

CLM: Asymptotic Assumptions • Last semester, to get asymptotic results for OLS, we presented a new set of assumptions for the CLM: (A1) DGP: $y = X \beta + \varepsilon$. (A2') X stochastic, but $E[X'\varepsilon] = 0$ and $E[\varepsilon] = 0$. (A3) $Var[\varepsilon|X] = \sigma^2 I_T$ (A4') plim (X'X/T) = Q (p.d. matrix with finite elements, rank=k) • We studied the large sample properties of OLS: - b and s² are consistent - b $\stackrel{a}{\longrightarrow} N(\beta, (\sigma^2/T) Q^{-1})$ - *t-tests* asymptotically $\chi^2_{rank}(Var[m])$. - Small sample behavior: Do simulations and/or bootstrapping.

The IV Problem

We start with our CLM: y = Xβ + ε. (DGP) - Let's pre-multiplying the DGP by X' X' y = X' Xβ + X' ε. - We can interpret b as the solution obtained by first approximating X'ε by zero, and then solving the k equations in k unknowns X' y = X'X b (normal equations).
<u>Note</u>: What makes b consistent when X'ε/T ^p→ 0 is that approximating (X'ε/T) by 0 is reasonably accurate in large samples.
Now, we challenge this approximation. We relax the assumption that {x_i, ε_i} is a sequence of independent observations. That is,

plim $(\mathbf{X'}\mathbf{\epsilon}/T) \neq \mathbf{0}$. \Rightarrow This is the IV Problem!

The IV Problem: OLS is Inconsistent

• A correlation between $X \& \varepsilon$ is not rare in economics, especially in corporate finance, where endogeneity is pervasive.

<u>Endogenous in econometrics</u>: A variable is correlated with the error term.

Q: What is the implication of the violation of plim(X'ε/T) = 0? From the asymptotic CLM version, we keep (A1), (A3), and (A4'): (A1) y = Xβ + ε. (A3) Var[ε|X] = σ² I_T (A4') plim (X'X/T) = Q
Now, we assume (A2") plim(X'ε/T) ≠ 0.
Then, plim b = plim β + plim (X'X/T)⁻¹ plim (X'ε/T)

The IV Problem: OLS is Inconsistent

plim \mathbf{b} = plim $\boldsymbol{\beta}$ + plim $(\boldsymbol{X'X}/T)^{-1}$ plim $(\boldsymbol{X'\varepsilon}/T)$ = $\boldsymbol{\beta}$ + \mathbf{Q}^{-1} plim $(\boldsymbol{X'\varepsilon}/T) \neq \boldsymbol{\beta}$

Under the new assumption, **b** is not a consistent estimator of β .

<u>Note</u>: For finite samples, we could have challenged assumption (A2) $E[\boldsymbol{\epsilon} | \boldsymbol{X}] = 0$. Then, $Cov(\boldsymbol{X}, \boldsymbol{\epsilon}) \neq 0 \Rightarrow E[\boldsymbol{b} | \boldsymbol{X}] \neq \boldsymbol{\beta}$.

• Diagram with $\text{Cov}(\boldsymbol{X}, \boldsymbol{\varepsilon}) \neq 0$

$X \rightarrow Y$

IV: Conditions for Instruments

• The solution to the IV problem, in the "traditional IV" literature aims to "cure" the endogeneity problem by finding instruments **Z**, such that the instruments used are both valid and relevant/informative.

• That is, we look for \boldsymbol{Z} such that

(1) $\operatorname{Cov}(\boldsymbol{X}, \boldsymbol{Z}) \neq \boldsymbol{0}$	- relevance condition			
(2) $\operatorname{Cov}(\boldsymbol{Z},\boldsymbol{\varepsilon}) = \boldsymbol{0}$	- valid condition (exclusion restriction)			

• (1) \Rightarrow the Cov(X, Z) should be high enough to produce an \hat{X} (from the first stage) that has relevant informative about X.

(2) In an omitted variables problem, we can think of (2) as broken into two parts: (a) Z is uncorrelated to ε, & (b) only affects y through X



• (2) \Rightarrow **Z** is not only uncorrelated to ε , but only affects **y** through **X** – after all, it is excluded from structural equation!

$$Z \to X \to y$$

From the 2nd part: Once I know the effect of **Z** on **X**, I can throw **Z**.

• Combining assumptions:



<u>Terminology</u>: As seen later, Z will be structured as a binary variable ("treated" or "not-treated") and y as the outcome.



IV Estimators: Properties – Consistency

• Properties of \mathbf{b}_{IV} (1) Consistent $\mathbf{b}_{IV} = (\mathbf{Z}'\mathbf{X})^{-1}\mathbf{Z}' \ \mathbf{y} = (\mathbf{Z}'\mathbf{X})^{-1}\mathbf{Z}' \ (\mathbf{X}\mathbf{\beta} + \mathbf{\epsilon})$ $= (\mathbf{Z}'\mathbf{X}/T)^{-1} \ (\mathbf{Z}'\mathbf{X}/T) \ \mathbf{\beta} + (\mathbf{Z}'\mathbf{X}/T)^{-1}\mathbf{Z}' \ \mathbf{\epsilon}/T$ Under assumptions: $\text{plim}(\mathbf{b}_{IV}) = \mathbf{Q}_{zx}^{-1} \ \mathbf{Q}_{zx} \ \mathbf{\beta} + \mathbf{Q}_{zx}^{-1} \ \text{plim}(\mathbf{Z}' \ \mathbf{\epsilon}/T)$ $= \mathbf{\beta} + \mathbf{Q}_{zx}^{-1} \ \text{plim}(\mathbf{Z}' \ \mathbf{\epsilon}/T) \xrightarrow{p} \mathbf{\beta}$ <u>Note:</u> - Under the context of Lecture 7 -i.e., (A2') $\text{plim}(\mathbf{X}'\mathbf{\epsilon}/T) = \mathbf{0}$ -, **b** is consistent. But, \mathbf{b}_{IV} is also consistent (though, not efficient)!

- Under the context of this Lecture –i.e., (A2') $plim(X'\epsilon/T) \neq 0$ –, *only* the IV estimator is consistent, **b** is not.

IV Estimators: Properties – Asym. Normality • Properties of \mathbf{b}_{IV} (2) Asymptotic normality $\sqrt{T}(\mathbf{b}_{IV} - \mathbf{\beta}) = \sqrt{T}(\mathbf{Z}'\mathbf{X})^{-1}\mathbf{Z}'\mathbf{\epsilon}$ $= (\mathbf{Z}'\mathbf{X}/T)^{-1}\sqrt{T}(\mathbf{Z}'\mathbf{\epsilon}/T)$ Using the Lindberg-Feller CLT $\sqrt{T}(\mathbf{Z}'\mathbf{\epsilon}/T) \xrightarrow{d} N(0, \sigma^2 \mathbf{Q}_{zz})$ Then, $\sqrt{T}(\mathbf{b}_{IV} - \mathbf{\beta}) \xrightarrow{d} N(0, \sigma^2 \mathbf{Q}_{zx}^{-1} \mathbf{Q}_{zz} \mathbf{Q}_{xz}^{-1})$



IV Estimators: Example

Simplest case: Linear model, two endogenous variables, one IV. $y_1 = y_2 \beta + \varepsilon$ $-\varepsilon \sim N(0, \sigma_{\varepsilon\varepsilon})$ $y_2 = z \pi + v$ $-v \sim N(0, \sigma_{VV})$ with reduced form: $y_1 = z \pi \beta + v \beta + \varepsilon = z \gamma + \xi$. The parameter of interest is $\beta (= \gamma/\pi)$. • We estimate β with IV: $b_{IV} = \frac{\frac{1}{T} \sum_{i}^{T} (y_{1,i} - \overline{y_1})(z_i - \overline{z})}{\frac{1}{T} \sum_{i}^{T} (y_{2,i} - \overline{y_2})(z_i - \overline{z})}$

IV and Causal Inference

• Recall that finding a good IV is very difficult. The last three decades have seen a different perspective of IV, emphasizing causal inference, in the context of randomized experiments. The instrument reflects an "experiment," with treated and non-treated units.

