(Week 5) Causality and hypothesis testing (Week 6) Quant. analysis: strengths & pitfalls

Research design

6 & 13 November, 2017

Andrew Eggers

The “credibility revolution”: from job training to political science

A story about program evaluation

National Supported Work Demonstration (1975-1979): ex-offenders, drug addicts, etc. receive 12-18 months of subsidized employment in 10 US cities.

MDRC implementing NSW in 1970s

Does it work? Of 6,600 eligible participants, some randomly assigned to **control group** (no subsidized employment).

	Treatment	Control
Avg earnings after program	\$4,670	\$3,819

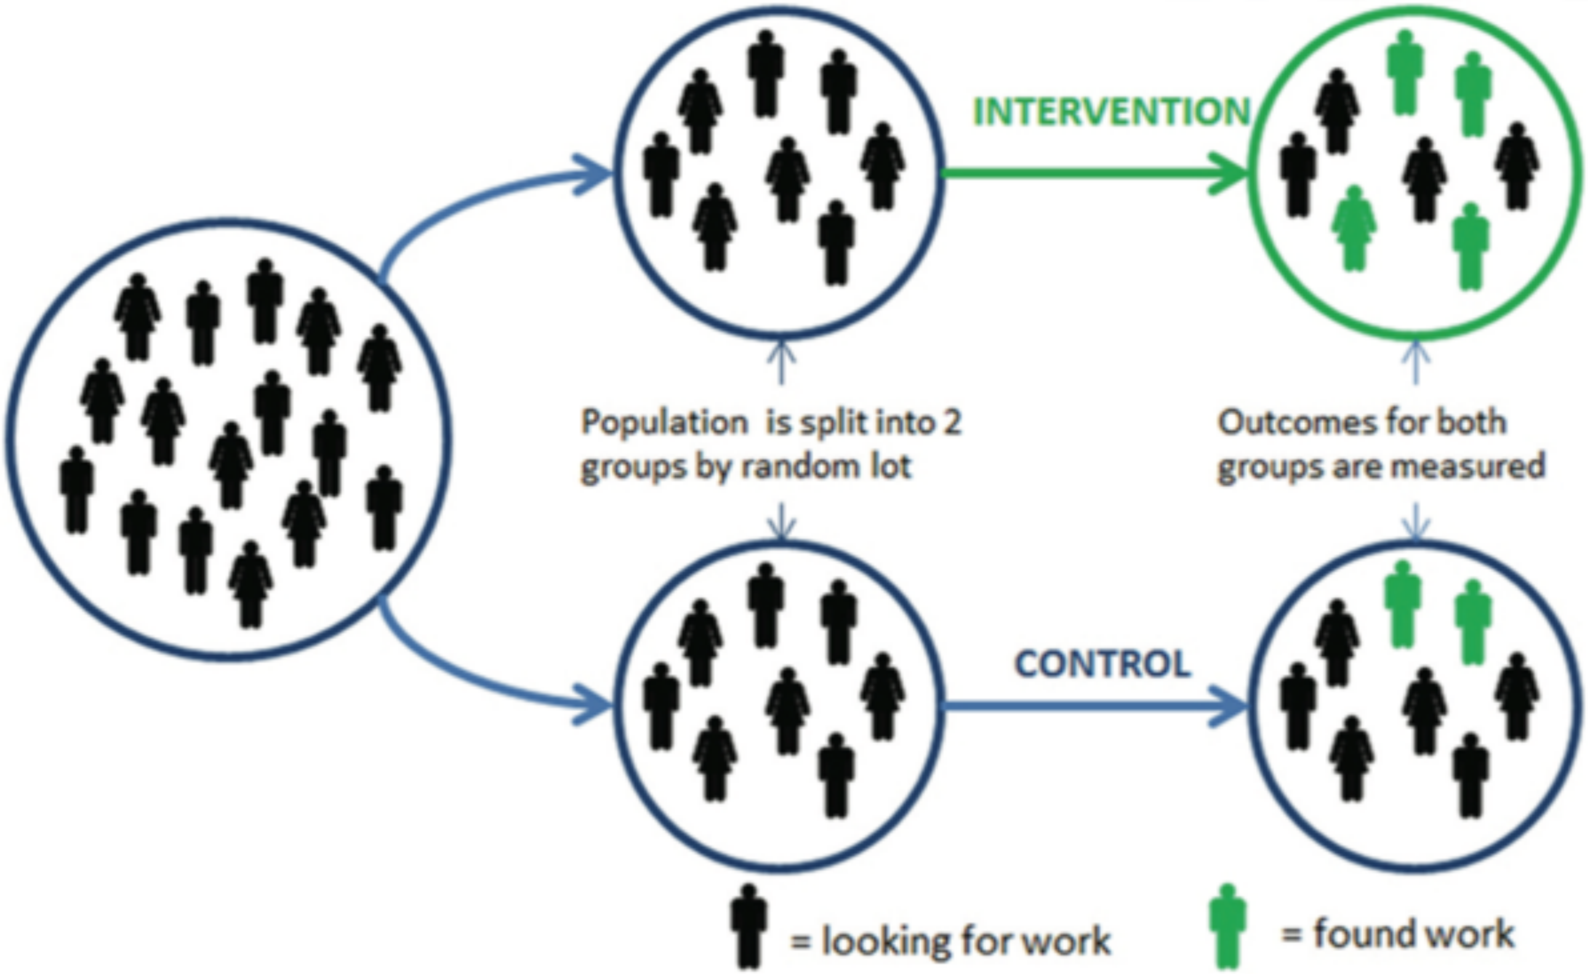

Figure 1. The basic design of a randomised controlled trial (RCT), illustrated with a test of a new 'back to work' programme.

Lalonde (1986): "Evaluating the econometric evaluations of training programs with experimental data"

Robert Lalonde, University of Chicago

Idea: Ignore the experimental control group; use standard economic surveys instead.

How close to the experimental benchmark do we get with standard econometric approaches?

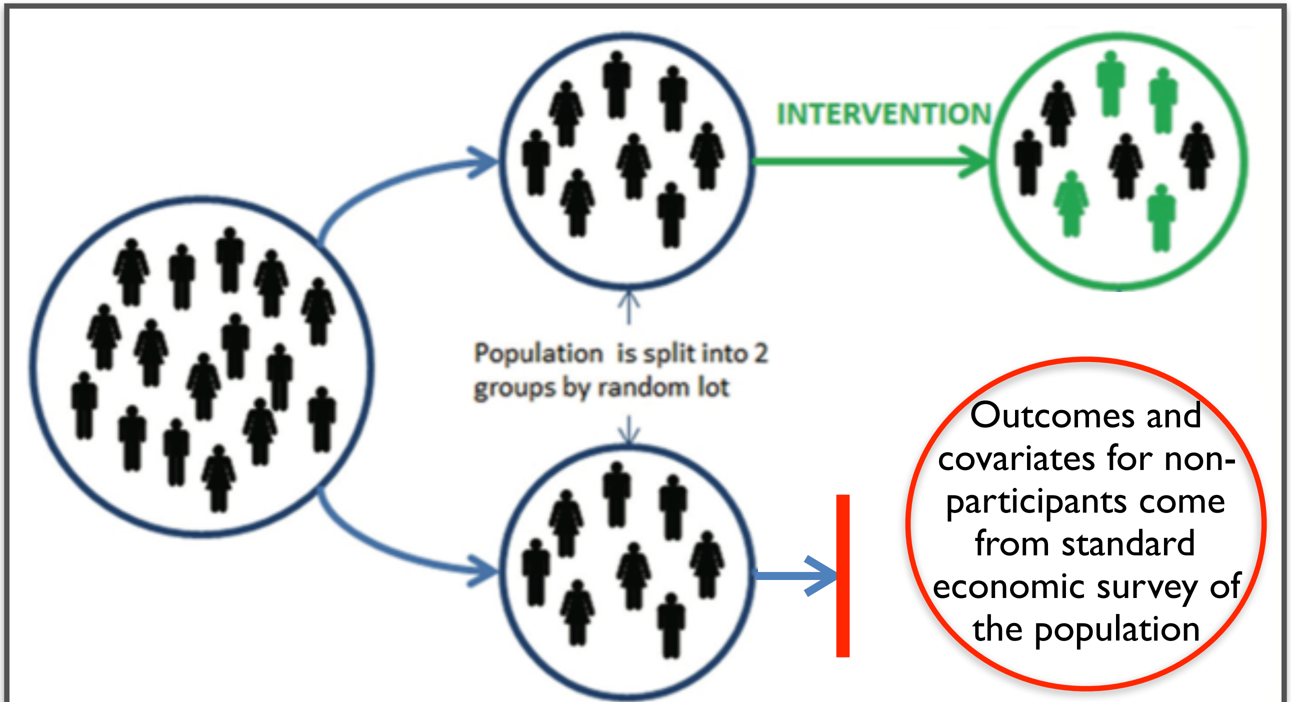

Lalonde (1986): "Evaluating the econometric evaluations of training programs with experimental data"

How close to the experimental benchmark do we get by applying standard econometric approaches to non-experimental data? Not very close!

"Policymakers should be aware that the available non-experimental evaluations of employment and training programs may contain large and unknown biases resulting from specification errors." (p. 617)

Fundamental problem of causal inference

What we want:

Outcome if individual **did participate** in program	minus	Outcome if individual **did not participate** in program
$y_i(1)$	-	$y_i(0)$

Fundamental problem of causal inference is that we never observe both *potential outcomes* for any individual.

