Experimental Analysis of Public Policies

Session 1

Foundations of Policy Evaluation

Etienne Dagorn - Universite de Lille

Session Objectives

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

- → translate a substantive policy question into unit, treatment, outcome, population, horizon, and decision margin;
- → define ATE, ATT, and ATU in words and notation, with ITT and SATE/PATE flagged as previews;
- → use potential outcomes to state the missing counterfactual precisely;
- → explain when a treatment effect is not well defined because of vague treatment versions, timing, or interference;
- → decompose a naive comparison into a causal component and selection bias;
- → explain why random assignment is the benchmark for internal validity and why external validity is a separate question.

Roadmap

From Policy Questions to Causal Estimands

Potential Outcomes and the Missing Counterfactual

When Is a Treatment Effect Well Defined?

Selection, Identification, and the Experimental Benchmark

From Experiments to Real-World Policy Evaluation

Roadmap

From Policy Questions to Causal Estimands

Potential Outcomes and the Missing Counterfactual

When Is a Treatment Effect Well Defined?

Selection, Identification, and the Experimental Benchmark

From Experiments to Real-World Policy Evaluation

From Last Time to Today

Last time

- → association is not enough for causation;
- → causal claims require counterfactuals;
- → identification depends on variation and assumptions.

Today

- → what causal effect are we estimating?
- → for whom?
- → under which treatment definition?
- → over what time horizon?

Shift: from causal intuition to the formal design of a policy evaluation.

Opening Activity: Headline to Estimand

Two minutes

The mayor says: “Free tutoring raised scores, so we should scale it.”

1. What is the unit?
2. Is the policy lever an offer, receipt, or hours?
3. What outcome and horizon matter?
4. Which pupils define the target population?

Output

Write one design sentence before choosing a method.

Stage 1: A City Introduces Tutoring

- → The city funds free after-school tutoring for eligible pupils.
- → The mayor asks: “Does tutoring work?”
- → The education department has test scores, school records, and attendance logs.

First design reaction

“Does tutoring work?” is a policy motivation, not yet a causal estimand.

Why: the answer changes if we mean an offer, actual attendance, 30 hours of tutoring, short-run scores, graduation, participants, or all eligible pupils.

The Policy-Evaluation Objects

Object	Tutoring example
Unit of analysis	pupil, classroom, school, or cohort
Treatment	tutoring offer, participation, or hours attended
Outcome	test score, grades, attendance, graduation, later earnings
Target population	eligible pupils in the city, in pilot schools, or nationally
Time horizon	end of semester, one year, or longer term
Policy decision	continue, scale, retarget, redesign, or stop
Causal estimand	the average effect that matches the decision

Discipline

Before choosing a method, write these objects in one design paragraph.

Changing One Object Changes the Question

Change	Different evaluation question
Outcome: test score vs graduation	learning this year or educational trajectory?
Population: pilot schools vs all eligible pupils	programme performance or scale-up?
Horizon: semester vs two years	immediate gains or persistent gains?
Treatment: offer vs attendance	policy access or treatment receipt?
Unit: pupil vs school	individual tutoring or school-level programme?

Key point

A causal effect is always attached to a defined intervention, population, outcome, and time.

Treatment Is Not One Word

$$Z_i = \text{being offered tutoring}$$

$$D_i = \text{actually participating in tutoring}$$

$$H_i = \text{number of tutoring hours attended}$$

→ Z_i is the offer or assignment that the policy can control.

→ D_i is treatment received.

→ Each can answer a legitimate but different policy question.

→ A design is ambiguous if it calls all three simply “treatment.”

Interactive: What Is the Treatment?

Tutoring scenario

The city randomly offers tutoring to eligible pupils, but only 60% of pupils offered a place attend at least one session.

Question: what exactly is the treatment?

Answer

Interactive: What Is the Treatment?

Tutoring scenario

The city randomly offers tutoring to eligible pupils, but only 60% of pupils offered a place attend at least one session.

