Experimental Analysis of Public Policies

Session 7

Regression Discontinuity

Etienne Dagorn - Universite de Lille

Session Objectives

By the end of this session, students should be able to:

- → audit an institutional rule by naming the running variable, cutoff, and treatment margin;
- → read an RDD graph before reading a coefficient table;
- → define the causal cutoff estimand before the observed discontinuity;
- → state continuity and no-manipulation logic in potential-outcomes language;
- → distinguish sharp RDD from fuzzy RDD;
- → interpret bandwidth, local fitting, diagnostics, and inference as design-supporting choices;
- → explain why RDD estimates are local and how fuzzy RDD reconnects to IV and LATE.

Roadmap

Roadmap

Institutional Rule and Local Comparison

Graph-First Intuition

Sharp RDD: Assumptions and Estimation

Diagnostics and Threats

Fuzzy RDD and External Validity

From Time Designs to Thresholds

- → DiD and event studies used timing and comparison groups to build counterfactual trends.
- → IV identified a local effect for people moved by an instrument.
- → RDD is local in a different way: units just around an institutional cutoff.

Central question

Can an institutional threshold create a locally credible comparison?

Course link: we again ask where the identifying variation comes from, which assumption makes it causal, and which population the estimate describes.

The RDD Idea

- → A rule changes treatment, eligibility, or treatment probability when a score crosses a cutoff.
- → Units very close to the cutoff should be similar in the running variable.
- → Under continuity and no other discontinuous change at the cutoff, an outcome jump can be interpreted as causal.
- → The estimand is local to the cutoff, not an average over the full score distribution.

In plain words: RDD compares just-missing and just-qualifying units, not all treated and untreated units.

RDD Workshop: Audit the Rule

Scenario: pupils scoring at least 70/100 receive an automatic scholarship.

1. What is the running variable?
2. What is the cutoff?
3. Is the design sharp or fuzzy?
4. Which pupils define the local comparison?
5. What manipulation or retesting behaviour would worry you?

Debrief

The credibility of RDD comes from the institutional rule and local comparability, not from fitting a flexible curve.

Expected output: one paragraph naming the running variable, cutoff, local population, and manipulation risk.

Threshold Designs in Practice

Institutional rule	Running variable	Possible outcome
Scholarship eligibility	test score	enrolment or graduation
Age eligibility	age	employment or health
Election victory	vote share	policy spending
Class-size rule	cohort enrolment	achievement
Income threshold	reported income	benefit take-up

Historical anchor

Maimonides' Rule works pedagogically because the institution creates a visible threshold in class size.

Running Variable and Cutoff

$$X_i = \text{running variable}, \qquad c = \text{cutoff}.$$

Sharp rule

$$D_i = 1[X_i \geq c].$$

- → In this session, X_i denotes the running variable; other covariates can be labelled W_i.
- → In a sharp RDD, the rule determines treatment received.
- → In a fuzzy RDD, the cutoff changes $Pr(D_i = 1 \mid X_i)$ rather than deterministically assigning everyone.
- → The comparison is local around c.

Roadmap

Read the RDD Graph

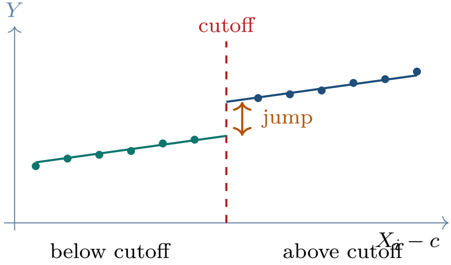

- → Are there enough observations near the cutoff?
- → Are fitted lines separate on each side?
- → Is the discontinuity at the cutoff visible?
- → Does anything else seem to jump there?

Reading order: graph, rule, density, covariates, bandwidth sensitivity, coefficient table.

Local Counterfactual

What is missing?

For units just above the cutoff, the missing counterfactual is what their outcome would have been just above c without treatment.