Overcoming FPOCI by comparing units?

What about

$$E[y_i(1) - y_i(0)]$$

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

Overcoming FPOCI by comparing units?

What about

Average outcomes of individuals who **did participate** in program

$$E[y_i(1)|t_i = 1]$$

$$E[y_i(1) - y_i(0)]$$

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

Overcoming FPOCI by comparing units?

What about

Average outcomes of individuals who **did participate** in program

minus

$$E[y_i(1)|t_i = 1]$$

-

$$E[y_i(1) - y_i(0)]$$

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

Overcoming FPOCI by comparing units?

What about

Average outcomes of individuals who **did participate** in program	minus	Average outcomes of individuals who **did not participate** in program
$E[y_i(1) \mid t_i = 1]$	-	$E[y_i(0) \mid t_i = 0]$

$$E[y_i(1) - y_i(0)]$$

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

Overcoming FPOCI by comparing units?

What about

Average outcomes of individuals who **did participate** in program	minus	Average outcomes of individuals who **did not participate** in program
$E[y_i(1)\vert t_i = 1]$	-	$E[y_i(0)\vert t_i = 0]$

Is this the same as the average treatment effect?

$$E[y_i(1) - y_i(0)]$$

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

Overcoming FPOCI by comparing units?

What about

Average outcomes of individuals who **did participate** in program	minus	Average outcomes of individuals who **did not participate** in program
$E[y_i(1) \mid t_i = 1]$	-	$E[y_i(0) \mid t_i = 0]$

Is this the same as the average treatment effect?

i.e. $E[y_i(1) - y_i(0)]$

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

Overcoming FPOCI by comparing units?

What about

Average outcomes of individuals who **did participate** in program	minus	Average outcomes of individuals who **did not participate** in program
$E[y_i(1)\|t_i = 1]$	-	$E[y_i(0)\|t_i = 0]$

Is this the same as the average treatment effect?

i.e. average of the individual treatment effects

i.e. $E[y_i(1) - y_i(0)]$

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

Overcoming FPOCI by comparing units? (2)

Overcoming FPOCI by comparing units? (2)

Difference in average outcomes between participants and non-participants reflects (by definition) two things:

Overcoming FPOCI by comparing units? (2)

Difference in average outcomes between participants and non-participants reflects (by definition) two things:

- Average effect of treatment on treated units **(ATT)**

Overcoming FPOCI by comparing units? (2)

Difference in average outcomes between participants and non-participants reflects (by definition) two things:

- Average effect of treatment on treated units **(ATT)**
- Differences between treated units and untreated units **(selection bias)**

Overcoming FPOCI by comparing units? (2)

Difference in average outcomes between participants and non-participants reflects (by definition) two things:

- Average effect of treatment on treated units **(ATT)**
- Differences between treated units and untreated units **(selection bias)**

Basically, correlation ≠ causation.

Overcoming FPOCI by comparing units? (2)

Difference in average outcomes between participants and non-participants reflects (by definition) two things:

- Average effect of treatment on treated units **(ATT)**
- Differences between treated units and untreated units **(selection bias)**

Basically, correlation ≠ causation.

You: "Of course! This is why we control for age, education, etc rather than just compare 'treated' and 'untreated' units."

Lalonde: "This didn't work for measuring the effect of a job training program. Selection bias didn't go away."

Overcoming FPOCI by comparing units? (2)

Difference in average outcomes between participants and non-participants reflects (by definition) two things:

- Average effect of treatment on treated units **(ATT)**
- Differences between treated units and untreated units **(selection bias)**

Basically, correlation ≠ causation.

You: "Of course! This is why we control for age, education, etc rather than just compare 'treated' and 'untreated' units."

Lalonde: "This didn't work for measuring the effect of a job training program. Selection bias didn't go away."

Do we think it works in political science?

What's going wrong?

What's going wrong?

We want to understand the effect of X on Y.

But variation in X is caused by/associated with variation in $Z_1, Z_2, Z_3, \ldots,$ which also affect Y.

What's going wrong?

We want to understand the effect of X on Y.

But variation in X is caused by/associated with variation in $Z_1, Z_2, Z_3, \ldots$, which also affect Y.

(Examples in job training? Effects of electoral system? Democratic peace?)

What's going wrong?

We want to understand the effect of X on Y.

But variation in X is caused by/associated with variation in $Z_1, Z_2, Z_3, \ldots,$ which also affect Y.

(Examples in job training? Effects of electoral system? Democratic peace?)

How do we know

What's going wrong?

We want to understand the effect of X on Y.

But variation in X is caused by/associated with variation in $Z_1, Z_2, Z_3, \ldots,$ which also affect Y.

(Examples in job training? Effects of electoral system? Democratic peace?)

How do we know

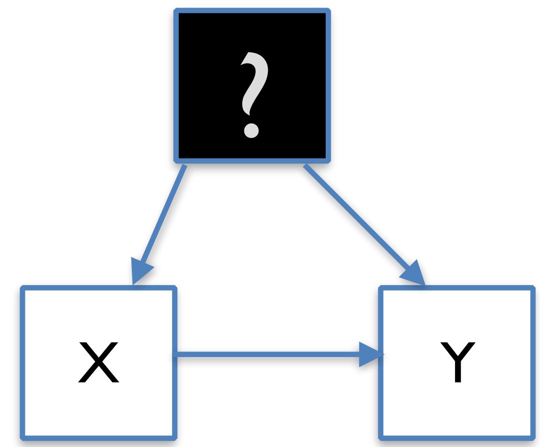

What's going wrong?

We want to understand the effect of X on Y.

But variation in X is caused by/associated with variation in $Z_1, Z_2, Z_3, \ldots,$ which also affect Y.

(Examples in job training? Effects of electoral system? Democratic peace?)

How do we know

- what confounders are important?

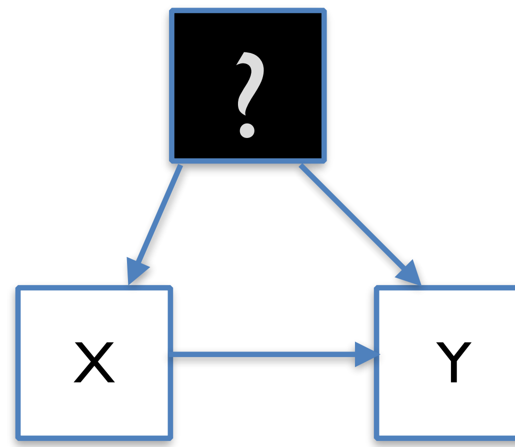

What's going wrong?

We want to understand the effect of X on Y.

But variation in X is caused by/associated with variation in $Z_1, Z_2, Z_3, \ldots,$ which also affect Y.

(Examples in job training? Effects of electoral system? Democratic peace?)

How do we know

- what confounders are important?
- how to measure them?
- how exactly they are related to X or Y?

The **credibility revolution** (Angrist & Pischke 2010) is due to a rising suspicion that **we don't** know.

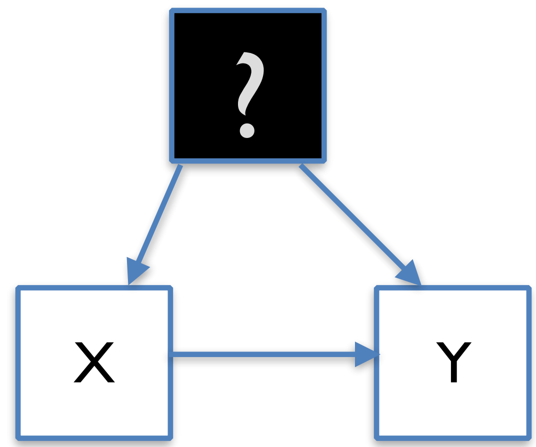

What's going wrong?

Basically, **selection bias** due to unobserved/unobservable characteristics or misspecification errors, i.e. **confounding**, **endogeneity**.

We want to understand the effect of X on Y.

But variation in X is caused by/associated with variation in $Z_1, Z_2, Z_3, \ldots,$ which also affect Y.

(Examples in job training? Effects of electoral system? Democratic peace?)

How do we know

- what confounders are important?
- how to measure them?
- how exactly they are related to X or Y?

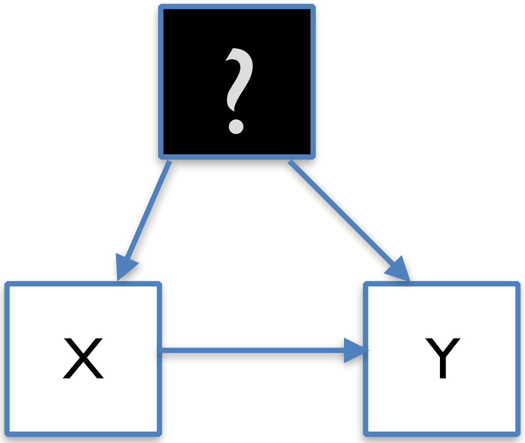

The **credibility revolution** (Angrist & Pischke 2010) is due to a rising suspicion that **we don't know**.

Credibility revolution (?)

Credibility revolution (?)