Question: what exactly is the treatment?

→ Z_i: being offered a tutoring place.

Answer

Interactive: What Is the Treatment?

Tutoring scenario

The city randomly offers tutoring to eligible pupils, but only 60% of pupils offered a place attend at least one session.

Question: what exactly is the treatment?

- → Z_i: being offered a tutoring place.
- → D_i: attending any tutoring.

Answer

Interactive: What Is the Treatment?

Tutoring scenario

The city randomly offers tutoring to eligible pupils, but only 60% of pupils offered a place attend at least one session.

Question: what exactly is the treatment?

- → Z_i: being offered a tutoring place.
- → D_i: attending any tutoring.
- → H_i: attending a given number of hours.

Answer

Interactive: What Is the Treatment?

Tutoring scenario

The city randomly offers tutoring to eligible pupils, but only 60% of pupils offered a place attend at least one session.

Question: what exactly is the treatment?

- $\rightarrow$ Z_i: being offered a tutoring place.
- $\rightarrow$ D_i: attending any tutoring.
- $\rightarrow$ H_i: attending a given number of hours.

Answer

The randomised policy lever is the offer. Actual participation and hours attended require additional assumptions if we want to interpret their effects causally.

From Question to Estimand

Working definition

A causal estimand is the causal quantity we want to learn before looking at the data.

- → It is not the regression coefficient by default.
- → It is chosen by the policy decision and the population of interest.
- → It must specify the comparison between potential outcomes.

Example: among all eligible pupils, what is the average effect of being offered tutoring this school year on end-of-year maths scores?

Stage 1 Design Sentence

Fill the blanks

For ____________ pupils, estimate the effect of ______________ on ______________ measured ____________, in order to decide whether to ____________.

Stage 1 Design Sentence

Fill the blanks

For ____________ pupils, estimate the effect of _____________ on _____________ measured ___________, in order to decide whether to ____________.

One possible answer

For eligible pupils, estimate the effect of being offered tutoring on end-of-year test scores measured at the end of the school year, in order to decide whether to scale the programme.

Stage 1 Design Sentence

Fill the blanks

For ____________ pupils, estimate the effect of _____________ on _____________ measured ___________, in order to decide whether to ____________.

One possible answer

For eligible pupils, estimate the effect of being offered tutoring on end-of-year test scores measured at the end of the school year, in order to decide whether to scale the programme.

Different answers are allowed, but each answer implies a different estimand.

Roadmap

From Policy Questions to Causal Estimands

Potential Outcomes and the Missing Counterfactual

When Is a Treatment Effect Well Defined?

Selection, Identification, and the Experimental Benchmark

From Experiments to Real-World Policy Evaluation

Potential Outcomes

$$Y_i(1) = \text{outcome for pupil } i \text{ if treated}$$
$$Y_i(0) = \text{outcome for pupil } i \text{ if not treated}$$

- → Each pupil has two potential outcomes for the same outcome and horizon.
- → The notation forces us to name the counterfactual state.
- → The treatment definition must already be clear before we write $Y_i(1)$ and $Y_i(0)$.

Tutoring translation: after we define the treatment, $Y_i(1)$ is the score under that treatment and $Y_i(0)$ is the score without it, for the same pupil and school year.

Individual Treatment Effects

$$\tau_i = Y_i(1) - Y_i(0)$$

- → τ_i is the individual causal effect.
- → It compares two potential outcomes for the same pupil.
- → It is conceptually simple and empirically impossible to observe directly.

Fundamental problem

For the same pupil at the same time, we never observe both $Y_i(1)$ and $Y_i(0)$.

Observed Outcome

$$Y_i = D_i Y_i(1) + (1 - D_i) Y_i(0)$$

If treated

$$Y_i = Y_i(1)$$

The untreated outcome $Y_i(0)$ is missing.

If untreated

$$Y_i = Y_i(0)$$

The treated outcome $Y_i(1)$ is missing.