- → Units just below the cutoff help approximate that missing untreated outcome.
- → The comparison is credible only if untreated potential outcomes would move smoothly through c.
- → The running variable can predict outcomes smoothly; the identifying content is the discontinuous change at the rule.

Scholarship translation: marginal qualifiers are compared with marginal non-qualifiers, not with pupils far from the threshold.

Causal Cutoff Estimand

$$\tau_{SRD}(c) = E[Y_i(1) - Y_i(0) \mid X_i = c].$$

- → This is the sharp-RDD causal target: the treatment effect at the cutoff.
- → With a continuous running variable, read $X_i = c$ as the limiting population around the cutoff.
- → It is not the ATE for all eligible units and not the ATT for all treated units.
- → Define this causal object before writing the observed jump.

Policy interpretation: the estimand is often exactly the margin relevant for changing the threshold.

Observed Discontinuity and Identification

$$\Delta_Y(c) = \lim_{x \downarrow c} E[Y_i \mid X_i = x] - \lim_{x \uparrow c} E[Y_i \mid X_i = x].$$

Identification statement

Under the sharp rule and continuity of potential outcomes at c,

$$\Delta_Y(c) = \tau_{SRD}(c).$$

- $\Delta_Y(c)$ is the observed discontinuity.
- $\tau_{SRD}(c)$ is the causal estimand.
- The assumption links the observed jump to the missing counterfactual.

Roadmap

Sharp RDD

- Treatment status changes deterministically at the cutoff: $D_i = 1[X_i \geq c]$.
- Everyone just above the cutoff is treated; everyone just below is untreated.
- If potential outcomes are continuous at c, the outcome jump identifies the local treatment effect.
- The comparison deliberately ignores distant observations when they require too much extrapolation.

Example

Pupils above a score threshold receive a scholarship automatically.

Continuity Assumption

Continuity

In the absence of treatment, expected potential outcomes would be smooth at the cutoff:

$$\lim_{x \downarrow c} E[Y_i(d) \mid X_i = x] = \lim_{x \uparrow c} E[Y_i(d) \mid X_i = x], \quad d \in \{0, 1\}.$$

- → Units just below and above c should not differ discontinuously for other reasons.
- → The running variable may affect outcomes; it should not create a separate jump at c absent treatment.
- → Diagnostics can support this story, but they cannot prove the counterfactual smoothness.

No Other Jump at the Rule

- → The cutoff should not bundle multiple policies unless the estimand is the bundle.
- → Administrative appeals, retesting, or reclassification can create discontinuous selection.
- → A simultaneous rule for another programme can contaminate the outcome jump.

Causal language

Say "the jump is interpreted as causal under continuity and no other discontinuous change", not simply "the jump is the effect".

Bandwidth

- → The bandwidth defines how close to the cutoff observations must be.
- → Narrow bandwidth: better local comparison, less precision.
- → Wide bandwidth: more precision, more functional-form risk.
- → Report sensitivity to alternative bandwidths.

Tradeoff: bandwidth is where the local-randomisation intuition meets the need for enough data.

Bandwidth: Bias-Precision Debrief

	Narrow bandwidth	**Wide bandwidth**
Credibility	more local	more extrapolation
Precision	fewer observations	more observations
Risk	noisy estimate	functional-form bias

Reflex

Report whether the conclusion survives reasonable bandwidth changes. A single bandwidth is rarely persuasive alone.

Local Linear Regression

$$Y_i = \alpha + \tau D_i + \beta_1(X_i - c) + \beta_2 D_i(X_i - c) + u_i.$$

- → Centre the running variable at the cutoff.
- → Allow different slopes on each side.
- → Estimate using observations near the cutoff.
- → In a sharp RDD, τ estimates the jump at $X_i = c$, not the average treated-control difference.

Polynomial Warning

Avoid overfitting

High-order global polynomials can create artificial jumps and unstable estimates.

- → Prefer local methods.
- → Show sensitivity to bandwidth and functional form.
- → Let the graph and institutional rule guide interpretation.