Critiques like Lalonde's have had a major impact on social sciences — not just evaluation of job training programs.

Credibility revolution (?)

Critiques like Lalonde's have had a major impact on social sciences — not just evaluation of job training programs.

Before: big questions, cross-country regressions with observational data

Credibility revolution (?)

Critiques like Lalonde's have had a major impact on social sciences — not just evaluation of job training programs.

Before: big questions, cross-country regressions with observational data

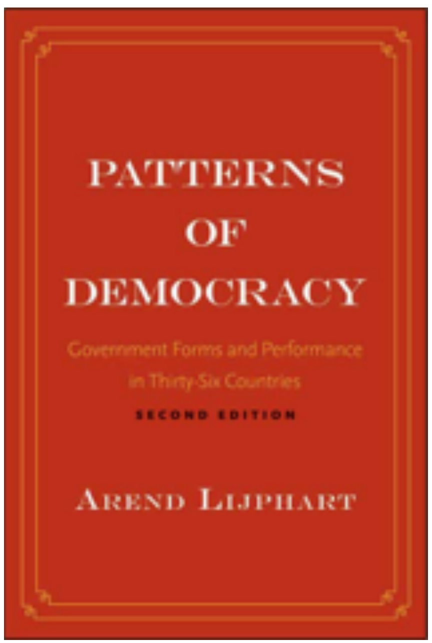

How does government form affect performance?

Credibility revolution (?)

Critiques like Lalonde's have had a major impact on social sciences — not just evaluation of job training programs.

Before: big questions, cross-country regressions with observational data

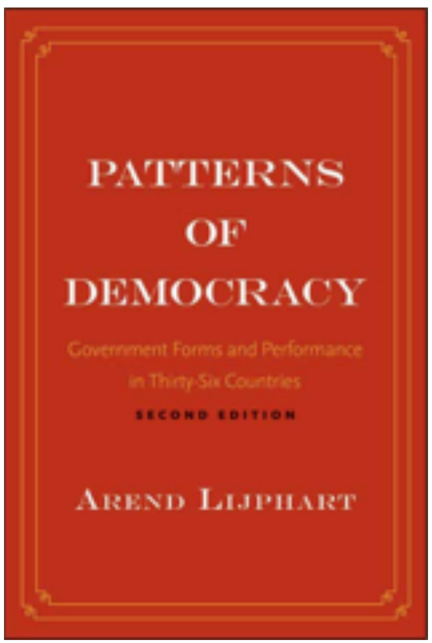

How does government form affect performance?

After: narrow questions, experiments and quasi-experiments, "exploiting variation" within one country

Credibility revolution (?)

Critiques like Lalonde's have had a major impact on social sciences — not just evaluation of job training programs.

Before: big questions, cross-country regressions with observational data

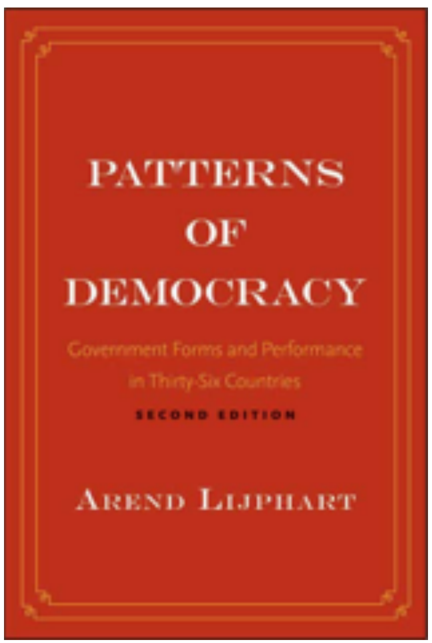

How does government form affect performance?

After: narrow questions, experiments and quasi-experiments, "exploiting variation" within one country

(Many talks in Nuffield Politics seminar, IR colloquium)

The rise of the “identification strategy”

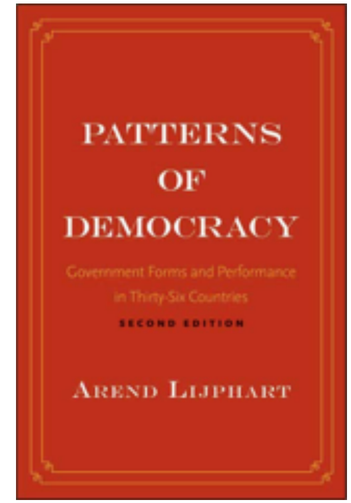

How does government form affect performance?

The rise of the “identification strategy”

Suppose you want to know about the effect of X on Y in population P.

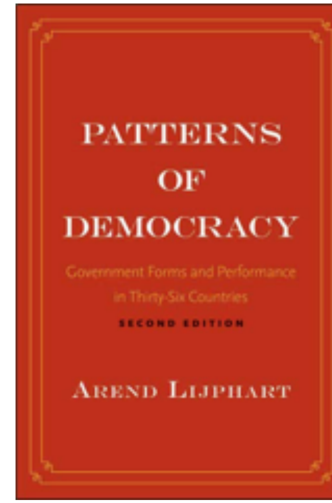

How does government form affect performance?

The rise of the “identification strategy”

Suppose you want to know about the effect of X on Y in population P.

Before: You measure X and Y and some controls Z in (a sample from) population P and run some (possibly complicated) analysis.

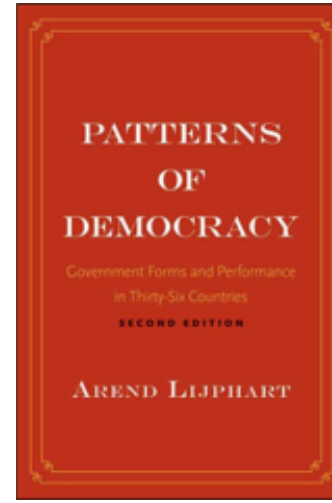

How does government form affect performance?

The rise of the "identification strategy"

Suppose you want to know about the effect of X on Y in population P.

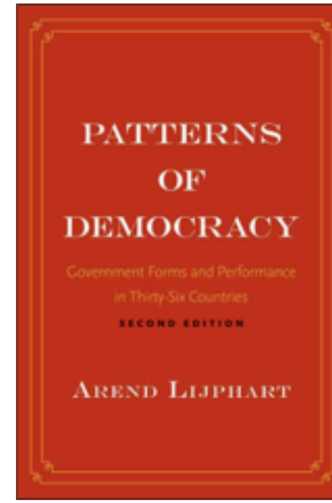

How does government form affect performance?

Before: You measure X and Y and some controls Z in (a sample from) population P and run some (possibly complicated) analysis.

After:

- You run an experiment in which you vary X (or something like X) in a sample that may not look like P; measure difference in means

The rise of the "identification strategy"

Suppose you want to know about the effect of X on Y in population P.

Before: You measure X and Y and some controls Z in (a sample from) population P and run some (possibly complicated) analysis.

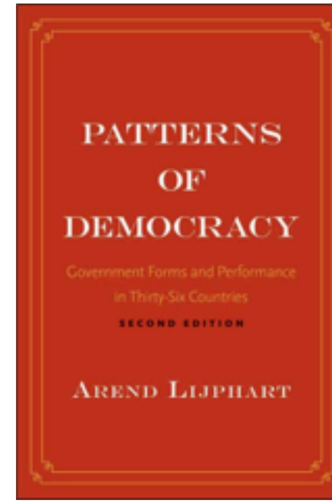

How does government form affect performance?

After:

- You run an experiment in which you vary X (or something like X) in a sample that may not look like P; measure difference in means
- You find a sample in which X varies for (conditionally) arbitrary reasons; measure difference in means

The rise of the "identification strategy"

Suppose you want to know about the effect of X on Y in population P.

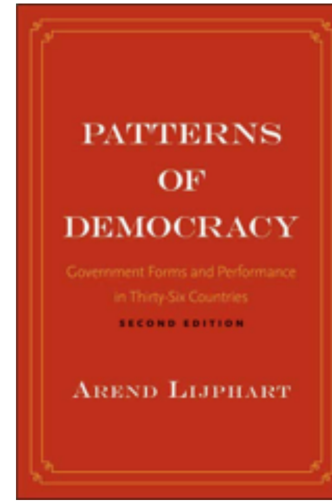

How does government form affect performance?

Before: You measure X and Y and some controls Z in (a sample from) population P and run some (possibly complicated) analysis.

After:

- You run an experiment in which you vary X (or something like X) in a sample that may not look like P; measure difference in means
- You find a sample in which X varies for (conditionally) arbitrary reasons; measure difference in means

Your "identification strategy" is what makes your design more credible than a cross-sectional regression in a sample from P.

Design vs statistical control

Key feature of design-based approach: choosing or creating settings where **statistical control** is **less necessary**.

Research question	Statistical control approach	(Non-experimental) design approach
What is effect of job training program?	Gather data on a bunch of people including participants and non-participants. Regress wages on participation indicator and controls.	Locate job program that was over-subscribed; compare outcomes for successful and unsuccessful applicants.
What is effect of PR (compared to plurality) on turnout?	Gather data on turnout from various countries. Regress turnout on electoral system indicator and controls.	Compare French cities just above and below population cutoff that determines electoral system.