• The traditional IV setting recognizes that units (individuals, firms) actively influence or select the level of the treatment they receive. Thus, units receiving the treatment are different from those receiving the "non-treatment." This makes the treatment potentially **endogenous**, creating **selection problems** -Heckman (1979).

Example: Individual i optimally selects her level of education (say, H or L). Individuals choosing higher levels of education may have higher skills, which lead to higher wages (w) for given levels of education.

 \Rightarrow Difficult to compute returns to education.

IV and Causal Inference: RCT

• It would be easier to measure returns to education if the level of education can be seen as randomly assigned.

• Randomized controlled trial (**RCT**) is considered the "gold standard" in experimental research because it eliminates selection problems.

Example: In the medical field, a drug and a placebo may be randomly given to patients, who do not know what they are taking. The researcher in charge is also uninformed regarding the drug-placebo assignment. This is called "**a double blinding**."

Patients are expected to be alike (and behave similarly). If individuals comply with random assignment, this feature allows inference to the average causal effect of the drug on a health outcome for the population in the experiment.

IV and Causal Inference: RCT

• Under some assumptions –i.e., "an ideal experiment"–, an RCT provides an unbiased estimate of the **average treatment effect** (**ATE**). Also, ideally, the results can be generalized to other populations and settings.

• But, poorly designed RCTs can create biased results –i.e., ATE may not be well estimated. Researchers should convinced their audience that the RCT is well-designed:

1) Understanding potential outcomes -i.e., with no "surprises."

2) Clear definition of Hypothesis and Experiment.

3) Understanding of Assignment Mechanism.

4) Define covariates (which ones to control, which ones to avoid).

5) Parameter(s) to be tested.

Causal Inference: RCT

- Basic Concepts/Terminology:
- Unit: A physical object at a particular point in time.
- **Treatment**: The intervention. This creates the dummy variable, Z_i :

 $Z_i = 1$ (*i* treated), $Z_i = 0$ (*i* control).

We want to assess the effect on the units relative to no intervention.

- **Potential outcomes**: The values to be observed of a unit's measurements of interest after treatment and non-treatment. For example, y_i (treatment=1) & y_i (control=0).

- **Causal Effect** or **TE:** For each unit, we compare the potential outcome under both scenarios (treatment & non-treatment):

 $TE_i = y_i$ (treatment=1) - y_i (control=0).

Causal Inference: Treatment Effect (TE)

- **Causal Effect** or **TE** (Treatment effect): For each unit, we compare the potential outcome under treatment & non-treatment:

 $TE_i = y_i$ (treatment=1) - y_i (control=0).

<u>Remark</u>: We define causality in terms of "potential outcomes" for each unit *i*. The parameter of interest ("**estimand**") is a "*within*" quantity.

• <u>Note</u>: Different units may have different effects, that is, TE_i may be heterogeneous.

• Using the definition of Z_i , we decompose observed outcome as: $y_i = y_i(1) * Z_i + y_i(0) * (1 - Z_i)$ $= y_i(0) + [y_i(1) - y_i(0)] * Z_i$ $= \text{baseline} + TE_i * Z_i$

Causal Inference: ATE, ATT & ATU

The **Average Treatment Effect** (**ATE**) is the population average over all *i*:

 $\mathbf{ATE} = \mathbf{E}[\boldsymbol{\delta}] = \mathbf{E}[\boldsymbol{y}(1) - \boldsymbol{y}(0)] = \mathbf{E}[\boldsymbol{y}(1)] - \mathbf{E}[\boldsymbol{y}(0)].$

Similarly, we can also define an ATE on the Treated (ATT) and an ATE on the Untreated (ATU):

ATT = $E[\delta | Z_i = 1] = E[y(1) | Z_i = 1] - E[y(0) | Z_i = 1].$ ATU = $E[\delta | Z_i = 0] = E[y(1) | Z_i = 0] - E[y(0) | Z_i = 0].$

• <u>Fundamental Problem of Causal Inference</u>: We only observe one outcome, either $y_i(1)$ or $y_i(0)$. The "**counterfactuals**," say, $y_i(0|Z_i = 1)$, are always missing. Thus, ATE, ATT & ATU cannot be computed. We can think of this problem as a "**missing data problem**."

Causal Inference: ATE, ATT & ATU

The **Average Treatment Effect** (**ATE**) is the population average over all *i*:

ATE =
$$E[\delta] = E[y(1) - y(0)] = E[y(1)] - E[y(0)].$$

Similarly, we can also define an ATE on the Treated (ATT) and an ATE on the Untreated (ATU):

$$AT^{T} = E[\delta | Z_{i} = 1] = E[y(1) | Z_{i} = 1] - E[y(0) | Z_{i} = 1].$$

$$ATU = E[\delta | Z_{i} = 0] = E[y(1) | Z_{i} = 0] - E[y(0) | Z_{i} = 0].$$

• <u>Fundamental Problem of Causal Inference</u>: We only observe one outcome, either $y_i(1)$ or $y_i(0)$. The "**counterfactuals**," say, $y_i(0|Z_i = 1)$, are always missing. Thus, ATE, ATT & ATU cannot be computed. We can think of this problem as a "**missing data problem**."

Causal Inference: Sample Causal Effects

• We compare (observed) outcomes –in the taxi driver example, $y_i(\text{loan}=1) \& y_i(\text{no loan}=0)$ on a subset of units, the *i*'s. For example,

Average =
$$\frac{1}{N}\sum_{i}^{N}(y_i(1) - y_i(0))$$

• If we have some information about unit *i*, the "**covariates**," then, we can look at the effects by covariates. That is, we allow for **heterogeneous effects**. For example, suppose we have $x_i =$ Education (High school education, and College education), then

Average for College educated =
$$\frac{1}{N_{coll}} \sum_{i}^{N_{coll}} (y_i(1) - y_i(0))$$

Causal Inference: ATE, ATT & ATU - Example

Example: A taxi driver may receive a loan ("treatment") from a bank, with the idea that the driver will buy a car and, thus, increase income ("potential outcome") and generate more business with the bank. If the driver receives a loan, the taxi driver becomes the treated unit. Suppose the bank has the following ideal data:

Unit	No loan	Loan	Diff
1	10	8	-2
2	3	3	0
3	6	10	4
4	9	10	1
5	12	6	-6
6	9	11	2
7	6	4	-2
Average	7.8571	7.4286	-0.375

The bank computes ATE = -0.375. That is, if every taxi drives receives a bank loan, on average, the bank loan reduces income by USD 375.



