

Unpacking the Black Box of Causality: Learning about Causal Mechanisms from Experimental and Observational Studies

Kosuke Imai
Princeton University

Talk at the Center for Lifespan Psychology
Max Planck Institute for Human Development

April 28, 2015

Joint work with
Luke Keele Dustin Tingley Teppei Yamamoto

Project References (click the article titles)

- **General:**

- Unpacking the Black Box of Causality: Learning about Causal Mechanisms from Experimental and Observational Studies. *American Political Science Review*

- **Theory:**

- Identification, Inference, and Sensitivity Analysis for Causal Mediation Effects. *Statistical Science*

- **Extensions:**

- A General Approach to Causal Mediation Analysis. *Psychological Methods*
- Experimental Designs for Identifying Causal Mechanisms. *Journal of the Royal Statistical Society, Series A (with discussions)*
- Identification and Sensitivity Analysis for Multiple Causal Mechanisms: Revisiting Evidence from Framing Experiments. *Political Analysis*

- **Software:**

- mediation: R Package for Causal Mediation Analysis. *Journal of Statistical Software*

Identification of Causal Mechanisms

- Causal inference is a central goal of scientific research
- Scientists care about causal **mechanisms**, not just about causal effects
- Randomized experiments often only determine **whether** the treatment causes changes in the outcome
- Not **how** and **why** the treatment affects the outcome
- Common criticism of experiments and statistics:

black box view of causality

- Question: How can we learn about causal mechanisms from experimental and observational studies?

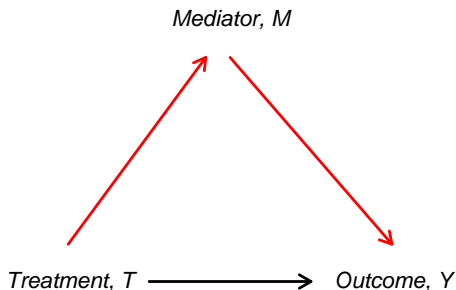
Overview of the Talk

Present a general framework for statistical analysis and research design strategies to understand causal mechanisms

- 1 Show that the **sequential ignorability** assumption is required to identify mechanisms even in experiments
- 2 Offer a flexible **estimation strategy** under this assumption
- 3 Introduce a **sensitivity analysis** to probe this assumption
- 4 Illustrate how to use **statistical software** `mediation`
- 5 Consider **research designs** that relax sequential ignorability

Causal Mediation Analysis

- Graphical representation

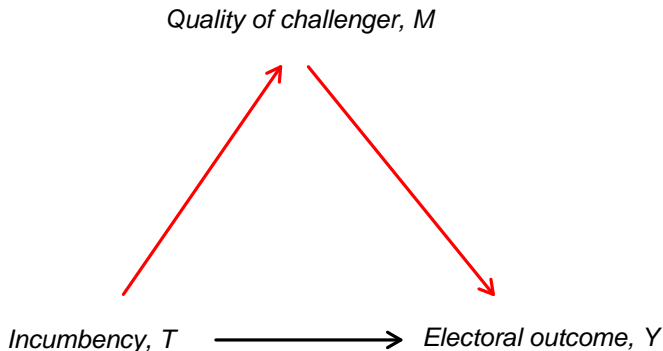


- Goal is to decompose total effect into direct and indirect effects
- Alternative approach: decompose the treatment into different components
- Causal mediation analysis as **quantitative process tracing**

Decomposition of Incumbency Advantage

- Incumbency effects: one of the most studied topics in American politics
- Consensus emerged in 1980s: incumbency advantage is positive and growing in magnitude
- New direction in 1990s: Where does incumbency advantage come from?
- **Scare-off/quality effect** (Cox and Katz): the ability of incumbents to deter high-quality challengers from entering the race
- Alternative causal mechanisms: name recognition, campaign spending, personal vote, television, etc.

Causal Mediation Analysis in Cox and Katz




- How much of incumbency advantage can be explained by scare-off/quality effect?
- How large is the mediation effect relative to the total effect?