Looking under a lamppost

Internal validity and external validity — see Cyrus Samii "Causal Empiricism"

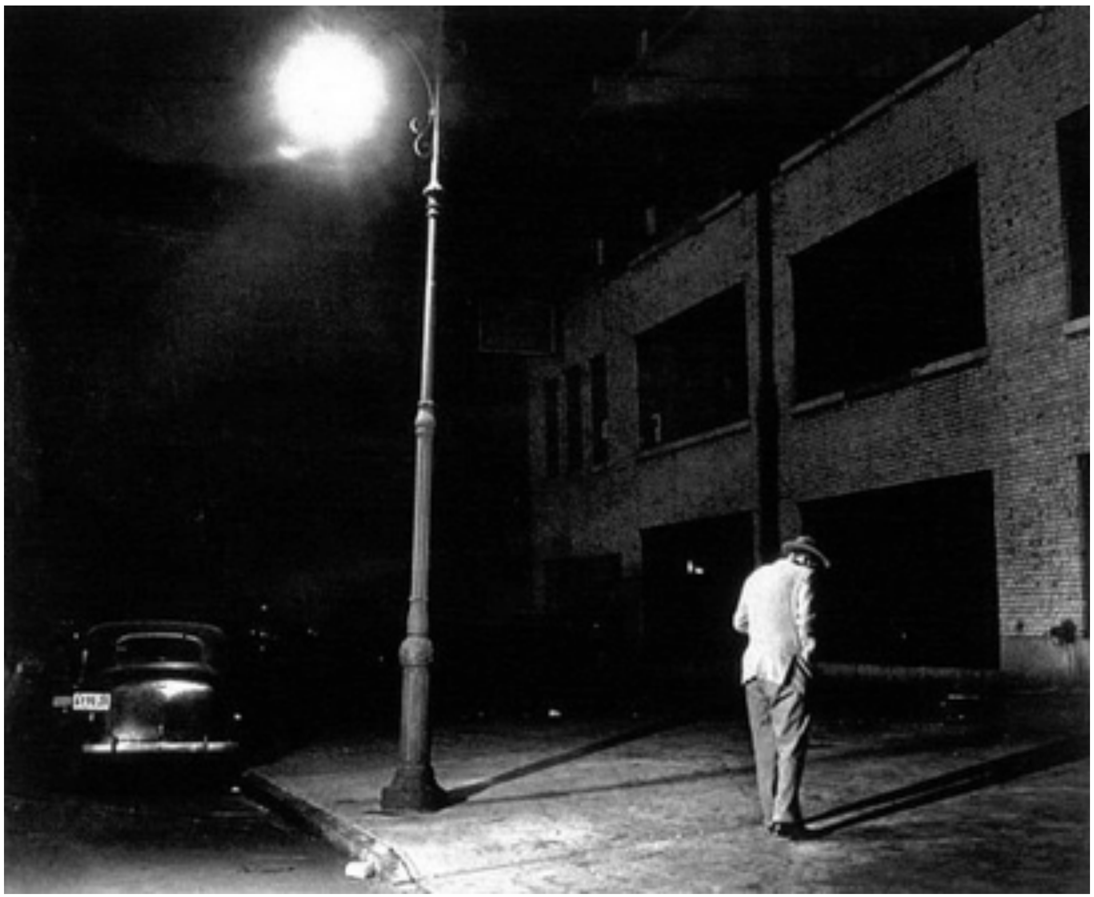

What does this mean for you?

Let’s back up: what are we trying to do anyway?

Let's back up: what are we trying to do anyway?

When we do research, we are trying to address a (conceptual) **problem**: we are trying to address confusion, ignorance, disagreement about something important.

Let's back up: what are we trying to do anyway?

When we do research, we are trying to address a (conceptual) **problem**: we are trying to address confusion, ignorance, disagreement about something important.

A scholar's job is to identify problems and figure out how to address/resolve them.

Let's back up: what are we trying to do anyway?

When we do research, we are trying to address a (conceptual) **problem**: we are trying to address confusion, ignorance, disagreement about something important.

A scholar's job is to identify problems and figure out how to address/resolve them.

The problem usually appears in the introduction of a paper, near a word like "but":

Comparative political economists generally agree that social democratic parties are the defenders of labor. The persistence of widespread unemployment witnessed under social democratic governments since the early 1970s, **however**, powerfully conflicts with this assumption. (Rueda 2005, APSR)

Let's back up: what are we trying to do anyway?

When we do research, we are trying to address a (conceptual) **problem**: we are trying to address confusion, ignorance, disagreement about something important.

A scholar's job is to identify problems and figure out how to address/resolve them.

The problem usually appears in the introduction of a paper, near a word like "but":

Comparative political economists generally agree that social democratic parties are the defenders of labor. The persistence of widespread unemployment witnessed under social democratic governments since the early 1970s, **however**, powerfully conflicts with this assumption. (Rueda 2005, APSR)

Also known as "motive": why are you making me read this?

Some types of conceptual problems

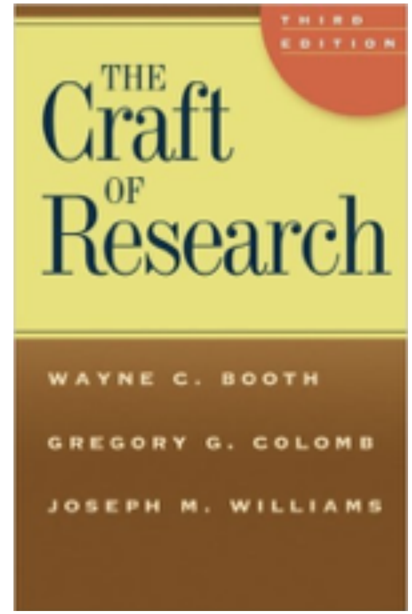

Some types of conceptual problems

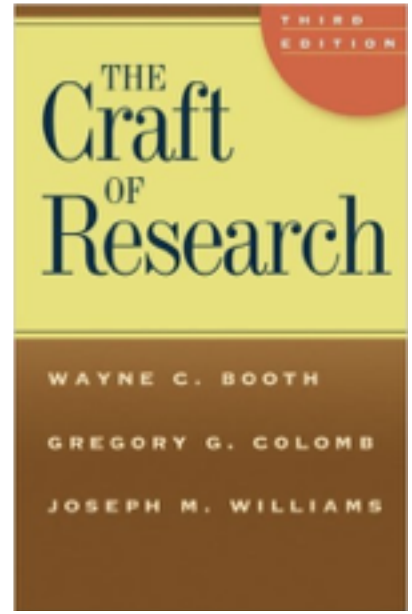

Some types of conceptual problems

- Data does not seem to fit with theory or conventional wisdom about relationship between X and Y

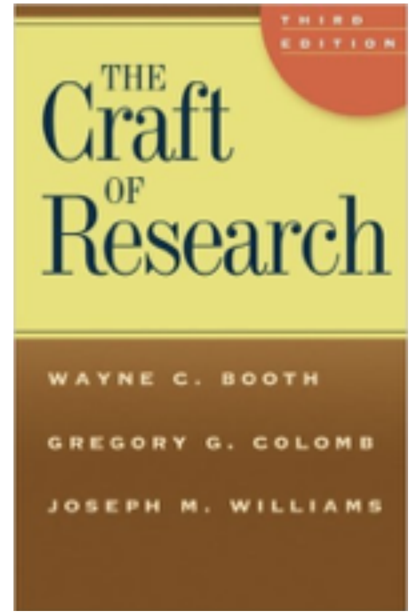

Some types of conceptual problems

- Data does not seem to fit with theory or conventional wisdom about relationship between X and Y
- There are conflicting explanations for Y

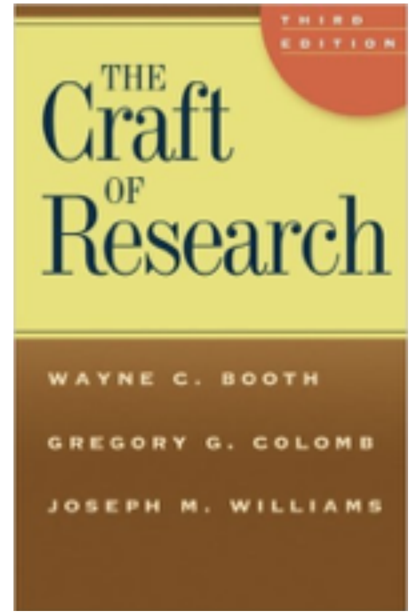

Some types of conceptual problems

- Data does not seem to fit with theory or conventional wisdom about relationship between X and Y
- There are conflicting explanations for Y
- We don't know how X affects Y ("program evaluation")

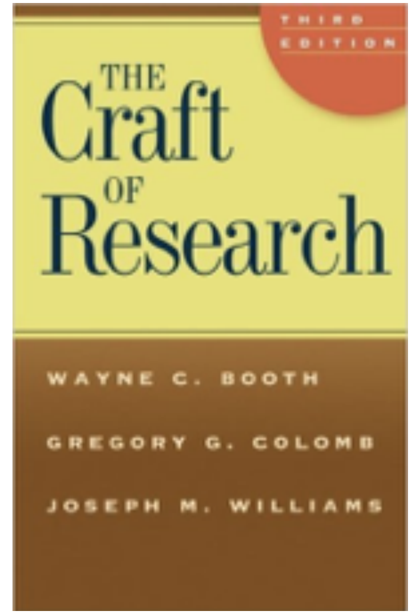

Some types of conceptual problems

- Data does not seem to fit with theory or conventional wisdom about relationship between X and Y
- There are conflicting explanations for Y
- We don't know how X affects Y ("program evaluation")
- We don't know how Y varies with X (descriptive)