The Missing Counterfactual in a Table

Pupil	$\boldsymbol{Y_i(1)}$	$\boldsymbol{Y_i(0)}$	$\boldsymbol{\tau_i}$	$\boldsymbol{D_i}$	$\boldsymbol{Y_i}$
Alice	78			1	78
Benoit		62		0	62
Chloe	84			1	84
David		71		0	71

Read the table

Observed outcomes are real data.

The Missing Counterfactual in a Table

Pupil	$Y_i(1)$	$Y_i(0)$	τ_i	D_i	Y_i
Alice	78	?		1	78
Benoit	?	62		0	62
Chloe	84	?		1	84
David	?	71		0	71

Read the table

Observed outcomes are real data. The missing cells are counterfactuals.

The Missing Counterfactual in a Table

Pupil	$\boldsymbol{Y_i(1)}$	$\boldsymbol{Y_i(0)}$	$\boldsymbol{\tau_i}$	$\boldsymbol{D_i}$	$\boldsymbol{Y_i}$
Alice	78	?	?	1	78
Benoit	?	62	?	0	62
Chloe	84	?	?	1	84
David	?	71	?	0	71

Read the table

Observed outcomes are real data. The missing cells are counterfactuals. Individual effects require both cells, so we estimate average effects using design.

Average Treatment Effect

$$ATE = E[Y_i(1) - Y_i(0)]$$

→ The ATE averages individual effects over the target population.

→ It answers a scale-up style question: what would happen, on average, if eligible pupils were treated rather than untreated?

→ The target population must be named for the ATE to be interpretable.

Tutoring example: the average effect of tutoring for all pupils eligible under the city's rule.

ATT and ATU

$$ATT = E[Y_i(1) - Y_i(0) \mid D_i = 1]$$
$$ATU = E[Y_i(1) - Y_i(0) \mid D_i = 0]$$

- → ATT: effect for pupils who actually receive treatment.
- → ATU: effect untreated pupils would have if treated.
- → They can differ when treatment effects are heterogeneous or selection is systematic.

Notation note: LATE is deliberately not introduced here; it requires an instrument and compliance behaviour.

Estimands as Policy Questions

Policy question	Relevant estimand
What is the effect on current participants?	ATT
What would be the effect if eligible pupils were treated?	ATE
What would happen to current non-participants if they took up tutoring?	ATU
What is the effect of offering tutoring as a policy lever?	ITT, introduced later

Choice rule

The estimand should match the decision margin: evaluate current performance, expand coverage, target non-participants, or change eligibility.

Interactive: Which Estimand?

1. The city asks whether tutoring helped pupils who enrolled.
2. The city asks whether every eligible pupil should be guaranteed a place.
3. The city asks whether non-participants should receive stronger encouragement.

Think first

Name the estimand before choosing any regression, experiment, or graph.

Interactive: Which Estimand? Debrief

Answers

1. ATT.
2. ATE of the offer/access policy; under full compliance, this coincides with a treatment-receipt ATE.
3. ATU or an offer effect, depending on the policy lever.

Heterogeneous Effects

$$\tau_i = Y_i(1) - Y_i(0) \quad \text{may differ across pupils.}$$

Group	**Untreated risk**	**Effect**
High baseline score	low	small
Middle baseline score	medium	moderate
Low baseline score	high	large

Preview

Average effects can differ across populations because individual effects can differ across pupils.

- → X_i is a pre-treatment characteristic.
- → Effects in one group need not equal effects in another.

Implication: ATE, ATT, ATU, and effects in another population may differ for substantive reasons.

