DegreeLive

Bachelor of Commerce (B2001) | Level 3 • Part B Major | Monash University

Applied econometrics

Course Code: ETC3410 • Module 3 Notes

Module 3: Regression Discontinuity Designs & Instrumental Variables in Practice

3.1 Regression Discontinuity Design (RDD): Core Intuition and Setup

In Australian public administration, education, and healthcare, policy eligibility is frequently governed by crisp numerical thresholds. Examples include Australian Tertiary Admission Rank (ATAR) cutoffs for university degrees, Medicare Extended Safety Net out-of-pocket spending thresholds, Centrelink income-test tapers, and Australian Taxation Office (ATO) turnover cutoffs for small-business instant asset write-offs. When treatment assignment changes discontinuously at an arbitrary known threshold c of a continuous **running (forcing) variable** X_i, **Regression Discontinuity Design (RDD)** exploits local near-random variation just above and below c to identify causal effects without global randomisation.

Let $X_i \in \mathrm{R}$ denote the observed running variable for unit i (centred at the cutoff as $\widetilde{X}_i = X_i - c$), and let $Z_i = \mathbf{1}(X_i \geq c)$ be the eligibility indicator jumping from 0 to 1 at c. RDD distinguishes two canonical regimes:

- **Sharp RDD:** Treatment receipt D_i is a deterministic step function of the running variable: $D_i = \mathbf{1}(X_i \geq c)$. Every unit at or above c receives treatment ($P\left(D_i = 1 \mid X_i \geq c\right) = 1$), and no unit below c receives treatment ($P\left(D_i = 1 \mid X_i < c\right) = 0$).
- **Fuzzy RDD:** Crossing the cutoff c raises the *probability* of treatment discontinuously, but compliance is imperfect ($0 < \lim_{x \downarrow c} P\left(D_i = 1 \mid X_i = x\right) - \lim_{x \uparrow c} P\left(D_i = 1 \mid X_i = x\right) < 1$). For instance, attaining an ATAR above 80.00 grants eligibility for a Monash scholarship, yet some eligible students decline enrolment while some below 80.00 enter via special consideration pathways.

3.2 Sharp RDD Identification and Continuity Assumptions

In a Sharp RDD, there is zero overlap across the cutoff at the same value of X_i: all units with $X_i \geq c$ are treated and all with $X_i < c$ are controls. Identification therefore relies on extrapolating across the boundary c via the **Continuity Assumption** (Hahn, Todd & Van der Klaauw, 2001):

> $$E[Y_i(0) \mid X_i = x] \text{ and } E[Y_i(1) \mid X_i = x] \text{ are continuous in } x \text{ at } x = c$$

Where:

- $E[Y_i(0) \mid X_i = x]$ is the conditional expectation function (CEF) of the untreated potential outcome as a function of the running variable x.
- $E[Y_i(1) \mid X_i = x]$ is the conditional expectation function of the treated potential outcome.
- Continuity requires that all other determinants of Y_i—such as innate ability, family wealth, or prior firm productivity—evolve smoothly across $X_i = c$ without jumping abruptly at the threshold.

Under Continuity, taking right-hand and left-hand limits as x approaches c identifies the **Sharp RDD Local Average Treatment Effect at the cutoff**:

Degree Live

$$\tau_{\text{SRD}} = \lim_{x \downarrow c} E[Y_i \mid X_i = x] - \lim_{x \uparrow c} E[Y_i \mid X_i = x] = E[Y_i(1) - Y_i(0) \mid X_i = c]$$

Where:

- $\lim_{x \downarrow c}$ denotes the right-hand limit approaching c from above ($x \geq c$).
- $\lim_{x \uparrow c}$ denotes the left-hand limit approaching c from below ($x < c$).
- Crucially, τ_{SRD} is a **local estimand** valid specifically for marginal units situated right at $X_i = c$; extrapolating τ_{SRD} to individuals far from the cutoff requires additional homogeneity or structural assumptions.

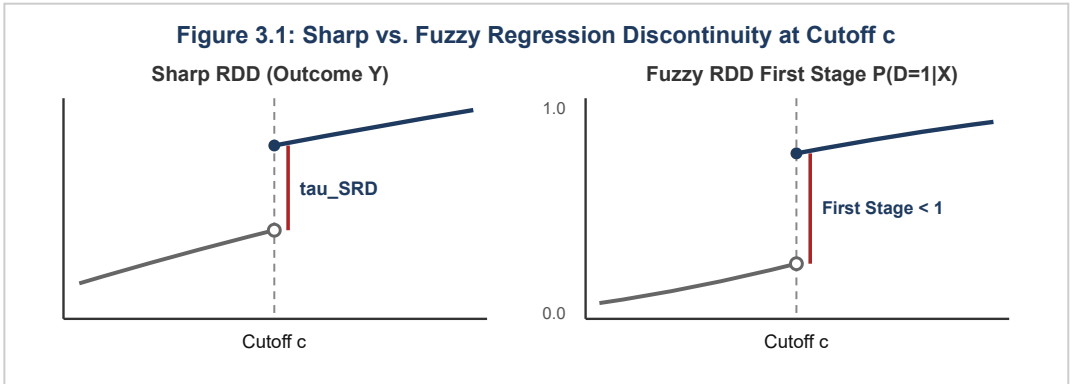
3.3 Fuzzy RDD and Local Instrumental Variables Equivalence

