Transfer Learning Between U.S. Presidential Elections

How much can we learn from a 2020 ad campaign to inform 2024 elections?

Xinran Miao, Jiwei Zhao, Hyunseung Kang

University of Wisconsin-Madison

American Causal Inference Conference (ACIC)
May 15, 2024, Seattle, Washington



Motivation: 2020 Digital Ad Campaign [Aggarwal et al. (2023) Nature Human Behavior]

Motivation: 2020 Digital Ad Campaign [Aggarwal et al. (2023) Nature Human Behavior]

2020 Digital Ad Campaign in U.S. Presidential Election: How would online ads against Donald Trump affect voter turnout in five battleground states: AZ, MI, NC, PA, and WI?

Motivation: 2020 Digital Ad Campaign [Aggarwal et al. (2023) Nature Human Behavior]

2020 Digital Ad Campaign in U.S. Presidential Election: How would online ads against Donald Trump affect voter turnout in five battle-ground states: AZ, MI, NC, PA, and WI?

- A stratified randomized experiment from Feb. 2020 to Nov. 2020 on nearly 2 million voters
    - Treatment group: an average of 754 ads against Trump by Acronym
    - Control group: no ads from Acronym
    - Outcome: voted in 2020 U.S. election? (**binary**)

Motivation: 2020 Digital Ad Campaign [Aggarwal et al. (2023) Nature Human Behavior]

2020 Digital Ad Campaign in U.S. Presidential Election: How would online ads against Donald Trump affect voter turnout in five battle-ground states: AZ, MI, NC, PA, and WI?

- A stratified randomized experiment from Feb. 2020 to Nov. 2020 on nearly 2 million voters
  - Treatment group: an average of 754 ads against Trump by Acronym
  - Control group: no ads from Acronym
  - Outcome: voted in 2020 U.S. election? (**binary**)
- Effect: difference in voter turnout between treated and control groups
- Estimate: $\widehat{\text{ATE}} = -0.06\%$, $\widehat{\text{SE}} = 0.12\%$ (**insignificant**)

Motivation: 2020 Digital Ad Campaign [Aggarwal et al. (2023) Nature Human Behavior]

2020 Digital Ad Campaign in U.S. Presidential Election: How would online ads against Donald Trump affect voter turnout in five battleground states: AZ, MI, NC, PA, and WI?

- A stratified randomized experiment from Feb. 2020 to Nov. 2020 on nearly 2 million voters
  - Treatment group: an average of 754 ads against Trump by Acronym
  - Control group: no ads from Acronym
  - Outcome: voted in 2020 U.S. election? (**binary**)
- Effect: difference in voter turnout between treated and control groups
- Estimate: $\widehat{\text{ATE}} = -0.06\%$, $\widehat{\text{SE}} = 0.12\%$ (**insignificant**)

*"One reasonable question... is how well **our findings would generalize... to other electoral contexts**... it could be that the 2020 election was exceptional because of COVID... perhaps digital advertising would have larger effects in more typical settings..."*

This Talk: Would the Ad Against Trump Remain Ineffective in 2024?

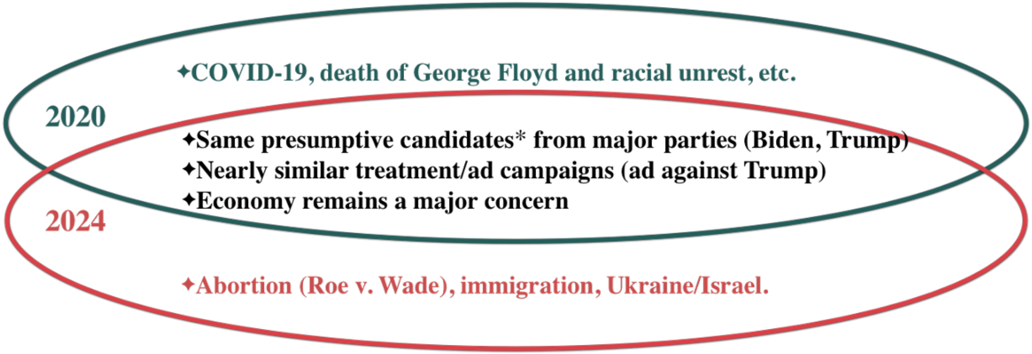

Figure. Similarities and differences between U.S. presidential elections in 2024 and 2020 (source: Pew Research Center).

*: Last time a rematch occurred was in 1956 (Eisenhower and Stevenson).

Our Setup

- Estimand: Ad effect in 2024,

$$\theta = \mathrm{E}[Y(1) - Y(0) | 2024 \text{ (i.e. target)}]$$

- Design
  - Source: 2020 RCT from [Aggarwal et al. 2023]
  - Target: 2024 voters in Pennsylvania ($\sim$4.8 million from 67 counties as of Apr. 15)
  - 2024 covariates $\subset$ 2020 covariates
- Allows
  - Shift in voter demographics between 2024 and 2020 (i.e., covariate shift)
  - Shift in voter turnout between 2024 and 2020 (sensitivity analysis):

$$\underbrace{\mathrm{pr}(Y(a) \mid \textbf{target (2024)}, \text{ voter demographics})}_{\text{unobserved}} \neq \mathrm{pr}(Y(a) \mid \textbf{source (2020)}, \text{voter demographics})$$

Our Approach: Transfer Learning With Sensitivity Analysis

- Step I: Conduct inference under [Robins, Rotnitzky, and Scharfstein 2000]'s sensitivity model
- Step II: Find a plausible range of sensitivity parameters (i.e., calibration)

Step I: Conduct Inference Under The Sensitivity Model

- (γ_0, γ_1) quantifies the shift in $(Y(0), Y(1))$ between 2024 and 2020 via an odds ratio model

$$\exp(\gamma_a) = \frac{\text{Odd}(Y(a) \mid \textbf{target (2024)}, \text{ voter demographics})}{\text{Odd}(Y(a) \mid \textbf{source (2020)}, \text{ voter demographics})}, \quad a = 0, 1 \tag{1}$$

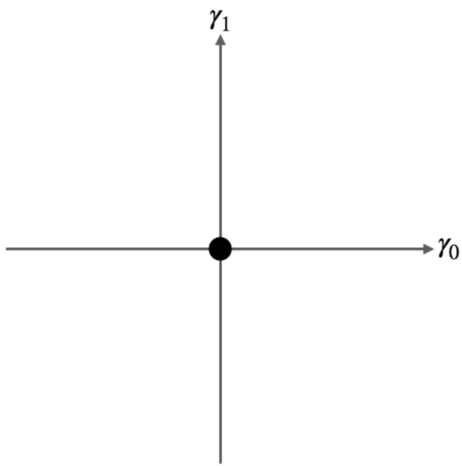

γ_0	γ_1
0	0

2024 ≈ 2020

Transported ATE (Ad Effect 2024)

$\theta(\gamma_0, \gamma_1)$

● 0.38

0 (%)