Psychological Study of Media Effects

- Large literature on how media influences public opinion
- A media framing experiment of Brader *et al.*:
 - ① (White) Subjects read a mock news story about immigration:
 - Treatment: Hispanic immigrant in the story
 - Control: European immigrant in the story
 - ② Measure attitudinal and behavioral outcome variables:
 - Opinions about increasing or decrease immigration
 - Contact legislator about the issue
 - Send anti-immigration message to legislator
- Why is group-based media framing effective?: role of emotion
- Hypothesis: Hispanic immigrant increases anxiety, leading to greater opposition to immigration
- The primary goal is to examine how, not whether, media framing shapes public opinion

Causal Mediation Analysis in Brader *et al.*



- Does the media framing shape public opinion by making people anxious?
- An alternative causal mechanism: change in beliefs
- Can we identify mediation effects from randomized experiments?

The Standard Estimation Method

- Linear models for mediator and outcome:

$$Y_i = \alpha_1 + \beta_1 T_i + \xi_1^\top X_i + \epsilon_{1i}$$

$$M_i = \alpha_2 + \beta_2 T_i + \xi_2^\top X_i + \epsilon_{2i}$$

$$Y_i = \alpha_3 + \beta_3 T_i + \gamma M_i + \xi_3^\top X_i + \epsilon_{3i}$$

where X_i is a set of pre-treatment or control variables

- 1 Total effect (ATE) is β_1
 - 2 Direct effect is β_3
 - 3 Indirect or mediation effect is $\beta_2\gamma$
 - 4 **Effect decomposition**: $\beta_1 = \beta_3 + \beta_2\gamma$.
- Some motivating questions:
 - 1 What should we do when we have interaction or nonlinear terms?
 - 2 What about other models such as logit?
 - 3 In general, under what conditions can we interpret β_1 and $\beta_2\gamma$ as causal effects?
 - 4 What do we really mean by causal mediation effect anyway?

Potential Outcomes Framework of Causal Inference

- Observed data:
 - Binary treatment: $T_i \in \{0, 1\}$
 - Mediator: $M_i \in \mathcal{M}$
 - Outcome: $Y_i \in \mathcal{Y}$
 - Observed pre-treatment covariates: $X_i \in \mathcal{X}$
- Potential outcomes model (Neyman, Rubin):
 - Potential mediators: $M_i(t)$ where $M_i = M_i(T_i)$
 - Potential outcomes: $Y_i(t, m)$ where $Y_i = Y_i(T_i, M_i(T_i))$

- **Total causal effect:**

$$\tau_i \equiv Y_i(1, M_i(1)) - Y_i(0, M_i(0))$$

- **Fundamental problem of causal inference:** only one potential outcome can be observed for each i

Back to the Examples

- $M_i(1)$:
 - ① Quality of her challenger if politician i is an incumbent
 - ② Level of anxiety individual i would report if he reads the story with Hispanic immigrant
- $Y_i(1, M_i(1))$:
 - ① Election outcome that would result if politician i is an incumbent and faces a challenger whose quality is $M_i(1)$
 - ② Immigration attitude individual i would report if he reads the story with Hispanic immigrant and reports the anxiety level $M_i(1)$
- $M_i(0)$ and $Y_i(0, M_i(0))$ are the converse

Causal Mediation Effects

- Causal mediation (Indirect) effects:

$$\delta_i(t) \equiv Y_i(t, M_i(1)) - Y_i(t, M_i(0))$$

- Causal effect of the change in M_i on Y_i that would be induced by treatment
- Change the mediator from $M_i(0)$ to $M_i(1)$ while holding the treatment constant at t
- Represents the mechanism through M_i
- Zero treatment effect on mediator \implies Zero mediation effect
- Examples:
 - 1 Part of incumbency advantage that is due to the difference in challenger quality induced by incumbency status
 - 2 Difference in immigration attitudes that is due to the change in anxiety induced by the treatment news story

Total Effect = Indirect Effect + Direct Effect

- **Direct effects:**

$$\zeta_i(t) \equiv Y_i(1, M_i(t)) - Y_i(0, M_i(t))$$

- Causal effect of T_i on Y_i , holding mediator constant at its potential value that would realize when $T_i = t$
- Change the treatment from 0 to 1 while holding the mediator constant at $M_i(t)$
- Represents all mechanisms other than through M_i
- Total effect = mediation (indirect) effect + direct effect:

$$\tau_i = \delta_i(t) + \zeta_i(1 - t) = \frac{1}{2} \{(\delta_i(0) + \zeta_i(0)) + (\delta_i(1) + \zeta_i(1))\}$$

