

# UC Santa Barbara

## UC Santa Barbara Electronic Theses and Dissertations

### Title

Essays on the Economics of Extremism in the United States

### Permalink

<https://escholarship.org/uc/item/5xw550w7>

### Author

Klinenberg, Daniel Jacob

### Publication Date

2023

Peer reviewed|Thesis/dissertation

University of California  
Santa Barbara

# Essays on the Economics of Extremism in the United States

A dissertation submitted in partial satisfaction  
of the requirements for the degree

Doctor of Philosophy  
in  
Economics

by

Daniel Klinenberg

Committee in charge:

Professor Richard Startz, Chair  
Professor Eli Berman, University of California San Diego  
Professor Gonzalo Vazquez-Bare

September 2023

The Dissertation of Daniel Klinenberg is approved.

---

Professor Eli Berman, University of California San Diego

---

Professor Gonzalo Vazquez-Bare

---

Professor Richard Startz, Committee Chair

August 2023

Essays on the Economics of Extremism in the  
United States

Copyright © 2023

by

Daniel Klinenberg

*If not you, then who? If not now, when?*

- Hillel

## Acknowledgements

This dissertation has been the accumulation of support and care from my friends, family, and colleagues. I'd like to start by thanking my family and friends. I am forever grateful to my parents, Andrea and Ed, who have supported me throughout this process. You were and always will be there for me with unconditional love, support, and unsolicited advice. I'd also like to thank my close friends, Matt, Brian, Henry, Shane, Connor, Kyle, and Curtis, for staying with me through the many years of stress, missed parties, and long distance. I look forward to a future with more time together. Finally, I'd like to thank my girlfriend Sarah. I can say with near certainty you are the reason I made it through the program. We have lived not just through a Ph.D. together, but also a pandemic! You are my best friend and I can't wait to see what comes next for us.

Next, I'd like to thank my colleagues at UCSB. I have been lucky to have a great group of colleagues who have become friends. I'd like to recognize Kite, Nir, Allegra, Antoine, Jaime, Ryan, and Michael. Thank you for half a decade of memories.

I want to thank my committee, Gonzalo Vazquez-Bare, and Eli Berman. To Gonzalo, thank you for indulging in my metrics questions, and last-minute emails asking for meetings. To Eli, thank you for taking a chance with me. Your advice and guidance down this path of research are undoubtedly one of the pivotal experiences in my career that I will carry with me throughout my life.

Finally, I want to thank Dick. Your advice transcends work. The way in which you advised me holistically has shown me there is more to research than words and equations. It's about people. If I can be half the advisor for someone else as you have been to me, then I will be far better than I could have ever imagined. Thank you for teaching me how to be a better economist and, more importantly, a better person.

# Curriculum Vitæ

## Daniel Klinenberg

### Education

2023	Ph.D. in Economics (Expected), University of California, Santa Barbara
2019	M.A. in Economics, University of California, Santa Barbara
2018	M.S. in Quantitative Economics, California Polytechnic State University, San Luis Obispo
2017	B.S. in Quantitative Economics, California Polytechnic State University, San Luis Obispo

### Research Interests

Applied microeconomics, Political economy, Bayesian applied econometrics

### Working Papers

"Selling Violent Extremism"

"Mostly Deterred: An Episodic Analysis of The Israel-Gaza Conflict" (*with Eli Berman, Esteban Klor, and John Powell*)

### Publications

Klinenberg, D. (2023). Does Deplatforming Work? *Journal of Conflict Resolution*, 0(0). <https://doi.org/10.1177/00220027231188909>

Danny Klinenberg (2022) Synthetic Control with Time Varying Coefficients: A State Space Approach with Bayesian Shrinkage, *Journal of Business and Economic Statistics*, DOI: 10.1080/07350015.2022.2102025

Danny Klinenberg & Richard Startz (2022) Covid, colleges, and classes, *Applied Economics*, DOI: 10.1080/00036846.2022.2091110

### Other Work

Advanced Gradescope autograder tutorial using R

Data wrangling for economists bookdown notes (*with Michael Topper*)

### Fellowships and Award

Research quarter fellowship

Winter 2021, Spring 2022

Robert T. Deacon Graduate Fellowship Winter 2020  
Outstanding Undergraduate TA 2020-2021

### **Research Assistance**

Eli Berman Summer 2022-Present

### **Selected Teaching Experience**

Economics 241B-Econometrics (Ph.D. first year) Winter 2022  
Econ 245-Data Wrangling for Economists (Ph.D.) Summer-Fall 2021  
Econ 145-Data Wrangling for Economists 2021

### **Presentations**

Bridging Divides Initiative Summer Summit 2023  
California Polytechnic State Economics Department Seminar Series (Presenter) 2022  
Empirical Study of Conflict Annual Meeting (Presenter and Panelist for Innovative Data Collection) 2022  
Western Economic Association Annual Conference (Presenter) 2022  
International Association of Applied Econometrics conference - Machine Learning Session (Presenter) 2020

### **Refereeing**

Empirical Economics, Journal of Policy Analysis and Management

### **Other Activities**

UCSB Economics Forecast Project Summer 2019  
United States Department of State Virtual Student Foreign Service - Millennium Challenge Corporation 2019-2020



## **Permissions and Attributions**

1. The content of Chapter 2 and Appendix B is accepted in the Journal of Conflict Resolution. A pre-print is used in this dissertation.
2. The content of Chapter 3 and Appendix C is accepted in the Journal of Business and Economics Statistics. This is an Author's Accepted Manuscript of an article published by Taylor & Francis Group in the Journal of Business and Economic Statistics, is available online:

<https://www.tandfonline.com/10.1080/07350015.2022.2102025>.

## Abstract

Essays on the Economics of Extremism in the  
United States

by

Daniel Klinenberg

This dissertation consists of two essays studying extremism in the United States using economic tools and theory and one essay extending a well-known econometric method.

In the first chapter, I study how domestic extremist organizations in the United States attract new members. Such organizations have become a significant threat to Western democracies, with some groups attracting tens of thousands of members. I study the impact of three tactics—membership discounts, armed events, and advertisements—on national recruitment by the Oath Keepers, America's largest paramilitary organization. Using a synthetic control framework, I find that discounts cause new member signups to increase by 144% and armed events by 170%; however, advertisements decrease it. Finally, I fail to find strong evidence that economic inequality drives the inflow of new members during any tactic.

In Chapter 2, I study the efficacy of social media policy designed to curb extremist activity. While governments deliberate on how to regulate, some social media companies have removed creators of offensive content—deplatforming. I estimate the effects of deplatforming on revenue and viewership, using variations in the timing of removals across two video-streaming companies—YouTube, and its far-right competitor, Bitchute. Being deplatformed on YouTube results in a 30% *increase* in weekly Bitcoin revenue and a 50% increase in viewership on Bitchute. This increase in Bitchute activity is less than that on YouTube, meaning that deplatforming works in decreasing a content creator's

overall views and revenue.

Finally, Chapter 3 extends the synthetic control methodology to account for more scenarios. Synthetic control methods are a popular tool for measuring the effects of policy interventions on a single treated unit. In practice, researchers create a counterfactual using a linear combination of untreated units that closely mimic the treated unit. Oftentimes, creating a synthetic control is not possible due to untreated units' dynamic characteristics such as integrated processes or a time varying relationship. These are cases in which viewing the counterfactual estimation problem as a cross-sectional one fails. In this paper, I investigate a new approach to estimate the synthetic control counterfactual incorporating time varying parameters to handle such situations. This is done using a state space framework and Bayesian shrinkage. The dynamics allow for a closer pre-treatment fit leading to a more accurate counterfactual estimate. Monte Carlo simulations are performed showcasing the usefulness of the proposed model in a synthetic control setting. I then compare the proposed model to existing approaches in a classic synthetic control case study.

# Contents

<b>Curriculum Vitae</b>	<b>vi</b>
<b>Abstract</b>	<b>ix</b>
<b>1 Selling Violent Extremism</b>	<b>1</b>
1.1 Introduction . . . . .	1
1.2 Background . . . . .	5
1.3 Data . . . . .	6
1.4 Econometric specification . . . . .	12
1.5 Results . . . . .	15
1.6 Mechanism . . . . .	21
1.7 Conclusion . . . . .	23
1.8 Tables and Graphs . . . . .	25
<b>2 Does Deplatforming Work?</b>	<b>37</b>
2.1 Introduction . . . . .	37
2.2 Background . . . . .	40
2.3 Theoretical Effects of Deplatforming . . . . .	43
2.4 Data . . . . .	48
2.5 Model and Estimation . . . . .	53
2.6 Change in Bitcoin Revenue . . . . .	56
2.7 Change in Viewership . . . . .	60
2.8 Mechanisms . . . . .	61
2.9 Conclusion . . . . .	62
2.10 Tables and Graphs . . . . .	64
<b>3 Synthetic Control with Time Varying Parameters</b>	<b>71</b>
3.1 Introduction . . . . .	71
3.2 Setup . . . . .	75
3.3 Estimation of Parameters and Counterfactual . . . . .	80
3.4 The Posterior Estimation (MCMC) . . . . .	82

3.5	Simulation Studies . . . . .	86
3.6	Empirical Results: Revisiting California Tobacco Tax . . . . .	92
3.7	Discussion . . . . .	93
3.8	Conclusion . . . . .	94
3.9	Algorithms, Tables and Graphs . . . . .	95
<b>A</b>	<b>Appendix for “Selling Violent Extremism”</b>	<b>102</b>
A.1	Descriptive map of Oath Keepers joining . . . . .	103
A.2	Plotted donor pool . . . . .	104
A.3	Counterfactual Robustness . . . . .	112
A.4	County level economic inequality analysis . . . . .	125
<b>B</b>	<b>Appendix for “Does Deplatforming Work?”</b>	<b>131</b>
B.1	Graphs . . . . .	132
B.2	Theoretical Effects of Deplatforming under Constant Elasticity of Substitution (CES) Utility . . . . .	132
B.3	Data collection funnel chart . . . . .	135
B.4	Topic Analysis . . . . .	136
B.5	Robustness Tests . . . . .	141
B.6	Mechanisms Continued . . . . .	143
B.7	Additional Pretrends . . . . .	146
B.8	Leave-One-Out TWFE Estimation . . . . .	149
B.9	Recreation of Main Results using Inverse Hyperbolic Sine . . . . .	151
B.10	Main Specification Estimation with Omitting Never-Treated Channels . . . . .	152
<b>C</b>	<b>Appendix for “Synthetic Control with Time Varying Coefficients: A state-space approach with Bayesian shrinkage”</b>	<b>155</b>
C.1	Additional Distributions . . . . .	155
C.2	Additional Figures for TVP simulation . . . . .	156
C.3	Empirical Monte Carlo Simulation Example - Tennessee . . . . .	156
C.4	Simulation Tables - Empirical Monte Carlo . . . . .	158

# Chapter 1

## Selling Violent Extremism

### 1.1 Introduction

Domestic violent extremism has recently migrated from the fringes of American society to a major security threat (of Homeland Security 2020). Incidents of domestic terrorism in the United States, a subcategory of domestic violent extremism, increased by 357% from 2010 to 2021 with 42 states experiencing at least one incident (U. S. Government Accountability Office 2023). Individuals who participate in extremist activities are rarely unaffiliated. Nearly 70% of extremists surveyed in the *Profiles of Individual Radicalization in the United States* belong to a formal or informal extremist organization or movement (START 2023). Additionally, 60% of felony criminal cases involving illegal political violence in the United States are filed against individuals affiliated with a known extremist group (Loadenthal et al. 2023). The prevalence of organizational affiliation in extremist activity raises two crucial questions: i) What recruitment tactics do domestic violent extremist organizations employ? and ii) How effective are these strategies in recruiting new members?

I study these questions by leveraging a leaked membership database and internal

chat forum detailing the activities of the Oath Keepers, the largest militia in the United States and an instrumental actor in multiple armed government standoffs, including the insurrection on January 6<sup>th</sup> at the United States Capitol. The database begins with the inception of the organization in 2009 and includes the date and reported location of nearly 25,000 people who first signed up and paid membership dues online, the predominant method for joining.

Surprisingly, the Oath Keepers sold memberships similarly to a firm selling a product: They offered discounts on yearly membership dues, sponsored events by planning and participating in armed government standoffs, and practiced traditional advertising by funding a NASCAR driver. Using a synthetic control framework, I estimate that membership discounts lead to an 114% increase in new members during the discount and armed Oath Keepers events lead to an 160% increase during the events. The five membership discounts studied caused over 520 new members to sign up during the discount, equating to approximately \$15,000 in revenue while the two armed government standoffs studied increased new membership by 1,000 people, or approximately \$42,000, during the standoffs. The effects of both tactics only last until shortly before or after the conclusion of the events. The NASCAR sponsorship caused a decrease in new membership. This counterintuitive finding may be driven by potential new members responding to advertising that showcases a group's ideology, such as armed standoffs, more so than attempts to increase name recognition in the short run. Together, these results suggest that armed conflict was a successful tactic to attract new members immediately.

I do not find strong evidence that the increase in membership during events is caused by inter-county nor intra-county wealth inequality. I analyze the difference in new enrollments between the 25% richest counties and the 25% poorest independently for all events, where wealth is measured using the median household income and the income

inequality ratio.<sup>1</sup> When measuring inequality by median household income, the estimates are insignificant regardless of weighting, though somewhat noisy. When measuring by the income inequality ratio, only the effect for two sales and the callout events are statistically significant depending on the weighting scheme suggesting a weak relationship at best between recruitment and economic inequality during events. These findings support past descriptive research that economic factors are not significant predictors of right-wing extremism in the United States (e.g. Piazza 2016). Policies aimed at addressing domestic violent extremism through blanket economic interventions may target the wrong areas ineffectively.

My findings contribute to the initial steps of crafting effective policy to mitigate the spread of violent extremism by providing new insights into the recruitment of both violent and nonviolent members. While only some individuals who joined the Oath Keepers during the tactics went on to perpetrate violent acts,<sup>2</sup> simply belonging to such an organization can increase the propensity to become radicalized through the group (Carvalho & Sacks 2022) or through a radicalization pipeline (Nilsson 2022). A nonviolent member may also increase extremism by recruiting others with a higher propensity for violence because of the contagion-like growth of these movements (Youngblood 2020). Therefore, effective policy should focus on mitigating recruitment of both violent and nonviolent potential new members.

The Oath Keepers' firm like behavior suggests the predominant modeling of extremist organizations (Iannaccone 1992, Berman & Laitin 2008, Berman 2009, Morales et al. 2018) may not apply to this group. Rather than being hypersensitive to leaks, infiltration,

---

<sup>1</sup>Following the Federal Reserve, economic inequality is measured as the ratio of the mean income in the highest quintile of earners divided by the mean income in the lowest quintile of earners in a particular county (Bureau 2010).

<sup>2</sup>Kenneth Harrelson, one of the members in the database, first joined the Oath Keepers during a membership discount. He was sentenced to four years of prison for crimes related to the January 6<sup>th</sup> insurrection (Department of Justice 2023b). I find that 75% of the individuals that joined the day Harrelson joined did so because of the discount.



and defection, as is common for groups perpetrating political violence, I find a financial motivation for the national organization’s ostensibly ideological activity. These findings are also supported by firsthand accounts of former Oath Keepers employees and individuals close to the leader (Levine 2022, Wilson 2022). Of the small literature that models violent groups as having firm-like behaviors (Fryer & Levitt 2012, Shapiro 2013), none have quantified the efficacy of tactics as is done in this paper. Even though court rulings imply that the Oath Keepers could be viewed as a terrorist organization (Department of Justice 2023a), they appear to operate differently from our previous understanding of such groups.

Finally, my findings add a new granularity to the study of domestic violent extremism. Obtaining information on domestic violent groups, even as general as the number of members per year, is difficult because of the culture of skepticism toward outsiders (Crost 2021, Williams et al. 2022). Researchers instead rely on the Southern Poverty Law Centers’ count of extremist organization chapters per state or Global Terrorism Database’s reporting on domestic terrorist acts as dependent variables (Jefferson & Pryor 1999, Mulholland 2010, Piazza 2016, Mulholland 2012, Savage & Wimmer 2023). A notable exception is Van Dijke & Wright (2021), which studies determinants of participation in the January 6<sup>th</sup> insurrection using geolocated cell phone data. My study is the first economic analysis to utilize the leaked Oath Keepers’ member database, which allows me to study organizational strategies at a specificity previously infeasible to researchers.

The remainder of the paper proceeds as follows: Section 1.2 provides additional background on the Oath Keepers. Section 1.3 summarizes the data used in both the causal analysis and county-level analysis. Section 1.4 introduces the econometric model for the causal analysis and investigates the plausibility of the assumptions. The results and robustness checks are presented and discussed in Section 1.5. Section 1.6 studies the differential effects of inflows across intra-county and inter-county wealth inequality. Finally,

Section 1.7 concludes by summarizing the main findings.

## 1.2 Background

The Oath Keepers is a far-right paramilitary organization founded on April 19, 2009 (SPLC 2022). The organization rose to prominence within the far-right community reporting 35,000<sup>3</sup> dues-paying members and being involved in the January 6 insurrection at the United States Capitol. They focus on recruiting former first responders and veterans, although anyone who can pay membership dues can join. The organization has been referred to as “. . . *exemplify[ing] a style of American politics that views violence as a legitimate means to achieve political goals, at least under certain conditions*” (Jackson 2020) and “. . . *loosely structured, lack[s] a rigid ideological focus, and [is] united by things [it] opposes. . . rather than any central tenet*” (Valasik & Reid 2021).

Testimonies by former members and employees of the Oath Keepers suggest the national organization was profit motivated (Kalen Hill 2021, Levine 2022, Wilson 2022). The Oath Keepers raise funds through membership dues, donations, and selling merchandise. They offer a monthly, yearly, and lifetime membership, with the majority of members choosing to buy the annual membership. These funding strategies differ from most politically violent organizations studied, which tend to fund themselves through illegal activities.

Prior to 2013, the organization did not receive much public attention. Their first major event occurred on April 2014 when Oath Keepers traveled to Bundy Ranch in Bunkerville, NV to support Clive Bundy in his conflict with the Bureau of Land Management (Jackson 2020). The Oath Keepers participated in two additional armed disputes with federal law enforcement similar to Bundy Ranch: one in southern Oregon and

---

<sup>3</sup>While the Oath Keepers reported 35,000 individuals on their roster, independent estimates suggest they never had more than 5,000 active members at a single point in time (Jackson 2020).

another in Montana. The organization became more active after the election of Donald Trump, serving as security forces for other far-right extremist groups and political candidates (Cheney 2022).

The organization was the focus of discussion at the United States House of Representatives Select Committee to Investigate the January 6<sup>th</sup> Attack on the Capital on July 12, 2022.<sup>4</sup> As of August 8, 2022, The Anti-Defamation League had identified 81 individuals currently holding or running for public office in 2022, 373 law enforcement employees and 117 individuals currently serving in the U.S. military who are included in the Oath Keepers membership database (ADL 2022).

The organization’s founder, Stewart Rhodes, had near-complete control over the national organization. Rhodes and other Oath Keepers have been convicted of seditious conspiracy related to their involvement with the January 6, 2021 insurrection at the United States Capitol. The court determined Rhodes’ conduct was terrorism and sentenced him to 18 years in prison (Department of Justice 2023a).

## 1.3 Data

Data for the main analysis comes from a leaked internal membership roster and Google Trends. In September of 2021, the organization’s internal membership records were hacked and made available to academics through *DDOSecrets*, a journalist 501(c)(3) nonprofit focused on publishing leaked data (*Distributed Denial of Secrets* n.d.). Each row of the roster contains an individual’s Oath Keepers membership ID, membership type, name, physical and email address, and join date. The columns were not labeled. While columns such as name and address were obvious, the join date column was identified by comparing the date individuals said they joined in the forum to the recorded date in the

---

<sup>4</sup>A transcript of the hearing can be found at <https://www.npr.org/2022/07/12/1111123258/jan-6-committee-hearing-transcript>.

roster.

In total, there are 37,976 rows in the database. Each individual in the database paid dues at least once (ADL 2022). I limit the analysis to individuals who bought an annual membership and that have a recorded join date. This was the predominant membership type, with 86.5% of the sample buying an annual membership. Due to recording errors, some members do not have their join dates recorded or have join dates before the organization began. Those individuals are dropped from the analysis. This process removes 12,920 individuals leaving 25,056 new Oath Keepers with join dates. All data is aggregated to the daily level. Unfortunately, the Oath Keepers' membership roster does not include recurring payments, which means individual's tenure in the organization cannot be studied. Therefore, all results pertain to the effect of discounts on *inflow* of new members, not retention.

The leaked documents also include all messages on an internal forum. Based on forum posts, the recording system for the Oath Keepers was formalized around 2013. In addition, the Oath Keepers updated their site starting November 2018, which caused significant back-end issues, such as members not receiving their welcome packages, having trouble accessing online sources, and recording errors in their database. To ensure that the findings aren't driven by recording issues, I limit my analysis to between January 1, 2013 and November 1, 2018. In total, there are 20,447 new Oath Keepers and join dates in the dataset in this time frame.

Membership expanded across the whole country with 79.2% of counties having at least one individual sign up for the Oath Keepers during this time period. Counties in the Southwest and Northwest had the largest membership. The top three most populous Oath Keepers counties measured in new recruits per capita were in Idaho, Oregon and Montana while the top three most populous Oath Keeper counties measured in total number of Oath Keepers were in Arizona and Nevada. A map of Oath Keepers membership

by county is presented in the appendix.

### 1.3.1 Selection of Tactics

I limit the tactics to Oath Keepers events that have explicit start and end dates, were supported by the leader, and follow two weeks during which no other Oath Keepers events occurred. These selection rules include most types of events the Oath Keepers organized. Notable exceptions are emergency response callout events and nonviolent gatherings, which tend to have a national call to action to initiate, but no official end dates. Examples include a call to action on January 17, 2017 to help tornado victims in Georgia and one on April 12, 2018 to attend a second amendment rally at any state capitol. Details on these types of events, including the level of participation, are ambiguous due to a lack of reporting on the internal forum and by outside sources.

Table 1.1 reports the start dates and end dates of all the events used in this analysis.

[Table 1.1 about here]

#### *Membership Discounts*

Annual dues cost \$40. A membership discount is when the organization temporarily reduced the cost of joining by approximately 25%. Seven membership discounts were identified. Each discount start and end date was first identified in the forum and corroborated by looking for changes in price on their website via the Wayback Machine.<sup>5</sup> Two discounts were dropped due to having less than a two week pre-period between events. The discounts tend to cluster around patriotic holidays including Veteran's Day and Memorial Day.

---

<sup>5</sup>The Wayback machine saves instances of a website at a given point in time. A directory of instances saving the Oath Keepers' main page is at: [https://web.archive.org/web/20091201000000\\*/oathkeepers.org](https://web.archive.org/web/20091201000000*/oathkeepers.org)

*Armed Callout Events*

An armed callout is when the Oath Keepers organized and participated in armed standoffs with federal agents. From 2013-2018, the Oath Keepers participated in armed standoffs at Bundy Ranch and Big Sky. Both events involved the Oath Keepers occupying territory. A third armed standoff occurred at Sugar Pine Mine in Josephine County, Oregon. Although the local Oath Keepers chapter voted unanimously to participate in the standoff (Tatenhove 2023), the Oath Keepers national leadership was opposed to the activity (Jackson 2020). I omit the standoff because the official callout on the internal forum was not made by Rhodes.

The Oath Keepers also sponsored two national armed security events without a national gathering. Both were in response to shootings, one at a school and the other at a military recruiting center. Military recruiters were instructed to treat the armed civilians, such as the Oath Keepers, as a threat and call the police (Tritten 2015). These events are omitted because of ambiguity in the length of the events and intensity of participation.

Finally, I do not include the Oath Keepers' activity during the civil unrest in Ferguson, Missouri, because it was not an official callout event.

*Traditional Advertisement*

In 2013, the Oath Keepers sponsored NASCAR driver Jeffrey Earnhardt. The Oath Keepers logo was prominently displayed on the hood of his car for four races beginning May 4, 2013 and ending July 13, 2013.<sup>6</sup>

The Oath Keepers' also ran a billboard campaign starting in 2013.<sup>7</sup> Unfortunately, the details about the length of time and location of billboards is provided for some, but not all, via the forum records. I omit the campaign from the analysis due to this

---

<sup>6</sup>See <https://www.racing-reference.info/sponsor-search/>.

<sup>7</sup>See <https://web.archive.org/web/20130822072521/http://oathkeepers.org/oath/billboard/>.

ambiguity.

Figure 1.1 plots the number of new Oath Keepers per day two weeks leading up to, during, and two weeks after the events. The number of new Oath Keepers increases during membership discounts and callout events, with over 150 new members signing up on one day during the Big Sky callout event. Events vary in length from as short as three days to months long.

[Figure 1.1 about here]

### 1.3.2 Control groups

The econometric specification described in Section 1.4 relies on creating a counterfactual from organizations similar to the Oath Keepers. I use three sources of information to identify these groups. The first is the *Center for International Security and Cooperation's Global Right-Wing Extremism Map* (for International Security & Cooperation 2022), which maps how far right extremist groups are related to one another. The second is the Southern Poverty Law Center's yearly summary of active patriot groups in the United States. I limit the sample to organizations that have a similar presence to the Oath Keepers across the country. Finally, similar groups are also found using the Armed Conflict Location & Event Data Project (ACLED) report on right-wing militia groups in the United States (Raleigh et al. 2020). The report identifies nine large, cross-state right wing militia groups including the Oath Keepers. I then collect **Google Trends** data on the organizations to use in identifying the causal effect of membership discounts.

Table 1.2 shows the organizations used as a donor pool in each tactic. An "X" signifies if the organization is included in the donor pool for that specific event. The John Birch Society is an antigovernment movement first started in the 1950s known for spreading

conspiracy theories (SPLC 2021). The Eagle Forum and We are Change are also listed as antigovernment movements (SPLC n.d.). The Three Percenters are another militia group that International Security and Cooperation (for International Security & Cooperation 2022) reports as an ally to the Oath Keepers. Finally, the Proud Boys is described as “*a nation-wide right-wing street movement*” (Raleigh et al. 2020). The leader of the Proud Boys and four other members were also indicted for seditious conspiracy related to their involvement at the capital on January 6<sup>th</sup> (Department of Justice 2022). The remainder of the groups were analyzed in Raleigh et al. (2020).<sup>8</sup>

[Table 1.2 about here]

*Eagle Forum*, and *We are Change* are only included in the donor pool when studying earlier tactics because the organizations significantly shrunk by mid-2017 and were no longer viable choices for the donor pool. The Proud Boys was founded in 2017.

### 1.3.3 County-level demographics

Section 1.6 studies the differential effects of tactics on the wealthiest and poorest counties. Wealth is measured using the median income and the income inequality metric, measured as the ratio of the mean income for the highest quintile divided by the mean income for the lowest quintile for each county. Both metrics for each county per year comes from the FRED database. Finally, population estimates come from the Census Bureau.

---

<sup>8</sup>The SPLC also includes the American Contingency, Light Foot Militia and Civilian Defense Forces in their list (SPLC n.d.). The American Contingency was founded in 2020, after the last tactic studied. Light Foot Militia and Civilian Defense Forces are omitted due to a lack of information on the groups’ founding dates.



## 1.4 Econometric specification

### 1.4.1 Setup

I estimate the causal effect of events on membership using a synthetic control framework. The synthetic control framework creates a counterfactual for Oath Keepers' new members using untreated times series. The causal effect of events on new Oath Keepers membership at time  $t$  can be represented as:

$$\tau_{i,t} = y_{i,t}(1) - y_{i,t}(0)$$

where  $y_{i,t}(1)$  is the number of new members joining organization  $i$ , in this case the Oath Keepers, had an event occurred and  $y_{i,t}(0)$  is the number of new members in the absence of an event. The goal is to estimate  $y_{i,t}(0)$  as a function of time series correlated with membership but independent of discounts:  $y_{i,t}(0) = g(\mathbf{y}_t)$  where  $\mathbf{y}_t = [y_{1,t}, \dots, y_{i-1,t}, y_{i+1,t}, \dots, y_{N,t}]$  are the Google Trends for similar organizations. The trends capture changes in overall popularity of similar organizations while being unaffected by the Oath Keepers' events.

I estimate the counterfactual for each event independently using the Brodersen et al. (2015) model, which builds off of Bayesian Structural Time Series (Scott & Varian 2013). The approach does not rely on asymptotic results, a major benefit with a short pre-treatment period, and was designed to identify causal effects using Google Trends data. I assume the following state space framework:

$$g_{\mathbf{v}}(\mathbf{y}_t) = \beta_{0,t} + \sum_{j \neq i} \beta_j y_j + \epsilon_t \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2) \quad (1.1)$$

$$\beta_{0,t} = \beta_{0,t-1} + \eta_t \quad \eta_t \sim \mathcal{N}(0, \theta) \quad (1.2)$$

where  $\mathbf{v} = [\beta_1, \dots, \beta_{i-1}, \beta_{i+1}, \dots, \beta_N, \theta, \sigma^2]$  are the parameters. To avoid overfitting the data, regularization is induced on the coefficients through the Bayesian shrinkage prior slab and spike. The approach induces sparsity in the coefficients similar to Abadie et al. (2010) approach without requiring the treated unit to be within the convex hull nor the coefficients to be positive and sum to one. In a standard synthetic control setting, where the donor pool and outcome time series are in the same units, weights outside of the zero to one interval may be viewed as problematic because the estimator is relying on extrapolation. Allowing the weights to be greater than one or less than zero in this paper's setting is both necessary and beneficial because the outcome time series is in a different unit than the donor pool meaning there is no reason to impose the convex hull. Finally, I include a local-level trend for the intercept,<sup>9</sup> and  $\theta$  and  $\sigma^2$  follow an inverse gamma distribution.

Creating the counterfactual consists of three steps. In the first step, the posterior distribution of the parameters is estimated via Gibbs sampling. The number of iterations is set to 20,000 with 1,000 burn-in iterations. After the burn-in, each draw from the Gibbs sampler is used to predict the missing potential outcome in the post-treatment periods. In the final step, the difference between observed and predicted values is recorded. Each step of the Gibbs sampler creates an estimated treatment effect, generating an empirical posterior distribution.

---

<sup>9</sup>This specification is the preset for the R package of Brodersen et al. (2015), though the paper discusses a local-linear trend.

## 1.4.2 Identifying assumptions

When constructing a counterfactual in the synthetic control framework, there is a balance between having a long enough pre-treatment to estimate the parameters but not too long such that the plausibility of the data generating process comes into question. Concerns include the relationship between the treated and explanatory times series varying over time, political activity outside the data generating process that could cause structural breaks in the time series, or different regional/national Oath Keepers activity.<sup>10</sup> To mitigate these concerns, I limit my window to two weeks before and after each tactic. After surveying the time periods surrounding the tactics, I failed to identify major political activity or additional Oath Keepers events occurring two weeks prior to each tactic. I relax this restriction in Section 1.5.1.

Identification in the synthetic control framework requires the proposed data generating process to be a “good” fit. I investigate this claim by artificially moving the treatment date forward into the pre-treatment window and compare the estimated treatment effect to zero, the true treatment effect in the placebo period.

Figure 1.2 artificially moves the date of treatment forward seven days.<sup>11</sup> Seven days are used to fit the counterfactual model and seven days to test it. The grey line with dots is the inflow of new Oath Keepers while the line without dots is the constructed counterfactual with 95% credibility intervals.<sup>12</sup> The first vertical dashed line is the beginning of the placebo test (seven days prior to treatment). Excluding the Flash Discount, these results suggest the estimation approach is accurately capturing the latent factors of Oath

---

<sup>10</sup>For example, the Oath Keepers performed a different national event 16 days prior to the Big Sky armed standoff.

<sup>11</sup>Using half of the pre-treatment periods to conduct a placebo-in-time follows from the analysis performed in Abadie et al. (2014).

<sup>12</sup>A credibility interval is a range of parameters that account for a certain portion of the posterior distribution. In this setting, it is the parameters that correspond to 95% of the posterior distribution of the treatment effect.

Keepers membership and the effect can be interpreted as causal. In the case of the Flash Discount, the estimates should be viewed as suggestive rather than causal.

[Figure 1.2 about here]

The other concern is the relevance of the donor pool. If the **Google Trends** data on similar organizations is not contributing to the donor pool, then the counterfactual estimate collapses to a local-level state space model. Figure 1.3 plots the inclusion probability and coefficient values. Organizations with missing bars were not used in the counterfactual estimate. The intercept is rarely included, suggesting the counterfactual is not being driven by a local-level model.

[Figure 1.3 about here]

## 1.5 Results

### 1.5.1 Effect of tactics on inflow of new members

Figure 1.4 plots the observed Oath Keepers' membership and the constructed counterfactual. The average number of individuals signing up for the Oath Keepers increases during the tactics, except for the NASCAR advertisement. Although the recruitment is in general higher during these periods, there are sporadic days of lower or similar membership. The counterfactual spikes during the Constitution Day Discount sale. This is most likely driven by the Unite the Right rally in Charlottesville, Virginia from August 11 - August 12, 2017.

After the discount ends, the inflow of new members returns to pre-discount levels. This supports the claim that the discounts are the main driver of increased membership

in the time period. Initially, callout events have no detectable effects on increased membership. After a few days, new membership increases drastically. The callout events do not lead to a permanent change in the inflow of new members: the significant effects dissipate either before or immediately after the callout event ends.

[Figure 1.4 about here]

Table 1.3 presents the effect of tactics on new Oath Keepers membership. The results are presented as the percent increase in recruits due to the tactic, the average increase per day during the duration of the tactic, and the total increase in new members attributed to the tactic. Excluding the Flash discount and Constitution Day discount, the Oath Keepers recruited an additional 527.98 annual members due to all the discounts which equates to an additional \$14,598 from initial signups. Similarly, the two callout events led to an additional 1,051 annual members equating to \$42,050 additional revenue from initial signups. The NASCAR advertisement campaign caused an average decrease of about eight members per day, effectively balancing out the gains in membership made during the discounts.

[Table 1.3 about here]

The large effects for callout events can be rationalized through the *complementary* view of advertising (Bagwell 2007). Following this perspective, advertising directly enters a consumer's preferences and complements the product, which is buying an Oath Keepers membership in this situation. If potential members value being in organizations that embody certain ideologies, then advertisements can further publicize those behaviors serving as a signal. This may be why armed government standoffs, a strong signal of specific tenets the group wants to associate with, were successful in recruiting new members.

Conversely, advertising on a NASCAR more aligns with the *informative* view of advertising (Bagwell 2007). This action would be most successful in increasing name recognition rather than signaling ideological principles. My findings suggest a lack of name recognition was not hampering the growth of the organization during this event. If anything, the sponsorship inadvertently signaled something to consumers that dissuaded them from joining the organization.

The findings suggest the inflow of new Oath Keepers is highly sensitive to price and reacts positively to violent extremism. The large increase in new members from the callout events provides an alternative explanation for ostensibly ideological activities: The Oath Keepers can profit off of violent extremist activity. It also generates more revenue than advertising and membership discounts combined.

## Robustness tests

### *Pre-treatment length*

My preferred specification limits the pre-treatment to two weeks to mitigate the risk of contamination in the construction of the counterfactual. Specifically, two weeks ensures there are no other Oath Keepers national events nor major political announcements occurring during the pre-treatment window. For example, there appears to be a structural break inflow of new Oath Keepers four weeks prior to the Memorial Day discount<sup>13</sup> and a different callout event occurred 16 days prior to the Big Sky Callout. As a robustness check, I extend the pre-treatment window to three and four weeks to investigate the sensitivity of the results. A table of treatment effects is provided in the online appendix.

The increased pre-treatment period suggests that the Constitution Day discount did lead to a positive significant increase in new membership, but is masked by the Unite the Right rally. The treatment effect for the Constitution Day discount using only the

---

<sup>13</sup>See the online appendix for the graph.

days leading up to the Unite the Right rally is on average 4.5 new members per day. The increased pre-treatment period had little to no effect on all other tactics.

#### *Alternative estimation strategies*

I compare the findings to the following alternative synthetic control estimation strategies: Ferman & Pinto (2021), Carvalho et al. (2018), Xu (2017), and Klinenberg (2022). Ferman & Pinto (2021) restricts the weights to be nonnegative and sum to one, as first suggested by ?, with the addition of an intercept. Carvalho et al. (2018) and Xu (2017) both create a counterfactual of treated units to estimate the treatment effect with inference derived from asymptotic results, not prior distributions. Finally, Klinenberg (2022) builds off of Brodersen et al. (2015) by allowing for time varying parameters. Additional details are provided in the time-varying robustness subsection below.

I compare the models based on mean squared forecast error using the following placebo specification: I fit the models on days 8-14 prior to treatment, then compare the mean squared forecast error for days one to seven.<sup>14</sup>

[Table 1.4 about here]

Table 1.4 presents the mean squared forecast error for the seven days leading up to a tactic. The main specification and Carvalho et al. (2018) have the smallest mean squared forecast errors, and tend to be similar to one another. Ferman & Pinto (2021) and Xu (2017) have a large forecast error in comparison to the other models. Klinenberg (2022) creating estimates with similar MSFE as the preferred specification suggests that the additional flexibility of the model will lead to more model uncertainty, as seen in wider credibility intervals, with negligible reductions in bias. Based on this, I rerun the analysis using only Carvalho et al. (2018).

---

<sup>14</sup>The results hold fitting the models on weeks two and three (days 8-21) and weeks two through four (days 8-28) prior to the tactic beginning. See the appendix for further details.

Table 1.5 presents the average treatment effects for the main specification from Table 1.3 along with estimates for Carvalho et al. (2018). The point estimate and significance levels are similar to the preferred specification, with two notable differences. First, I find evidence of an increase in the inflow of new members during the Constitution Day discount using Carvalho et al. (2018). The effect is comparable to the treatment effect using the main specification prior to the Unite the Right rally. Second, Carvalho et al. (2018) produces similar point estimates to the main specification when studying the Big Sky Callout, but with larger confidence intervals.

[Table 1.5 about here]

### *Endogeneity of tactics*

Another concern may be that discounts occur during already popular membership recruitment periods, callout events are strategically planned around seasonal trends, and NASCAR happens to occur during a time when individuals have a lower propensity for membership. I investigate these concerns by rerunning the synthetic control estimation for all tactics between 2013 and 2018. A counterfactual for Veteran’s Day discount, Constitution Day discount, and Big Sky callout in 2018 cannot be calculated because no new Oath Keepers joined in the pre-treatment period.

[Table 1.6 about here]

Table 1.6 displays the treatment effects for the placebo years. In general, the treatment effects are small and insignificant. None of the placebo years are significant for Veteran’s Day and the Big Sky callout. One placebo test is significant for Bundy Ranch, but the point estimate is far smaller than the actual treatment effect. The placebo effects for the Christmas discount are not consistently positive nor negative suggesting



there was not a seasonal component driving membership. In conclusion, I fail to find strong evidence of seasonal endogeneity occurring during the time of the discounts.

#### *Time-varying relationship between Oath Keepers' membership and Google Trends*

The preferred specification models the relationship between the Google Trends and Oath Keepers membership inflow data as time-invariant. While relationship between the times series may change over time, the method in which Brodersen et al. (2015) allows for time-varying parameters leads to overspecification.