Step I: Conduct Inference Under The Sensitivity Model

- (γ_0, γ_1) quantifies the shift in $(Y(0), Y(1))$ between 2024 and 2020 via an odds ratio model

$$\exp(\gamma_a) = \frac{\text{Odd}(Y(a) \mid \textbf{target (2024)}, \text{ voter demographics})}{\text{Odd}(Y(a) \mid \textbf{source (2020)}, \text{ voter demographics})}, \quad a = 0, 1 \tag{1}$$

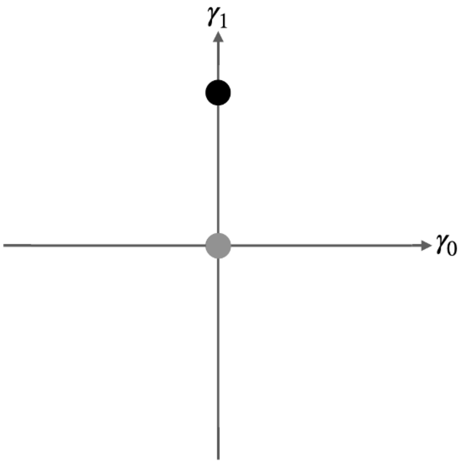

Transported ATE (Ad Effect 2024)

γ_0	γ_1		$\theta(\gamma_0, \gamma_1)$
0	0	2024 $\approx$ 2020	0.38
0	0.04	2024 $\neq$ 2020	1.16

0 (%)

Step I: Conduct Inference Under The Sensitivity Model

- (γ_0, γ_1) quantifies the shift in $(Y(0), Y(1))$ between 2024 and 2020 via an odds ratio model

$$\exp(\gamma_a) = \frac{\text{Odd}(Y(a) \mid \textbf{target (2024)}, \text{ voter demographics})}{\text{Odd}(Y(a) \mid \textbf{source (2020)}, \text{ voter demographics})}, \quad a = 0, 1 \tag{1}$$

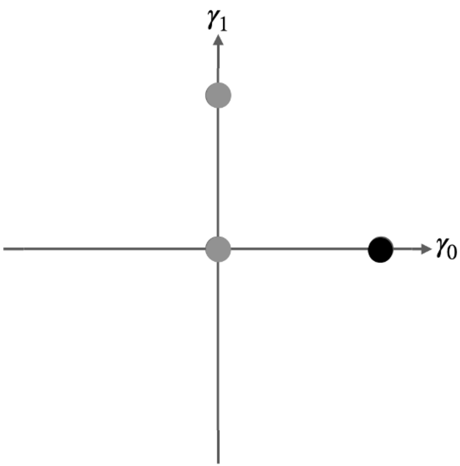

γ_0	γ_1		Transported ATE (Ad Effect 2024) $\theta(\gamma_0, \gamma_1)$
0	0	$2024 \approx 2020$	0.38
0	0.04	$2024 \neq 2020$	1.16
0.04	0	$2024 \neq 2020$	−0.39

0 (%)

Step I: Conduct Inference Under The Sensitivity Model

- (γ_0, γ_1) quantifies the shift in $(Y(0), Y(1))$ between 2024 and 2020 via an odds ratio model

$$\exp(\gamma_a) = \frac{\text{Odd}(Y(a) \mid \textbf{target (2024)}, \text{ voter demographics})}{\text{Odd}(Y(a) \mid \textbf{source (2020)}, \text{ voter demographics})}, \quad a = 0, 1 \tag{1}$$

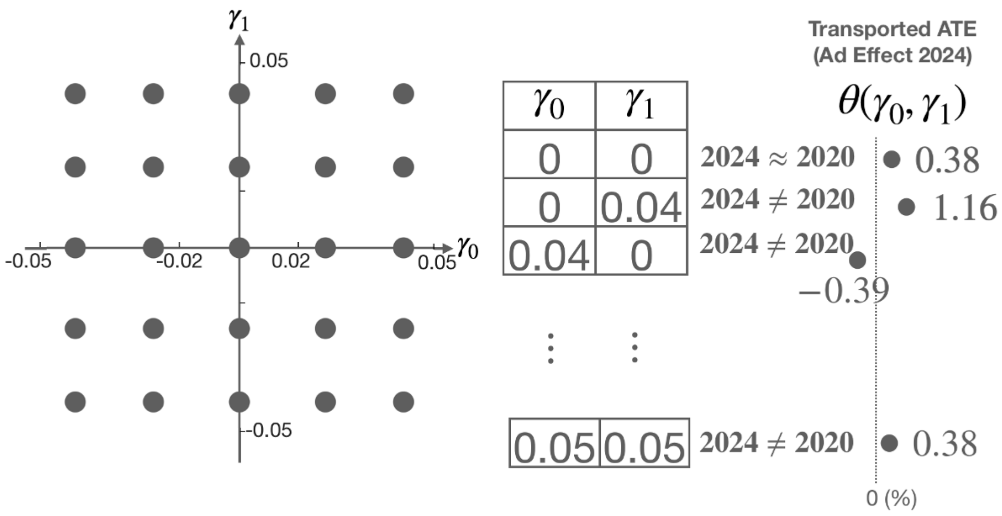

Step I: Conduct Inference Under The Sensitivity Model

- (γ_0, γ_1) quantifies the shift in $(Y(0), Y(1))$ between 2024 and 2020 via an odds ratio model

$$\exp(\gamma_a) = \frac{\text{Odd}(Y(a) \mid \textbf{target (2024)}, \text{ voter demographics})}{\text{Odd}(Y(a) \mid \textbf{source (2020)}, \text{ voter demographics})}, \quad a = 0, 1 \tag{1}$$

- Calculate p-values using the proposed plug-in estimator (nonparametric & efficient) with bootstrap

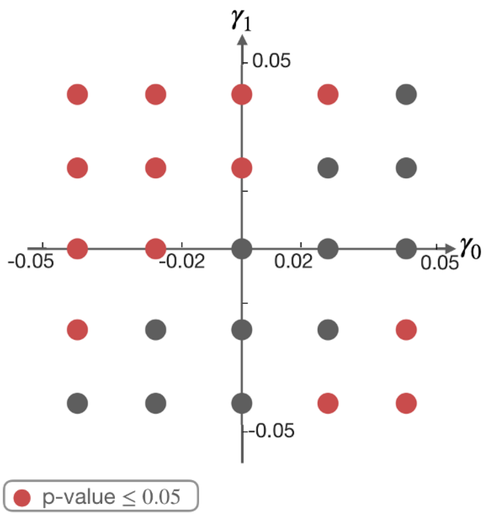

γ_0	γ_1
0	0
0	0.04
0.04	0
⋮	⋮
0.05	0.05

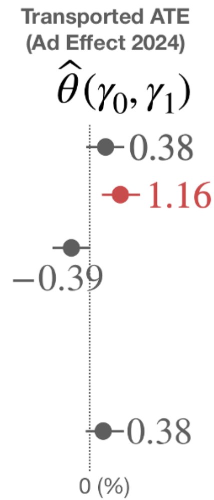

Step I: Conduct Inference Under The Sensitivity Model

- (γ_0, γ_1) quantifies the shift in $(Y(0), Y(1))$ between 2024 and 2020 via an odds ratio model

$$\exp(\gamma_a) = \frac{\text{Odd}(Y(a) \mid \textbf{target (2024)}, \text{ voter demographics})}{\text{Odd}(Y(a) \mid \textbf{source (2020)}, \text{ voter demographics})}, \quad a = 0, 1 \tag{1}$$

