# The Effect of State Data Breach Notification Laws on Medical Identity Theft

Aniket Kesari
JD/PhD Student, Jurisprudence & Social Policy
Yale University/The University of California, Berkeley

October 10, 2020

**Abstract**

As the number of data breaches in the United States grows each year, cybersecurity has become an increasingly important policy area. The primary mechanism for regulating and deterring data breaches is the "data breach notification law." Every U.S. state now has such a law that mandates that certain organizations disclose data breaches to their data subjects. Despite the popularity of these laws, there is relatively little evidence about their effectiveness at deterring breaches, and therefore reducing identity theft. Using medical identity theft panel data collected from the Consumer Financial Protection Bureau (CFPB), this study implements an augmented synthetic control approach to analyze the effect of certain data breach notification standards on medical identity theft.

## 1 Introduction

Identity theft and data breaches are becoming more common, and consequently, cybersecurity is quickly coming to the forefront of policy discourse. Data breaches can compromise consumer information related to things such as transaction history, payment information, health data, and personally identifiable information (PII). The financial and reputational harms that stem from these kinds of losses can be large in magnitude, but also difficult to detect because the consequences of identity theft do not materialize immediately.

Despite the growing incidents of data breaches and identity theft, there is no single federal law that regulates how an organization must respond to a data breach once it has been discovered. There are sectoral regulations that affect certain industries, however. One such sectoral regulation is the Health Insurance Portability and Accountability Act (HIPAA) that regulates healthcare information. One of its provisions requires that medical providers and health insurance companies provide notice to the department of Health and Human Services (HHS) and their data subjects when unencrypted data is breached.

States may adopt stricter requirements on top of the federal requirements. Beginning in the mid-2000s, nearly every state adopted data breach notification laws for all organizations that maintain unencrypted data from various sectors. In 2016, California amended its existing data breach notification law to mandate disclosure of breaches even when medical data is encrypted. Between 2016 and 2019, Illinois, Nebraska, New York, Rhode Island, Maryland, Delaware, and New Mexico explicitly included protected health information ("PHI") in their definitions of personal information. Some states are now moving toward creating private causes of action following data breaches.

Despite the popularity and continued adoption of state data breach notification laws, there is relatively little evidence on their efficacy. The two main theories underlying data breach notification laws are that they will encourage organizations to invest in better cybersecurity practices so they can avoid making damaging disclosures, and that they will allow consumers the opportunity to guard against identity theft once a disclosure is made. Concretely, the goal of these laws is to minimize identity theft, yet the exercise of determining whether they work is fraught with serious methodological challenges. A simple research design that looks at the number of breaches before and after the passage of a data breach notification law would be intractable because the pre-treatment figure is difficult to estimate prior to the passage of such a law. Analyzing the effect of laws on large corporations with a nationwide presence becomes tricky when such organizations are likely to default to the strictest state's standards rather than tailor different notices to citizens of different states (Bradford, 2020). [1] Identity theft is also notoriously underreported, and many victims may not discover the crime until months or years after it occurs.

To address these challenges, I use a panel dataset containing medical identity theft complaints, and analyze the effect of breach notification requirements using an augmented synthetic control approach. Specifically, I examine the effect of California's 2015 amendments that expanded breach notification requirements to include encrypted medical data. In adopting this expanded breach notification standard, California considerably raised the bar for organizations that hold critical health data. More broadly, the effects these laws have on healthcare organizations could potentially be generalized to other types of data, including financial, educational, and consumer data.

# 2 Overview of Data Breach Notification Laws

## 2.1 Law & Economics of Data Breaches

The classic law and economics theory of crime suggests that the law should minimize the social cost of a crime, which is equal to the sum of the harm the criminal activity causes, and the costs of preventing that activity (Cooter and Ulen, 2016). This economic rule provides insights into

[1] As with many other regulatory areas, privacy law may be subject to a "California Effect." The California Effect (also called the Brussels Effect when discussing the European Union) essentially describes the phenomenon where a large market is able to unilaterally impose its regulatory preferences on other markets. In this case, California privacy laws that regulate large businesses will change privacy regulation across U.S. states. See Anu Bradford's, "The Brussels Effect" for more on the general theory.

the optimal level of punishment that the state should set to deter crime. In contrast to tort law, criminal sanctions are justified primarily as a mechanism for deterring criminal behavior, rather than forcing perpetrators and victims to internalize the costs of their own actions.