Mechanisms

- **Indirect effects:** $\delta_i(t) \equiv Y_i(t, M_i(1)) - Y_i(t, M_i(0))$
- Counterfactuals about treatment-induced mediator values

Manipulations

- **Controlled direct effects:** $\xi_i(t, m, m') \equiv Y_i(t, m) - Y_i(t, m')$
- Causal effect of directly manipulating the mediator under $T_i = t$

Interactions

- **Interaction effects:** $\xi(1, m, m') - \xi(0, m, m')$
- The extent to which controlled direct effects vary by the treatment

What Does the Observed Data Tell Us?

- Recall the Brader *et al.* experimental design:
 - ① randomize T_i
 - ② measure M_i and then Y_i
- Among observations with $T_i = t$, we observe $Y_i(t, M_i(t))$ but not $Y_i(t, M_i(1 - t))$ unless $M_i(t) = M_i(1 - t)$
- But we want to estimate

$$\delta_i(t) \equiv Y_i(t, M_i(1)) - Y_i(t, M_i(0))$$

- For $t = 1$, we observe $Y_i(1, M_i(1))$ but not $Y_i(1, M_i(0))$
- Similarly, for $t = 0$, we observe $Y_i(0, M_i(0))$ but not $Y_i(0, M_i(1))$
- We have the **identification problem** \implies Need assumptions or better research designs

Counterfactuals in the Examples

1 Incumbency advantage:

- An incumbent ($T_i = 1$) faces a challenger with quality $M_i(1)$
- We observe the electoral outcome $Y_i = Y_i(1, M_i(1))$
- We also want $Y_i(1, M_i(0))$ where $M_i(0)$ is the quality of challenger this incumbent politician would face if she is not an incumbent

2 Media framing effects:

- A subject viewed the news story with Hispanic immigrant ($T_i = 1$)
- For this person, $Y_i(1, M_i(1))$ is the observed immigration opinion
- $Y_i(1, M_i(0))$ is his immigration opinion in the counterfactual world where he still views the story with Hispanic immigrant but his anxiety is at the same level as if he viewed the control news story

In both cases, we can't observe $Y_i(1, M_i(0))$ because $M_i(0)$ is not realized when $T_i = 1$

Sequential Ignorability Assumption

- Proposed identification assumption: **Sequential Ignorability** (SI)

$$\{Y_i(t', m), M_i(t)\} \perp\!\!\!\perp T_i \mid X_i = x, \quad (1)$$

$$Y_i(t', m) \perp\!\!\!\perp M_i(t) \mid T_i = t, X_i = x \quad (2)$$

- In words,

- T_i is (as-if) randomized conditional on $X_i = x$
- $M_i(t)$ is (as-if) randomized conditional on $X_i = x$ and $T_i = t$

- Important limitations:

- In a standard experiment, (1) holds but (2) may not
- X_i needs to include all confounders
- X_i must be pre-treatment confounders \implies post-treatment confounder is not allowed
- Randomizing M_i via manipulation is not the same as assuming $M_i(t)$ is as-if randomized

Sequential Ignorability in the Standard Experiment

Back to Brader *et al.*:

- Treatment is randomized \implies (1) is satisfied
- But (2) may not hold:
 - ① Pre-treatment confounder or X_i : state of residence
those who live in AZ tend to have higher levels of perceived harm and be opposed to immigration
 - ② Post-treatment confounder: alternative mechanism
beliefs about the likely negative impact of immigration makes people anxious
- Pre-treatment confounders \implies measure and adjust for them
- Post-treatment confounders \implies adjusting is not sufficient

Nonparametric Identification

Under SI, both ACME and average direct effects are **nonparametrically identified** (can be consistently estimated without modeling assumption)

- ACME $\bar{\delta}(t)$

$$\int \int \mathbb{E}(Y_i | M_i, T_i = t, X_i) \{dP(M_i | T_i = 1, X_i) - dP(M_i | T_i = 0, X_i)\} dP(X_i)$$

- Average direct effects $\bar{\zeta}(t)$

$$\int \int \{\mathbb{E}(Y_i | M_i, T_i = 1, X_i) - \mathbb{E}(Y_i | M_i, T_i = 0, X_i)\} dP(M_i | T_i = t, X_i) dP(X_i)$$