- Calculate p-values using the proposed plug-in estimator (nonparametric & efficient) with bootstrap

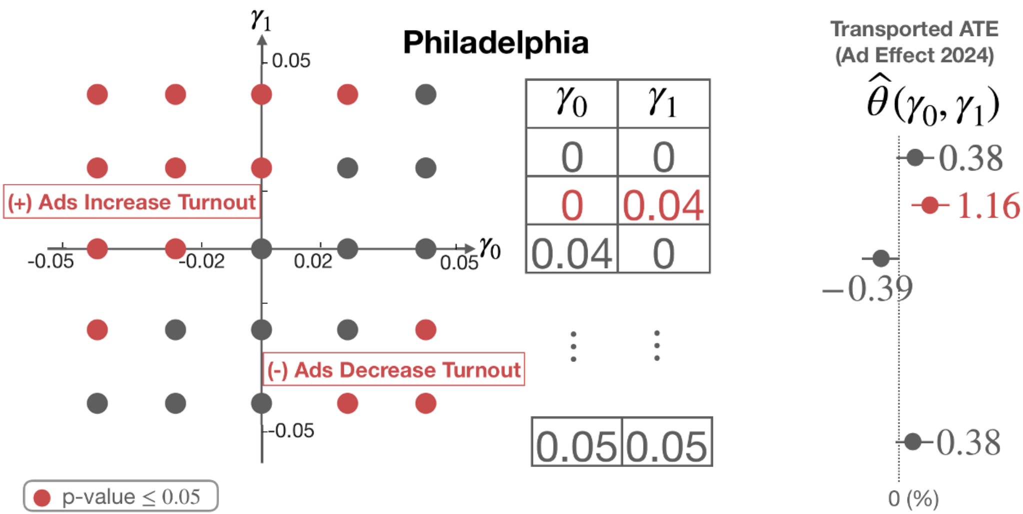

Step II: Find A Plausible Range of Sensitivity Parameters (i.e., Calibration)

Key idea: Find (γ_0, γ_1) that shifts blue collar states' $Y(a)$'s to non-blue collar states' $Y(a)$'s in 2020

Step II: Find A Plausible Range of Sensitivity Parameters (i.e., Calibration)

Key idea: Find (γ_0, γ_1) that shifts blue collar states' $Y(a)$'s to non-blue collar states' $Y(a)$'s in 2020

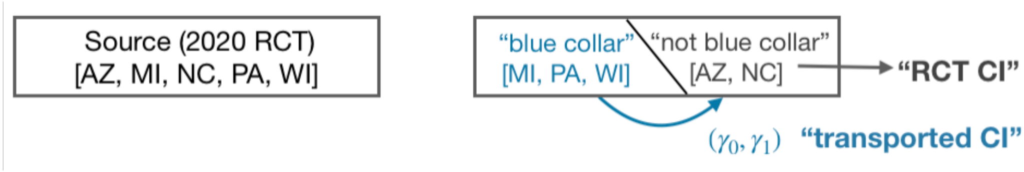

- (γ_0, γ_1) pairs such that the **"transported CI"**s overlap with the **"RCT CI"** consititute a plausible range

Step II: Find a Plausible Range of Sensitivity Parameters (i.e., Calibration)

Key idea: Find (γ_0, γ_1) that shifts blue collar states' $Y(a)$'s to non-blue collar states' $Y(a)$'s in 2020

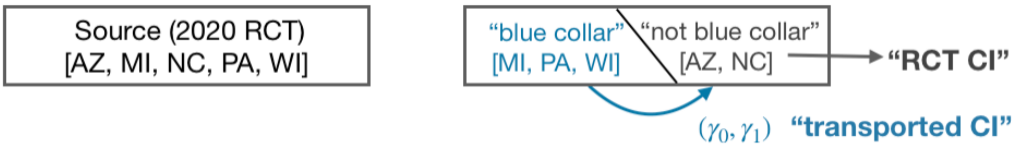

- (γ_0, γ_1) pairs such that the **"transported CI"**s overlap with the **"RCT CI"** consititute a plausible range

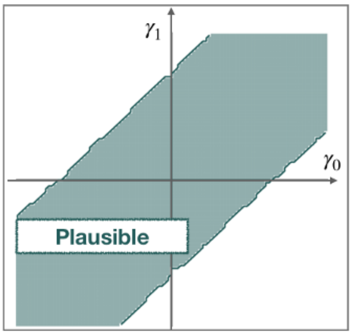

Step II: Find a Plausible Range of Sensitivity Parameters (i.e., Calibration)

Key idea: Find (γ_0, γ_1) that shifts blue collar states' $Y(a)$'s to non-blue collar states' $Y(a)$'s in 2020

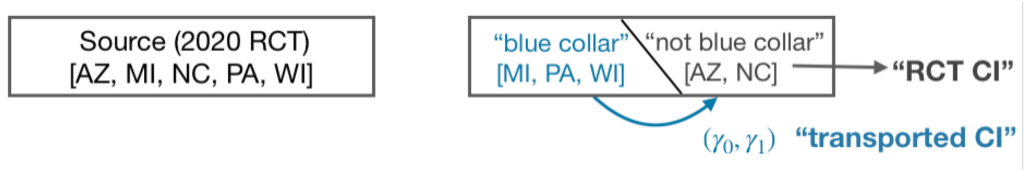

- (γ_0, γ_1) pairs such that the **"transported CI"**s overlap with the **"RCT CI"** consititute a plausible range
  - For Philadelphia, this calibration procedure rules out negative ad effects in 2024