A Simple Interaction Picture

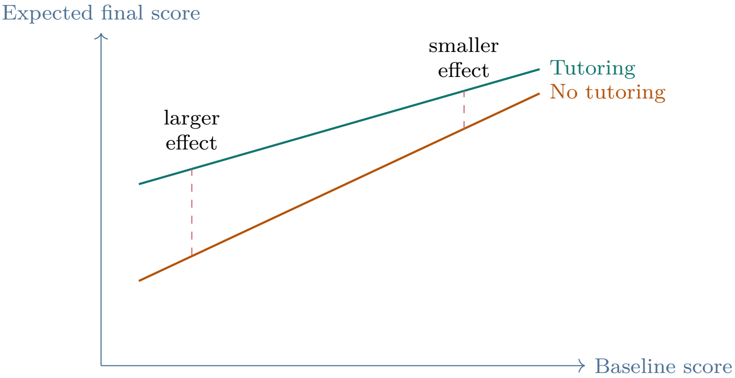

Reading

Heterogeneity is not a technical detail. It determines which pupils, schools, or regions a policy should prioritize.

Roadmap

From Policy Questions to Causal Estimands

Potential Outcomes and the Missing Counterfactual

When Is a Treatment Effect Well Defined?

Selection, Identification, and the Experimental Benchmark

From Experiments to Real-World Policy Evaluation

Before Identification: Define the Treatment

Conceptual order

First ask whether the causal effect is well defined. Then ask whether it is identified.

- → What exactly happens to a treated pupil?
- → Can untreated pupils receive something similar?
- → Can one pupil's treatment change another pupil's outcome?
- → Are treatment, covariates, and outcomes ordered correctly in time?

Reason: vague treatments produce vague potential outcomes.

Consistency

Consistency

If pupil i actually receives treatment level d, then the observed outcome equals the potential outcome under that level:

$$Y_i = Y_i(d).$$

- → This sounds obvious only when d is well defined.
- → "Tutoring" must correspond to a clear intervention.
- → Otherwise $Y_i(1)$ bundles together many different versions of treatment.

Multiple Versions of Treatment

Version called "tutoring"	Why it may matter
1 hour online each month	low intensity, limited personalization
30 hours face-to-face	much higher exposure
Small-group tutoring	peer interaction and lower cost
One-to-one tutoring	more individualized attention
Exam-preparation tutoring	short-run test strategy vs learning

Interpretation

If different versions have different effects, the estimate describes the implemented package, not an abstract word.

No Interference

Standard notation

$$Y_i(d_i)$$

Pupil i's outcome depends only on pupil i's treatment.

More general world

$$Y_i(d_1, \ldots, d_N)$$

Pupil i's outcome may depend on classmates' treatment too.

- → Treated pupils may share materials with untreated classmates.
- → Teachers may change classroom instruction because some pupils receive tutoring.
- → Programme capacity may change school resources.

Classroom Spillovers

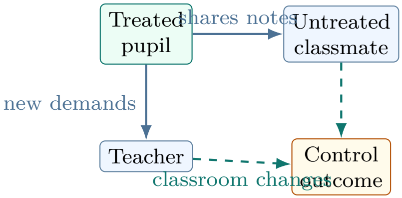

Design question

Should the evaluation randomize pupils, classrooms, or schools if peer spillovers are likely?

Timing

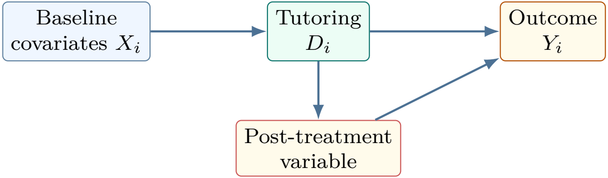

→ Specify when treatment starts.

→ Specify when the outcome is measured.

→ Check that control variables are truly pre-treatment.

Bridge to later material: controlling for a post-treatment variable can change the estimand or create bias.

Well-Defined Effect Checklist

1. Is the treatment version specific enough to define $Y_i(1)$?
2. Is the comparison condition specific enough to define $Y_i(0)$?
3. Is the unit of analysis clear?
4. Are outcomes measured after treatment?
5. Could one unit's treatment affect another unit's outcome?
6. Does the estimand match the policy implementation?

Stage 2

Before debating selection bias, revise the tutoring protocol so another team could implement the same treatment.

