Causality you can believe in?

The identification challenge and taking assumptions seriously

Chad Hazlett
SICSS 2026, UCLA
Department of Political Science
Department of Statistics and Data Science

The plan

1. Tell you things I think are the main lessons but will sound too abstract and that you won't believe or remember on a first hearing.

The plan

1. Tell you things I think are the main lessons but will sound too abstract and that you won't believe or remember on a first hearing.

2. Backup: preliminary concepts about causality and identification

The plan

1. Tell you things I think are the main lessons but will sound too abstract and that you won't believe or remember on a first hearing.

2. Backup: preliminary concepts about causality and identification

3. Provide a cursory look at the main identification strategies out there.

The plan

1. Tell you things I think are the main lessons but will sound too abstract and that you won't believe or remember on a first hearing.

2. Backup: preliminary concepts about causality and identification

3. Provide a cursory look at the main identification strategies out there.

4. Repeat 1 and see if you believe me.

1. Major points I want to convince you of

1. Major points I want to convince you of

You care about causal inference, really.

1. Major points I want to convince you of

You care about causal inference, really.

Credible causal inference is NOT about more observations, more variables, stronger (or deterministic) effects, smaller p-values.

1. Major points I want to convince you of

You care about causal inference, really.

Credible causal inference is NOT about more observations, more variables, stronger (or deterministic) effects, smaller p-values.

Causal claims depend on assumptions that cannot be verified on the data. Thus,

1. Major points I want to convince you of

You care about causal inference, really.

Credible causal inference is NOT about more observations, more variables, stronger (or deterministic) effects, smaller p-values.

Causal claims depend on assumptions that cannot be verified on the data. Thus,

- the credibility (and transparency) of a causal claim depends on bringing those assumptions to life and reasoning about them.

1. Major points I want to convince you of

You care about causal inference, really.

Credible causal inference is NOT about more observations, more variables, stronger (or deterministic) effects, smaller p-values.

Causal claims depend on assumptions that cannot be verified on the data. Thus,

- the credibility (and transparency) of a causal claim depends on bringing those assumptions to life and reasoning about them.
- there are not "causal methods"; there are causal assumptions you are either clear about or not, and either defend or don't.

1. Major points I want to convince you of

You care about causal inference, really.

Credible causal inference is NOT about more observations, more variables, stronger (or deterministic) effects, smaller p-values.

Causal claims depend on assumptions that cannot be verified on the data. Thus,

- the credibility (and transparency) of a causal claim depends on bringing those assumptions to life and reasoning about them.
- there are not "causal methods"; there are causal assumptions you are either clear about or not, and either defend or don't.
- in other words, no result can be labeled "causal" solely because of your choice of estimator.

1. Major points I want to convince you of

You care about causal inference, really.

Credible causal inference is NOT about more observations, more variables, stronger (or deterministic) effects, smaller p-values.

Causal claims depend on assumptions that cannot be verified on the data. Thus,

- the credibility (and transparency) of a causal claim depends on bringing those assumptions to life and reasoning about them.
- there are not "causal methods"; there are causal assumptions you are either clear about or not, and either defend or don't.
- in other words, no result can be labeled "causal" solely because of your choice of estimator.

There are many choices of assumptions, some not yet exploited, that offer different paths to support a claim. We will explore what is currently relatively standard.

What is causal inference?

This is too basic, but helpful to avoid some misunderstandings.

What is causal inference?

This is too basic, but helpful to avoid some misunderstandings.

You or a loved one probably cares about something causal. Some flavors:

What is causal inference?

This is too basic, but helpful to avoid some misunderstandings.

You or a loved one probably cares about something causal. Some flavors:

- "what is the effect of some cause"

What is causal inference?

This is too basic, but helpful to avoid some misunderstandings.

You or a loved one probably cares about something causal. Some flavors:

- "what is the effect of some cause"
- "what was the cause of some effect"

What is causal inference?

This is too basic, but helpful to avoid some misunderstandings.

You or a loved one probably cares about something causal. Some flavors:

- “what is the effect of some cause”
- “what was the cause of some effect”
- any other query relating to how things influence each other in a possibly complex system

What is causal inference?

This is too basic, but helpful to avoid some misunderstandings.

You or a loved one probably cares about something causal. Some flavors:

- "what is the effect of some cause"
- "what was the cause of some effect"
- any other query relating to how things influence each other in a possibly complex system

Caring about "description" only doesn't always save you:

- Sometimes, descriptives are a trojan horse
- Even purely descriptive queries can turn out to depends on assumptions of underlying causal structure. (e.g. due to sample selection, missingness, attrition).

What is causal inference?

This is too basic, but helpful to avoid some misunderstandings.

You or a loved one probably cares about something causal. Some flavors:

- "what is the effect of some cause"
- "what was the cause of some effect"
- any other query relating to how things influence each other in a possibly complex system

Caring about "description" only doesn't always save you:

- Sometimes, descriptives are a trojan horse
- Even purely descriptive queries can turn out to depends on assumptions of underlying causal structure. (e.g. due to sample selection, missingness, attrition).

We will limit ourselves to simple "effects of a cause"; the other settings are mostly harder.

Proof that you won't escape assumptions

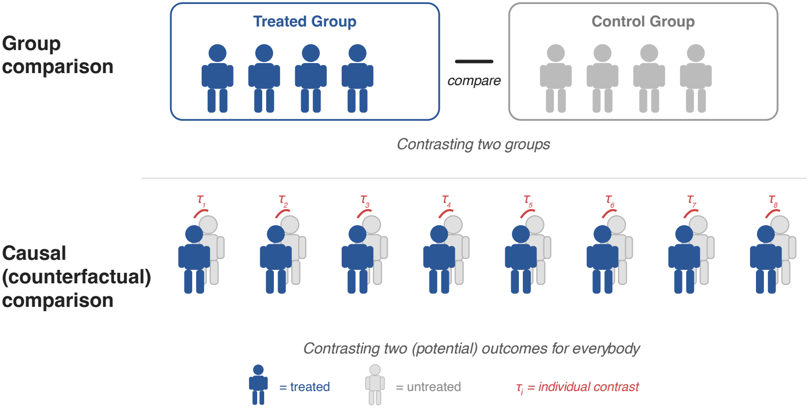

It is half imaginary!

For the causal comparison you would need **both** figures behind each person — the blue (treated) *and* the gray (untreated).

It is half imaginary!

For the causal comparison you would need **both** figures behind each person — the blue (treated) *and* the gray (untreated).

But in your data you only ever see **one of the two**. The other is never observed — for *anybody*.

It is half imaginary!

For the causal comparison you would need **both** figures behind each person — the blue (treated) *and* the gray (untreated).

But in your data you only ever see **one of the two**. The other is never observed — for *anybody*.

There is no escaping it: to make the causal comparison you must **assume** something that lets the half you *do* see stand in for the half you *don't*.

It is half imaginary!

For the causal comparison you would need **both** figures behind each person — the blue (treated) *and* the gray (untreated).

But in your data you only ever see **one of the two**. The other is never observed — for *anybody*.

There is no escaping it: to make the causal comparison you must **assume** something that lets the half you *do* see stand in for the half you *don't*.

Assumptions are the whole game...you can't defend a causal claim without being clear about the assumptions on which it rests.

How then is causal inference possible?

We will invoke sets of assumptions that allow us to say what we *would see* for those missing counterfactuals, one way or another.

How then is causal inference possible?

We will invoke sets of assumptions that allow us to say what we *would see* for those missing counterfactuals, one way or another.

Under such assumptions you can connect something you want to know about counterfactuals, with something you can know from your data.

How then is causal inference possible?

We will invoke sets of assumptions that allow us to say what we *would see* for those missing counterfactuals, one way or another.

Under such assumptions you can connect something you want to know about counterfactuals, with something you can know from your data.

This is called the identification problem:

- Absent identification, any estimate is consistent with many causal quantities...knowing the estimate says nothing about what you wish to know.
- The assumptions we use to achieve this bridge are called identification assumptions.
- Different set of assumptions and procedures that can achieve this are called identification strategies.

Features of identification

- These assumptions have to do with the unseen, therefore the data cannot be enough ... you must always import information or claims from outside the data.

Features of identification

- These assumptions have to do with the unseen, therefore the data cannot be enough … you must always import information or claims from outside the data.
- That is to say they are *not testable* on the data! (but they may be directly falsifiable or cast-doubt-uponable.)

Features of identification

- These assumptions have to do with the unseen, therefore the data cannot be enough … you must always import information or claims from outside the data.
- That is to say they are *not testable* on the data! (but they may be directly falsifiable or cast-doubt-uponable.)
- One trick to distinguish identification from "statistical" assumptions is that making $n \to \infty$ does not make them hold.

Specific identification approaches

We will *briefly* cover a non-exhaustive list of the well-known identification approaches.

Specific identification approaches

We will *briefly* cover a non-exhaustive list of the well-known identification approaches.

You don't get to just pick one...

- the structure of your problem limits which of these apply
- most importantly, each will have different, demanding assumptions about how the world works outside the data, it is about alignment of what you must assume with what is arguable.

Specific identification approaches

We will *briefly* cover a non-exhaustive list of the well-known identification approaches.

You don't get to just pick one...

- the structure of your problem limits which of these apply
- most importantly, each will have different, demanding assumptions about how the world works outside the data, it is about alignment of what you must assume with what is arguable.

$\therefore$ There is no "best" or "most credible" approach – just the question of how well you can argue for the assumptions required in your case.

Specific identification approaches

We will *briefly* cover a non-exhaustive list of the well-known identification approaches.

You don't get to just pick one...

- the structure of your problem limits which of these apply
- most importantly, each will have different, demanding assumptions about how the world works outside the data, it is about alignment of what you must assume with what is arguable.

$\therefore$ There is no "best" or "most credible" approach – just the question of how well you can argue for the assumptions required in your case.

We can sort these into two main groups reflecting an evolution in the field:

- Point identification: idealistic, aspirational
- Partial identification, sensitivity analysis: more honest, often unsatisfying.

Point identification: main ideas

Strategy 1: Argue for no unobserved confounding

- Random treatment (experiments)
- Natural/quasi-experiments
- No unobserved confounding given covariates

Point identification: main ideas

Strategy 1: Argue for no unobserved confounding

- Random treatment (experiments)
- Natural/quasi-experiments
- No unobserved confounding given covariates

Point identification: main ideas

Strategy 1: Argue for no unobserved confounding

- Random treatment (experiments)
- Natural/quasi-experiments
- No unobserved confounding given covariates

Strategy 2: Leveraging small changes

- Regression discontinuity
- Stretch: Interrupted surveys or time series

Point identification: main ideas

Strategy 1: Argue for no unobserved confounding

- Random treatment (experiments)
- Natural/quasi-experiments
- No unobserved confounding given covariates

Strategy 2: Leveraging small changes

- Regression discontinuity
- Stretch: Interrupted surveys or time series

Point identification: main ideas

Strategy 1: Argue for no unobserved confounding

- Random treatment (experiments)
- Natural/quasi-experiments
- No unobserved confounding given covariates

Strategy 2: Leveraging small changes

- Regression discontinuity
- Stretch: Interrupted surveys or time series

Strategy 3: Equiconfounding/ parallel-trends

- Diff-in-diffs, two-way fixed effects and its staggered variants.