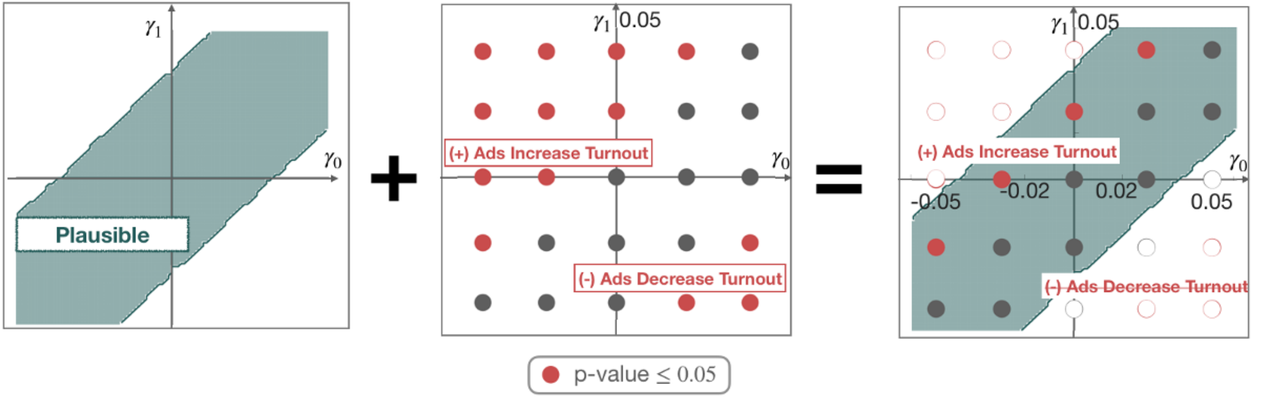

Calibrated Result for Pennsylvania: Positive Ad Effect (i.e., Increased Turnout)

- $\gamma_0 = \gamma_1 = 0$ (i.e., 2024 $\approx$ 2020) $\Rightarrow$ **ad effect remains insignificant in all counties**
- $\gamma_0 \gamma_1 \neq 0$ (i.e., 2024 $\neq$ 2020) $\Rightarrow$ **positive ad effect in 9 counties**
  - They are mostly urban-ish (Philadelphia, Pittsburgh, suburbs of Philadelphia) or college towns
  - Biden won all of the 9 counties in 2020

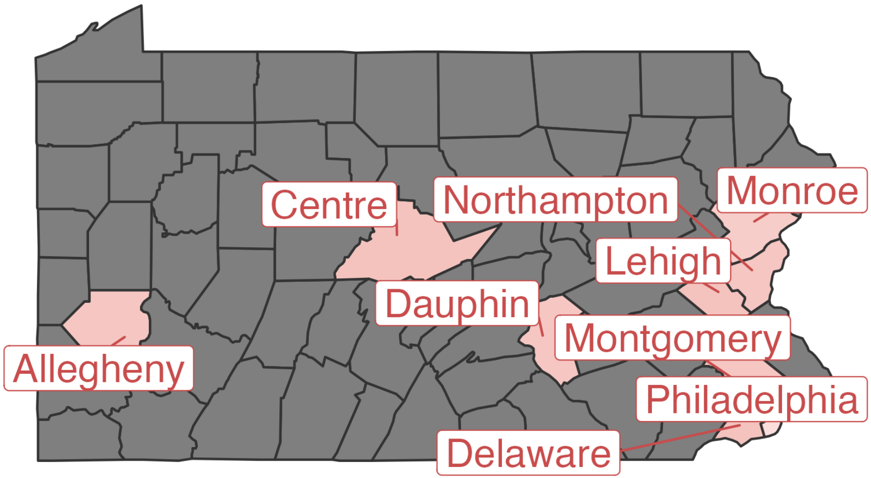

Figure. Gray: all significant (γ_0, γ_1) pairs are implausible.

Contact & Acknowledgements

- Want to learn more?
  - Come by my poster this afternoon :)
  - Reach out at xinran.miao@wisc.edu

- Acknowledgements
  - UW-Madison: **Xiaobin Zhou***, **Elaine Chiu***, **Xindi Lin***, **Ang Yu***, Sameer Deshpande, Jingqi Duan, Steven Moen, Ajinkya Hemant Kokandakar, Ben Teo, Kwangmoon Park, Jiaxin Hu, and statistics student seminar participants on Apr. 29
  - **Chan Park*** (UPenn), **Melody Huang*** (Harvard), Ying Jin (Stanford), and OCIS seminar participants on Apr. 30

*: people here at ACIC today

FAQs About Data

- Is the voting data self-reported?
  No. See [Aggarwal et al. 2023] for details.
- Does the data contain which candidate the voter voted for?
  No.
- Why is your voter data discrete?
  We're not sure. Perhaps, this is done to preserve some privacy?
- Is party registration measured accurately?
  Yes and no. [Aggarwal et al. 2023] and current voter registration data documentation discuss some reasons for errors.
- How was the treatment randomized?
  The randomization was stratified within gender, race, and age groups with the intention of increasing the propensity for women, black, and young people. The average treatment probability was 85.6%.
- Was the randomization done through Facebook/Meta?
  No. Our understanding from Aggarwal et al. 2023 is that the participants were randomized before the advertising company delivered the ads.
- Were ads delivered?
  Yes and no. 60% of the treatment group participants were identified and served ads. The analysis was intention-to-treat. See [Aggarwal et al. 2023] for details.

FAQs About Analysis

- What about treating this data as longitudinal?
  Excellent idea, but this requires measuring same voter over time. This data is not easy to get.
- Is our calibration reasonable?
  Great question! We're also experimenting with the "right" way to assess whether a given γ value is extreme or not. Note that our calibration is sensitive to the source population's study design.
- Is your conclusion sensitive to data quality from 2024 voter registration data (i.e. target data)?
  Yes. Unfortunately, high quality target data is expensive.
- Why is there a high proportion of treated individuals?
  We're not sure. Perhaps Acronym wanted to deliver the ads against Trump to as many voters as possible?
- Do you plan to validate your results for 2024?
  Great question! How to validate these results is a bit trick and we would be happy to talk to you more about this

FAQs About Framework and Assumptions

- Is SUTVA violated?
  Great question! It is possible, especially if
  - Different doses of ads: $Y(700 \text{ ads}) \neq Y(800 \text{ ads}) \neq Y(1)$
  - Different ads in 2020 and 2024: $Y(\text{ad in } 2020) \neq Y(\text{ad in } 2024)$
  - Voters talk to each other due to ads: $Y(\text{ my trt }, \text{ your trt}) \neq Y(\text{ my trt })$
  - There are carry-over effects from 2020 ad campaign into 2024

  But, we also picked 2020 and 2024 to minimize SUTVA violations as the candidates are identical between the two years.
- Is the data from 2020 independent from 2024?
  Excellent question! This question is a bit tricky to answer, especially if 2024 is a fixed, census-level data. We're happy to talk more about this.
- Can the sensitivity model depend on covariates?
  Yes. While this introduces more complexity in interpreting the sensitivity parameters, it could be useful if there is a priori knowledge about how the ad effects in 2024 and 2020 differ with respect to measured covariates. For example, for $A = 0$, we can define

$$\exp(\gamma_{\text{Democrat?}} I(\text{Democrat?})), \quad \exp(\gamma_{\text{Republican?}} I(\text{Republican?}))$$

  if we believe the change between 2020 and 2024 is different for Democrats and Republicans. We call this a **local sensitivity model.**