To investigate the plausibility of a time invariant relationship between the time series, I implement the method proposed in Klinenberg (2022). Klinenberg (2022) decomposes the time varying parameter into a time-varying and time invariant component following  $\beta_j = \alpha_j + \gamma_j$ ,<sup>15</sup>. The relationships between time series are thus allowed to be time varying with a nonzero mean, time-invariant with a nonzero mean, time varying with zero mean, and time-invariant with zero mean. While Klinenberg (2022) does improve on Brodersen et al. (2015) specification with time varying parameters, the additional flexibility of the model produces more uncertainty in the posterior estimates compared to a time-invariant model when the true data generating process does not include time varying parameters. Therefore, this model should only be used if the researcher suspects the true data generating process does include time varying relationships.

I investigate the plausibility of time varying parameters by calculating Klinenberg (2022) and plotting the parameters, with graphs for two, three, and four week pre-treatment periods provided in the appendix. Under this model specification, I fail to find evidence of a statistically significant time varying parameter, suggesting that a time invariant model is a plausible approximation for this setting.

#### *Additional robustness tests*

<sup>15</sup>Technically,  $\beta_j$  is replaced by  $\beta_j + \sqrt{\theta_j} \tilde{\beta}_{j,t}$  where  $\tilde{\beta}_{j,t} = \tilde{\beta}_{j,t-1} + \eta_j$  and  $\eta_j \sim \mathcal{N}(0, 1)$ . See Klinenberg (2022) for more details.

I rerun the analysis including monthly and lifetime membership signups. The results are statistically and substantively similar. The findings are presented in the appendix.

## 1.6 Mechanism

Does economic inequality have an effect on recruitment during tactics? Some findings suggest that poverty and unemployment have a strong influence on domestic terrorism (Enders & Hoover 2012) internationally and contributed to the rise of anti-democratic extremism in the United States (Croft 2021), while others find that these factors are not significant predictors of right-wing extremism in the United States (Piazza 2016). However, a lack of data has left the relationship between recruitment and economic inequality is an unresolved question.

I answer this by comparing the inflow of Oath Keepers in the 25% richest counties to the 25% poorest counties, measured using median household income and economic inequality,<sup>16</sup> before and during events. Studying median income levels provides insights into the effect of inter-county wealth inequality on recruitment while the economic inequality measure provides insights on the effects of intra-county inequality. The top 25% richest counties have an average median income between \$13,000-\$15,000 more than the average median income of the 25% poorest counties, depending on the year in which an event occurred. The economic inequality ratio is about 11 for the bottom quartile and 15 for the top quartile.

Economic inequality affecting recruitment inflows during the tactics supports arguments that economic conditions are related to domestic extremism, in this case through the inflow of new members. Conversely, a precisely estimated null effect would sup-

---

<sup>16</sup>Following the Federal Reserve, economic inequality is measured as the ratio of the mean income in the highest quintile of earners divided by the mean income in the lowest quintile of earners in a particular county (Bureau 2010).

port past research arguing that economic conditions, specifically economic inequality, are unrelated to domestic extremism.

I compare the change in recruitment between the bottom and top quartile of income before and during each tactic following a two-way fixed effects specification:

$$y_{i,t} = \beta D_{i,t} + \gamma_t + \gamma_i + \epsilon_i \quad (1.3)$$

where  $y_{i,t}$  are the number of Oath Keepers to join per 100,000 per county on day  $t$ .  $D_{i,t}$  is an indicator equal to one if county  $i$  is in the top 25% and  $t$  is after the start of the event. Each event is estimated separately and standard errors are clustered at the county level.  $\gamma_t$  and  $\gamma_i$  are the time and county fixed effects.  $\beta_j$  estimates the differential effects between the top 25% richest counties and bottom 25% poorest counties, meaning counties in the middle quantiles are dropped from the analysis.

Borrowing from the difference-in-differences terminology, the “pre-period” is the fourteen days leading up to the start of an event while the “post-period” is the duration of the event. The “treated” units in this setting are counties that are in the 25<sup>th</sup> quartile. The analysis does not include days after the events end. I fail to find evidence of diverging trends in the pre-event periods, meaning that the (lack of) difference in recruitment rates during the events can be attributed to the difference in income. The event studies are supplied in the online appendix.

Following Solon et al. (2015)2015, Table 1.7 presents the regression estimates unweighted and weighted by population. Panels A and B use the median household income of a county as a measure of wealth while Panels C and D use the income inequality ratio. Panels A and C are unweighted estimates and Panels B and D weight the regression by county population.

[Table 1.7 about here]

I fail to find evidence of inter-county economic inequality affecting recruitment when measuring wealth using median household income. The effects are precisely estimated for the callout events regardless of weighting. The precise nulls suggests that economic inequality did not cause inter-county differences in recruitment. While I fail to find evidence of inter-county inequality affecting recruitment during discounts and the Nasacr sponsorship, only the Constitution Day Discount and Flash Discount are precisely estimated.

I also fail to find consistent evidence of intra-county economic inequality affecting recruitment when measuring wealth using the economic inequality metric. However, the estimates are imprecise when unweighted, and I find sporadic significance when weighted by county level population. More economically equal counties receive more Oath Keepers during the Veteran's Day Discount, Constitution Day Discount, and Bundy Ranch Callout when weighting by population, but not unweighted. Conversely, I less economically equal countries receive more Oath Keepers during the Big Sky callout when estimated unweighted, but not weighted. The lack of consistent evidence suggests a weak relationship at best between recruitment and economic inequality during events.

These findings suggest that county level economic inequality is not consistently causing differential effects in recruitment. The estimates are robust to measuring inequality continuously and limiting the sample to counties that have at least one individual recorded as joining the Oath Keepers in the data. See the appendix for further details.

## 1.7 Conclusion

The resurgence of domestic violent extremism across Western nations has become a significant area of concern among academic scholars and policy makers alike. A critical

step to crafting policies to mitigate their harmful activity is to understand the efficacy of various recruitment tactics. Drawing upon leaked membership data on the Oath Keepers, I estimate the effects of advertising, membership discounts, and sponsored events on the inflow of new members.

Offering discounts for membership and sponsoring violent events both increase the inflow of new members, while traditional advertising had a negative effect. While sponsoring a NASCAR led to short term reductions in membership, the increased exposure may have normalized the Oath Keepers to a larger audience, leading to more members joining post-sponsorship. Such longer-term effects are left for future research.

I fail to find evidence inter-county income inequality effects the inflow of new members during any event. I also fail to find consistent evidence that intra-county income inequality effects recruitment during events. Together, these findings provide new insights into the efficacy of commonly employed business practices in facilitating the discount of violent extremism.

## 1.8 Tables and Graphs




Figure 1.1: Oath Keepers events. The graphs depict the number of new sign-ups per day. Each pre/post period is 14 days long with the exception of the flash discount post period. Notice the axes differ for each panel.

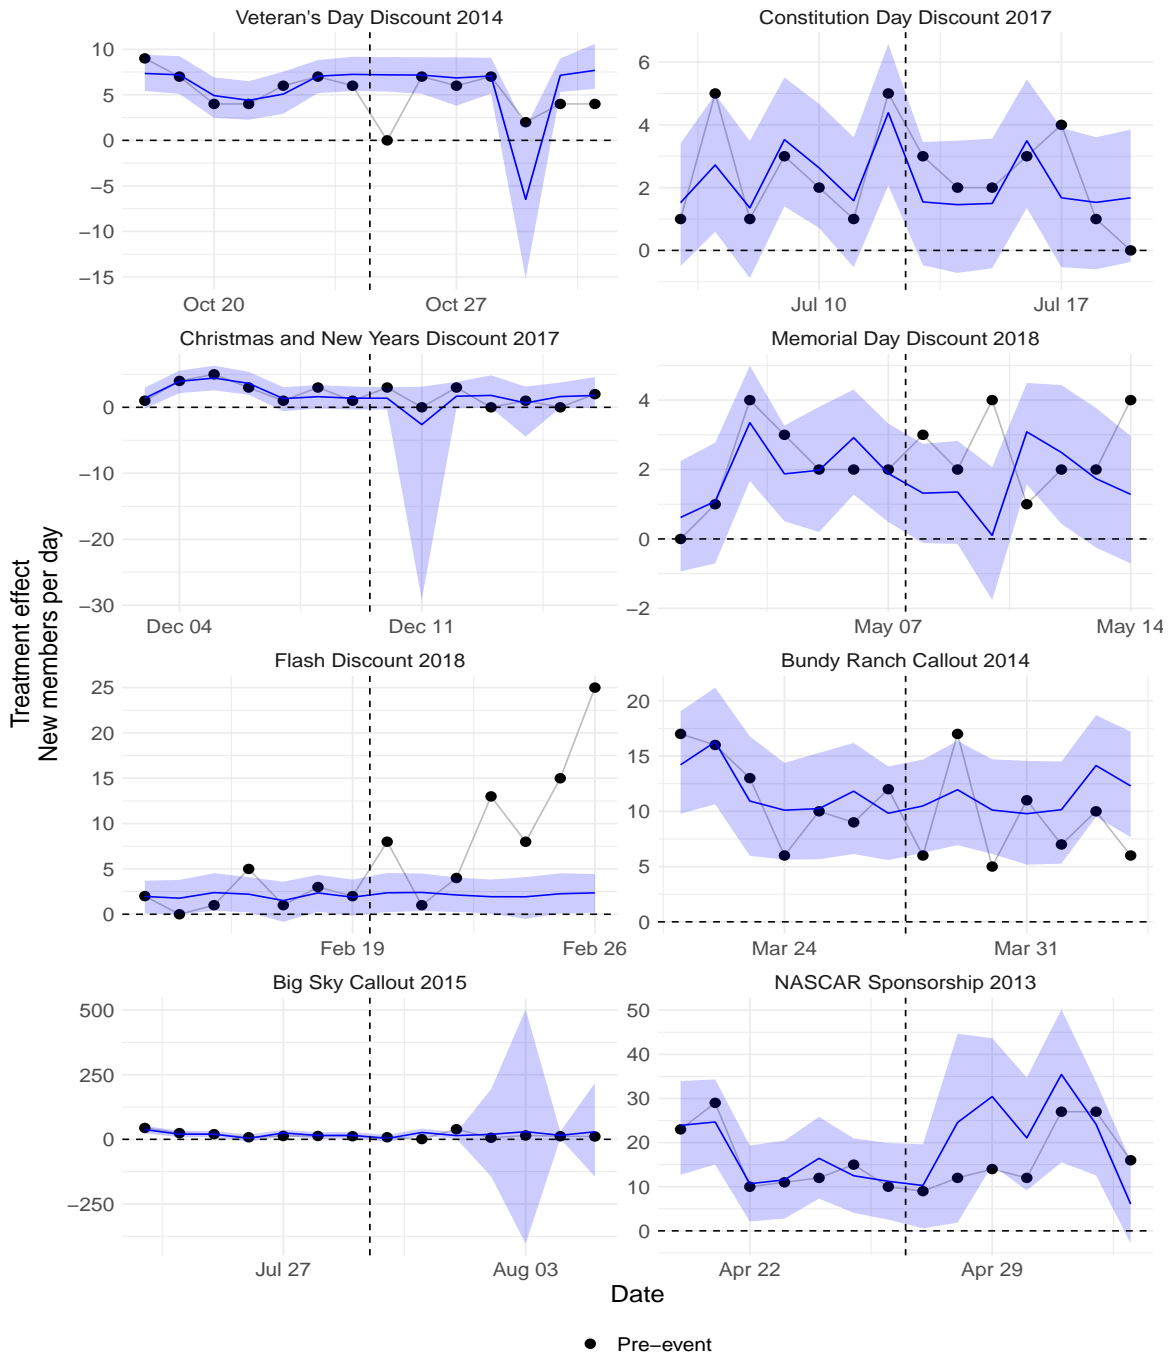


Figure 1.2: Placebo test of Oath Keepers discounts on new membership. The dashed lines show the beginning and end of the placebo period. Window limited to pre-discount. The blue line is the constructed counterfactual with 95% credibility intervals.




Figure 1.3: Contribution of donor pool to each counterfactual.






Figure 1.4: Effect of Oath Keepers tactics on new membership. Blue line is the constructed counterfactual with 95% credibility intervals. Notice the axes differ for each panel.

Table 1.1: Oath Keepers tactics.

	Start Date	End Date	Notes
<i>Membership Discounts</i>			
Veteran's Day Discount	2014-11-01	2014-11-11	25% membership discount from \$40 to \$29.
Constitution Day Discount	2017-07-20	2017-09-17	25% membership discount from \$40 to \$30 and gun giveaway.
Christmas/New Years Discount	2017-12-17	2018-02-09	25% membership discount from \$40 to \$29.95 and gun giveaway.
Flash Discount	2018-02-27	2018-03-01	25% membership discount from \$40 to \$29.95 and gun giveaway.
Memorial Day Discount	2018-05-15	2018-05-23	25% membership discount from \$40 to \$29.95 and gun giveaway.
<i>Armed Callout Events</i>			

Table 1.1: Oath Keepers tactics. (*continued*)

	Start Date	End Date	Notes
Bundy Ranch	2014-04-04	2014-04-27	The event takes place at Bundy Ranch in Clark County, Nevada. It begins with the first arrest and ends when the Oath Keepers left due to fears of a drone strike.
Big Sky	2015-08-06	2015-09-03	The event takes place at White Hope mine in Lincoln, Montana. It begins with the official callout video and ends when the Oath Keepers began another callout.
<i>Traditional Advertisement</i>			
NASCAR Sponsorship	2013-05-04	2013-07-13	Oath Keepers sponsor NASCAR driver Jeffrey Earnhardt for four races in the Xfinity Series with their logo on the hood of the car.

Table 1.2: Donor pool per tactic.

	3 Percenters	John Birch Society	Eagle Forum	We Are Change	Proud Boys	Patriot Prayer
<i>Membership Discounts</i>						
Veteran's Day Discount 2014	X	X	X	X		
Constitution Day Discount 2017	X	X	X		X	X
Christmas and New Years Discount 2017	X	X	X		X	X
Memorial Day Discount 2018	X	X			X	X
Flash Discount 2018	X	X			X	X
<i>Armed Callout Events</i>						
Bundy Ranch Callout 2014	X	X	X	X		
Big Sky Callout 2015	X	X	X	X		
<i>Traditional Advertisement</i>						
NASCAR Sponsorship 2013	X	X	X	X		

*Note:*

Data used in the analysis is Google Trends data. All tactics by organization Google Trends are exported independently. An X signifies an organization is included in the donor pool.

Table 1.3: Effect of tactics on Oath Keepers' recruitment.

	Relative Effect (%)	Average Effect	Cumulative Effect
<i>Membership Discounts</i>			
Veteran's Day Discount 2014	429.43 [326.58, 561.86]	22.88 [21.64, 24]	251.65 [238.09, 264.01]
Constitution Day Discount 2017	3.44 [-29.51, 59.47]	-0.14 [-3.55, 3.16]	-8.5 [-213.06, 189.82]
Christmas and New Years Discount 2017	137.74 [79.62, 228.24]	2.33 [1.81, 2.84]	128.32 [99.74, 156.45]
Memorial Day Discount 2018	164.19 [105.12, 255.44]	3.62 [3.02, 4.23]	32.59 [27.16, 38.09]
Flash Discount 2018	2397.97 [242.61, 3885.66]	38.47 [32.95, 44.3]	115.42 [98.84, 132.9]
<i>Armed Callout Events</i>			
Bundy Ranch Callout 2014	273.73 [209.04, 355.29]	28.21 [26.15, 30.17]	677.1 [627.71, 724.17]
Big Sky Callout 2015	84.17 [39.17, 153.37]	12.9 [8.14, 17.51]	374.14 [236.15, 507.87]
<i>Traditional Advertisement</i>			
NASCAR Sponsorship 2013	-49.67 [-59.63, -35.62]	-8.11 [-11.82, -4.43]	-576.13 [-838.88, -314.23]

*Note:*

The relative effect is in terms of percent change. The average effect is the average number of new Oath Keepers per day due to the discount during the discount while the cumulative effect is the total number of new Oath Keepers due to the discount during the discount. Brackets are 95% credibility intervals. The placebo test for the Flash discount suggests the estimated counterfactual did not accurately approximate the underlying data generating process. The results for the Flash discount should be interpreted as suggested, not causal.

Table 1.4: Mean squared forecast error of alternative models using first seven days prior to tactic. Counterfactual estimates are fitted to days 8-14 prior to a tactic.

	Main Specification	Ferman and Pinto (2021)	Carvalho et al. (2018)	Xu (2017)	Klinenberg (2022)
<i>Membership Discounts</i>					
Veteran's Day Discount 2014	21.1	1393.9	9.4	775.2	17.0
Constitution Day Discount 2017	1.6	3.4	1.7	24.9	3.5
Christmas and New Years Discount 2017	2.5	866.3	2.8	80.0	4.1
Memorial Day Discount 2018	4.3	7.4	1.4	7.9	6.7
Flash Discount 2018	124.5	91.3	128.0	83.1	81.2
<i>Armed Callout Events</i>					
Bundy Ranch Callout 2014	20.0	153.4	24.3	226.6	75.6
Big Sky Callout 2015	287.7	738.9	165.0	1778.6	184.1
<i>Traditional Advertisement</i>					
NASCAR Sponsorship 2013	98.5	756.7	92.2	1385.7	162.1

*Note:*

See the appendix for results fitting the models on weeks 2-3 and weeks 2-4.

Table 1.5: Average effect of Oath Keepers tactics using alternative estimation strategies.

	Main Specification	Carvahlo et al. (2018)
<i>Membership Discounts</i>		
Veteran's Day Discount 2014	22.88 [21.64, 24]	12.31 [4.48 20.13]
Constitution Day Discount 2017	-0.14 [-3.55, 3.16]	5.02 [3.45 6.6]
Christmas and New Years Discount 2017	2.33 [1.81, 2.84]	2.18 [1.13 3.23]
Memorial Day Discount 2018	3.62 [3.02, 4.23]	2.89 [0.58 5.2]
Flash Discount 2018	38.47 [32.95, 44.3]	13.13 [2.9 23.35]
<i>Armed Callout Events</i>		
Bundy Ranch Callout 2014	28.21 [26.15, 30.17]	18.34 [6.99 29.7]
Big Sky Callout 2015	12.9 [8.14, 17.51]	8.11 [-7.46 23.68]
<i>Traditional Advertisement</i>		
NASCAR Sponsorship 2013	-8.11 [-11.82, -4.43]	-8.13 [-11.25 -5.01]

*Note:*

Brackets are 95% credibility/confidence intervals.

Table 1.6: Average effect of Oath Keepers tactics using alternative years.

Discount	2013	2014	2015	2016	2017	2018
<i>Membership Discounts</i>						
Veteran's Day Discount 2014	-5.35	<b><u>22.88</u></b>	0.71	2.23	-0.98	-
	[-13.09, 2.29]	<b><u>[21.8, 24.02]</u></b>	[-0.67, 2.19]	[-0.04, 4.44]	[-1.71, -0.22]	-
Constitution Day Discount 2017	-1.48	-0.26	3.01	4.13	<b><u>-0.14</u></b>	-
	[-2.91, -0.1]	[-2.38, 1.7]	[-1.36, 7.05]	[-0.65, 8.67]	<b><u>[-3.54, 3.09]</u></b>	-
Christmas and New Years Discount 2017	27.5	-2.41	-15.05	-0.15	<b><u>2.36</u></b>	2.2
	[26.48, 28.65]	[-3.73, -1.1]	[-23.3, -5.35]	[-1.16, 0.87]	<b><u>[1.67, 3.05]</u></b>	[1.62, 2.81]
Memorial Day Discount 2018	-3.57	-2.48	-1.13	1.85	-0.38	<b><u>3.62</u></b>
	[-6.29, -0.58]	[-5.14, 0.36]	[-21.01, 23.51]	[0.84, 2.83]	[-1.47, 0.73]	<b><u>[2.96, 4.27]</u></b>
Flash Discount 2018	-4.35	7.62	5.93	2.13	2.69	<b><u>38.47</u></b>
	[-12.19, 3.82]	[5.11, 10.09]	[-1.42, 13.77]	[-0.38, 4.87]	[0.63, 4.72]	<b><u>[33.14, 44.59]</u></b>
<i>Armed Callout Events</i>						
Bundy Ranch Callout 2014	1.8	<b><u>28.21</u></b>	-5.78	-0.19	0.61	-0.53
	[0, 3.62]	<b><u>[26.08, 30.24]</u></b>	[-25.1, 12.68]	[-2.43, 2.6]	[-1.62, 3.62]	[-3.36, 2.05]
Big Sky Callout 2015	-0.06	0.29	<b><u>12.9</u></b>	0.47	-18.71	-
	[-1.24, 1.19]	[-0.81, 1.38]	<b><u>[7.84, 17.72]</u></b>	[-0.23, 1.17]	[-40.02, 6.66]	-
<i>Traditional Advertisement</i>						
NASCAR Sponsorship 2013	<b><u>-9.56</u></b>	-17.89	2.72	3.5	0.1	-3.87
	<b><u>[-11.69, -7.36]</u></b>	[-21.75, -14.48]	[-84.62, 85.5]	[1.86, 5.04]	[-1.75, 1.9]	[-8.46, 0.02]

*Note:*

Brackets are 95% credibility/confidence intervals. Bolded and underlined values are the treatment effects. All estimates use a two week pre-treatment period.



Table 1.7: Two-way fixed effects analysis between counties in the top and bottom quartiles.

	Veteran's Day Discount 2014	Constitution Day Discount 2017	Christmas/New Years Discount 2017	Memorial Day Discount 2018	Flash Discount 2018	Bundy Ranch Callout 2014	Big Sky Callout 2015	Nascar Sponsorship 2013
<i>Panel A: Median household income unweighted</i>								
I(Top quartile) X I(During event)	0.003 (0.005)	-0.0006 (0.001)	-0.001 (0.001)	-0.002 (0.007)	0.0005 (0.0007)	-0.0001 (0.006)	0.004 (0.004)	0.004 (0.004)
Std.Errors	by: county	by: county	by: county	by: county	by: county	by: county	by: county	by: county
FE: Day	X	X	X	X	X	X	X	X
FE: County	X	X	X	X	X	X	X	X
Number of Clusters	1569	1568	1568	1568	1568	1569	1568	1568
Outcome Average	0.007579	0.002805	0.001969	0.005457	0.001126	0.015216	0.010886	0.003975
<i>Panel B: Median household income weighted by county population</i>								
I(Top quartile) X I(During event)	-0.003 (0.003)	-0.0007 (0.0008)	-0.0008** (0.0004)	-0.005 (0.005)	0.0004 (0.0008)	0.0001 (0.002)	0.001 (0.002)	-0.001 (0.001)
Std.Errors	by: county	by: county	by: county	by: county	by: county	by: county	by: county	by: county
FE: Day	X	X	X	X	X	X	X	X
FE: County	X	X	X	X	X	X	X	X
Number of Clusters	1569	1568	1568	1568	1568	1569	1568	1568
Outcome Average	0.007579	0.002805	0.001969	0.005457	0.001126	0.015216	0.010886	0.003975
<i>Panel C: Income inequality unweighted</i>								
I(Top quartile) X I(During event)	-0.007 (0.011)	-0.0007 (0.001)	-0.00002 (0.0009)	-0.007 (0.011)	-0.0008 (0.0006)	-0.0001 (0.006)	0.011** (0.005)	0.002 (0.002)
Std.Errors	by: county	by: county	by: county	by: county	by: county	by: county	by: county	by: county
FE: Day	X	X	X	X	X	X	X	X
FE: County	X	X	X	X	X	X	X	X
Number of Clusters	1564	1564	1564	1564	1564	1564	1564	1564
Outcome Average	0.007588	0.002808	0.001972	0.005464	0.001127	0.015235	0.0109	0.00398
<i>Panel D: Income inequality weighted by population</i>								
I(Top quartile) X I(During event)	-0.009*** (0.002)	-0.001** (0.0006)	-0.0002 (0.0004)	-0.004 (0.004)	0.00006 (0.001)	-0.006*** (0.002)	-0.00007 (0.002)	0.0008 (0.001)
Std.Errors	by: county	by: county	by: county	by: county	by: county	by: county	by: county	by: county
FE: Day	X	X	X	X	X	X	X	X
FE: County	X	X	X	X	X	X	X	X
Number of Clusters	1564	1564	1564	1564	1564	1564	1564	1564
Outcome Average	0.007588	0.002808	0.001972	0.005464	0.001127	0.015235	0.0109	0.00398

Note:

\* p < 0.1, \*\* p < 0.05, \*\*\* p < 0.01

Outcome average is the average number of Oath Keepers per 100,000 of both top and lower quartiles before and after the callout event. I(Top Quartile) is an indicator if a county is in the top quartile. The reference group is the bottom quartile. I(During Event) is an indicator equal to one if the day is during an event and zero if the day is before the event. The middle quartiles and days after the event concludes are dropped from the analysis. Standard errors are clustered at the county level.

# Chapter 2

## Does Deplatforming Work?

### 2.1 Introduction

In July 2022, a jury convicted Alex Jones, founder of the far-right radio/internet conspiracy website Infowars, of defamation. One claim Jones made during the trial was his revenue plunged after being banned from major social media platforms (Hsu 2022). However, Hsu (2022) reports that Jones' revenue, which amounts to tens of millions of dollars annually, increased after being banned from social media platforms. The first revelation suggests there is a strong incentive for online content creators operating as profit-maximizing individuals to produce misinformation in the current, unregulated market. The second revelation suggests that deplatforming, a popular policy in which companies remove creators (and their content) when they violate terms and conditions, could further incentivize this behavior. While illuminating in the case of Alex Jones, the findings from the trial beg two additional questions: i) do far-right content creators, in general, receive additional revenue from alternative platforms after being banned? and ii) if so, is the change in net revenue due to being banned positive or negative?

I examine the net effect of deplatforming across social media sites, known as cross-

platform effects. I study far-right content creators who have accounts on YouTube and the far-right alternative Bitchute before being banned on YouTube. When these content creators are deplatformed from YouTube, they experience an increase in revenue and views on their Bitchute channels. However, the increases on Bitchute are exceeded by the losses on YouTube. In other words, deplatforming succeeded in decreasing total views and donations across YouTube and Bitchute.

The increase in activity on Bitchute and decrease on Youtube shows deplatforming can have two, simultaneous and opposite effects. I rationalize this finding as viewers deciding between two substitution options when a content creator is banned: follow the content creator to a new website, in effect substituting platforms, or stay on the current site and watch new content creators. Which effect will dominate depends on if the viewer is more willing to substitute content or the platform they use. My results show that even when a policy like deplatforming is overall successful in limiting the spread and financial feasibility of content creators, a smaller, opposite effect can still be present.

Identifying individuals' online activity and revenue streams across platforms requires innovation in data collection. I introduce a new method to match social media activity across platforms. I match far-right content creators' banned YouTube channels to their Bitchute channels and use direct donations via bitcoin as a measure of revenue on Bitchute.<sup>1</sup> Using this novel matching strategy, I link 79 Bitchute channels to YouTube accounts and bitcoin wallets. The sample encompasses channels linked to violent extremist movements including white nationalist outlets Red Ice and Vdare, and neo-Nazi Mark Collett. Ideas promoted by these channels have been the inspiration for far-right terror attacks in Pittsburgh, Pennsylvania; El Paso, Texas; and Christchurch, New Zealand.

---

<sup>1</sup>Although content creators can earn revenue through many avenues including direct subscriptions, sponsorships, and merchandise sales Hua et al. (2022), cryptocurrencies became a popular form of payment among far-right content creators after financial companies like Paypal and Patreon began deplatforming them (Keatinge et al. (2019); ISD (2020); Bogle (2021)).

I use a two-way fixed effects regression (TWFE) analysis at the channel week level to estimate the effect of deplatforming on revenue, by exploiting variation in ban dates. A YouTube ban leads to a statistically significant 29% increase in weekly bitcoin revenue. Using the average weekly bitcoin earnings of untreated and pre-ban treated channels (\$73), the point estimate can be interpreted as a \$21 increase in weekly bitcoin revenue. This is below the lower bound of my estimate of how much YouTube ad revenue the channels could have earned if they had been monetized, suggesting that the ban caused a net loss in revenue. A YouTube ban also leads to a statistically significant increase in average views per week on Bitchute. Similarly, the increase in viewership on Bitchute accounts for 5.9% of the estimated loss from YouTube.

The implication of the results extend well beyond the current unregulated market in the United States and can help inform the effects of government regulation. For example, the Network Enforcement Act in Germany, a recent social media regulation, applies only to platforms above two million users (Wagner et al. 2020). Such a policy would regulate YouTube but not Bitchute.<sup>2</sup> Because of this, the results of YouTube’s self-censorship provides useful insights to the effects of proposed and implemented regulations.<sup>3</sup>

From a policy perspective, the overall effects suggest deplatforming does work. While banned content creators receive additional views and revenue on Bitchute after being banned from YouTube, the increase in activity is small compared to lost views and potential revenue from YouTube. There are two main takeaways for policy makers. First, removing content can serve as an effective tool to limit extremists’ online presence. Second, banning content on one platform can affect activity on another platform. Cross-platform effects should be included when evaluating the efficacy of deplatforming policies.

---

<sup>2</sup>Trujillo et al. (2022) finds there were only 61,000 accounts on Bitchute between June 2019 to December 2021, well below the cutoff necessary for regulation.

<sup>3</sup>Most misinformation isn’t targeted at a small, focused constituency who feel strongly enough to sue, as was the case in Alex Jones’ misinformation about the Sandy Hook shooting. This implies the threat of lawsuit does not act as a deterrent for most disinformation by profit-seekers.

To the best of my knowledge, this is the first paper to focus on the financial effects of being deplatformed. Previous literature focuses on sentiment analysis, word usage, views, and activity (Chandrasekharan et al. (2017); Mathew et al. (2020); Jhaver et al. (2021)). I focus instead on the financial effects of being deplatformed. Knowing what affects revenue (and how) is essential to understanding how policies might influence extremist content through incentives. It is possible that being banned on one site could leave a content creator better off financially if a back-fire effect is large enough on alternative platforms. My results show in the case of far right content creators on YouTube and Bitchute, the back-fire effect on Bitchute is small compared to the overall losses on YouTube, meaning deplatforming was successful in limiting the overall amount of views and revenue.

This work contributes to the general literature on extremism and the media. Before social media, interest focused on the effect of traditional media coverage on extremist activities (Tokgoz 2012, Balcells & Torrats-Espinoso 2018, Jetter 2017, 2019). This includes how extremist groups spread their message and recruit through the internet (Forest 2005) and how terrorist organizations use the media to increase their publicity (Frey 2004). A recent branch of literature studies how ISIS propaganda relates to violent outcomes (Cremin & Popescu 2021) and pro-ISIS statements on social media (Mitts 2018). This paper contributes to that strand of literature by focusing on the effectiveness of policies designed to mitigate the influence, measured in profitability, of extremist activity on social media, a business model not yet well understood (Zhuravskaya et al. 2020).

## 2.2 Background

YouTube is a video-sharing website. As of January 2021, it is the second most visited site in the world, behind only `google.com` (*Alexa - Top sites* n.d.), and is considered the

largest social networking site in the United States. Users share, post, and discuss videos on the platform. Videos are either pre-recorded and uploaded or live-streamed and saved directly on the site. Creating an account on YouTube is free and requires a valid email address. After creating an account, individuals create channels where they upload videos. Accounts can have multiple channels for multiple topics (e.g., beauty, business, comedy).

YouTube allows individuals to monetize their channels within the site. These methods include advertising revenue, channel membership, merchandise sales directly on YouTube, and donations while live streaming. Monetization requires a minimum number of subscribers and the production of original content within the YouTube terms and conditions. Beginning in 2019, YouTube increasingly removed ads from political videos to minimize the risk of advertisers being associated with extremist content (Munger & Phillips 2020). This process is referred to as *demonetizing* a channel because the content creator no longer earns revenue from YouTube.

In part because of demonetization, individuals augmented their revenue with third-party sources including in-video advertisements and donations from viewers. Short descriptions are added to videos where the creator can provide a summary of the video, their personal information, and links to where donations are accepted. These include Patreon, Paypal, and SubscribeStar accounts, cryptocurrency wallets, and advertisements for products. Individuals refer the viewer to the links and ask for a donation or provide endorsements for products at some point throughout the video. These revenue streams tend to be advertised across multiple streaming platforms. The approach is popular on YouTube with an estimated 61% of channels employing some form of alternative monetization. Channels considered alt-right and alt-light are more likely to utilize these revenue sources compared to the general YouTube population (Hua et al. 2022).

Social media companies increased content moderation with the first major implementation on YouTube involving the removal of content associated with the Islamic State in

2015 (Mitts 2021). The policies aim to enforce the community guidelines agreed to by all users of the platform.

When YouTube bans an account, no one can log into the account. This means all content on the account is no longer accessible. YouTube has outlined a set of rules with regards to an account being removed. Typically, a first violation is a warning from YouTube. However, YouTube explicitly states they terminate channels of severe abuse without warning (*Community Guidelines strike basics - YouTube Help* n.d.). If the abuse does not warrant immediate termination, YouTube implements a “three-strike” policy. The first two strikes limit functionality on YouTube including monetization. The third strike results in the termination of the account.

As a response, multiple alternatives to YouTube were formed. Bitchute was the first to focus on video content rather than live streaming or messaging (Trujillo et al. 2020). Like YouTube, Bitchute is organized around the idea of channels. An individual creates a channel where they post videos and a brief description of the purpose of the channel. Examples include entertainment, education, and news and politics. Other users can then freely watch the videos, like, comment, share, and upvote/downvote. Like YouTube, owners of the channel can provide a description in each video including links to other social media sites. Owners of the channel can also add donation options including cryptocurrency, PayPal, and subscriptions such as SubscribeStar and Patreon. One reason Bitchute has increased in popularity is the ability to sync a YouTube channel. By syncing the YouTube channel, all previously uploaded YouTube videos are backloaded onto the Bitchute channel and all newly uploaded YouTube videos are automatically loaded onto Bitchute.

The Anti-Defamation League (League n.d. *a*) and Southern Poverty Law Center (Center n.d. *a*) have documented **Bitchute** and its users in multiple reports. As YouTube has taken steps to enforce its terms and conditions through moderation and content removal,

Bitchute experienced increased popularity due to its “free speech” message and limited content moderation (Trujillo et al. 2020). Content on the site focuses on the dissemination of news and political commentary.

## 2.3 Theoretical Effects of Deplatforming

I present a simple model to explain the potential effects of deplatforming on overall revenue for a content creator, drawing from Chan & Gillingham (2015) modeling of rebound effects from environmental economics. The model’s concept is rooted in consumer theory, where viewers have the option to either follow the content creator to a different social media platform or stay on the existing platform to consume alternative content. The viewer’s decision to either follow the content creator or remain on the current platform is dependent on the level of substitutability between their activities on each platform, as well as other possible activities or content creators.

Suppose there is one representative consumer who can watch a specific content creator on two social media sites: Site 1 and Site 2. The representative agent is endowed with  $\$I$  and can spend the money three ways: i) donating  $\$d_1$  to the content creator on site 1, ii) donating  $\$d_2$  to the same content creator on Site 2, iii) spending  $\$x$  on some other activity/content creator.<sup>4</sup> Suppose Site 1 can choose to ban the content creator, forcing  $d_1 = 0$ . Let  $b = 1$  if the content creator is banned from Site 1 and 0 otherwise.

The consumer seeks to maximize their twice continuously differentiable, strictly increasing, and strictly quasiconcave utility function subject to a budget constraint and ban status of the content creator:<sup>5</sup>

<sup>4</sup>The model can also be applied to views by relabeling donations as time spent and endowed income as time.

<sup>5</sup>See the online appendix for similar results derived under a constant elasticity of substitution utility function.



$$\begin{aligned}
\max_{d_1, d_2, x} \quad & U(d_1, d_2, x|b) \quad \text{sbj. to } d_1 + d_2 + x = I \\
& d_1, d_2, x, I \geq 0 \\
& b \in \{0, 1\}
\end{aligned} \tag{2.1}$$

The solution to the maximization problem, as a function of  $b$ , is:

$$\begin{aligned}
(d_1^*(b), d_2^*(b), x^*(b)) &= \arg \max_{d_1, d_2, x} U(d_1, d_2, x|b) \\
&\text{sbj. to } d_1 + d_2 + x = I
\end{aligned} \tag{2.2}$$

To see how a policy change at Site 1 may affect donations across both sites and outside activity, first notice that:

$$d_1^*(0) + d_2^*(0) + x^*(0) = I = 0 + d_2^*(1) + x^*(1) \tag{2.3}$$

$d_1^*(1) = 0$  because the content creator is banned on Site 1, meaning that they can no longer accept donations on the site. Assuming  $d_1^*(0) \neq 0$  and the budget constraint is binding,<sup>6</sup> line 2.3 can be rearranged as:

$$\frac{\Delta d_2^*}{\Delta d_1^*} + \frac{\Delta x^*}{\Delta d_1^*} = -1 \tag{2.4}$$

where  $\Delta y = y(1) - y(0)$  for  $y \in \{d_1^*, d_2^*, x^*, \pi^*\}$ . The content creator's profit is

---

<sup>6</sup>If  $d_1^*(0) = 0$ , then the content creator was not receiving donations on Site 1 prior to be deplatformed. I focus instead on the case where the content creator was earning revenue on Site 1 prior to being banned.

$\pi^*(b) = d_1^*(b) + d_2^*(b)$ , assuming that they have no costs. The effect of a change in policy on Site 1 on the content creator's profit is:

$$\Delta\pi^* = \Delta d_1^* + \Delta d_2^* \quad (2.5)$$

Dividing both sides by  $\Delta d_1^*$  and substituting in line 2.4 yields:

$$\frac{\Delta\pi^*}{\Delta d_1^*} = -\frac{\Delta x^*}{\Delta d_1^*} \quad (2.6)$$

Intuitively, line 2.4 states that  $\Delta d_2^*$  and  $\Delta x^*$  must sum to  $d_1^*(0)$  if the budget constraint is binding. If a content creator is banned, then the portion of income spent on  $d_1^*$  must be reallocated to  $d_2^*$  and  $x^*$ . How it is reallocated depends on the substitutability between donations on Site 2 and donations on Site 1, and donations on Site 1 and outside spending. In addition,  $d_2^*$  or  $x^*$  could decrease with a decrease in  $d_1^*$  if the goods are complements (e.g.  $\frac{\Delta d_2^*}{\Delta d_1^*} > 0$ ) or remain unchanged if goods are independent (e.g.  $\frac{\Delta d_2^*}{\Delta d_1^*} = 0$ ). Line 2.6 states that the content creator's profit will decrease if  $x^*$  and  $d_1^*$  are substitutes, it will increase if they are complements, and remain unchanged if the goods are independent of each other.

Assume prior to a ban, there exists a unique interior solution that maximizes consumer utility. Figure 2.1 plots  $(d_1^*(0), d_2^*(0), x^*(0))$ , the optimal bundle when there is no ban, along with potential optimal bundles when the content creator is banned, referred to as A, B, C, D and E.