Point identification: main ideas

Strategy 1: Argue for no unobserved confounding

- Random treatment (experiments)
- Natural/quasi-experiments
- No unobserved confounding given covariates

Strategy 2: Leveraging small changes

- Regression discontinuity
- Stretch: Interrupted surveys or time series

Strategy 3: Equiconfounding/ parallel-trends

- Diff-in-diffs, two-way fixed effects and its staggered variants.

Strategy 4: Isolating random variation in the treatment: instrumental variables

Point identification: main ideas

Strategy 1: Argue for no unobserved confounding

- Random treatment (experiments)
- Natural/quasi-experiments
- No unobserved confounding given covariates

Strategy 2: Leveraging small changes

- Regression discontinuity
- Stretch: Interrupted surveys or time series

Strategy 3: Equiconfounding/ parallel-trends

- Diff-in-diffs, two-way fixed effects and its staggered variants.

Strategy 4: Isolating random variation in the treatment: instrumental variables

Strategy 5: Learning from prior outcomes/ allowing time-varying confounding:

- Synthetic control, lagged DV, related factor models

Startegy 1: Arguing about confounding

A useful mental model to establish here is "I see some association, but why?"

Startegy 1: Arguing about confounding

A useful mental model to establish here is "I see some association, but why?"

No matter how fancy your estimator looks, it is a way of quantifying the association between your treatment (D) and outcome (Y).

- e.g. a difference in means, a regression coefficient, a matching estimate, AIPW, DML, TMLE, g-comp
- will also apply to the other ID strategies but not always as obviously.

Startegy 1: Arguing about confounding

A useful mental model to establish here is "I see some association, but why?"

No matter how fancy your estimator looks, it is a way of quantifying the association between your treatment (D) and outcome (Y).

- e.g. a difference in means, a regression coefficient, a matching estimate, AIPW, DML, TMLE, g-comp
- will also apply to the other ID strategies but not always as obviously.

Whatever association your approach picks up, we worry about where it is coming from: beyond the $D \rightarrow Y$ you hope to see, it might be due to other *causal* features,

- **omitted variable bias**/ unobserved confounding
- reverse causation
- others, e.g. selection into sample, analytical mistakes

Startegy 1: Arguing about confounding

A useful mental model to establish here is "I see some association, but why?"

No matter how fancy your estimator looks, it is a way of quantifying the association between your treatment (D) and outcome (Y).

- e.g. a difference in means, a regression coefficient, a matching estimate, AIPW, DML, TMLE, g-comp
- will also apply to the other ID strategies but not always as obviously.

Whatever association your approach picks up, we worry about where it is coming from: beyond the $D \to Y$ you hope to see, it might be due to other *causal* features,

- **omitted variable bias**/ unobserved confounding
- reverse causation
- others, e.g. selection into sample, analytical mistakes

Startegy 1: Arguing about confounding

A useful mental model to establish here is "I see some association, but why?"

No matter how fancy your estimator looks, it is a way of quantifying the association between your treatment (D) and outcome (Y).

- e.g. a difference in means, a regression coefficient, a matching estimate, AIPW, DML, TMLE, g-comp
- will also apply to the other ID strategies but not always as obviously.

Whatever association your approach picks up, we worry about where it is coming from: beyond the $D \to Y$ you hope to see, it might be due to other *causal* features,

- **omitted variable bias**/ unobserved confounding
- reverse causation
- others, e.g. selection into sample, analytical mistakes

The job of "ignorability" assumptions is to formally write down the conditions under which the association you see is just due to $D \to Y$.

Randomized experiments

Identification assumption

- If D assigned at random, $\{Y_i(0), Y_i(1)\} \perp\!\!\!\perp D_i$
- This makes mean$(Y(0), Y(1))$ about the same for treated and untreated.
- NB: This assumption comes from outside the data $(p(Y, D, X))$! Where exactly?

Randomized experiments

Identification assumption

- If D assigned at random, $\{Y_i(0), Y_i(1)\} \perp\!\!\!\perp D_i$
- This makes mean$(Y(0), Y(1))$ about the same for treated and untreated.
- NB: This assumption comes from outside the data ($p(Y, D, X)$)! Where exactly?

Estimation:

- can be as simple as difference in means ($\hat{\mathbb{E}}[Y|D=1] - \hat{\mathbb{E}}[Y|D=0]$)
- you may want to verify decent balance on observables ($p(X|D=1) \approx p(X|D=0)$).
- but even for unobservables Z_i you anticipate you have balance
- go to still control for X only improve efficiency (real and reported)

Randomized experiments

Identification assumption

- If D assigned at random, $\{Y_i(0), Y_i(1)\} \perp\!\!\!\perp D_i$
- This makes mean($Y(0), Y(1)$) about the same for treated and untreated.
- NB: This assumption comes from outside the data ($p(Y, D, X)$)! Where exactly?

Estimation:

- can be as simple as difference in means ($\hat{\mathbb{E}}[Y|D = 1] - \hat{\mathbb{E}}[Y|D = 0]$)
- you may want to verify decent balance on observables ($p(X|D = 1) \approx p(X|D = 0)$).
- but even for unobservables Z_i you anticipate you have balance
- go to still control for X only improve efficiency (real and reported)

Threats:

- If experiment went wrong or selective attrition or cross-over etc. broke it.
- Controlling for a post-treatment outcome, asking a question not justified by the randomization (e.g. mediation; effect among those who...)

Natural/Quasi-Experiments

Roughly, situations where a natural mechanism rather than deliberate design does the randomization for you.

Natural/Quasi-Experiments

Roughly, situations where a natural mechanism rather than deliberate design does the randomization for you.

Identification assumption: $\{Y_i(0), Y_i(1)\} \perp\!\!\!\perp D_i$, "ignorability", aka "as-if random"

- This is the same "ignorability" assumption as in randomized experiments.
- Also called "as-if random" or "no confounders" or "exogeneous".

Natural/Quasi-Experiments

Roughly, situations where a natural mechanism rather than deliberate design does the randomization for you.

Identification assumption: $\{Y_i(0), Y_i(1)\} \perp\!\!\!\perp D_i$, "ignorability", aka "as-if random"

- This is the same "ignorability" assumption as in randomized experiments.
- Also called "as-if random" or "no confounders" or "exogeneous".

You will have to argue on the basis of knowledge not in the data, why the things determining D are independent of $Y(d)$.

- Described where randomization is coming from; why for two similar units, one gets treated and the other doesn't.
- Seek and describe potential counterexamples.
- Balance tests (for pre-treatment X), see if $p(X|D=1) \approx p(X|D=0)$) can hope to falsify.

Mistakes to avoid

Failing to argue why ignorability holds even though you didn't get to randomize/ failure to suggest counterexamples.

- Not providing a clear story for why the treatment assignment is independent of the potential outcomes.
- Calling it a "natural experiment" or "quasi-experiment" in your title doesn't actually do anything!

Mistakes to avoid

Failing to argue why ignorability holds even though you didn't get to randomize/ failure to suggest counterexamples.

- Not providing a clear story for why the treatment assignment is independent of the potential outcomes.
- Calling it a "natural experiment" or "quasi-experiment" in your title doesn't actually do anything!

Treating balance tests as proof of ignorability

- Balance tests can only provide evidence against ignorability, not proof of it.
- Hopefully you feel like this is obvious right now... but it's a common mistake to treat balance tests the object of argumentation as if they are sufficient or reveal the key threat. Please don't!

Conditional ignorability: Use and abuse

Imagine you know and measure all the pre-treatment variables X that both affect treatment assignment(D) and the outcome (Y), or really the potential outcomes $Y(d)$. Let's call these confounders.

Conditional ignorability: Use and abuse

Imagine you know and measure all the pre-treatment variables X that both affect treatment assignment(D) and the outcome (Y), or really the potential outcomes $Y(d)$. Let's call these confounders.

Put differently, we hope that for units with the same X values, the reason one took treatment and the other didn't is something like randomization, but more specifically, unrelated to the (potential) outcomes. This is the key idea behind conditional ignorability.

$$\{Y_i(1), Y_i(0)\} \perp\!\!\!\perp D_i \mid X_i = x, \quad \forall x \in \mathcal{X}$$

Conditional ignorability: Use and abuse

Imagine you know and measure all the pre-treatment variables X that both affect treatment assignment(D) and the outcome (Y), or really the potential outcomes $Y(d)$. Let's call these confounders.

Put differently, we hope that for units with the same X values, the reason one took treatment and the other didn't is something like randomization, but more specifically, unrelated to the (potential) outcomes. This is the key idea behind conditional ignorability.

$$\{Y_i(1), Y_i(0)\} \perp\!\!\!\perp D_i \mid X_i = x, \quad \forall x \in \mathcal{X}$$

aka conditional ignorability, selection on observables, no-unmeasured-confounders, satisfying backdoor/adjustment criterion on DAG,...

Conditional ignorability: Use and abuse

Imagine you know and measure all the pre-treatment variables X that both affect treatment assignment(D) and the outcome (Y), or really the potential outcomes $Y(d)$. Let's call these confounders.

Put differently, we hope that for units with the same X values, the reason one took treatment and the other didn't is something like randomization, but more specifically, unrelated to the (potential) outcomes. This is the key idea behind conditional ignorability.

$$\{Y_i(1), Y_i(0)\} \perp\!\!\!\perp D_i \mid X_i = x, \quad \forall x \in \mathcal{X}$$

aka conditional ignorability, selection on observables, no-unmeasured-confounders, satisfying backdoor/adjustment criterion on DAG,...

Estimators use various tricks to either ensure X is similarly distributed across treatment groups or to adjust for the relationship between X and $Y(d)$ in each arm.

How to argue for conditional ignorability?

First, understand it! Find the interpretation that works for you, e.g. "no unobserved confounders", or "experiments within strata of X."

How to argue for conditional ignorability?

First, understand it! Find the interpretation that works for you, e.g. "no unobserved confounders", or "experiments within strata of X." Second, state the assumption clearly

but also bring it to life!. Two useful tools:

How to argue for conditional ignorability?

First, understand it! Find the interpretation that works for you, e.g. "no unobserved confounders", or "experiments within strata of X." Second, state the assumption clearly

but also bring it to life!. Two useful tools:

- **Negative arguments**. Surface counterexamples! This helps to (i) clarify what these assumptions really mean, and (ii) adopt a more self-adversarial strategy to finding problems. Even (especially!) if you can dismiss them.

How to argue for conditional ignorability?

First, understand it! Find the interpretation that works for you, e.g. "no unobserved confounders", or "experiments within strata of X." Second, state the assumption clearly

but also bring it to life!. Two useful tools:

- **Negative arguments**. Surface counterexamples! This helps to (i) clarify what these assumptions really mean, and (ii) adopt a more self-adversarial strategy to finding problems. Even (especially!) if you can dismiss them.
- **Positive arguments**. What is your story for why, within strata of X, treatment is unrelated to potential outcomes? If you are lucky you might:
  - Know enough about treatment assignments to know it depends only on X you have, or on things that arguably do not related to $Y(d)$
  - Know enough about Y to know that there is little room for mystery factors that might also influence D

How to argue for conditional ignorability?