Notation and Causal Assumptions

- Population type: $S \in \{0, 1\}$ where
  - $S = 1$ is source (e.g. 2020)
  - $S = 0$ is target (e.g. 2024)
- Outcome: $Y \in \{0, 1\}$ where $Y = 1$ is voted
- Treatment: $A \in \{0, 1\}$ where $A = 1$ is ad against Trump
- Covariates: $\mathbf{X} \in \mathbb{R}^p$ where
  - Source covariates: $\mathbf{X}$
  - Target covariates: $\mathbf{V} \subset \mathbf{X}$
- Potential outcomes: $Y(a) \in \{0, 1\}$, $a \in \{0, 1\}$
  - $Y(1)$: voted if, contrary to fact, voter got negative ad
  - $Y(0)$: voted if, contrary to fact, voter did not get negative ad
- Causal estimand: $\theta = \mathbb{E}[Y(1) - Y(0) \mid S = 0]$

Data Table for Our Setup

		$\mathbf{X}$					
	S	$\mathbf{V}$	$\mathbf{X} \setminus \mathbf{V}$	A	$Y(1)$	$Y(0)$	Y
	1	$\checkmark$	$\checkmark$	1	$\checkmark$		$\checkmark$
	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$		$\vdots$
Source RCT (i.e. 2020)	1	$\checkmark$	$\checkmark$	1	$\checkmark$		$\checkmark$
	1	$\checkmark$	$\checkmark$	0		$\checkmark$	$\checkmark$
	$\vdots$	$\vdots$	$\vdots$	$\vdots$		$\vdots$	$\vdots$
	1	$\checkmark$	$\checkmark$	0		$\checkmark$	$\checkmark$
	0	$\checkmark$					
Target (i.e. 2024)	$\vdots$	$\vdots$					
	0	$\checkmark$					

The goal is to identify and estimate $\theta = \mathbb{E}[Y(1) - Y(0) \mid S = 0]$

Review of Transportability in RCT

Causal assumptions on the source:

(A1) SUTVA: $Y = Y(A)$ if $S = 1$

(A2) Randomized treatment: $A \perp Y(1), Y(0) \mid \mathbf{X}, S = 1$

(A3) Overlap of A: $0 < \pi(\mathbf{x}) = \mathrm{pr}(A = 1 \mid \mathbf{X} = \mathbf{x}, S = 1) < 1$

(A1)-(A3) are usually satisfied in RCTs

Transportation assumptions:

(A4) Overlap of S: $0 < \mathrm{pr}(S = 1 \mid \mathbf{V} = \mathbf{v}) < 1$ for all $\mathbf{v}$

(A5) Transportability: $Y(1), Y(0) \perp S \mid \mathbf{V}$

(A4) can be checked with data while (A5) cannot be checked with data

See [Tipton and Peck 2017],[Dahabreh, Robins, Haneuse, and Hernán 2019], [Egami and Hartman 2023], and [Degtiar and Rose 2023] for details

Review: Identification Under (A1)-(A5)

Let $\mu_a = \mathbb{E}[Y \mid \mathbf{X}, A = a, S = 1]$. Under (A1)-(A5), the target ATE θ is identified:

$$\begin{aligned}\theta &= \mathbb{E}[Y(1) - Y(0) \mid S = 0] \\ &= \underbrace{\mathbb{E}[\underbrace{\mathbb{E}[\underbrace{\mu_1(\mathbf{X}) - \mu_0(\mathbf{X})}_{\text{CATE}(\mathbf{X}) \text{ in source}} \mid \mathbf{V}, S = 1]}_{\text{CATE}(\mathbf{V}) \text{ in source}} \mid S = 0]}_{\text{Reweigh CATE}(\mathbf{V}) \text{ to target}}\end{aligned}$$

For reference, when $\mathbf{X} = \mathbf{V}$, we have

$$\theta = \mathbb{E}[\mathbb{E}[\mu_1(\mathbf{X}) - \mu_0(\mathbf{X}) \mid S = 0].$$

In other words, when source and target covariates differ, we have more nuisance parameters (i.e. projection of CATE($\mathbf{X}$) onto $\mathbf{V}$); see recent work by [Zeng et al. 2023].

What If Transportability (A5) Fails?
Sensitivity Analysis

Suppose transportability (A5) does not hold:

$$\underbrace{\mathrm{pr}(Y(1), Y(0) \mid \mathbf{V}, S=0)}_{\text{2024 (i.e., target)}} \neq \underbrace{\mathrm{pr}(Y(1), Y(0) \mid \mathbf{V}, S=1)}_{\text{2020 (i.e., source)}}$$

For each $Y(a)$, we measure the deviation between the two probabilities via odds ratios:

$$\exp(\gamma_a) = \frac{\mathrm{Odd}(Y(a) \mid \mathbf{v}, S=0)}{\mathrm{Odd}(Y(a) \mid \mathbf{v}, S=1)}, \ \gamma_a \in (-\infty, \infty) \quad (2)$$

$$\mathrm{Odd}(Y(a) \mid \mathbf{v}, s) = \frac{\mathrm{pr}(Y(a)=1 \mid \mathbf{V}=\mathbf{v}, S=s)}{1-\mathrm{pr}(Y(a)=1 \mid \mathbf{V}=\mathbf{v}, S=s)}, \ s \in \{0,1\} \quad (3)$$

Broadly speaking, the unobservable part (target) differs from the observable part (source) by $\exp(\gamma_a)$.

- Since the red part is unobserved, γ_a cannot be estimated. Instead, γ_a is chosen to quantify the difference between the target and the source
- In general, a large $|\gamma_a| \Rightarrow$ large difference between 2020 and 2024

Interpreting Sensitivity Parameter γ_a

$$\exp(\gamma_a) = \frac{\text{Odd}(Y(a) \mid \mathbf{v}, S = 0)}{\text{Odd}(Y(a) \mid \mathbf{v}, S = 1)}$$

- When $\gamma_a = 0$, transportability (A5) holds
- When $\gamma_a \neq 0$, large $|\gamma_a|$ $\Rightarrow$ large difference between 2024 and 2020
- Positive γ_1 $\Rightarrow$ more turnout in 2024 after receiving ads against Trump compared to that in 2020
- Negative γ_1 $\Rightarrow$ less turnout in 2024 after receiving ads against Trump compared to that in 2020