[Figure 2.1 about here]

The solid black line is the partial budget constraint when the content creator is not

banned, subtracting out  $x^*$ , spending on the alternative activity, while the other lines are potential budget constraints removing  $x^*$  after the change in policy. The potential new equilibrium are on the y-axis because if the content creator is banned, then the representative consumer can only donate  $d_1^*(1) = 0$  to the content creator on Site 1. The amount spent on donations shifts in or out depending on how  $x^*$  changes after the ban. Below, I describe each potential equilibrium and the sign of the derivatives.

$$\textit{Point A: } \frac{\Delta x^*}{\Delta d_1^*} = 0, \frac{\Delta d_2^*}{\Delta d_1^*} < 0$$

The consumer donates the same total amount of money to the content creator by increasing their donations to Site 2 after the ban. This means the consumer completely substitutes donations on Site 1 to Site 2, supporting theories of displacement (Keatinge et al. 2019). However, there is no change in the overall donation amount. A shift to point A implies consumers maintained their overall pre-ban level of donations. The deplatforming had no effect on the content creator's profit.

$$\textit{Point B: } \frac{\Delta x^*}{\Delta d_1^*} < 0, \frac{\Delta d_2^*}{\Delta d_1^*} = 0$$

If there are no *cross-platform* effects of deplatforming, then the content creator will see no change in revenue on Site 2. The viewer substitutes all of the money once spent on donations on Site 1 to outside activities after the content creator is banned on Site 1. Observing this equilibrium suggests policy changes on one social media site do not effect another site. In other words, donations on Site 2 are not affected by a policy change on Site 1 and overall profits for the content creator decrease.

$$\textit{Point C: } \frac{\Delta x^*}{\Delta d_1^*} < 0, \frac{\Delta d_2^*}{\Delta d_1^*} < 0$$

There could be a partial substitution to Site 2. If the representative consumer re-allocates  $d_1^*(0)$  between other activities/content creators and donations on Site 2, then the content creator will see both increased revenue on Site 2 and a decrease in overall

profits. From a policy perspective, this suggests deplatforming partially worked: while the content creator experiences an *overall* decrease in support, as measured in revenue, there is a partial recovery of support in increased donations on Site 2. The magnitude of the deplatforming effect depends on the substitutability between donations on Site 1 and other activities/content creators.

$$\textit{Point D: } \frac{\Delta x^*}{\Delta d_1^*} < 0, \frac{\Delta d_2^*}{\Delta d_1^*} > 0$$

Observing the content creator being banned could potentially lead the consumer to discontinue support for the content creator. In this case, the representative consumer substitutes all of their donations from Site 1 to other activities/content creators and decreases their donations on Site 2. Graphically, the optimal bundle will move to Point D. This suggests deplatforming policies produce the intended effects: overall support for the content creator goes down across both platforms when they are banned. In the most extreme cases, Point D is at the origin, meaning that the representative consumer stops donating completely and the content creator loses all revenue across the two platforms after being banned.

$$\textit{Point E: } \frac{\Delta x^*}{\Delta d_1^*} > 0, \frac{\Delta d_2^*}{\Delta d_1^*} < 0$$

Finally, the consumer may increase the amount of donations after the ban, meaning the content creator's profit *increases* from being banned. This is represented in Figure 2.1 as a shift from the pre-banned equilibrium to point E. Moving to point E means consumers not only substitute all their donations from Site 1 to Site 2, but also increase their *overall* donations to the content creator. Such perverse effects have been well documented in the economic literature and are referred to as Peltzman Effects (Peltzman 1975). Examples of similar effects include social media bans causing more hateful activity on alternative sites (Mitts 2021); (Ali et al. 2021) and increases in anti-vaccine content (Mitts et al. 2022). Potential reasons for such counter-intuitive effects include curiosity about the

banned content. Jansen & Martin (2015) refer to this phenomenon as the “Streisand Effect” and document it in other contexts.<sup>7</sup>

Deplatforming may work, backfire, or have partial success. In the following sections, I use reduced form estimation strategies to measure the overall change in donations on YouTube ( $\Delta d_1^*$ ) and change in revenue on Bitchute ( $\Delta d_2^*$ ) after content creators are banned from YouTube, which implies the signs of  $\frac{\Delta d_2^*}{\Delta d_1^*}$  and  $\frac{\Delta x^*}{\Delta d_1^*}$ . I find  $\Delta d_1^* < 0$  and  $\Delta d_2^* > 0$ , which implies  $\frac{\Delta d_2^*}{\Delta d_1^*} < 0$ ,  $\frac{\Delta x^*}{\Delta d_1^*} < 0$  and  $\frac{\Delta \pi^*}{\Delta d_1^*} > 0$ . Since  $\Delta d_1^* < 0$ ,  $\Delta \pi^* < 0$  meaning deplatforming decreases overall profits. These results suggest consumers are more willing to substitute content creators and continue using YouTube rather than follow the banned actor to Bitchute.

## 2.4 Data

### 2.4.1 Data Collection

Social media companies, including YouTube, do not provide publicly available lists of banned channels. Past efforts to study bans have systematically scraped social media sites over the course of months identifying bans in real-time (Mitts (2021); Rauchfleisch & Kaiser (2021)). This paper focuses instead on channels reported as banned on `altcensored.com`. The site’s goal is to provide a platform for content that is “neither illegal nor violates YouTube’s terms and conditions or community guidelines” but was still banned or at risk of being banned from YouTube. `altcensored.com` receives suggestions for which channels to archive. A channel is uploaded to `altcensored.com` if the channel meets one of the following requirements (*Altensored.com FAQ* n.d.):

1. Videos are placed in a limited state or arbitrarily age-restricted by YouTube.

<sup>7</sup>Hagenbach & Koessler (2017) rationalize the Streisand Effect using a signal model.

2. Videos are removed or demonetized.
3. The channel is demonetized.

This focuses the population on channels that have incurred some moderation. By focusing on this group, the results of this paper capture the effect of a ban conditional on prior moderation.

As of September 2021, `altcensored.com` recorded 9,370 YouTube channels of which 2,287 are banned. While all channels have a name, created date, and banned date, only some channels include videos. 4,351 of the 9,370 channels have at least one video uploaded.

Like Bitchute, `altcensored.com` syncs YouTube channels to their site. This means that the unique identifier in the video URL is the same between YouTube, `altcensored.com` and Bitchute allowing for linkage between YouTube and Bitchute accounts. For example, if the URL for a YouTube video is `https://www.youtube.com/watch?v=1234`, where “1234” is the unique identifier of the video, then the URL for the same video on `altcensored.com` is `https://altcensored.com/watch?v=1234`. If the YouTube account was synced to a Bitchute account, then the exact same video on YouTube will be on Bitchute with the unique URL `https://www.bitchute.com/video/1234`. Because the syncing uses the same unique part of the URL, YouTube and Bitchute accounts can be matched. `altcensored.com` serves as a repository of banned YouTube channels that allows retrospective analysis despite all the banned channels being unavailable on YouTube. Finally, syncing a YouTube and Bitchute account requires having the username and password for both accounts. This is the first paper to utilize the synced URLs for tracking activity across social media platforms. This method of linking YouTube and Bitchute accounts can be applied to other Bitchute datasets such as Horne et al. (2022).

There are 9,370 channels and 184,849 YouTube videos on `altcensored.com`. 2,287 of

the channels have a reported ban date. 18,792 videos are synced to a video on Bitchute. The 18,363 videos are linked to 280 unique Bitchute channels. I then scrape each individual video description from the 280 Bitchute channels identified for a bitcoin wallet and accompanying transactions using “WalletExplorer: Smart Bitcoin Block Explorer” (*WalletExplorer: Smart Bitcoin Block Explorer* n.d.). All bitcoin values are converted to USD using the publicly available opening bitcoin-USD conversion on the day of transfer. In total, 79 of the 280 channels advertise a bitcoin wallet. Forty-six of the Bitchute channels are synced to at least one banned YouTube channel while the other 33 Bitchute channels are synced to YouTube channels that have not been reported banned. The data is analyzed at the channel-week level. A figure of the data collection process is provided in the online appendix.

Defining a far-right content creator can be a subjective process. For the purposes of this paper, a far-right content creator is defined as an individual who has a social media presence on known far-right social media platforms. This is akin to defining members of an organization as those who show up to organizational events. This allows individuals to self-select into the category removing potential biases in classifying content. *Altcensored.com* and *Bitchute* are used to classify content creators because both were reported as far-right sites by the SPLC.<sup>8</sup>

This approach does impose limitations. The population of interest excludes content creators who were not registered on both sites and YouTube. Some popular far-right extremists host content on their own websites exclusively. The results will not apply to those individuals.

The results also will not apply to individuals who did not post their content on *Bitchute* or *altcensored* prior to being banned. In other words, an individual must have synced their YouTube account to *Bitchute* and *altcensored* prior to a YouTube ban.

---

<sup>8</sup>Cryptocurrency Report | Southern Poverty Law Center ([splcenter.org](http://splcenter.org))

Syncing material and having multiple outlets to post content is a popular strategy among content creators. Those who do not have multiple accounts are expected to be a small, less active subpopulation. Nevertheless, these results will not apply to such individuals.

Each channel is observed beginning with the first reported bitcoin wallet and ending the week of December 19, 2021. Because Bitchute allowed users to sync their YouTube channels and back-download all old videos, some start dates for channels are before Bitchute was created. Figure 2.2 plots the number of observed and banned channels by date.

[Figure 2.2 about here]

### 2.4.2 Description of the Sample

The sample includes 4,035 donations totaling \$1,788,555 spanning from the week of January 1, 2017 until the week of December 19, 2021 with scraping completed on December 25, 2021. This consists of 79 Bitchute channels. Forty-six of the Bitchute channels are synced to at least one banned YouTube channel (treated) while the other 33 Bitchute channels are synced to YouTube channels that have not been reported banned (never-treated). Topics of the channels include vaccine hesitancy, the Qanon conspiracy, and support for Donald Trump. This is common for most Bitchute channels (Trujillo et al. 2020).

Channels in the sample have been identified by other sources as potential hate speech or “alt-right”. For example, the Anti-Defamation League identified *Red Ice TV* and *The Red Elephants with Vincent James* as anti-Semitic channels (League n.d.b) and Stefan Molyneux has been reported by the Southern Poverty Law Center (Center n.d.b). These individuals are included in this sample. Several individuals are included in Squire (Squire 2021) analysis of right-wing streamers’ revenue on Dlive and the SPLC list of bitcoin



wallets affiliated with far-right organizations and individuals. Trujillo et al. (Trujillo et al. 2020) highlights *highimpactflix* (another channel in the sample) as promoting hyper-partisan political commentary and being in the top ten Bitchute channels for views and comments.

Bitcoin is highly referenced in the sample. Figure 2.3 plots the ten most linked domain names in the video descriptions and bitcoin. The most linked-to website is twitter.com followed by internal links within bitchute.com. Bitcoin is referenced more often in the sample than the top referenced websites.

[Figure 2.3 about here]

Table 2.1 compares average features of Bitchute channels linked to never-banned YouTube channels and Bitchute channels linked to banned YouTube channels prior to the ban. On average, Bitchute channels earned \$46 in weekly bitcoin revenue prior to the YouTube ban compared to \$74 for channels linked to never-banned YouTube channels. The distribution for both groups is dispersed with 13% of the channel-weeks earning more than \$0 in bitcoin revenue. On average, both groups received less than 1 donation per week. Focusing only on weeks where a channel received donations, Bitchute channels linked to banned YouTube channels earned \$397 while Bitchute channels linked to never-banned YouTube channels earned \$568.

Topics are similar between Bitchute channels linked to banned and never-banned YouTube channels. A visualization of the 50 most popular words used in Bitchute titles for channels associated with banned and never-banned YouTube channels is presented in the online appendix. Twenty-four of the words overlap and the ranking of the 24 overlap words for the banned and never-banned groups. As found in Trujillo et al. (2020), the content tends to revolve around political events and political figures. Finally, a chart is provided in the online appendix of the top 10 video categories for Bitchute

channels affiliated with banned and never-banned YouTube channels. The categories for a video are chosen by the creator of the video. Similar to the video title content, there is a strong overlap in the video categories. Together, this evidence supports the claim Bitchute videos affiliated with banned and never-banned channels had relatively similar content.

[Table 2.1 about here]

Figure 2.4 plots the average weekly bitcoin revenue in USD. On average, channels experienced a large spike in revenue the week of the ban compared to prior and again a large spike four weeks post ban. Graphically, this suggests a potential positive increase in bitcoin donation revenue as a reaction to a YouTube ban. On average, channels experience a large increase in viewership on Bitchute the weeks of the ban which appears persistent.

[Figure 2.4 about here]

## 2.5 Model and Estimation

I first estimate the change in bitcoin revenue using a two-way fixed effects (TWFE) framework:

$$Y_{i,t} = \alpha + \beta_1 D_{i,t} + X_{i,t} \beta + \gamma_i + \gamma_t + \epsilon_{i,t} \quad (2.7)$$

where  $Y_{i,t}$  is logged weekly bitcoin revenue in USD,<sup>9</sup>  $D_{i,t}$  is an indicator equal to 1 if the channel was banned from YouTube in week  $t$  and 0 otherwise,  $\gamma_i$  are channel fixed

---

<sup>9</sup>Estimation is performed at the channel week level. All analysis is redone using the inverse hyperbolic sine function. Please see the online appendix.

effects and  $\gamma_t$  are week fixed effects.  $X_{i,t}$  includes the following time varying covariates: number of videos per week, indicator if Patreon was referenced at least once in a video description by streamer  $i$  in week  $t$ , indicator if Paypal was referenced at least once in a video description by streamer  $i$  in week  $t$ , indicator if SubscribeStar was referenced at least once in a video description by streamer  $i$  in week  $t$ , and an indicator if Venmo was referenced at least once in a video description by streamer  $i$  in week  $t$ . In the appendix, I show there is no evidence the controls are affected by a YouTube ban.

Similar to YouTube, payment companies do not produce a public list of banned accounts and dates. As a proxy for additional revenue streams (and potential bans), I include references to popular payment platforms from the video description. If a streamer no longer advertises a specific donation option (like a paypal account), it is likely they no longer accept donations from that source.

I assume once a YouTube channel is banned, it remains banned.<sup>10</sup> Errors are clustered at the channel level.

A causal interpretation requires the *parallel trends* assumption to hold. One way to explore the plausibility of this assumption is through an event study. Due to the small sample size, I estimate the event study grouping observations at the monthly level. I calculate the event study using the following equation:

$$Y_{i,t} = \alpha + \sum_{k=-12}^{12} \beta_k D_{i,t}^k + X_{i,t}\beta + \gamma_i + \gamma_t + \epsilon_{i,t} \quad (2.8)$$

where  $D_{i,t}^k$  is an indicator if channel  $i$  at week  $t$  was banned  $k$  months ago.<sup>11</sup> All leads(lags) beyond 12 are binned at -13(13) and omitted from the event study graph

<sup>10</sup>This assumption is investigated further in Section 2.6.1.

<sup>11</sup>For example, if a channel was banned in February, set  $k = -1$  for all weeks in January for channel  $i$ .

(Schmidheiny & Siegloch (2019)). Standard errors are clustered at the channel level.  $\beta_{-1}$  is standardized to 0. Figure 2.5 presents the event study.

[Figure 2.5 about here]

The pre-period coefficients do not exhibit signs of pre-trends. A joint significance test on the pre-periods fails to reject all coefficients are equal to zero at conventional levels. This supports the credibility of parallel trends.

Recent advances in the difference-in-differences literature demonstrate potential biases in using TWFE with staggered treatment adoption (Callaway & Sant’Anna (2020); de Chaisemartin & D’Haultfœuille (2020); Goodman-Bacon (2021)). The TWFE estimate can be decomposed into a weighted sum of average treatment effects in each channel by week. Concerns of bias arise when treatment effects are heterogeneous. In extreme cases, this can lead the TWFE estimate to be of opposite sign of the average treatment effect, commonly referred to as the “negative weights problem”.

de Chaisemartin & D’Haultfœuille (2020) provides a metric to understand how serious the issue is in application. In this setting, 9.7% of the average treatment effects have negative weights where the negative weights sum to -0.01. I also perform the analysis using the proposed de Chaisemartin & D’Haultfœuille (2020) estimator (DIDM), which is robust to heterogeneous treatment effects. Due to efficiency concerns with DIDM, I focus on comparing the point estimates from DIDM and TWFE. Unless noted, the DIDM point estimates fall within the 95% confidence interval of the TWFE estimate.<sup>12</sup>

Cluster heterogeneity can cause size distortions leading to under-rejection of the null hypothesis (Carter et al. (2017)). The differences in observation length between channels and frequency of donations suggest this setting is susceptible to such issues. While there are 79 channels in the sample producing 79 clusters, there are 69 effective clusters due

---

<sup>12</sup>See the online appendix for more details.

to cluster heterogeneity. This is well above the minimum threshold suggested by Carter et al. (2017). As a robustness check, I report the wild cluster bootstrap 95% confidence interval with 10,000 replications.

## 2.6 Change in Bitcoin Revenue

Table 2.2 presents the results using the standard TWFE at the channel week-level with fixed effects for channel and week. standard errors are clustered at the channel level. The outcome is bitcoin revenue in logged USD. The average effect of a YouTube ban on weekly logged bitcoin revenue is 0.26. In percent terms, this equates to an increase of 29%. In terms of USD, the average increase in weekly bitcoin revenue is \$21. These results show content creators experienced an *increase* in bitcoin revenue after being banned from YouTube.

[Table 2.2 about here]

I next estimate the overall change in profits between bitcoin revenue and YouTube revenue comparing the above effect to potentially foregone YouTube ad revenue. I use `socialblade.com` to estimate forgone revenue on YouTube. The site tracks daily YouTube channel activity and estimates the revenue from ads based on the number of views (referred to as clicks) and general YouTube payment rules of thumb.<sup>13</sup>

I collect the average weekly video views for the YouTube channels in sample as of January 24, 2022. `Socialblade.com` found most channels earn between \$0.25 USD to \$4.00 USD per 1,000 clicks. Anecdotal evidence suggests the average rate is \$2.00 per 1,000 clicks. These rough estimates paired with the average number of views are

---

<sup>13</sup>Tracking YouTube earnings is very difficult due to many factors. How much revenue a channel earns depends on where the viewer is watching whether the content is being consumed on a home desktop or mobile device, if the viewer skipped through the ads, if the viewer uses an adblocker, season variations (e.g. holiday spikes) and the content of the channel (e.g. business, gaming, news, etc.).

used to calculate bounds on YouTube ad revenue, assuming all of the YouTube channels are monetized. To be archived on `altensored.com`, two criteria require some form of demonetization. This means it is unlikely most of the channels were receiving any income through YouTube monetization. The YouTube revenue estimates should be thought of as the opportunity cost of producing content that results in deplatforming. It should not be thought of as an estimate of foregone revenue.

If all the channels were monetized and earned \$0.25 per 1,000 clicks, then the average weekly earnings from YouTube ad revenue is \$32. If instead, they earned \$4.00 per 1,000 clicks, then the average ad revenue is \$507. Both estimates are above the estimated weekly increase in bitcoin revenue due to a YouTube ban (\$21).

This shows that overall profits decreased while revenue on Bitchute increased. There are two opposite effects occurring simultaneously: a decrease in potential revenue on YouTube but an increase on Bitchute. The overall decrease on YouTube is larger than the increase on Bitchute leading to an overall decrease in profits, on average, for the content creators. The partial recovery of lost revenue on YouTube suggests viewers, on average, are more willing to substitute content and remain on YouTube, rather than follow banned actors to Bitchute.

## 2.6.1 Robustness Checks

### Potential Sample Issues

This subsection identifies and address five potential issues to the sample: i) channels run by the same individual, ii) inactive channels, iii) unbanned channels, iv) duplicate channel and v) observation window. A table of results is provided in the online appendix

**Channels Run by the Same Individual** Potential concerns with the sample are multiple channels owned by the same individual. Bitchute has a *profile* category in addition to channels. There are 77 unique profiles because two profiles manage multiple channels. I rerun the analysis clustering by profile instead of channel. The results are near-identical to the main specification.

**Inactive Channels** There are 11 channels that were banned from YouTube after their last Bitchute video post. In the most extreme case, one channel was banned 51 weeks after the last posted video on Bitchute. The average time between the last post and ban for these 11 channels is 13.9 weeks with a median difference of 9 weeks. I remove the 11 channels and rerun the main specification. This focuses the subsample on content creators still active on Bitchute rather than those who ever had a Bitchute account. The TWFE is about the same magnitude as the main specification. Standard errors increase due to a smaller sample size.

**Duplicate Channel** Some Bitchute channels advertised the same bitcoin wallet consistently. While they were linked to different YouTube channels, this suggests double counting revenue. Channels that consistently advertised the same bitcoin wallet were combined into one unit observation by averaging over the channels each week. Removing the channel does not substantively change the magnitude of the results. However, the results are now significant only at the 10% level.

**Unbanned Channels** The length of a YouTube ban is not available. One concern is a channel is banned for a specified amount of time (e.g., 30 days) and reinstated. To investigate this, I check if any video on Bitchute uploaded post-ban is synced to a YouTube video. One banned channel has its most recent video synced to YouTube meaning its channel was unbanned at some point. I rerun the analysis dropping the

channel. With the channel removed, the effect of a YouTube ban on revenue remains the same as the main specification.

### **Additional Investigation into the Parallel Trends Assumption**

The choice of 12 months of leads and lags was made because it was long enough to observe a pre-trend while considering sample size limitations. To further investigate the validity of the parallel trends, I rerun the event study increasing the number of leads and lags to 24 months. The event study is plotted with 24 leads and 24 lags in the online appendix. It casts doubt on the validity of parallel trends over a year prior to a YouTube ban. I rerun the analysis limiting the treated units to one year before and after the YouTube ban. The TWFE estimates are similar in magnitude and significance. The results are presented in the online appendix.

### **Additional Robustness Tests**

Equation 2 is re-estimated leaving out one treated unit. This is to ensure the results are not driven by one influential channel. The results are insensitive to changes with the point estimates ranging from 0.22 to 0.31. The difference between the smallest and largest leave-one-out estimate is 9.5%, or \$6.6. The leave-one-out estimates compared to the main specification are provided in the online appendix.

Finally, I re-run the analysis omitting the never-treated channels. Omitting the never-treated channels leads to a substantially larger share of negative weights (25%). The point estimate is like the DIDM estimate. The loss of sample size leads to an increase in the standard errors. However, the estimates still overlap with the main specification. The TWFE estimates are presented in the online appendix.



## 2.7 Change in Viewership

A content creator may be motivated by the number of views, referred to as reach, of content rather than financial gains. Similarly, a social planner's objective may be to limit the audience of the content creator. In such cases, a change in views would be the primary outcome of interest.

I estimate the effect of a YouTube ban on *logged average views per week* using Equation 2. This is calculated using the number of views for videos produced in each week. I implicitly assume all views for the channel occur through newly created videos. While an imperfect assumption, past research suggests most views for videos on Bitchute occur within the first week of being posted (Trujillo et al. (2020)). Figure 2.6 provides the events study. The event study suggests parallel trends are plausible in this setting.

[Figure 2.6 about here]

Table 2.3 presents the TWFE estimates using Equation 1 specifications. Content creators banned by YouTube experience an average increase in the average number of views per week of 67%. Using the average views per week of unbanned channels and pre-banned channel weeks (12,146), a channel experiences an average increase in viewership of 7,483 per week. Rauchfleisch & Kaiser (2021) concludes a similar magnitude but the effect is short-lived while Figure 2.6 suggests a persistent change in viewership on Bitchute.

[Table 2.3 about here]

I compare the increase in Bitchute viewership due to a YouTube ban to lost viewership on YouTube. Similar to the bitcoin analysis, I use `socialblade.com` data to estimate weekly viewership for banned channels. Figure 2.7 graphically shows the effects

of a YouTube ban on estimated weekly revenue and views for YouTube and Bitchute. YouTube channels affiliated with Bitchute channels in the sample receive 125,991 weekly views on average. This estimate implies streamers recovered about 5.9% of lost YouTube viewership through increased Bitchute viewership.

[Figure 2.7 about here]

As with revenue, I find an overall decrease in viewership across the two platforms, but an increase in views on Bitchute.

## 2.8 Mechanisms

What's driving this change in bitcoin revenue? Three potential mechanisms are investigated: i) increased size of donations, ii) increased number of donations, and iii) more content being produced.

If there are heterogenous viewers, then this equilibrium could be reached by those who already donate giving more (intensive margin), viewers start donating because of the ban (extensive margin), or both. An increase in the average value of bitcoin donations suggests those who donate are increasing their intensity of support. The ban itself serves as a galvanizing tool to increase the support of the followers.<sup>14</sup> However, a lack of evidence of a change in bitcoin donation size paired with an increase in the number of donations suggests the ban actually drew more attention to the banned material (as discussed by Rogers (2020) and Keatinge et al. (2019)).

A third mechanism for an observed increase in bitcoin revenue is an increase in the amount of content being produced. If content creators change the amount of content being created on Bitchute, the increase in bitcoin revenue may be a mechanical result - more content leads to more money.

---

<sup>14</sup>This is akin to arguments of political alienation leading to terrorist activity (Krueger (2008)).

To investigate the viability of the potential mechanisms, I rerun the main analysis using five outcomes: i) *average bitcoin donation size* (logged USD +1), ii) *number of donations* (logged +1), iii) *the number of videos* (logged +1), and iv) *whether a channel week received at least one donation* (0 or 1). Outcomes i), ii), and iv) provide insight into whether there is an extensive or intensive effect driving the increase in revenue. Outcome iii) is used to better understand if a change in the supply of videos is driving the increase. Like the main analysis, there is little concern for negative weights biasing the TWFE estimate. A table and graph of estimates and additional discussion is provided in the online appendix.

The results suggest that the increase in bitcoin revenue is from an increase in the probability and number of donations. However, average donations and changes in the amount of content cannot be ruled out as potential mechanisms due to imprecise estimates and a lack of evidence of parallel trends holding. These findings suggest the increase revenue through bitcoin was caused by *more* viewers donating.

## 2.9 Conclusion

How to manage content creators on social media is a problem governments and firms face. The current approach to mitigating such activity is an unregulated market where firms can choose when and how to manage. A popular enforcement technique is deplatforming. The intuition of the policy is straightforward: the threat of removal should act as a deterrent causing content creators to alter behavior or risk being deplatformed. The results of this paper show that while content creators may recover some lost revenue and views through increased activity on alternative sites, deplatforming leads to an overall decrease in revenue and views. In this setting, a YouTube ban led to a 29% increase in weekly bitcoin donations on Bitchute. The change is most likely driven by increases in

the probability of receiving a donation and the number of donations. Both are suggestive of more individual viewers supporting deplatformed content creators on Bitchute. However, this behavior appears to represent a small fraction of viewers on YouTube. Content creators in my sample experienced a decrease in both revenue and viewership across Bitchute and YouTube after being banned from YouTube, suggesting most viewers prefer to stay on YouTube and substitute to other content creators, rather than follow the banned content creator to Bitchute. Based on these results, deplatforming was, overall, successful.

While I find deplatforming has the expected overall effects, the presence of two, opposite effects suggests it can have perverse unintended consequences. The overall net negative effect may be driven by Bitchute's small size compared to YouTube. A smaller site may struggle to attract viewers to switch from a larger platform due to limited content or lack of advertising. Another mechanism that could affect the substitution effects is curiosity about the deplatformed content. Jansen & Martin (2015) provide several examples of the "Streisand Effect" outside of deplatforming, where censorship led to more media coverage, in turn leading to more interest in the content. If the political outrage is large enough, a viewer may be willing to follow a content creator to a different platform, meaning the policy backfired. Finally, the type of content may lead to viewers having different levels of loyalty to content creators. Viewers may be more willing to follow content creators discussing different topics than far-right ideology to alternative sites. If and when these mechanisms can change the effect of deplatforming is beyond the scope of this paper.

Extremist organizations have effectively been utilizing the internet to spread propaganda, recruit, and coordinate attacks throughout the 21<sup>st</sup> century (Frey (2004); Cremin & Popescu (2021)). Deplatforming is a simple policy governments and companies can implement that is effective at limiting the overall number of views and revenue of banned

content creators across sites.

## 2.10 Tables and Graphs




Figure 2.1: Graphical representation of deplatforming effects. Bundles A, B, C, D, and E are potential optimal bundles after the content creator is banned.

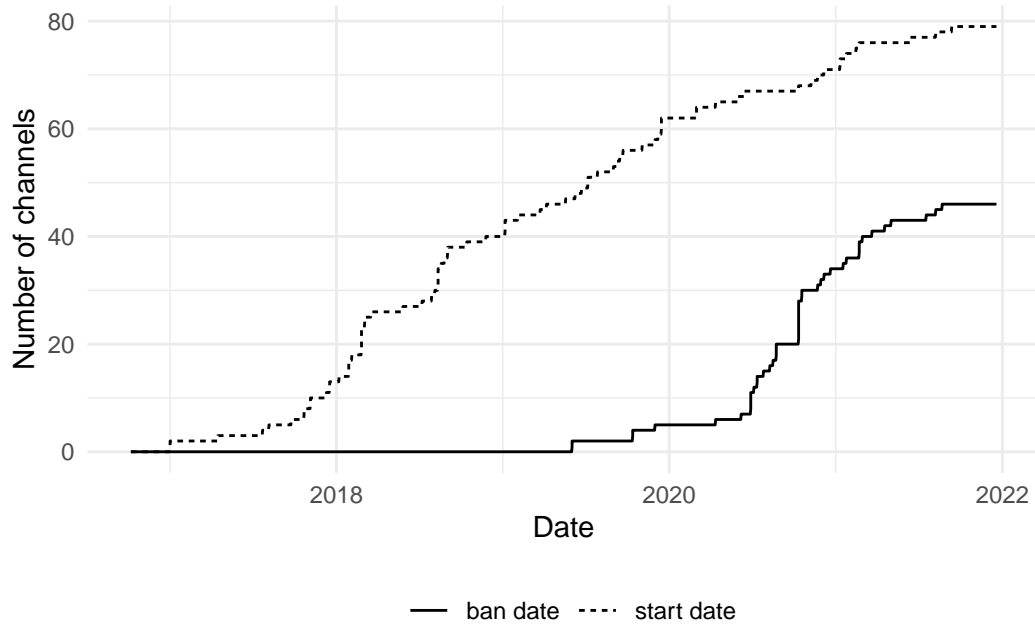


Figure 2.2: Number of Bitchute channels observed at each week by start and ban dates for the sample. Sample includes 79 Bitchute channels. 46 of the channels are linked to at least one banned YouTube account (treated) while 33 are linked to never-banned YouTube accounts (never-treated).

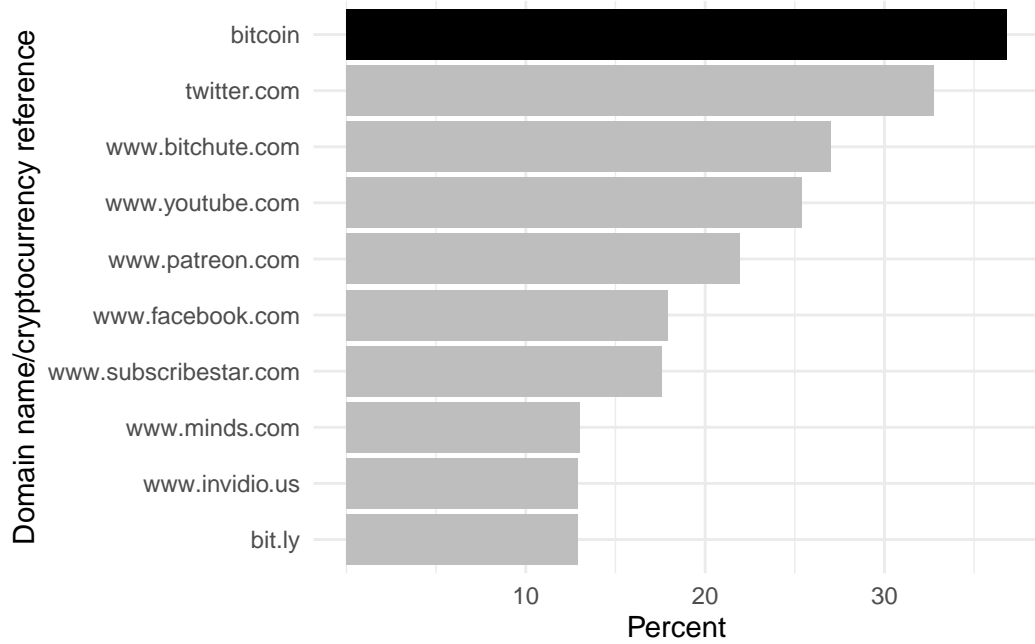


Figure 2.3: Percent of videos in sample of 79 Bitchute channels that reference domain name/cryptocurrency at least once.




Figure 2.4: Outcome variables for treated units with respect to treatment. The x-axis is the average outcome per channel-week.




Figure 2.5: Event study of effect of YouTube ban on bitcoin revenue. Study includes 12 leads and lags with all other leads and lags binned at +/- 12 and omitted from the graph. One period before the ban is omitted. Right hand axis is the logged estimates multiplied by the pre-ban average bitcoin revenue.

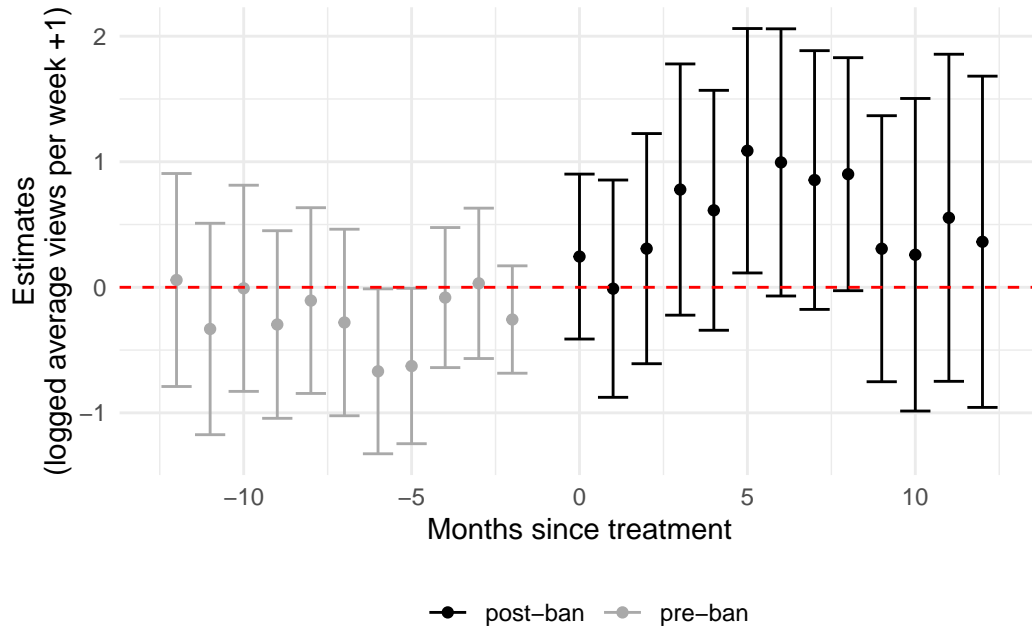


Figure 2.6: Event study of effect of YouTube ban on weekly average video views on Bitchute. Study includes 12 leads and lags with all other leads and lags binned at +/- 13 and omitted from the graph. One month before the ban is omitted. The right hand axis is the logged estimates multiplied by the pre-ban average views per video on Bitchute.




Figure 2.7: Average weekly effect of YouTube ban on revenue and views for YouTube and Bitchute.



Table 2.1: Comparison of banned and never-banned Bitchute channels prior to a YouTube ban.

	Banned (N=4146)		Never-Banned (N=4744)		Diff. in Means	Std. Error
	Mean	Std. Dev.	Mean	Std. Dev.		
Weekly bitcoin revenue (USD)	45.836	600.015	96.708	2336.331	50.872	35.177
Weekly number bitcoin donations	0.270	1.087	0.311	1.455	0.042	0.027
Weekly number of views	7868.940	20114.237	15884.298	55871.358	8015.358	869.250
Weekly number of likes	244.684	697.023	1254.662	5082.913	1009.979	74.587
Weekly number of dislikes	7.074	22.216	20.992	68.080	13.918	1.047
Average start date	2019-01-17 17:13:02		2019-03-25 16:43:38			
Average length of channel (weeks)	137		128		16	
Number of Channels	46		33		13	

*Note:*

The table compares channel weeks of banned channels prior to the ban to channel weeks of never-banned channels throughout the sample.

N refers to the number of channel week observations.

Table 2.2: Effect of YouTube ban on weekly bitcoin revenue on Bitchute (logged USD).

	1	2
YouTube Ban	0.250** (0.119)	0.257** (0.118)
Number Videos		0.001 (0.002)
Std. Errors	by: channel_name	by: channel_name
FE: Channel	X	X
FE: Date	X	X
I(Referenced Patreon)		X
I(Referenced Paypal)		X
I(Referenced Subscribestar)		X
I(Referenced Venmo)		X
Number clusters	79	79
Effective number clusters	69	69
Wild bootstrap 95% confidence interval	[0.01,0.49]	[0.02,0.5]
Average bitcoin revenue pre-ban/never-banned (USD)	72.98	72.98

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:*

Analysis is conducted at the channel week level using a TWFE strategy. There are 79 channels - 46 are treated and 33 are never treated. Date refers to the week. Average bitcoin revenue pre-ban/never banned takes the average channel weekly bitcoin revenue in USD for all weeks before a ban. All never-banned channel weeks are included. Wild bootstrap confidence interval between 2.5 and 97.5 percentiles with 10,000 replications clustered at the channel level.

Table 2.3: Effect of a YouTube ban on average views per video per week on Bitchute (Average Views per Week (logged + 1)).

	1	2	3
YouTube Ban	0.379	0.523***	0.523***
	(0.613)	(0.132)	(0.132)
Number Videos			0.342
			(74320.180)
Std. Errors	by: channel_name	by: channel_name	by: channel_name
FE: Channel	X	X	X
FE: Date	X	X	X
I(Referenced Patreon)		X	X
I(Referenced Paypal)		X	X
I(Referenced Subscribestar)		X	X
I(Referenced Venmo)		X	X
Number clusters	79	79	79

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:*

Analysis is conducted at the channel week level using a TWFE strategy. There are 79 channels - 46 are treated and 33 are never treated. Date refers to the week.

# Chapter 3

## Synthetic Control with Time Varying Parameters

### 3.1 Introduction

In this paper, I consider the problem of estimating the causal effect of an intervention on an outcome of interest when there is one treated unit. A common approach to this problem is the synthetic control framework. The goal is to construct a counterfactual for the treated unit as a linear combination of untreated units. The synthetic control method has been used in many areas of economics including (but not limited to) the effects of terrorism Abadie & Gardeazabal (2003), trade policies Billmeier & Nannicini (2013), and social issues Powell (2021).

Abadie et al. (2010) show that if there exists some linear combination of untreated units such that a perfect pretreatment estimate of the treated unit exists, then the asymptotic bias of the estimated treatment is zero given a linear factors data generating process. Their approach simplifies the time series estimation problem to a cross sectional one. Each time period in the pretreatment is an “observation” used to estimate miss-

ing “observations” in the post treatment. This approach limits the scope of the tool to constant relationships between treatment and control units. If heterogeneous relationships exist and are ignored, the test of no intervention effect is extremely oversized (Carvalho et al. (2016); Masini & Medeiros (2020)). An example of such heterogeneity is nonstationary data including integrated processes. Abadie et al. (2010), Brodersen et al. (2015), and Xu (2017) acknowledge this limitation and explicitly warn against the uses of synthetic control when an accurate counterfactual cannot be constructed.