In a Fuzzy RDD, eligibility $Z_i = \mathbf{1}(X_i \geq c)$ shifts treatment uptake D_i imperfectly. Define potential treatment indicators $D_i(z)$ for assignment $z \in \{0, 1\}$ near the cutoff. Assuming local continuity of $E[Y_i(0) \mid X_i = x]$, $E[Y_i(1) \mid X_i = x]$, local monotonicity ($D_i(1) \geq D_i(0)$ near c), and a non-zero first-stage discontinuity in treatment probability:

$$\tau_{\text{FRD}} = \frac{\lim_{x \downarrow c} E[Y_i \mid X_i = x] - \lim_{x \uparrow c} E[Y_i \mid X_i = x]}{\lim_{x \downarrow c} E[D_i \mid X_i = x] - \lim_{x \uparrow c} E[D_i \mid X_i = x]} = E[Y_i(1) - Y_i(0) \mid X_i = c, \text{ Complier}]$$

Where:

- The numerator is the **reduced-form (Intention-to-Treat)** jump in the outcome Y_i at cutoff c.
- The denominator is the **first-stage** jump in treatment take-up probability $P(D_i = 1 \mid X_i)$ at cutoff c.
- τ_{FRD} recovers the Local Average Treatment Effect specifically for **cutoff compliers**—units situated at $X_i = c$ whose treatment status is induced to switch from 0 to 1 by crossing the threshold.
- Algebraically, estimating τ_{FRD} via local linear regression on both sides of c within bandwidth h is numerically equivalent to local Weighted Two-Stage Least Squares (2SLS) using $Z_i = \mathbf{1}(X_i \geq c)$ as the excluded instrument for D_i, controlling for $\widetilde{X}_i$ and $Z_i \cdot \widetilde{X}_i$.



DegreeLive

3.4 Bandwidth Selection, Local Polynomials, and Kernel Weights

Because the true functional form of $E[Y_i \mid X_i = x]$ away from c is unknown, empirical RDD estimation faces a fundamental bias-variance tradeoff between global parametric fitting and local non-parametric approximation.

The Gelman and Imbens (2019) Critique of High-Order Global Polynomials

Early RDD papers frequently fitted global 4th- or 5th-degree polynomials across the entire support of X_i interacted with Z_i. Gelman and Imbens (2019) demonstrated that high-order global polynomials should be avoided in RDD for three reasons: (1) noisy implicit weights that assign extreme leverage to observations far from the cutoff; (2) severe sensitivity of estimated jumps to polynomial degree (Runge's boundary phenomenon); and (3) artificially narrow standard errors that yield high false-positive rates in placebo cutoff tests.

Local Linear Regression and Kernel Weighting

Modern applied practice estimates **local linear regressions** ($p = 1$)—or at most local quadratic ($p = 2$) as a robustness check—exclusively on observations falling inside a narrow symmetric bandwidth $X_i \in [c - h, c + h]$:

$$\min_{\alpha_-, \alpha_+, \beta_-, \beta_+} \sum_{i:|X_i - c| \leq h} K\left(\frac{X_i - c}{h}\right) \left[Y_i - \alpha_- - \beta_-(X_i - c) - \tau Z_i - \delta Z_i (X_i - c)\right]^2$$

Where:

- $h > 0$ is the bandwidth governing the width of the estimation window around c.
- $Z_i = \mathbf{1}(X_i \geq c)$ captures the intercept shift $\hat{\tau}_{\text{SRD}}$ at $X_i = c$ (where $X_i - c = 0$).
- β_- is the slope below the cutoff and $\beta_+ = \beta_- + \delta$ is the slope above the cutoff, permitting different linear trends on each side.
- $K(u)$ is a symmetric kernel function assigning higher weight to units closest to the cutoff. Standard choices are the **Triangular kernel** $K(u) = (1 - |u|) \cdot \mathbf{1}(|u| \leq 1)$ (MSE-optimal for boundary estimation) and the **Uniform kernel** $K(u) = \frac{1}{2}\mathbf{1}(|u| \leq 1)$ (unweighted OLS inside $[c - h, c + h]$).

MSE-Optimal Bandwidths and Robust Bias-Corrected Inference (CCT)

Choosing h involves trading off asymptotic squared bias (which grows at rate $O(h^4)$ for local linear fits due to boundary curvature $m''(c)$) against sampling variance (which shrinks at rate $O((Nh)^{-1})$). Imbens and Kalyanaraman (2012) and Calonico, Cattaneo, and Titiunik (2014, **CCT**) derive the **Mean Squared Error (MSE)-optimal bandwidth**:

$$h_{\text{MSE}} = C_K \cdot N^{-1/5}$$

DegreeLive

Where:

- C_K depends on the kernel, outcome variance at c, density $f_X(c)$, and second derivatives $m''_+(c) - m''_-(c)$ on either side of the cutoff.
- Crucially, because h_{MSE} balances bias and variance rather than undersmoothing, standard confidence intervals centred at $\hat{\tau}(h_{\text{MSE}})$ suffer from non-negligible smoothing bias, causing under-coverage. CCT (2014) resolve this via **robust bias-corrected (RBC) inference**: they estimate the leading bias term using a local quadratic fit (with pilot bandwidth b), subtract it from $\hat{\tau}$, and adjust the standard error to reflect the extra variability of bias estimation.