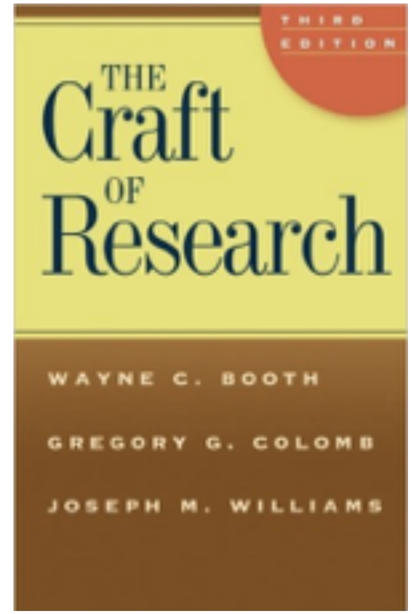

Some types of conceptual problems

- Data does not seem to fit with theory or conventional wisdom about relationship between X and Y
- There are conflicting explanations for Y
- We don't know how X affects Y ("program evaluation")
- We don't know how Y varies with X (descriptive)

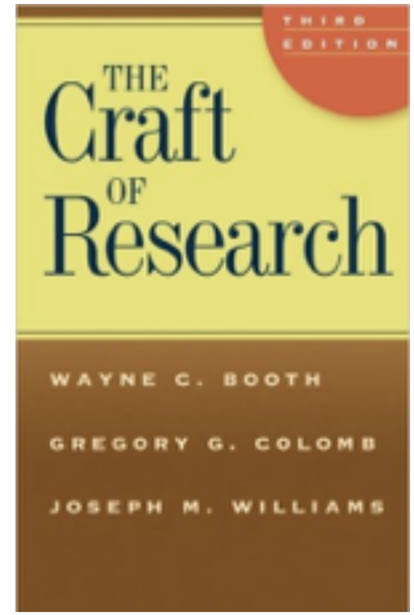

Some types of conceptual problems

- Data does not seem to fit with theory or conventional wisdom about relationship between X and Y
- There are conflicting explanations for Y
- We don't know how X affects Y ("program evaluation")
- We don't know how Y varies with X (descriptive)

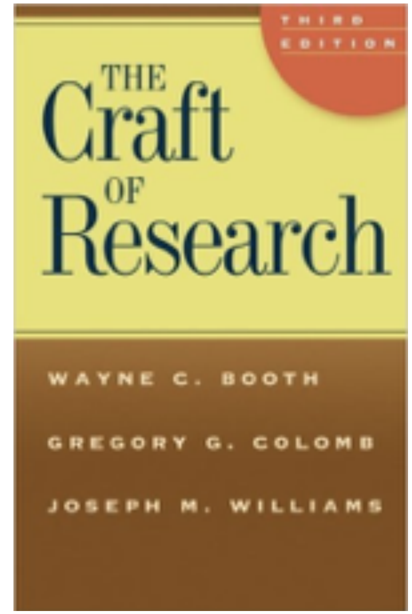

Notes:

Some types of conceptual problems

- Data does not seem to fit with theory or conventional wisdom about relationship between X and Y
- There are conflicting explanations for Y
- We don't know how X affects Y ("program evaluation")
- We don't know how Y varies with X (descriptive)

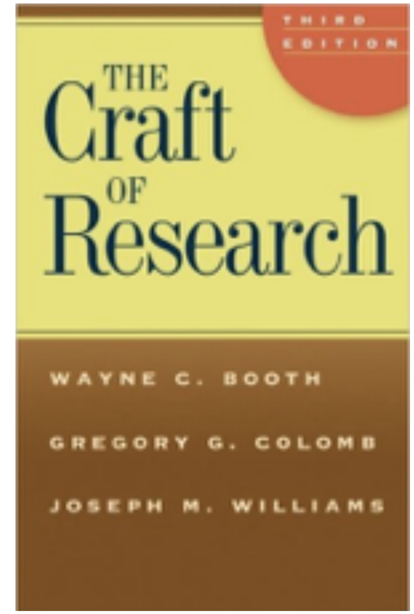

Notes:

- You can start with a **topic** (e.g. "the Left in France", "globalization and preferences"), but eventually you need a **problem**.

Some types of conceptual problems

- Data does not seem to fit with theory or conventional wisdom about relationship between X and Y
- There are conflicting explanations for Y
- We don't know how X affects Y ("program evaluation")
- We don't know how Y varies with X (descriptive)

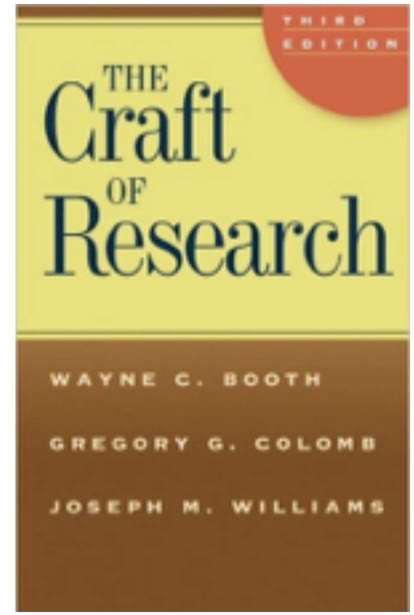

Notes:

- You can start with a **topic** (e.g. "the Left in France", "globalization and preferences"), but eventually you need a **problem**.
- Identifying a problem requires reading the literature. (Gary King: "Whose mind does this change about what?")

Some types of conceptual problems

- Data does not seem to fit with theory or conventional wisdom about relationship between X and Y
- There are conflicting explanations for Y
- We don't know how X affects Y ("program evaluation")
- We don't know how Y varies with X (descriptive)

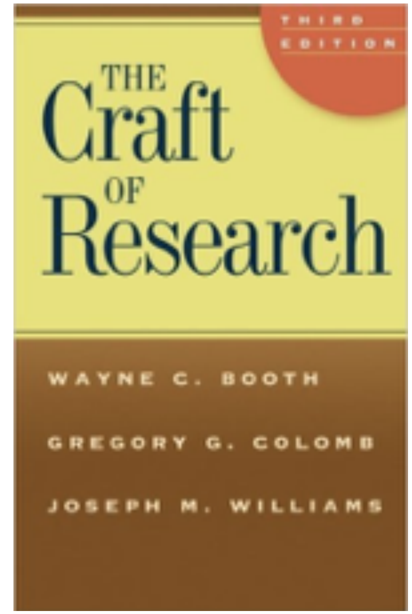

Notes:

- You can start with a **topic** (e.g. "the Left in France", "globalization and preferences"), but eventually you need a **problem**.
- Identifying a problem requires reading the literature. (Gary King: "Whose mind does this change about what?")
- You also need to convince us that it is **important** to address this problem.

So what does this have to do with causal inference and the credibility revolution?

So what does this have to do with causal inference and the credibility revolution?

The ingredients of good empirical research:

So what does this have to do with causal inference and the credibility revolution?

The ingredients of good empirical research:

- a problem (real and consequential)

So what does this have to do with causal inference and the credibility revolution?

The ingredients of good empirical research:

- a problem (real and consequential)
- a solution: a piece of analysis that addresses the problem

So what does this have to do with causal inference and the credibility revolution?

The ingredients of good empirical research:

- a problem (real and consequential)
- a solution: a piece of analysis that addresses the problem

So what does this have to do with causal inference and the credibility revolution?

The ingredients of good empirical research:

- a problem (real and consequential)
- a solution: a piece of analysis that addresses the problem

What has changed? The way people assess proposed solutions that rely on empirical measurements of effects.

So what does this have to do with causal inference and the credibility revolution?

The ingredients of good empirical research:

- a problem (real and consequential)
- a solution: a piece of analysis that addresses the problem

What has changed? The way people assess proposed solutions that rely on empirical measurements of effects.

So what does this have to do with causal inference and the credibility revolution?

The ingredients of good empirical research:

- a problem (real and consequential)
- a solution: a piece of analysis that addresses the problem

What has changed? The way people assess proposed solutions that rely on empirical measurements of effects.

One response: start by looking for research designs that *randomnistas* will find credible, then see if it addresses a problem.

So what does this have to do with causal inference and the credibility revolution?

The ingredients of good empirical research:

- a problem (real and consequential)
- a solution: a piece of analysis that addresses the problem

What has changed? The way people assess proposed solutions that rely on empirical measurements of effects.

One response: start by looking for research designs that *randomnistas* will find credible, then see if it addresses a problem.

- Is ballot order randomized in California?

So what does this have to do with causal inference and the credibility revolution?

The ingredients of good empirical research:

- a problem (real and consequential)
- a solution: a piece of analysis that addresses the problem

What has changed? The way people assess proposed solutions that rely on empirical measurements of effects.

One response: start by looking for research designs that *randomnistas* will find credible, then see if it addresses a problem.

- Is ballot order randomized in California?
- Does the municipal electoral system depend on an arbitrary population cutoff?

So what does this have to do with causal inference and the credibility revolution?