First, understand it! Find the interpretation that works for you, e.g. "no unobserved confounders", or "experiments within strata of X." Second, state the assumption clearly

but also bring it to life!. Two useful tools:

- **Negative arguments**. Surface counterexamples! This helps to (i) clarify what these assumptions really mean, and (ii) adopt a more self-adversarial strategy to finding problems. Even (especially!) if you can dismiss them.
- **Positive arguments**. What is your story for why, within strata of X, treatment is unrelated to potential outcomes? If you are lucky you might:
  - Know enough about treatment assignments to know it depends only on X you have, or on things that arguably do not related to $Y(d)$
  - Know enough about Y to know that there is little room for mystery factors that might also influence D

Sound hard and unlikely? It is! Skepticism is good... sensitivity analysis formalizes it...another day.

Some traps to avoid

1. Failing to clearly state and defend as-if-randomization/ absence of omitted confounders.

- Yes, we see this mistake in print all the time, but you'll want to be transparent about your assumptions and offer clear arguments/ be honest about what you have.

Some traps to avoid

1. Failing to clearly state and defend as-if-randomization/ absence of omitted confounders.

- Yes, we see this mistake in print all the time, but you'll want to be transparent about your assumptions and offer clear arguments/ be honest about what you have.

2. Being fooled by the apparent sophistication of a new estimator. There is long history of estimators here such as

- e.g. stratification/sub-classification, regression (simple or fancy), DML, AIPW, TMLE, g-computation/imputation, matching (many flavors), propensity score weighting, balancing weights (many flavors)...
- These choices *might* mitigate misspecification concerns but no choice between these conditioning technologies will make your identification story better!

Some traps to avoid

1. Failing to clearly state and defend as-if-randomization/ absence of omitted confounders.

- Yes, we see this mistake in print all the time, but you'll want to be transparent about your assumptions and offer clear arguments/ be honest about what you have.

2. Being fooled by the apparent sophistication of a new estimator. There is long history of estimators here such as

- e.g. stratification/sub-classification, regression (simple or fancy), DML, AIPW, TMLE, g-computation/imputation, matching (many flavors), propensity score weighting, balancing weights (many flavors)...
- These choices *might* mitigate misspecification concerns but no choice between these conditioning technologies will make your identification story better!

Some traps to avoid

1. Failing to clearly state and defend as-if-randomization/ absence of omitted confounders.

   - Yes, we see this mistake in print all the time, but you'll want to be transparent about your assumptions and offer clear arguments/ be honest about what you have.

2. Being fooled by the apparent sophistication of a new estimator. There is long history of estimators here such as

   - e.g. stratification/sub-classification, regression (simple or fancy), DML, AIPW, TMLE, g-computation/imputation, matching (many flavors), propensity score weighting, balancing weights (many flavors)...
   - These choices *might* mitigate misspecification concerns but no choice between these conditioning technologies will make your identification story better!

Hopefully it is now clear why you would not say "it's just a regression so it's not causal, but if we use matching or IPW..."

More traps to avoid

3. Over-reliance on balance tests and suggestive results

    - Balance tests are sometimes misunderstood here... your estimator should give you balance on the things you balanced on, that's not evidence of anything but the estimator doing what it should.
    - Balance on observables you didn't explicitly balance is nice in that failure to get it can raise concerns.
    - But ultimately ignorability is about unobservables, not balance on observables.

More traps to avoid

3. Over-reliance on balance tests and suggestive results
    - Balance tests are sometimes misunderstood here... your estimator should give you balance on the things you balanced on, that's not evidence of anything but the estimator doing what it should.
    - Balance on observables you didn't explicitly balance is nice in that failure to get it can raise concerns.
    - But ultimately ignorability is about unobservables, not balance on observables.

So, life is not easy. And yet, we don't want to make it so that only applications with persuasively perfect identification (absence of unobserved confounding) are worth pursuing or publishing...that will bring us to instead engage with transparent discussion...see below.

Strategy 2. Exploit sudden change in $Pr(D)$

Leading example: regression discontinuity designs (RDD)

- suppose each unit has some covariate Z (e.g. test score, vote share)
- if $Z > c$ you have a much higher probability of being treated than if $Z < c$.

Strategy 2. Exploit sudden change in $Pr(D)$

Leading example: regression discontinuity designs (RDD)

- suppose each unit has some covariate Z (e.g. test score, vote share)
- if $Z > c$ you have a much higher probability of being treated than if $Z < c$.

Identification in three flavors,

- assume "for units with Z near enough to c, treatment is random", or more broadly,
- assume that $\mathbb{E}[Y(d)|Z]$ does not jump discontinuously as Z crosses c, or
- comparing units with $Z \in \{c - \epsilon, c\}$ to those with $Z \in \{c, c + \epsilon\}$, very small bias compared to potential treatment effect.

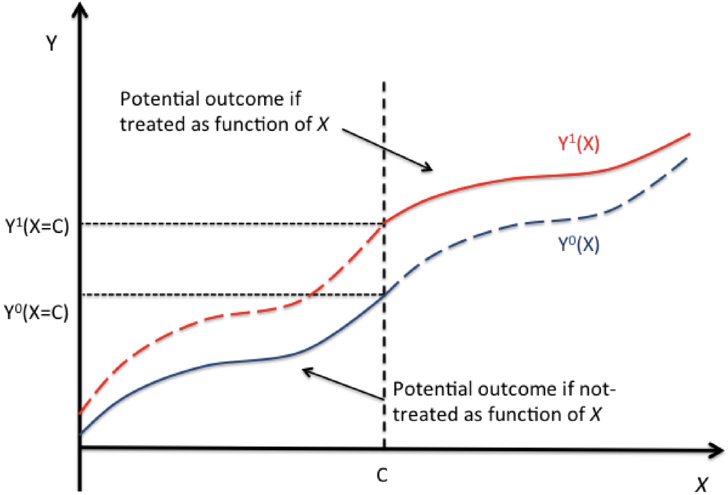

Source: Yona Rubinstein

Strategy 2. Small changes: RDD

Estimation: Essentially model $\mathbb{E}[Y(0)|Z]$ from the left $(Z < c)$ to get a $\widehat{\mathbb{E}}[Y(0)|Z = c]$ and from the right $(Z > c)$ to get $\widehat{\mathbb{E}}[Y(1)|Z = c]$, then compare the two.

- actually choosing and fitting these models is fraught.
- present practice is roughly to tie your hands and rely on something reasonable (`rdrobust`), though in smaller samples I suggest using a Gaussian Process (GP) to do this (`gpss`).

Strategy 2. Small changes: RDD

Estimation: Essentially model $\mathbb{E}[Y(0)|Z]$ from the left ($Z < c$) to get a $\widehat{\mathbb{E}}[Y(0)|Z = c]$ and from the right ($Z > c$) to get $\widehat{\mathbb{E}}[Y(1)|Z = c]$, then compare the two.

- actually choosing and fitting these models is fraught.
- present practice is roughly to tie your hands and rely on something reasonable (`rdrobust`), though in smaller samples I suggest using a Gaussian Process (GP) to do this (`gpss`).

Threats: Overall this is a pretty friendly assumption. The main threat is **other treatment at the same cutoff**.

Strategy 2. Small changes: RDD

Estimation: Essentially model $\mathbb{E}[Y(0)|Z]$ from the left $(Z < c)$ to get a $\widehat{\mathbb{E}}[Y(0)|Z = c]$ and from the right $(Z > c)$ to get $\widehat{\mathbb{E}}[Y(1)|Z = c]$, then compare the two.

- actually choosing and fitting these models is fraught.
- present practice is roughly to tie your hands and rely on something reasonable (`rdrobust`), though in smaller samples I suggest using a Gaussian Process (GP) to do this (`gpss`).

Threats: Overall this is a pretty friendly assumption. The main threat is **other treatment at the same cutoff**.

- if something *else* changes at $Z = c$ and affects Y, it is making $Y(0)$ jump at $Z = c$.

Strategy 2. Small changes: RDD

Estimation: Essentially model $\mathbb{E}[Y(0)|Z]$ from the left $(Z < c)$ to get a $\widehat{\mathbb{E}}[Y(0)|Z = c]$ and from the right $(Z > c)$ to get $\widehat{\mathbb{E}}[Y(1)|Z = c]$, then compare the two.

- actually choosing and fitting these models is fraught.
- present practice is roughly to tie your hands and rely on something reasonable (`rdrobust`), though in smaller samples I suggest using a Gaussian Process (GP) to do this (`gpss`).

Threats: Overall this is a pretty friendly assumption. The main threat is **other treatment at the same cutoff**.

- if something *else* changes at $Z = c$ and affects Y, it is making $Y(0)$ jump at $Z = c$.
- E.g. common administrative cutoffs (e.g. poverty line); geographic boundaries (!!)

Strategy 2b. (Not so) small changes: ITS

The **interrupted time-series**. Everybody is "unexposed" then a surprise event occurs and everybody is "exposed". Let Z be time and c be when the event occurred.

- In ITS you will estimate a model for $\mathbb{E}[Y(0)|Z]$ using $Z < c$, then carry it forward and compare to actual (or fitted) outcomes in $Z > c$.
- Contrast to RD: RD would be perfect if you only wanted to know the effect at time $Z \approx c$
- But often we need time for an effect so asked about some time later.
- Identification is harder, in short
  - need model for $\mathbb{E}[Y(0)|Z]$ learned from $Z < c$ to extrapolate, and note and extrapolated to $Z > c$— must be good enough to use as $Y(0)$ post-treatment.
  - other new things that happen after time c will alias the treatment.

Strategy 3. Equiconfounding/ parallel-trends

Difference in Difference (DiD): Consider the simple two group, two period setting, where both groups start out untreated and then individuals in one of the groups are treated between periods.

Strategy 3. Equiconfounding/ parallel-trends

Difference in Difference (DiD): Consider the simple two group, two period setting, where both groups start out untreated and then individuals in one of the groups are treated between periods.

Identification

- "Parallel trends": absent treatment, the additive change in the two groups would have been the same.
- i.e. $\mathbb{E}[Y_0(post) - Y_0(pre)|G = 1] = \mathbb{E}[Y_0(post) - Y_0(pre)|G = 0]$.
- To write it differently, letting $\Delta Y(0)$ be the the over-time change in $Y(0)$,

$$\Delta Y(0) \perp G$$

- So, just ignorability again but on the trend rather than the level.

Strategy 3. Equiconfounding/ parallel-trends

- Positive arguments?: How is D assigned such that, even though D may be associated with $Y_(0)(pre)$, it is independent of the trend in $Y(0)$?
- Counterexamples?: Anything affecting both the trend in $Y(0)$ and who takes the treatment.

Strategy 3. Equiconfounding/ parallel-trends

- Positive arguments?: How is D assigned such that, even though D may be associated with $Y_(0)(pre)$, it is independent of the trend in $Y(0)$?
- Counterexamples?: Anything affecting both the trend in $Y(0)$ and who takes the treatment.

Strategy 3. Equiconfounding/ parallel-trends

- Positive arguments?: How is D assigned such that, even though D may be associated with $Y_(0)(pre)$, it is independent of the trend in $Y(0)$?
- Counterexamples?: Anything affecting both the trend in $Y(0)$ and who takes the treatment.

