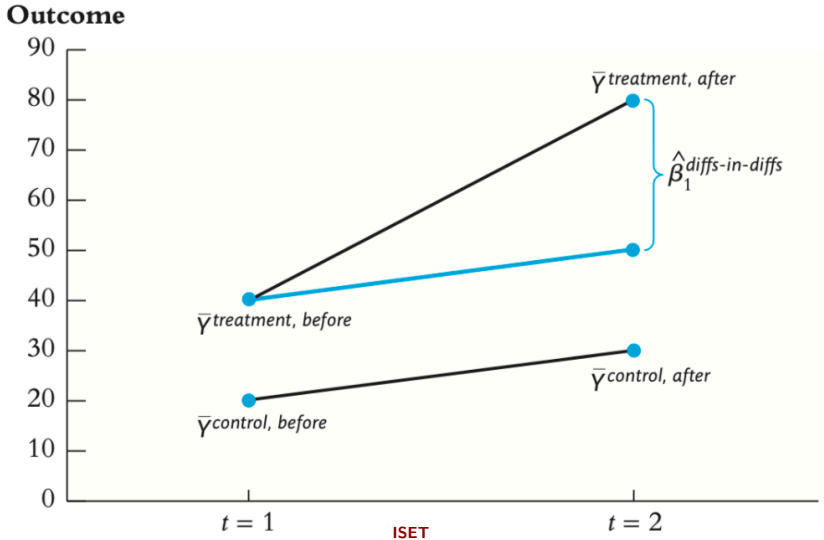

DiD – The Picture (2/2)

- Naive post-difference $= 80 - 30 = 50$, but groups differed by $40 - 20 = 20$ at baseline.
- $\hat{\beta}_1^{\text{diffs-in-diffs}} = (80 - 40) - (30 - 20) = 50 - 20 = 30$: the change for treated **minus** the change for controls, assuming **parallel trends**.

DiD as a Regression – The Interaction Model (1/2)

- The differenced regression needs **panel data** (the same units twice). The equivalent, more general specification runs on the **levels**, with a group dummy, a time dummy, and their interaction:

$$Y_{it} = \beta_0 + \beta_1 \,\text{Treat}_i + \beta_2 \,\text{Post}_t + \beta_3 \,(\text{Treat}_i \times \text{Post}_t) + u_{it}, \tag{6}$$

  where $\text{Treat}_i = 1$ for the treatment group and $\text{Post}_t = 1$ in the after period.

- Each group $\times$ period mean is a combination of coefficients; the **difference of differences** is the interaction coefficient β_3:

Group	**Before**	**After**	**After − Before**
Control	β_0	$\beta_0 + \beta_2$	β_2
Treatment	$\beta_0 + \beta_1$	$\beta_0 + \beta_1 + \beta_2 + \beta_3$	$\beta_2 + \beta_3$
Treat − Control	β_1	$\beta_1 + \beta_3$	β_3

DiD as a Regression – The Interaction Model (2/2)

Note

β_3 in (6) equals $\hat{\beta}_1^{\text{diffs-in-diffs}}$ from (5). Controls W enter additively just as before. Crucially, (6) needs only group and period **labels** – so it works with **repeated cross-sections**, where the differenced form cannot.

DiD as a Regression – Two-Way Fixed Effects

- With many groups and many periods, generalize (6) by replacing the single group and time dummies with full sets of **entity and time fixed effects**:

$$Y_{it} = \alpha_i + \lambda_t + \beta_1 D_{it} + u_{it}, \tag{7}$$

  where α_i are entity effects, λ_t are time effects, and $D_{it} = 1$ once entity i is treated. This is the **two-way fixed effects (TWFE)** estimator (Chapter 10 panel methods).

- For two groups and two periods, (7) reduces exactly to the interaction model (6): α_i absorbs Treat_i, λ_t absorbs Post_t, and $\beta_1 = \beta_3$.