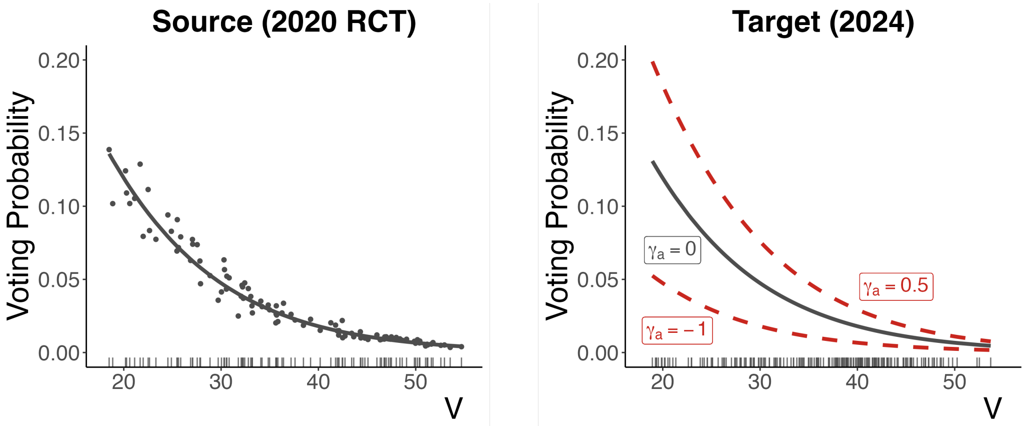

Toy example: $\text{pr}(Y(a) = 1 \mid V, S = 1) = \text{expit}(-0.1V)$

Two-Parameter Sensitivity Model and Some Remarks

To jointly characterize $Y(1)$, $Y(0)$, we use the following model

$$\text{pr}(Y(1) = y_1, Y(0) = y_0 \mid \mathbf{V} = \mathbf{v}, S = 0) \propto \exp(\gamma_1 y_1 + \gamma_0 y_0) \cdot \text{pr}(Y(1) = y_1, Y(0) = y_0 \mid \mathbf{V} = \mathbf{v}, S = 1)$$

Some remarks:

- The sensitivity model does not place any observable restrictions on the data [Robins, Rotnitzky, and Scharfstein 2000; Franks, D'Amour, and Feller 2019]
- A pseudo-R^2 version of γ_a is in Proposition 3 of Franks, D'Amour, and Feller 2019.
- We can also reparametrize the sensitivity model in terms of $P(S = 1 \mid Y(1), Y(0), \mathbf{V} = \mathbf{v})$; see Appendix and [Carroll et al. 1997].
- The sensitivity model can depend on covariates (e.g. $\exp(\gamma_{\mathbf{v}}^{\intercal}\mathbf{v} + \ldots)$; "local" sensitivity analysis)
- Some works that use this model: [Robins, Rotnitzky, and Scharfstein 2000; Franks, D'Amour, and Feller 2019; Scharfstein et al. 2021; Dahabreh, Robins, Haneuse, Robertson, et al. 2022]
- There is a **long and healthy** discussion about what constitutes a "good" model for sensitivity analysis [Robins 2002; Rosenbaum 2002]

Alternate Parametrization of The Sensitivity Model

- The sensitivity model $\exp(\gamma_a) = \dfrac{\mathrm{Odd}(Y(a) \mid \mathbf{v}, S=0)}{\mathrm{Odd}(Y(a) \mid \mathbf{v}, S=1)}$ implies the following partially linear logistic regression model [Carroll et al. 1997]:

$$
\begin{aligned}
\mathrm{pr}(S=1 \mid Y(a)=y, \mathbf{V}=\mathbf{v}) &= \mathrm{expit}\left(-\gamma_a y - \eta_a(\mathbf{v})\right) \\
\eta_a(\mathbf{v}) &= \log\left(\frac{\mathrm{pr}(S=0)}{\mathrm{pr}(S=1)} \frac{w(\mathbf{v})}{\mathrm{E}\{\exp(\gamma_a Y(a)) \mid \mathbf{v}, S=1\}}\right) \\
w(\mathbf{V}) &= p(\mathbf{V} \mid S=0)/p(\mathbf{V} \mid S=1)
\end{aligned}
$$

- The joint sensitivity model implies the following partially linear logistic regression model:

$$
\begin{aligned}
\mathrm{pr}(S=1 \mid Y(1)=y_1, Y(0)=y_0, \mathbf{V}=\mathbf{v}) &= \mathrm{expit}\left(-\gamma_1 y_1 - \gamma_0 y_0 - \eta(\mathbf{v})\right) \\
\eta(\mathbf{v}) &= \log\left(\frac{\mathrm{pr}(S=0)}{\mathrm{pr}(S=1)} \frac{w(\mathbf{v})}{\mathrm{E}\{\exp(\gamma_1 Y(1) + \gamma_0 Y(0)) \mid \mathbf{V}, S=1\}}\right)
\end{aligned}
$$

Identification Under (A1)-(A4) + Sensitivity Model

Again, let $\mu_a(\mathbf{X}) = \mathrm{E}(Y \mid \mathbf{X}, A = a, S = 1)$. Under (A1)-(A4) and the sensitivity model, we have

$$\mathrm{E}[Y(a) \mid S = 0] = \mathrm{E}\left[\frac{\mathrm{E}\{\exp(\gamma_a)\mu_a(\mathbf{X}) \mid \mathbf{V}, S = 1\}}{\mathrm{E}\{\exp(\gamma_a)\mu_a(\mathbf{X}) + 1 - \mu_a(\mathbf{X}) \mid \mathbf{V}, S = 1\}}\middle| S = 0\right].$$

- It is an exponential tilt of $\rho_a(\mathbf{V}) = \mathbb{E}[\mu_a(\mathbf{X}) \mid \mathbf{V}, S = 1]$
- If $\gamma_a = 0$ (i.e. transportability (A5) holds), we return to the previous result:

$$\mathbb{E}[Y(a) \mid S = 0] = \mathbb{E}[\mathbb{E}[\mu_a(\mathbf{X}) \mid \mathbf{V}, S = 1] \mid S = 0]$$

Estimation: Simple, Plug-In Estimator

$$\mathrm{E}[Y(a) \mid S=0] = \mathrm{E}\left[\frac{\mathrm{E}\{\exp(\gamma_a)\mu_a(\mathbf{X}) \mid \mathbf{V}, S=1\}}{\mathrm{E}\{\exp(\gamma_a)\mu_a(\mathbf{X}) + 1 - \mu_a(\mathbf{X}) \mid \mathbf{V}, S=1\}} \middle| S=0\right]$$

Identification leads to a simple, plug-in estimator:

1. Estimate $\rho_a(\mathbf{V}) = \mathbb{E}[\mu_a(\mathbf{X}) \mid \mathbf{V}, S=1]$ from source data
   - An example: for $a=1$, regress $Y/\widehat{\pi}(\mathbf{X})$ on $\mathbf{V}$ to get $\widehat{\rho}_1(\mathbf{V})$
   - Because source is an RCT, $\rho_a(\mathbf{V})$ can be consistently estimated
2. Average the exponentially tilted $\rho_a(\mathbf{V})$ among target sample

$$\widehat{\mathbb{E}}[Y(a) \mid S=0] = \frac{1}{n_t} \sum_{i \in \text{Target}} \frac{\exp(\gamma_a)\widehat{\rho}_a(\mathbf{v}_i)}{\exp(\gamma_a)\widehat{\rho}_a(\mathbf{v}_i) + 1 - \widehat{\rho}_a(\mathbf{v}_i)}$$

$$\widehat{\theta}(\gamma_1, \gamma_0) = \widehat{\mathbb{E}}[Y(1) \mid S=0] - \widehat{\mathbb{E}}[Y(0) \mid S=0]$$