How often have you seen such an argument?

Strategy 3. Equiconfounding/ parallel-trends

- Positive arguments?: How is D assigned such that, even though D may be associated with $Y_(0)(pre)$, it is independent of the trend in $Y(0)$?
- Counterexamples?: Anything affecting both the trend in $Y(0)$ and who takes the treatment.

How often have you seen such an argument?

Seeing parallel trends over prior periods does not test parallel trends "when it counts"...and remember the units decided to take treatment when they did for a reason.

Strategy 3. Equiconfounding/ parallel-trends

- Positive arguments?: How is D assigned such that, even though D may be associated with $Y_(0)(pre)$, it is independent of the trend in $Y(0)$?
- Counterexamples?: Anything affecting both the trend in $Y(0)$ and who takes the treatment.

How often have you seen such an argument?

Seeing parallel trends over prior periods does not test parallel trends "when it counts"...and remember the units decided to take treatment when they did for a reason.

Abuses/threats

- Failing to argue why parallel trends should hold.
- Saying you "tested parallel trends" and it holds.

A little more on DiD

When you have more than two-periods, i.e. staggered treatment adoption,

- traditional two-way fixed effect models suffers becuase they sometimes use already-treated units as controls, which only works if treatment effects are homogeneous.
- a raft of new estimators (e.g. Callaway & Sant'Anna 2021, Sub & Abraham 2021, de Chaisemartin-D'Haultfœuille, Borusyak-Jaravel-Spiess, Wooldridge 2021) side-step the TWFE issue, estimating separate clean two-period DiD models and choosing how to average them together.

A little more on DiD

When you have more than two-periods, i.e. staggered treatment adoption,

- traditional two-way fixed effect models suffers becuase they sometimes use already-treated units as controls, which only works if treatment effects are homogeneous.
- a raft of new estimators (e.g. Callaway & Sant'Anna 2021, Sub & Abraham 2021, de Chaisemartin-D'Haultfœuille, Borusyak-Jaravel-Spiess, Wooldridge 2021) side-step the TWFE issue, estimating separate clean two-period DiD models and choosing how to average them together.

A little more on DiD

When you have more than two-periods, i.e. staggered treatment adoption,

- traditional two-way fixed effect models suffers becuase they sometimes use already-treated units as controls, which only works if treatment effects are homogeneous.
- a raft of new estimators (e.g. Callaway & Sant'Anna 2021, Sub & Abraham 2021, de Chaisemartin-D'Haultfœuille, Borusyak-Jaravel-Spiess, Wooldridge 2021) side-step the TWFE issue, estimating separate clean two-period DiD models and choosing how to average them together.

If you have this staggered treatment adoption you should not use TWFE, use one of these approaches but remember these solve a self-inflicted estimation problem, they do not touch the underlying parallel trends assumption.

Strategy 4. Extracting randomness: Instrumental variables

The quickest way to talk about instrumental variables (IV) is probably:

Strategy 4. Extracting randomness: Instrumental variables

The quickest way to talk about instrumental variables (IV) is probably:

- D is confounded with Y, but Z is not

Strategy 4. Extracting randomness: Instrumental variables

The quickest way to talk about instrumental variables (IV) is probably:

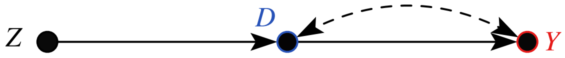

- D is confounded with Y, but Z is not
- For some group ("compliers") Z fully determines D, so for them the effect of Z on Y is the effect of D on Y.

Strategy 4. Extracting randomness: Instrumental variables

The quickest way to talk about instrumental variables (IV) is probably:

- D is confounded with Y, but Z is not
- For some group ("compliers") Z fully determines D, so for them the effect of Z on Y is the effect of D on Y.
- You can recover the effect in that group by taking the overall ("diluted") effect of Z on Y and dividing by the proportion of compliers

Strategy 4. Extracting randomness: Instrumental variables

The quickest way to talk about instrumental variables (IV) is probably:

- D is confounded with Y, but Z is not
- For some group ("compliers") Z fully determines D, so for them the effect of Z on Y is the effect of D on Y.
- You can recover the effect in that group by taking the overall ("diluted") effect of Z on Y and dividing by the proportion of compliers
- Put another way, Z isolates a "random component of D" you can utilize.

Strategy 4: Extracting randomness: Instrumental variables

This comes with its own demanding identification assumptions:

- Z only affects Y through causing movement in D (exclusion restriction)
- Z-Y relationship not itself confounded (exogeneity)
- Z affects D, only in one direction (relevance, monotonicity)

Strategy 4: Extracting randomness: Instrumental variables

This comes with its own demanding identification assumptions:

- Z only affects Y through causing movement in D (exclusion restriction)
- Z-Y relationship not itself confounded (exogeneity)
- Z affects D, only in one direction (relevance, monotonicity)

Use and abuse

- Best when Z itself is randomized, but even then, exclusion restriction needs to be defended with claims from outside the data (and is not testable)
- Often, Z is not randomized and exclusion restriction is indefensible, so big trouble.
- Absent a solid defense of these arguments, there is no reason to regard IV as more credible than anything else.

Strategy 5. Learning from prior outcomes

If parallel trends is a bust, what to do with longitudinal data?

Strategy 5. Learning from prior outcomes

If parallel trends is a bust, what to do with longitudinal data?

With multiple (T) pre-treatment periods, it is tempting to think there useful information in the pre-treatment outcomes, $y_i^{\text{pre}} \equiv [Y_i^{t=1}, Y_i^{t=2} ..., Y_i^{t=T-1}]$.

Strategy 5. Learning from prior outcomes

If parallel trends is a bust, what to do with longitudinal data?

With multiple (T) pre-treatment periods, it is tempting to think there useful information in the pre-treatment outcomes, $y_i^{\text{pre}} \equiv [Y_i^{t=1}, Y_i^{t=2}..., Y_i^{t=T-1}]$.

Specifically you might think confounding factors should "show up" in y_i^{pre}, so let's condition on y_i^{pre} (as if we compare treated and control units with identical y_i^{pre}).

Strategy 5. Learning from prior outcomes

If parallel trends is a bust, what to do with longitudinal data?

With multiple (T) pre-treatment periods, it is tempting to think there useful information in the pre-treatment outcomes, $y_i^{\text{pre}} \equiv [Y_i^{t=1}, Y_i^{t=2}..., Y_i^{t=T-1}]$.

Specifically you might think confounding factors should "show up" in y_i^{pre}, so let's condition on y_i^{pre} (as if we compare treated and control units with identical y_i^{pre}).

- *Synthetic control (synth)*: weight control units to get average Y_{pre} to look the same as for the treated group/unit. Follow that weighted control group forward as our $Y(0)$ prediction to compare with treated group post-tx.

Strategy 5. Learning from prior outcomes

If parallel trends is a bust, what to do with longitudinal data?

With multiple (T) pre-treatment periods, it is tempting to think there useful information in the pre-treatment outcomes, $y_i^{\text{pre}} \equiv [Y_i^{t=1}, Y_i^{t=2}..., Y_i^{t=T-1}]$.

Specifically you might think confounding factors should "show up" in y_i^{pre}, so let's condition on y_i^{pre} (as if we compare treated and control units with identical y_i^{pre}).

- *Synthetic control (synth)*: weight control units to get average Y_{pre} to look the same as for the treated group/unit. Follow that weighted control group forward as our $Y(0)$ prediction to compare with treated group post-tx.

Strategy 5. Learning from prior outcomes

If parallel trends is a bust, what to do with longitudinal data?

With multiple (T) pre-treatment periods, it is tempting to think there useful information in the pre-treatment outcomes, $y_i^{\text{pre}} \equiv [Y_i^{t=1}, Y_i^{t=2}..., Y_i^{t=T-1}]$.

Specifically you might think confounding factors should "show up" in y_i^{pre}, so let's condition on y_i^{pre} (as if we compare treated and control units with identical y_i^{pre}).

- *Synthetic control (synth)*: weight control units to get average Y_{pre} to look the same as for the treated group/unit. Follow that weighted control group forward as our $Y(0)$ prediction to compare with treated group post-tx.
- *Lagged dependent variable* (LDV) regression, i.e. $Y_T \sim D + Y_{T-1} + Y_{T-2}...$ similarly "adjust for pre-treatment outcomes"
- While synth is cool and LDV is practically illegal in some circles, they are virtually the same (e.g. Shen et al 2023)

California smoking example

Omitting details, a picture helps: California cigarette sales before and after 1988 (prop 99) sales tax increase (Abadie, Diamond, Hainmueller 2007)

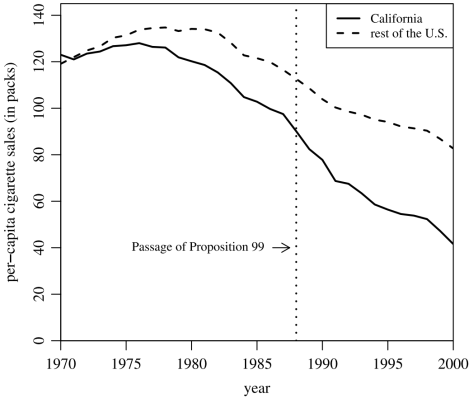

Figure 1. Trends in per-capita cigarette sales: California vs. the rest of the United States.

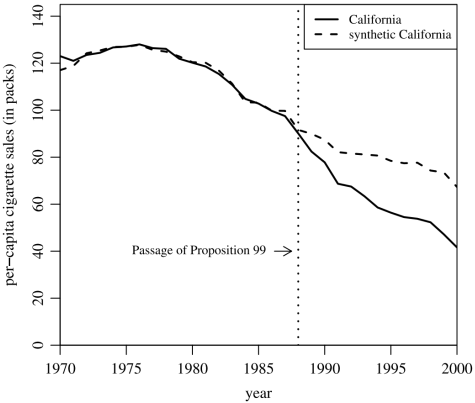

Figure 2. Trends in per-capita cigarette sales: California vs. synthetic California.

What to worry about?

Honestly, the identification challenges with synth (and friends) have been **under-appreciated**.

What to worry about?

Honestly, the identification challenges with synth (and friends) have been **under-appreciated**.

It is seductive: it *feels* like you are holding the counterfactual in your hand.

What to worry about?

Honestly, the identification challenges with synth (and friends) have been **under-appreciated**.

It is seductive: it *feels* like you are holding the counterfactual in your hand.

- Wildly popular — Athey & Imbens call it "arguably the most important innovation in the policy evaluation literature in the last 15 years"; the original ADH paper has ~9,000 citations.

What to worry about?

Honestly, the identification challenges with synth (and friends) have been **under-appreciated**.

It is seductive: it *feels* like you are holding the counterfactual in your hand.

- Wildly popular — Athey & Imbens call it "arguably the most important innovation in the policy evaluation literature in the last 15 years"; the original ADH paper has ∼9,000 citations.
- Yet you almost **never** see the identification assumptions stated, let alone defended, in an application.

What to worry about?

Honestly, the identification challenges with synth (and friends) have been **under-appreciated**.

It is seductive: it *feels* like you are holding the counterfactual in your hand.