Causal Inference: ATE, ATT & ATU - Example

• Interpretation:

ATE: It measures the average effect of treatment across all units –i.e., giving a bank loan to every taxi driver.

ATT: It measures the average effect of treatment on the units already treated –i.e., the effect of a loan on the taxi drivers who received the loan. A useful quantity if the bank plans to remove the loan program.

ATU: It measures the average effect of treatment on units untreated – i.e., the effect of a loan on the taxi drivers who have not received the loan. A useful quantity if the bank plans to expand the loan program.

See Greifer and Stuart (2023) for a discussion and applications of these effects.

Causal Inference: SDO

• We observe mean outcomes for the treated (E[$y(1) | Z_i = 1$]) and the untreated E[$y(0) | Z_i = 0$]. After some algebra, we compute **SDO** (Simple difference in mean outcomes) as:

SDO = E[$y(1) | Z_i = 1$] - E[$y(0) | Z_i = 0$] = E[y(1)] - E[y(0)] + E[$y(0) | Z_i = 1$] - E[$y(0) | Z_i = 0$] + + (1 - π) (ATT - ATU)

where π is the share of units in the treatment group.

• Interpretation of terms:

 $- \operatorname{ATE} = \operatorname{E}[\boldsymbol{y}(1)] - \operatorname{E}[\boldsymbol{y}(0)]$

- Selection Bias: $E[y(0) | Z_i = 1] E[y(0) | Z_i = 0]$
- -Heterogeneous Treatment Effect Bias: (1π) (ATT ATU):

Causal Inference: SDO & Selection Bias

• Usually, the selection bias is addressed by:

(1) Modeling it directly and remove it (Heckman).

(2) By design: Experiments, randomization (Rubin).

• We will deal with (1) later on in the semester. In this lecture, we focus on (2).

• Selection bias in causal inference is when one or both mean potential outcome differ by treatment status. The source of the bias is caused by the why people get treated, or the "*treatment assignment mechanism*."

Causal Inference: SUTVA & Assignments

• More Concepts:

- **SUTVA** (Stable Unit Treatment Value Assumption): There is no interference among units –i.e., each unit's potential outcomes are independent of what happens to other units. In addition, for each unit, for each unit there is only one form of treatment and non-treatment.

Example: A taxi driver buys a car with the loan if another driver does not buy a car with the loan. This is a violation of SUTVA.

- Assignment Mechanism: The rules that determine which units receive treatment and which receive control.

Example: A loan officer evaluates a taxi driver personality and gives a loan to the one that she believes will make better use of the loan –i.e., increase income.

Causal Inference: Assignment Mechanism

• To draw conclusions from the causal effects, we need to understand the assignment mechanism of treatment.

Example: Recall the potential outcomes from the loan experiment are (in USD 1,000):

Unit	No loan	Loan	Diff
1	10	8	-2
2	3	3	0
3	6	10	4
4	9	10	1
5	12	6	-6
6	9	11	2
7	6	4	-2
	7.8571	7.4286	-0.375

ATE = $-0.375 \Rightarrow$ If all taxi drivers receives the loan (the treatment), their income would get reduced by USD 375.

Example: S the loan to t	he taxi dri	e loan officer ver that is goi	("Clairvoy ng to use i	vant Loan Officer") t in a positive way:
		Z	No loan	Loan
	Unit	(assignment)	$y_i(0)$	$y_i(1)$
	1	0	10	X
	2	0	Х	х
	3	1	Х	10
	4	1	х	10
	5	0	12	х
	6	1	х	11
	7	0	6	Х
			9.3333	10.3333
The obse r ve	ed average	causal effect i	s 1 > -0.3 '	75.



Causal Inference: Randomization & ATE • Implications for ATE. Let's look at the SDO decomposition: SDO = ATE + Selection Bias + Heterogeneous Treatment Effect Bias • Now, under independence assumption: - Selection Bias: $E[y(0) | Z_i = 1] - E[y(0) | Z_i = 0] = 0$. Heterogeneous Treatment Effect Bias: $(1 - \pi) (ATT - ATU) = 0$. Check last term: $ATT - ATU = (E[y(1) | Z_i = 1] - E[y(0) | Z_i = 1]) - (E[y(1) | Z_i = 0] - E[y(0) | Z_i = 0]) = 0$ Then, with randomized treatment assignment, the SDO is an unbiased estimator of ATE.

Causal Inference: Randomization & Regression

• We can also estimate ATE using a regression.

• We assume that treatment effects are constant: $\delta = y_i(1) - y_i(0)$. Then, after some algebra, we get:

$$y_{i} = y_{i}(1) Z_{i} + (1 - Z_{i}) y_{i}(0)$$

= $y_{i}(0) + (y_{i}(1) - y_{i}(0)) Z_{i}$
= $y_{i}(0) + \delta Z_{i}$
= $y_{i}(0) + \delta Z_{i} + E[y_{i}(0)] - E[y_{i}(0)]$
= $E[y_{i}(0)] + \delta Z_{i} + \varepsilon_{i}$
= $\alpha + \delta Z_{i} + \varepsilon_{i}$

where ε_i is the error component of $y_i(0)$. This is what we presented last semester, for example in Card's (1990) paper. This framework allows easy incorporation of covariates ("control variables").



Causal Inference: Confounders

• So far, we have only paid attention to y (outcome) & Z (treatment), implicitly, in a directional way:

 $Z \rightarrow y$.

• Now, we incorporate information about unit i, the covariates, x_i , into our analysis. The relevant covariates are the "**confounders**," the i variables that affect both Z & y. In graphic form:



• The confounders are a source of selection bias. We will use *X* to get "**stratified**" conditional estimates, say, by gender or income brackets, and then, take a weighted average.

Causal Inference: Matching Matching: A treated unit is "matched" to a comparison unit that is identical on known and quantified relevant confounders. Example: A taxi driver that receives a loan will be matched with another that didn't receive a loan. Matching is based on the *k* relevant covariates: age, experience, income, family members, and city. Question: If we can use a regression to estimate ATE, why matching? Two main reasons: Matching does not assume a functional form. If functional form is incorrect, the estimates are inconsistent. Matching only uses the untreated units "similar" (in covariates) to the treated. In a regression, all the untreated unit.

Causal Inference: Exact Matching

• If we can find an exact counterpart, we have **exact matching**, otherwise, we have **inexact matching**.

• Suppose we can do **exact matching**, that is, in the previous example, we find a non-treated taxi driver with exactly the same covariates as treated taxi driver *i*. The matched non-treated driver has outcome $y_{j(i)}(0)$.

Under exact matching we can estimate the treatment effect for *i*:

$$\hat{\delta_i} = y_i(1) - y_{j(i)}(0)$$

• Now, by matching all treated units in the experiment, we can compute the ATT.

Causal Inference: Exact Matching

• We do the exact matching for every treated taxi driver in the sample (suppose we have N_T treated taxi drivers). Then, the ATT estimate:

$$\widehat{ATT} = \frac{1}{N_T} \sum_{i}^{N_T} (y_i(1) - y_{j(i)}(0))$$

where $y_{j(i)}$ is the outcome of the *j*-th unit matched to the *i*-th unit based on exact matching.