RD Basics

Causal Inference – The Core Challenge

- The fundamental problem: we never observe the same unit under both treatment and control

- We have seen several strategies for finding a valid **counterfactual**:
  - Literal random assignment (RCT)
  - As-good-as-random assignment (quasi-experiments, SSIV)
  - Parallel trends over time (DiD)
  - Excludable and relevant instrument (IV/2SLS)

- Next: **Regression Discontinuity (RD)** – a strategy that arises when treatment is (partially) determined by whether a score crosses a threshold

The RD Idea

- Many real-world policies assign treatment based on a **threshold rule**:
  - Households below a poverty score receive welfare benefits
  - Students above a GPA cutoff are admitted to a program
  - Villages above a risk score are selected for bombing
- Key insight: units just **above** and just **below** the threshold are likely to be very similar in all respects – except treatment status
- This local comparability is what RD exploits for causal identification
- RD compares outcomes just above vs. just below the cutoff to recover a **causal effect at the threshold**

Motivating Example – Georgian Health Insurance

- Hu and Chao (2008): does lack of health insurance prevent poor people from getting medical care?
- **Why not just compare insured vs. uninsured?** – confounding: the insured tend to be richer, which affects health independently
- **Setting:** Republic of Georgia, 2006. Each household received a poverty score from 80 variables. Households with score $\leq 100{,}000$ received government health insurance
- **RD logic:** households just above 100,000 are very similar to households just below 100,000 – except the latter received insurance
- **Outcome:** probability of utilizing acute surgeries / in-patient services

Georgian Health Insurance – The Plot

- The x-axis is the welfare score; the vertical line marks the 100,000 cutoff
- Left of the line: **eligible** (insured)
- Right of the line: **ineligible** (uninsured)
- There is a sharp **jump** in surgery probability at the cutoff
- Estimated treatment effect: $0.009 - 0.001 = 0.008$

Figure 3: The Effect of MAP on Utilization of Acute Surgeries/In-Patient Services

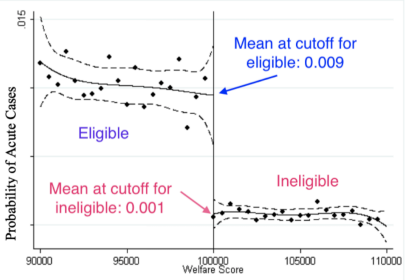

Note: This figure plots probability of utilization of acute surgeries/inpatient services against welfare scores. Each dot is the average probability within 500 intervals of welfare scores. Solid lines are fitted values from 4th order polynomial regressions on either side of the discontinuity. Dotted lines are 95% confidence intervals.

Formalizing Sharp RD

- Let $D_i \in \{0, 1\}$ be treatment; R_i be the **running variable** (score)
- In a **sharp RD**, treatment is a deterministic function of R_i:

$$D_i = \mathbf{1}[R_i \geq c] \tag{8}$$

- Unlike DiD or IV: D_i is **deterministic** in a likely confounder – no selection-on-observables story is possible

Definition: RD Estimand

The RD estimand compares the limits of the conditional expectation function (CEF) around the cutoff c:

$$\tau_{RD} = \lim_{r \downarrow c} E[Y_i \mid R_i = r] - \lim_{r \uparrow c} E[Y_i \mid R_i = r]$$

When Does τ_{RD} Identify a Causal Effect?

- Below the cutoff ($R_i < c$): $D_i = 0$, so $\lim_{r \uparrow c} E[Y_i \mid R_i = r] = \lim_{r \uparrow c} E[Y_i(0) \mid R_i = r]$
- Above the cutoff ($R_i \geq c$): $D_i = 1$, so $\lim_{r \downarrow c} E[Y_i \mid R_i = r] = \lim_{r \downarrow c} E[Y_i(1) \mid R_i = r]$

Key Assumption (Continuity)

$f_d(r) = E[Y_i(d) \mid R_i = r]$ is continuous at c for $d = 0, 1$.