- Wildly popular — Athey & Imbens call it "arguably the most important innovation in the policy evaluation literature in the last 15 years"; the original ADH paper has ~9,000 citations.
- Yet you almost **never** see the identification assumptions stated, let alone defended, in an application.

There are **several distinct** challenges. Let's name them — first under a friendly model, then in general.

A handy DGP: the linear factor model

Suppose untreated potential outcomes follow a linear factor model:

$$Y_{it}(0) = \alpha_i + F_t^\top U_i + \varepsilon_{it}$$

A handy DGP: the linear factor model

Suppose untreated potential outcomes follow a linear factor model:

$$Y_{it}(0) = \alpha_i + F_t^\top U_i + \varepsilon_{it}$$

In words:

- U_i: unit i's time-*invariant* "type" (factor loadings)
- F_t: time-*varying* common factors (e.g. health consciousness, anti-tobacco sentiment)
- the confounding lives in $F_t^\top U_i$ — it is time-varying, which is exactly what parallel trends / DiD cannot handle.

Why it could work: the identification chain

The hope rests on a chain:

$$\text{balance on } y_i^{\text{pre}} \;\rightarrow\; \text{balance on } U_i \;\rightarrow\; \text{balance on } Y(0)^{\text{post}} \;\rightarrow\; \text{ATT}$$

Why it could work: the identification chain

The hope rests on a chain:

$$\text{balance on } y_i^{\text{pre}} \;\rightarrow\; \text{balance on } U_i \;\rightarrow\; \text{balance on } Y(0)^{\text{post}} \;\rightarrow\; \text{ATT}$$

If the pre-treatment trajectory is rich enough, units with similar y_i^{pre} must have similar U_i — and similar U_i means similar $Y(0)$ going forward.

Why it could work: the identification chain

The hope rests on a chain:

$$\text{balance on } y_i^{\text{pre}} \;\rightarrow\; \text{balance on } U_i \;\rightarrow\; \text{balance on } Y(0)^{\text{post}} \;\rightarrow\; \text{ATT}$$

If the pre-treatment trajectory is rich enough, units with similar y_i^{pre} must have similar U_i — and similar U_i means similar $Y(0)$ going forward.

The lucky special case — sequential ignorability — is "y_i^{pre} is the *only* confounder." If you truly believed that, you would be done. Usually heroic.

Why it could work: the identification chain

The hope rests on a chain:

$$\text{balance on } y_i^{\text{pre}} \;\to\; \text{balance on } U_i \;\to\; \text{balance on } Y(0)^{\text{post}} \;\to\; \text{ATT}$$

If the pre-treatment trajectory is rich enough, units with similar y_i^{pre} must have similar U_i — and similar U_i means similar $Y(0)$ going forward.

The lucky special case — sequential ignorability — is "y_i^{pre} is the *only* confounder." If you truly believed that, you would be done. Usually heroic.

More realistically, this chain will go through a kind of "recoverablity" logic as follows.

The two main residual biases

Even granting the factor model, balancing on y_i^{pre} is *not* the same as conditioning on U_i:

The two main residual biases

Even granting the factor model, balancing on y_i^{pre} is *not* the same as conditioning on U_i:

1. Errors-in-variables (EIV). y_i^{pre} is a *noisy* view of U_i; balancing on the noise leaves residual imbalance in the true U_i.

- Shrinks like $1/\sqrt{T_{\text{pre}}}$ — a finite-sample problem. More pre-periods eventually fix it.

The two main residual biases

Even granting the factor model, balancing on y_i^{pre} is *not* the same as conditioning on U_i:

1. Errors-in-variables (EIV). y_i^{pre} is a *noisy* view of U_i; balancing on the noise leaves residual imbalance in the true U_i.

- Shrinks like $1/\sqrt{T_{\text{pre}}}$ — a finite-sample problem. More pre-periods eventually fix it.

2. Recoverability. Even with *no* noise, balance on y_i^{pre} buys balance on U_i only if every post-relevant factor leaves a **distinguishable footprint** before treatment.

- In the factor model literature this is baked in as a "rank condition", but it is important to see the content of it.
- Informally there are two main failure modes: (i) few pre-periods than $dim(U)$, and a "sleeping giant": a factor silent before treatment that wakes up after.
- This is an identification-level concern: **more data does not fix it**.

More generally: where does the bias come from?

Drop the factor model. Assume only latent ignorability $Y(0) \perp\!\!\!\perp D \mid U$, a weak DGP $Y(0) = g(U) + \varepsilon$, and that you balance on *some* features $\phi(y_i^{\text{pre}})$.

More generally: where does the bias come from?

Drop the factor model. Assume only latent ignorability $Y(0) \perp\!\!\!\perp D \mid U$, a weak DGP $Y(0) = g(U) + \varepsilon$, and that you balance on *some* features $\phi(y_i^{\text{pre}})$.

The bias of a synth-type estimator then decomposes (Asher, Hazlett & Xu, in progress):

$$\widehat{\tau} - \tau \approx \underbrace{r_{\text{bal}}}_{\text{balancing}} + \underbrace{r_{\text{eiv}}}_{\text{errors-in-variables}} + \underbrace{r_{\text{rec}}}_{\text{recoverability}} + \underbrace{\varepsilon_{T+1}}_{\text{post-period noise}}$$

More generally: where does the bias come from?

Drop the factor model. Assume only latent ignorability $Y(0) \perp\!\!\!\perp D \mid U$, a weak DGP $Y(0) = g(U) + \varepsilon$, and that you balance on *some* features $\phi(y_i^{\text{pre}})$.

The bias of a synth-type estimator then decomposes (Asher, Hazlett & Xu, in progress):

$$\widehat{\tau} - \tau \approx \underbrace{r_{\text{bal}}}_{\text{balancing}} + \underbrace{r_{\text{eiv}}}_{\text{errors-in-variables}} + \underbrace{r_{\text{rec}}}_{\text{recoverability}} + \underbrace{\varepsilon_{T+1}}_{\text{post-period noise}}$$

- r_{bal} how well your weights actually balance the chosen features $\phi(y_i^{\text{pre}})$; shrinks with N.
- r_{eiv} y_i^{pre} is a *noisy* view of U; shrinks as $T_{\text{pre}} \to \infty$, but real in finite T.
- r_{rec} confounding not captured by $\phi(y_i^{\text{pre}})$ (sleeping giants); does not shrink with data.
- ε_{T+1} mean-zero: no bias, but adds variance.

More generally: where does the bias come from?

Drop the factor model. Assume only latent ignorability $Y(0) \perp\!\!\!\perp D \mid U$, a weak DGP $Y(0) = g(U) + \varepsilon$, and that you balance on *some* features $\phi(y_i^{\text{pre}})$.

The bias of a synth-type estimator then decomposes (Asher, Hazlett & Xu, in progress):

$$\widehat{\tau} - \tau \;\approx\; \underbrace{r_{\text{bal}}}_{\text{balancing}} + \underbrace{r_{\text{eiv}}}_{\text{errors-in-variables}} + \underbrace{r_{\text{rec}}}_{\text{recoverability}} + \underbrace{\varepsilon_{T+1}}_{\text{post-period noise}}$$

- r_{bal} how well your weights actually balance the chosen features $\phi(y_i^{\text{pre}})$; shrinks with N.
- r_{eiv} y_i^{pre} is a *noisy* view of U; shrinks as $T_{\text{pre}} \to \infty$, but real in finite T.
- r_{rec} confounding not captured by $\phi(y_i^{\text{pre}})$ (sleeping giants); does not shrink with data.
- ε_{T+1} mean-zero: no bias, but adds variance.

How often do we see arguments about these — especially r_{rec} — in applied work?

Changing the game:

That covers most of the well-known identification strategies; there are some others.

Changing the game:

That covers most of the well-known identification strategies; there are some others.

But that whole game was about "point identification", which is to say "can you construct an estimate you know is unbiased for the target causal effect?"

Changing the game:

That covers most of the well-known identification strategies; there are some others.

But that whole game was about “point identification”, which is to say “can you construct an estimate you know is unbiased for the target causal effect?”

We can instead ask questions like:

- How much can I relax these assumptions and still reach this conclusion?
- What would I have to be willing to assume for a particular conclusion to hold?
- If I adopt some clever, reasonable assumptions, can I narrow down the valid range of conclusions?

A name for the goal: safe inference

Call this alternative posture safe inference:

A name for the goal: safe inference

Call this alternative posture safe inference:

- Make conclusions that are **guaranteed correct under clearly stated assumptions** you are willing to defend

A name for the goal: safe inference

Call this alternative posture safe inference:

- Make conclusions that are **guaranteed correct under clearly stated assumptions** you are willing to defend
- ...or equivalently **tell you what you would have to believe about some contestable assumption to support a given research conclusion**.

A name for the goal: safe inference

Call this alternative posture safe inference:

- Make conclusions that are **guaranteed correct under clearly stated assumptions** you are willing to defend
- ...or equivalently **tell you what you would have to believe about some contestable assumption to support a given research conclusion**.
- Treat "*we can't say*" (under those assumptions) as a **legitimate, informative** answer — not a failure.

A name for the goal: safe inference

Call this alternative posture safe inference:

- Make conclusions that are **guaranteed correct under clearly stated assumptions** you are willing to defend
- ...or equivalently **tell you what you would have to believe about some contestable assumption to support a given research conclusion**.
- Treat "*we can't say*" (under those assumptions) as a **legitimate, informative** answer — not a failure.

This invokes two broad conventional categories of tools:

- **Sensitivity analysis**: how do conclusions change as we relax a point assumption?
- **Partial identification**: for a set of assumptions we will entertain, what *can* we conclude?

Results may look "weaker," but they are honest, and they protect against over-confidence in conclusions the data + defensible assumptions cannot support.

Example: sensitivity for conditional ignorability

Conditional ignorabilitiy remains the most common by far, so let's demonstrate with this.

Example: sensitivity for conditional ignorability

Conditional ignorabilitiy remains the most common by far, so let's demonstrate with this.

The point identification policing perspective:

Example: sensitivity for conditional ignorability

Conditional ignorabilitiy remains the most common by far, so let's demonstrate with this.

The point identification policing perspective:

- Any estimate with any sign is consistent with any true effect...you cannot even say the result is "suggestive" without assumption.

Example: sensitivity for conditional ignorability

Conditional ignorabilitiy remains the most common by far, so let's demonstrate with this.

The point identification policing perspective:

- Any estimate with any sign is consistent with any true effect…you cannot even say the result is "suggestive" without assumption.

The good(ish) news from sensitivity analysis

- Don't argue there is no confounding, ask "how stong would confounding need to be to change the result"
- Interrogate whether confoudning can arguable be limited to an informative leve at present.