Worked Numerical Example 3.1: Sharp and Fuzzy RDD Estimation

A Monash research team evaluates an Australian Medicare Safety Net threshold at $c = \text{A\$2,400}$ of annual out-of-pocket medical spending (X_i). Let $\widetilde{X}_i = X_i - 2,400$ be centred spending within a uniform-kernel local linear bandwidth $h = \text{A\$200}$. OLS estimation on $\widetilde{X}_i \in [-200, +200]$ produces two local linear regressions:

- **Reduced-Form (Annual Specialist Visits Y_i):**

$$\hat{Y}_i = 4.20 + 0.0015\widetilde{X}_i + 0.96Z_i + 0.0004Z_i\widetilde{X}_i$$

- **First-Stage (Subsidised Safety Net Rebate Take-Up D_i):**

$$\hat{D}_i = 0.08 + 0.0002\widetilde{X}_i + 0.80Z_i - 0.0001Z_i\widetilde{X}_i$$

Step 1 — Evaluate Left and Right Limits at the Cutoff ($\widetilde{X}_i = 0$):

$$\lim_{x \uparrow 2,400} \hat{Y}(x) = 4.20, \qquad \lim_{x \downarrow 2,400} \hat{Y}(x) = 4.20 + 0.96 = 5.16$$

$$\lim_{x \uparrow 2,400} \hat{D}(x) = 0.08, \qquad \lim_{x \downarrow 2,400} \hat{D}(x) = 0.08 + 0.80 = 0.88$$

Where:

- The reduced-form outcome jump (ITT at the cutoff) is $5.16 - 4.20 = 0.96$ specialist visits per year.
- The first-stage jump in actual Safety Net rebate coverage is $0.88 - 0.08 = 0.80$ (80 percentage points).

Step 2 — Compute the Fuzzy RDD Causal Estimate $\hat{\tau}_{\text{FRD}}$:

$$\hat{\tau}_{\text{FRD}} = \frac{0.96}{0.80} = 1.20 \text{ specialist visits per year}$$

Where:

- Among marginal patients at the A$2,400 spending threshold who comply with Safety Net enrolment (80% of the local population), receiving the lower co-payment increases specialist utilisation by 1.20 visits annually.

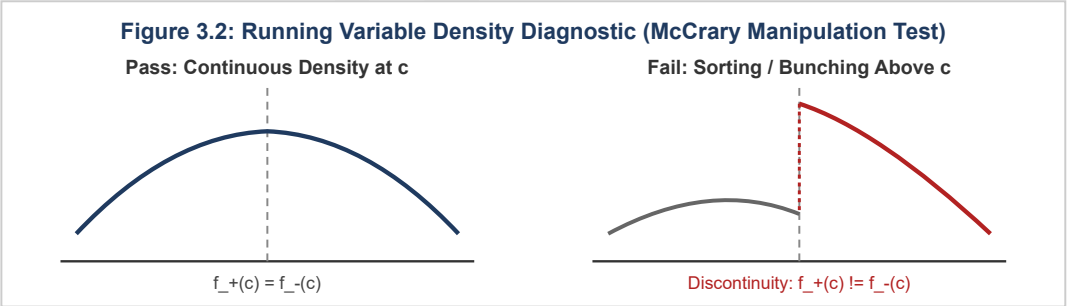
3.5 RDD Validation, Manipulation, and Falsification Diagnostics

Because the continuity assumption is fundamentally untestable on potential outcomes, credible RDD studies present a rigorous battery of falsification checks:

McCrary (2008) Density Test & Cattaneo-Jansson-Ma (2020) Test

If individuals or firms can precisely manipulate their running variable X_i to sort onto the favourable side of the cutoff—such as small businesses artificially suppressing reported ATO revenue to stay below a tax threshold, or students gaming test re-scoring—the continuity assumption fails. Manipulation leaves a visible signature: a **discontinuity in the density** $f_X(x)$ at $x = c$ (excess bunching just above or below c, accompanied by a missing mass on the opposite side).

- **McCrary (2008):** Bins X_i into a finely spaced histogram, smooths bin heights separately on each side via local linear regression, and tests $H_0 : \ln f_+(c) - \ln f_-(c) = 0$.
- **Cattaneo, Jansson & Ma (2020):** Avoids arbitrary pre-binning by applying local polynomial density estimation directly to the empirical CDF, providing higher power and clean bias-corrected inference.



Additional RDD Falsification Tests

- **Pre-Treatment Covariate Smoothness:** Run the exact RDD local linear specification using predetermined baseline covariates W_i (e.g., parental income, age, gender, prior medical history) as placebo outcomes. Statistically significant jumps at c cast doubt on continuity.
- **Placebo Cutoff Tests:** Estimate local linear jumps at artificial cutoffs $c' \neq c$ inside strictly treated or strictly control regions (e.g., median of left/right subsamples). Finding zero effects at placebo cutoffs strengthens confidence that the true jump at c reflects policy rather than curvature noise.

- **Bandwidth Sensitivity & Donut-Hole RDD:** Plot $\hat{\tau}(h)$ across multiples of h_{MSE} ($0.5h, 0.75h, h, 1.5h, 2h$). If discrete heaping (e.g., rounded salary numbers at A$50,000) or localised sorting contaminates observations right at c, a **Donut-Hole RDD** excludes observations within a small radius $|X_i - c| < \epsilon$ to test stability.

3.6 Instrumental Variables (IV) and Two-Stage Least Squares (2SLS)

When treatment D_i is endogenous ($\text{Cov}(D_i, u_i) \neq 0$) due to omitted variables, simultaneity (such as price and quantity in supply-demand systems), or classical measurement error in D_i, OLS yields inconsistent estimates. **Instrumental Variables (IV)** resolves endogeneity by introducing an excluded variable Z_i that shifts D_i without directly affecting Y_i.