• If we find more than one unit with identical covariates, say M matching unites, then we use an average of the matching units to estimate ATT:

$$\widehat{ATT} = \frac{1}{N_T} \sum_{i}^{N_T} (y_i(1) - \left[\frac{1}{M} \sum_{i}^{M} y_{jm(i)}(0)\right])$$

<u>Note</u>: When matching we do not use all the data, only the relevant data that allows us to estimate ATT.

Causal Inference: Exact Matching Issues

Continuous covariates

In practice, exact matching with continuous covariates, X_i , is almost impossible. We have to, somewhat, discretize the data. Usually, we rely on ranges of the X_i . We call this process *coarsening* or *binning* of X_i .

Example: We match on income if between \$25,000 and \$27,500.

• When we do binning, the matching is called **coarsened exact** matching (**CEM**).

Common support

To do exact matching, we need **common support**, that is, for each value of X_i , there exist an observation in the treated group and in the control group. If there is no common support, there are units without matches, and ATE cannot be estimated.

Causal Inference: Inexact Matching

• It is not easy to find exact matches for all treated units. One solution is to use the "**closest**" value among all of the control observations.

• We can use the Euclidean distance or the Manhattan distance to define closeness. It is common to normalized distances to make the metric invariant to the units of the covariates. For example, the Mahalanobis distance:

$$||\mathbf{X}_i - \mathbf{X}_j|| = sqrt\left((\mathbf{X}_i - \mathbf{X}_j)'\widehat{\mathbf{V}}^{-1}(\mathbf{X}_i - \mathbf{X}_j)\right)$$

where \hat{V} is the variance-covariance matrix of **X**.

• As *k* increases, we run into the "**curse of dimensionality**." Suppose k = 6 and each X_k is binary, then, we have $64 (=2^6)$ possible combinations to match. Not easy. Usual advice, be greedy. Match only using the variables that make sense to control for, no more, no less.

Causal Inference: Matching Bias

• Matching biases will appear in the estimates, though asymptotically, matching estimators tend to be consistent.

• A Matching bias arises because of the effect of matching discrepancies between $\mu^{0}(X_{i})$ and $\mu^{0}(X_{j(i)})$. Lack of **common support** increases as *k* increases.

• We can minimize matching discrepancies by:

- Using a small M (say, M = 1). Larger values of M produce large matching discrepancies.

- Matching with replacement. Because matching with replacement can use untreated units as a match more than once, matching with replacement produces smaller matching discrepancies.

- Trying to match better covariates with a large effect on μ^0 .

Causal Inference: Matching Bias & Corrections

• We know the data. Then, we observe matching discrepancies. If discrepancies are large, it is always possible to use bias correction techniques. Abadie and Imbens (2011) propose:

$$\widehat{\delta_{ATT}}^{BC} = \frac{1}{N_T} \sum_{i}^{N_T} [(y_i(1) - y_{j(i)}(0) - (\widehat{\mu^0}(X_i) - \widehat{\mu^0}(X_{j(i)}))]$$

where $\widehat{\mu^0}(X_i)$ is an estimate of $\mathbb{E}[y(0)|X = x, Z_i = 0]$, say, using OLS.

• For example, under unconfoundedness:

 $E[y(0) | X = x, Z_i = 1] = E[y(0) | X = x, Z_i = 0].$

That is, the relationship between y(0) & X in the control group is the same in the treatment group, even though we do not observe y(0) in the treatment group.

• Then, fit a regression of y(0) against X in the Z = 0 group only to predict y(0) in both groups.

Causal Inference: Matching & Propensity Score

• Rosenbaum and Rubin (1983) propose an alternative to the distance methods: Use a "**propensity score (PS)**" to summarize all the covariates into a single scalar and then weight or match on it.

Definition: Propensity score

A unit specific probability bounded between 0 and 1 measuring the probability of treatment assignment given a set of covariates:

p(X) = P(Z = 1 | X).

• Think of PS as a useful dimension reduction tool. The definition is not specific about how to compute p(X) (logit? Non-parametric?) or what variables are in X. In practice, X should only include variables that affect both treated & non-treated (*confounders*).



Causal Inference: Matching & Propensity Score

• There are many tests of the balancing property. Usually, we look at bins of our data with similar PS and check if there are significant differences in their covariates.

• We still need common support, which should be easier to achieve than in the case with several covariates.

• Usually, we check common support by looking for an overlap of the histograms of PS for treated and non-treated units. For ATE, the distributions should completely overlap. For ATT, we only need a partial overlap (the treated distribution should be a subset of the non-treated distribution).

Causal Inference: Matching & Propensity Score

• A popular use of PS is to use the inverse to weight observations. This is called **inverse probability of treatment weighting (IPTW**).

• With IPTW, we give a higher weight to the units treated that had a low PS –i.e., a priori, a unit unlikely to be treated – and, viceversa, a lower weight to the units with higher PS.

• To compute ATE, we compute the difference of the weighted groups, to compute ATT, we only weight the untreated group.

Causal Inference: Matching - Remarks

• Matching is a useful tool that avoids the usual assumptions behind regression models, for example, a functional form.

• Matching requires that the researcher observes and uses all the variables that affect both participation and outcomes.

• Matching can be interpreted as reorganizing the data from an observational study in such a way that the assumptions from a randomized experiment hold, at least approximately.

• Matching may be inexact, systematic differences in pre-exposure variables across the matched pairs may remain but can be subsequently be adjusted in the analysis stage.

Natural Experiments

• A **natural experiment** is an event or a situation, that is not under the control of the units under study, which generates variation in the variable of interest that is as if it had been randomly assigned.

• That is, a natural experiment is exogenous to a structural model. Like the previous external examples, the exclusion condition is met.

Examples: Weather shocks, policy changes, administrative rules, immigration shocks, birth dates, birth weight, etc.

• "Natural" points out that us (the researchers) did not design the episode to be analyzed, but can use it to identify causal relationships...

Natural Experiments & IV Estimation

• Back to the IV framework:

The key is to find a natural experiment (defining Z) that is correlated with X and has no direct effect on y –the impact on y is through X.

Z is an exogenous event \Rightarrow resulting values of **X** induced by **Z** may be considered **randomized**, like in an RCT.

<u>Key feature</u>: **Randomization of treatment**. We need to show that the two groups are comparable along all dimensions relevant for the outcome variable (age, gender, previous health, etc.) except treatment.

Natural Experiments: RCT Substitution

• Recall that in the CEO compensation model, we want to test the causal effect of networking on compensation, but an omitted variable – the CEO's unobserved skills– creates endogeneity.

A solution to the omitted variables problem is to assign networking (*x*) randomly: we have two similar groups of CEOs (with similar skills!) & randomly assign them values (say, large network & small network).

• Of course, this randomized experiment is not possible.

• But, suppose we have a natural event, **Z**, unrelated to CEO compensation, which randomly assigns networking, **x**, to two groups (immigration? lack of U.S. education?). Then, we can test causality, without the omitted variable problem.

Natural Experiments: Steps

• Steps of a natural experiment:

(1) Experiment defines an IV: $Z_i = 1$ (*i* treated),

 $Z_i = 0$ (*i* control).