Implies the general **mediation formula** under any statistical model

Traditional Estimation Methods: LSEM

- **Linear structural equation model (LSEM):**

$$\begin{aligned}M_i &= \alpha_2 + \beta_2 T_i + \xi_2^\top X_i + \epsilon_{i2}, \\Y_i &= \alpha_3 + \beta_3 T_i + \gamma M_i + \xi_3^\top X_i + \epsilon_{i3}.\end{aligned}$$

- Fit two least squares regressions separately
- Use **product of coefficients** ($\hat{\beta}_2 \hat{\gamma}$) to estimate ACME
- Use asymptotic variance to test significance (Sobel test)
- Under SI and the **no-interaction assumption** ($\bar{\delta}(1) \neq \bar{\delta}(0)$), $\hat{\beta}_2 \hat{\gamma}$ consistently estimates ACME
- Can be extended to LSEM with interaction terms
- Problem: Only valid for the simplest LSEM

Popular Baron-Kenny Procedure

- The procedure:
 - ① Regress Y on T and show a significant relationship
 - ② Regress M on T and show a significant relationship
 - ③ Regress Y on M and T , and show a significant relationship between Y and M

- The problems:
 - ① First step can lead to false negatives especially if indirect and direct effects in opposite directions
 - ② The procedure only anticipates simplest linear models
 - ③ Don't do star-gazing. Report quantities of interest

Proposed General Estimation Algorithm

- 1 Model outcome and mediator
 - Outcome model: $p(Y_i | T_i, M_i, X_i)$
 - Mediator model: $p(M_i | T_i, X_i)$
 - These models can be of **any form** (linear or nonlinear, semi- or nonparametric, with or without interactions)
- 2 Predict mediator for both treatment values ($M_i(1), M_i(0)$)
- 3 Predict outcome by first setting $T_i = 1$ and $M_i = M_i(0)$, and then $T_i = 1$ and $M_i = M_i(1)$
- 4 Compute the average difference between two outcomes to obtain a consistent estimate of ACME
- 5 Monte-Carlo or bootstrap to estimate uncertainty

Example: Binary Mediator and Outcome

- Two logistic regression models:

$$\begin{aligned}\Pr(M_i = 1 \mid T_i, X_i) &= \text{logit}^{-1}(\alpha_2 + \beta_2 T_i + \xi_2^\top X_i) \\ \Pr(Y_i = 1 \mid T_i, M_i, X_i) &= \text{logit}^{-1}(\alpha_3 + \beta_3 T_i + \gamma M_i + \xi_3^\top X_i)\end{aligned}$$

- Can't multiply β_2 by γ
- Difference of coefficients $\beta_1 - \beta_3$ doesn't work either

$$\Pr(Y_i = 1 \mid T_i, X_i) = \text{logit}^{-1}(\alpha_1 + \beta_1 T_i + \xi_1^\top X_i)$$

- Can use our algorithm (example: $\mathbb{E}\{Y_i(1, M_i(0))\}$)
 - 1 Predict $M_i(0)$ given $T_i = 0$ using the first model
 - 2 Compute $\Pr(Y_i(1, M_i(0)) = 1 \mid T_i = 1, M_i = \widehat{M}_i(0), X_i)$ using the second model

Sensitivity Analysis

- Standard experiments require sequential ignorability to identify mechanisms
- The sequential ignorability assumption is often too strong
- Need to assess the robustness of findings via sensitivity analysis
- **Question:** How large a departure from the key assumption must occur for the conclusions to no longer hold?
- Parametric sensitivity analysis by assuming

$$\{Y_i(t', m), M_i(t)\} \perp\!\!\!\perp T_i \mid X_i = x$$

but not

$$Y_i(t', m) \perp\!\!\!\perp M_i(t) \mid T_i = t, X_i = x$$

- Possible existence of unobserved *pre-treatment* confounder

Parametric Sensitivity Analysis

- **Sensitivity parameter:** $\rho \equiv \text{Corr}(\epsilon_{i2}, \epsilon_{i3})$
- Sequential ignorability implies $\rho = 0$
- Set ρ to different values and see how ACME changes

- **Result:**

$$\bar{\delta}(0) = \bar{\delta}(1) = \frac{\beta_2 \sigma_1}{\sigma_2} \left\{ \tilde{\rho} - \rho \sqrt{(1 - \tilde{\rho}^2)/(1 - \rho^2)} \right\},$$

where $\sigma_j^2 \equiv \text{var}(\epsilon_{ij})$ for $j = 1, 2$ and $\tilde{\rho} \equiv \text{Corr}(\epsilon_{i1}, \epsilon_{i2})$.