Consider the linear structural equation with endogenous scalar regressor D_i, exogenous covariates $\mathbf{W}_i$ (including a constant), and excluded instruments $\mathbf{Z}_i$ ($L \geq 1$ instruments):

$$Y_i = \beta D_i + \mathbf{W}_i' \boldsymbol{\gamma} + u_i$$

$$D_i = \mathbf{Z}_i' \boldsymbol{\pi} + \mathbf{W}_i' \boldsymbol{\delta} + v_i$$

Where:

- Y_i is the outcome variable and D_i is the endogenous explanatory variable.
- $\mathbf{W}_i$ is the vector of included exogenous control variables.
- $\mathbf{Z}_i$ is the $L \times 1$ vector of excluded instrumental variables.
- u_i is the structural error term and v_i is the first-stage error term.

Valid IV identification requires two core conditions:

1. **Instrument Relevance (First Stage):** $\boldsymbol{\pi} \neq \mathbf{0}$ (after partialling out $\mathbf{W}_i$, $\mathbf{Z}_i$ has non-zero partial correlation with D_i).
2. **Exclusion Restriction & Independence:** $\text{Cov}(\mathbf{Z}_i, u_i) = \mathbf{0}$ (and $\text{Cov}(\mathbf{W}_i, u_i) = \mathbf{0}$). The instrument must influence Y_i *exclusively* through its effect on D_i, sharing no common unobserved confounders with Y_i.

The 2SLS Estimator and Wald Formula

Let $\mathbf{X}_i = (D_i, \mathbf{W}_i')'$ and $\widetilde{\mathbf{Z}}_i = (\mathbf{Z}_i', \mathbf{W}_i')'$. When $L = 1$ (just-identified), the IV estimator reduces to the scalar ratio of sample covariances (after partialling out $\mathbf{W}_i$ via FWL), and for a binary instrument $Z_i \in \{0, 1\}$ without covariates, it becomes the **Wald Estimator**:

$$\hat{\beta}_{\text{IV}} = \frac{\widehat{\text{Cov}}(\widetilde{Z}_i, \widetilde{Y}_i)}{\widehat{\text{Cov}}(\widetilde{Z}_i, \widetilde{D}_i)}, \qquad \hat{\beta}_{\text{Wald}} = \frac{\bar{Y}_{Z=1} - \bar{Y}_{Z=0}}{\bar{D}_{Z=1} - \bar{D}_{Z=0}}$$

DegreeLive

Where:

- $\widetilde{Z}_i, \widetilde{Y}_i, \widetilde{D}_i$ are residuals from projecting Z_i, Y_i, D_i onto $\mathbf{W}_i$.
- For overidentified models ($L > 1$), **Two-Stage Least Squares (2SLS)** regresses D_i on $(\mathbf{Z}_i, \mathbf{W}_i)$ in Stage 1 to obtain fitted values $\hat{D}_i$, and then regresses Y_i on $(\hat{D}_i, \mathbf{W}_i)$ in Stage 2:

$$\hat{\boldsymbol{\theta}}_{2\text{SLS}} = (\mathbf{X}'\mathbf{P}_Z\mathbf{X})^{-1}(\mathbf{X}'\mathbf{P}_Z\mathbf{Y}), \qquad \mathbf{P}_Z = \widetilde{\mathbf{Z}}(\widetilde{\mathbf{Z}}'\widetilde{\mathbf{Z}})^{-1}\widetilde{\mathbf{Z}}'$$

Where:

- $\mathbf{P}_Z$ is the orthogonal projection matrix onto the column space of all exogenous variables $(\mathbf{Z}, \mathbf{W})$.
- **Computational Warning:** Never compute 2SLS standard errors manually by running two separate OLS commands, because Stage 2 OLS uses residuals $Y_i - \hat{\beta}\hat{D}_i$ instead of true structural residuals $Y_i - \hat{\beta}D_i$, producing incorrect variance estimates.

Overidentification Testing (Sargan / Hansen J-Test)

When $L > 1$ (more instruments than endogenous variables), researchers can test whether the overidentifying exclusion restrictions are mutually consistent via the **Sargan (homoskedastic) / Hansen J (heteroskedasticity- or cluster-robust) test**. Let $\hat{u}_i = Y_i - \mathbf{X}_i'\hat{\boldsymbol{\theta}}_{2\text{SLS}}$ be the 2SLS structural residuals. Under homoskedasticity, regressing $\hat{u}_i$ on $(\mathbf{Z}_i, \mathbf{W}_i)$ gives:

$$S = N \cdot R^2_{\hat{u}} \overset{d}{\sim} \chi^2_{L-1}$$

Where:

- N is sample size and $R^2_{\hat{u}}$ is the uncentred R^2 from regressing 2SLS residuals on all instruments and exogenous controls.
- $L - 1$ is the number of overidentifying restrictions (instruments minus endogenous regressors).
- **Critical Limitation:** The Sargan/Hansen J-test only verifies whether different instruments produce statistically compatible estimates of β. If *all* instruments share the same bias (invalid exclusion restriction in the same direction), the J-test will fail to reject; conversely, if treatment effects are heterogeneous, two valid instruments targeting different complier groups can produce different LATEs and trigger rejection!

3.7 Local Average Treatment Effect (LATE) Theorem

In classical linear IV with constant treatment effects ($Y_i(1) - Y_i(0) = \beta$ for every i), 2SLS identifies $\text{ATE} = \beta$. In reality, causal responses vary across individuals. Imbens and Angrist (1994) established exactly what IV identifies under **heterogeneous treatment effects** with a binary instrument $Z_i \in \{0, 1\}$ and binary treatment $D_i \in \{0, 1\}$.