(2) Identify two groups:

- treated (all i with $Z_i = 1$) with observations: $\mathbf{y}(1), \mathbf{X}(1)$

- control (all *i* with $Z_i = 0$) with observations: $\mathbf{y}(0), \mathbf{X}(0)$

(3) We analyze differences between $(\mathbf{y}(1), \mathbf{X}(1)) \& (\mathbf{y}(0), \mathbf{X}(0))$.

<u>Note</u>: Like in RCT, it is impossible to observe $y_i(1) \& y_i(0)$ for the same individual *i*. Thus, we focus on the ATE between groups, which under some assumptions can be estimated by SDO. Recall:

 $ATE = E[\boldsymbol{y}(1)] - E[\boldsymbol{y}(0)].$

Natural Experiments: Steps

• <u>Remarks</u>: Steps 1-3 will be treated like a lab experiment if we show that the treatment is in fact randomly assigned. We need to show that two groups are comparable except for the treatment.

• This is the key for the experiment to be valid. We need to convince the audience that the we have a *quasi-random* treatment.

• Heterogeneity. With heterogeneous treatment effects, ATE may vary with different groups: men/women, immigrants/non-immigrants, etc.

• Typical problem: Selection bias

- Individual *i* select treatment or not.

- Treatment is assigned to the individual i with the highest chance of being successful (assignment is not independent of potential outcome).

Natural Experiments: Selection Bias

• Treatment should affect the outcome. If treated individual *i* would have achieved desired outcome without treatment, we do not correctly estimate ATE.

• The problem? Treatment is not independent of outcome. That is,

 $E[\mathbf{y}(0) | Z_i = 1] \neq E[\mathbf{y}(0) | Z_i = 0].$

Example: Companies pay to show ads after an individual conducts an internet search for a particular product. Consumers who click on the ads are likely informed and with high likelihood would have found the product independent of the ad.

Blake, Nosko and Tadelis (Econometrica, 2015) find that nonexperimental measures of returns on paid search ads are huge (over 1600%), but experimental measures of returns are close to zero.

Natural Experiments: Heterogeneity

Example: Angrist (1990) use the Vietnam-era draft lottery, a randomized draw of birth dates, to estimate the effect of military service on earnings later in life.

The draft (natural experiment) is the instrument for military service. But, most of the individuals who served in Vietnam were volunteers. Thus,, the draft lottery only affected individuals who would not have served voluntarily in the military.

Implication: ATE is likely not representative of those who volunteered for service in the Vietnam War.

• In most situations, responses are **heterogeneous**. When treatment effects vary across individuals and people make choices, there is likely to be **incomplete compliance** with the (natural) experiment.

Natural Experiments: Compliance

A&K (1991) found a causal return to schooling that is between 8% and 10%, a higher return than the OLS implied return. But, this is not the end of the story.

• The quasi-experimental variation mainly affected those with a high probability of dropping out of school as soon as possible. It may well be that the returns to schooling in this part of the population are not representative of the overall population.

• In the language of controlled experiments, those who were unaffected by the natural experiment are "**non-compliers**," and their returns to schooling are potentially different than among the "**compliers**," because of heterogeneous treatment effects..

Natural Experiments: Compliance

• Compliance terminology:

- **Compliers**: the units whose behaviors are affected by instrument as expected –i.e., late birth data \Rightarrow finish HS.

- **Defiers**: the units whose behaviors are affected by instrument in an unexpected way –i.e., late birth data \Rightarrow HS drop out.

- Always takers and Never takers: units whose behaviors are not affected by instrument (finish/drop-out HS regardless of date of birth).

Example: In Angrist (1990), always-takers serve in the military no matter the lottery number (with, likely, no direct effect on earnings). Never-takers do not serve in the military no matter their lottery number. Very likely, there are no defiers.

Natural Experiments: Identifying ATE

• The combination of treatment heterogeneity and incomplete compliance poses a problem for causal analysis. Under these elements, ATE identification is complicated:

- Heterogeneity: treatment effects vary by unit.

- Incomplete compliance: *treatment status* \neq *treatment eligibility*.

• In general, treatment status depends on individual treatment effects. Thus, additional assumptions are needed to identify ATE:

- Units do not know effect of treatment Heckman (1997).

- One-sided non-compliance –i.e., the probability of participating is zero for individuals who are not eligible for treatment Bloom (1984).

• We think of these assumptions as strong assumptions.

Identifying Causal Effects

• Angrist and Imbens (1994) set a general framework to identify causal effects. The key is the link of the assignment mechanism to an instrument (defined by RCT or natural experiment).

• We start with a binary treatment. In A&K (1991), y_i denotes earnings, with treatment indicator, d_i , denoting finishing high school (=1) or not (=0). The instrument, Z_i , is the birth date (Z_i =1 if born July-Dec, Z_i =0 if born Jan-Jun).

• Since the treatment indicator (completing high school or not) is endogenous, we think of it in terms of potential outcomes, $y_i(d_i, Z_i)$, which are four.

Identifying Causal Effects: Assumptions

• A good instrument, *Z*, should satisfy: (1) Randomly assigned: $[\{y_i(d_i, Z_i) \lor d, z\}, d_i(1), d_i(0)] \perp Z_i$ In the A&K example, earnings are independent of the dates on which individuals are born. Under this assumption, the regression of *Y* on *Z* (reduced form) identifies the causal effect of instrument on outcome: $E[y_i|Z_i=1] - E[y_i|Z_i=0] = E[y_i(1, d_i(1))|Z_i=1] - E[y_i(0, d_i(0))|Z_i=0]$ $= E[y_i(1, d_i(1))] - E[y_i(0, d_i(0))]$ = Causal effect of being born late on earnings.(last step follows from random assignment). Similarly, $E[d_i|Z_i=1] - E[d_i|Z_i=0] = E[d_i(1)|Z_i=1] - [d_i(0)|Z_i=0]$

$$= E[d_i(1)] - [d_i(0)]$$

= Causal effect of being born late on
likelihood of finishing HS

Identifying Causal Effects: Assumptions

(2) Relevant: $E[d_i(1)] - E[d_i(0)] \neq 0$

We require a stronger response to the possibility of dropping out of high school among those who are born early relative to those born late.

(3) Exclusion: $y_i(1, d_i) = y_i(0, d_i) = y_i(d_i)$

We need this assumption since $Z \neq d$. The only way the instrument (birth date) affects the outcome is through the treatment.

Example: In Angrist's (1990) draft lottery paper, the exclusion restriction requires that potential earnings with and without military service be independent of the lottery number.

• Now, we can identify the effect of the treatment (finishing HS) on the outcome (earnings).

Identifying Causal Effects: Assumptions

• Notice that $\{d_i(1) - d_i(0)\}$ can take three values: 1, 0, -1. Then, $E[y_i|Z_i=1] - E[y_i|Z_i=0] =$ $= \sum_k kP[d_i(1) - d_i(0) = k] * E[y_i(1) - y_i(0) | \{d_i(1) - d_i(0)\} = k]$ $= P[(d_i(1) - d_i(0) = 1] * E[y_i(1) - y_i(0) | \{d_i(1) - d_i(0)\} = 1]$ $- P[(d_i(1) - d_i(0) = -1] * E[y_i(1) - y_i(0) | \{d_i(1) - d_i(0)\} = -1]$ \Rightarrow units who do not respond to the instrument -with $d_i(1) = d_i(0)$ -

 \Rightarrow units who do not respond to the instrument –with $d_i(1) = d_i(0)$ – do not contribute to identification. Impossible to estimate a causal effect for the units that do not change behavior.