Sensitivity analysis to omitted variables

One overview of this change in perspective and a tool to achieve it is:

From "Is It Unconfounded?" to "How Much Confounding Would It Take?": Applying the Sensitivity-Based Approach to Assess Causes of Support for Peace in Colombia

Chad Hazlett, University of California, Los Angeles
Francesca Parente, Christopher Newport University

Sensitivity analysis to omitted variables

One overview of this change in perspective and a tool to achieve it is:

From "Is It Unconfounded?" to "How Much Confounding Would It Take?": Applying the Sensitivity-Based Approach to Assess Causes of Support for Peace in Colombia

Chad Hazlett, University of California, Los Angeles
Francesca Parente, Christopher Newport University

The more technical work behind that is in:

> *Cinelli, Carlos, and Chad Hazlett. "Making sense of sensitivity: Extending omitted variable bias." Journal of the Royal Statistical Society Series B: Statistical Methodology 82.1 (2020): 39-67.*

Sensitivity analysis to omitted variables

In the briefest form I can manage:

Sensitivity analysis to omitted variables

In the briefest form I can manage:

- You have a regression coefficient from a regression that you admit was likely missing important confounders, Z

Sensitivity analysis to omitted variables

In the briefest form I can manage:

- You have a regression coefficient from a regression that you admit was likely missing important confounders, Z
- How would the coefficient have changed had you been able to include those Z?

Sensitivity analysis to omitted variables

In the briefest form I can manage:

- You have a regression coefficient from a regression that you admit was likely missing important confounders, Z
- How would the coefficient have changed had you been able to include those Z?
- The answer depends *only* on how strongly Z linearly predicts D (given included observables X) and how strongly Z linearly predict Y (given X and D).

Sensitivity analysis to omitted variables

In the briefest form I can manage:

- You have a regression coefficient from a regression that you admit was likely missing important confounders, Z
- How would the coefficient have changed had you been able to include those Z?
- The answer depends *only* on how strongly Z linearly predicts D (given included observables X) and how strongly Z linearly predict Y (given X and D).
- You can postulate such strengths and see what it does to the estimate, but also produce some simple sensitivity statistics that summarize how strong confounding would have to be to alter the conclusion.

Sensitivity analysis to omitted variables

In the briefest form I can manage:

- You have a regression coefficient from a regression that you admit was likely missing important confounders, Z
- How would the coefficient have changed had you been able to include those Z?
- The answer depends *only* on how strongly Z linearly predicts D (given included observables X) and how strongly Z linearly predict Y (given X and D).
- You can postulate such strengths and see what it does to the estimate, but also produce some simple sensitivity statistics that summarize how strong confounding would have to be to alter the conclusion.
- These may be directly valuable, especially for revealing very fragile claims. And they can be interpreted with the help of benchmarking exercises that compare problematic confounding strengths to those of observables you may know to be important.

Bias formulas

$$|\widehat{\text{bias}}| = \text{se}(\hat{\tau}_{\text{res}})\sqrt{\frac{R^2_{Y\sim Z|\mathbf{X},D}\; R^2_{D\sim Z|\mathbf{X}}}{1 - R^2_{D\sim Z|\mathbf{X}}}(\text{df})} \tag{1}$$

$$\text{se}(\hat{\tau}) = \text{se}(\hat{\tau}_{\text{res}})\sqrt{\frac{1 - R^2_{Y\sim Z|\mathbf{X},D}}{1 - R^2_{D\sim Z|\mathbf{X}}}\left(\frac{\text{df}}{\text{df} - 1}\right)} \tag{2}$$

Augmenting the standard regression table

Omitting details and context, you can produce results like this,

Augmenting the standard regression table

Omitting details and context, you can produce results like this,

Outcome: *Peace Index*

Treatment:	Est.	SE	t-value	$R^2_{Y\sim D\mid\mathbf{X}}$	RV	$\text{RV}_{\alpha=0.05}$
Directly Harmed	0.097	0.023	4.18	2.2%	13.9%	7.6%

df = 783

- RV: Confounding that explains at least 13.9% of residual variance in both treatment and outcome would eliminate effect estimate; confounding that explains less (of both) cannot.
- $RV_{\alpha=0.05}$: Confounding the explains at least 7.6% of residual variance in treatment and outcome can reduce effect to barely reject null with $p = 0.05$.
- $R^2_{Y\sim D|X}$: If confounding explained 100% of residual variance in outcome, it still has to explain 2.2% of treatment to explain away effect.

Contour plots with bounds

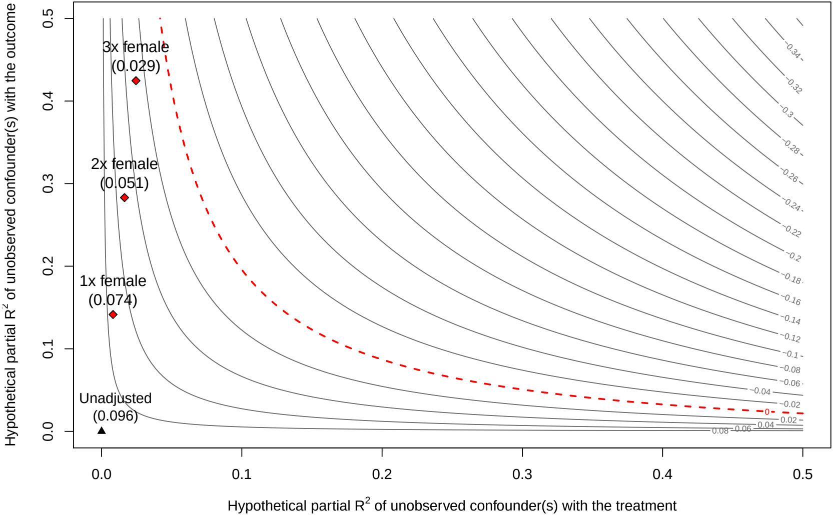

Qualifying claims and complaints

1. The best case: **you make a reasonable defense of a claim**. E.g. here, "A causal interpretation of the research conclusion may be defended by claiming that, given the way injuries (the "treatment") occurred, the scope for targeting particular types of individuals was quite limited... aircraft dropped makeshift and unguided bombs and other objects over villages, and militia raided without concern for who they would attack—the only known major exception to this, due to sexual assaults, was targeting gender, which is also one of the most visually apparent characteristics of an individual. Thus, a confounder twice as strong as female would be surprising."

Qualifying claims and complaints

1. The best case: **you make a reasonable defense of a claim**. E.g. here, "A causal interpretation of the research conclusion may be defended by claiming that, given the way injuries (the "treatment") occurred, the scope for targeting particular types of individuals was quite limited... aircraft dropped makeshift and unguided bombs and other objects over villages, and militia raided without concern for who they would attack—the only known major exception to this, due to sexual assaults, was targeting gender, which is also one of the most visually apparent characteristics of an individual. Thus, a confounder twice as strong as female would be surprising."

2. In other settings the **"to believe research conclusion X, you need to believe X about the strength of confounding"** format tells us where we stand and what not to yet believe.

Qualifying claims and complaints

1. The best case: **you make a reasonable defense of a claim**. E.g. here, "A causal interpretation of the research conclusion may be defended by claiming that, given the way injuries (the "treatment") occurred, the scope for targeting particular types of individuals was quite limited... aircraft dropped makeshift and unguided bombs and other objects over villages, and militia raided without concern for who they would attack—the only known major exception to this, due to sexual assaults, was targeting gender, which is also one of the most visually apparent characteristics of an individual. Thus, a confounder twice as strong as female would be surprising."

2. In other settings the **"to believe research conclusion X, you need to believe X about the strength of confounding"** format tells us where we stand and what not to yet believe.

3. **Not all complaints are sufficient.** For the causal conclusion to be persuasively dismissed, it does not suffice to argue that "*some* confounding might exist". Helpful skepticism articulates why a confounder is strong enough to overturn the concllusion (here, it must explains more than twice of the variation of the treatment assignment than the covariate female is plausible).

What do sensitivity and partial ID really buy us?

You get an honest, (hopefully) transparent statement of what we are prepared to conclude or not, or what you'd have to assume to reach a conclusion.

What do sensitivity and partial ID really buy us?

You get an honest, (hopefully) transparent statement of what we are prepared to conclude or not, or what you'd have to assume to reach a conclusion.

Gets us away from point identification policing

- reduces incentive to oversell your identification assumption
- shows that the blanket "it is not identified" critique is not enough to disqualify...raises the specificity of debate.
- allows us to seek causal claims under imperfect identification, without the fiction.

What do sensitivity and partial ID really buy us?

You get an honest, (hopefully) transparent statement of what we are prepared to conclude or not, or what you'd have to assume to reach a conclusion.

Gets us away from point identification policing

- reduces incentive to oversell your identification assumption
- shows that the blanket "it is not identified" critique is not enough to disqualify...raises the specificity of debate.
- allows us to seek causal claims under imperfect identification, without the fiction.

Not a panecea:

- Sometimes you get lucky and you can defend a conclusion
- More often, get comfortable with disappointment and ignorance.

What do sensitivity and partial ID really buy us?

You get an honest, (hopefully) transparent statement of what we are prepared to conclude or not, or what you'd have to assume to reach a conclusion.

Gets us away from point identification policing

- reduces incentive to oversell your identification assumption
- shows that the blanket "it is not identified" critique is not enough to disqualify...raises the specificity of debate.
- allows us to seek causal claims under imperfect identification, without the fiction.

Not a panecea:

- Sometimes you get lucky and you can defend a conclusion
- More often, get comfortable with disappointment and ignorance.

Great scope and creativity to discover new partial ID ideas that might prove informative. Also apply to IV, ITS, DID at various stages of development.

Safe inteference for Strategy 3: Beyond equiconfounding/parallel trends

The same type of approach applies to other ID settings. For example, for challenging parallel trends, many ideas have been proposed:

Safe inteference for Strategy 3: Beyond equiconfounding/parallel trends

The same type of approach applies to other ID settings. For example, for challenging parallel trends, many ideas have been proposed:

- Allow specified non-parallelness in trends, either directly (subtract!) or scaled by observed features, such as pre-treatment trend differences (e.g. Rambachan and Roth, 2023)

Safe inteference for Strategy 3: Beyond equiconfounding/parallel trends

The same type of approach applies to other ID settings. For example, for challenging parallel trends, many ideas have been proposed:

- Allow specified non-parallelness in trends, either directly (subtract!) or scaled by observed features, such as pre-treatment trend differences (e.g. Rambachan and Roth, 2023)
- Another idea, that also applies in settings without persistent groups: For two cohorts, make an assumption on drift in $\mathbb{E}[Y(0)]$ over time rather than parallel trends (stability controlled quasi-experiment, SCQE) — see below.

Safe inteference for Strategy 3: Beyond equiconfounding/parallel trends

The same type of approach applies to other ID settings. For example, for challenging parallel trends, many ideas have been proposed:

- Allow specified non-parallelness in trends, either directly (subtract!) or scaled by observed features, such as pre-treatment trend differences (e.g. Rambachan and Roth, 2023)
- Another idea, that also applies in settings without persistent groups: For two cohorts, make an assumption on drift in $\mathbb{E}[Y(0)]$ over time rather than parallel trends (stability controlled quasi-experiment, SCQE) — see below.
- Treat prior outcomes as placebos; allow non-equiconfounding...very new area.

Safe inteference for Strategy 3: Beyond equiconfounding/parallel trends

The same type of approach applies to other ID settings. For example, for challenging parallel trends, many ideas have been proposed:

- Allow specified non-parallelness in trends, either directly (subtract!) or scaled by observed features, such as pre-treatment trend differences (e.g. Rambachan and Roth, 2023)
- Another idea, that also applies in settings without persistent groups: For two cohorts, make an assumption on drift in $\mathbb{E}[Y(0)]$ over time rather than parallel trends (stability controlled quasi-experiment, SCQE) — see below.
- Treat prior outcomes as placebos; allow non-equiconfounding...very new area.
- Use $Y_t - Y_{t-1}$ as the outcome, turns DiD into a regression, use the OVB above.