The ingredients of good empirical research:

- a problem (real and consequential)
- a solution: a piece of analysis that addresses the problem

What has changed? The way people assess proposed solutions that rely on empirical measurements of effects.

One response: start by looking for research designs that *randomnistas* will find credible, then see if it addresses a problem.

- Is ballot order randomized in California?
- Does the municipal electoral system depend on an arbitrary population cutoff?
- Is there detailed weather data?

Analyze first, ask questions (i.e. decide on a problem) **later?**

Analyze first, ask questions (i.e. decide on a problem) **later?**

This actually works: there are more interesting questions than credible research designs, so why not start with the design?

Analyze first, ask questions (i.e. decide on a problem) **later?**

This actually works: there are more interesting questions than credible research designs, so why not start with the design?

Analyze first, ask questions (i.e. decide on a problem) **later?**

This actually works: there are more interesting questions than credible research designs, so why not start with the design?

And it can work for other types of research too:

Analyze first, ask questions (i.e. decide on a problem) **later?**

This actually works: there are more interesting questions than credible research designs, so why not start with the design?

And it can work for other types of research too:

- obtain some new data, document some patterns…

Analyze first, ask questions (i.e. decide on a problem) **later?**

This actually works: there are more interesting questions than credible research designs, so why not start with the design?

And it can work for other types of research too:

- obtain some new data, document some patterns…
- adapt a formal model to a different setting, study its features…

Analyze first, ask questions (i.e. decide on a problem) **later?**

This actually works: there are more interesting questions than credible research designs, so why not start with the design?

And it can work for other types of research too:

- obtain some new data, document some patterns…
- adapt a formal model to a different setting, study its features…

find out if there is something interesting that addresses confusion/ignorance/disagreement in literature.

Analyze first, ask questions (i.e. decide on a problem) **later?**

This actually works: there are more interesting questions than credible research designs, so why not start with the design?

And it can work for other types of research too:

- obtain some new data, document some patterns…
- adapt a formal model to a different setting, study its features…

find out if there is something interesting that addresses confusion/ignorance/disagreement in literature.

Analyze first, ask questions (i.e. decide on a problem) **later?**

This actually works: there are more interesting questions than credible research designs, so why not start with the design?

And it can work for other types of research too:

- obtain some new data, document some patterns…
- adapt a formal model to a different setting, study its features…

find out if there is something interesting that addresses confusion/ignorance/disagreement in literature.

But the method- or data-driven approach often leads to trivial questions, scattered research profile.

Analyze first, ask questions (i.e. decide on a problem) **later?**

This actually works: there are more interesting questions than credible research designs, so why not start with the design?

And it can work for other types of research too:

- obtain some new data, document some patterns…
- adapt a formal model to a different setting, study its features…

find out if there is something interesting that addresses confusion/ignorance/disagreement in literature.

But the method- or data-driven approach often leads to trivial questions, scattered research profile.

And causal inference-driven approach excludes other types of questions: descriptive, explanatory, conceptual/theoretical.

The credibility/importance tradeoff

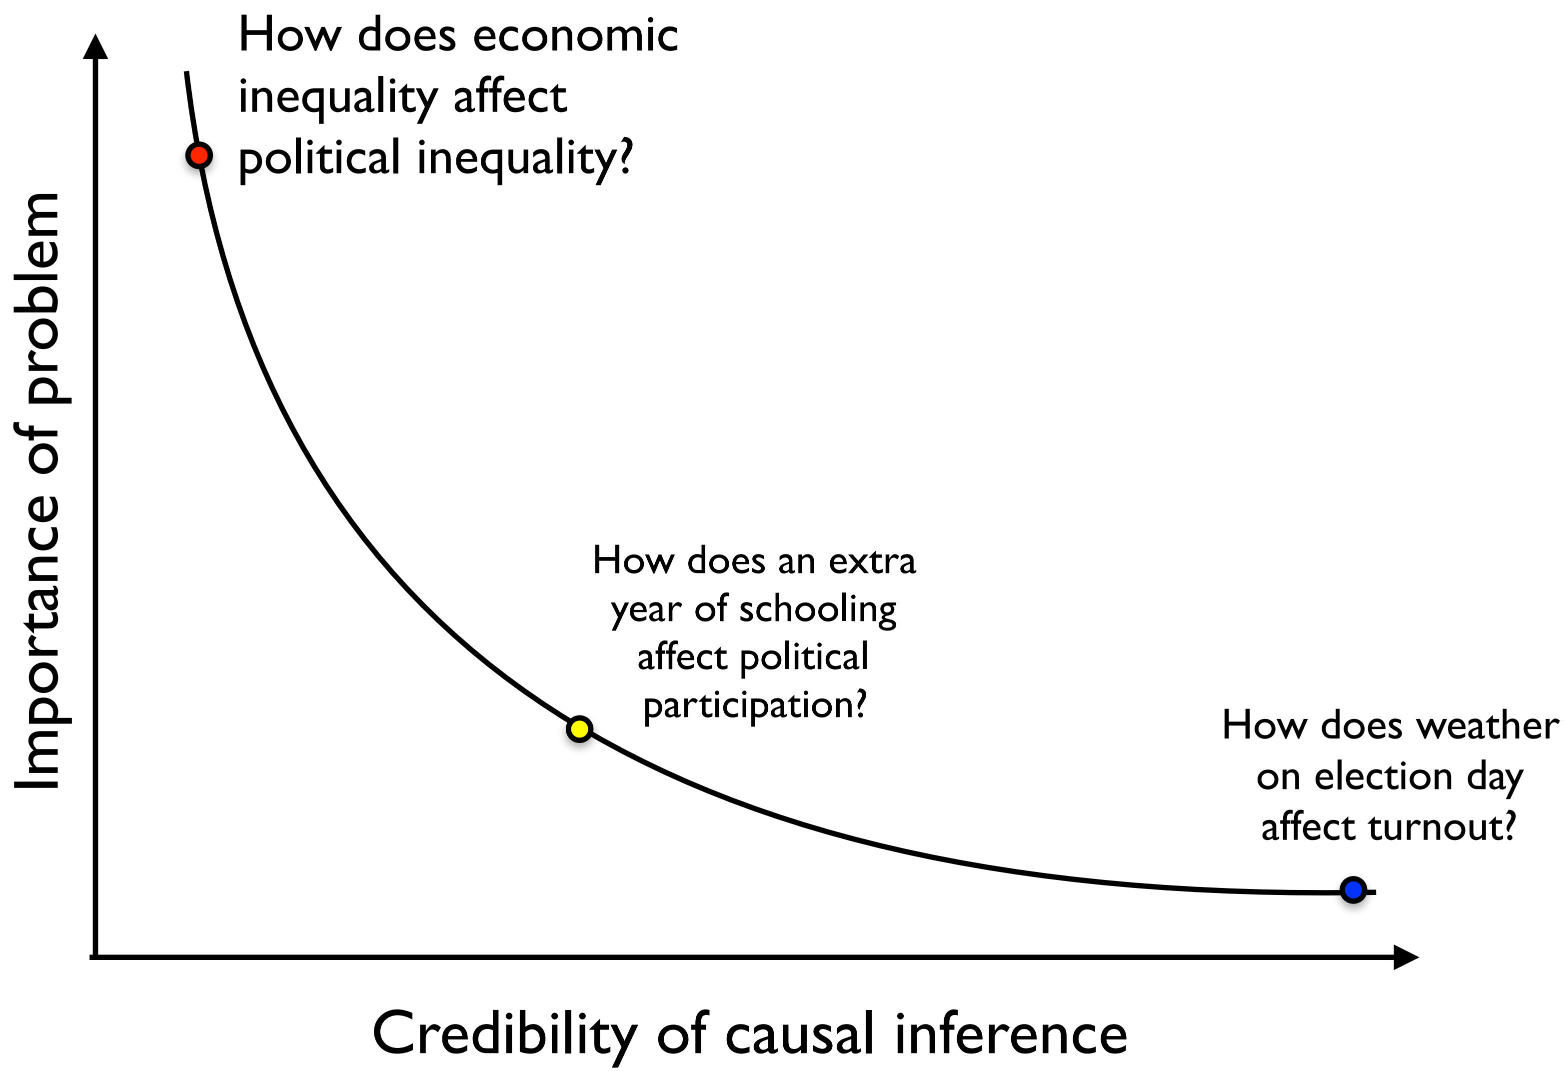

The credibility/importance tradeoff

Some underlying reasons:

- Internal validity requires special circumstances (thus less external validity)
- Many interesting "treatments" have very wide effects
- Ethically and practically, can't experiment on many interesting things
- Interesting treatments aren't determined arbitrarily