<u>Note</u>: The reduced form can be negative! Why? The treatment effect for those who shift from nonparticipation to participation (or from "0" to "1") when Z is switched from 0 to 1 ("*compliers*") can be cancelled out by the treatment effect of those who shift from 1 to 0 ("*defiers*").

 \Rightarrow we need more assumptions.

Identifying Causal Effects: Monotonicity

(4) Monotonicity: $d_i(1) - d_i(0) \ge 0$ (or ≤ 0)

Monotonicity requires that the instrument operates in the same direction on all individual units. Anyone affected by the instrument is affected *in the same direction* (positively or negatively, but not both).

Examples: In A&K, we have 4 cases:

- Compliers:	$d_i(1) = 1, \ d_i(0) = 0 \Rightarrow d_i(1) > d_i(0)$	$_{i}(0)$
--------------	--	-----------

Defiers:	$d_i(1)=0,$	$d_i(0) = 1 \Rightarrow$	$d_i(1) < d_i(1)$	$d_{i}(0)$
Defiers:	$a_i(1)=0,$	$a_i(0) = 1 \Rightarrow$	$a_i(1) < $	$a_i(0)$

- Always takers: $d_i(1) = 1$, $d_i(0) = 1 \Rightarrow d_i(1) = d_i(0)$
- Never takers: $d_i(1) = 0, \ d_i(0) = 0 \Rightarrow d_i(1) = d_i(0)$

Monotonicity requires no defiers. Recall that always takers and nevertakers play no role in the computation of $E[y_i|Z_i=1] - E[y_i|Z_i=0]$.

Identifying Causal Effects: Monotonicity

Examples (continuation): In Angrist (1990) monotonicity requires that someone who would serve in the military with lottery number m would also serve in the military with lottery number $n \ (m > n)$, which is plausible.

Dobbie et al. (2018) used the detention tendencies of (quasi)-randomly assigned bail judges to estimate the causal effect of pretrial detention on subsequent outcomes. Monotonicity "requires that individuals released by a strict judge would also be released by a more lenient judge, and that individuals detained by a lenient judge would also be detained by a stricter judge," which may not be true.

• Monotonicity implies $P[(d_i(1) - d_i(0)) = -1] = 0$. Then, the causal effect of instrument on treatment equals:

 $E[d_i(1) - d_i(0)] = P[(d_i(1) - d_i(0)) = 1]$

Identifying Causal Effects: IVE & ITT • With assumptions (1)-(4), the causal effect of instrument on outcome (reduced form estimation, also called intention-to-treat or ITT): $E[y_i | Z_i = 1] - E[y_i | Z_i = 0] =$ $= P[(d_i(1) - d_i(0) = 1] * E[y_i(1) - y_i(0) | \{d_i(1) - d_i(0)\} = 1]$ \Rightarrow LHS is the causal effect of Z on y. First component of the RHS is the causal effect of Z on d (treatment).

• The ratio is the IV estimator:

 $\frac{E[y_i|Z_i=1] - E[y_i|Z_i=0]}{P[(d_i(1) - d_i(0) = 1]} = E[y_i(1) - y_i(0)|\{d_i(1) - d_i(0)\} = 1]$

This represents the average causal effect for the population that changed their treatment status in accordance with the change in the value of the instrument. This is an ATE for a "*local*" group: compliers.

Identifying Causal Effects: LATE

• Angrist and Imbens (1994) called this effect **local average treatment effect** (**LATE**). Since this estimator involves compliers, LATE is also called the **complier average causal effect** (**CACE**).

Example: In Angrist (1990), IV estimates the average effect of military service on earnings for the subpopulation who enrolled in military service because of the draft but would not have served otherwise.

LATE does not tell us what the causal effect of military service was for volunteers or those who were exempted from military service for medical reasons.

<u>Remark</u>: IV estimates the average causal effect for those units affected by the instrument (i.e., complier causal effects only).

 \Rightarrow LATE is not ATE.

Identifying Causal Effects: LATE

• For never-takers, the exclusion restriction might be problematic: going to College or moving abroad can have a direct impact on earnings.

• The existence of defiers in the sample biases the results. The bias can be big if we have a significant proportion of defiers. In Angrist (1990), very likely there are very few defiers.

Causal Effects & Experiments: Inference Issues

Clustering

Following Moulton (1986), researchers clusters (or "grouping") structures in the data to analyze results. Recall that if units in the same group are exposed to the same variation, one should take the correlation across individuals within group into account.

This is a big concern for Difference-in-differences studies, which use variation across groups over time for identification. See Bertrand, et al. (2004), Donald and Lang (2007), and Hansen (2007).

• Weak Instruments

IV studies rely on strong instruments (Nelson and Startz 1990; Staiger and Stock 1997). But, as seen last semester, it is not easy to find them. If the instruments are weak, IV estimates can be severely biased. See Andrews et al. (2019), Keane and Neal (2021), and Young (2020).

Causal Effects & Experiments: Inference Issues

Heteroscedasticity

Recall that Stock, Wright and Yogo (2002) suggest that an F-stat > 10 from the 1st-stage regression indicates that the instruments are not weak. This F-stat is called *non-robust F*, since it is not robust in the presence of heteroscedasticity.

Kleibergen and Paap (2006) use a "*robust*" F-stat, based on clustering SE and/or HAC SE.

Andrews, et al. (2019) suggest the Olea and Pflueger (2013) "*effective*" 1st-stage F-stat, which is equal to the *non-robust* F times a correction factor for non-homoscedasticity. With this effective F-stat, the rule of thumb of "effective F-stat" > 10 applies.

Causal Effects & Experiments: Inference Issues

• Publication bias (also called p-hacking)

Card and Krueger (1995) conducted a meta-analysis of the prior literature on the minimum wage and concluded that it suffered from publication bias.

• More recently, Brodeur, Lé, Sangnier, and Zylberberg (2016) suggested that there is publication bias because there is excess mass of estimates having a p-value just below 0.05 than just above 0.05.

• Brodeur, Cook, and Heyes (2020) focused in particular on the methods associated with the design-based approach. They concluded that p-hacking is more common for DiD and IV methods than for RCTs and RD designs. The last two methods are more "structured" and, thus, more difficult to "cheat."

Identifying Causal Effects: Generalizations

• We know our data. We can use it to gauge heterogeneity and analyze if there is a biased LATE. In many situations, heterogeneity may work in our favor. For example, if the policy-target is low-income individuals, and low-income individuals are overrepresented among compliers, then LATE should be highly relevant for policy purposes.

• With more than one instrument, we can test the heterogeneous responses to the instrument (Angrist, Lavy, and Schlosser, 2010). If heterogeneity is not substantial, LATE estimates may generalize.

• Alternatively, we can fill in the gaps in the data using auxiliary assumptions. For example, Heckman et al. (2001, 2003) and Angrist (2004) used parametric latent-index models to identify causal effects. Brinch et al. (2017) used a structural model. Chamberlain (2010) used Bayesian techniques.