Roadmap

From Policy Questions to Causal Estimands

Potential Outcomes and the Missing Counterfactual

When Is a Treatment Effect Well Defined?

Selection, Identification, and the Experimental Benchmark

From Experiments to Real-World Policy Evaluation

Stage 2: Participation Is Voluntary

- → The city opens tutoring places to eligible pupils.
- → Some families enrol; others do not.
- → Schools promote the programme with different intensity.

Tempting comparison

Compare participants and non-participants at the end of the year.

Design question: do non-participants represent what participants would have scored without tutoring?

The Naive Comparison

$$\Delta^{naive} = E[Y_i \mid D_i = 1] - E[Y_i \mid D_i = 0]$$

→ Easy to compute.

→ Useful as a descriptive starting point.

→ Causal only if untreated potential outcomes are comparable across groups.

Ask

Before seeing any estimate, should we expect participants to have higher or lower $Y_i(0)$ than non-participants?

Selection-Bias Decomposition

$$\begin{aligned} E[Y_i \mid D_i = 1] - E[Y_i \mid D_i = 0] &= E[Y_i(1) \mid D_i = 1] - E[Y_i(0) \mid D_i = 0] \\ &= E[Y_i(1) - Y_i(0) \mid D_i = 1] \\ &\quad + E[Y_i(0) \mid D_i = 1] - E[Y_i(0) \mid D_i = 0]. \end{aligned}$$

ATT

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

Selection bias

$$\begin{aligned} &E[Y_i(0) \mid D_i = 1] \\ &\quad - E[Y_i(0) \mid D_i = 0] \end{aligned}$$

What the Decomposition Teaches

- → The naive comparison is not "wrong" mathematically.
- → It answers a descriptive question unless the selection-bias term is zero.
- → Selection bias is about untreated potential outcomes, $Y_i(0)$.
- → Its sign depends on how pupils enter or are targeted into tutoring.

Identification

An identification argument explains why an observed comparison equals the intended causal estimand under explicit assumptions.

Interactive Sign of Bias: Scenario A

Scenario A

Highly motivated families are more likely to enrol in tutoring.

Question:

$$E[Y_i(0) \mid D_i = 1] \quad \text{is greater or smaller than} \quad E[Y_i(0) \mid D_i = 0]?$$

Answer

Implication

Interactive Sign of Bias: Scenario A

Scenario A

Highly motivated families are more likely to enrol in tutoring.

Question:

$$E[Y_i(0) \mid D_i = 1] \quad \text{is greater or smaller than} \quad E[Y_i(0) \mid D_i = 0]?$$

Answer

Likely greater: participants would have scored higher even without tutoring.

Implication

Interactive Sign of Bias: Scenario A

Scenario A

Highly motivated families are more likely to enrol in tutoring.

Question:

$$E[Y_i(0) \mid D_i = 1] \quad \text{is greater or smaller than} \quad E[Y_i(0) \mid D_i = 0]?$$

Answer

Likely greater: participants would have scored higher even without tutoring.

Implication

The naive participant–non-participant comparison may overstate the effect.

Interactive Sign of Bias: Scenario B

Scenario B

Schools target tutoring at pupils who are already struggling.

Question:

$$E[Y_i(0) \mid D_i = 1] \quad \text{is greater or smaller than} \quad E[Y_i(0) \mid D_i = 0]?$$

Answer

Implication

Interactive Sign of Bias: Scenario B

Scenario B

Schools target tutoring at pupils who are already struggling.

Question:

$$E[Y_i(0) \mid D_i = 1] \quad \text{is greater or smaller than} \quad E[Y_i(0) \mid D_i = 0]?$$

Answer

Likely smaller: treated pupils would have scored lower without tutoring.

Implication

Interactive Sign of Bias: Scenario B

Scenario B

Schools target tutoring at pupils who are already struggling.