This paper proposes incorporating time varying coefficients for the control units into the synthetic control framework to address such situations. An immediate concern with time varying coefficients is model misspecification and the risk of overfitting. Recent advances in macroeconometric forecasting developed methods to address such concerns (Dangl & Halling 2012, Bitto & Frhwirth-Schnatter 2019, Belmonte et al. 2014). These methods rely on two ideas: non-centered state space modeling and Bayesian shrinkage. Non-centered state space modeling decomposes time varying parameters into a time varying component and a time invariant component. This allows the researcher to observe which relationships between treatment and control units are time varying. Bayesian shrinkage techniques can then be applied to the decomposed relationships “shrinking” irrelevant parameters towards zero allowing the model to automatically choose if a relationship should be time varying, time invariant, or null. I incorporate these two techniques in a synthetic control framework allowing the proposed model to perform as well as a static-coefficient model when the true data generating process involves only static coefficients and superior otherwise.

The purpose of this paper is to introduce time varying parameters to the synthetic control setting and investigate their usefulness. I contribute to the literature in two ways: incorporating advances in macroeconometric forecasting to a popular counterfactual estimation process and analyzing the performance on small sample sizes and short

pre-treatment periods. Testing of this class of model has focused on high frequency data including stock prices and inflation rates (Dangl & Halling 2012, Belmonte et al. 2014, Bitto & Frhwirth-Schnatter 2019). This differs from a synthetic control setting where data tends to be yearly or monthly with few pre and post periods and potentially many controls. In addition, these scenarios are representative of synthetic control settings applied researchers tend to encounter.

This paper is part of the general literature on estimation strategies for the synthetic control counterfactual. Recent literature proposed recasting the estimation problem as a prediction problem leading to the utilization of machine learning techniques. For example, past works have utilized various machine learning methods such as elastic net penalized regression Doudchenko & Imbens (2016), various implementations of lasso (Li & Bell 2017, Kinn 2018, Carvalho et al. 2018, Hollingsworth & Wing 2020), matrix completion methods (Athey et al. 2020), and random forests (Mhlbach 2020). Similar to these works, the proposed approach incorporates machine learning methods to construct a synthetic control counterfactual.

Another strand of the literature studies the inferential challenges of counterfactual estimation with non-stationary processes (see, for example Carvalho et al. 2016, Li 2019, Masini & Medeiros 2020, 2021, Cattaneo et al. 2021, Chernozhukov et al. 2021). The authors identify the counterfactual estimate is not guaranteed to be consistent when integrated processes are present and have non-standard distributions. Proposed solutions include first-differencing the data, subsampling methods, and simulation based approaches. However, this strand of literature does not explicitly address time varying parameters.

Bayesian methods have also been used to construct the counterfactual in a synthetic control setting (some examples include Amjad et al. 2018, Kim et al. 2020, Samartsidis et al. 2020). A popular Bayesian approach is Brodersen et al. (2015). The authors model the counterfactual using a combination of spike and slab priors and linear Gaussian state

space modeling. Unlike other synthetic control approaches, the authors create counterfactuals from untreated time series - not necessarily untreated units. For example, they advocate for the use of Google trends data when untreated units are not available. In their model, the authors allow the coefficients of the untreated units to be constant or dynamic. However, they warn of the dangers of overfitting and implausibly large credibility intervals with dynamic coefficients. These issues occur because of model misspecification. The author's implementation of dynamic coefficients forces every coefficient to be dynamic when it is more plausible that some coefficients are dynamic while others are static.

This paper solves the issues Brodersen et al. (2015) faced when incorporating dynamic coefficients. This is done through three changes to their model. First, the proposed model incorporates the decomposition of time varying coefficients. This allows for individual shrinkage to occur on the time varying and time invariant portion of the coefficient. Now, the dynamic components can be shrunk towards zero to better resemble static relationships when appropriate. Second, the model uses a different set of priors to create the Bayesian Lasso. Adding the parameter decomposition and Bayesian Lasso allows for the use of time varying parameters without implausibly large credibility intervals. Finally, the local linear trend is replaced with the local trend. This is done to limit the risk of overfitting. Although their paper presents the model with a local linear trend, Brodersen et al. (2015) accompanying R package also defaults to a local trend.

Pang et al. (2021) also propose a Bayesian method to explicitly estimate the latent factors serving as a Bayesian extension of ?. Rather than estimating the relationship between treated and untreated units directly, the method uses the untreated units to estimate the factor loadings which are then used to construct the counterfactual. If explanatory variables are included (i.e. controls that aren't untreated units), the relationships are modeled as time varying utilizing non-centered parameterization proposed

by Frhwirth-Schnatter & Wagner (2010). This differs from the proposed method of using the directly estimated relationship between treated and control units to construct a counterfactual. I compare Pang et al. (2021) to the proposed model with both time varying and time-invariant data generating processes in simulations.

First, I propose a time varying parameter model based on recent macroeconomic advances to construct a counterfactual. Second, I compare popular synthetic control approaches, including Brodersen et al. (2015) and Abadie et al. (2010), to my model in two simulation studies. The first is with a time varying data generating process and the second is an empirical monte carlo simulation based on the classic California Tobacco Tax dataset. Finally, I revisit the effect of California passing a tobacco tax on cigarette sales.

## 3.2 Setup

### 3.2.1 Potential Outcomes

I define the treatment as an intervention or policy change. In following with Abadie & Gardeazabal (2003) and Brodersen et al. (2015), suppose only one unit is treated in period  $T_0$ . Once treated, the unit remains treated indefinitely.

Let  $(y_t(0), y_t(1))$  represent potential outcomes in the presence and absence of a treatment with  $t = 1, \dots, T_0 - 1, T_0, T_0 + 1, \dots, T$  and the period of intervention as  $T_0$ . Define the treatment status as  $d_t = I(t \geq T_0)$ . Define  $y_t = (1 - d_t)y_t(0) + d_t y_t(1)$ . Suppose there are  $J + 1$  untreated covariates observed. Let  $\mathbf{x}_t = [x_{1,t}, \dots, x_{J+1,t}]$  be a vector of covariates in period  $t$ . This can include untreated units, as in Abadie et al. (2010), as well as additional controls such as Google Trends used in Brodersen et al. (2015). Define the pre-treatment outcome vector  $\mathbf{y} = [y_1(0), \dots, y_{T_0-1}(0)]$  and the vector matrix of



pre-treatment covariates as  $\mathbf{x}^{\mathbf{T}} = [\mathbf{x}_1 \dots \mathbf{x}_{T_0-1}]_{(J+1) \times T_0-1}$ .

Suppose in each period, we observe  $(y_t, d_t, \mathbf{x}_t)$ . The goal is to estimate the sample treatment effect for the treated unit in the post-intervention period (i.e.  $t \geq T_0$ ):

$$\tau_{T_0+i} = y_{T_0+i} - y_{T_0+i}(0) = y_{T_0+i}(1) - y_{T_0+i}(0)$$

for  $i \in \{0, \dots, T - T_0\}$ .  $y_{T_0+i}(1)$ ,  $d_{T_0+i}$ , and  $\mathbf{x}_{T_0+i}$  are observed while  $y_{T_0+i}(0)$  is unobserved.  $y_{T_0+i}(0)$  is a missing value treated as a random variable.

In order to draw causal inference, I make two common assumptions in the synthetic control literature:

Conditional Independence on Observed Outcomes:

$$y_{T_0+i}(0) \perp d_{T_0+i} | \mathbf{y}, \mathbf{x}, \mathbf{x}_{T_0+i} \quad (3.1)$$

for  $i \in \{0, \dots, T - T_0\}$ .

No Spillovers: Suppose the covariate  $x_{k,t} = f_k(\mathbf{z})$  where  $f$  is some functional form and  $\mathbf{z}$  is a combination of observables and parameters such that  $d_t \notin \mathbf{z} \forall t$ . Then  $f_k(\mathbf{z}) = f_k(\mathbf{z}, d_1, \dots, d_T)$  for all  $i$ .

The conditional independence assumption uses the full set of pretreatment outcomes and untreated units to proxy for potential unobserved confounders. This allows for modeling the unobserved potential outcome in the post period. This assumption suggests a close fit in the pre-treatment periods accurately captures the underlying data generating process yielding a valid counterfactual. If unobserved confounders affect the potential outcome of interest that aren't accounted for in the pre-period, this assumption may not hold and the estimate may be subject to bias.

Assumption 2 guarantees the covariates are not affected by the treatment. If the covariates were affected, the counterfactual estimation could be biased in either direction depending on the spillover effects.

Given these two assumptions, estimating  $\tau_t$  for  $t \geq T_0$  is a problem of estimating the missing potential outcome. Let  $g$  be a function with parameters  $\mathbf{v}$  such that  $y_t(0) = g_{\mathbf{v}}(\mathbf{y}, \mathbf{x}, \mathbf{x}_t)$ . The estimand of interest can be rewritten as:

$$\tau_{T_0+i} = y_{T_0+i} - g_{\mathbf{v}}(\mathbf{y}, \mathbf{x}, \mathbf{x}_{T_0+i}) \quad (3.2)$$

$$\Delta_{\tau} = \frac{1}{T - T_0} \sum_{i=0}^{T-T_0} \tau_{T_0+i} \quad (3.3)$$

### 3.2.2 Functional Form

The functional form  $g$  with parameters  $\mathbf{v}$  is defined using a state space framework:

$$g_{\mathbf{v}}(\mathbf{y}, \mathbf{x}, \mathbf{x}_t) = \sum_{j=1}^{J+1} \beta_{j,t} x_{j,t} + \epsilon_t \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2) \quad (3.4)$$

$$\beta_{j,t} = \beta_{j,t-1} + \eta_{j,t} \quad \eta_{j,t} \sim \mathcal{N}(0, \theta_j) \quad \forall j \quad (3.5)$$

$$\beta_{j,0} \sim \mathcal{N}(\beta_j, \theta_j P_{jj}) \quad \forall j \quad (3.6)$$

where  $x_{J+1,t} = 1$ . The parameters of the model are  $\mathbf{v} = \{\sigma^2, \beta_1, \dots, \beta_{J+1}, \theta_1, \dots, \theta_{J+1}\}$  with  $\epsilon_t$  and  $\eta_{j,t}$  assumed independent of all other unknowns. This specification follows from Brodersen et al. (Brodersen et al. (2015)) with a local trend in place of a local linear trend and dynamic coefficients. The functional form explicitly incorporates  $\mathbf{x}_t$  and implicitly incorporates  $\mathbf{y}$  and  $\mathbf{x}$  through the calculations of  $\beta_{j,t}$ .

The coefficients are modeled as random walks, a common choice in state space lit-

erature (Belmonte et al. 2014, Bitto & Frhwirth-Schnatter 2019). The random walks allows for the coefficients to quickly learn changes in underlying relations. As noted by Dagl & Halling (2012), modeling time varying parameters as random walks is commonly used because of the superior empirical predictive performance of the specification rather than theoretical justification. In addition, the random walk specification allows for a coefficient decomposition to improve counterfactual estimation.

The variances,  $\theta_j$ , are traditionally defined by the inverse gamma distribution, as is the case for Brodersen et al. (2015) implementation of dynamic coefficients. However, the inverse gamma does not allow for effective shrinkage given its nonnegative support. Frhwirth-Schnatter & Wagner (2010) provide an in-depth argument for the use of the normal distribution as an alternative. Briefly, the inverse gamma prior performs poorly in terms of shrinkage due to 0 being an extreme value in the distribution, limiting the amount of mass which can be placed at 0 and in turn limiting the amount of shrinkage. This becomes problematic when many parameters are believed to be equal to zero. In this context, using an inverse gamma to define  $\theta_j$  forces all the coefficients to have some time varying aspect. In application, a researcher may expect some coefficients to be time varying while others are static. Forcing all coefficients to be time varying can cause overfitting, leading to poor counterfactual estimation and poor inference. However, forcing all coefficients to be static can also lead to poor counterfactual estimation because of model misspecification. To account for this, I decompose equation (3.5) into a time varying and constant component:

$$\beta_{j,t} = \beta_j + \tilde{\beta}_{j,t} \sqrt{\theta_j} \quad (3.7)$$

$$\tilde{\beta}_{j,t} = \tilde{\beta}_{j,t-1} + \tilde{\eta}_{j,t} \quad \tilde{\eta}_{j,t} \sim N(0, 1) \quad (3.8)$$

$$\tilde{\beta}_{j,0} \sim N(0, P_{jj}) \quad (3.9)$$

Priors are provided in Section 3.3.1.  $\beta_j$  can now be interpreted as the time invariant component of  $\beta_{j,t}$  and  $\sqrt{\theta_j} \tilde{\beta}_{j,t}$  the time varying component.  $\sqrt{\theta_j}$  is defined as the root of  $\theta_j$  and allowed to take both positive and negative values. Defining  $\sqrt{\theta_j}$  in this manner allows 0 to be an interior point in the prior distribution. This is a desirable feature when performing Bayesian shrinkage (Bitto & Frhwirth-Schnatter (2019)). The absolute value of  $\sqrt{\theta_j}$  is the standard deviation of the time varying coefficient. Substituting the reformulation back into the original equation yields the proposed state space model.

The Time Varying Parameter Bayesian Lasso ( BL-TVP ) takes the following functional form:

$$g_{\mathbf{v}}(\mathbf{y}, \mathbf{x}, \mathbf{x}_t) = \sum_{j=1}^{J+1} \left( \beta_j + \tilde{\beta}_{j,t} \sqrt{\theta_j} \right) x_{j,t} + \epsilon_t \quad \epsilon_t \sim N(0, \sigma^2) \quad (3.10)$$

$$\tilde{\beta}_{j,t} = \tilde{\beta}_{j,t-1} + \tilde{\eta}_{j,t} \quad \tilde{\eta}_{j,t} \sim N(0, 1) \quad (3.11)$$

$$\tilde{\beta}_{j,0} \sim N(0, P_{jj}) \quad (3.12)$$

$$\pi(\mathbf{v}) \quad (3.13)$$

$\pi(\mathbf{v})$  represents the prior distribution of the parameters with the specific distributions defined in Section 3.3.1. Equations (3.10) - (3.13) constitute the model. This setup is commonly known as the *non-centered parameterization of state space models*. This

formulation allows estimation of the time varying and time invariant component of the coefficients individually. The relationship between each control unit and the treated unit can be summarized into one of the four categories: (i) time varying non-zero, (ii) time invariant non-zero, (iii) time varying centered at zero, and (iv) time invariant zero coefficients (irrelevant).

Notice if  $\sqrt{\theta_j} = 0$  for all  $j$ , the model is a Bayesian version of the Lasso estimator discussed in Kinn (2018). Setting  $\sqrt{\theta_j} = 0$  and restricting  $\beta$  such that  $\beta_j \in [0, 1]$  and  $\sum_j \beta_j = 1$  yields a simplified parametric version of Abadie & Gardeazabal (2003) synthetic control model. Similarly, if the data generating process does not include untreated units (e.g.  $\beta_j = \sqrt{\theta_j} = 0$  for all  $j$ ), then (3.10) - (3.13) collapses to a local level model. Setting  $\sqrt{\theta_j} = 0$  and using the Horseshoe prior for  $\beta$  yields Kim et al. (2020).

### 3.3 Estimation of Parameters and Counterfactual

Model 1 follows a standard state space formulation and can be solved with such methods. A Bayesian approach is used to estimate the parameters of interest. This approach lends itself well to the synthetic control setting with few pre-treatment periods.

To estimate the posterior parameter distribution of the coefficients  $\beta_j$  and  $\sqrt{\theta_j}$ , I incorporate the hierarchical Bayesian Lasso (Park & Casella 2008, Belmonte et al. 2014, Bitto & Frhwirth-Schnatter 2019, Pang et al. 2021).

A drawback of the Bayesian Lasso is the approach lacks sparsity - all coefficients are biased towards zero but no coefficients is set exactly to zero. In application, this means there is no easy-to-interpret inclusion probability as is seen with priors such as slab and spike. The benefits are two-fold. First, Bayesian Lasso does not suffer from the same computation issues as slab-and-spike (Bitto & Frhwirth-Schnatter 2019). Second, recent studies in macroeconomic forecasting found evidence economic data tends to be

dense rather than sparse supporting the use of hierarchical shrinkage priors over sparse estimation strategies like spike-and-slab (Cross et al. 2020, Giannone et al. 2021).

### 3.3.1 Bayesian Shrinkage Priors

I set up the prior distribution for coefficients  $\beta = [\beta_1, \beta_2, \dots, \beta_{J+1}]$  with variances  $\alpha^2 = [\alpha_1^2, \alpha_2^2, \dots, \alpha_{J+1}^2]$ :

$$\beta | \alpha^2 \sim \mathcal{N}_{J+1}(0_{J+1}, \text{diag}[\alpha_1^2, \dots, \alpha_{J+1}^2]) \quad (3.14)$$

$$\alpha_j^2 | \lambda^2 \sim \exp\left(-\frac{\lambda^2}{2}\right) \quad (3.15)$$

$$\lambda^2 \sim \text{Gamma}(z_1, z_2) \quad (3.16)$$

$$z_1, z_2 \geq 0 \quad (3.17)$$

The hierarchical formulation of  $\beta$  and  $\alpha^2$  are then identical to a priori independent Laplace priors. (Park & Casella 2008) show this choice of priors leads to posterior performance similar to the frequentist machine learning approach LASSO (Tibshirani 1996).

Similar to  $\beta$ , let  $\sqrt{\theta} = [\sqrt{\theta_1}, \sqrt{\theta_2}, \dots, \sqrt{\theta_{J+1}}]$  with variances  $\xi^2 = [\xi_1^2, \xi_2^2, \dots, \xi_{J+1}^2]$  have the following prior:

$$\sqrt{\theta}|\xi^2 \sim \mathcal{N}_{J+1}(0_{J+1}, \text{diag}[\xi_1^2, \dots, \xi_{J+1}^2]) \quad (3.18)$$

$$\xi_j^2 | \kappa^2 \sim \exp\left(\frac{\kappa^2}{2}\right) \quad (3.19)$$

$$\kappa^2 \sim \text{Gamma}(z'_1, z'_2) \quad (3.20)$$

$$z'_1, z'_2 \geq 0 \quad (3.21)$$

$\sigma^2$  is defined as  $\frac{1}{\sigma^2} \sim \text{Gamma}(a_1, a_2)$  with *shape* hyperparameter  $a_1$  and *scale* parameter  $a_2 \sim \text{Gamma}(g_0, G_0)$  and  $\frac{1}{P_{jj}} \sim \text{Gamma}(20, 19)$ . Following Bitto & Frhwirth-Schnatter (2019), set  $c_0 = 2.5$ ,  $g_0 = 5$ , and  $G_0 = \frac{g_0}{\mathbb{E}[\sigma^2](c_0-1)}$  where  $\mathbb{E}[\sigma^2]$  being a best guess of  $\sigma^2$ .

### 3.4 The Posterior Estimation (MCMC)

In order to draw the counterfactual, the posterior distribution must be calculated. With values drawn from the posterior distribution,  $g_{\hat{\nu}}(\mathbf{y}, \mathbf{x}, \mathbf{x}_t)$  can then be estimated for  $T_0 + i \geq T_0$ . A closed form does not exist for the posterior. Therefore, I implement Gibbs sampling. After a sufficiently large initial sample, or burn in, the draws from the conditional posterior will be simulations of the joint posterior.

Various MCMC algorithms have been proposed to estimate such models (Belmonte et al. 2014, Bitto & Frhwirth-Schnatter 2019). I implement the algorithm proposed in Bitto & Frhwirth-Schnatter (2019). Two key features of the algorithm are the implementation of the *all without a loop* algorithm (McCausland et al. 2011) and *ancillarity-sufficiency interweaving strategy* (ASIS) (Yu & Meng 2011). The *All without a loop* algorithm smooths states more efficiently than Kalman filter based methods adding efficiency to the MCMC strategy. ASIS interweaves sampling from both the centered and

non-centered parameterizations to boost efficiency compared to MCMC algorithm using only centered or non-centered parameterizations. This addresses MCMC convergence issues associated with centered and non-centered state space models (Roberts & Sahu (1997)).

The steps of the algorithm are briefly summarized below:

1. Sample the state using AWOL proposed by McCausland et al. (2011) and implemented in Bitto & Frhwirth-Schnatter (2019).
2. Block draw  $\beta$  and  $\sqrt{\theta}$  from the normal conditional posterior:

$$\mathcal{N}_{2(J+1)}((\tilde{\mathbf{x}}^T \tilde{\mathbf{x}} + \sigma^2 V^{-1})^{-1} \tilde{\mathbf{x}}^T \mathbf{y}, \sigma^2 (\tilde{\mathbf{x}}^T \tilde{\mathbf{x}} + \sigma^2 V^{-1})^{-1}) \quad (3.22)$$

Where:

$$\tilde{\mathbf{x}} = \begin{pmatrix} x_{1,1} & x_{2,1} & \dots & x_{J+1,1} & \tilde{\beta}_{1,1} x_{1,1} & \tilde{\beta}_{2,1} x_{2,1} & \dots & \tilde{\beta}_{J+1,1} x_{J+1,1} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_{1,T_0-1} & x_{2,T_0-1} & \dots & x_{J+1,T_0-1} & \tilde{\beta}_{1,T_0-1} x_{1,T_0-1} & \tilde{\beta}_{2,T_0-1} x_{2,T_0-1} & \dots & \tilde{\beta}_{J+1,T_0-1} x_{J+1,T_0-1} \end{pmatrix} \quad (3.23)$$

$$V = \text{diag} [\alpha_1^2, \alpha_2^2, \dots, \alpha_{J+1}^2, \xi_1^2, \xi_2^2, \dots, \xi_{J+1}^2] \quad (3.24)$$

3. Perform *ASIS* (Yu & Meng 2011). Namely, the model is transformed to the centered parameterization.  $\beta_j$  and  $\sqrt{\theta_j}$  are drawn from centered posterior distributions



described in the online appendix. The state process  $\tilde{\beta}_{j,t}$  is then updated with the newly drawn  $\beta_j$  and  $\sqrt{\theta_j}$  from the centered parameterization.  $\beta_j$  and  $\sqrt{\theta_j}$  are then converted back to the non-centered parameterization.

4. Sample the remaining parameters from the following posterior conditional distributions:

$$\alpha_j^2 \sim GIG\left(\frac{1}{2}, \lambda^2, \beta_j^2\right) \quad (3.25)$$

$$\lambda^2 \sim Gamma\left(z_1 + \frac{1}{J+1}, z_2 + \frac{1}{2} \sum_{j=1}^{J+1} \alpha_j^2\right) \quad (3.26)$$

$$\xi_j^2 \sim GIG\left(\frac{1}{2}, \kappa^2, \theta_j\right) \quad (3.27)$$

$$\kappa^2 \sim Gamma\left(z'_1 + \frac{1}{J+1}, z'_2 + \frac{1}{2} \sum_{j=1}^{J+1} \xi_j^2\right) \quad (3.28)$$

$$\sigma^2 \sim InverseGamma\left(a_1 + \frac{T_0 - 1}{2}, a_2 + \frac{\sum_{t=1}^{T_0-1} \left(y_t - \sum_{j=1}^{J+1} (\beta_j + \tilde{\beta}_{j,t} \sqrt{\theta_j}) x_{j,t}\right)^2}{2}\right) \quad (3.29)$$

$$a_2 \sim Gamma\left(g_0 + a_1, G_0 + \frac{1}{\sigma^2}\right) \quad (3.30)$$

$$\frac{1}{P_{0,jj}} \sim Gamma\left(20 + \frac{1}{2}, 19 + \frac{\tilde{\beta}_{j0}^2}{2}\right) \quad (3.31)$$

where *GIG* is the generalized inverse Gaussian distribution. Frhwirth-Schnatter & Wagner (2010) note an identification problem arises when using the non-centered parameterization. There is no way to distinguish between  $\sqrt{\theta_j} \tilde{\beta}_{j,t}$  and  $(-\sqrt{\theta_j})(-\tilde{\beta}_{j,t})$ . This problem is referred to as *label switching problem*. This issue is a common occurrence in Bayesian estimation when a distribution is multi-modal, as is the case with the square

root of a variance. To solve this identification problem, Frhwirth-Schnatter & Wagner (2010) suggest a random sign change at the end of each iteration of the Gibbs sampler. With 50% chance, the signs on  $\tilde{\beta}$  and  $\sqrt{\theta}$  are switched. Both Belmonte et al. (2014) and Bitto & Frhwirth-Schnatter (2019) employ this method.

A final note of interest is the formulation of  $\lambda^2$  (and  $\kappa^2$ ). The conditional distribution of  $\lambda^2$  relies on  $\sum_{j=1}^{J+1} \alpha_j^2$  where each posterior  $\alpha_j^2$  relies on  $\beta_j$ . This direct reliance on  $\beta_j$  in the conditional distributions can lead to scaling issues. Data bigger in magnitude can dominate the distribution of  $\lambda^2$ . The issue of scaling is common in both parametric and nonparametric shrinkage estimation. To account for this, all control units (except the intercept) are scaled to mean zero variance one prior to analysis.

### 3.4.1 Sample of $g_{\hat{\nu}}(\mathbf{y}, \mathbf{x}, \mathbf{x}_{\mathbf{T}_0+i})$ for $i \geq 0$ .

After a sufficiently large burn in period, use the proceeding draws to calculate  $g_{\hat{\nu}}(\mathbf{y}, \mathbf{x}, \mathbf{x}_{\mathbf{T}_0+i})$  for  $T_0 + i \geq T_0$ . Namely, perform the following steps:

- (1) Simulate  $\tilde{\beta}_{j,t} = \tilde{\beta}_{j,t-1} + \tilde{\eta}_{j,t}$  for all  $j$ . Use  $\tilde{\beta}_{j,T_0-1}$  simulated in section 4 as an initial value. Each iteration of the Gibbs sampler will create a new  $\tilde{\beta}_{j,T_0-1}$ .
- (2) Using the simulated  $\tilde{\beta}_{j,t}$ , predict  $g_{\hat{\nu}}(\mathbf{y}, \mathbf{x}, \mathbf{x}_{\mathbf{T}_0+i})$  as:

$$g_{\hat{\nu}}(\mathbf{y}, \mathbf{x}, \mathbf{x}_{\mathbf{T}_0+i}) = \sum_{j=1}^{J+1} \left( \beta_j + \tilde{\beta}_{j,T_0+i} \sqrt{\theta_j} \right) x_{j,T_0+i} + \epsilon_{T_0+i}$$

drawing  $\epsilon_{T_0+i} \sim N(0, \sigma^2)$ . Each iteration of the Gibbs sampler will produce new parameter and state values.

### 3.4.2 Estimate $\hat{\tau}_{T_0+i}$ and $\hat{\Delta}_{\tau}$

The sample average treatment effect on the treated is:

$$\hat{\tau}_{T_0+i} = y_{T_0+i} - g_{\hat{\nu}}(\mathbf{y}, \mathbf{x}, \mathbf{x}_{T_0+i}) \quad (3.32)$$

for  $i = \{0, \dots, T - T_0\}$ . The estimated value of the treatment effect in each simulation creates an empirical distribution of the effect. This allows for valid inference even in the presence of small sample size. The sample average treatment effect on the treated in the post period is then calculated as:

$$\hat{\Delta}_{\tau} = \frac{1}{T - T_0} \sum_{i=0}^{T-T_0} \hat{\tau}_{T_0+i} \quad (3.33)$$

Sampling from the Gibbs sampler creates an empirical distribution for the treatment effects. Statistical testing can then be performed with the distribution. A major benefit of this approach is valid credibility intervals.

### 3.5 Simulation Studies

To test the proposed model's performance, I develop a simulation study comparing seven synthetic control approaches: 1) the original Brodersen et al. (2015) model (*CI*), 2) Brodersen et al. (2015) with time varying coefficients as presented in the `CausalImpact` R documentation (*CI-TVP*), 3) Abadie et al. (2010) synthetic control (*SC*), 4) The proposed model (*BL-TVP*), 5) Carvalho et al. (2018) (*ArCo*), 6) Pang et al. (2021) (*DM-LFM*), and 7) Kim et al. (2020) with horseshoe priors (*BSCM-Horseshoe*). These methods include popular synthetic control estimation approaches from both the frequentist and Bayesian perspectives.

The simulations focus on three metrics of interest: the mean squared forecast error

(MSFE), credibility/confidence interval spread for  $\hat{\Delta}_\tau$  (CI Spread), and coverage. MSFE, CI Spread, and coverage are calculated as:

$$MSFE = \frac{1}{T - T_0} \sum_{i=0}^{T-T_0} (y_{T_0+i} - \hat{y}_{T_0+i})^2 \quad (3.34)$$

$$CI \text{ Spread} = \hat{\Delta}_\tau^{.975} - \hat{\Delta}_\tau^{.025} \quad (3.35)$$

$$Coverage = \frac{1}{T - T_0} \sum_{i=0}^{T-T_0} \mathbb{I} \left( y_{T_0+i} \in [\hat{\Delta}_\tau^{.025}, \hat{\Delta}_\tau^{.975}] \right) \quad (3.36)$$

where  $\hat{\Delta}_\tau^x$  is the  $x^{th}$  percentile from the simulated distribution. The mean squared forecast error is used to identify the closeness of the fit. Lower mean squared forecast error suggests a better counterfactual measurement. However, a low mean squared forecast error with implausibly large credibility intervals will lead to poor inference. The second metric is used to compare the models inferential capabilities.

### 3.5.1 Time Varying Parameter DGP

The first data generating process is based on simulations performed in Brodersen et al. (2015) and the Abadie et al. (2010) data. I define 4 untreated time series  $x_{1,t}$ ,  $x_{2,t}$ ,  $x_{3,t}$  and  $x_{4,t}$  where  $t \in \{1, 2, \dots, 31\}$ . Treatment occurs in  $t = 19$  and treatment  $\tau_t = 0 \forall t$ . The untreated variables are generated from the following processes:

$$x_{1,t} = \sin \left( \frac{\pi t}{12} \right)$$

$$x_{2,t} = \cos \left( \frac{\pi t}{6} \right)$$

$$x_{3,t} = 3 + \sin(t)$$

$$x_{4,t} = \cos \left( \frac{t}{2} \right)$$

$y_t$  is a linear combination of  $x_{1,t}$ ,  $x_{2,t}$ , and  $x_{3,t}$ . I then define  $y_t$  as:

$$\begin{aligned}
 y_t &= \mu_t + \sum_{i=1}^4 \left( \tilde{\beta}_{i,t} \sqrt{\theta_i} + \beta_i \right) x_{i,t} + \epsilon_t & \epsilon_t &\sim \mathcal{N}(0, 1) \\
 \tilde{\beta}_{i,t+1} &= \tilde{\beta}_{i,t} + \eta_{i,t+1} & \eta_{i,t+1} &\sim \mathcal{N}(0, 1) \quad i \in \{1, 2, 3, 4\} \\
 \mu_{t+1} &= \mu_t + \eta_{\mu,t+1} & \eta_{\mu,t+1} &\sim \mathcal{N}(0, .1^2) \\
 \beta_{1,0} &= \beta_{2,0} = \beta_3 = \mu_0 = 0
 \end{aligned}$$

Where the parameters  $\left[ \left( \beta_1, \sqrt{\theta_1} \right), \left( \beta_2, \sqrt{\theta_2} \right), \left( \beta_3, \sqrt{\theta_3} \right), \left( \beta_4, \sqrt{\theta_4} \right) \right] = [(1, 2), (1, 0), (0, 3), (0, 0)]$ . This produces all possible parameter combinations: non-zero time varying ( $x_1$ ), non-zero time invariant ( $x_2$ ), mean zero time varying ( $x_3$ ), and irrelevant ( $x_4$ ). A graph of one of the simulation iterations is presented in the online appendix.

BL-TVP, SC, ArCo, BSCM-Horseshoe, DM-LFM, CI, and CI-TVP are compared in terms of MSFE, CI Spread and coverage. The simulation is run 1000 times. Table 3.1 displays the results from the simulation. BL-TVP is the only model to obtain above 95% coverage and has the lowest reported MSFE. CI-TVP produces a similarly sized credibility spread to BL-TVP but does not achieve 95% coverage.

[Table 3.1 about here]

Figure 3.1 plots the CI Spread, coverage rate, and MSFE for the seven methods in the post-period averaged over 1000 simulations. BL-TVP is the only method that maintains at least 95% coverage every period. It also produces the lowest MSFE in almost every period and comparable CI Spread to CI-TVP.

[Figure 3.1 about here]

### 3.5.2 Empirical Monte Carlo: California Tobacco Tax

To test the proposed model’s performance with misspecification, I develop an empirical Monte Carlo simulation. Empirical Monte Carlo studies fit a data generating process to well-known datasets allowing for a simulation study based in real world data (see Ben-Michael et al. 2021, Kellogg et al. 2021, for additional examples in the synthetic control literature). The empirical Monte Carlo is based on data used in Abadie et al. (2010). The authors studied the impacts of a major tobacco tax implemented only in California. The outcome variable was per capita cigarette sales (in packs). The analysis begins in 1970 and ends in 2000 with California as the single treated unit. The policy was passed in January 1989, leading to 19 years of pre-treatment. 38 other states were used as control units.

Algorithm 1 describes the empirical Monte Carlo exercise. A placebo is created for each untreated unit using a Bayesian regression. The model is then used to create 1000 simulated placebos. Each synthetic control approach is then performed on every simulated placebo assuming the date of treatment is 1989 and  $\tau_t = 0$  for all  $T$  yielding 38,000 simulations in the empirical Monte Carlo. A current issues researchers face applying Brodersen et al. (2015) with time varying parameters is the implausibly large credibility intervals. If a researcher is unsure of the presence of time varying parameters, they have one of two options: (1) run a model that assumes all parameters are invariant and risk model misspecification or (2) run Brodersen et al. (2015) with time varying parameters and have credibility intervals unable to reject most treatment effects.

[Algorithm 1 about here]

This specification induces model misspecification for all models except DM-LFM. Since CI-TVP and BL-TVP include time varying parameters and a time varying intercept, both experience the same level of model misspecification. ArCo, CI, and SC induce

sparsity in the estimation processes creating misspecification because the data generating processes is dense. In addition, SC requires the covariates to sum to one while the data generating process does not have such a restriction.

The data generating process is estimated using the `bayesreg` package and setting the priors to lasso (Makalic & Schmidt 2016). This does not put restrictions on any of the coefficients. In addition, the posterior estimation creates an empirical distribution from each draw allowing each iteration simulation to have a unique relationship between treated and controls. the online appendix provides an example of the Monte Carlo process.

Figure 3.2 compares the ratio of credibility spread and MSFE between the models and BL-TVP. Each dot represents the average ratio between BL-TVP and an alternative model per state over 1000 simulations. If a dot is in the top right quadrant, then BL-TVP produced both a smaller MSFE and credibility spread compared to the alternative model. The top left quadrant signifies BL-TVP produced a lower MSFE than the alternative model but larger credibility interval spread. The bottom left quadrant signifies BL-TVP produced both larger MSFE and credibility interval spread and the bottom right quadrant signifies BL-TVP produced a smaller credibility spread but larger MSFE. The difference in credibility spread between SC and BL-TVP is set to 0 due to SC not producing confidence intervals.

[Figure 3.2 about Here]

BL-TVP maintains a comparable MSFE to CI-TVP for most of the simulations. The MSFE tends to be lower for BL-TVP compared to CI. Of the 38 states, BL-TVP produces a lower MSFE for 28 of the states compared to CI, 18 of the states compared to CI-TVP, 22 of the states compared to SC, 25 compared to ArCo, 24 compared to DM-LFM, 19 compared to BSCM-Horseshoe.

The added flexibility of BL-TVP leads to larger credibility intervals than CI, ArCo, and DM-LFM. This is expected since the data generating process only includes fixed covariates. However, BL-TVP produces credibility intervals between CI than CI-TVP. CI-TVP produces inflated credibility intervals due to model misspecification. BL-TVP demonstrates an ability to balance the additional flexibility of time varying parameters without the risk of implausibly large credibility intervals due to model misspecification.

CI produced smaller credibility intervals for most states compared to BL-TVP. On average, CI produced 46.3% smaller credibility spread. BL-TVP produces credibility intervals 233.2% larger than ArCo confidence intervals on average, 139.5% larger credibility intervals than DM-LFM on average and 188.8% larger credibility intervals than BSCM-Horseshoe on average. In contrast, CI-TVP produced 221.7% larger credibility spread than BL-TVP. The added uncertainty is due to model misspecification. CI-TVP forces all covariates to be time varying. The model can't collapse to the time-invariant case. The simulation table is presented in the online appendix.

The results of the simulation show BL-TVP creates a “best of both worlds” situation compared to CI-TVP. BL-TVP creates a lower MSFE similar to CI-TVP without the implausibly large credibility spread. BL-TVP also creates a similarly sized credibility spread to CI but with a smaller MSFE.

Even with additional model misspecification, BL-TVP is able to produce similarly sized credibility spread to CI with lower MSFE. Compared to CI-TVP, BL-TVP produces comparable MSFE and smaller credibility spread. I also provide simulation results for coverage. Due to space constraints, the results are presented in the online appendix.



## 3.6 Empirical Results: Revisiting California Tobacco Tax

Abadie et al. (2010) studied the impacts of a major tobacco tax implemented only in California. The outcome variable was per capita cigarette sales (in packs). The analysis focuses on the time period 1970-2000 with California as the single treated unit. The policy was passed in January 1989 leading to 19 years of pre-treatment. 38 other states were used as control units. This equates to  $T_0 = 1989 - 1970 = 19$  and  $T = 2000 - 1970 = 30$ . Data was drawn using Cunningham (2021) Github repository. Figure 3.3 presents the synthetic control estimates for each of the models compared to California. California's per capita cigarette sales are normalized to be 0 in each period.

[Figure 3.3 Here]

All the models follow comparable paths. Notice CI has significantly larger credibility intervals in the pre-period compared to BL-TVP suggesting only including time invariant coefficients leads to additional model misspecification. CI-TVP creates a near perfect fit in the pre-period with credibility intervals spanning about 0.65 each period. The tight credibility interval paired with inflated post-period credibility intervals suggest overfitting (Mhlbach 2020). ArCo and DM-LFM also follow comparable paths in the pre-treatment both concluding a significant change in the treatment.

BL-TVP produces tighter intervals in the pre-treatment than CI. Both BL-TVP and CI closely follow the original SC estimate. By utilizing the Bayesian Lasso and non-centered parameterization, the BL-TVP is able to discern similar inference with tighter credibility intervals than CI-TVP. However, BL-TVP produces larger credibility intervals compared to the credibility intervals in DM-LFM, BSCM-Horseshoe, and CI and confidence interval of ArCo due to the added flexibility of the model.

Table 3.2 provides the average reduction in California’s per capita cigarette sales. While the point estimates are similar to Abadie et al. (2010), the proposed estimate suggests the effect was not statistically significant. Other flexible methods have failed to find significant effects. For example, Arkhangelsky et al. (2021) found a similar insignificant point estimate using their approach. One explanation why Abadie et al. (2010) found significant results while other methods, including CI, BL-TVP, and CI-TVP, did not could be due to structural breaks in the control group (Linden & Yarnold 2016). Structural breaks would be a violation of the underlying linear factors model invalidating the synthetic control estimate.