- Under continuity, the limits identify potential outcome means at c:

$$\tau_{RD} = E[Y_i(1) \mid R_i = c] - E[Y_i(0) \mid R_i = c] = E[Y_i(1) - Y_i(0) \mid R_i = c]$$

- This is the **average treatment effect for units at the cutoff** (a CATE)

Graphical Intuition

- The dashed curves show the counterfactual potential outcome schedules $E[Y_0 \mid R = r]$ and $E[Y_1 \mid R = r]$
- Below c: we observe $E[Y_0 \mid R = r]$
- Above c: we observe $E[Y_1 \mid R = r]$
- **Continuity** allows us to extrapolate each schedule to the cutoff
- The **gap at** c is the CATE $E[Y_1 - Y_0 \mid R = c]$

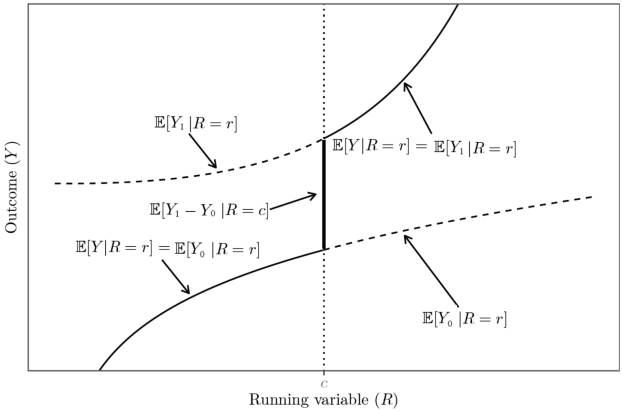

When Does Continuity Fail?

- The continuity assumption $E[Y_i(d) \mid R_i = r]$ continuous at c can fail in two main ways:

- **Failure #1: Confounding factors jump at the cutoff**
  - If other policies or characteristics also change discontinuously at the threshold, the jump in outcomes reflects more than treatment
  - Example: Holmes (1998) – right-to-work (RTW) laws change at state borders, but so do many other state policies

- **Failure #2: Manipulation of the running variable**
  - If agents can sort just above or below the threshold, units on either side are no longer comparable
  - Example: NYC Regents exams – teachers bump students over the 55/65 scoring thresholds, creating a non-random composition near the cutoff

From Identification to Estimation

- Under continuity we have identification:

$$\tau_{RD} = \lim_{r \downarrow c} E[Y_i \mid R_i = r] - \lim_{r \uparrow c} E[Y_i \mid R_i = r] = E[Y_i(1) - Y_i(0) \mid R_i = c]$$

- To estimate τ_{RD}, we need to estimate the CEF $E[Y_i \mid R_i = r]$ and evaluate it at the limit from each side

- **Tool:** OLS regression – fit separate functions of R_i on each side, then read off the predicted values as $r \to c$

- Two questions:
  - What functional form should we use?
  - Which observations should we use?

From Identification to Estimation

- For implementation see
  - Huntington-Klein, N., *The Effect: An Introduction to Research Design and Causality*, Chapter 20: "Regression Discontinuity." Available online: https://theeffectbook.net/ch-RegressionDiscontinuity.html

Threats to Quasi-Experiments

Internal Validity of Quasi-Experiments

- The threats that we have already discussed carry over, with modifications. The central issue: **is the as-if randomization really random?**

- **Failure of randomization.** If the as-if assignment is not random, OLS (or IV) is biased / inconsistent.
  - **Partial test:** regress X (or Z) on the W's and test for systematic differences. But X/Z may still relate to **unobservables** – untestable, so judgment is required.

- **Partial influence on treatment** is the quasi-experimental analogue of partial compliance $\Rightarrow$ use IV.

- **Attrition** $\Rightarrow$ selection bias, as before. **Experimental effects** are usually **absent** – subjects do not know they are "in" an experiment.