Let $D_i(1)$ and $D_i(0)$ represent unit i's potential treatment take-up when assigned $Z_i = 1$ versus $Z_i = 0$. The population partitions into four mutually exclusive latent compliance types:

Compliance Group	$D_i(0)$	$D_i(1)$	Behavioural Description	Response to Instrument Z_i
Always-Takers (AT)	1	1	Takes treatment regardless of Z_i	$D_i(1) - D_i(0) = 0$
Never-Takers (NT)	0	0	Refuses treatment regardless of Z_i	$D_i(1) - D_i(0) = 0$
Compliers (C)	0	1	Takes treatment if and only if $Z_i = 1$	$D_i(1) - D_i(0) = +1$
Defiers (DF)	1	0	Does the exact opposite of Z_i assignment	$D_i(1) - D_i(0) = -1$

The Four LATE Identification Assumptions

1. **Independence:** $(Y_i(1), Y_i(0), D_i(1), D_i(0)) \perp Z_i$ (the instrument is as good as randomly assigned).
2. **Exclusion Restriction:** Potential outcomes depend solely on D_i, not directly on Z_i: $Y_i(d, 1) = Y_i(d, 0) = Y_i(d)$.
3. **First Stage Relevance:** $P(D_i(1) = 1) \neq P(D_i(0) = 1)$ (there is a non-zero share of compliers).
4. **Monotonicity (No Defiers):** $D_i(1) \geq D_i(0)$ for all i (meaning $P(\text{Defier}) = 0$). Crossing Z_i from 0 to 1 can only induce units into treatment, never push any unit out of treatment.

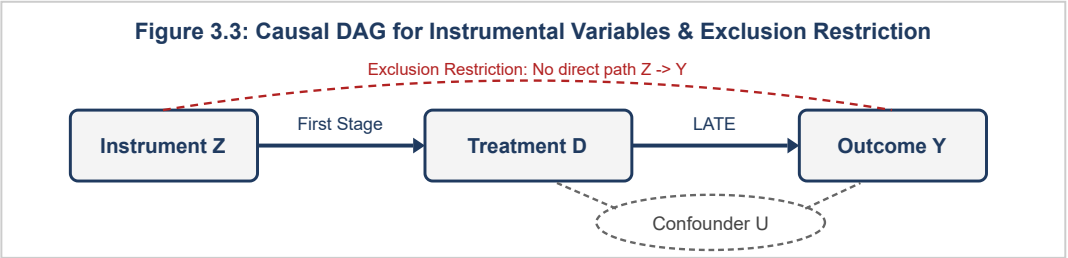
Under these four assumptions, the Wald / IV estimand identifies the **Local Average Treatment Effect (LATE)**—the average causal gain exclusively among **Compliers**:

$$\beta_{\text{IV}} = \frac{E[Y_i \mid Z_i = 1] - E[Y_i \mid Z_i = 0]}{E[D_i \mid Z_i = 1] - E[D_i \mid Z_i = 0]} = E[Y_i(1) - Y_i(0) \mid D_i(1) > D_i(0)] = \text{LATE}$$

Where:

- The denominator $E[D_i \mid Z_i = 1] - E[D_i \mid Z_i = 0] = P(\text{Complier})$ equals the population share of Compliers.
- Because Never-Takers and Always-Takers exhibit $D_i(1) = D_i(0)$, their outcomes cancel out of the numerator $E[Y_i \mid Z_i = 1] - E[Y_i \mid Z_i = 0]$, leaving zero information about their causal effects τ_{AT} or τ_{NT}.
- When an RCT experiences *one-sided noncompliance* (zero control cross-over, so $D_i(0) = 0$ for everyone and no Always-Takers exist), Compliers coincide exactly with Treated units ($D_i = 1$), so $\text{LATE} = \text{ATT} = \text{TOT}$.

DegreeLive



Characterising Compliers (Abadie's κ-Weighting)

Although we cannot label individual unit i as a Complier (since $D_i(1)$ and $D_i(0)$ are not simultaneously observed), Abadie (2003) showed that we can identify the population shares of Always-Takers, Never-Takers, and Compliers, as well as the average covariate profile $E[X_i \mid \text{Complier}]$:

- $P(\text{Always-Taker}) = P(D_i = 1 \mid Z_i = 0)$
- $P(\text{Never-Taker}) = P(D_i = 0 \mid Z_i = 1)$
- $P(\text{Complier}) = P(D_i = 1 \mid Z_i = 1) - P(D_i = 1 \mid Z_i = 0)$

And for any covariate X_i, the Complier average is recovered by weighting observations by $\kappa_i = 1 - \frac{D_i(1-Z_i)}{1-P(Z_i=1)} - \frac{(1-D_i)Z_i}{P(Z_i=1)}$ (or equivalently computing the ratio of first-stage-weighted covariates).

Worked Numerical Example 3.2: Compliance Decomposition, LATE, and Overidentification

Suppose the Victorian Government trials an opt-in digital energy-efficiency audit ($D_i \in \{0, 1\}$) across $2,000$ Melbourne households, randomly sending priority fee-waiver letters ($Z_i = 1$) to $1,000$ households and standard brochures ($Z_i = 0$) to $1,000$ households. Observed audit completion rates and subsequent annual electricity bills (AUD) are:

- Assigned priority waiver ($Z_i = 1$): 650 households complete the audit ($\bar{D}_1 = 0.65$), and mean annual electricity bill is $\bar{Y}_1 = \text{A\$1,740}$.
- Assigned standard brochure ($Z_i = 0$): 150 households complete the audit ($\bar{D}_0 = 0.15$), and mean annual electricity bill is $\bar{Y}_0 = \text{A\$1,920}$.

Step 1 — Decompose the Population into Compliance Subgroups (Assuming Monotonicity):

- **Always-Takers:** $P(\text{AT}) = \bar{D}_0 = 0.15$ (15%).
- **Never-Takers:** $P(\text{NT}) = 1 - \bar{D}_1 = 1 - 0.65 = 0.35$ (35%).
- **Compliers:** $P(\text{C}) = \bar{D}_1 - \bar{D}_0 = 0.65 - 0.15 = 0.50$ (50%).