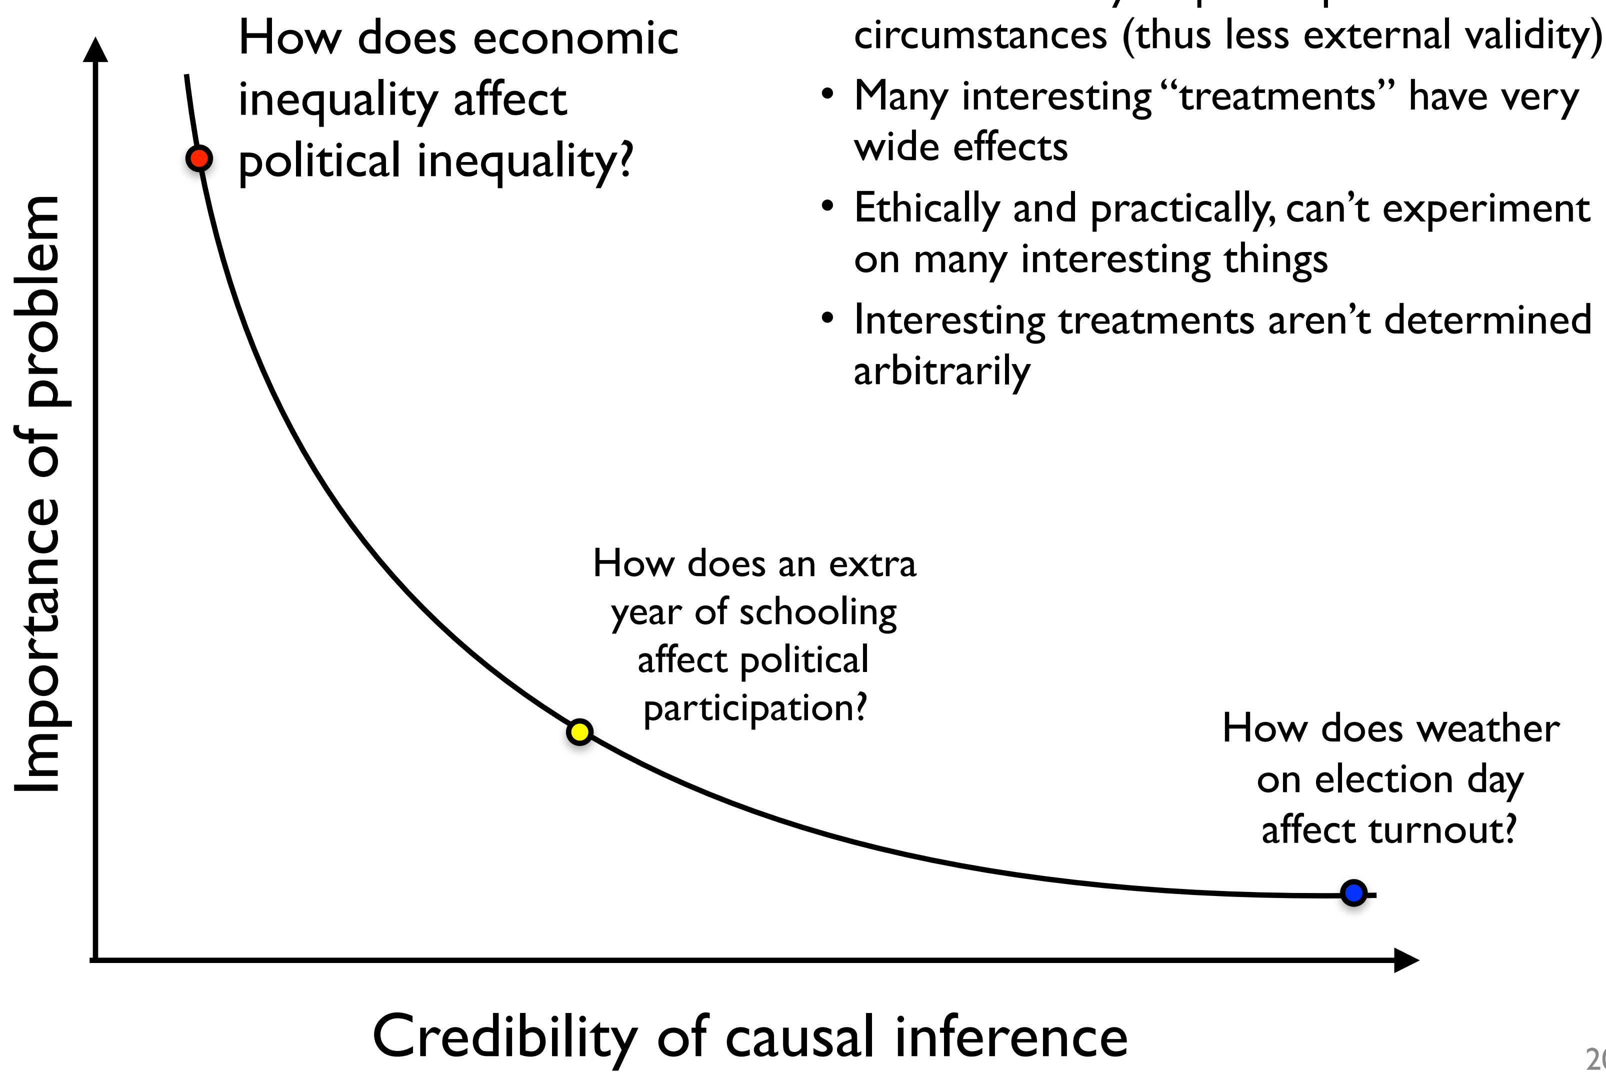

Please remember this advice if your research involves any claims about effects

Suppose there were no constraints (time, money, ethics, the number of countries). What is the most informative experiment I could run to measure the effect I want to study?

Benefits: Clarifies to you (and reader)

- what you are trying to study
- what challenges you face
- what feasible designs actually exist

Some big questions about design-based inference and the "credibility revolution"

- Does a study on elections in French villages tell us anything about national elections? (external validity)
- What about explanation? What does the study in French villages tell us about why turnout is higher in PR countries (the puzzle to be explained)?
- What about "theory-testing"? What theory is tested when the setting for our analysis was carefully chosen?
- What about "effects of causes" questions that can't be answered this way: what is the effect of globalization?

Design-based research and hypothesis testing

The basic hypothesis testing framework

- Generate a hypothesis and a test statistic (e.g. regression coefficient)
- Derive sampling distribution for test statistic under the null hypothesis (e.g. if true regression coefficient is zero)
- p-value indicates probability of getting a test statistic as extreme as observed estimate if null hypothesis actually true
- Convention: reject null hypothesis if p-value < .05

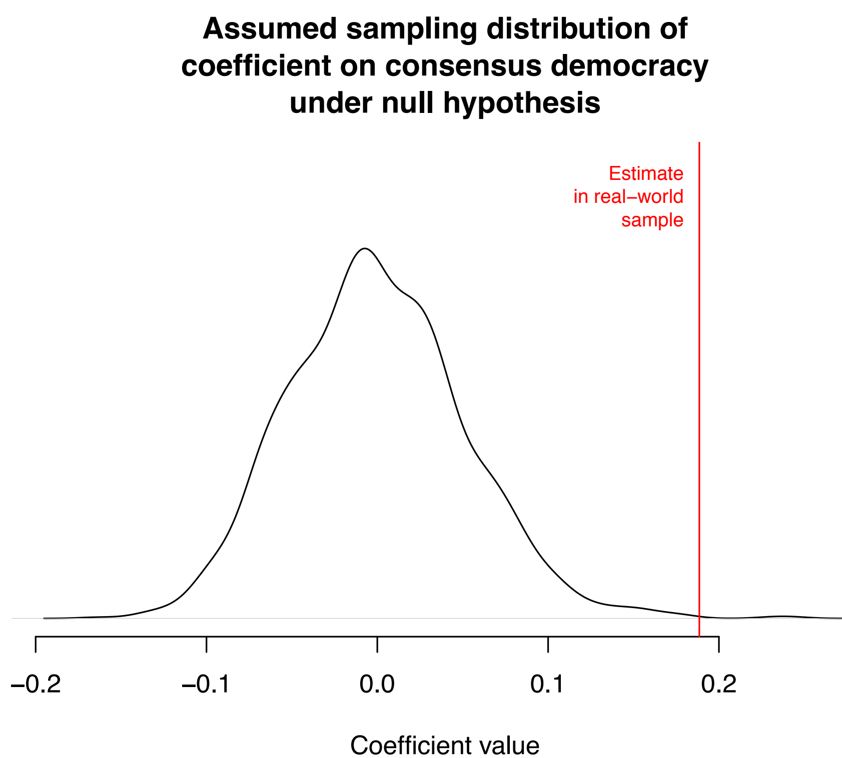

$$\alpha \equiv \Pr(\text{Reject null} | \text{Null is true}) = .05$$

Should we believe empirical claims in published research? (1)

What's the probability that the null hypothesis is *actually* false, given that the author rejects the null hypothesis in a statistical test?

Remember Bayes Theorem?

$$\Pr(a|b) = \frac{\Pr(b|a)\Pr(a)}{\Pr(b)}$$

$$\Pr(\text{Null is false}|\text{Reject null}) = \frac{\Pr(\text{Reject null}|\text{Null is false})\Pr(\text{Null is false})}{\Pr(\text{Reject null})}$$

Should we believe empirical claims in published research? (2)

What's the probability that the null hypothesis is *actually* false, given that the author rejects the null hypothesis in a statistical test?

$$\Pr(\text{Null is false}|\text{Reject null}) = \frac{\text{Power} \times \Pr(\text{Null is false})}{\text{Power} \times \Pr(\text{Null is false}) + \alpha(1 - \Pr(\text{Null is false}))}$$

Power: probability of **correctly** rejecting null

alpha: probability of **incorrectly** rejecting null

Implication: we should believe claim of **statistical significance** when

- alpha is low (standard *goal* is .05)
- Pr(Null is false) is high (i.e. the rejection is not surprising)

Should we believe empirical claims in published research? (3)

First challenge: Editors and reviewers require surprising results.