[Table 3.2 Here]

To test the robustness of the results, I perform a placebo test on California. The treatment date is artificially moved up four periods. BL-TVP, CI-TVP, SC, and DM-LFM both track California closely in the psuedo treatment period. BSCM-Horseshoe, CI and ArCo suggests divergence beginning in the placebo period. This suggests BSCM-Horseshoe, CI and ArCo are not appropriate model specifications for California.

[Figure 3.4 Here]

## 3.7 Discussion

The additional flexibility of BL-TVP comes with additional restrictions. These restrictions include parametric assumptions on the errors and prior distributions on the parameters. There are no safeguards against misspecified priors. In a Bayesian setting, this may become a point of concern. I recommend varying the prior distribution of the parameters  $\beta_j$  and  $\sqrt{\theta_j}$  from Bayesian Lasso to other shrinkage methods like the double

gamma or horseshoe prior to test for model sensitivity. This method would be in addition to the in-state and in-time placebo tests common to the literature.

While the BL-TVP focuses on modeling the unobserved potential outcome with observed data, the state space framework allows for modeling additional time series components. For example, a local linear trend can be substituted for the local trend, seasonality and cycles can be incorporated as well as structural breaks Durbin & Koopman (2012). The framework allows the researcher to incorporate case-specific information into the modeling process.

### 3.8 Conclusion

This paper proposes a Bayesian alternative to estimating a counterfactual in a synthetic control setting. BL-TVP builds on recent advances in macroeconomic research to allow the inclusion of time varying parameters while mitigating model misspecification. I do this by applying shrinkage to the time invariant and time varying portion of the coefficient separately allowing for better approximation of static and null coefficients. Simulations and empirical Monte Carlo results suggest the additional flexibility allows for smaller mean squared forecast errors compared to CI without the inflated credibility intervals when using CI-TVP. If an empirical researcher suspects the presence of heterogeneous relationships in their synthetic control analysis, BL-TVP provides an alternative approach that captures the time varying relationship without the risk of implausibly large credibility intervals that plague current methods.

### 3.9 Algorithms, Tables and Graphs

---

**Algorithm 1:** Empirical Monte Carlo steps
 

---

**Data:** Abadie, Diamond, and Hainmueller (2010) cigarette sales per capita  
excluding California

```

1 states  $\equiv$  {States excluding California}
2 models  $\equiv$  {SC, BL-TVP, CI-TVP, CI, ArCo, DM-LFM, BSCM-Horseshoe}
3 for i in states do
4   set  $y_t \equiv$  state i
5   set  $\mathbf{x}_t \equiv$  states excluding state i
6   fit  $y_t = \beta_0 + \mathbf{x}_t\beta + \epsilon_t$  with  $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$  using Bayesian Lasso priors
   (implemented using Makalic and Schmidt (2016)) over all 31 periods
7   for j in 1:1000 do
8     Simulate  $y_t$  from the fitted model generating the  $j^{th}$  monte carlo iteration
     for state i
9     for k in models do
10      construct the synthetic control counterfactual for simulated  $y_t$  as a
      function of  $\mathbf{x}_t$  using model k with treatment occurring in 1989 and
       $\tau_t = 0 \forall t$ 
11      Calculate and save MSFE, CI Spread and coverage rate
  
```

---

Note: CI, CI-TVP, and BL-TVP are set to 20,000 iterations with 1,000 burn in. DM-LFM iteration and burn in is the default package values. BSM-Horseshoe is set to 2,000 iterations with 1,000 burn in.

Table 3.1: TVP simulation results for each model

Model	MSFE	CI Spread	Coverage Rate
ArCo	1167.52	54.31	0.47
BL-TVP	641.97	104.39	0.96
BSCM-Horseshoe	1153.91	56.57	0.63
CI	1219.27	41.14	0.50
CI-TVP	1719.03	106.36	0.82
DM-LFM	1341.04	70.37	0.65
SC	2032.63	NA	NA

*Note:*

Simulation run 1000 times using the specification outlined above. BL-TVP: Bayesian Lasso with Time Varying Parameters. CI: Causal Impact.. CI-TVP: Causal Impact with Time Varying Parameters. SC: Synthetic Control. DM-LFM: Dynamic Multilevel Latent Factor Model. ArCo: Artificial Counterfactual. BSCM-Horseshoe: Kim, Clarence, and Sachin 2020.

Table 3.2: Average Reduction in California’s Per Capita Cigarette Sales (Post Intervention)

Model	Average Decrease	2.5th Percentile Decrease	97.5th Percentile Decrease
ArCo	17.3	11.7	22.8
BL-TVP	17.7	-16.1	51.7
BSCM-Horseshoe	13.1	3.1	22.6
CI	20.7	-3.6	41.0
CI-TVP	12.3	-107.8	129.8
DM-LFM	9.6	1.9	17.2
SC	18.1	NA	NA

*Note:*

BL-TVP: Bayesian Lasso with Time Varying Parameters. CI: Causal Impact. CI-TVP: Causal Impact with Time Varying Parameters. SC: Synthetic Control. DM-LFM: Dynamic Multilevel Latent Factor Model. ArCo: Artificial Counterfactual. BSCM-Horseshoe: Kim, Clarence, and Sachin 2020.

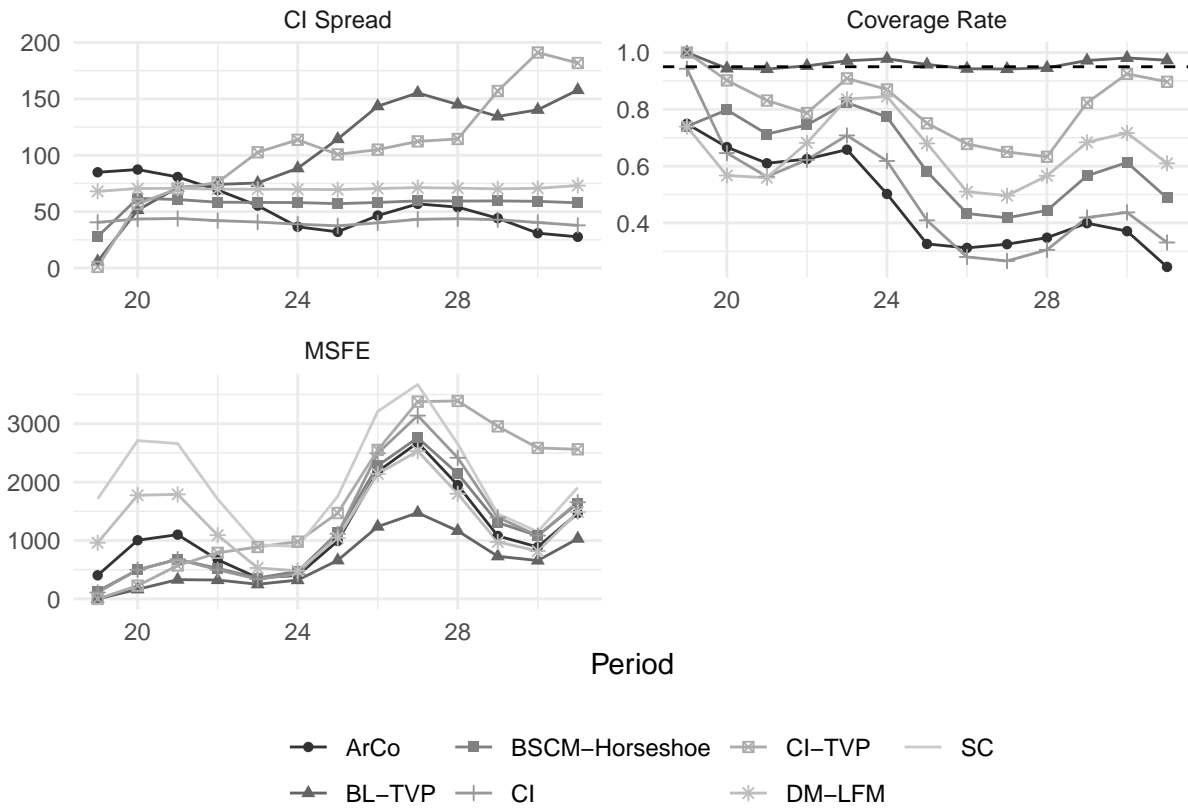


Figure 3.1: Average per post-treatment period per model over 1000 simulations. Treatment begins in period 19. The dotted line is 95 coverage. SC only included in MSFE analysis.




Figure 3.2: Log Ratio Comparison Alternative Model to BL-TVP for time invariant empirical monte carlo simulations. Black lines are at 0. Each dot represents one state. The top left quadrant means BLTVP has a larger credibility/confidence interval spread but smaller MSFE. The top right quadrant means BL-TVP has a smaller credibility/confidence spread and MSFE than the alternative model. The bottom left quadrant means BL-TVP has a larger credibility/confidence spread and MSFE than the alternative model. The bottom right quadrant means BL-TVP has a smaller credibility/confidence spread but larger MSFE than the alternative model.  $\log(\text{CI spread SC}) - \log(\text{CI spread BL-TVP})$  is imputed to 0 due to SC not creating confidence intervals.






Figure 3.3: California (Dashed) vs Synthetic California (Full) with 95% credibility/confidence intervals (when applicable). Treatment start is dotted line. CI-TVP credibility interval extends beyond the graph dimensions. CI-TVP and BSCM-Horseshoe produce tight credibility intervals in the pre-period that do not appear in the graph.




Figure 3.4: Placebo Test: California (Dashed) vs Synthetic California (Full) with 95% credibility intervals (when applicable). Placebo Treatment starts at leftmost dotted line. Actual treatment starts at rightmost dotted line. CI-TVP credibility interval extends beyond the graph dimensions. CI-TVP and BSCM-Horseshoe produce tight credibility intervals in the pre-period that do not appear in the graph.

# Appendix A

## Appendix for “Selling Violent

# Extremism”

## A.1 Descriptive map of Oath Keepers joining

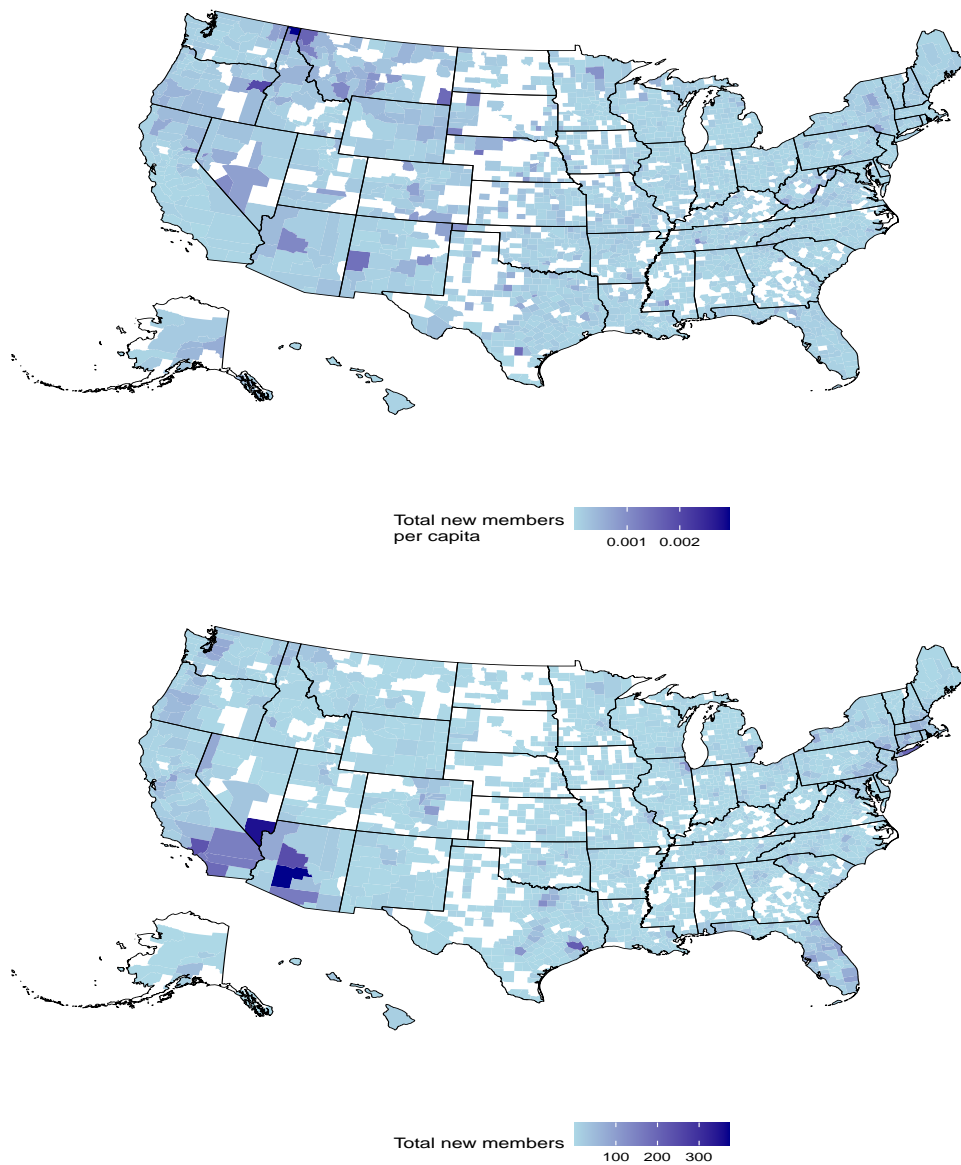


Figure A.1: Total number of Oath Keepers signups between 2009-2018 per capita by county using 2018 population estimates. Counties colored white received no Oath Keepers.

## A.2 Plotted donor pool




Figure A.2: Number Oath Keepers per day (solid black) versus Google trends for given group (dashed blue) during the Veteran’s Day sale.

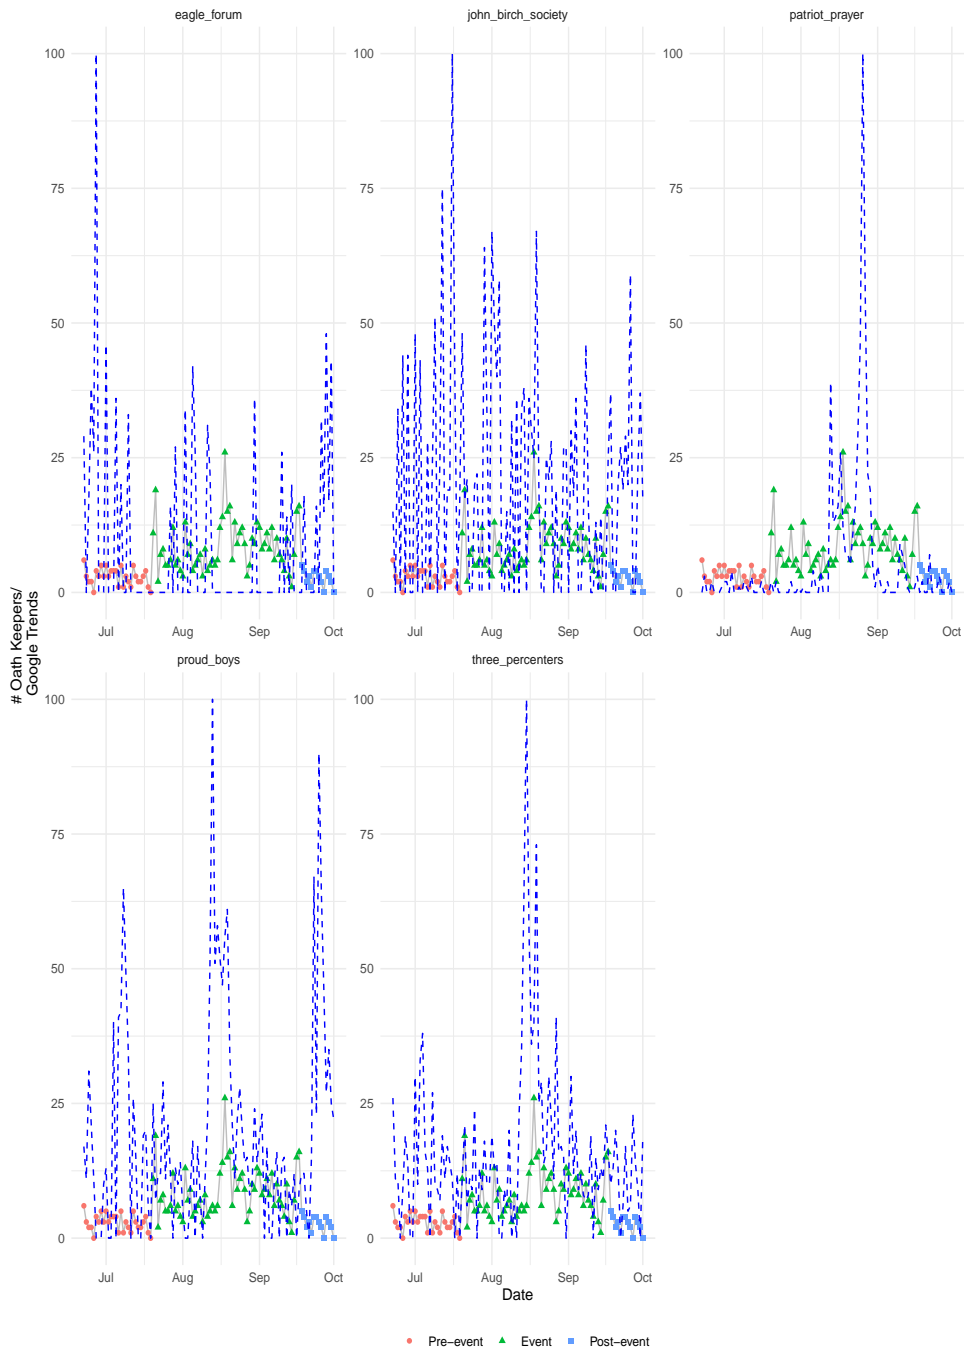


Figure A.3: Number Oath Keepers per day (solid black) versus Google trends for given group (dashed blue) during the Constitution Day Day sale.




Figure A.4: Number Oath Keepers per day (solid black) versus Google trends for given group (dashed blue) during the Christmas/New Years sale.

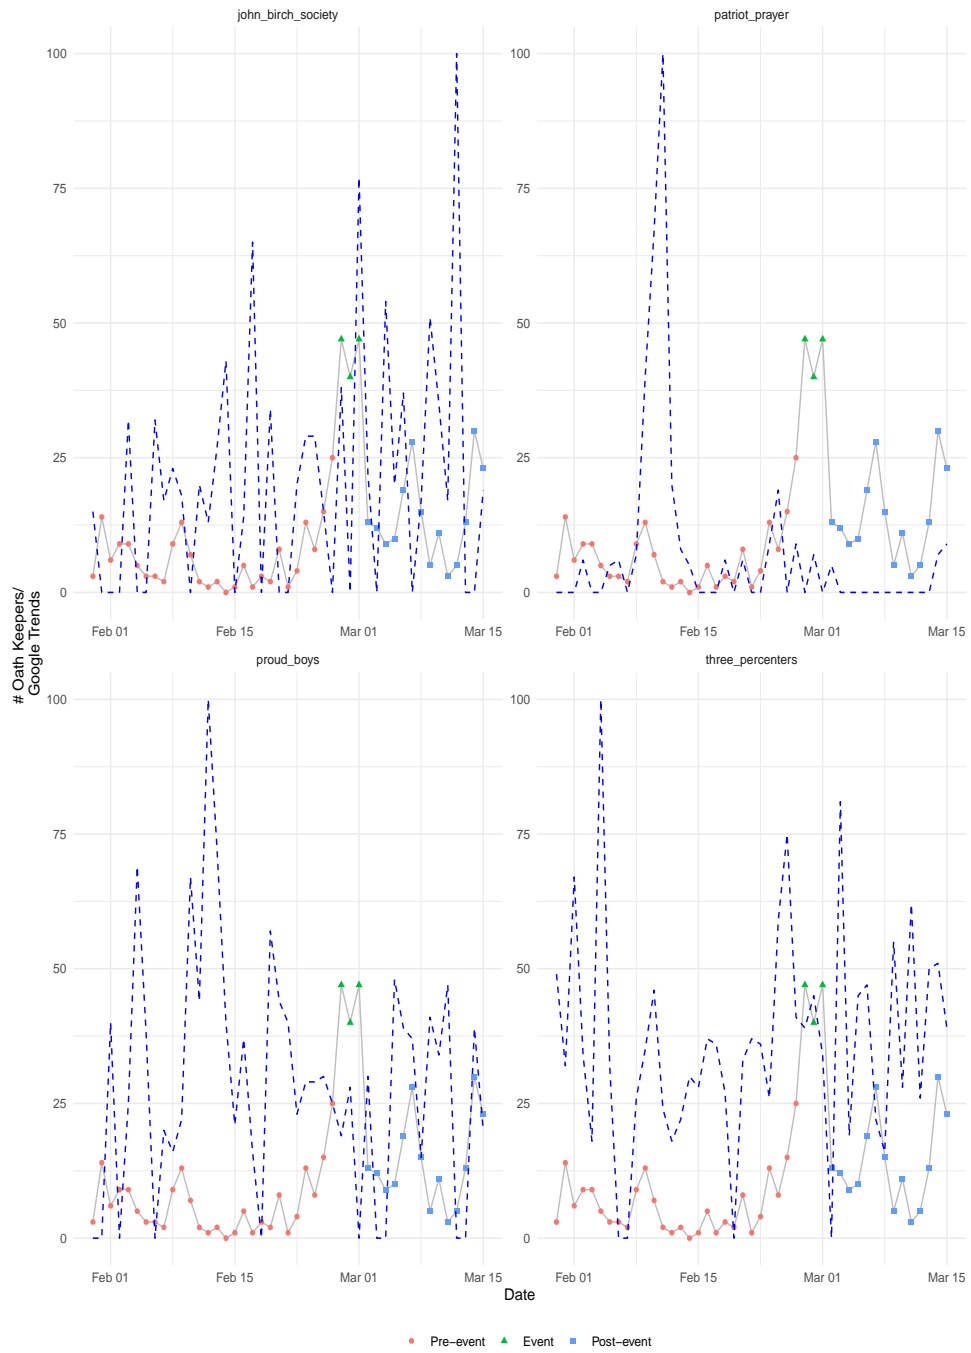


Figure A.5: Number Oath Keepers per day (solid black) versus Google trends for given group (dashed blue) during the Flash sale.



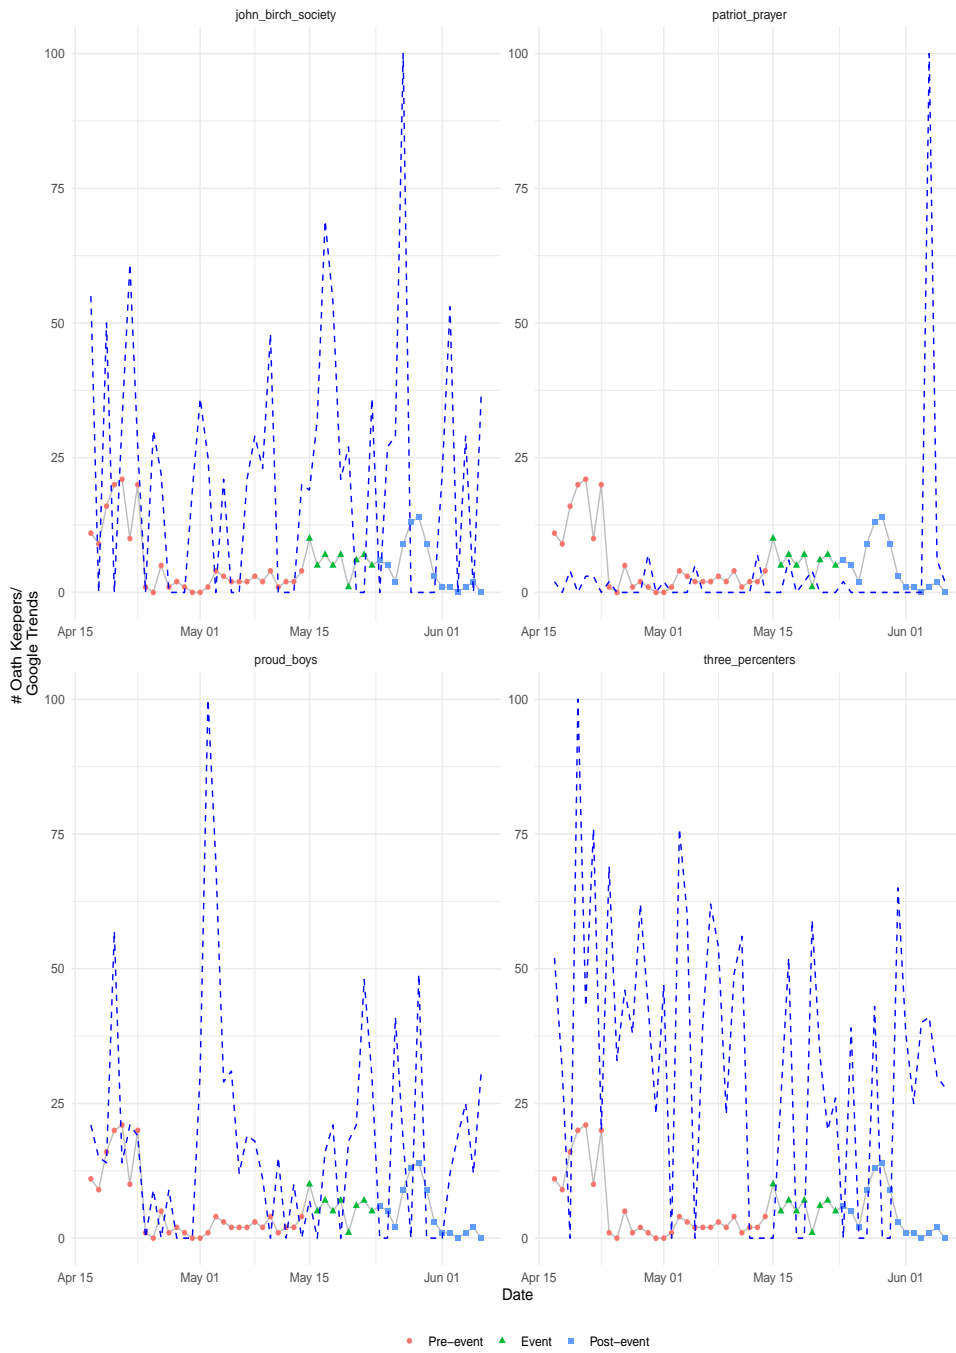


Figure A.6: Number Oath Keepers per day (solid black) versus Google trends for given group (dashed blue) during the Memorial Day sale.

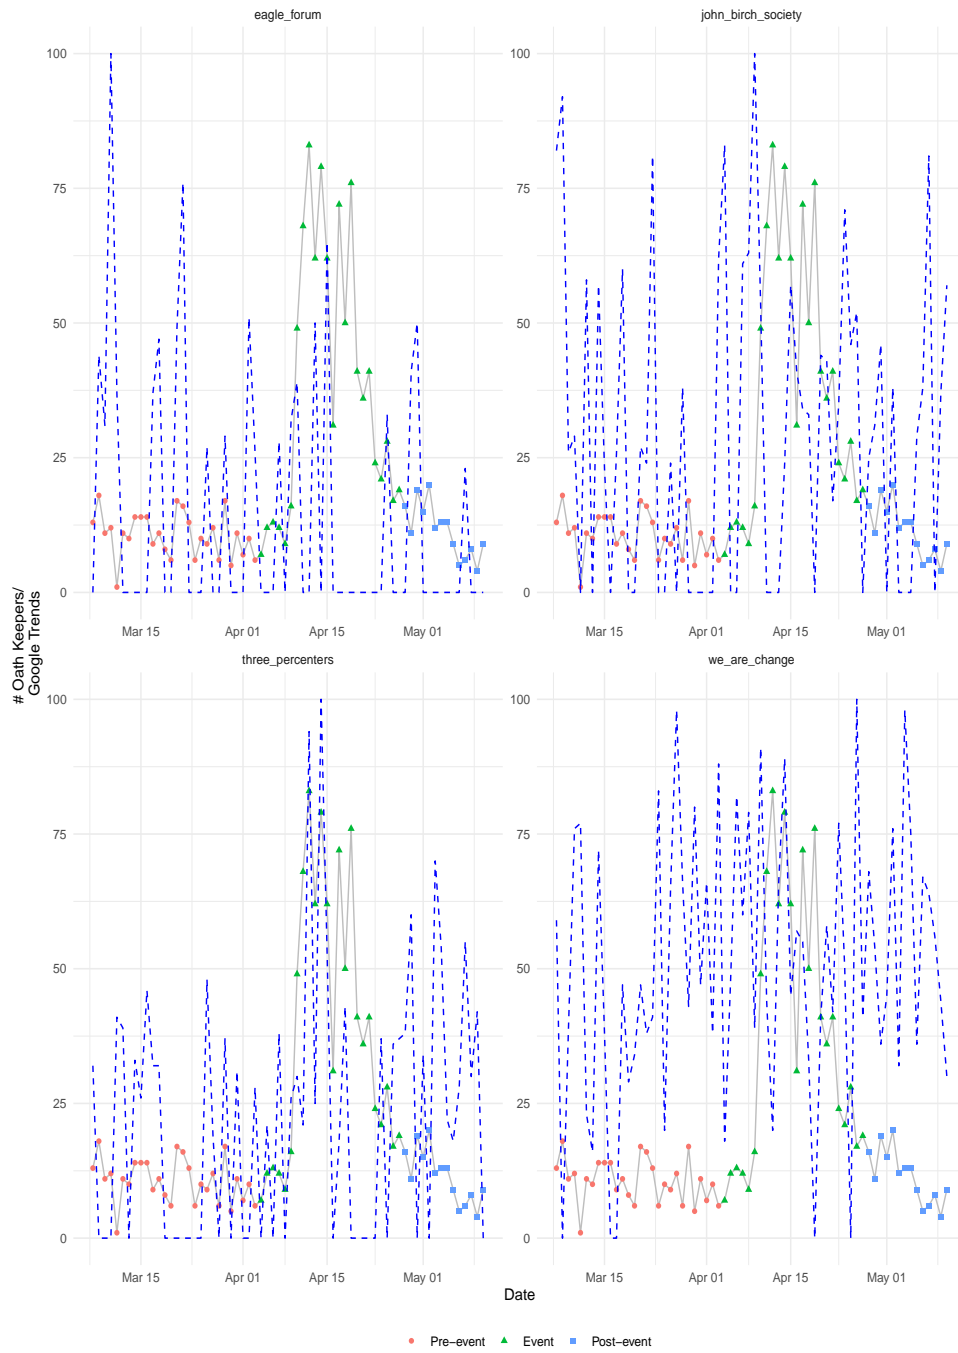


Figure A.7: Number Oath Keepers per day (solid black) versus Google trends for given group (dashed blue) during the Bundy Ranch callout event.

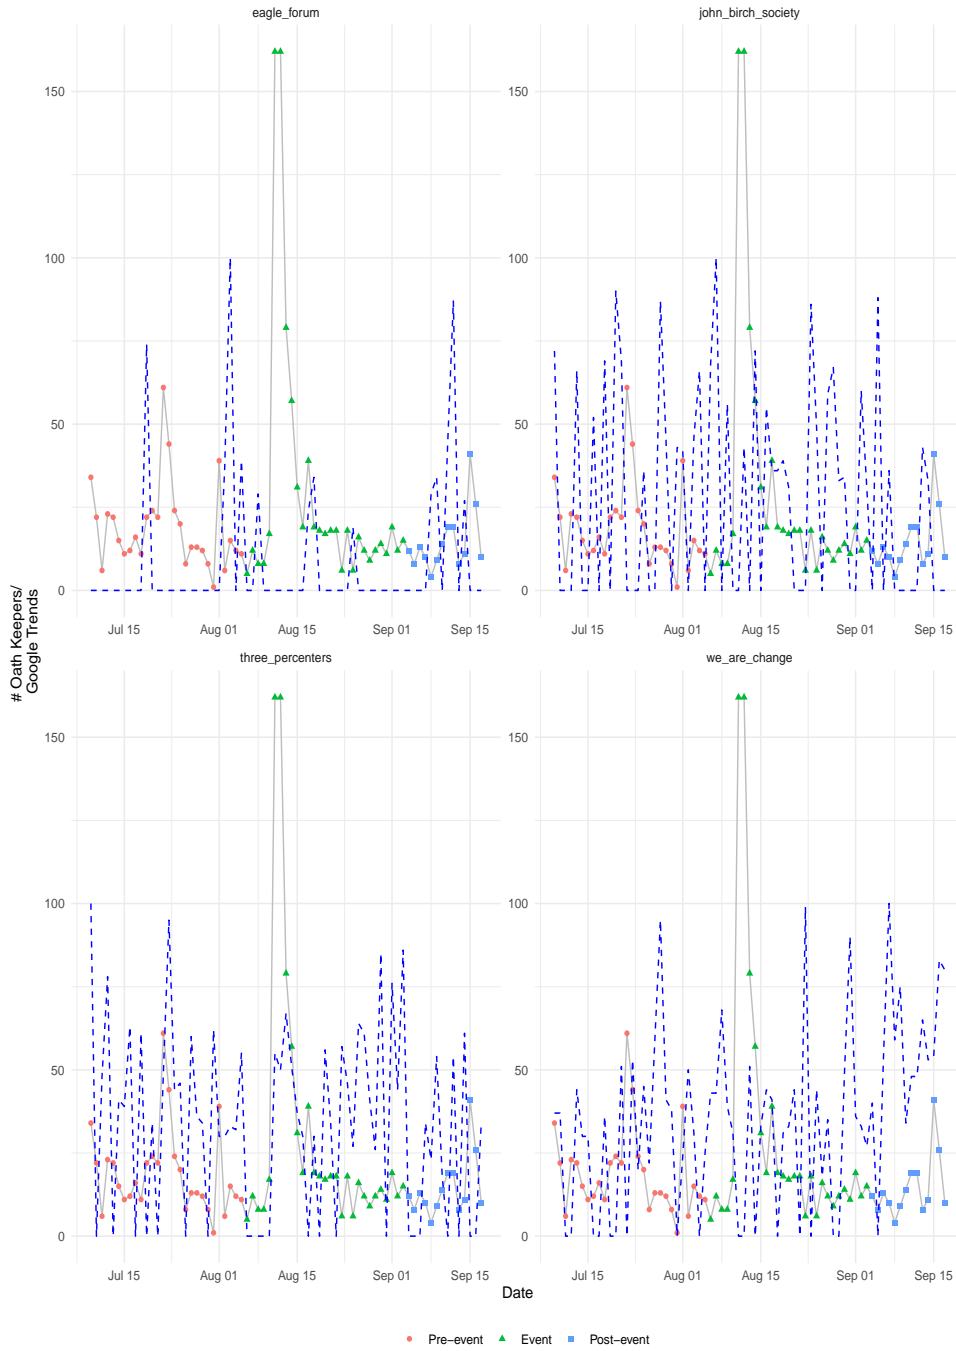


Figure A.8: Number Oath Keepers per day (solid black) versus Google trends for given group (dashed blue) during the Big Sky callout event.




Figure A.9: Number Oath Keepers per day (solid black) versus Google trends for given group (dashed blue) during the NASCAR advertising.

## **A.3 Counterfactual Robustness**

### **A.3.1 Brodersen et al. (2015))**

Table A.1: Average effect of tactics on Oath Keepers' recruitment using different pre-treatment windows.

	Two Weeks Pre-treatment	Three Weeks Pre-treatment	Four Weeks Pre-treatment
<b>Membership Discounts</b>			
Veteran's Day Discount 2014	22.88 [21.64, 24]	23.25 [22.07, 24.42]	23.5 [22.3, 24.8]
Constitution Day Discount 2017	-0.14 [-3.55, 3.16]	1.76 [-0.93, 5.47]	4.57 [1.35, 6.08]
Christmas and New Years Discount 2017	2.33 [1.81, 2.84]	2.24 [1.73, 2.76]	2.48 [1.99, 3]
Memorial Day Discount 2018	3.62 [3.02, 4.23]	3.9 [3.1, 4.6]	0.1 [-3.2, 3.73]
Flash Discount 2018	38.47 [32.95, 44.3]	37.44 [32.18, 42.88]	38.57 [33.34, 43.58]
<b>Armed Callout Events</b>			
Bundy Ranch Callout 2014	28.21 [26.15, 30.17]	28.14 [26.51, 29.65]	27.99 [26.31, 29.62]
Big Sky Callout 2015	12.9 [8.14, 17.51]	10.31 [4.88, 15.55]	9.91 [5.11, 14.45]
<b>Traditional Advertisement</b>			
NASCAR Sponsorship 2013	-8.11 [-11.82, -4.43]	-9.56 [-11.77, -7.59]	-8.28 [-10.53, -6.16]

*Note:*

The effect is the average number of new Oath Keepers per day due to the tacting during the tactic. Brackets are 95% credibility intervals.

Table A.2: Effect of sales on Oath Keepers' recruitment using all membership types.

	Relative Effect (%)	Average Effect	Cumulative Effect
<b>Membership Sales</b>			
Veteran's Day Discount 2014	442.11	27.22	299.45
	[352.07, 560.95]	[26.05, 28.39]	[286.6, 312.32]
Constitution Day Discount 2017	11.05	0.47	28.5
	[-24.38, 72.4]	[-2.93, 3.82]	[-176.02, 229.29]
Christmas and New Years Discount 2017	179.96	3.04	167.32
	[112.14, 281.35]	[2.54, 3.54]	[139.56, 194.77]
Memorial Day Discount 2018	61.16	2.08	18.72
	[15.18, 151.19]	[0.78, 3.54]	[6.99, 31.9]
Flash Discount 2018	1283.83	47.81	143.42
	[294.15, 4507.2]	[41.64, 53.47]	[124.91, 160.42]
<b>Armed Callout Events</b>			
Bundy Ranch Callout 2014	261.29	32.96	790.94
	[203.77, 344.47]	[30.69, 35.46]	[736.54, 850.97]
Big Sky Callout 2015	82.88	12.67	367.29
	[35.17, 154.69]	[7.59, 17.72]	[220.11, 513.83]
<b>Traditional Advertisement</b>			
NASCAR Sponsorship 2013	-50.05	-10.2	-724.45
	[-58.97, -35.9]	[-14.29, -5.57]	[-1014.7, -395.46]

*Note:*

The relative effect is in terms of percent change. The average effect is the average number of new Oath Keepers per day due to the sale during the sale while the cumulative effect is the total number of new Oath Keepers due to the sale during the sale. Brackets are 95% credibility intervals. The placebo test for the Flash Sale suggests the estimated counterfactual did not accurately approximate the underlying data generating process. The results for the Flash Sale should be interpreted as suggested, not causal.