Safe inteference for Strategy 3: Beyond equiconfounding/parallel trends

The same type of approach applies to other ID settings. For example, for challenging parallel trends, many ideas have been proposed:

- Allow specified non-parallelness in trends, either directly (subtract!) or scaled by observed features, such as pre-treatment trend differences (e.g. Rambachan and Roth, 2023)
- Another idea, that also applies in settings without persistent groups: For two cohorts, make an assumption on drift in $\mathbb{E}[Y(0)]$ over time rather than parallel trends (stability controlled quasi-experiment, SCQE) — see below.
- Treat prior outcomes as placebos; allow non-equiconfounding...very new area.
- Use $Y_t - Y_{t-1}$ as the outcome, turns DiD into a regression, use the OVB above.

Mainly, I'm trying to point out that you can be creative about the assumptions you rely on in order to choose those you can plausibly work with.

Other approaches: Stability-Controlled Quasi-Experiments (SCQE)

Once we focus on the goal of "finding assumptions you can bound to get bounded inferences", you find new opportunities for safe inference. E.g., the stability-controlled quasi-experiment (SCQE) approach.

Other approaches: Stability-Controlled Quasi-Experiments (SCQE)

Once we focus on the goal of “finding assumptions you can bound to get bounded inferences”, you find new opportunities for safe inference. E.g., the stability-controlled quasi-experiment (SCQE) approach.

Briefly:

- Suppose you have two groups or cohorts, with very different baseline treatment rates (let's say a “no-use” and “high-use” group).
- Treatment assignment within group is highly non-random.
- The two groups aren't comparable but you are willing to bound how much $\mathbb{E}[Y(0)]$ changes between them (e.g. over time):
  $\delta \equiv \mathbb{E}[Y(0) \mid \text{G=high-use}] - \mathbb{E}[Y(0) \mid \text{G=low-use}]$
- Any range of δ that you cannot rule out $\rightarrow$ range of ATT effects you cannot rule out.

SCQE: Dexamethasone and COVID mortality before RCTs

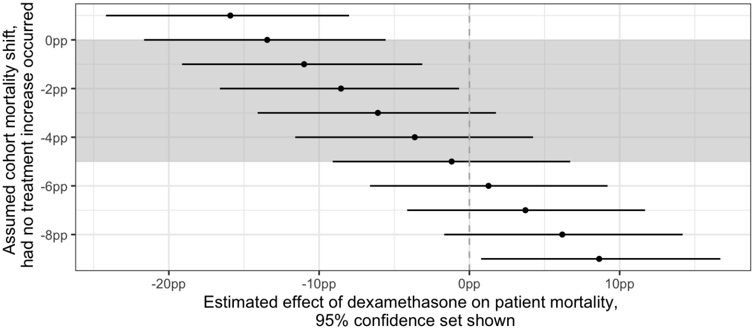

- Beneficial ($p<0.05$) if baseline mortality increased, flat, or fell by up to 2.3pp (-21%)
- Non-harmful point estimate until drop of 5pp (46%)
- Harmful ($p<0.05$) only if baseline mortality fell by 6.8pp (93%), leaving mortality of just 0.4%.

Argumentation approach

You can then make transparent claims such as:

Argumentation approach

You can then make transparent claims such as:

> *"We consider [the drop in baseline mortality required to support a harmful effect] to be extremely unlikely, particularly without a known cause that escaped attention of the physicians on the team and does not appear in the observable differences in cohorts."*

Argumentation approach

You can then make transparent claims such as:

> *"We consider [the drop in baseline mortality required to support a harmful effect] to be extremely unlikely, particularly without a known cause that escaped attention of the physicians on the team and does not appear in the observable differences in cohorts."*
>
> *"Thus, we have considerable confidence that dexamethasone was at least directionally, and possibly statistically significantly beneficial, and that it was not significantly harmful."*

Beyond these particular approaches...

There is nothing exhaustive about the list we've gone through here!

Beyond these particular approaches...

There is nothing exhaustive about the list we've gone through here!

There are other solutions to the problems we've raised, and other types of questions, e.g. those around mediation/mechanisms, spillover/interference, and much more.

Beyond these particular approaches...

There is nothing exhaustive about the list we've gone through here!

There are other solutions to the problems we've raised, and other types of questions, e.g. those around mediation/mechanisms, spillover/interference, and much more.

But I hope you get a sense of the way in which causal identification is at the core of all these issues, tells you what to worry about, and forms the basis for any credible claim.

Returing to step 1: I hope you now believe...

1. Credible causal inference is NOT about more observations, more variables, stronger (or deterministic) effects, smaller p-values.

Returing to step 1: I hope you now believe...

1. Credible causal inference is NOT about more observations, more variables, stronger (or deterministic) effects, smaller p-values.
2. Causal claims depend on assumptions that cannot be verified on the data, so:

Returing to step 1: I hope you now believe...

1. Credible causal inference is NOT about more observations, more variables, stronger (or deterministic) effects, smaller p-values.
2. Causal claims depend on assumptions that cannot be verified on the data, so:
    - You should (and can) bring causal assumptions to life, using positive and negative arguments.

Returing to step 1: I hope you now believe...

1. Credible causal inference is NOT about more observations, more variables, stronger (or deterministic) effects, smaller p-values.
2. Causal claims depend on assumptions that cannot be verified on the data, so:
   - You should (and can) bring causal assumptions to life, using positive and negative arguments.
   - There are not "causal methods"; there are causal assumptions you are either clear about or not, and either defend or don't.

Returing to step 1: I hope you now believe...

1. Credible causal inference is NOT about more observations, more variables, stronger (or deterministic) effects, smaller p-values.
2. Causal claims depend on assumptions that cannot be verified on the data, so:
   - You should (and can) bring causal assumptions to life, using positive and negative arguments.
   - There are not "causal methods"; there are causal assumptions you are either clear about or not, and either defend or don't.
   - Advances in solving estimation problems typically don't touch causal identification problems.

Returing to step 1: I hope you now believe...

1. Credible causal inference is NOT about more observations, more variables, stronger (or deterministic) effects, smaller p-values.
2. Causal claims depend on assumptions that cannot be verified on the data, so:
   - You should (and can) bring causal assumptions to life, using positive and negative arguments.
   - There are not "causal methods"; there are causal assumptions you are either clear about or not, and either defend or don't.
   - Advances in solving estimation problems typically don't touch causal identification problems.
   - Newness, sophistication, and popularity of a method are poor proxies for credibility.

Returing to step 1: I hope you now believe...

1. Credible causal inference is NOT about more observations, more variables, stronger (or deterministic) effects, smaller p-values.

2. Causal claims depend on assumptions that cannot be verified on the data, so:
   - You should (and can) bring causal assumptions to life, using positive and negative arguments.
   - There are not "causal methods"; there are causal assumptions you are either clear about or not, and either defend or don't.
   - Advances in solving estimation problems typically don't touch causal identification problems.
   - Newness, sophistication, and popularity of a method are poor proxies for credibility.

3. Between finding opportunities for point identification and employing sensitivity analysis/ partial ID, you can do careful and honest causal work on interesting questions in your field.

Thank you

Extra slides: Sensitivity for omitted variables

Bias formulas

$$|\widehat{\text{bias}}| = \text{se}(\hat{\tau}_{\text{res}})\sqrt{\frac{R^2_{Y\sim Z|\mathbf{X},D}\ R^2_{D\sim Z|\mathbf{X}}}{1-R^2_{D\sim Z|\mathbf{X}}}(\text{df})} \tag{3}$$

$$\text{se}(\hat{\tau}) = \text{se}(\hat{\tau}_{\text{res}})\sqrt{\frac{1-R^2_{Y\sim Z|\mathbf{X},D}}{1-R^2_{D\sim Z|\mathbf{X}}}\left(\frac{\text{df}}{\text{df}-1}\right)} \tag{4}$$

The landscape (more features to come)

Figure: Contour plot with partial R^2

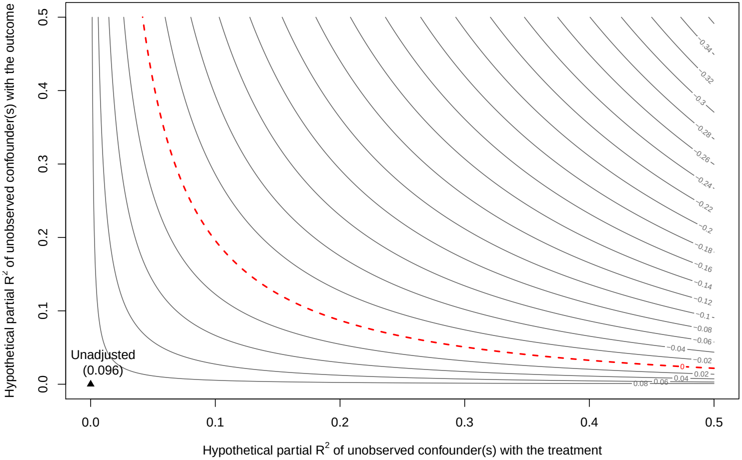

Simple statistics for reporting sensitivity

Robustness Value (RV): what is the point with equal association to treatment and outcome that would eliminate the effect?

Simple statistics for reporting sensitivity

Robustness Value (RV): what is the point with equal association to treatment and outcome that would eliminate the effect?

- RV_q: What confounder, with equal association...would bring down your estimate by some proportion q?

Simple statistics for reporting sensitivity

Robustness Value (RV): what is the point with equal association to treatment and outcome that would eliminate the effect?

- RV_q: What confounder, with equal association...would bring down your estimate by some proportion q?
- RV_α: What confounder, with equal association...would bring down your estimate so that t-statistic corresponds to some α level?

Simple statistics for reporting sensitivity

Robustness Value (RV): what is the point with equal association to treatment and outcome that would eliminate the effect?

- RV_q: What confounder, with equal association...would bring down your estimate by some proportion q?
- RV_α: What confounder, with equal association...would bring down your estimate so that t-statistic corresponds to some α level?

$R^2_{Y\sim D|\mathbf{X}}$: partial variance explained by treatment

Simple statistics for reporting sensitivity

Robustness Value (RV): what is the point with equal association to treatment and outcome that would eliminate the effect?

- RV_q: What confounder, with equal association...would bring down your estimate by some proportion q?
- RV_α: What confounder, with equal association...would bring down your estimate so that t-statistic corresponds to some α level?

$R^2_{Y\sim D|\mathbf{X}}$: partial variance explained by treatment

- *if confounding explains all residual variance of the outcome, you would need* $R^2_{D\sim Z|X}$ *of* $R^2_{Y\sim D|\mathbf{X}}$ *to explain away all of treatment effect estimate.*

Augmenting the standard regression table

Outcome: *Peace Index*

Treatment:	Est.	SE	t-value	$R^2_{Y\sim D\mid\mathbf{X}}$	RV	$\text{RV}_{\alpha=0.05}$
Directly Harmed	0.097	0.023	4.18	2.2%	13.9%	7.6%
df = 783						

- RV: Confounding that explains at least 13.9% of residual variance in both treatment and outcome would eliminate effect estimate; confounding that explains less (of both) cannot.
- $RV_{\alpha=0.05}$: Confounding the explains at least 7.6% of residual variance in treatment and outcome can reduce effect to barely reject null with $p = 0.05$.
- $R^2_{Y\sim D|X}$: If confounding explained 100% of residual variance in outcome, it still has to explain 2.2% of treatment to explain away effect.