- *This gene causes turnout!*
- *The outcome of football games affects incumbent vote share!*
- *The disease environment during colonization affects current economic development (through institutions)!*

This makes published results (especially in top journals) less believable.

What can we do?

Should we believe empirical claims in published research? (4)

Second challenge: What is true value of alpha (probability of incorrectly rejecting the null hypothesis)?

We reject the null when the p-value is below .05. Is this the same as alpha = .05?

Consider:

- You run 20 regressions to pick your "preferred specification" (specification search)
- You don't pursue projects where, even with 20 regressions, you still get null results (file-drawer problem)

In practice, alpha might be **much** higher than .05!

Four kinds of "search" to worry about

Specification search: Having chosen an X and Y of interest and a setting, try various control variables, functional forms, etc until you find a significant relationship between X and Y

Outcome search: Having found a setting where X is quasi-randomly assigned, try various outcome variables Y until find a significant relationship

Treatment search: Having chosen a Y of interest and a setting, run a regression and choose your hypothesis based on what coefficients turn out to be significant/interesting

Subgroup search: Having found a setting where X is quasi-randomly assigned, try various subgroups (e.g. young Asian men) until find a significant relationship

Which of these is better with "credibility revolution"?
Which is worse?

Replication movement and DA-RT

http://www.dartstatement.org/

Petition to delay DA-RT implementation

Petition to Delay DA-RT Implementation

November 3, 2015 [list includes those who signed of November 8 5:15 pm EST]]

Dear Colleagues,

We write as concerned members of the American Political Science Association to urge an important amendment to the statement, "Data Access and Research Transparency (DA-RT): A Joint Statement by Political Science Journal Editors." In the joint statement, dated October 6, 2014, journal editors committed their respective journals to a set of principles, to be implemented by January 15, 2016.

DA-RT organizers have made many efforts over the past five years to reach out to members of the profession through various symposia and meetings. However, these issues began to gain widespread attention only when the journal editors signed the statement of October 6, 2014 and panels at the 2015 annual meeting of the American Political Science Association brought the issue to the attention of many scholars who had not realized the possible implications of that statement for their own research, despite the previous outreach activities. Conversations at the panels, roundtables, section business meetings, and other venues at the recent annual meeting demonstrated that members of the Association have only just begun to grapple with the

Pre-registration movement and EGAP

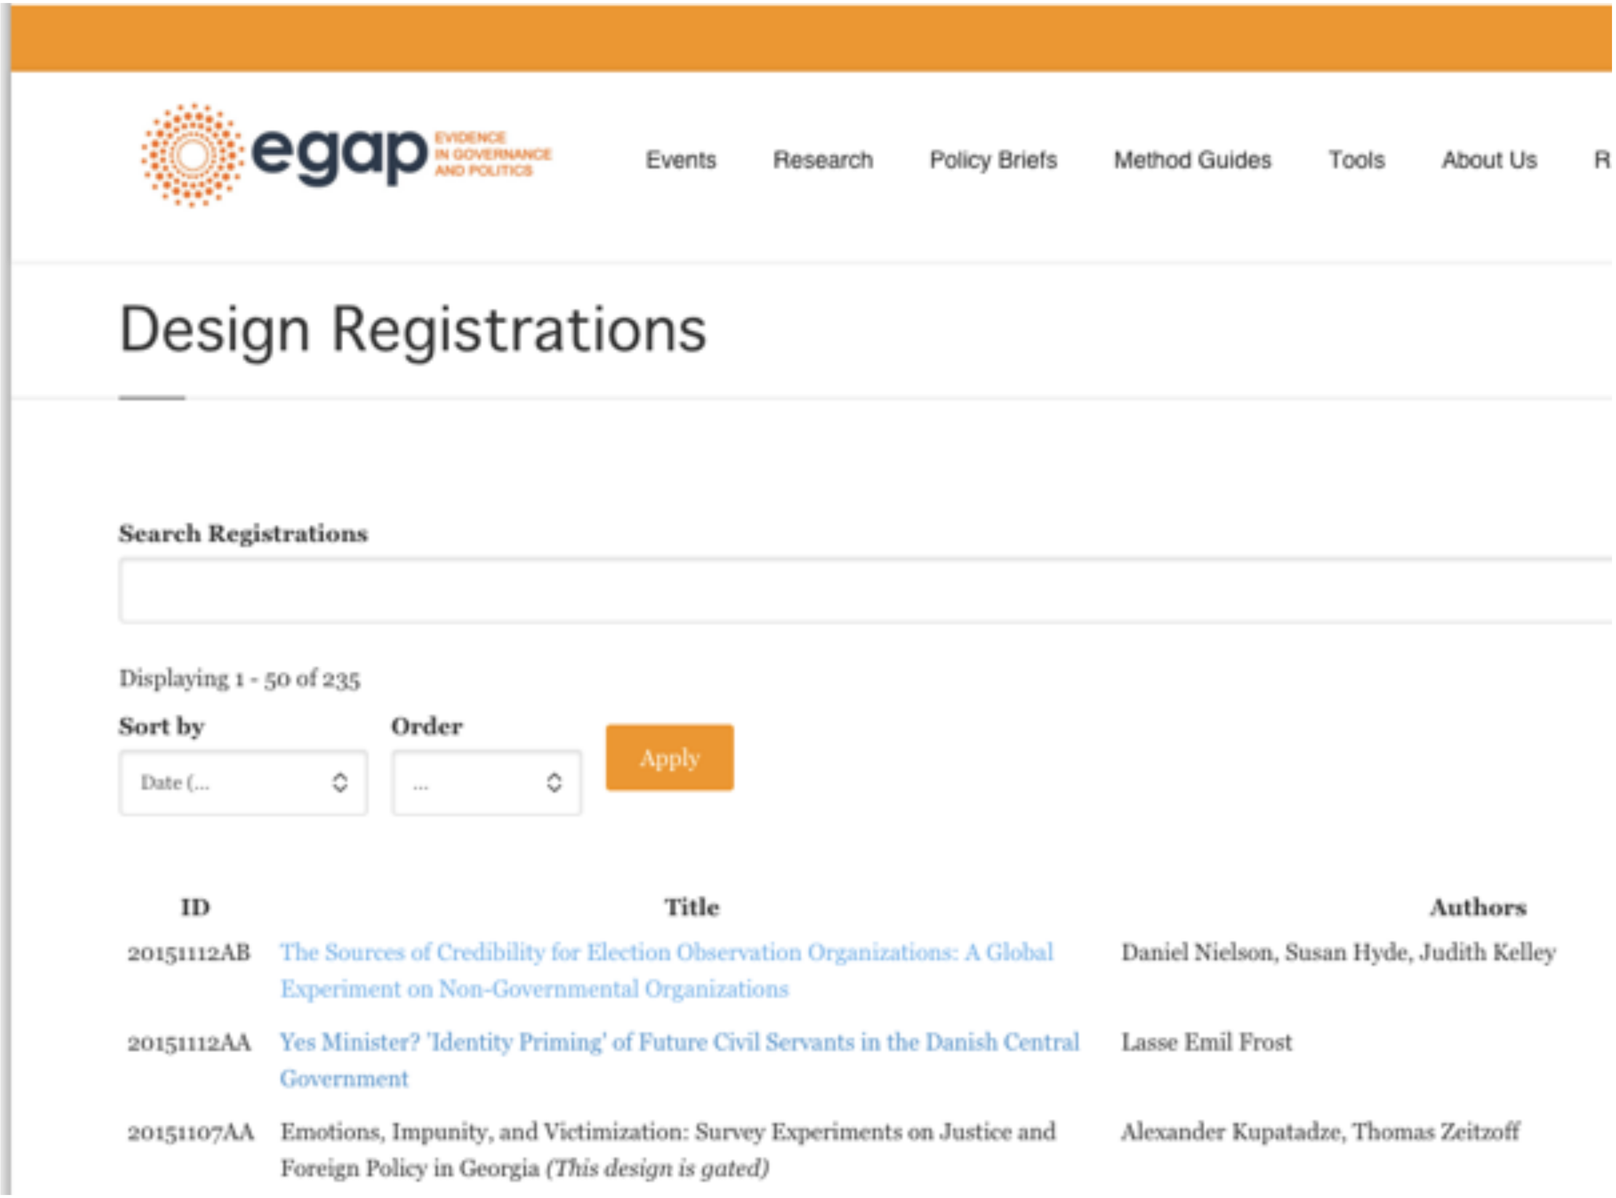

egap EVIDENCE IN GOVERNANCE AND POLITICS

Events Research Policy Briefs Method Guides Tools About Us R

Design Registrations

Search Registrations

Displaying 1 - 50 of 235

Sort by Date (...

Order ...

Apply

ID	Title	Authors
20151112AB	The Sources of Credibility for Election Observation Organizations: A Global Experiment on Non-Governmental Organizations	Daniel Nielson, Susan Hyde, Judith Kelley
20151112AA	Yes Minister? 'Identity Priming' of Future Civil Servants in the Danish Central Government	Lasse Emil Frost
20151107AA	Emotions, Impunity, and Victimization: Survey Experiments on Justice and Foreign Policy in Georgia *(This design is gated)*	Alexander Kupatadze, Thomas Zeitzoff