Table A.3: Effect of sales on Oath Keepers' recruitment with four week pre-treatment window

	Relative Effect (%)	Average Effect	Cumulative Effect
<b>Membership Sales</b>			
Veteran's Day Discount 2014	505.49	23.5	258.53
	[373.53, 713.73]	[22.3, 24.8]	[245.32, 272.78]
Constitution Day Discount 2017	143.05	4.57	273.97
	[18.84, 253.18]	[1.35, 6.08]	[80.7, 364.88]
Christmas and New Years Discount 2017	162.14	2.48	136.54
	[94.44, 273.78]	[1.99, 3]	[109.28, 164.8]
Memorial Day Discount 2018	15.76	0.1	0.93
	[-35.46, 170.72]	[-3.2, 3.73]	[-28.81, 33.61]
Flash Discount 2018	667.77	38.57	115.71
	[265.69, 2838.67]	[33.34, 43.58]	[100.03, 130.74]
<b>Armed Callout Events</b>			
Bundy Ranch Callout 2014	263.69	27.99	671.73
	[212.97, 327.57]	[26.31, 29.62]	[631.49, 710.96]
Big Sky Callout 2015	54.49	9.91	287.26
	[21.43, 99.8]	[5.11, 14.45]	[148.08, 419.07]
<b>Traditional Advertisement</b>			
NASCAR Sponsorship 2013	-50.66	-8.28	-588.2
	[-56.83, -43.5]	[-10.53, -6.16]	[-747.75, -437.28]

*Note:*

The relative effect is in terms of percent change. The average effect is the average number of new Oath Keepers per day due to the sale during the sale while the cumulative effect is the total number of new Oath Keepers due to the sale during the sale. Brackets are 95% credibility intervals. The placebo test for the Flash Sale suggests the estimated counterfactual did not accurately approximate the underlying data generating process. The results for the Flash Sale should be interpreted as suggested, not causal.



### A.3.2 (Klinenberg (2022)2022)

#### Two week pre-treatment coefficients

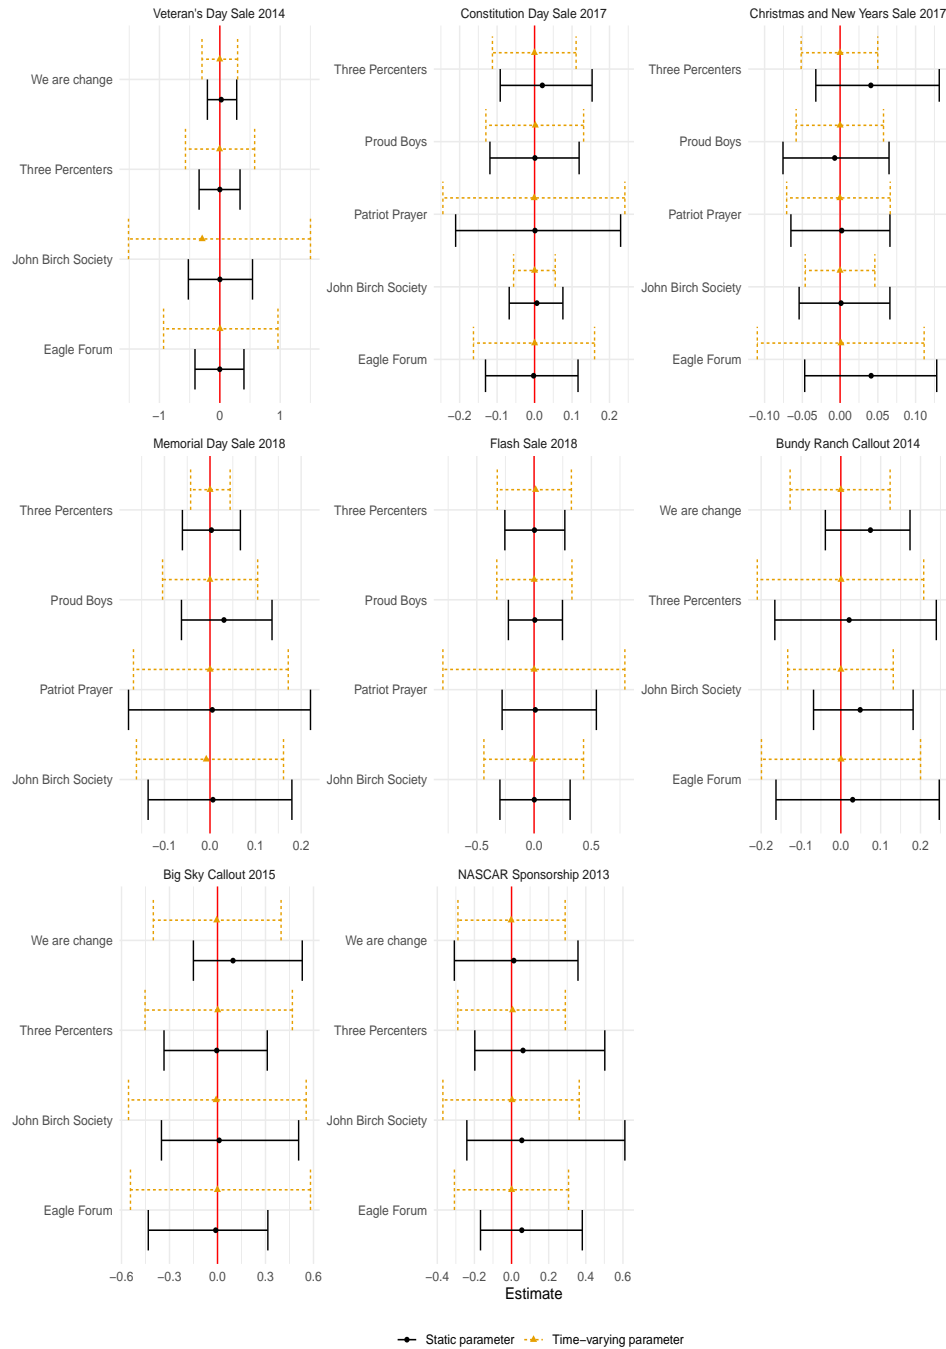


Figure A.10: Time-invariant and time-varying parameters estimated using Klinenberg (2022) on a two week pre-treatment period with 95% credibility intervals.

Three week pre-treatment coefficients

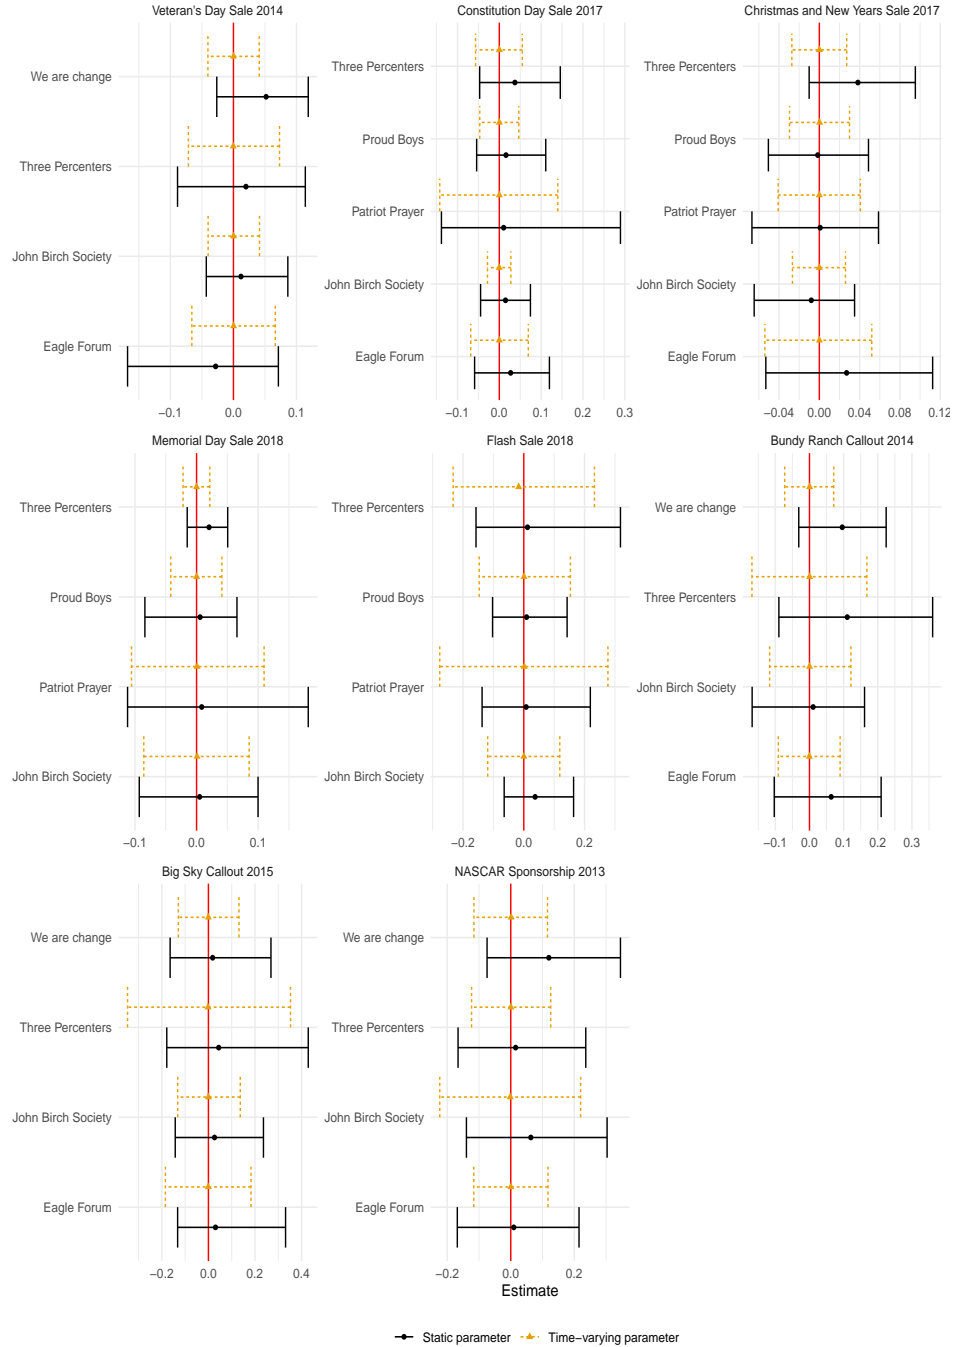


Figure A.11: Time-invariant and time-varying parameters estimated using Klinenberg (2022) on a three week pre-treatment period.

Four week pre-treatment coefficients

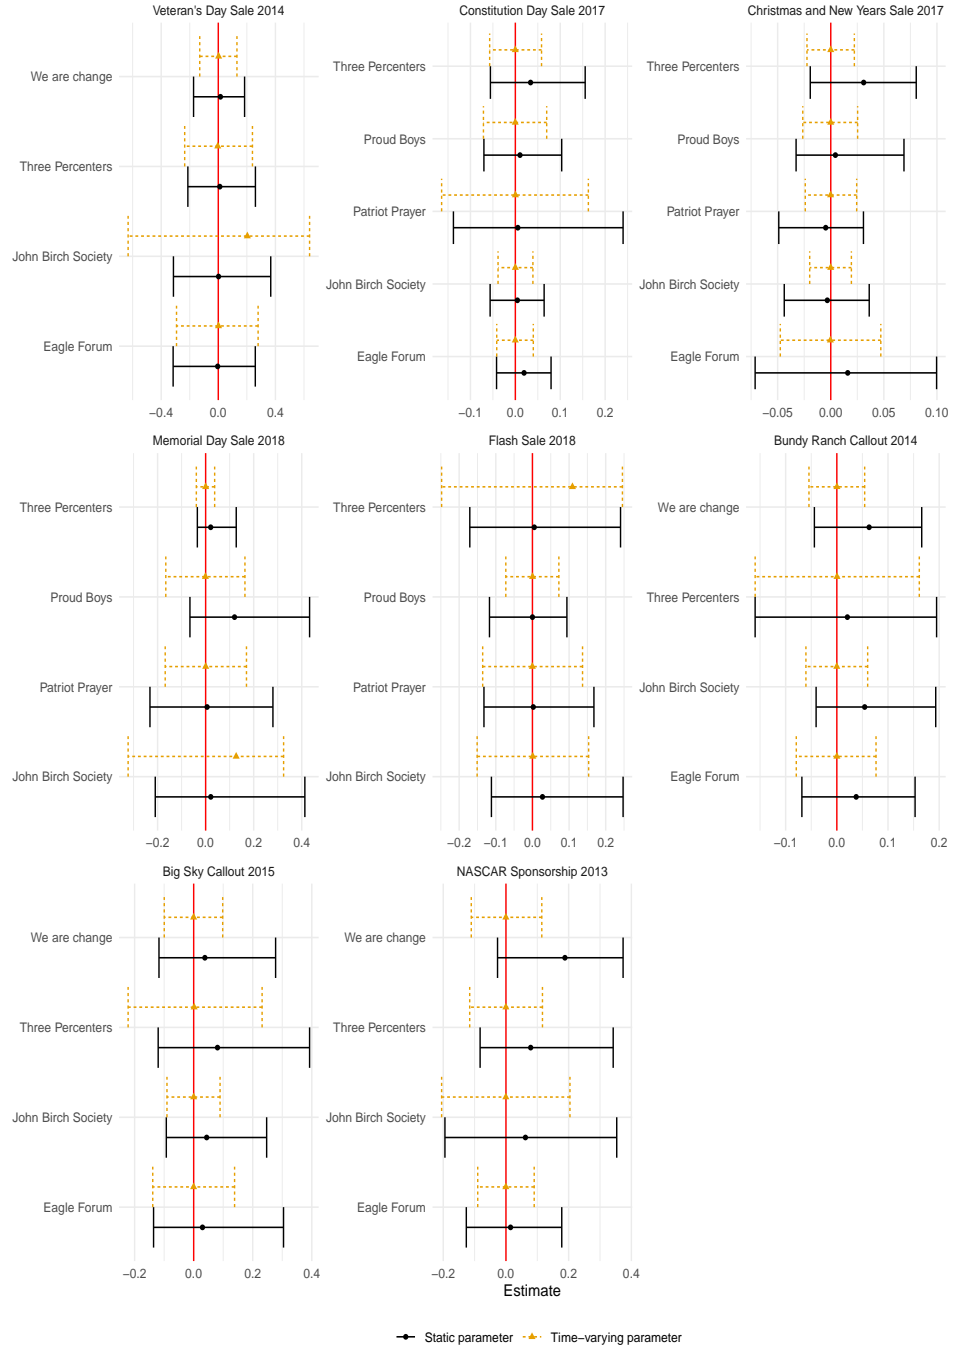


Figure A.12: Time-invariant and time-varying parameters estimated using Klinenberg (2022) on a three week pre-treatment period.

### **A.3.3 Horse Race**

In the paper, I compare the mean squared forecast errors on the week before each tactic begins. The model is fitted to days 8-14 prior to the tactic. I now fit the model on days 8-21 and days 8-28 prior to the tactic beginning, then calculate the mean squared forecast error on days 1-7 prior to the tactic. The results do not change.

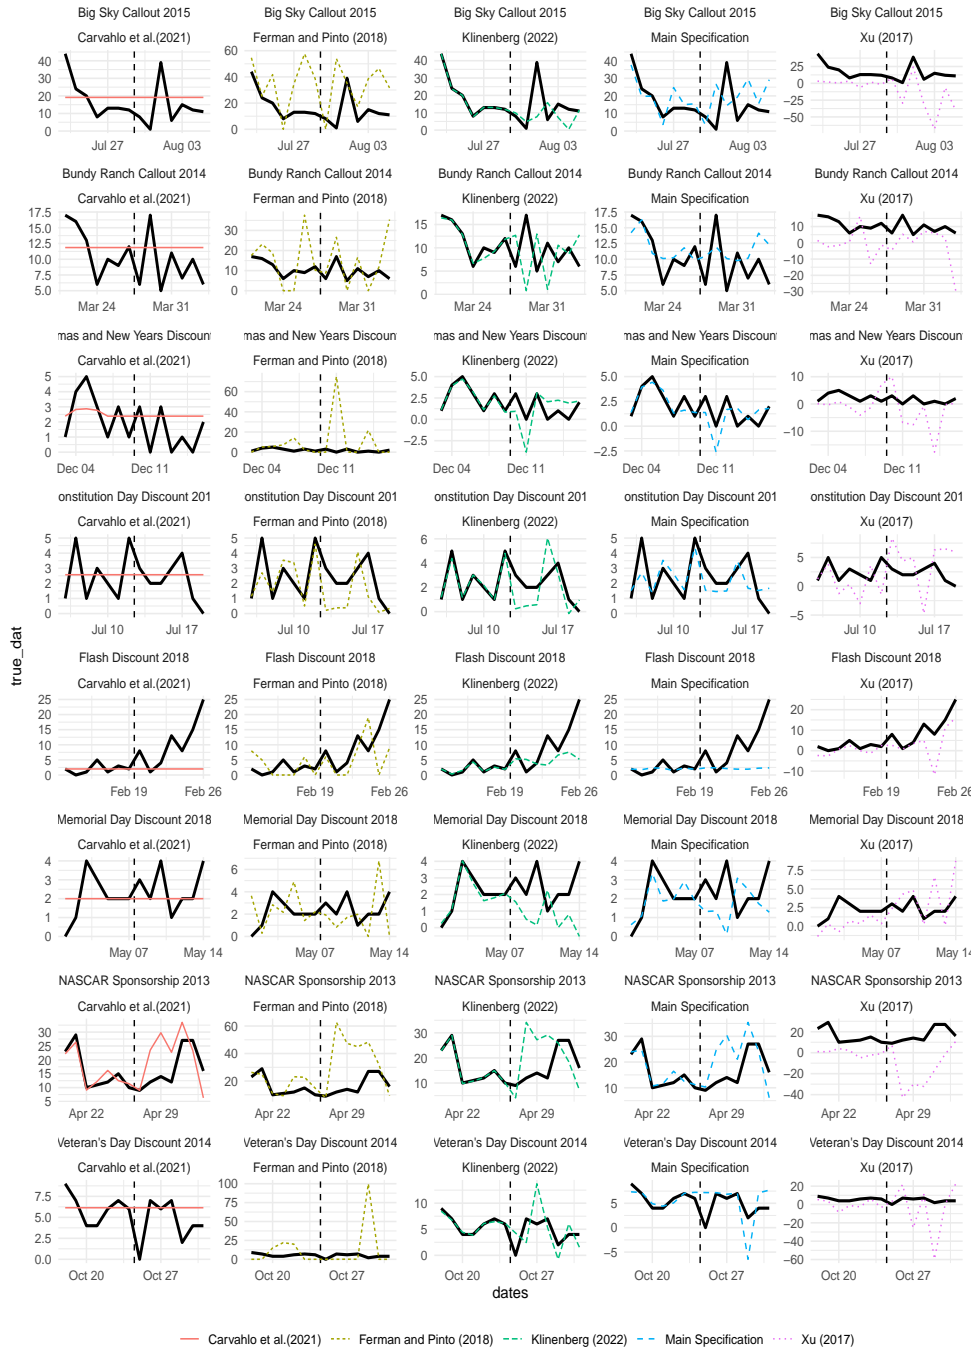


Figure A.13: Plotted raw data (black) and estimated counterfactual (other line). The dashed vertical line is the beginning of the placebo test period. Models are fit to days 8-14.

Table A.4: Mean squared forecast error of alternative models using first seven days prior to tactic. Counterfactual estimates are fitted to days 8-21 prior to a tactic.

	Main Specification	Ferman and Pinto (2021)	Carvalho et al. (2018)	Xu (2017)	Klinenberg (2022)
<b>Membership Sales</b>					
Veteran's Day Discount 2014	5.7	1378	6.9	1330.5	16.0
Constitution Day Discount 2017	2.4	3	2.9	4.2	1.9
Christmas and New Years Discount 2017	1.9	292	2.5	52.0	2.1
Memorial Day Discount 2018	2.2	7	1.8	5.9	6.1
Flash Discount 2018	91.5	264	102.5	220.0	80.1
<b>Armed Callout Events</b>					
Bundy Ranch Callout 2014	21.6	149	21.5	267.4	43.7
Big Sky Callout 2015	207.4	868	200.0	685.3	253.6
<b>Traditional Advertisement</b>					
NASCAR Sponsorship 2013	36.8	673	46.8	1049.0	134.9

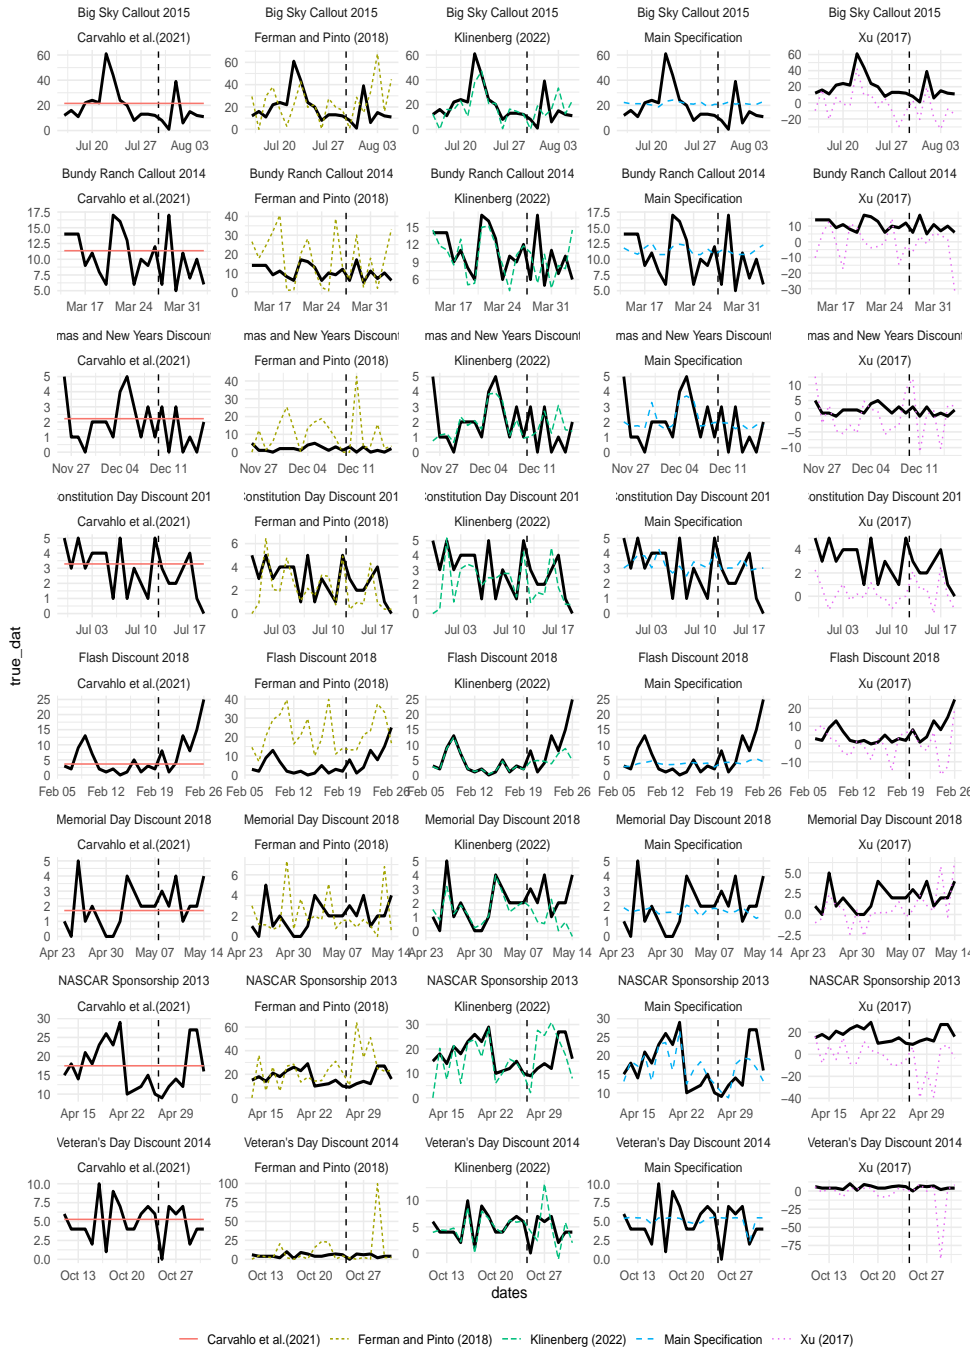


Figure A.14: Plotted raw data (black) and estimated counterfactual (other line). The dashed vertical line is the beginning of the placebo test period. Models are fit to days 8-21.

Table A.5: Mean squared forecast error of alternative models using first seven days prior to tactic. Counterfactual estimates are fitted to days 8-28 prior to a tactic.

	Main Specification	Ferman and Pinto (2021)	Carvalho et al. (2018)	Xu (2017)	Klinenberg (2022)
<b>Membership Sales</b>					
Veteran's Day Discount 2014	10.1	1339.2	6.4	1417	12.9
Constitution Day Discount 2017	2.5	2.8	2.5	3	1.7
Christmas and New Years Discount 2017	1.6	272.4	2.0	108	2.4
Memorial Day Discount 2018	12.9	9.2	14.6	30	8.9
Flash Discount 2018	86.3	135.8	88.3	72	78.6
<b>Armed Callout Events</b>					
Bundy Ranch Callout 2014	25.2	140.0	20.7	212	44.9
Big Sky Callout 2015	190.7	779.3	186.3	707	260.9
<b>Traditional Advertisement</b>					
NASCAR Sponsorship 2013	51.8	722.4	47.1	1099	171.3



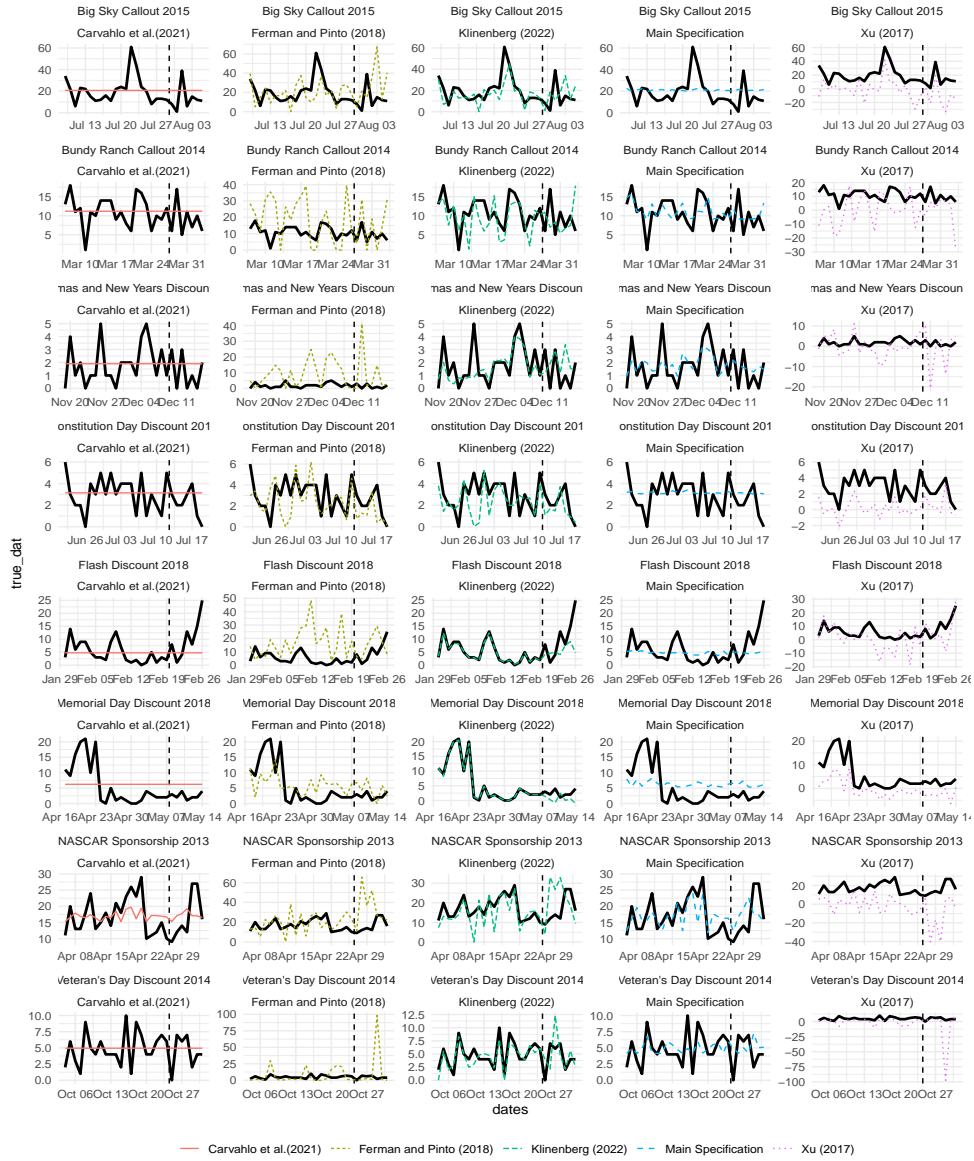


Figure A.15: Plotted raw data (black) and estimated counterfactual (other line). The dashed vertical line is the beginning of the placebo test period. Models are fit to days 8-28.

## A.4 County level economic inequality analysis

### A.4.1 Main specification event study

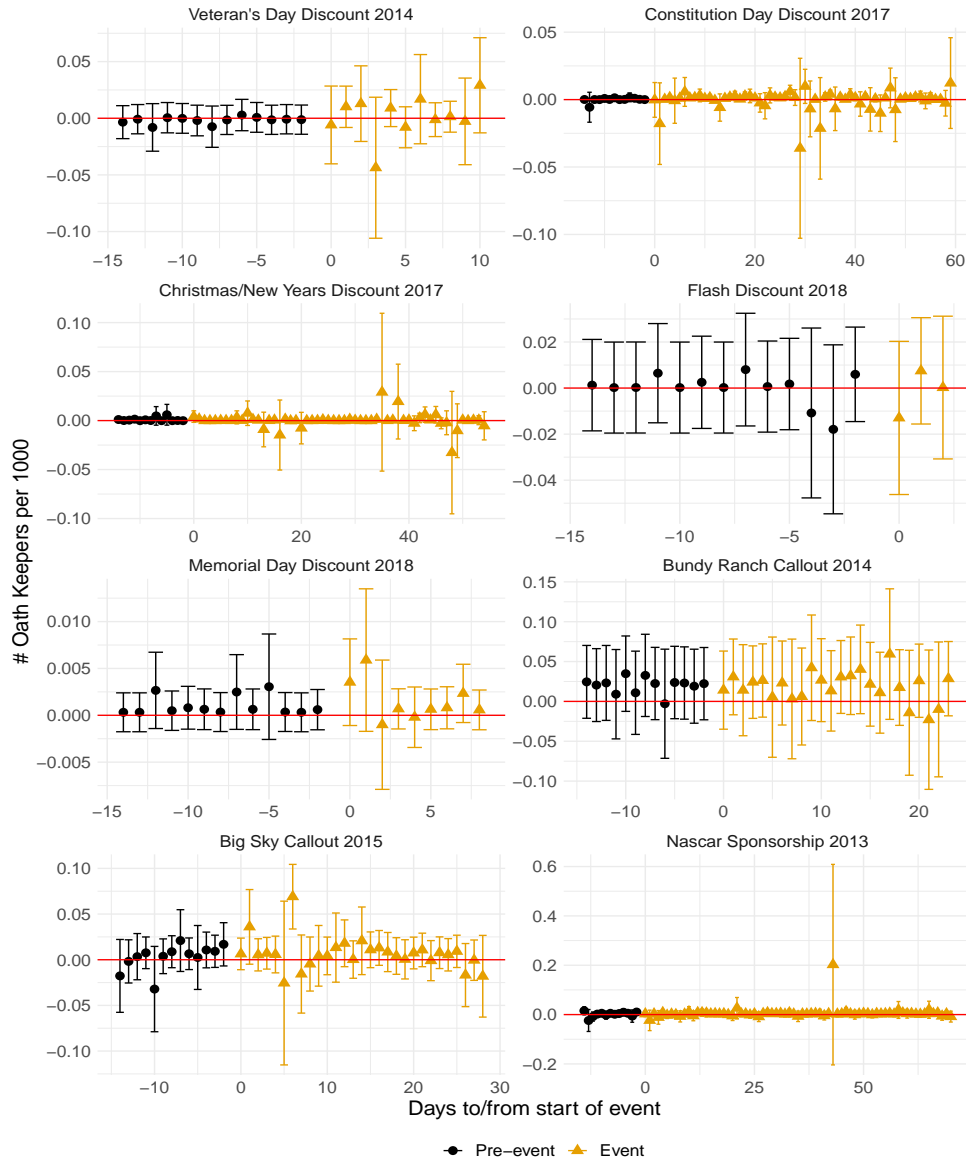


Figure A.16: Event study between the top the 25% richest counties and bottom the 25% poorest counties based on median household income.

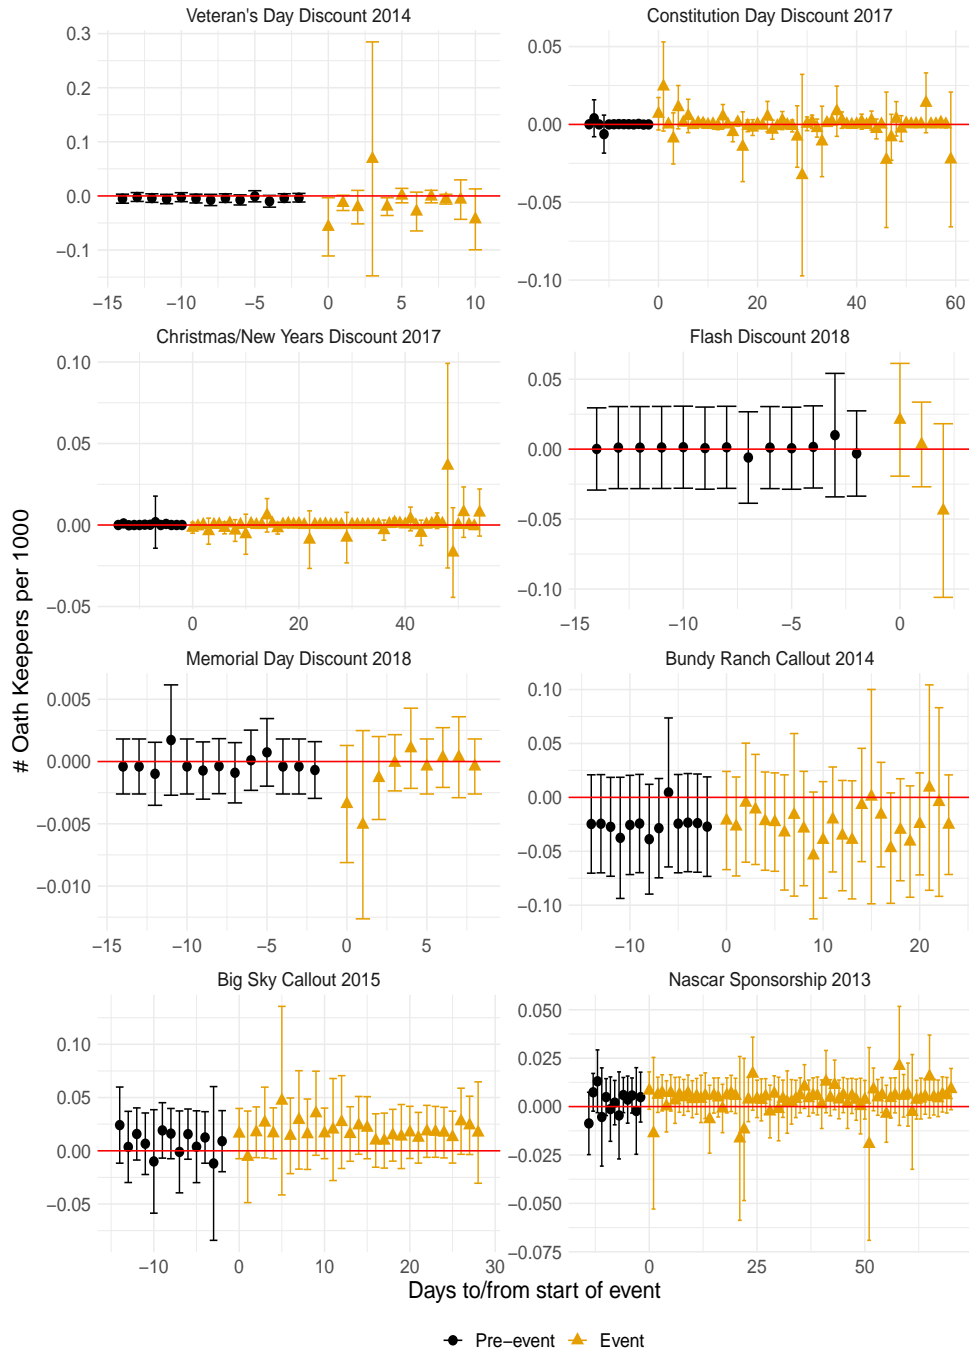


Figure A.17: Event study between the top the 25% richest counties and bottom the 25% poorest counties based on economic inequality.

### A.4.2 Alternative specifications of differential effects

Table A.6: Oath Keepers per 1,000 on economic conditions unweighted.

	Veteran's Day Discount 2014	Constitution Day Discount 2017	Christmas/New Years Discount 2017	Memorial Day Discount 2018	Flash Discount 2018	Bundy Ranch 2014	Big Sky 2015	Nascar Sponsorship 2013
<i>Panel A: Median Household Income</i>								
log Median Income	-0.00006 (0.001)	-0.0009 (0.0008)	0.001** (0.0005)	-0.0002 (0.002)	0.0005** (0.0003)	-0.002 (0.003)	-0.001 (0.004)	0.0009 (0.004)
log Median Income X Event	-0.001 (0.007)	0.001 (0.001)	-0.001 (0.001)	-0.002 (0.012)	-0.001 (0.001)	-0.007 (0.009)	-0.002 (0.006)	0.003 (0.005)
Std.Errors	by: county	by: county	by: county	by: county	by: county	by: county	by: county	by: county
Number of Clusters	3136	3136	3136	3136	3136	3136	3136	3136
Outcome Average	0.007579	0.002805	0.001969	0.005457	0.001126	0.015216	0.010886	0.003975
FE: Day	X	X	X	X	X	X	X	X
<i>Panel B: Income Inequality</i>								
log Income Inequality	0.0002 (0.0007)	-0.0005 (0.0007)	0.0003 (0.0006)	0.0001 (0.003)	-0.0001 (0.0003)	-0.002 (0.004)	-0.011** (0.004)	-0.005 (0.003)
log Income Inequality X Event	-0.011 (0.013)	-0.0004 (0.001)	-0.0001 (0.001)	-0.017 (0.011)	-0.0004 (0.001)	-0.002 (0.007)	0.012** (0.006)	0.002 (0.003)
Number of Clusters	3132	3132	3132	3132	3132	3132	3132	3132
Outcome Average	0.007588	0.002808	0.001972	0.005464	0.001127	0.015235	0.0109	0.00398
Std.Errors	by: county	by: county	by: county	by: county	by: county	by: county	by: county	by: county
FE: Day	X	X	X	X	X	X	X	X

*Note:*\* p  $\leq$  0.1, \*\* p  $\leq$  0.05, \*\*\* p  $\leq$  0.01

Outcome average is the average number of Oath Keepers per 1000 of both top and lower quantiles before and after the callout event. Median household income and income inequality is measured continuously and all counties are used in the analysis. Standard errors are clustered at the county level.