Step 2 — Compute the Intention-to-Treat (ITT) and LATE (Wald Estimate):

$$\text{ITT} = \bar{Y}_1 - \bar{Y}_0 = 1,740 - 1,920 = -\$180$$

DegreeLive

$$\widehat{\text{LATE}} = \frac{\bar{Y}_1 - \bar{Y}_0}{\bar{D}_1 - \bar{D}_0} = \frac{-180}{0.50} = -\$360 \text{ per year}$$

Where:

- Sending the priority waiver reduces average household electricity bills by A$180/year (ITT), while actually completing the energy audit reduces electricity bills by A$360/year specifically among Complier households induced by the fee waiver.

Step 3 — Sargan Overidentification Illustration: Suppose a second independent instrument Z_{2i} (SMS reminder) is added ($L = 2$ instruments for 1 endogenous variable), and regressing the 2SLS residuals $\hat{u}_i$ on $(1, Z_{1i}, Z_{2i})$ across $N = 2,000$ households yields $R^2_{\hat{u}} = 0.0012$. The Sargan statistic is $S = 2,000 \times 0.0012 = 2.40$. Comparing 2.40 against $\chi^2_{1,0.95} = 3.84$ fails to reject H_0, indicating no statistical conflict between the two instruments.

3.8 Weak Instruments and Modern Robust Inference

When the excluded instruments $\mathbf{Z}_i$ explain very little variation in D_i—meaning the first-stage correlation is near zero—the IV estimator suffers from **Weak Instrument Bias** and severe non-normal sampling distributions, even in large samples ($N = 10,000$). Why? Because the denominator $\widehat{\text{Cov}}(\widetilde{Z}_i, \widetilde{D}_i)$ approaches zero and is noisy; finite-sample 2SLS becomes biased toward the inconsistent OLS estimator, while standard asymptotic t-tests dramatically over-reject.

First-Stage F-Statistic and Stock-Yogo Critical Values

In a single endogenous regressor model, instrument strength is diagnosed by the **first-stage F-statistic** testing $H_0 : \boldsymbol{\pi} = \mathbf{0}$ (using cluster- or heteroskedasticity-robust Kleibergen-Paap / Montiel Olea-Pflueger effective F when errors are non-spherical):

- **Staiger & Stock (1997) Rule of Thumb:** Historically, $F > 10$ indicated that maximal relative bias of 2SLS compared to OLS is under 10%.
- **Stock & Yogo (2005):** Tabulates exact critical values for both relative bias and **actual Wald test size distortion** (for $L = 1$ instrument, keeping true 5% test size below 10% requires $F > 16.38$).
- **Lee, McCrary, Moreira & Porter (2022) tF Procedure:** Proved that conventional t-ratio critical value 1.96 remains invalid unless $F > 104.7$! For moderate first-stage $F \in [10, 104.7]$, researchers should inflate the t-statistic critical value using Lee et al.'s tF adjustment table or bypass 2SLS t-tests altogether using weak-instrument-robust confidence sets.

Anderson-Rubin (AR) Weak-IV-Robust Inference

The gold-standard solution in just-identified ($L = 1$) models with potentially weak instruments is the **Anderson-Rubin (1949) test**, which retains exact asymptotic size regardless of first-stage strength—even when $F = 0$. To test $H_0 : \beta = \beta_0$, construct the transformed variable $Y_i - \beta_0 D_i$ and regress it directly on $(\mathbf{Z}_i, \mathbf{W}_i)$:

$$Y_i - \beta_0 D_i = \mathbf{Z}_i' \boldsymbol{\lambda}(\beta_0) + \mathbf{W}_i' \boldsymbol{\psi} + \eta_i$$

Where:

- Under $H_0 : \beta = \beta_0$ and valid exclusion $\text{Cov}(\mathbf{Z}_i, u_i) = 0$, the reduced-form coefficient $\boldsymbol{\lambda}(\beta_0) = \mathbf{0}$ identically!
- Testing $H_0 : \boldsymbol{\lambda}(\beta_0) = \mathbf{0}$ via a standard robust F-test (or Wald χ^2) yields a valid p-value free of denominator instability.
- Inverting the AR test across a grid of candidate values β_0 produces the 95% **Anderson-Rubin confidence set** $\{\beta_0 : AR(\beta_0) \leq \chi^2_{L,0.95}\}$—which automatically widens (or becomes unbounded) when instruments carry no identifying signal, faithfully reflecting genuine uncertainty.

Worked Numerical Example 3.3: Weak Instrument Bias and Anderson-Rubin Inversion

An analyst tests the effect of weekly social media hours (D_i) on teenage academic scores (Y_i) across $N = 1,000$ Australian students using a single weak policy instrument Z_i (local council broadband rollout speed) yielding first-stage $F = 4.5$. Suppose OLS gives $\hat{\beta}_{\text{OLS}} = -1.20$, while 2SLS gives $\hat{\beta}_{\text{2SLS}} = -2.80$ with naive standard error $\text{SE}_{\text{2SLS}} = 1.15$ ($t = -2.43$, seemingly significant at 5%).

Step 1 — Diagnose Relative Finite-Sample Bias: With $L = 1$ instrument and $F = 4.5 < 10$, the approximate finite-sample 2SLS bias ratio toward OLS is roughly $1/F \approx 22\%$, and Stock-Yogo size distortion exceeds 25% (the nominal 5% test actually rejects true nulls over 25% of the time).