- When do my results go away completely?
- $\bar{\delta}(t) = 0$ if and only if $\rho = \tilde{\rho}$
- Easy to estimate from the regression of Y_i on T_i :

$$Y_i = \alpha_1 + \beta_1 T_i + \epsilon_{i1}$$

Interpreting Sensitivity Analysis with R squares

- Interpreting ρ : how small is too small?
- An unobserved (pre-treatment) confounder formulation:

$$\epsilon_{i2} = \lambda_2 U_i + \epsilon'_{i2} \quad \text{and} \quad \epsilon_{i3} = \lambda_3 U_i + \epsilon'_{i3}$$

- How much does U_i have to explain for our results to go away?
- Sensitivity parameters: **R squares**
 - 1 Proportion of **previously unexplained variance** explained by U_i

$$R_M^{2*} \equiv 1 - \frac{\text{var}(\epsilon'_{i2})}{\text{var}(\epsilon_{i2})} \quad \text{and} \quad R_Y^{2*} \equiv 1 - \frac{\text{var}(\epsilon'_{i3})}{\text{var}(\epsilon_{i3})}$$

- 2 Proportion of **original variance** explained by U_i

$$\tilde{R}_M^2 \equiv \frac{\text{var}(\epsilon_{i2}) - \text{var}(\epsilon'_{i2})}{\text{var}(M_i)} \quad \text{and} \quad \tilde{R}_Y^2 \equiv \frac{\text{var}(\epsilon_{i3}) - \text{var}(\epsilon'_{i3})}{\text{var}(Y_i)}$$

- Then reparameterize ρ using (R_M^{2*}, R_Y^{2*}) (or $(\tilde{R}_M^2, \tilde{R}_Y^2)$):

$$\rho = \text{sgn}(\lambda_2 \lambda_3) R_M^* R_Y^* = \frac{\text{sgn}(\lambda_2 \lambda_3) \tilde{R}_M \tilde{R}_Y}{\sqrt{(1 - R_M^2)(1 - R_Y^2)}},$$

where R_M^2 and R_Y^2 are from the original mediator and outcome models


- $\text{sgn}(\lambda_2 \lambda_3)$ indicates the direction of the effects of U_i on Y_i and M_i
- Set (R_M^{2*}, R_Y^{2*}) (or $(\tilde{R}_M^2, \tilde{R}_Y^2)$) to different values and see how mediation effects change

Reanalysis: Estimates under Sequential Ignorability

- Original method: **Product of coefficients** with the **Sobel test**
 - Valid only when both models are linear w/o T - M interaction (which they are not)
- Our method: Calculate ACME using our general algorithm


Outcome variables	Product of Coefficients	Average Causal Mediation Effect (δ)
Decrease Immigration $\bar{\delta}(1)$.347 [0.146, 0.548]	.105 [0.048, 0.170]
Support English Only Laws $\bar{\delta}(1)$.204 [0.069, 0.339]	.074 [0.027, 0.132]
Request Anti-Immigration Information $\bar{\delta}(1)$.277 [0.084, 0.469]	.029 [0.007, 0.063]
Send Anti-Immigration Message $\bar{\delta}(1)$.276 [0.102, 0.450]	.086 [0.035, 0.144]

Reanalysis: Sensitivity Analysis w.r.t. ρ



- ACME > 0 as long as the error correlation is less than 0.39 (0.30 with 95% CI)


Reanalysis: Sensitivity Analysis w.r.t. \tilde{R}_M^2 and \tilde{R}_Y^2




- An unobserved confounder can account for up to 26.5% of the variation in both Y_i and M_i before ACME becomes zero

Open-Source Software “Mediation”

Model-Based Inference



Design-Based Inference



Implementation Examples

- 1 Fit models for the mediator and outcome variable and store these models

```
> m <- lm(Mediator ~ Treat + X)
> y <- lm(Y ~ Treat + Mediator + X)
```