Question:

$$E[Y_i(0) \mid D_i = 1] \quad \text{is greater or smaller than} \quad E[Y_i(0) \mid D_i = 0]?$$

Answer

Likely smaller: treated pupils would have scored lower without tutoring.

Implication

The naive comparison may understate the effect.

Stage 3: Capacity Is Limited

- → Eligible demand exceeds available tutoring places.
- → The city randomly assigns offers among eligible pupils.
- → Random assignment is decided before outcomes are realised.
- → For this clean benchmark, suppose every offered pupil attends: $D_i = Z_i$.

Experimental benchmark

$$Z_i \perp (Y_i(1), Y_i(0))$$

Assignment is independent of the potential outcomes relevant to the treatment contrast.

Interpretation: assigned arms are comparable in expectation because the offer does not use expected outcomes.

Why Random Assignment Works

If

$$Z_i \perp (Y_i(1), Y_i(0)),$$

then

$$E[Y_i(0) \mid Z_i = 1] = E[Y_i(0) \mid Z_i = 0]$$

and

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

- → The selection-bias term is zero in expectation.
- → Treatment and control groups can stand in for each other's missing potential outcomes.
- → In the full-compliance benchmark, $D_i = Z_i$, so assigned arms are also receipt groups.

What the Difference in Means Identifies

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

→ With full compliance, no interference, and observed outcomes, assigned-arm means identify the average treatment-receipt effect in the randomised study sample.

→ For a fixed experimental sample, the sample difference estimates SATE in expectation over the randomisation.

→ If the randomised sample represents a larger target population, it can also inform PATE.

→ Randomisation solves baseline selection; it does not by itself solve attrition, spillovers, or vague treatments.

Roadmap

From Policy Questions to Causal Estimands

Potential Outcomes and the Missing Counterfactual

When Is a Treatment Effect Well Defined?

Selection, Identification, and the Experimental Benchmark

From Experiments to Real-World Policy Evaluation

Stage 4: Offers Are Randomised, Attendance Is Not

$Z_i =$ randomised offer or assignment $\qquad D_i =$ actual participation

- → The city controls who receives an offer.
- → Pupils and families decide whether to attend.
- → Suppose only 60% of pupils offered tutoring attend.

Question

What effect does the experiment identify immediately: the effect of the offer or the effect of attending?

Preview: Intention-to-Treat

$$ITT = E[Y_i(1_Z) - Y_i(0_Z)]$$

$$\text{Under randomised } Z: \quad ITT = E[Y_i \mid Z_i = 1] - E[Y_i \mid Z_i = 0]$$

→ Here potential outcomes are indexed by the offer Z, not by receipt D.

→ The causal ITT answers: what is the effect of being offered the policy?

→ The observed assigned-arm contrast identifies it under random assignment, no interference, and observed outcomes.

→ It includes take-up behaviour and implementation frictions.

Preview only

If we later want the causal effect of treatment receipt using randomised assignment, instrumental variables can help under additional assumptions.

Stage 5: Pupils Share Materials

- → Treated pupils share worksheets with untreated classmates.
- → Teachers adjust classroom examples after seeing tutoring content.
- → Control pupils may no longer represent "no exposure to tutoring."

Question

What does the treatment-control difference estimate now?

Lesson

Real-world policy evaluation depends on the treatment environment, not only the assignment rule.

Stage 6: Some Outcomes Are Missing

Scenario

Twenty percent of pupils do not take the final test.

- → Random assignment solved baseline treatment selection.
- → Missing outcomes can create a new selection problem after assignment.
- → The threat is sharper if attrition differs by assignment or expected outcome.

Question

Is random assignment enough if outcomes are missing selectively?

What Randomisation Does Not Automatically Solve

Issue	Why it matters
Non-compliance	assignment differs from treatment receipt
Attrition	observed outcomes may be selected after assignment
Spillovers	control outcomes may be affected by treated units
Multiple treatment versions	$Y_i(1)$ may mix different interventions
Measurement	outcome changes may reflect measurement changes
External validity	the effect may not travel to another context

Course

arc: later sessions study designs and diagnostics for these threats one by one.