Table A.7: Oath Keepers per 1,000 on economic conditions weighted by county population.

	Veteran's Day Discount 2014	Constitution Day Discount 2017	Christmas/New Years Discount 2017	Memorial Day Discount 2018	Flash Discount 2018	Bundy Ranch 2014	Big Sky 2015	Nascar Sponsorship 2013
<i>Panel A: Median Household Income</i>								
log Median Income	-0.00006 (0.001)	-0.0009 (0.0008)	0.001** (0.0005)	-0.0002 (0.002)	0.0005** (0.0003)	-0.002 (0.003)	-0.001 (0.004)	0.0009 (0.004)
log Median Income X Event	-0.001 (0.007)	0.001 (0.001)	-0.001 (0.001)	-0.002 (0.012)	-0.001 (0.001)	-0.007 (0.009)	-0.002 (0.006)	0.003 (0.005)
Std.Errors	by: county	by: county	by: county	by: county	by: county	by: county	by: county	by: county
Number of Clusters	3136	3136	3136	3136	3136	3136	3136	3136
Outcome Average	0.007579	0.002805	0.001969	0.005457	0.001126	0.015216	0.010886	0.003975
FE: Day	X	X	X	X	X	X	X	X
<i>Panel B: Income Inequality</i>								
log Income Inequality	0.0002 (0.0007)	-0.0005 (0.0007)	0.0003 (0.0006)	0.0001 (0.003)	-0.0001 (0.0003)	-0.002 (0.004)	-0.011** (0.004)	-0.005 (0.003)
log Income Inequality X Event	-0.011 (0.013)	-0.0004 (0.001)	-0.0001 (0.001)	-0.017 (0.011)	-0.0004 (0.001)	-0.002 (0.007)	0.012** (0.006)	0.002 (0.003)
Number of Clusters	3132	3132	3132	3132	3132	3132	3132	3132
Outcome Average	0.007588	0.002808	0.001972	0.005464	0.001127	0.015235	0.0109	0.00398
Std.Errors	by: county	by: county	by: county	by: county	by: county	by: county	by: county	by: county
FE: Day	X	X	X	X	X	X	X	X

*Note:*\* p  $\leq$  0.1, \*\* p  $\leq$  0.05, \*\*\* p  $\leq$  0.01

Outcome average is the average number of Oath Keepers per 1000 of both top and lower quantiles before and after the callout event. Median household income and income inequality is measured continuously and all counties are used in the analysis. Standard errors are clustered at the county level.

Table A.8: Two-way fixed effects analysis between counties in the top and bottom quartiles limiting to counties that ever received an Oath Keeper.

	Veteran's Day Discount 2014	Constitution Day Discount 2017	Christmas/New Years Discount 2017	Memorial Day Discount 2018	Flash Discount 2018	Bundy Ranch 2014	Big Sky 2015	Nascar Sponsorship 2013
<i>Panel A: Median household income</i>								
I(Top quartile) X I(During event)	0.001 (0.006)	-0.002 (0.001)	-0.002 (0.002)	-0.005 (0.009)	0.0005 (0.0008)	-0.003 (0.008)	0.004 (0.005)	0.005 (0.004)
Std.Errors	by: county	by: county	by: county	by: county	by: county	by: county	by: county	by: county
FE: Day	X	X	X	X	X	X	X	X
FE: County	X	X	X	X	X	X	X	X
Number of Clusters	1258	1276	1276	1261	1261	1258	1260	1254
Outcome Average	0.009336	0.003453	0.002424	0.006718	0.001386	0.018743	0.013402	0.004898
<i>Panel B: Income inequality</i>								
I(Top quartile) X I(During event)	-0.009 (0.014)	-0.001 (0.001)	-0.00004 (0.001)	-0.010 (0.014)	-0.001 (0.0007)	0.0001 (0.007)	0.015** (0.006)	0.003 (0.002)
Std.Errors	by: county	by: county	by: county	by: county	by: county	by: county	by: county	by: county
FE: Day	X	X	X	X	X	X	X	X
FE: County	X	X	X	X	X	X	X	X
Number of Clusters	1241	1216	1216	1208	1208	1241	1229	1230
Outcome Average	0.009336	0.003453	0.002424	0.006718	0.001386	0.018743	0.013402	0.004898

*Note:*

\* p < 0.1, \*\* p < 0.05, \*\*\* p < 0.01

Outcome average is the average number of Oath Keepers per 1000 of both top and lower quartiles before and after the callout event. I(Top Quartile) is an indicator if a county is in the top quartile. The reference group is the bottom quartile. The middle quartiles are dropped from the analysis. Standard errors are clustered at the county level.

Table A.9: Two-way fixed effects analysis between counties in the top and bottom quartiles limiting to counties that ever received an Oath Keeper weighted by population.

	Veteran's Day Discount 2014	Constitution Day Discount 2017	Christmas/New Years Discount 2017	Memorial Day Discount 2018	Flash Discount 2018	Bundy Ranch 2014	Big Sky 2015	Nascar Sponsorship 2013
<i>Panel A: Median household income</i>								
I(Top quartile) X I(During event)	-0.005 (0.003)	-0.0009 (0.0009)	-0.0009** (0.0004)	-0.007 (0.006)	0.0003 (0.0009)	-0.0009 (0.003)	0.0009 (0.003)	-0.001 (0.001)
Std.Errors	by: county	by: county	by: county	by: county	by: county	by: county	by: county	by: county
FE: Day	X	X	X	X	X	X	X	X
FE: County	X	X	X	X	X	X	X	X
Number of Clusters	1258	1276	1276	1261	1261	1258	1260	1254
Outcome Average	0.009336	0.003453	0.002424	0.006718	0.001386	0.018743	0.013402	0.004898
<i>Panel B: Income inequality</i>								
I(Top quartile) X I(During event)	-0.009*** (0.002)	-0.001** (0.0006)	-0.0002 (0.0004)	-0.005 (0.004)	0.00003 (0.001)	-0.006*** (0.002)	-0.0002 (0.002)	0.0009 (0.001)
Std.Errors	by: county	by: county	by: county	by: county	by: county	by: county	by: county	by: county
FE: Day	X	X	X	X	X	X	X	X
FE: County	X	X	X	X	X	X	X	X
Number of Clusters	1241	1216	1216	1208	1208	1241	1229	1230
Outcome Average	0.009336	0.003453	0.002424	0.006718	0.001386	0.018743	0.013402	0.004898

*Note:*

\* p  $\leq$  0.1, \*\* p  $\leq$  0.05, \*\*\* p  $\leq$  0.01

Outcome average is the average number of Oath Keepers per 1000 of both top and lower quartiles before and after the callout event. I(Top Quartile) is an indicator if a county is in the top quartile. The reference group is the bottom quartile. The middle quartiles are dropped from the analysis. Standard errors are clustered at the county level.

## Appendix B

### Appendix for “Does Deplatforming Work?”



## B.1 Graphs

## B.2 Theoretical Effects of Deplatforming under Constant Elasticity of Substitution (CES) Utility

Suppose the consumer follows a constant elasticity of substitution (CES) utility function:

$$\begin{aligned} \max_{d_1, d_2, x} U(d_1, d_2, x) &= (d_1^\rho + \alpha_2 d_2^\rho + \alpha_x x^\rho)^{\frac{1}{\rho}} \\ \text{sbj. to } d_1 + d_2 + x &= I \\ d_1, d_2, x, I &\geq 0 \\ \alpha_2, \alpha_x &\geq 0 \end{aligned} \tag{B.1}$$

Without loss of generality, I assume the content creator has no costs. The content creator’s profit is  $\pi = d_1 + d_2$ . Site 1 can ban the content creator. When a content creator is banned,  $d_1 = 0$  because the content creator is no longer allowed to use the site. The viewer is now maximizing over  $d_2$  and  $x$  only, meaning the new equilibrium will be a corner solution with  $d_1 = 0$ . How  $\pi$  and  $d_2$  change when a content creator is banned depends on the model parameters. I provide theoretical predictions for  $\rho = 1$  and  $\rho = 0$ .

### Case 1: Perfect Substitutes ( $\rho = 1$ )

If  $\rho = 1$ , the consumer’s utility function simplifies to one of perfect substitutes:  $U(d_1, d_2, x) = d_1 + \alpha_2 d_2 + \alpha_x x$ . When the content creator is not banned, the consumer spends all their income on donation to Site 1 if  $1 > \alpha_2, \alpha_x$ . However, if the content creator is banned, the consumer will substitute their spending to donations on Site 2 or

outside spending.

If  $\alpha_2 > 1$  and or  $\alpha_x > 1$ , then there is no change to the consumer’s optimal bundle of goods. The optimal bundle already sets  $d_1^* = 0$ .

If  $1 > \alpha_2 > \alpha_x$ , then the consumer will donate the same amount as before, namely  $I$ . This means the consumer completely substitutes donations on Site 1 to Site 2, supporting theories of displacement (Keatinge et al. (2019)). Such perverse effects have been well documented in the economic literature and are referred to as Peltzman Effects (Peltzman (1975)). Examples of similar effects include social media bans causing more hateful activity on alternative sites (Mitts (2021); Ali et al. (2021)) and increases in anti-vaccine content (Mitts et al. (2022)). Potential reasons for such counter-intuitive effects include curiosity about the banned content. Jansen and Martin (Jansen & Martin (2015)) refer to this phenomena as the “Streisand Effect” and document it in other contexts.

If  $1 > \alpha_x > \alpha_2$ , then the consumer will completely stop donating to the content creator. Observing the content creator being banned could potentially lead the consumer to discontinue support for the content creator. This suggests deplatforming policies produce the intended effects: overall support for the content creator goes down across both platforms when they are banned.

### Case 2: Cobb-Douglass Utility ( $\rho \rightarrow 1$ )

As  $\rho \rightarrow 1$ , the log-linear CES utility function represents the same preferences as the generalized log-linear Cobb-Douglas utility function:  $U(d_1, d_2, x) = d_1^{\frac{1}{1+\alpha_2+\alpha_x}} d_2^{\frac{\alpha_2}{1+\alpha_2+\alpha_x}} x^{\frac{\alpha_x}{1+\alpha_2+\alpha_x}}$ .

When the content creator is not banned, the equilibrium values are

$$(d_1^*, d_2^*, x^*) = \left( \frac{1}{1+\alpha_2+\alpha_x} I, \frac{\alpha_2}{1+\alpha_2+\alpha_x} I, \frac{\alpha_x}{1+\alpha_2+\alpha_x} I \right), \text{ and the content creator earns:}$$

$$\pi(d_1^*) = d_1^* + d_2^* \tag{B.2}$$

$$= \frac{1 + \alpha_2}{1 + \alpha_2 + \alpha_x} I \tag{B.3}$$

When the content creator is banned, profit decrease to  $\pi(0) = I \frac{\alpha_2}{\alpha_2 + \alpha_x}$ , but the amount of donations on Site 2 increase to  $I \frac{\alpha_2}{\alpha_2 + \alpha_x}$ .  $d_1$  is being reallocated between Site 2 and outside activities based on  $\alpha_2$  and  $\alpha_x$ . From a policy perspective, this suggests deplatforming partially worked: while the content creator experiences an *overall* decrease in support, as measured in revenue, there is a partial recovery of support in increased donations on Site 2. The magnitude of the deplatforming effect depends on the substitutability between donations on Site 1 and other activities/content creators.

Deplatforming may work, backfire, or have partial success under different parameter values. In the following sections, I use reduced form estimation strategies to measure the overall change in profit ( $\pi$ ) and change in revenue on Bitchute ( $d_2$ ) after content creators are banned from YouTube.

### B.3 Data collection funnel chart




Figure B.1: Funnel chart of data collection process. The grey square is the final sample.



Table B.1: Word overlap in top 50 most common words in video titles. Banned channels refer to Bitchute channels linked to banned YouTube channels. Never-banned refers to Bitchute channels linked to never-banned YouTube channels. The ranking refers to the popularity of the word in video titles with 1 being the most popular word. Twenty-four of the top 50 words for Banned and Never-banned video titles overlapped.”

Words	Banned Channels Rank	Never-Banned Channels Rank
biden	14	2
coronavirus	29	28
covid	11	22
election	35	29
joe	40	27
just	31	31
live	15	12
man	36	50
media	22	5
new	5	6
news	13	17
now	34	14
people	18	20
police	16	34
says	19	25
speech	43	45
state	39	46
trump	3	1
truth	20	23
us	4	9
video	28	11
war	25	13
white	12	18
world	8	43

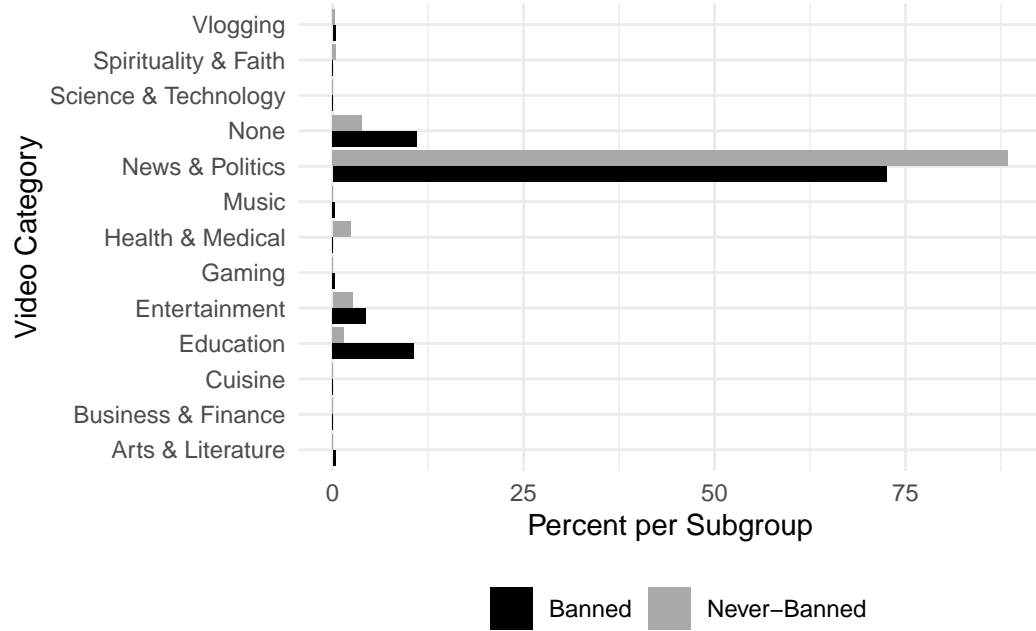


Figure B.3: Frequency of video topic breakdown. Topic category is chosen by the video’s creator and saved in the Bitchute metadata. The graph includes the top 10 categories for Bitchute channels affiliated with banned YouTube channels and the top 10 categories for Bitchute channels affiliated with never-banned YouTube channels. The top 10 categories account for almost 100% of the videos in the sample.



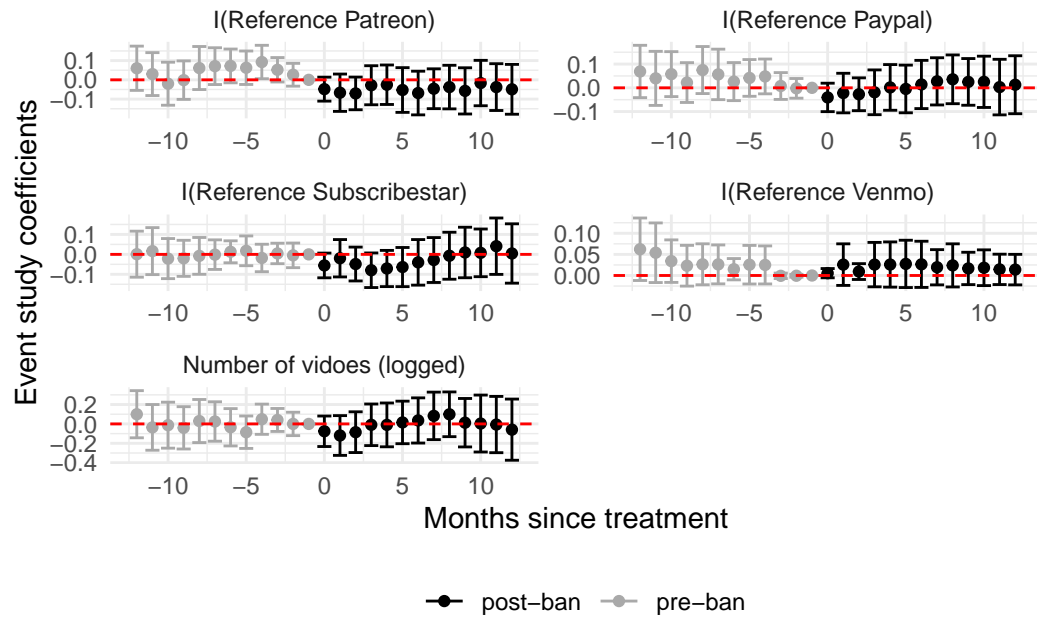


Figure B.4: Event study for control units. Study includes 12 leads and lags with all other leads and lags binned at +/- 13 and omitted from the graph. One month before the ban is omitted.

## **B.5 Robustness Tests**

Table B.2: Effect of a YouTube ban on weekly bitcoin revenue on Bitchute robustness checks.

	Profile cluster	Omit unused channels	Potential Duplicate Channel	Omit reinstated banned channels
YouTube Ban	0.257** (0.118)	0.246* (0.144)	0.210* (0.124)	0.263** (0.120)
Number Videos	0.001 (0.002)	0.000 (0.002)	0.000 (0.002)	0.001 (0.002)
Std. Errors	by: profile_name	by: channel_name	by: channel_name	by: channel_name
FE: Channel		X	X	X
FE: Profile	X			
FE: Date	X	X	X	X
I(Referenced Patreon)	X	X	X	X
I(Referenced Paypal)	X	X	X	X
I(Referenced Subscribestar)	X	X	X	X
I(Referenced Venmo)	X	X	X	X
Number clusters	77	68	73	78

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

*Note:*

Analysis is conducted at the channel week level using a TWFE strategy. In model 1, standard errors are clustered at the profile instead of the channel. In model 2, channels are removed due to inactivity prior to their YouTube ban. Model 3 removes potential duplicate channels. Model 4 removes one banned channel that was unbanned at some point. All specifications were checked for negative weights. Each specification is robust to heterogeneous treatment effects. Date refers to the week.

## B.6 Mechanisms Continued

Table B.3 and Figure B.5 presents the TWFE estimates for the five outcomes. *Average donation size* and *total time* suggests significant differences in pre-trends implying the estimates do not identify a causal effect. Notably, *number of donations* and the probability of receiving at least one donation per week both increase significantly after the YouTube ban. A Bitchute channel is 6.3 percentage points more likely to receive at least one donation after the YouTube ban compared to Bitchute channels linked to never-banned YouTube channels. The channels receive 5.1% more donations per week. In absolute terms, this is an increase from an average of .26 donations per week to .27 donations per week.

Figure B.5 plots the event study graphs using Equation 2 omitting *number of videos* from the right-hand side. Standard errors are clustered at the channel level. There appears to be no trend in the pre-periods for *number of donations*, *number of videos*, and *probability of receiving at least one donation*. *Average donations* suggests parallel trends are implausible. This is supported by a joint significance test for the pre-ban coefficients for each outcome.

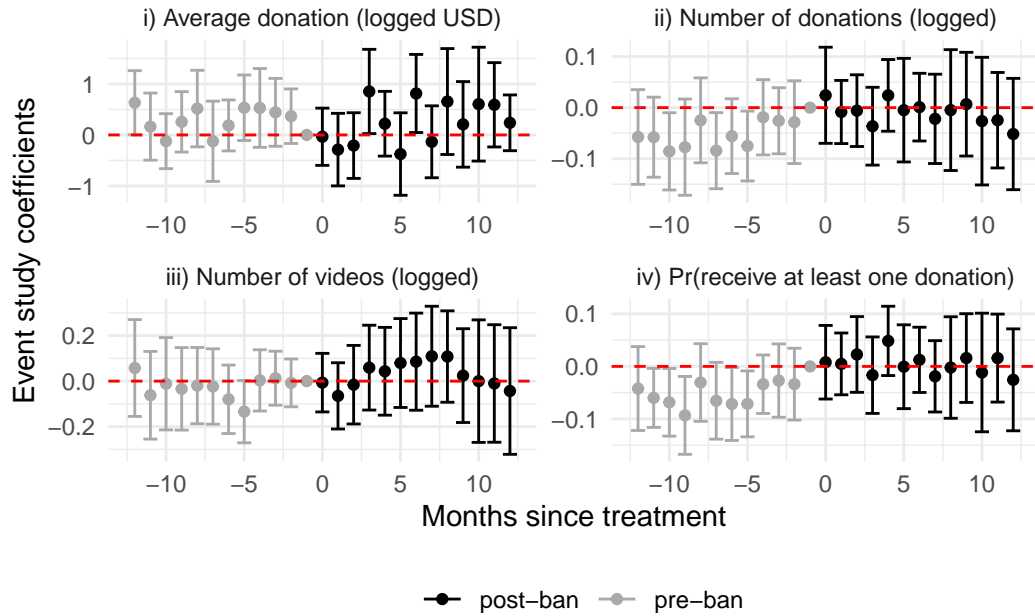


Figure B.5: Additional outcomes to study mechanism. Outcomes i), ii) and iv) investigate intensive versus extensive effects of a YouTube ban on donations. Outcomes iii) and v) investigate if changes in the supply of content on Bitchute was affecting the change in revenue. The outcome is above each graph with units in paranthese.

Table B.3: Effect of a YouTube ban on additional outcomes on Bitchute

	Average donation size (logged USD +1)	Number of Bitcoin donations (logged number +1)	Number of videos (logged +1)	Pr(receive at least one donation)
YouTube Ban	0.016 (0.166)	0.058 (0.032)	0.050 (0.088)	0.067 (0.027)
Std. Errors	by: channel_name	by: channel_name	by: channel_name	by: channel_name
FE: Channel	X	X	X	X
FE: Date	X	X	X	X
I(Referenced Patreon)	X	X	X	X
I(Referenced Paypal)	X	X	X	X
I(Referenced Subscribestar)	X	X	X	X
I(Referenced Venmo)	X	X	X	X

*Note:*

Additional outcomes to motivate potential mechanisms using a TWFE strategy.

Average donation size and changes in the amount of content suggest pre-trends.

The coefficients do not have a causal interpretation. All other outcomes suggest no pre-trend.

## **B.7 Additional Pretrends**




Figure B.6: Event study with 24 leads and 24 lags. All specifications are identical to equation 2 with the exception of 25 leads and lags instead of 18.



Table B.4: TWFE estimate of a YouTube ban on on weekly bitcoin revenue with sample limited to one year before and after a ban.

	Bitcoin Revenue (logged USD +1)	Bitcoin Revenue (Arcsinh(USD))
YouTube Ban	0.238 (0.107)	0.203 (0.093)
Number Videos	-0.002 (0.006)	-0.002 (0.005)
Std. Errors	by: channel_name	by: channel_name
FE: Channel	X	X
FE: Date	X	X
I(Referenced Patreon)	X	X
I(Referenced Paypal)	X	X
I(Referenced Subscribestar)	X	X
I(Referenced Venmo)	X	X

*Note:*

Treated units are limited to 12 months before and after the month of treatment using a TWFE strategy.

Analysis is performed as described in the main specification (Equation 1).

## B.8 Leave-One-Out TWFE Estimation

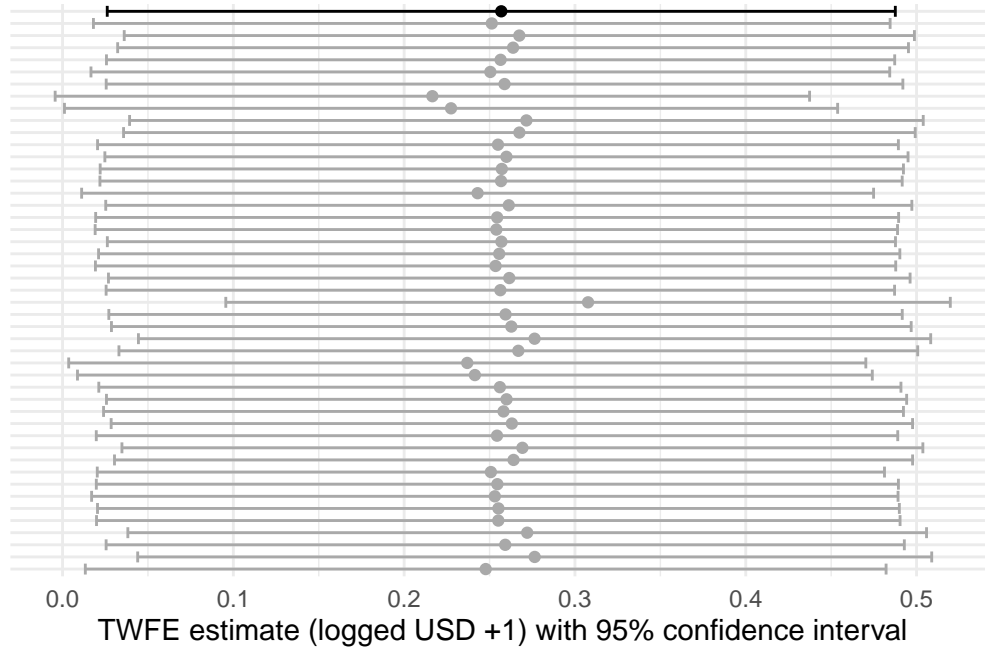


Figure B.7: Leave-one-out sensitivity analysis of Equation 2. Equation 2 is re-estimated sequentially omitting one of the banned YouTube channels. The dark black line is the estimation from the main specification (no channel left out).

Table B.5: Effect of a YouTube ban on weekly bitcoin revenue on Bitchute.

Bitcoin Revenue (Arcsinh(USD))	
YouTube Ban	0.297 (0.135)
Number Videos	0.001 (0.002)
Std. Errors	by: channel_name
FE: Channel	X
FE: Date	X
I(Referenced Patreon)	X
I(Referenced Paypal)	X
I(Referenced Subscribestar)	X
I(Referenced Venmo)	X
Number clusters	75
Effective number clusters	69

*Note:*

Analysis is conducted at the channel week level using a TWFE strategy. There are 75 channels - 45 are treated and 30 are never treated. Date refers to the week. All analysis uses the same specification as Table 2.

<sup>1</sup> The point estimate is 30.17%.

## B.9 Recreation of Main Results using Inverse Hyperbolic Sine




Figure B.8: Recreation of Figure 4 using  $\text{Arcsinh}(\text{USD})$  as the outcome.

## B.10 Main Specification Estimation with Omitting Never-Treated Channels

Table B.6: Effect of a YouTube ban on weekly bitcoin revenue using only treated units.

	Bitcoin Revenue (logged USD +1)
YouTube Ban	0.103 (0.152)
Number Videos	-0.000 (0.002)
Std. Errors	by: channel_name
FE: Channel	X
FE: Date	X
I(Referenced Patreon)	X
I(Referenced Paypal)	X
I(Referenced Subscribestar)	X
I(Referenced Venmo)	X

*Note:*

Analysis is conducted at the channel week level using a TWFE strategy. There are 45 treated channels. Date refers to the week.

The main analysis is recreated using de Chaisemartin and D’Haultfeuille (de Chaisemartin & D’Haultfoeuille (2020)) proposed estimation method herein referred to as the DIDM estimator. While robust to heterogeneous treatment effects, this estimator is inefficient leading to imprecise estimates. Since the major concern with regards to hetero-

geneous treatment effects is that the TWFE estimator is of different sign and magnitude than the average treatment effect, I focus on comparing the magnitudes. Standard errors are calculated through a bootstrap procedure. The procedure is performed 100 iterations.

Figure B.9 recreates Figure 2.5 using logged weekly bitcoin earnings and the inverse hyperbolic sine. The lead terms follow a similar pattern and do not show signs of a pre-trend, though the estimates are noisy. Similarly to Figure B.9, there is an increase in revenue the week of the ban followed by no statistically significant changes in the following 18 weeks.

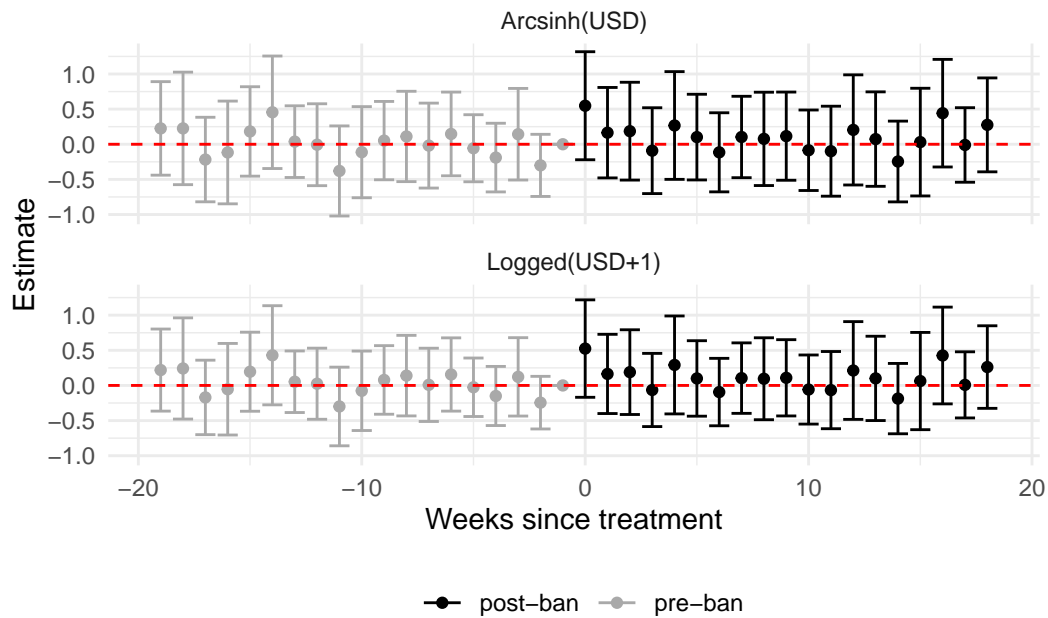


Figure B.9: Recreation of main results using DIDM estimator.

Table B.7 recreates the main results for both logged and inverse hyperbolic sine outcomes. The point estimates are within the TWFE confidence intervals. All specifications are the same (i.e. standard errors clustering, fixed effects) as the main model except estimation is performed using DIDM. The DIDM confidence interval is significantly larger due to the estimation method being less efficient.

Table B.7: DIDM estimation a YouTube ban on weekly bitcoin revenue.

	Logged (USD +1)	Arcsinh(USD)
YouTube Ban	0.16	0.27
95% Bootstrap CI	[-1.38, 1.71]	[-1.80, 2.33]
Std. Errors	Clustered (channel)	Clustered (channel)
FE: Channel	X	X
FE: Date	X	X
FE: Referenced Patreon	X	X
FE: Referenced Paypal	X	X
FE: Subscribestar	X	X
FE: Venmo	X	X
Number Clusters	75	75

*Note:*

Analysis is conducted at the channel week level. There are 75 channels - 45 are treated and 30 are never treated Date refers to the week.

# Appendix C

## Appendix for “Synthetic Control with Time Varying Coefficients: A state-space approach with Bayesian shrinkage”

### C.1 Additional Distributions

The centered distributions for  $\beta_j$  and  $\theta_j$  are (Bitto and Frhwirth-Schnatter 2019):

$$\theta_j \sim GIG \left( \frac{-T_0 - 1}{2}, \frac{1}{\xi_j^2}, \sum_{t=1}^{T_0-1} (\beta_{jt} - \beta_{j,t-1})^2 + \frac{(\beta_{j,0} - \beta_j)^2}{P_{0,jj}} \right) \quad (\text{C.1})$$

$$\beta_j \propto \mathcal{N} \left( \frac{\beta_{j,0} \alpha_j^2}{\alpha_j^2 + \theta_j P_{0,jj}}, \frac{\alpha_j^2 \theta_j P_{0,jj}}{\alpha_j^2 + \theta_j P_{0,jj}} \right) \quad (\text{C.2})$$

where  $GIG$  is the generalized inverse Gaussian distribution.






Figure C.1: Example Simulation data. Period 19 (dotted vertical line) is the hypothetical treatment period. The treatment effect is 0 for all simulations.

## C.2 Additional Figures for TVP simulation

## C.3 Empirical Monte Carlo Simulation Example - Tennessee

To further explain the Monte Carlo simulation, I provide an example using a randomly drawn state - Tennessee. Tennessee would be fit from 1970-2000 using the 37 other untreated units. For the time-invariant data generating process, the parameter estimates from Makalic and Schmidt (2016) are used to generate 1000 simulated treated units. This is done to create simulations that still have the underlying relationship between




Figure C.2: Example Empirical Monte Carlo: Actual Tennessee (black) versus 50 simulated Tennessee (grey).

treated and control units. Using a Bayesian data-generating process, each simulation will have different parameter values but maintain the underlying relationship. The approach creates similar estimates to the actual data. Figure C.2 displays 50 of the simulations for Tennessee:

Each simulated unit is fitted with the synthetic control approaches in the pre-period using the untreated units as controls, then predicted into the post period. The process is performed for every control unit. Modeling the data generating process with time invariant coefficients tests the ability of BL-TVP to shrink irrelevant coefficients. If BL-TVP isn't able to shrink irrelevant coefficients, the credibility interval in the post-period should be inflated, similarly to CI-TVP. If it does shrink, then the credibility interval should be of similar size to CI.

## C.4 Simulation Tables - Empirical Monte Carlo

Table C.1: Empirical MCMC Results for Each State in the Donor Pool

State	Average Treatment Effect Credibility Spread						SC
	ArCo	BL-TVP	BSCM-Horseshoe	CI	CI-TVP	DM-LFM	
Alabama	18.10	74.98	23.67	40.01	212.28	15.16	NA
Arkansas	24.00	77.15	21.39	46.47	229.54	15.67	NA
Colorado	45.48	79.16	28.61	37.34	284.22	29.24	NA
Connecticut	23.79	82.15	26.06	29.36	89.99	28.12	NA
Delaware	34.04	86.97	31.35	39.79	134.74	28.88	NA
Georgia	12.34	77.29	17.45	17.86	155.39	27.97	NA
Idaho	26.44	83.10	39.38	45.64	293.15	33.98	NA
Illinois	27.99	76.42	24.20	32.97	196.56	31.94	NA
Indiana	20.89	80.40	30.06	53.36	257.89	33.54	NA
Iowa	13.11	72.03	15.85	30.79	184.90	32.53	NA
Kansas	16.96	74.55	20.71	31.92	218.40	33.70	NA
Kentucky	39.89	100.42	52.37	58.39	518.20	44.71	NA

Table C.1: Empirical MCMC Results for Each State in the Donor Pool (*continued*)

State	Average Treatment Effect Credibility Spread						
	ArCo	BL-TVP	BSCM-Horseshoe	CI	CI-TVP	DM-LFM	SC
Louisiana	17.38	74.48	20.31	36.75	249.26	33.72	NA
Maine	20.73	78.12	24.61	29.86	166.92	34.32	NA
Minnesota	18.94	76.08	21.05	29.11	158.48	33.35	NA
Mississippi	16.56	73.07	21.47	36.40	191.12	33.98	NA
Missouri	10.54	73.21	12.24	23.56	133.98	32.65	NA
Montana	27.47	74.66	25.23	41.87	269.97	33.26	NA
Nebraska	15.04	72.97	13.91	29.63	147.90	33.48	NA
Nevada	49.00	100.07	65.48	78.45	548.56	42.23	NA
New Hampshire	77.43	124.93	83.31	156.98	821.77	51.94	NA
New Mexico	16.79	70.41	15.36	31.20	186.12	33.62	NA
North Carolina	43.55	112.16	71.82	122.82	675.43	45.24	NA
North Dakota	18.44	76.38	23.79	41.52	300.04	34.38	NA
Ohio	18.17	77.86	17.37	27.31	96.52	33.60	NA

Table C.1: Empirical MCMC Results for Each State in the Donor Pool (*continued*)

State	Average Treatment Effect Credibility Spread						
	ArCo	BL-TVP	BSCM-Horseshoe	CI	CI-TVP	DM-LFM	SC
Oklahoma	30.90	85.31	36.32	42.83	322.50	36.83	NA
Pennsylvania	12.70	72.84	15.58	29.91	150.69	33.41	NA
Rhode Island	23.73	78.69	27.90	37.55	181.22	31.72	NA
South Carolina	17.85	76.14	20.44	39.91	207.44	34.03	NA
South Dakota	16.73	71.60	20.33	29.85	217.11	33.09	NA
Tennessee	24.31	77.96	24.71	44.98	219.61	34.94	NA
Texas	22.48	78.13	27.03	38.99	229.86	32.55	NA
Utah	12.59	67.96	13.34	24.62	146.73	33.38	NA
Vermont	28.72	83.55	33.62	53.39	355.51	34.84	NA
Virginia	16.73	75.74	18.27	29.33	239.41	32.45	NA
West Virginia	14.51	73.77	20.10	29.35	163.98	32.64	NA
Wisconsin	11.04	72.14	13.78	18.48	122.49	33.36	NA

Table C.1: Empirical MCMC Results for Each State in the Donor Pool (*continued*)

State	Average Treatment Effect Credibility Spread						
	ArCo	BL-TVP	BSCM-Horseshoe	CI	CI-TVP	DM-LFM	SC
Wyoming	29.33	84.93	36.66	68.26	371.61	38.05	NA
<b>Average</b>	<b>24.07</b>	<b>80.20</b>	<b>27.77</b>	<b>43.07</b>	<b>253.93</b>	<b>33.49</b>	<b>NA</b>
<b>Median</b>	<b>19.83</b>	<b>76.78</b>	<b>23.73</b>	<b>37.05</b>	<b>214.69</b>	<b>33.45</b>	<b>NA</b>

Artificial Counterfactual (ArCo), Bayesian Lasso with Time Varying Parameters (BL-TVP), Bayesian Synthetic Control with Horseshoe (BSCM-Horseshoe), CausalImpact (CI), CausalImpact with time varying parameters (CI-TVP), DM-LFM, and Synthetic Control (SC) are compared using two metrics. The Average Treatment Effect Credibility Spread is defined as the length of the 95% credibility interval for the average treatment effect averaged over the 1000 simulations.