Since our voter data is discrete, this plug-in estimator is **nonparametric** and **efficient**; see Theorem 1 of [Chamberlain 1987]

Bootstrapping for Transfer Learning

In general, bootstrapping is easy and convenient for estimating SEs or confidence intervals (CIs)

Here, we lay out one (theoretically valid) bootstrap for transfer learning with sensitivity analysis

In each bootstrap iteration $b \in \{1, \cdots, B\}$:

1. Resample source data with replacement of size n_s, obtain data $\mathcal{D}^*_{\mathcal{S}}$
2. Resample target data with replacement of size n_t, obtain data $\mathcal{D}^*_{\mathcal{T}}$
3. With $\mathcal{D}^*_{\mathcal{S}}$ and $\mathcal{D}^*_{\mathcal{T}}$, construct the ATE estimator $\widehat{\theta}^*_b$ from above

Theorem: If $\rho(\mathbf{V})$ is smooth enough and Donsker condition holds, the above procedure yields a valid $(1-\alpha)$ CI of θ.

The smoothness $+$ Donsker conditions hold for our discrete voter data.

Estimator Based on Efficient Influence Function

For a given γ_a, we have the efficient influence function (EIF) of θ.

- The EIF is **very messy** because (a) $\mathbf{V} \neq \mathbf{X}$ and (b) sensitivity analysis; see next page.
- Our EIF recovers Zeng et al. 2023's EIF when $\gamma_a = 0$.

Practically, an EIF-based estimator is useful if $\mathbf{V}$ is continuous.

- Four nuisance functions: (i) propensity score $\pi(\mathbf{X})$, (ii) outcome regression $\mu_a(\mathbf{X})$, (iii) projection of outcome regression $\rho(\mathbf{V})$, and (iv) weights between source and target $w(\mathbf{V}) = p(\mathbf{V} \mid S = 0)/p(\mathbf{V} \mid S = 1)$.
- To avoid Donsker conditions, we need cross-fitting in source data.
- The estimator is not doubly robust for $\gamma_a \neq 0$.
- Also, the estimator does not reduce "plug-in bias" from $\rho_a(\mathbf{V})$.

EIF-Based Estimator

$$
\begin{aligned}
\widehat{\theta}_{\text{EIF},a} =& \frac{1}{n_s} \sum_{i \in \text{Source}} \widehat{w}(\mathbf{v}_i) \left(\left\{ \frac{a_i}{\widehat{\pi}(\mathbf{x}_i)} + \frac{1-a_i}{1-\widehat{\pi}(\mathbf{x}_i)} \right\} \left[\frac{e^{\gamma_a y_i} y_i}{e^{\gamma_a}\widehat{\rho}(\mathbf{v}_i) + 1 - \widehat{\rho}(\mathbf{v}_i)} - \frac{e^{\gamma_a}\widehat{\mu}_a(\mathbf{x}_i)}{e^{\gamma_a}\widehat{\rho}(\mathbf{v}_i) + 1 - \widehat{\rho}(\mathbf{v}_i)} \right.\right. \\
& \left. - \frac{e^{\gamma_a y_i}\widehat{\rho}_a(\mathbf{v}_i)}{[e^{\gamma_a}\widehat{\rho}_a(\mathbf{v}_i) + 1 - \rho_a(\mathbf{v}_i)]^2} + \frac{\{e^{\gamma_a}\widehat{\mu}_a(\mathbf{x}_i) + 1 - \widehat{\mu}_a(\mathbf{x}_i)\} e^{\gamma_a}\widehat{\rho}(\mathbf{v}_i)}{[e^{\gamma_a}\widehat{\rho}_a(\mathbf{v}_i) + 1 - \rho_a(\mathbf{v}_i)]^2} \right] \\
& \left. + \frac{\widehat{\mu}_a(\mathbf{x}_i)\{e^{\gamma_a}\widehat{\rho}(\mathbf{v}_i) + 1 - \widehat{\rho}(\mathbf{v}_i)\} - \widehat{\rho}_a(\mathbf{v}_i)\{e^{\gamma_a}\widehat{\mu}_a(\mathbf{x}_i) + 1 - \widehat{\mu}_a(\mathbf{x}_i)\}}{[e^{\gamma_a}\widehat{\rho}_a(\mathbf{v}_i) + 1 - \rho_a(\mathbf{v}_i)]^2} \right) \\
& + \frac{1}{n_t} \sum_{i \in \text{Target}} \frac{e^{\gamma_a}\widehat{\rho}_a(\mathbf{v}_i)}{e^{\gamma_a}\widehat{\rho}_a(\mathbf{v}_i) + 1 - \widehat{\rho}_a(\mathbf{v}_i)}
\end{aligned}
$$

Our EIF recovers [Zeng et al. 2023]'s EIF when $\gamma_a = 0$:

$$
\begin{aligned}
=& \frac{1}{n_s} \sum_{i \in \text{Source}} \widehat{w}(\mathbf{v}_i) \left(\frac{a_i}{\widehat{\pi}(\mathbf{x}_i)} + \frac{1-a_i}{1-\widehat{\pi}(\mathbf{x}_i)} \right) [y_i - \widehat{\mu}_a(\mathbf{x}_i)] + \\
& \frac{1}{n_s} \sum_{i \in \text{Source}} \widehat{w}(\mathbf{v}_i)[\widehat{\mu}_a(\mathbf{x}_i) - \widehat{\rho}_a(\mathbf{v}_i)] + \frac{1}{n_t} \sum_{i \in \text{Target}} \widehat{\rho}(\mathbf{v}_i)
\end{aligned}
$$

Voter Demographics Between Source and Target Population

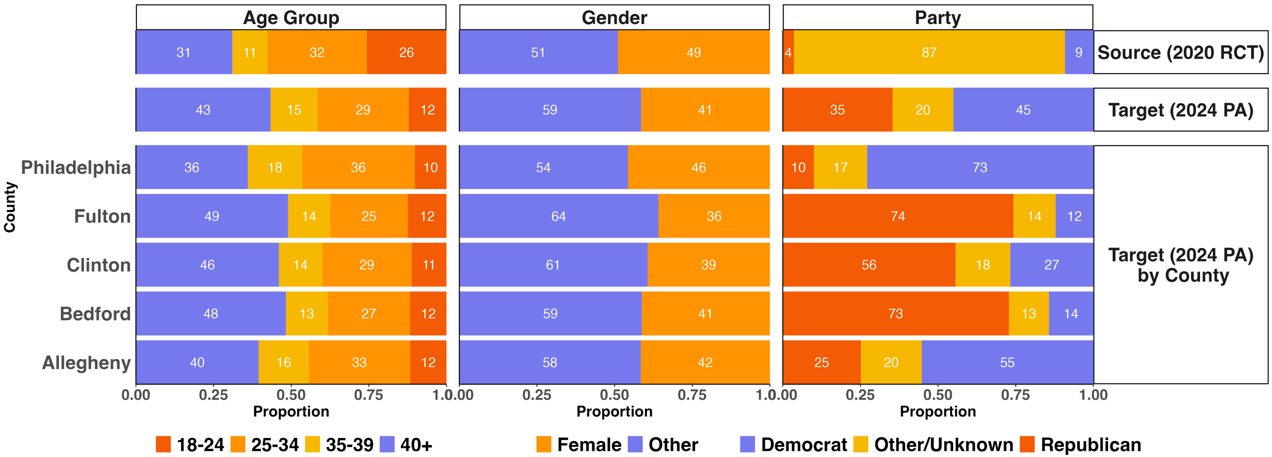

Figure. Registered voter demographics.