Instrument Validity Must Be Scrutinized

As-if random does not guarantee exogeneity

A randomly assigned instrument can still be **endogenous** if it affects Y through a channel other than treatment.

- **Angrist (1990) – draft lottery.** Lottery numbers were truly random. But a low number could trigger **draft-avoidance** behavior (extra schooling) that itself affects civilian earnings $\Rightarrow \text{corr}(Z_i, u_i) \neq 0$.
- **Cardiac catheterization – distance.** If patients living near a catheterization hospital are **healthier** (better access generally), distance correlates with omitted health factors.
- **External validity:** the special event creating the as-if randomness may also limit generalization (Mariel: Miami's Cuban community; Angrist: wartime service).

Heterogeneous Populations and LATE

Unobserved Heterogeneity

- We have allowed effects to vary with **observables**. Now effects vary with **unobservables** – a **heterogeneous population**.

- Each individual has their own effect $\beta_{1i} = Y_i(1) - Y_i(0)$, giving the random coefficients model (derived from potential outcomes in (3)):

$$Y_i = \beta_0 + \beta_{1i} X_i + u_i.$$

- β_{1i} is now a **random variable** reflecting unobserved variation (e.g. how much someone already knows about resume-writing). The target is the **average** causal effect $E[\beta_{1i}]$.

- **Key question:** what do our estimators recover when β_{1i} varies and is unobserved? Answer: **OLS** gets the ATE; **IV** generally does **not**.

OLS Recovers the ATE (Random Assignment)

- With X_i randomly assigned, $X_i \perp \beta_{1i}$ and $\text{cov}(u_i, X_i) = 0$. The OLS slope $\hat{\beta}_1 = s_{XY}/s_X^2 \xrightarrow{p} \sigma_{XY}/\sigma_X^2$:

$$\begin{aligned} \hat{\beta}_1 \xrightarrow{p} \frac{\sigma_{XY}}{\sigma_X^2} &= \frac{\text{cov}(\beta_0 + \beta_{1i}X_i + u_i,\ X_i)}{\sigma_X^2} \\ &= \frac{\text{cov}(\beta_{1i}X_i,\ X_i)}{\sigma_X^2} \\ &= E[\beta_{1i}]. \end{aligned} \tag{9}$$

- Second line: $\text{cov}(\beta_0, X_i) = 0$ and $\text{cov}(u_i, X_i) = 0$.
- Third line: $\beta_{1i} \perp X_i$ under random assignment.
- So under (as-if) random assignment of the **treatment**, the differences estimator is consistent for the ATE even with fully heterogeneous effects.

IV with Heterogeneous Effects: First Stage

- Now treatment is only **partly** driven by an instrument Z_i. Allow the first-stage effect to vary across individuals:

$$X_i = \pi_0 + \pi_{1i} Z_i + v_i, \tag{10}$$

  where π_{1i} – how much Z_i moves i's treatment – differs by individual.

- Assume Z_i is (as-if) randomly assigned, so $Z_i \perp (u_i, v_i, \pi_{1i}, \beta_{1i})$, and $E[\pi_{1i}] \neq 0$ (relevance on average).

Note

π_{1i} encodes who is a "**complier**": for whom does the instrument actually change treatment status? This is the hinge on which LATE turns.

IV Estimates a Weighted Average: LATE

- The TSLS estimator $\hat{\beta}_1^{\text{TSLS}} = s_{ZY}/s_{ZX} \xrightarrow{p} \sigma_{ZY}/\sigma_{ZX}$. Substituting (10) and the random-coefficients model (Appendix 13.2) gives

$$\hat{\beta}_1^{\text{TSLS}} \xrightarrow{p} \frac{\sigma_{ZY}}{\sigma_{ZX}} = \frac{E[\beta_{1i}\,\pi_{1i}]}{E[\pi_{1i}]}. \tag{11}$$

- This is a **weighted average** of the individual effects β_{1i}, with weights $\pi_{1i}/E[\pi_{1i}]$.