Step 2 — Apply the Anderson-Rubin Test at $H_0 : \beta_0 = 0$: Under $\beta_0 = 0$, $Y_i - 0 \cdot D_i = Y_i$, so the AR statistic equals the reduced-form F-statistic of Y_i on Z_i (controlling for $\mathbf{W}_i$). Suppose the reduced-form slope is $\hat{\lambda}(0) = -0.35$ with $\text{SE} = 0.25$, giving $F_{\text{AR}}(0) = (-0.35/0.25)^2 = 1.96 < 3.84$ ($p = 0.161$).

Conclusion: Whereas the naive 2SLS t-test ($t = -2.43$) falsely claimed statistical significance at the 5% level, the exact weak-instrument-robust Anderson-Rubin test ($p = 0.161$) fails to reject $H_0 : \beta = 0$, preventing a spurious policy conclusion.

3.9 Applied IV Designs: Policy, Institutional, and Mendelian Instruments

Finding credible excluded instruments is one of the hardest tasks in applied econometrics. Four major quasi-experimental IV strategies dominate modern empirical research in economics, finance, and health.

1. Institutional & Policy Cutoffs (School-Entry Age & Eligibility Rules)

Following Angrist and Krueger's (1991) quarter-of-birth design, Australian labour and education economists exploit state school-starting cutoff dates (e.g., turning 5 before 30 April in Victoria vs. 30 June in NSW) as instruments for years of compulsory schooling or relative age-in-grade when studying adult earnings, university completion, and mental health.

Degreelive

2. Judge / Caseworker Leniency ("Examiner Designs")

When individuals, disability claims, bankruptcy filings, or Centrelink job-seeker files are quasi-randomly assigned to decision-makers (judges, patent examiners, or case managers) who vary systematically in their propensity to grant approval, the examiner's **leave-one-out leniency rate** serves as a continuous instrument for approval D_i:

$$Z_{ij} = \frac{1}{N_j - 1} \sum_{k \in \mathrm{J}(j),\ k \neq i} D_k$$

Where:

- N_j is the total caseload handled by examiner j (excluding unit i to avoid own-observation mechanical correlation).
- Valid examiner IV requires conditional random assignment across caseworkers within office-by-month strata, exclusion (the caseworker affects Y_i solely via the binary decision D_i, not via mentoring or fines), and average monotonicity.

3. Shift-Share (Bartik) Instruments

Regional labour and trade economists construct predicted local industry shocks by interacting pre-period local industry employment shares $s_{r,k,0}$ with national (leave-one-region-out) industry growth rates $g_{-r,k,t}$:

$$B_{r,t} = \sum_{k=1}^{K} s_{r,k,0} \cdot g_{-r,k,t}$$

Where:

- $s_{r,k,0}$ is region r's initial employment share in industry k (e.g., mining share in WA vs. manufacturing in Victoria).
- $g_{-r,k,t}$ is national growth in industry k excluding region r.
- As Goldsmith-Pinkham, Sorkin, and Swift (2020) and Borusyak, Hull, and Jaravel (2022) clarified, shift-share IV validity relies either on exogenous initial local industry shares ($s_{r,k,0}$) or on a large number of quasi-randomly assigned national industry shocks (g_k).

4. Mendelian Randomisation (Genetic Instruments in Health & Labour Economics)

Mendelian Randomisation (MR) uses random inheritance of genetic variants—Single Nucleotide Polymorphisms (SNPs)—at conception as instrumental variables to estimate the causal effect of modifiable health phenotypes (such as Body Mass Index, alcohol consumption, or depression) on economic outcomes (such as employment, wages, and medical expenditure in biobank-linked panels).

While random parental allele segregation satisfies baseline independence against postnatal socioeconomic confounders, applied econometricians must guard against three severe violations of the IV assumptions:

- **Horizontal Pleiotropy:** A single genetic variant directly affects multiple biological pathways simultaneously (e.g., an allele influencing both BMI and cognitive executive function), violating the exclusion restriction ($Z_i \to Y_i$ outside D_i). Sensitivity estimators such as **MR-Egger regression** and **Weighted Median MR** relax strict zero-pleiotropy under weaker InSIDE assumptions.
- **Linkage Disequilibrium (LD):** Nearby SNPs on the same chromosome are co-inherited rather than independent, meaning an observed SNP may proxy for a neighbouring causal gene affecting Y_i directly.
- **Dynastic Effects & Population Stratification:** Parental genotypes shape the childhood home environment ("genetic nurture"), and ancestral subgroup differences across regions correlate both with allele frequencies and socioeconomic outcomes unless controlled via genetic principal components and sibling-difference designs.

3.10 Regression Kink Design (RKD) for Continuous Policy Schedules

In many Australian transfer and taxation systems—such as the HECS-HELP student loan repayment schedule, Childcare Subsidy taper rates, and marginal income tax brackets—benefit levels remain continuous at a threshold c while the *slope* (marginal rate) changes abruptly. Standard RDD fails when there is no jump in treatment levels at c. Instead, Card, Lee, Pei, and Weber (2015) formalised the **Regression Kink Design (RKD)**, which identifies causal effects from the change in the first derivative of $E[Y_i \mid X_i = x]$ relative to the change in the first derivative of the policy schedule $b(x)$ at $x = c$:

$$\tau_{\text{RKD}} = \frac{\lim_{x \downarrow c} \frac{dE[Y_i|X_i=x]}{dx} - \lim_{x \uparrow c} \frac{dE[Y_i|X_i=x]}{dx}}{\lim_{x \downarrow c} b'(x) - \lim_{x \uparrow c} b'(x)}$$

Where:

- $b(x)$ is the deterministic (Sharp RKD) or expected (Fuzzy RKD) policy benefit or tax schedule as a function of running variable x.
- The denominator $\Delta b'(c) = b'_+(c) - b'_-(c) \neq 0$ measures the sharp kink in marginal incentives at c.
- Identification requires smoothness not only of the density $f_X(x)$ at c, but also of its first derivative $f'_X(x)$ and covariate slopes $E[W_i \mid X_i = x]'$ across the kink. Because estimating changes in derivatives requires larger samples than estimating level jumps, RKD typically uses local quadratic polynomials ($p = 2$) with bias-corrected bandwidths.