Additional Results on Pennsylvania: Sensitivity Across Counties

- We calculate the smallest $|\gamma_1 - \gamma_0|$ that makes the target ATE significant
  - A small $|\gamma_1 - \gamma_0|$ $\Rightarrow$ a small shift can make the ad effect significant $\Rightarrow$ more sensitive

	County	$\|\gamma_1 - \gamma_0\|$	Estimate (95% CI)
Most Sensitive for Positive Effects	Philadelphia	0.018	0.71 (0.00, 1.43)
	Monroe	0.028	0.62 (0.02, 1.22)
Most Sensitive for Negative Effects	Bedford	0.002	−0.81 (−1.60, −0.01)
	Fulton	0.002	−0.88 (−1.72, −0.03)
Sensitive In Either Direction	Centre	0.032	0.59 (0.01, 1.16)
		0.028	−0.059 (−1.17, 0.01)
	Lehigh	0.034	0.53 (0.02, 1.23)
		0.032	−0.62 (−1.23, −0.02)
	Northampton	0.032	0.60 (0.01, 1.20)
		0.030	−0.59 (−1.18, 0.00)
Insensitive In Both Directions	Clinton	-	-

Table. The smallest $|\gamma_1 - \gamma_0|$ that makes the ad effect significant and the corresponding effect estimates with 95% confidence intervals. Ad effects are in the unit of percent point.

Additional Results on Pennsylvania: Negative Ad Effects (i.e., Decreased Turnout)

- Most counties are sensitive towards a negative effect (i.e., ads against Trump decrease turnout)
  - The most sensitive countes are Fulton and Bedford, which had the largest margin for Trump in 2024 (85.41% for Trump in Fulton; 84.49% for Trump in Bedford)
- Three counties are sensitive for both positive and negative effects: Centre, Lehigh, Northampton[1]

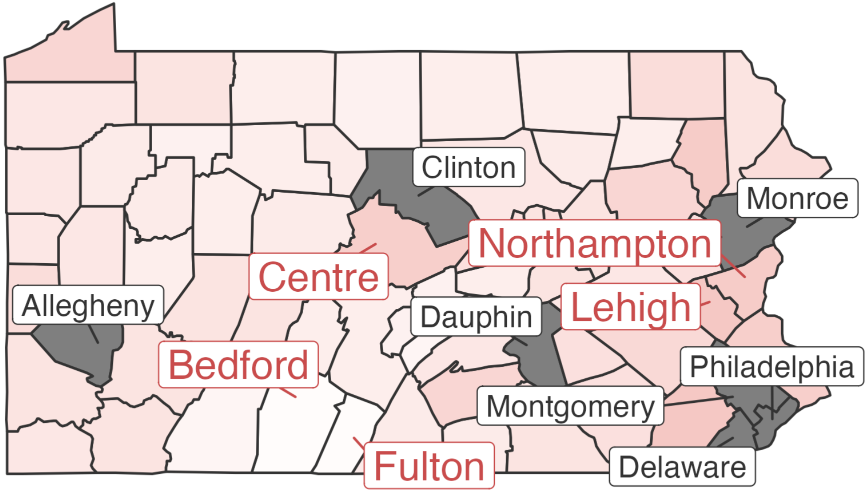

Figure. Black: insensitive to negative ad effects. Gray: all significant (γ_0, γ_1) pairs are implausible.

[1] Northampton is a "pivot county" identified by Ballopedia.

References I

- Aggarwal, Minali et al. (2023). “A 2 million-person, campaign-wide field experiment shows how digital advertising affects voter turnout”. In: *Nature Human Behaviour* 7.3, pp. 332–341.
- Carroll, Raymond J et al. (1997). “Generalized partially linear single-index models”. In: *Journal of the American Statistical Association* 92.438, pp. 477–489.
- Chamberlain, Gary (1987). “Asymptotic efficiency in estimation with conditional moment restrictions”. In: *Journal of econometrics* 34.3, pp. 305–334.
- Dahabreh, Issa J, James M Robins, Sebastien JP Haneuse, and Miguel A Hernán (2019). “Generalizing causal inferences from randomized trials: counterfactual and graphical identification”. In: *arXiv preprint arXiv:1906.10792*.
- Dahabreh, Issa J, James M Robins, Sebastien JP Haneuse, Sarah E Robertson, et al. (2022). “Global sensitivity analysis for studies extending inferences from a randomized trial to a target population”. In: *arXiv preprint arXiv:2207.09982*.
- Degtiar, Irina and Sherri Rose (2023). “A review of generalizability and transportability”. In: *Annual Review of Statistics and Its Application* 10, pp. 501–524.
- Egami, Naoki and Erin Hartman (2023). “Elements of external validity: Framework, design, and analysis”. In: *American Political Science Review* 117.3, pp. 1070–1088.
- Franks, AlexanderM, Alexander D’Amour, and Avi Feller (2019). “Flexible sensitivity analysis for observational studies without observable implications”. In: *Journal of the American Statistical Association*.

References II

- Robins, James M (2002). "[Covariance Adjustment in Randomized Experiments and Observational Studies]: Comment". In: *Statistical Science* 17.3, pp. 309–321.
- Robins, James M, Andrea Rotnitzky, and Daniel O Scharfstein (2000). "Sensitivity analysis for selection bias and unmeasured confounding in missing data and causal inference models". In: *Statistical models in epidemiology, the environment, and clinical trials*. Springer, pp. 1–94.
- Rosenbaum, Paul R (2002). "Covariance adjustment in randomized experiments and observational studies". In: *Statistical Science* 17.3, pp. 286–327.
- Scharfstein, Daniel O et al. (2021). "Semiparametric sensitivity analysis: Unmeasured confounding in observational studies". In: *arXiv preprint arXiv:2104.08300*.
- Tipton, Elizabeth and Laura R Peck (2017). "A design-based approach to improve external validity in welfare policy evaluations". In: *Evaluation review* 41.4, pp. 326–356.
- Zeng, Zhenghao et al. (2023). "Efficient Generalization and Transportation". In: *arXiv preprint arXiv:2302.00092*.