Benchmark bounds

One tool for transforming substantive knowledge to judgements about confounding:

- Suppose you can argue an observed (group of) covariate(s) explains more of treatment assignment or outcome than unobserved confounding could.

Benchmark bounds

One tool for transforming substantive knowledge to judgements about confounding:

- Suppose you can argue an observed (group of) covariate(s) explains more of treatment assignment or outcome than unobserved confounding could.
- Generically, there is always some k_D, k_Y s.t.:

$$k_D := \frac{R^2_{D\sim Z|\mathbf{X}_{-j}}}{R^2_{D\sim X_j|\mathbf{X}_{-j}}}, \qquad k_Y := \frac{R^2_{Y\sim Z|\mathbf{X}_{-j},D}}{R^2_{Y\sim X_j|\mathbf{X}_{-j},D}} \tag{5}$$

Benchmark bounds

One tool for transforming substantive knowledge to judgements about confounding:

- Suppose you can argue an observed (group of) covariate(s) explains more of treatment assignment or outcome than unobserved confounding could.
- Generically, there is always some k_D, k_Y s.t.:

$$k_D := \frac{R^2_{D\sim Z|\mathbf{X}_{-j}}}{R^2_{D\sim X_j|\mathbf{X}_{-j}}}, \qquad k_Y := \frac{R^2_{Y\sim Z|\mathbf{X}_{-j},D}}{R^2_{Y\sim X_j|\mathbf{X}_{-j},D}} \tag{5}$$

- For such a claim we can determine the logically implied bounds on the strength of confounding, and ths bias there!

Contour plots with bounds

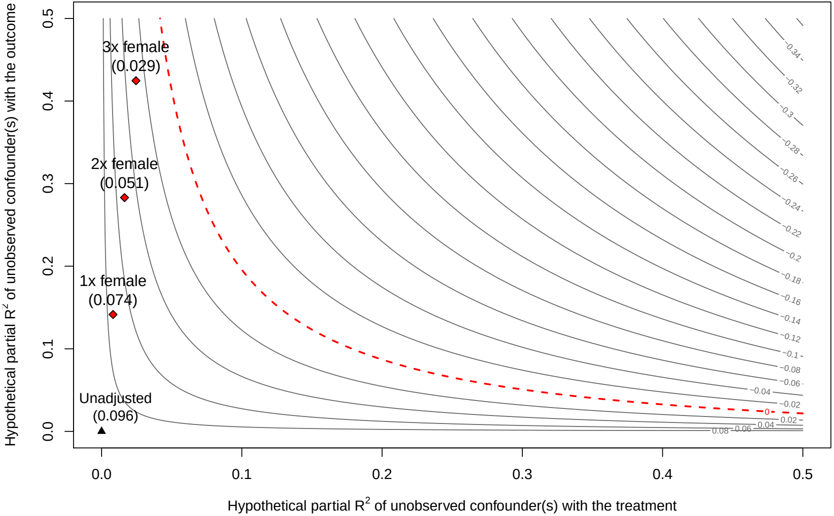

Contour plot of t-statistics

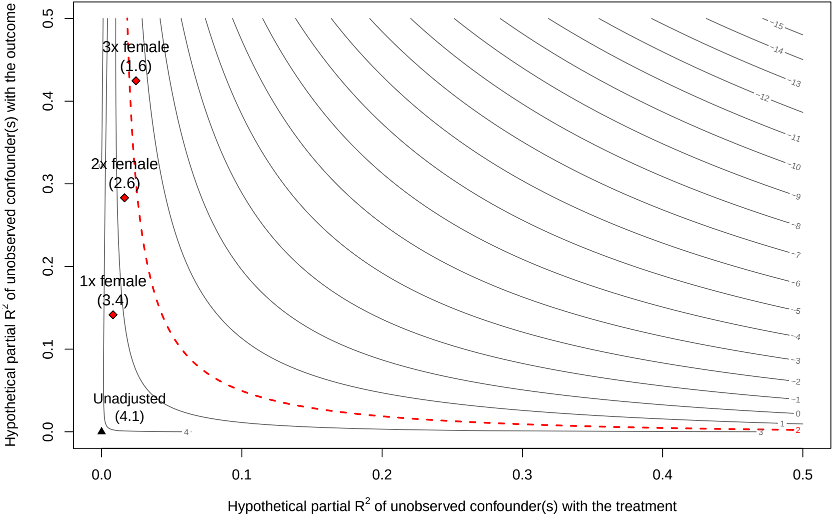

Benchmarks are only as good as your story

- Example: In medicine or law, you may not know all the influences into treatment but you know the main factor doctors/judges consider

Benchmarks are only as good as your story

- Example: In medicine or law, you may not know all the influences into treatment but you know the main factor doctors/judges consider
- In the case shown here: you may believe that nothing can explain more of treatment (and perhaps outcome) than gender, given outside information.

Benchmarks are only as good as your story

- Example: In medicine or law, you may not know all the influences into treatment but you know the main factor doctors/judges consider
- In the case shown here: you may believe that nothing can explain more of treatment (and perhaps outcome) than gender, given outside information.
- In a poorly controlled regression, you may not be able to defend any values of k_D and k_Y for any covariate or even all covariates.

Benchmarks are only as good as your story

- Example: In medicine or law, you may not know all the influences into treatment but you know the main factor doctors/judges consider
- In the case shown here: you may believe that nothing can explain more of treatment (and perhaps outcome) than gender, given outside information.
- In a poorly controlled regression, you may not be able to defend any values of k_D and k_Y for any covariate or even all covariates.
- Or to show hopelessness: if you cannot argue that confounding is limited to some k-times a given X_j, and yet such confounding would flip your answer, you know you are defenseless.

Hopefully it is now obvious, but...

- There cannot exist a suitable guideline for what a “good enough” RV or anything else should be. Why?

Hopefully it is now obvious, but...

- There cannot exist a suitable guideline for what a "good enough" RV or anything else should be. Why?
  - This is not a "test" for confounding; we are revealing what confounding would be problematic.

Hopefully it is now obvious, but...

- There cannot exist a suitable guideline for what a "good enough" RV or anything else should be. Why?
  - This is not a "test" for confounding; we are revealing what confounding would be problematic.
  - The idea is to refine debate about confounding in your situation, not "show robustness"

Hopefully it is now obvious, but...

- There cannot exist a suitable guideline for what a "good enough" RV or anything else should be. Why?
  - This is not a "test" for confounding; we are revealing what confounding would be problematic.
  - The idea is to refine debate about confounding in your situation, not "show robustness"

- Readers should demand good stories for bounding exercise
  - point is not to show "it survives a bound", but rather to convince us that confounding is not likely to be k-times worse than a given covariate.

What kind of conclusiosn does this lead to?

1. A causal interpretation of the research conclusion **may be defended by claiming that**, given the way injuries (the "treatment") occurred, the scope for targeting particular types of individuals was quite limited... aircraft dropped makeshift and unguided bombs and other objects over villages, and militia raided without concern for who they would attack—the only known major exception to this, due to sexual assaults, was targeting gender, which is also one of the most visually apparent characteristics of an individual. Thus, a confounder twice as strong as female would be surprising.

What kind of conclusiosn does this lead to?

1. A causal interpretation of the research conclusion **may be defended by claiming that**, given the way injuries (the "treatment") occurred, the scope for targeting particular types of individuals was quite limited... aircraft dropped makeshift and unguided bombs and other objects over villages, and militia raided without concern for who they would attack—the only known major exception to this, due to sexual assaults, was targeting gender, which is also one of the most visually apparent characteristics of an individual. Thus, a confounder twice as strong as female would be surprising.

2. Similarly, **for the causal conclusion to be persuasively dismissed, it does not suffice to argue that some confounding might exist**. Helpful skepticism must articulate why a confounder that explains more than twice of the variation of the treatment assignment than the covariate female is plausible. Otherwise, the putative confounder cannot logically account for all the observed association, even if it explains all or some large portion of the residual outcome variation.

Good practice:

1. For any regression, provide a table with RV, $\text{RV}_{\alpha=0.05}$, and $R^2_{Y\sim D|\mathbf{X}}$

Good practice:

1. For any regression, provide a table with RV, $\text{RV}_{\alpha=0.05}$, and $R^2_{Y\sim D|\mathbf{X}}$
2. Say what you can about whether confounders that would alter conclusions are likely, given design and knowledge
   - If informative, show bounds and explain why that covariate is likely stronger than confounder
   - Show contour plots to examine more elaborate arguments

Good practice:

1. For any regression, provide a table with RV, $\text{RV}_{\alpha=0.05}$, and $R^2_{Y\sim D|\mathbf{X}}$
2. Say what you can about whether confounders that would alter conclusions are likely, given design and knowledge
   - If informative, show bounds and explain why that covariate is likely stronger than confounder
   - Show contour plots to examine more elaborate arguments
3. These tools also empower you as a reader or reviewers: if researchers do not reveal sensitivity of their estimates, you can do much of this with just t and df.

Good practice:

1. For any regression, provide a table with RV, $\text{RV}_{\alpha=0.05}$, and $R^2_{Y\sim D|\mathbf{X}}$
2. Say what you can about whether confounders that would alter conclusions are likely, given design and knowledge
   - If informative, show bounds and explain why that covariate is likely stronger than confounder
   - Show contour plots to examine more elaborate arguments
3. These tools also empower you as a reader or reviewers: if researchers do not reveal sensitivity of their estimates, you can do much of this with just t and df.

Bigger picture: Moving away from "is it identified", towards more informative debate over "how much confounding would change the conclusion".

Strategy 4: IV, with confounding

While we are here, let's touch on sensitivity with IV.

Strategy 4: IV, with confounding

While we are here, let's touch on sensitivity with IV.

Violations of exclusion and exogeneity can also be cast as omitted variable problems,

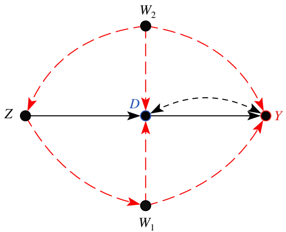

Strategy 4: IV, with confounding

While we are here, let's touch on sensitivity with IV.

Violations of exclusion and exogeneity can also be cast as omitted variable problems,

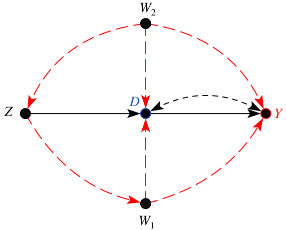

For now I'll say only

- you can perform similar analyses as above (with the *RV* and benchmarks etc.) for the IV case (paper is R&R, Biometrika)
- for purposes of asking what would make the effect zero, you can actually just use the analyses above but on the reduced form (effect of Z on Y)