3.11 Simultaneous Equations Models (SEM) and Market Equilibrium

A classic motivation for Instrumental Variables arises in **Simultaneous Equations Models (SEM)**, where two or more endogenous variables are jointly determined in equilibrium. Consider a competitive Australian wholesale market (such as regional electricity or agricultural produce) where observed equilibrium log quantity Q_i and log price P_i satisfy both demand and supply equations simultaneously:

Degreelive

$$\text{Demand: } Q_i = \alpha_0 + \alpha_1 P_i + \alpha_2 Y_i + u_{di}, \qquad \alpha_1 < 0$$

$$\text{Supply: } Q_i = \beta_0 + \beta_1 P_i + \beta_2 W_i + u_{si}, \qquad \beta_1 > 0$$

Where:

- Q_i and P_i are jointly determined endogenous variables (equilibrium log quantity and log price).
- Y_i is an exogenous demand shifter (e.g., regional household disposable income in Victoria) excluded from supply.
- W_i is an exogenous supply shifter (e.g., wholesale natural gas input cost or regional rainfall index) excluded from demand.
- u_{di} and u_{si} are structural demand and supply shocks.

Solving the system for (P_i, Q_i) in terms of the exogenous variables (Y_i, W_i) yields the **Reduced-Form Equations**:

$$P_i = \pi_{10} + \pi_{11} Y_i + \pi_{12} W_i + v_{1i}, \qquad Q_i = \pi_{20} + \pi_{21} Y_i + \pi_{22} W_i + v_{2i}$$

Where:

- Because equilibrium P_i depends directly on u_{di} ($extCov(P_i, u_{di}) = \frac{\text{Var}(u_{di}) - \text{Cov}(u_{di}, u_{si})}{\beta_1 - \alpha_1} \neq 0$), regressing Q_i on P_i by OLS fails to recover either the demand elasticity α_1 or the supply elasticity β_1 (**simultaneity bias**).

Order and Rank Conditions for Structural Identification

In a linear system of M structural equations with M endogenous variables, an individual equation is identifiable only if exogenous shifters excluded from that equation shift the other equations:

- **Order Condition (Necessary):** To identify a structural equation containing k_{endog} endogenous right-hand-side variables, the number of exogenous variables excluded from that equation (K_{excl}) must be at least as large as k_{endog} ($K_{\text{excl}} \geq k_{\text{endog}}$). If $K_{\text{excl}} = k_{\text{endog}}$, the equation is *just-identified*; if $K_{\text{excl}} > k_{\text{endog}}$, it is *overidentified*.
- **Rank Condition (Necessary & Sufficient):** The excluded exogenous variables must actually appear with non-zero coefficients in the remaining structural equations (forming a full-rank submatrix of reduced-form coefficients), ensuring a non-zero first stage.

Worked Numerical Example 3.4: 2SLS Estimation of Demand & Supply Elasticities

Suppose OLS reduced-form regressions for quarterly Australian citrus market data (Q_i = log tonnes, P_i = log price per kg) on exogenous household income Y_i and regional rainfall shock W_i produce:

DegreeLive

$$P_i = 1.20 + 0.40Y_i - 0.50W_i$$

$$Q_i = 4.80 + 0.60Y_i + 0.40W_i$$

Where:

- Y_i is excluded from the Supply equation (identifying Supply), and W_i is excluded from the Demand equation (identifying Demand). Both equations satisfy the Order Condition ($1 = 1$) and Rank Condition ($\pi_{12} = -0.50 \neq 0$ and $\pi_{11} = 0.40 \neq 0$).

Step 1 — Recover the Price Elasticity of Demand (α_1) via Indirect Least Squares / 2SLS:

Since W_i shifts Supply without shifting Demand, the ratio of the reduced-form response of Q_i to W_i over the reduced-form response of P_i to W_i equals α_1:

$$\hat{\alpha}_1 = \frac{\hat{\pi}_{22}}{\hat{\pi}_{12}} = \frac{0.40}{-0.50} = -0.80$$

Where:

- A 1% rainfall-induced rise in citrus price lowers quantity demanded by 0.80% along the demand curve.

Step 2 — Recover the Price Elasticity of Supply (β_1) and Income Elasticity of Demand (α_2):

Since Y_i shifts Demand without shifting Supply, the supply elasticity is:

$$\hat{\beta}_1 = \frac{\hat{\pi}_{21}}{\hat{\pi}_{11}} = \frac{0.60}{0.40} = +1.50$$

And using the demand reduced-form identity $\pi_{21} = \alpha_1\pi_{11} + \alpha_2$ gives the income elasticity of demand:

$$\hat{\alpha}_2 = \hat{\pi}_{21} - \hat{\alpha}_1\hat{\pi}_{11} = 0.60 - (-0.80)(0.40) = 0.60 + 0.32 = +0.92$$

Where:

- The estimated price elasticity of supply is $+1.50$ and the income elasticity of demand is $+0.92$, both economically consistent and free of simultaneity bias.

DegreeLive