Preview: Random Assignment vs Random Sampling

Random assignment

Who receives treatment?

Purpose: internal validity and causal identification.

Random sampling

Who enters the study?

Purpose: representativeness and generalisation.

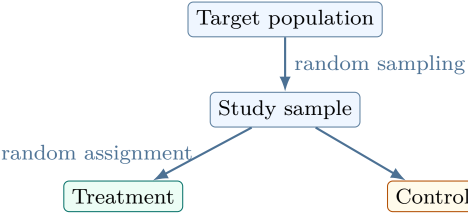

Preview: Four Possible Designs

Design	What it can support
Random assignment, no random sampling	strong internal validity, narrower generalisation
Random sampling, no random assignment	representative description, weak causal identification
Both random sampling and assignment	strong design for causality and population inference
Neither	requires a careful observational identification argument

Common confusion

Random sampling is not required for internal validity. Random assignment is not enough for external validity.

Preview: SATE and PATE

$$SATE = \frac{1}{N}\sum_{i=1}^{N}[Y_i(1) - Y_i(0)]$$

$$PATE = E[Y(1) - Y(0)]$$

- → SATE refers to the units in the experimental sample.
- → PATE refers to a target population.
- → Random assignment identifies effects in the study sample.
- → Generalising beyond that sample is a separate design and interpretation question.

Internal and External Validity in This Framework

Validity question	Framework translation
Internal validity	Does the design identify the intended causal effect in the study?
External validity	To what population, setting, implementation, or policy context does the effect generalise?

- → External validity is broader than sample representativeness.
- → It also depends on treatment-effect heterogeneity, implementation quality, institutional context, and target population.
- → A strong experiment can answer a narrow policy question very well.

Stage 7: Should the City Scale Nationally?

- → The pilot used trained tutors in motivated schools.
- → National rollout may involve different pupils, schools, budgets, and tutors.
- → The treatment version and target population may change.

External-validity question

Is the pilot effect informative for the national policy environment, or only for the pilot as implemented?

Policy interpretation: scaling a programme changes who receives it, who delivers it, and sometimes what the treatment actually is.

Evolving Tutoring Case: Full Arc

Stage	Policy situation	Design question
1	City introduces tutoring	What is the estimand?
2	Participation is voluntary	Why is selection a threat?
3	Offers are randomised	What does randomisation solve?
4	Only 60% attend	Is treatment offer or receipt?
5	Pupils share materials	What assumption may fail?
6	20% miss the final test	What new selection threat appears?
7	National expansion is considered	Where does the estimate apply?

Teaching point

The same policy generates different causal questions as implementation details become visible.

How to Read a Policy-Evaluation Paper

1. What policy decision motivates the study?
2. What exactly are Z, D, Y, X, the population, and the horizon?
3. What causal estimand is being targeted?
4. What is the missing counterfactual?
5. What creates identifying variation?
6. What assumption makes the comparison causal?
7. What diagnostics or threats should change our confidence?
8. To whom and where does the estimate apply?

Reading habit

Do not start by asking whether the table is significant. Start by asking what causal object the table is trying to estimate.

Final Design Logic

Question	What exactly are we trying to learn when evaluating a policy?
Ideal experiment	Randomly assign a clearly defined offer or treatment within the study population.
Counterfactual	The same units under the other treatment state.
Variation	The random assignment, or a comparison justified by a later design.
Assumption	The comparison units are exchangeable in potential outcomes for the intended estimand.
Threat	Vague treatment versions, spillovers, noncompliance, attrition, measurement, and limited external validity.

Transition to Session 2

- → We now know what causal effect we want.
- → We know what the ideal experiment would look like.
- → We know why selection creates bias in voluntary participation.

Next

What can we do when treatment is not randomised but we observe rich pre-treatment covariates?

Bridge: regression adjustment and matching will try to build a credible comparison group using observed baseline information.