Roadmap

RDD Diagnostic Workshop

Check	**What would worry us?**	**What it can show**
Density	bunching at c	possible manipulation
Covariates	jumps in W_i	comparability concern
Placebo cutoffs	jumps where no rule exists	functional-form concern
Bandwidths	sign changes under reasonable windows	fragile estimation
Institutional audit	another rule at c	bundled treatment

RDD Diagnostic Workshop: Activity

Visible output

For the scholarship rule, choose the two most important diagnostics and explain which assumption each one probes.

- → Name the diagnostic.
- → State the threat it probes.
- → Explain why the diagnostic cannot prove the identifying assumption.

Covariate Continuity

- → Pre-treatment covariates should not jump at the cutoff.
- → Check age, baseline outcomes, demographics, prior trends, or institutional variables.
- → Covariate discontinuities suggest the cutoff captures more than treatment.
- → Smooth covariates support the design but do not prove continuity of potential outcomes.

Example: if baseline test scores jump exactly at a scholarship cutoff beyond the mechanical threshold variable, another selection process may be operating.

Manipulation and Density

- → Units may manipulate the running variable to cross the threshold.
- → Schools may reclassify pupils; firms may adjust reported employment; households may time eligibility.
- → A density jump at c is most worrying when agents know the cutoff and can precisely affect X_i.
- → A smooth density supports, but does not prove, local comparability.

Threat

Precise manipulation breaks the idea that units just around the cut-off are comparable.

Placebo Cutoffs

- → Estimate discontinuities at thresholds where no treatment changes.
- → Large placebo jumps weaken credibility.
- → Placebo cutoffs should be chosen transparently.
- → They help detect functional-form artefacts or unstable relationships between X_i and Y_i.

Good placebo: choose fake cutoffs on both sides of the true cutoff, while keeping enough observations nearby.

Donut RDD

- Exclude observations extremely close to the cutoff.
- Useful for local heaping, appeals, retesting, or administrative quirks around c.
- It changes the local population and can move the estimand away from the cutoff.
- If manipulation changes who appears on each side, a donut check does not automatically repair the design.

Example: if scores can be appealed within one point of the cutoff, compare results with those observations removed, but still discuss selection.

Threats

Threat	Diagnostic or response
Manipulation	density and institutional checks
Other policy cutoff	rule audit and covariate continuity
Functional form	graph, bandwidth, and specification sensitivity
Noisy local estimate	confidence intervals and transparent bandwidth choice
Local external validity	describe units near cutoff

Roadmap

Institutional Rule and Local Comparison

Graph-First Intuition

Sharp RDD: Assumptions and Estimation

Diagnostics and Threats

Fuzzy RDD and External Validity

Fuzzy RDD

$$Z_i = 1[X_i \geq c], \qquad D_i \neq Z_i \text{ for some units.}$$

- → Treatment probability jumps at the cutoff but not from 0 to 1.
- → The cutoff indicator Z_i becomes a candidate instrument for treatment received D_i.
- → The estimand is a local IV effect for compliers at the cutoff.
- → A first-stage jump is necessary, but not enough for validity.

Connection: fuzzy RDD imports the IV logic from Session 4 into a threshold setting.

Wald Ratio at the Cutoff

$$\Delta_D(c) = \lim_{x \downarrow c} E[D_i \mid X_i = x] - \lim_{x \uparrow c} E[D_i \mid X_i = x].$$

$$\tau_{FRD}(c) = \frac{\Delta_Y(c)}{\Delta_D(c)}.$$

→ Numerator: reduced-form outcome jump at the cutoff.

→ Denominator: first-stage treatment-probability jump at the cutoff.

→ Interpretation mirrors LATE: effect for local compliers moved by crossing the threshold.