- 2 **Mediation analysis:** Feed model objects into the `mediate()` function. Call a summary of results

```
> m.out <- mediate(m, y, treat = "Treat",
                  mediator = "Mediator")
> summary(m.out)
```

- 3 **Sensitivity analysis:** Feed the output into the `medsens()` function. Summarize and plot

```
> s.out <- medsens(m.out)
> summary(s.out)
> plot(s.out, "rho")
> plot(s.out, "R2")
```

Data Types Available via **mediate**

<i>Mediator Model Types</i>	<i>Outcome Model Types</i>						
	Linear	GLM	Ordered	Censored	Quantile	GAM	Survival
Linear (<code>lm/lmer</code>)	✓	✓	✓*	✓	✓	✓*	✓
GLM (<code>glm/bayesglm/glmer</code>)	✓	✓	✓*	✓	✓	✓*	✓
Ordered (<code>polr/bayespolr</code>)	✓	✓	✓*	✓	✓	✓*	✓
Censored (<code>tobit</code> via <code>vglm</code>)	-	-	-	-	-	-	-
Quantile (<code>rq</code>)	✓*	✓*	✓*	✓*	✓*	✓*	✓
GAM (<code>gam</code>)	✓*	✓*	✓*	✓*	✓*	✓*	✓*
Survival (<code>survreg</code>)	✓	✓	✓*	✓	✓	✓*	✓

Types of Models That Can be Handled by `mediate`. Stars (*) indicate the model combinations that can only be estimated using the nonparametric bootstrap (i.e. with `boot = TRUE`).

Additional Features

- Treatment/mediator interactions, with formal statistical tests
- Treatment/mediator/pre-treatment interactions and reporting of quantities by pre-treatment values
- Factoral, continuous treatment variables
- Cluster standard errors/adjustable CI reporting/p-values
- Support for multiple imputation
- Multiple mediators
- Multilevel mediation (**NEW!**)

Please read our vignette file **here**.

Causal Mediation Analysis in Stata

Based on the same algorithm

Hicks, R, Tingley D. 2011. Causal Mediation Analysis. Stata Journal. 11(4):609-615.

```
ssc install mediation
```

More limited coverage of models (just bc. of time though!)

Syntax: medeff

```
medeff (equation 1) (equation 2) [if] [in] [[weight]] ,  
[sims(integer) seed(integer) vce(vcetype) Level(#)  
interact(varname)] mediate(varname) treat(varname)
```

Where “equation 1” or “equation 2” are of the form (For equation 1, the mediator equation):


```
probit M T x
```

or

```
regress M T x
```

Beyond Sequential Ignorability

- Without sequential ignorability, standard experimental design lacks identification power
- Even the sign of ACME is not identified
- Need to develop **alternative experimental designs** for more credible inference
- Possible when the mediator can be directly or indirectly manipulated
- All proposed designs preserve the ability to estimate the ACME under the SI assumption
- Trade-off: statistical power
- These experimental designs can then be extended to natural experiments in observational studies



- Must assume **no direct effect of manipulation** on outcome
- More informative than standard single experiment
- If we assume no $T-M$ interaction, ACME is point identified

Why Do We Need No-Interaction Assumption?


- Numerical Example:

Prop.	$M_i(1)$	$M_i(0)$	$Y_i(t, 1)$	$Y_i(t, 0)$	$\delta_i(t)$
0.3	1	0	0	1	-1
0.3	0	0	1	0	0
0.1	0	1	0	1	1
0.3	1	1	1	0	0

- $\mathbb{E}(M_i(1) - M_i(0)) = \mathbb{E}(Y_i(t, 1) - Y_i(t, 0)) = 0.2$, but $\bar{\delta}(t) = -0.2$
- The Problem: Causal effect heterogeneity
 - T increases M only *on average*
 - M increases Y only *on average*
 - $T - M$ interaction: Many of those who have a positive effect of T on M have a negative effect of M on Y (first row)
- A solution: sensitivity analysis (see Imai and Yamamoto, 2013)
- Pitfall of “mechanism experiments” or “causal chain approach”

Encouragement Design

- Direct manipulation of mediator is difficult in most situations
- Use an **instrumental variable** approach:



- Advantage: allows for unobserved confounder between M and Y
- Key Assumptions:
 - 1 Z is randomized or as-if random
 - 2 No direct effect of Z on Y (a.k.a. exclusion restriction)

Example: Social Norm Experiment on Property Taxes

- Lucia Del Carpio. “Are Neighbors Cheating?”
- Treatment: informing average rate of compliance
- Outcome: compliance rate obtained from administrative records
- Large positive effect on compliance rate \approx 20 percentage points
- Mediators:
 - ① social norm (not measured; direct effect)
 - ② M_1 : beliefs about compliance (measured)
 - ③ M_2 : beliefs about enforcement (measured)
- Instruments:
 - ① Z_1 : informing average rate of enforcement
 - ② Z_2 : payment-reminder
- Assumptions:
 - ① Z_1 affects Y only through M_1 and M_2
 - ② Z_2 affects Y only through M_1
- Results:
 - Average direct effect is estimated to be large
 - The author interprets this effect as the effect of social norm

Crossover Design

- Recall ACME can be identified if we observe $Y_i(t', M_i(t))$
- Get $M_i(t)$, then switch T_i to t' while holding $M_i = M_i(t)$
- **Crossover design:**
 - ① Round 1: Conduct a standard experiment
 - ② Round 2: Change the treatment to the opposite status but fix the mediator to the value observed in the first round
- Very powerful – identifies mediation effects for each subject
- Must assume **no carryover effect**: Round 1 must not affect Round 2
- Can be made plausible by design

Example: Labor Market Discrimination

EXAMPLE Bertrand & Mullainathan (2004, AER)

- Treatment: Black vs. White names on CVs
- Mediator: Perceived qualifications of applicants
- Outcome: Callback from employers

- Quantity of interest: Direct effects of (perceived) race
- Would Jamal get a callback if his name were Greg but his qualifications stayed the same?

- Round 1: Send Jamal's actual CV and record the outcome
- Round 2: Send his CV as Greg and record the outcome

- Assumption: their different names do not change the perceived qualifications of applicants
- Under this assumption, the direct effect can be interpreted as blunt racial discrimination

Designing Observational Studies

- Key difference between experimental and observational studies: treatment assignment
- Sequential ignorability:
 - ① Ignorability of treatment given covariates
 - ② Ignorability of mediator given treatment and covariates
- Both (1) and (2) are suspect in observational studies
- Statistical control: matching, propensity scores, etc.
- Search for quasi-randomized treatments: “natural” experiments
- How can we design observational studies?
- Experiments can serve as templates for observational studies

Cross-over Design in Observational Studies

EXAMPLE Back to incumbency advantage

- Use of cross-over design (Levitt and Wolfram)
 - ① 1st Round: two non-incumbents in an open seat
 - ② 2nd Round: same candidates with one being an incumbent
- Assume challenger quality (mediator) stays the same
- Estimation of direct effect is possible

- Redistricting as natural experiments (Ansolabehere et al.)
 - ① 1st Round: incumbent in the old part of the district
 - ② 2nd Round: incumbent in the new part of the district
- Challenger quality is the same but treatment is different
- Estimation of direct effect is possible

Multiple Mediators



- Quantity of interest = The average indirect effect with respect to M
- W represents the alternative observed mediators
- Left: Assumes **independence** between the two mechanisms
- Right: Allows M to be affected by the other mediators W
- Applied work often assumes the independence of mechanisms
- Under this independence assumption, one can apply the same analysis as in the single mediator case
- For causally dependent mediators, we must deal with the heterogeneity in the $T \times M$ interaction as done under the parallel design \implies sensitivity analysis

Concluding Remarks

- Even in a randomized experiment, a strong assumption is needed to identify causal mechanisms
- However, progress can be made toward this fundamental goal of scientific research with modern statistical tools
- A general, flexible estimation method is available once we assume sequential ignorability
- Sequential ignorability can be probed via sensitivity analysis
- More credible inferences are possible using clever experimental designs
- Insights from new experimental designs can be directly applied when designing observational studies
- Multiple mediators require additional care when they are causally dependent

The project website for papers and software:

<http://imai.princeton.edu/projects/mechanisms.html>

Email for questions and suggestions:

kimai@princeton.edu