Identifying Causal Effects: MTE

• To evaluate policy effects, many studies move from estimating LATE to "marginal treatment effects" (MTE).

• The literature started with Heckman and Vycatil (1999, 2001, 2005). The MTE is the expected effect of treatment conditional on covariates (observed and unobserved) at a margin.

• We compute MTE in the context of a structural (switching) model, with individuals shifting into (or out of) treatment by a *marginal change* in the instrument (the cost of the treatment, propensity score, etc.).

<u>Note</u>: MTE can be done only with a continuous instrument. Remember, MTE is a "marginal effect," we need a derivative! Thus, we model y_i as function of covariates, X_i , and the instrument, Z_i .

Identifying Causal Effects: MTE

• MTE: Gain from treatment for people who were shifted into (or out of) treatment status by a marginal change in the treatment

• When $y_i(1) \& y_i(0)$ are value outcomes, MTE is a willingness to pay measure –see Heckman and Vycatil (2005).

• Under some assumptions (conditional independence, monotonicity, additive separability of effects, etc.), ATE, ATT & ATU can be recovered using weighted averages over the MTE curve. For example, ATE is the equally weighted average over the entire MTE curve.

• Cornelissen et al. (2016) provides a review of the literature.

Causal Effects & Experiments: Criticism

• Deaton (2010) and Heckman and Urzua (2010) see an excessive (and inappropriate?) use of experimental methods in empirical work. Both focus their criticism on LATE. Heckman and Urzua (2010) say:

"Problems of identification and interpretation are swept under the rug and replaced by 'an effect' identified by IV that is often very difficult to interpret as an answer to an interesting economic question."

• Experimental studies emphasize "too much" credible identification, or on "internal validity as opposed to external validity.

• Imbens (2010) argues that it is useful to separate the assumptions needed to identify a causal effect in the sample studied from the assumptions needed to generalize an internally valid estimate to other populations. (See also DiNardo and Lee (2011).

Causal Effects & Experiments: Criticism

• In fact, quasi-experimental variation and structural models may usefully complement one another.

• For example, Card and Hyslop (2005) used experimental variation to aid the identification of a structural model for welfare participation.

• Design-based estimates can also be used to validate structural models, as done by Blundell (2013). Related to this, Kline and Walters (2019) showed that IV and selection-correction type of estimates (Heckman (1979)) of LATE are numerically equivalent.

• Under some conditions, the choice between these two estimators is unimportant for estimating treatment effects that are identified in the data.

Causal Effects: Remarks

• RCTs are not easily conducted in finance and economics. We usually rely with observational data outside of our control.

• With a natural experiment, treatment or instrument assignment is as good as random, almost like an RCT.

• We would like to compute ATE = E[y(1)] - E[y(0)]. But, the effect identified is **LATE**, the average causal effect among compliers:

The causal effect for the subset of the population that changed behavior because of the value of the instrument.

• Angrist and Imbens (1994) set a general framework where LATE is identified. This framework has become the dominant one for both quasi-experimental and experimental work.

Causal Effects: Other Methods

• There are other methods for causal inference exploiting a similar framework: we place units into a treated group and non-treated group:

- **Difference-in-differences** (DiD) design (Roth et al. (2023)). Similar to design framework, but we do not assume that units are the same.

- **Regression discontinuity design** (RD) (Hahn et al. (2001)). RD takes advantage of a jump in the likelihood of being treated, generated by an arbitrary threshold.

Regression kink design (RKD) (Nielsen et al. (2010) and Card, Lee, Pei, and Weber (2015)). RKD takes advantage of a change in the slope at the likelihood of being treated at a (kink) point.

Synthetic control method (Abadie (2021)). We replicate unobserved outcomes, $y_i(0)$, using observed outcomes from *donor* units $y_{j(i)}(0)$.

Causal Effects: Difference-in-Difference (DiD)

• **Difference-in-differences** (DiD) design (de Chaisemartin and D'Haultfoeuille (2020), Roth et al. (2023)). The canonical DiD model assumes two groups (treated and non-treated) and two periods (before & after treatment, or t=1 & t=2).

• The key identifying assumption is **parallel trends (D1)**: In the absence of treatment the unobserved differences between treatment & control groups are the same over the two periods:

$$\mathbf{E}[y_{i,2}(0) - y_{i,1}(0) \mid Z_i = 1] = \mathbf{E}[y_{i,2}(0) - y_{i,1}(0) \mid Z_i = 0].$$

• With an additional assumption, no anticipation effects (D2): –i.e., $y_{i,1}(0) = y_{i,1}(0)$ for all *i* with $Z_i = 1$ –, we can estimate ATT (δ):

$$\delta = \mathbb{E}[y_{i,2} - y_{i,1} | Z_i = 1] - \mathbb{E}[y_{i,2} - y_{i,1} | Z_i = 0],$$

using the observed sample mean of $y_{i,2} \& y_{i,1}$.

Causal Effects: Difference-in-Difference (DiD)

• We can also estimate δ with a regression, including covariates, x_i :

 $\Delta y_{i,t} = y_{i,2} - y_{i,1} = \delta_0 + \frac{\delta}{\delta} Z_i + (x_{i,2} - x_{i,1})' \beta + \varepsilon_{i,t}$

Note: This is a CLM: δ is consistent and asymptotically normal distributed (with *N* large & *T* fixed). We use Clustered SE for inference (the usual problems, discussed last semester, apply).

• Or, using levels and including individual, α_i , & time effects, $\theta_t = \gamma * Post_i$, (two-way FEM or **TWFE**), we estimate δ with:

 $y_{i,t} = \alpha_i + \gamma * Post_i + \delta Z_i * Post_i + x_{i,t}'\beta + \varepsilon_{i,t}$, where $Post_i$ is a t=2 dummy variable (=1 if t=2).

• We covered this topic (and the next example) last semester.

Causal Effects: Difference-in-Difference (DiD)

Example: Card (1990) studies the effect of a supply labor shock (the Mariel Liftboat influx of immigrants) in Miami, where we observe outcome $y_i(1)$. Card compares $y_{i=Miami}$ with observed outcomes, y_j , in four similar markets (j = Atlanta, Houston, LA & Tampa).

• Several extensions and issues:

- **Multiple periods** (**Staggered treatment**): Units receive treatment at different times. We have different treatment cohorts over time.

- Variation in treatment timing. There are heterogeneous effects over time across the cohort receiving treatment.

- Non-parallel trends: Assumption D1 is violated.

- **Small treated clusters**. As discussed last semester, the CLT may be difficult to invoke.

Causal Effects: DiD – Staggered Treatment

- Multiple periods (staggered treatment). Now, taxi drivers, taxi drivers can receive the bank loan at different times. The (static) TWFE works well when there is not heterogeneity in treatment effects across either time (this case) or units (next case). It turns out the TWFE estimator is a weighted average of the heterogeneous time effects (δ_s): $\delta = \sum_{s=1}^{T} w_s \delta_s$

where the weights can be negative, which can lead to a negative δ !