Fuzzy RDD Assumptions at the Cutoff

Assumption	Local question
Relevance	Does crossing c visibly change D_i?
Continuity	Would potential outcomes and potential treatment behaviour be smooth through c?
Exclusion	Does the cutoff affect Y_i only through treatment received?
Monotonicity	Is anyone less likely to receive treatment because they are just above c?

Same IV discipline

Do not treat a threshold first stage as proof of exclusion or monotonicity.

External Validity

- → Sharp RDD estimates effects for units near the cutoff.
- → Fuzzy RDD estimates effects for local compliers near the cutoff.
- → Effects may differ for units far above or below the threshold.
- → This local estimate can be highly policy-relevant when the decision is whether to move the cutoff.

Policy reflex: distinguish moving a threshold from universal expansion; they often require different estimands.

Inference and Reporting

- → Use standard errors compatible with local estimation and bandwidth choice.
- → Account for clustering if treatment or shocks vary by group.
- → Report confidence intervals, not only point estimates.
- → For fuzzy designs, report the treatment jump as carefully as the outcome jump.

Minimum reporting: graph, rule, bandwidth, observations on each side, main estimate, sensitivity checks, density test, and covariate checks.

Code Mini-Lab

- → Plot the running variable and outcome.
- → Estimate a local linear RDD.
- → Vary the bandwidth.
- → Check covariate continuity around the cutoff.
- → Interpret a fuzzy design by reporting both $\Delta_Y(c)$ and $\Delta_D(c)$.

Student deliverable: one graph and one table, with a sentence explaining how the estimate changes when the bandwidth changes.

Exercise Bridge

- → Theory sheet: `exercises/theory/session_07_theory.md`.
- → QCM: `assessments/qcm/session_07_qcm.md`.
- → Applied exercise: assess an eligibility-threshold design.
- → Exam skill: explain the local estimand and validity checks.

Exercise rhythm: identify the rule, state continuity, inspect diagnostics, then interpret the local effect.

Design Checklist

1. What is the running variable?
2. What is the cutoff rule?
3. Is treatment sharp or fuzzy?
4. What is the causal cutoff estimand?
5. What assumption links the observed jump to the causal effect?
6. Is manipulation plausible?
7. Do covariates or other rules jump at the cutoff?
8. How local is the estimand?

Transition to Session 8

- → RDD is powerful when a rule creates local variation.
- → Its credibility is often local: cutoff effects or cutoff-complier effects.
- → Some policies affect one city, region, or country with no lottery, no clean threshold, and few treated units.
- → Session 8 asks how to construct a counterfactual, assess heterogeneity and external validity, and stress-test credibility.

Bridge question: when there is no threshold and no many-treated-units panel, what comparison can we defend?

Key Takeaways

- → RDD starts from an institutional rule, not from curve fitting.
- → The causal cutoff estimand is distinct from the observed discontinuity.
- → Continuity, no other cutoff jump, and no precise manipulation carry the causal interpretation.
- → Bandwidth choice trades local credibility against precision.
- → Fuzzy RDD is a local IV design with IV assumptions at the cutoff.

One-sentence memory aid: RDD is a local comparison made credible by a discontinuous rule and smooth counterfactual outcomes.

Appendix: Extensions

- → Regression kink designs use changes in slope rather than jumps.
- → Geographic RDD uses borders as cutoffs.
- → Multiple cutoffs can improve external-validity discussion.
- → These extensions keep the same local-comparison logic.

Student focus: these extensions are optional; the core exam skill is the sharp/fuzzy cutoff design.

References and Further Reading

- → Angrist and Lavy (1999), "Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement".
- → Lee and Lemieux (2010), "Regression Discontinuity Designs in Economics".
- → Imbens and Lemieux (2008), "Regression Discontinuity Designs: A Guide to Practice".
- → Cattaneo, Idrobo, and Titiunik (2020), *A Practical Introduction to Regression Discontinuity Designs: Foundations*.
- → Calonico, Cattaneo, and Titiunik (2014), "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs".

Reading reflex

Start with the rule, graph, running-variable density, and covariate checks before interpreting the treatment jump.