Table C.2: Empirical MCMC Results for Each State in the Donor Pool

State	Mean Squared Forecast Error						
	ArCo	BL-TVP	BSCM-Horseshoe	CI	CI-TVP	DM-LFM	SC
Alabama	59.26	33.43	40.78	404.87	10.86	23.27	14.24
Arkansas	29.74	22.64	18.28	100.28	24.66	35.88	14.35
Colorado	61.58	14.22	35.64	134.54	41.29	212.76	59.69
Connecticut	455.16	197.36	312.60	553.18	197.19	228.64	226.86
Delaware	77.33	181.63	46.94	102.31	50.78	278.96	201.01
Georgia	216.51	42.57	148.71	257.02	87.52	213.30	94.80
Idaho	34.92	59.26	44.77	40.49	56.32	64.94	34.39
Illinois	227.06	149.76	94.11	292.43	27.45	134.07	78.40
Indiana	64.80	218.85	40.64	113.05	54.63	301.93	343.44
Iowa	13.67	24.95	19.11	9.34	43.85	29.93	26.50
Kansas	46.10	11.93	20.55	67.27	18.83	32.67	15.09
Kentucky	1479.98	292.95	1089.62	2271.76	388.68	152.16	1599.07



Table C.2: Empirical MCMC Results for Each State in the Donor Pool (*continued*)

State	Mean Squared Forecast Error						
	ArCo	BL-TVP	BSCM-Horseshoe	CI	CI-TVP	DM-LFM	SC
Louisiana	22.24	39.57	46.63	10.20	96.77	15.26	22.81
Maine	285.66	93.57	172.41	214.88	54.09	76.29	58.84
Minnesota	41.10	20.66	27.30	61.57	18.92	51.66	17.01
Mississippi	23.38	62.12	16.28	30.79	69.52	28.26	36.26
Missouri	4.57	88.51	7.66	4.44	14.86	46.62	119.36
Montana	153.67	49.58	59.69	152.20	112.31	94.56	45.19
Nebraska	10.92	8.54	7.73	16.82	10.75	13.39	32.11
Nevada	667.24	419.95	412.16	588.17	266.53	199.48	195.64
New Hampshire	200.57	156.70	207.86	280.82	320.70	391.96	144.03
New Mexico	54.96	27.74	15.04	48.03	8.47	10.45	19.15
North Carolina	106.82	109.81	241.03	123.04	204.08	82.75	54.03
North Dakota	28.62	7.05	31.54	25.38	140.47	19.11	65.85
Ohio	12.68	6.33	35.26	42.26	22.49	7.57	9.24

Table C.2: Empirical MCMC Results for Each State in the Donor Pool (*continued*)

State	Mean Squared Forecast Error						
	ArCo	BL-TVP	BSCM-Horseshoe	CI	CI-TVP	DM-LFM	SC
Oklahoma	225.36	207.96	270.90	215.20	175.59	129.92	259.71
Pennsylvania	14.00	4.69	20.84	28.87	15.23	12.73	5.24
Rhode Island	559.20	368.82	343.44	461.99	204.08	478.72	443.23
South Carolina	23.77	28.38	19.76	21.08	22.87	19.53	19.28
South Dakota	33.20	76.89	52.14	34.71	59.37	43.98	52.71
Tennessee	49.66	51.64	21.63	120.75	18.31	12.47	111.39
Texas	168.62	170.40	152.79	332.57	24.73	382.46	279.88
Utah	13.33	10.91	7.98	12.22	17.16	20.09	204.70
Vermont	122.31	92.10	63.53	98.36	63.28	297.03	332.58
Virginia	406.07	86.81	302.50	520.23	75.39	228.98	199.74
West Virginia	87.68	227.94	80.49	126.16	80.20	278.16	260.80
Wisconsin	59.22	5.83	18.45	58.72	9.34	10.84	52.32

Table C.2: Empirical MCMC Results for Each State in the Donor Pool (*continued*)

State	Mean Squared Forecast Error						
	ArCo	BL-TVP	BSCM-Horseshoe	CI	CI-TVP	DM-LFM	SC
Wyoming	229.23	59.89	229.34	267.10	323.25	109.20	38.78
<b>Average</b>	<b>167.64</b>	<b>98.21</b>	<b>125.69</b>	<b>216.92</b>	<b>90.28</b>	<b>125.53</b>	<b>152.31</b>
<b>Median</b>	<b>60.42</b>	<b>59.58</b>	<b>45.70</b>	<b>107.68</b>	<b>54.36</b>	<b>70.62</b>	<b>59.26</b>

Artificial Counterfactual (ArCo), Bayesian Lasso with Time Varying Parameters (BL-TVP), Bayesian Synthetic Control with Horseshoe (BSCM-Horseshoe), CausalImpact (CI), CausalImpact with time varying parameters (CI-TVP), DM-LFM, and Synthetic Control (SC) are compared using two metrics. The Mean Squared Forecast Error is defined as the average error in the post-treatment period averaged over the 1000 simulations. Each state was simulated 1000 times using BL-TVP with time invariant coefficients.

Table C.3: Empirical MCMC Results for Each State in the Donor Pool

State	Number Simulations Reject Null SATT=0 Out of 1000					
	ArCo	BL-TVP	BSCM-Horseshoe	CI	CI-TVP	DM-LFM
Alabama	719	0	11	1	0	6
Arkansas	36	0	12	0	0	61
Colorado	18	0	16	1	0	112
Connecticut	1000	0	875	957	0	306
Delaware	12	0	4	3	0	270
Georgia	855	0	774	929	0	353
Idaho	1	0	4	0	0	0
Illinois	746	0	141	234	0	1
Indiana	325	0	21	0	0	333
Iowa	8	0	35	0	0	0
Kansas	166	0	20	0	0	0
Kentucky	948	0	465	910	0	60

Table C.3: Empirical MCMC Results for Each State in the Donor Pool (*continued*)

State	Number Simulations Reject Null SATT=0 Out of 1000					
	ArCo	BL-TVP	BSCM-Horseshoe	CI	CI-TVP	DM-LFM
Louisiana	0	0	66	0	0	0
Maine	599	0	282	140	0	0
Minnesota	129	0	53	16	0	0
Mississippi	48	0	13	0	0	0
Missouri	0	0	5	0	0	0
Montana	15	0	50	0	0	0
Nebraska	86	0	10	0	0	0
Nevada	471	0	55	5	0	12
New Hampshire	0	0	5	0	0	126
New Mexico	340	0	17	0	0	0
North Carolina	5	0	4	0	0	1
North Dakota	9	0	21	0	0	0
Ohio	27	0	85	7	0	0

Table C.3: Empirical MCMC Results for Each State in the Donor Pool (*continued*)

State	Number Simulations Reject Null SATT=0 Out of 1000					
	ArCo	BL-TVP	BSCM-Horseshoe	CI	CI-TVP	DM-LFM
Oklahoma	38	0	74	0	0	0
Pennsylvania	22	0	47	0	0	0
Rhode Island	857	0	582	494	0	803
South Carolina	8	0	14	0	0	0
South Dakota	44	0	96	0	0	0
Tennessee	159	0	11	0	0	0
Texas	593	0	216	71	0	857
Utah	50	0	10	0	0	0
Vermont	21	0	16	0	0	217
Virginia	920	0	883	851	0	116
West Virginia	642	0	218	6	0	306
Wisconsin	458	0	67	26	0	0

Table C.3: Empirical MCMC Results for Each State in the Donor Pool (*continued*)

State	Number Simulations Reject Null SATT=0 Out of 1000					
	ArCo	BL-TVP	BSCM-Horseshoe	CI	CI-TVP	DM-LFM
Wyoming	88	0	172	0	0	37
<b>Average</b>	<b>275.34</b>	<b>0</b>	<b>143.42</b>	<b>122.39</b>	<b>0</b>	<b>104.66</b>
<b>Median</b>	<b>68</b>	<b>0</b>	<b>41</b>	<b>0</b>	<b>0</b>	<b>0</b>

\* All approaches use 95% confidence/credibility intervals.

# Bibliography

- Abadie, A., Diamond, A. & Hainmueller, J. (2010), ‘Synthetic control methods for comparative case studies: Estimating the effect of california’s tobacco control program’, *Journal of the American Statistical Association* **105**(490), 493–505.  
**URL:** <http://www.tandfonline.com/doi/abs/10.1198/jasa.2009.ap08746>
- Abadie, A., Diamond, A. & Hainmueller, J. (2014), ‘Comparative politics and the synthetic control method’, *American Journal of Political Science* **59**(2), 495–510.  
**URL:** <http://dx.doi.org/10.1111/ajps.12116>
- Abadie, A. & Gardeazabal, J. (2003), ‘The Economic Costs of Conflict: A Case Study of the Basque Country’, *American Economic Review* **93**(1), 113–132.  
**URL:** <http://pubs.aeaweb.org/doi/10.1257/000282803321455188>
- ADL (2022), ‘The oath keepers data leak: Unmasking extremism in public life — adl’.  
**URL:** <https://www.adl.org/resources/report/oath-keepers-data-leak-unmasking-extremism-public-life>
- Alexa - Top sites (n.d.).  
**URL:** <https://www.alexa.com/topsites>
- Ali, S., Saeed, M., Aldreabi, E., Blackburn, J., De Cristofaro, E., Zannettou, S. & Stringhini, G. (2021), Websci ’21: Websci ’21 13th acm web science conference 2021, ACM, Virtual Event United Kingdom, pp. 187–195.  
**URL:** <https://dl.acm.org/doi/10.1145/3447535.3462637>
- Altensored.com FAQ (n.d.).  
**URL:** <https://altensored.com/about/>
- Amjad, M., Shah, D. & Shen, D. (2018), ‘Robust synthetic control’, *Journal of Machine Learning Research* **19**(22), 1–51.  
**URL:** <http://jmlr.org/papers/v19/17-777.html>
- Arkhangelsky, D., Athey, S., Hirshberg, D., Imbens, G. & Wager, S. (2021), ‘Synthetic difference-in-differences’, *American Economic Review* **111**(12), 4088–4118.  
**URL:** <http://dx.doi.org/10.1257/aer.20190159>



- Athey, S., Bayati, M., Doudchenko, N., Imbens, G. & Khosravi, K. (2020), ‘Matrix Completion Methods for Causal Panel Data Models’, *arXiv:1710.10251 [econ, math, stat]*. arXiv: 1710.10251.  
**URL:** <http://arxiv.org/abs/1710.10251>
- Bagwell, K. (2007), *Chapter 28 The Economic Analysis of Advertising*, Elsevier, pp. 1701–1844.  
**URL:** [http://dx.doi.org/10.1016/S1573-448X\(06\)03028-7](http://dx.doi.org/10.1016/S1573-448X(06)03028-7)
- Balcells, L. & Torrats-Espinosa, G. (2018), ‘Using a natural experiment to estimate the electoral consequences of terrorist attacks’, *Proceedings of the National Academy of Sciences* **115**(42), 10624–10629.  
**URL:** <http://dx.doi.org/10.1073/pnas.1800302115>
- Belmonte, M. A., Koop, G. & Korobilis, D. (2014), ‘Hierarchical Shrinkage in Time-Varying Parameter Models: Hierarchical Shrinkage in Time-Varying Parameter Models’, *Journal of Forecasting* **33**(1), 80–94.  
**URL:** <http://doi.wiley.com/10.1002/for.2276>
- Ben-Michael, E., Feller, A. & Rothstein, J. (2021), ‘The augmented synthetic control method’, *Journal of the American Statistical Association* **116**(536), 1789–1803.  
**URL:** <http://dx.doi.org/10.1080/01621459.2021.1929245>
- Berman, E. (2009), *Radical, Religious, and Violent*, The MIT Press.  
**URL:** <http://dx.doi.org/10.7551/mitpress/7881.001.0001>
- Berman, E. & Laitin, D. D. (2008), ‘Religion, terrorism and public goods: Testing the club model’, *Journal of Public Economics* **92**(10-11), 1942–1967.  
**URL:** <http://dx.doi.org/10.1016/j.jpubeco.2008.03.007>
- Billmeier, A. & Nannicini, T. (2013), ‘Assessing Economic Liberalization Episodes: A Synthetic Control Approach’, *Review of Economics and Statistics* **95**(3), 983–1001.  
**URL:** [http://www.mitpressjournals.org/doi/10.1162/REST\\_a00324](http://www.mitpressjournals.org/doi/10.1162/REST_a00324)
- Bitto, A. & Frhwirth-Schnatter, S. (2019), ‘Achieving shrinkage in a time-varying parameter model framework’, *Journal of Econometrics* **210**(1), 75–97.  
**URL:** <https://linkinghub.elsevier.com/retrieve/pii/S0304407618302070>
- Bogle, A. (2021), ‘Policy brief: Buying and selling extremism: new funding opportunities in the right-wing extremist online ecosystem’, *International Cyber Policy Center* p. 36.
- Brodersen, K. H., Gallusser, F., Koehler, J., Remy, N. & Scott, S. L. (2015), ‘Inferring causal impact using Bayesian structural time-series models’, *The Annals of Applied Statistics* **9**(1), 247–274.  
**URL:** <http://projecteuclid.org/euclid.aos/1430226092>

- Bureau, U. C. (2010), ‘Income inequality in los angeles county, ca’. Publisher: FRED, Federal Reserve Bank of St. Louis.  
**URL:** <https://fred.stlouisfed.org/series/2020RATIO006037>
- Callaway, B. & Sant’Anna, P. (2020), ‘Difference-in-differences with multiple time periods’, *Journal of Econometrics* .  
**URL:** <https://www.sciencedirect.com/science/article/pii/S0304407620303948>
- Carter, A., Schnepel, K. & Steigerwald, D. (2017), ‘Asymptotic behavior of a t-test robust to cluster heterogeneity’, *The Review of Economics and Statistics* **99**(4), 698–709.  
**URL:** [http://dx.doi.org/10.1162/REST\\_a00639](http://dx.doi.org/10.1162/REST_a00639)
- Carvalho, C., Masini, R. & Medeiros, M. (2016), The perils of counterfactual analysis with integrated processes, Technical report, Rochester, NY. DOI: 10.2139/ssrn.2894065.  
**URL:** <https://papers.ssrn.com/abstract=2894065>
- Carvalho, C., Masini, R. & Medeiros, M. (2018), ‘Arco: An artificial counterfactual approach for high-dimensional panel time-series data’, *Journal of Econometrics* **207**(2), 352–380.  
**URL:** <https://www.sciencedirect.com/science/article/pii/S0304407618301544>
- Carvalho, J.-P. & Sacks, M. (2022), ‘Radicalization’.
- Cattaneo, M., Feng, Y. & Titiunik, R. (2021), ‘Prediction intervals for synthetic control methods’, *Journal of the American Statistical Association* pp. 1–16.  
**URL:** <http://dx.doi.org/10.1080/01621459.2021.1979561>
- Center, S. P. L. (n.d.a), ‘A problem of epik proportions’.  
**URL:** <https://www.splcenter.org/hatewatch/2019/01/11/problem-epik-proportions>
- Center, S. P. L. (n.d.b), ‘Stefan molyneux’.  
**URL:** <https://www.splcenter.org/fighting-hate/extremist-files/individual/stefan-molyneux>
- Chan, N. W. & Gillingham, K. (2015), ‘The microeconomic theory of the rebound effect and its welfare implications’, *Journal of the Association of Environmental and Resource Economists* **2**(1), 133–159.  
**URL:** <http://dx.doi.org/10.1086/680256>
- Chandrasekharan, E., Pavalanathan, U., Srinivasan, A., Glynn, A., Eisenstein, J. & Gilbert, E. (2017), ‘You can’t stay here: The efficacy of reddit’s 2015 ban examined through hate speech’, *Proceedings of the ACM on Human-Computer Interaction* **1**(CSCW), 1–22.  
**URL:** <https://dl.acm.org/doi/10.1145/3134666>

- Cheney, K. (2022), ‘Text message trove shows oath keepers discussing security details for trump associates’.  
**URL:** <https://www.politico.com/news/2022/04/18/oath-keepers-security-trump-jan6-00026157>
- Chernozhukov, V., Wuthrich, K. & Zhu, Y. (2021), ‘A t-test for synthetic controls’, *arXiv:1812.10820 [econ]*. arXiv: 1812.10820.  
**URL:** <http://arxiv.org/abs/1812.10820>
- Community Guidelines strike basics - YouTube Help* (n.d.).  
**URL:** <https://support.google.com/youtube/answer/2802032?hl=en>
- Cremin, M. R. & Popescu, B. G. (2021), ‘Sticks and stones? connecting insurgent propaganda with violent outcomes’, *Journal of Conflict Resolution* **66**(3), 504–528.  
**URL:** <http://dx.doi.org/10.1177/00220027211027291>
- Cross, J., Hou, C. & Poon, A. (2020), ‘Macroeconomic forecasting with large bayesian vars: global-local priors and the illusion of sparsity’, *International Journal of Forecasting* **36**(3), 899–915.  
**URL:** <https://www.sciencedirect.com/journal/international-journal-of-forecasting>
- Crost, B. (2021), ‘Economic conditions and the rise of anti-democratic extremism’, *Empirical Studies of Conflict ESOC Working Paper*(24).  
**URL:** <https://esoc.princeton.edu/WP24>
- Cunningham, S. (2021), *Causal inference: the mixtape*, Yale University Press, New Haven.
- Dangl, T. & Halling, M. (2012), ‘Predictive regressions with time-varying coefficients’, *Journal of Financial Economics* **106**(1), 157–181.  
**URL:** <https://linkinghub.elsevier.com/retrieve/pii/S0304405X12000633>
- de Chaisemartin, C. & D’Haultfœuille, X. (2020), ‘Two-way fixed effects estimators with heterogeneous treatment effects’, *American Economic Review* **110**(9), 2964–2996.  
**URL:** <https://www.aeaweb.org/articles?id=10.1257/aer.20181169>
- Department of Justice (2022), ‘Leader of proud boys and four other members indicted in federal court for seditious conspiracy and other offenses related to u.s. capitol breach’.  
**URL:** <https://www.justice.gov/opa/pr/leader-proud-boys-and-four-other-members-indicted-federal-court-seditious-conspiracy-and>
- Department of Justice (2023a), ‘Court sentences two oath keepers leaders to 18 years in prison on seditious conspiracy and other charges related to u.s. capitol breach’.  
**URL:** <https://www.justice.gov/usao-dc/pr/court-sentences-two-oath-keepers-leaders-18-years-prison-seditious-conspiracy-and-other>

Department of Justice (2023b), ‘Two additional oath keepers members sentenced on felony charges related to u.s. capitol breach’.

**URL:** <https://www.justice.gov/opa/pr/two-additional-oath-keepers-members-sentenced-felony-charges-related-us-capitol-breach>

*Distributed Denial of Secrets* (n.d.).

**URL:** [https://ddosecrets.com/wiki/DistributedDenial\\_ofSecrets](https://ddosecrets.com/wiki/DistributedDenial_ofSecrets)

Doudchenko, N. & Imbens, G. (2016), Balancing, Regression, Difference-In-Differences and Synthetic Control Methods: A Synthesis, Technical Report w22791, National Bureau of Economic Research, Cambridge, MA.

**URL:** <http://www.nber.org/papers/w22791.pdf>

Durbin, J. & Koopman, S. J. (2012), *Time series analysis by state space methods*, number 38 in ‘Oxford statistical science series’, 2nd ed edn, Oxford University Press, Oxford.

Enders, W. & Hoover, G. A. (2012), ‘The nonlinear relationship between terrorism and poverty’, *American Economic Review* **102**(3), 267–272.

**URL:** <http://dx.doi.org/10.1257/aer.102.3.267>

Ferman, B. & Pinto, C. (2021), ‘Synthetic controls with imperfect pretreatment fit’, *Quantitative Economics* **12**(4), 1197–1221.

**URL:** <http://dx.doi.org/10.3982/QE1596>

for International Security, C. & Cooperation, S. (2022), ‘Mmp: Oath keepers’.

**URL:** <https://cisac.fsi.stanford.edu/mappingmilitants/profiles/oath-keepers>

Forest, J. J. F. (2005), *The Making of a Terrorist: Recruitment, Training, and Root Causes*, illustrated edition edn, Praeger, Westport, Conn.

Frey, B. S. (2004), *Dealing with Terrorism – Stick or Carrot?*, illustrated edition edn, Edward Elgar Publishing, Cheltenham, UK ; Northampton, MA.

Fryer, R. & Levitt, S. D. (2012), ‘Hatred and profits: Under the hood of the ku klux klan\*’, *The Quarterly Journal of Economics* **127**(4), 1883–1925.

**URL:** <https://doi.org/10.1093/qje/qjs028>

Fhrwirth-Schnatter, S. & Wagner, H. (2010), ‘Stochastic model specification search for Gaussian and partial non-Gaussian state space models’, *Journal of Econometrics* **154**(1), 85–100.

**URL:** <https://linkinghub.elsevier.com/retrieve/pii/S0304407609001614>

Giannone, D., Lenza, M. & Primiceri, G. (2021), ‘Economic predictions with big data: The illusion of sparsity’, *Econometrica* **89**(5), 2409–2437.

**URL:** <http://dx.doi.org/10.3982/ECTA17842>

- Goodman-Bacon, A. (2021), ‘Difference-in-differences with variation in treatment timing’, *Journal of Econometrics* .  
**URL:** <https://www.sciencedirect.com/science/article/pii/S0304407621001445>
- Hagenbach, J. & Koessler, F. (2017), ‘The streisand effect: Signaling and partial sophistication’, *Journal of Economic Behavior & Organization* **143**, 1–8.  
**URL:** <http://dx.doi.org/10.1016/j.jebo.2017.09.001>
- Hollingsworth, A. & Wing, C. (2020), Tactics for design and inference in synthetic control studies: An applied example using high-dimensional data, Technical report, Rochester, NY. DOI: 10.2139/ssrn.3592088.  
**URL:** <https://papers.ssrn.com/abstract=3592088>
- Horne, B., Trujillo, M., Gruppi, M. & Buntain, C. (2022), ‘The MeLa BitChute Dataset’.  
**URL:** <https://doi.org/10.7910/DVN/KRD1VS>
- Hsu, T. (2022), ‘The sandy hook defamation cases have put alex jones’s finances under scrutiny.’, *The New York Times* .
- Hua, Y., Ribeiro, M. H., West, R., Ristenpart, T. & Naaman, M. (2022), ‘Characterizing alternative monetization strategies on youtube’.  
**URL:** <https://arxiv.org/abs/2203.10143>
- Iannaccone, L. R. (1992), ‘Sacrifice and stigma: Reducing free-riding in cults, communes, and other collectives’, *Journal of Political Economy* **100**(2), 271–291. Publisher: University of Chicago Press.  
**URL:** <https://www.jstor.org/stable/2138608>
- ISD (2020), ‘Bankrolling bigotry: An overview of the online funding strategies of american hate groups’.  
**URL:** <https://www.isdglobal.org/isd-publications/bankrolling-bigotry/>
- Jackson, S. (2020), *Oath Keepers: Patriotism and the Edge of Violence in a Right-Wing Antigovernment Group*, Columbia University Press. Pages: 240 Pages.
- Jansen, S. C. & Martin, B. (2015), ‘The streisand effect and censorship backfire’, *Faculty of Law, Humanities and the Arts - Papers (Archive)* pp. 656–671.  
**URL:** <https://ro.uow.edu.au/lhapapers/1884>
- Jefferson, P. N. & Pryor, F. L. (1999), ‘On the geography of hate’, *Economics Letters* **65**(3), 389–395.  
**URL:** <https://linkinghub.elsevier.com/retrieve/pii/S0165176599001640>
- Jetter, M. (2017), ‘The effect of media attention on terrorism’, *Journal of Public Economics* **153**, 32–48.  
**URL:** <http://dx.doi.org/10.1016/j.jpubeco.2017.07.008>

- Jetter, M. (2019), ‘The inadvertent consequences of al-qaeda news coverage’, *European Economic Review* **119**, 391–410.  
**URL:** <http://dx.doi.org/10.1016/j.euroecorev.2019.08.004>
- Jhaver, S., Boylston, C., Yang, D. & Bruckman, A. (2021), ‘Evaluating the effectiveness of deplatforming as a moderation strategy on twitter’, *Proceedings of the ACM on Human-Computer Interaction* **5**(CSCW2), 1–30.  
**URL:** <https://dl.acm.org/doi/10.1145/3479525>
- Kalen Hill, Z. (2021), ‘Oath keepers leader allegedly spent group’s money at sex shop, gun store’, *Newsweek* . Section: News.  
**URL:** <https://www.newsweek.com/leader-oath-keepers-allegedly-spent-some-groups-money-sex-shop-gun-store-1601765>
- Keatinge, T., Keen, F. & Izenman, K. (2019), ‘Fundraising for right-wing extremist movements’, *The RUSI Journal* **164**(2), 10–23. Publisher: Routledge .eprint: <https://doi.org/10.1080/03071847.2019.1621479>.  
**URL:** <https://doi.org/10.1080/03071847.2019.1621479>
- Kellogg, M., Mogstad, M., Pouliot, G. & Torgovitsky, A. (2021), ‘Combining matching and synthetic control to tradeoff biases from extrapolation and interpolation’, *Journal of the American Statistical Association* **0**(0), 1–13. Publisher: Taylor Francis .eprint: <https://doi.org/10.1080/01621459.2021.1979562>.  
**URL:** <https://doi.org/10.1080/01621459.2021.1979562>
- Kim, S., Lee, C. & Gupta, S. (2020), ‘Bayesian synthetic control methods’, *Journal of Marketing Research* **57**(5), 831–852.  
**URL:** <http://dx.doi.org/10.1177/0022243720936230>
- Kinn, D. (2018), ‘Synthetic Control Methods and Big Data’, *arXiv:1803.00096 [econ]* . arXiv: 1803.00096.  
**URL:** <http://arxiv.org/abs/1803.00096>
- Klinenberg, D. (2022), ‘Synthetic control with time varying coefficients: A state space approach with bayesian shrinkage’, *Journal of Business & Economic Statistics* pp. 1–26. Publisher: Taylor & Francis.  
**URL:** <https://www.tandfonline.com/doi/abs/10.1080/07350015.2022.2102025>
- Krueger, A. (2008), *What makes a terrorist: economics and the roots of terrorism: Lionel Robbins lectures*, 4. print., and 1. paperback print. with a new afterw. by the author edn, Princeton Univ. Press, Princeton, NJ. OCLC: 845390394.
- League, A.-D. (n.d.a), ‘Bitchute: A hotbed of hate’.  
**URL:** <https://www.adl.org/blog/bitchute-a-hotbed-of-hate>

- League, A.-D. (n.d.b), ‘Despite youtube policy update, anti-semitic, white supremacist channels remain’.  
**URL:** <https://www.adl.org/blog/despite-youtube-policy-update-anti-semitic-white-supremacist-channels-remain>
- Levine, M. (2022), ‘Ex-oath keepers spokesperson warns right-wing ‘propaganda’ is ‘more dangerous than bullets’’.  
**URL:** <https://web.archive.org/web/20220816222122/https://abcnews.go.com/US/oath-keepers-spokesperson-warns-wing-propaganda-dangerous-bullets/story?id=82094999>
- Li, K. (2019), ‘Statistical inference for average treatment effects estimated by synthetic control methods’, *Journal of the American Statistical Association* **115**(532), 2068–2083.  
**URL:** <http://dx.doi.org/10.1080/01621459.2019.1686986>
- Li, K. & Bell, D. (2017), ‘Estimation of average treatment effects with panel data: Asymptotic theory and implementation’, *Journal of Econometrics* **197**(1), 65–75.  
**URL:** <http://dx.doi.org/10.1016/j.jeconom.2016.01.011>
- Linden, A. & Yarnold, P. (2016), ‘Using machine learning to identify structural breaks in single-group interrupted time series designs: Machine learning to identify structural breaks’, *Journal of Evaluation in Clinical Practice* **22**(6), 855–859.  
**URL:** <http://doi.wiley.com/10.1111/jep.12544>
- Loadenthal, M., Donahoe, L., Weaver, M., Sarah, G. & Kathryn, B. (2023), ‘the prosecution project’.  
**URL:** <https://theprosecutionproject.org/>
- Makalic, E. & Schmidt, D. F. (2016), ‘High-Dimensional Bayesian Regularised Regression with the BayesReg Package’, *arXiv:1611.06649 [stat]*. arXiv: 1611.06649.  
**URL:** <http://arxiv.org/abs/1611.06649>
- Masini, R. & Medeiros, M. (2020), ‘Counterfactual analysis and inference with nonstationary data’, *Journal of Business Economic Statistics* pp. 1–13.  
**URL:** <https://www.tandfonline.com/doi/full/10.1080/07350015.2020.1799814>
- Masini, R. & Medeiros, M. (2021), ‘Counterfactual analysis with artificial controls: Inference, high dimensions, and nonstationarity’, *Journal of the American Statistical Association* **116**(536), 1773–1788.  
**URL:** <http://dx.doi.org/10.1080/01621459.2021.1964978>
- Mathew, B., Illendula, A., Saha, P., Sarkar, S., Goyal, P. & Mukherjee, A. (2020), ‘Hate begets hate: A temporal study of hate speech’, *Proceedings of the ACM on Human-Computer Interaction* **4**(CSCW2), 1–24.  
**URL:** <https://dl.acm.org/doi/10.1145/3415163>

- McCausland, W., Miller, S. & Pelletier, D. (2011), ‘Simulation smoothing for state–space models: A computational efficiency analysis’, *Computational Statistics Data Analysis* **55**(1), 199–212.  
**URL:** <https://www.sciencedirect.com/science/article/pii/S0167947310002823>
- Mitts, T. (2018), ‘From isolation to radicalization: Anti-muslim hostility and support for isis in the west’, *American Political Science Review* **113**(1), 173–194.  
**URL:** <http://dx.doi.org/10.1017/S0003055418000618>
- Mitts, T. (2021), ‘Banned: How deplatforming extremists mobilizes hate in the dark corners of the internet’, *Working Paper* p. 35.
- Mitts, T., Pisharody, N. & Shapiro, J. (2022), ‘Removal of anti-vaccine content impacts social media discourse’, *14th ACM Web Science Conference 2022* .  
**URL:** <http://dx.doi.org/10.1145/3501247.3531548>
- Morales, K., Raynold, P. & Li, J. (2018), ‘The empirical relationship between commitment enhancement devices and terrorism’, *Applied Economics* **50**(50), 5366–5380.  
**URL:** <http://dx.doi.org/10.1080/00036846.2018.1486991>
- Mulholland, S. E. (2010), ‘Hate fuel: On the relationship between local government policy and hate group activity’, *Eastern Economic Journal* **36**(4), 480–499.  
**URL:** <http://dx.doi.org/10.1057/ej.2009.38>
- Mulholland, S. E. (2012), ‘White supremacist groups and hate crime’, *Public Choice* **157**(1-2), 91–113.  
**URL:** <http://dx.doi.org/10.1007/s11127-012-0045-7>
- Munger, K. & Phillips, J. (2020), ‘Right-wing youtube: A supply and demand perspective’, *The International Journal of Press/Politics* p. 194016122096476.  
**URL:** <http://journals.sagepub.com/doi/10.1177/1940161220964767>
- Mhlbach, N. N. (2020), ‘Tree-based Synthetic Control Methods: Consequences of moving the US Embassy’, *arXiv:1909.03968 [econ]* . arXiv: 1909.03968.  
**URL:** <http://arxiv.org/abs/1909.03968>
- Nilsson, M. (2022), ‘“aren’t you tired of talking?” – priming men and women into violence through gateway organizations’, *Studies in Conflict Terrorism* pp. 1–22.  
**URL:** <http://dx.doi.org/10.1080/1057610X.2022.2123217>
- of Homeland Security, D. (2020), ‘Homeland threat assessment’, p. 26.  
**URL:** [https://www.dhs.gov/sites/default/files/publications/2020106\\_homeland\\_threat\\_assessment.pdf](https://www.dhs.gov/sites/default/files/publications/2020106_homeland_threat_assessment.pdf)
- Pang, X., Liu, L. & Xu, Y. (2021), ‘A bayesian alternative to synthetic control for comparative case studies’, *Political Analysis* pp. 1–20.  
**URL:** <http://dx.doi.org/10.1017/pan.2021.22>



- Park, T. & Casella, G. (2008), ‘The Bayesian Lasso’, *Journal of the American Statistical Association* **103**(482), 681–686.  
**URL:** <http://www.tandfonline.com/doi/abs/10.1198/016214508000000337>
- Peltzman, S. (1975), ‘The effects of automobile safety regulation’, *Journal of Political Economy* **83**(4), 677–725.  
**URL:** <http://dx.doi.org/10.1086/260352>
- Piazza, J. A. (2016), ‘The determinants of domestic right-wing terrorism in the usa: Economic grievance, societal change and political resentment’, *Conflict Management and Peace Science* **34**(1), 52–80.  
**URL:** <http://dx.doi.org/10.1177/0738894215570429>
- Powell, D. (2021), ‘Synthetic control estimation beyond comparative case studies: Does the minimum wage reduce employment?\*', *Journal of Business Economic Statistics* pp. 1–39.  
**URL:** <http://dx.doi.org/10.1080/07350015.2021.1927743>
- Raleigh, H., Stall, R. & Kishi, C. (2020), ‘Standing by: Right-wing militia groups and the united states election’.  
**URL:** <https://acleddata.com/2020/10/21/standing-by-militias-election/>
- Rauchfleisch, A. & Kaiser, J. (2021), Deplatforming the far-right: An analysis of youtube and bitchute, Technical report, Rochester, NY. DOI: 10.2139/ssrn.3867818.  
**URL:** <https://papers.ssrn.com/abstract=3867818>
- Roberts, G. & Sahu, S. (1997), ‘Updating schemes, correlation structure, blocking and parameterization for the gibbs sampler’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **59**(2), 291–317. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/1467-9868.00070>.  
**URL:** <https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-9868.00070>
- Rogers, R. (2020), ‘Deplatforming: Following extreme internet celebrities to telegram and alternative social media’, *European Journal of Communication* **35**(3), 213–229.  
**URL:** <http://journals.sagepub.com/doi/10.1177/0267323120922066>
- Samartsidis, P., Seaman, S. R., Montagna, S., Charlett, A., Hickman, M. & Angelis, D. D. (2020), ‘A Bayesian multivariate factor analysis model for evaluating an intervention by using observational time series data on multiple outcomes’, *Journal of the Royal Statistical Society: Series A (Statistics in Society)* p. rssa.12569.  
**URL:** <https://onlinelibrary.wiley.com/doi/abs/10.1111/rssa.12569>
- Savage, S. & Wimmer, B. S. (2023), ‘Local entry in the market for hate’, *SSRN Electronic Journal* .  
**URL:** <http://dx.doi.org/10.2139/ssrn.4360248>

- Schmidheiny, K. & Siegloch, S. (2019), ‘On event study designs and distributed-lag models: Equivalence, generalization and practical implications’, p. 28.
- Scott, S. & Varian, H. (2013), Bayesian variable selection for nowcasting economic time series, Technical report, Cambridge, MA. DOI: 10.3386/w19567.  
**URL:** <http://www.nber.org/papers/w19567.pdf>
- Shapiro, J. N. (2013), ‘The terrorist’s dilemma’.  
**URL:** <http://dx.doi.org/10.23943/princeton/9780691157214.001.0001>
- Solon, G., Haider, S. J. & Wooldridge, J. M. (2015), ‘What are we weighting for?’, *Journal of Human Resources* **50**(2), 301–316.  
**URL:** <http://dx.doi.org/10.3368/jhr.50.2.301>
- SPLC (2021), ‘Antigovernment movement’.  
**URL:** <https://www.splcenter.org/fighting-hate/extremist-files/ideology/antigovernment>
- SPLC (2022), ‘Oath keepers’.  
**URL:** <https://www.splcenter.org/fighting-hate/extremist-files/group/oath-keepers>
- SPLC (n.d.), ‘Rise above movement’.  
**URL:** <https://www.splcenter.org/fighting-hate/extremist-files/group/rise-above-movement>
- Squire, M. (2021), ‘Monetizing propaganda: How far-right extremists earn money by video streaming’, p. 13.
- START (2023), ‘Profiles of individual radicalization in the united states’.  
**URL:** <http://www.start.umd.edu/pirus>
- Tatenhove, J. V. (2023), *The Perils of Extremism: How I Left the Oath Keepers and Why We Should be Concerned about a Future Civil War*, Skyhorse.
- Tibshirani, R. (1996), ‘Regression Shrinkage and Selection via the Lasso’, *Journal of the Royal Statistical Society. Series B (Methodological)* **58**(1), 267–288.  
**URL:** <http://www.jstor.org/stable/2346178>
- Tokgoz, K. (2012), *Enhancing Cooperation in Defence against Terrorism - Volume 99 NATO Science for Peace and Security Series - E: Human and Societal Dynamics*, IOS Press, Washington, DC.
- Tritten, T. (2015), ‘Army to recruiters: Treat armed citizens as security threat’.  
**URL:** <https://www.stripes.com/theaters/us/army-to-recruiters-treat-armed-citizens-as-security-threat-1.359134>

- Trujillo, M., Gruppi, M., Buntain, C. & Horne, B. (2020), What is bitchute? characterizing the, HT '20, Association for Computing Machinery, New York, NY, USA, p. 139–140.  
**URL:** <https://doi.org/10.1145/3372923.3404833>
- Trujillo, M., Gruppi, M., Buntain, C. & Horne, B. D. (2022), ‘The mela bitchute dataset’.
- U. S. Government Accountability Office, T. (2023), ‘The rising threat of domestic terrorism in the u.s. and federal efforts to combat it — u.s. gao’.  
**URL:** <https://www.gao.gov/blog/rising-threat-domestic-terrorism-u.s.-and-federal-efforts-combat-it>
- Valasik, M. & Reid, S. E. (2021), ‘The alt-right movement and national security’, *The US Army War College Quarterly: Parameters* **51**(3).  
**URL:** <http://dx.doi.org/10.55540/0031-1723.3076>
- Van Dijcke, D. & Wright, A. L. (2021), ‘Profiling insurrection: Characterizing collective action using mobile device data’, *SSRN Electronic Journal* .  
**URL:** <https://www.ssrn.com/abstract=3776854>
- Wagner, B., Rozgonyi, K., Sekwenz, M.-T., Cobbe, J. & Singh, J. (2020), ‘Regulating transparency? facebook, twitter and the german network enforcement act’, p. 11.
- WalletExplorer: Smart Bitcoin Block Explorer* (n.d.).  
**URL:** <https://www.walletexplorer.com/>
- Williams, H. J., Matthews, L. J., Moore, P., DeNardo, M. A., Marrone, J. V., Jackson, B. A., Marcellino, W. & Helmus, T. C. (2022), Mapping white identity terrorism and racially or ethnically motivated violent extremism: A social network analysis of online activity, Technical report.  
**URL:** [https://www.rand.org/pubs/research\\_reports/RRA1841-1.html](https://www.rand.org/pubs/research_reports/RRA1841-1.html)
- Wilson, J. (2022), ‘Exclusive: Oath keepers leader stewart rhodes’ children speak’.  
**URL:** <https://www.splcenter.org/hatewatch/2022/05/12/exclusive-oath-keepers-leader-stewart-rhodes-children-speak>
- Xu, Y. (2017), ‘Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models’, *Political Analysis* **25**(1), 57–76.
- Youngblood, M. (2020), ‘Extremist ideology as a complex contagion: the spread of far-right radicalization in the united states between 2005 and 2017’, *Humanities and Social Sciences Communications* **7**(1).  
**URL:** <http://dx.doi.org/10.1057/s41599-020-00546-3>

Yu, Y. & Meng, X. (2011), ‘To center or not to center: That is not the question - an ancillarity-sufficiency interweaving strategy (asis) for boosting mcmc efficiency’, *Journal of Computational and Graphical Statistics* **20**(3), 531–570. Publisher: Taylor Francis .eprint: <https://doi.org/10.1198/jcgs.2011.203main>.

**URL:** <https://doi.org/10.1198/jcgs.2011.203main>

Zhuravskaya, E., Petrova, M. & Enikolopov, R. (2020), ‘Political effects of the internet and social media’, *Annual Review of Economics* **12**(1), 415–438. .eprint: <https://doi.org/10.1146/annurev-economics-081919-050239>.

**URL:** <https://doi.org/10.1146/annurev-economics-081919-050239>