If we assume that the dynamic effect of the bank loan expansion after *r* years is the same (on average) regardless of what year the taxi driver received the loan, a **dynamic TWFE** can identify ATT (with adaptations of assumptions **D1-D2**):

 $y_{i,t} = \alpha_i + \theta_t + \delta \sum_{r \neq 0} I[R_{i,t} = r] + x_{i,t}'\beta + \varepsilon_{i,t}, \neq$ where $R_{i,t} = t - G_i + 1$ -+is the time relative to treatment and G_i is the time of first treatment. (See Borusyak and Jaravel (2018)).

Causal Effects: DiD - Treatment Timing

- Variation in treatment timing. We have heterogeneous effects over time across the cohort receiving treatment. For example, the effect of the bank loan on the cohort that was treated –i.e., received a loan– at t may be different from the cohort treated at t'. The effect varies with the period treatment was received. Dynamic TWFE can produce very biased results.

Callaway and SantAnna (2021) propose a weighted average of ATT per period and Borusyak, Jaravel and Spiess (2021) propose an estimation using units that where never treated over time.

For example, BJS use averages over the untreated period (0 - t) to estimate $y_i(0)$ for i= treated (T) & never-treated (NT):

 $\hat{\delta}_{BJS} = (\hat{y}_{T,t}(1) - \hat{y}_{T,0-t}(0)) - (\hat{y}_{NT,t}(0) - \hat{y}_{NT,0-t}(0))$ using the observed sample means.

Causal Effects: DiD - Non-parallel Trends

- **Non-parallel trends**: Assumption **D1** can fail if the confounding factors are time-varying. For example, in the bank loan example, experience may have a time-varying effect on outcomes. If we condition on confounders, we can recover the parallel trend:

 $\mathbf{E}[y_{i,2}(0) - y_{i,1}(0) | Z_i = 1, X_i] = \mathbf{E}[y_{i,2}(0) - y_{i,1}(0) | Z_i = 0, X_i].$

With this assumption, along a "strong overlap" one, we estimate ATT:

$$\delta_x = \mathbb{E}[y_{i,2} - y_{i,1} | Z_i = 1, X_i = x] - \mathbb{E}[y_{i,2} - y_{i,1} | Z_i = 0, X_i = x]$$

The unconditional ATT can then be identified by averaging δ_x over the distribution of X_i in the treated population.

• Abadie (2005) proposes a two-step strategy to deal with this issue:

1) Estimate the PS based on observed covariates and compute PS

2) Estimate ATT, using IPW.

Causal Effects: DiD – Small Clusters

- **Small treated clusters**. As discussed last semester, it is a difficult theoretical problem, the CLT does not apply. Model-based solutions (Donald and Lang (2007) and bootstrapping have been proposed (Canay et al. (2021)).

• Model-based solutions: We start with a structural equation:

 $y_{i,j,t} = \alpha_j + \theta_t + \delta Z_{i,j} * Post_i + v_{j,t} + \varepsilon_{i,t},$

where *j* is a cluster and $v_{j,t}$ is the cluster-level error term. With few clusters, the averages of the cluster level shocks $\Delta v_{j,t}$ among treated & untreated clusters will tend not to be approximately by a normal.

We need assumptions: For example, the $v_{j,t}$'s are normally distributed (Donald and Lang (2007)), or that treated and non-treated cluster (a larger cluster) follow the same distribution (Conley and Taber (2011) and Ferman and Pinto (2019), allowing for hereteroscedasticity).

Causal Effects: Regression Discontinuity (RD)

- **Regression discontinuity design** (RD) (Hahn, Todd, and van der Klaauw, 2001). There is a jump in the likelihood of being treated, generated by an arbitrary threshold. There are two practical cases:

- **Sharp RD**: There is a deterministic rule that sets the jump into treatment.

- **Fuzzy RD**: There is an encouragement (incentives) to get treatment. We can think of the probability of treatment increasing at threshold.

Example: In the U.S., to be a bank regulated as a *systemic risk bank* you need to have more than a certain amount of assets (USD 50B). This creates a Sharp RD. In many states, some "small" business have access to subsidized loans. This creates a Fuzzy RD.

In some situations, a Sharp RD can create endogeneity issues: a bank can manage the amount of assets not to be treated (regulated).

Causal Effects: Regression Discontinuity (RD)

In the Sharp RD case, we estimate E[y_i(1) - y_i(0)| x_i = c], where *c* is the arbitrary cutoff point for receiving treatment. We compare units very close (& on both sides) of *c*. We compute the estimand by:
1) Two potential regressions:
(Treated) E[y_i(1) | x_i = c] = lim_{x↓c} E[y_i(1) | x_i = x]
(Control) E[y_i(0) | x_i = c] = lim_{x↓c} E[y_i(0) | x_i = x]
2) A linear regression (using a window around *c*): y_i = α + β (x_i - c) + ε_i
the bandwith (x_i - c) can be selected in some optimal way, say, by minimizing the MSE, as suggested by Imbens and Kalyanaraman (2012). Non-linear functional forms are also OK.

Causal Effects: Regression Kink Design (RKD)

- **Regression kink design** (RKD) (Nielsen et al. (2010) and Card, Lee, Pei, and Weber, 2015). Similar to RD, RKD takes advantage of a change in the slope at the likelihood of being treated at a (kink) point, resulting in a discontinuity in the first-derivative of the assignment function.

If individuals on either side of the kink threshold are "similar," any kink in the outcome can be attributed to the effect of treatment.

Example: In the usual government unemployment insurance scheme run by states, insurance payments cover up to some percentage of the previous wage, up to a maximum amount (sometimes with a floor).

Causal Effects: Synthetic Control Method

- Synthetic control method (Abadie and Gardeazabal (2003) and Abadie, Diamond, and Hainemuller, 2010). We replicate unobserved outcomes, usually $y_i(0)$, using weighted observed outcomes, usually from similar units, or "donor units," $y_{j(i)}(0)$:

$$\hat{y}_{i=N}(0) = \sum_{j=1}^{N-1} w_i y_{j(i)}(0)$$

where we have N - 1 donor units (N is the total sample).

Ideally, we pick the non-negative weights optimally. For example:

$$w_j = \underset{w_j}{\operatorname{argmin}} \sum_{i=1}^{T} \{ y_i(1) - \sum_{j=1}^{N-1} w_j y_{j(i)}(0) \}^2.$$

Note: the argmin is $\hat{\delta}_{SC}$! Also note the weights can be computed using a constrained regression ($w_i \ge 0, \sum_{i=1}^{N-1} w_i = 1$, and no constant).

• We use this method for aggregated units (countries, states, cities, etc.)

Causal Effects: Synthetic Control

• Behind this idea, there is an underling model, with observable (x) and unobservable (u) factors:

$$y_{i,t}(0) = \alpha_i + \gamma_t + \delta_t u_i + x_{i,t}' \beta + \varepsilon_{i,t}.$$

• Under some assumptions, the weights reconstruct $u_{i=N} \& x_{i=N,t}$ just using donor variables:

$$\sum_{j=1}^{N-1} w_j x_{j,t} = x_{i=N,t}$$

$$\sum_{j=1}^{N-1} w_j u_j = u_{i=N}$$

Example: Abadie et al. (2015) estimate the effect of German reunification on GDP per capita using data from five donor countries.

• Abadie (2021) has a review article on JEL.