Local average treatment effect (LATE)

IV consistently estimates the **LATE** – a weighted average of β_{1i} that places the **most weight on individuals whose treatment is most influenced by the instrument** (the compliers). "Local" = to the compliers.

Sketch of the LATE Derivation

- With $Z_i \perp (u_i, v_i, \pi_{1i}, \beta_{1i})$ and $E[Z_i - \mu_Z] = 0$:

$$\sigma_{ZX} = E[(Z_i - \mu_Z)X_i] = E[\pi_{1i}]\,\sigma_Z^2$$

  using $E[\pi_{1i} Z_i (Z_i - \mu_Z)] = E[\pi_{1i}]\,E[Z_i(Z_i - \mu_Z)] = E[\pi_{1i}]\sigma_Z^2$ and $\operatorname{cov}(Z_i, v_i) = 0$.

- Similarly, substituting $Y_i = \beta_0 + \beta_{1i}(\pi_0 + \pi_{1i} Z_i + v_i) + u_i$,

$$\sigma_{ZY} = E[(Z_i - \mu_Z)Y_i] = E[\beta_{1i}\pi_{1i}]\,\sigma_Z^2.$$

- Taking the ratio, the σ_Z^2 cancels:

$$\widehat{\beta}_1^{\mathsf{TSLS}} \xrightarrow{p} \frac{\sigma_Z^2\, E[\beta_{1i}\pi_{1i}]}{\sigma_Z^2\, E[\pi_{1i}]} = \frac{E[\beta_{1i}\pi_{1i}]}{E[\pi_{1i}]},$$

  which is (11).

When Does LATE = ATE?

- The weighted average (11) equals the ATE $E[\beta_{1i}]$ in **three** special cases:
- **(1) Homogeneous treatment effect:** $\beta_{1i} = \beta_1$. Then $\dfrac{E[\beta_1 \pi_{1i}]}{E[\pi_{1i}]} = \beta_1$.
- **(2) Instrument affects everyone equally:** $\pi_{1i} = \pi_1$. Then $\dfrac{E[\beta_{1i}]\pi_1}{\pi_1} = E[\beta_{1i}]$.
- **(3)** β_{1i} **and** π_{1i} **uncorrelated:** $\text{cov}(\beta_{1i}, \pi_{1i}) = 0$, so $E[\beta_{1i}\pi_{1i}] = E[\beta_{1i}]E[\pi_{1i}]$, giving $E[\beta_{1i}]$.

i Note

The problem case is exactly when people **select into treatment based on their own gains** – so β_{1i} and π_{1i} move together. Then LATE $\neq$ ATE.

Conclusion

- The **potential outcomes framework** $\{Y_i(1), Y_i(0)\}$ defines causal effects cleanly and exposes the **fundamental problem**: one potential outcome is always missing.
- **Randomization** (or as-if randomization) makes treatment independent of potential outcomes, killing selection bias and identifying the ATE – and the regression model is just **renamed potential outcomes**.
- **Quasi-experiments** extend the logic to observational data via DiD, RD, and IV. Their great virtue is a **transparent** source of randomness; their great risk is that as-if randomness may not be truly exogenous.
- With **heterogeneous effects**, OLS under random assignment recovers the ATE, but **IV recovers a LATE** – weighting compliers most. Knowing *whose* effect you estimated is as important as the estimate itself.

Required Reading

- Stock, J. H. and Watson, M. W., *Introduction to Econometrics*, 4th edition (Global Edition), Chapter 13: "Experiments and Quasi-Experiments" (Sections 13.1–13.7).

- Especially Appendix 13.3, "The Potential Outcomes Framework for Analyzing Data from Experiments," and Appendix 13.2 on IV estimation with heterogeneous causal effects.