Cybercrime vexes policymakers because cybercriminals' benefits are large, the costs of perpetrating cybercrimes are low, and the costs of punishment are high. The economics of crime suggests that a rational criminal will choose to commit crime when the utility of the crime exceeds the utility of not committing the crime. Formally:

$$(1-p)\mu_s + p\mu_f \geq \mu_{nc}$$

Where:

- $\mu_s$ is the utility of successfully committing a crime
- $\mu_f$ is the utility of failing to commit a crime
- is the probability of being caught and punished
- $\mu_{nc}$ is the utility of not committing a crime

In this framework, the probability of being caught and punished, $p$, is vanishingly small. Cybercriminals frequently operate anonymously, outside of the local jurisdiction, and in complex networks. These facts make it difficult for law enforcement to identify, much less apprehend, cybercriminals. Thus $p\mu_f$ becomes very small because $p$ is close to 0, meaning even large punishments ($\mu_f$) will be ineffective deterrents. Deterrence through punishment is therefore an ineffective strategy, even if the social costs of cybercriminal activity are high.

Cardenas et. al. provide a typology and economic analysis of these dynamics in "An Economic Map of Cybercrime" (Cardenas et al., 2010). They explicate various types of cybercrime techniques including malware, botnet herding, phishing, Distributed Denial of Service (DDoS), and identity theft, among others. They also argue that estimating the social costs of cybercrime is difficult because there is little reliable data about the costs borne by companies that are the victims of cybercrime. More disclosure could help alleviate this problem. Similarly, estimating the benefits accrued by cybercriminals is difficult because of the uncertainty surrounding the monetization of cybercrime.

Identity theft illustrates some of the difficulty of estimating the benefits of cybercrime. Creating the necessary infrastructure to properly impersonate someone for these purposes is quite laborious, and mitigates the expected benefit of identity theft. Conducting identity theft at scale becomes extremely difficult because of the elaborate process involved with impersonating even one person. Thus only a fraction of the records compromised in a data breach will be successfully used for conducting identity theft. Although the benefits to criminals are ameliorated by the complexity of the apparatus involved, the problem of estimating the benefits is still difficult because of the lack of consistent data.

These dynamics make the law and economics of cybercrime slightly different than conventional accounts of the law and economics of crime in that law focuses more on the benefits than the costs of the crime. In particular, with such a low probability of punishment, estimating the optimal level of punishment to achieve deterrence would lead to an implausibly high figure for legal sanctions. In practice, cybercrime law instead focuses on deterrence by denial, specifically by reducing the benefits of cybercrime. For instance, if breached companies provide their consumers with identity protection services, the benefits of identity theft will be reduced, thus making cybercriminal activity less profitable.

This focus on reducing benefits motivates data breach notification laws. Because sanctions against perpetrators are ineffective, regulatory attention is instead placed on organizations that collect and hold data. By requiring disclosures from breached companies, policymakers aim to nudge organizations to invest in cybersecurity and give consumers adequate opportunity to safeguard their identities. These goals both serve to reduce the potential benefits that might be realized by a cybercriminal by making stolen credentials useless for impersonating someone.

## 2.2 Legal Background

As of 2018, every U.S. state now has a data breach notification law. These laws share several similarities, but there are also some key variations that are important to note. Generally, data breach notification laws contain the following elements:

- **Definition of Personal Information**: Examples of personal information that is covered by the law. Some example include email addresses and passwords, health information, driver's licenses, federal identification numbers, and biometric data.

- **Covered Entities**: Entities that must comply with the law. Typically, all government agencies, businesses, and non-profits are covered.

- **Encryption Safe Harbor**: Whether organizations are exempt from disclosure requirements if the breached data was encrypted. Almost every state provides a safe harbor for encrypted data.

- **Notification Trigger**: The threshold that triggers a notification obligation. Examples include "substantial harm to individuals," "reasonable likelihood of harm," or "awareness of breach." Under the first two standards, organizations only need to provide notice if they believe there will be harm to their data subjects, whereas standards closer to "awareness of breach" remove this discretion.

- **Content**: The content of the breach notification. Most states do not mandate any specific content, while others regulate the information that must be provided, including things like descriptions of the incident, types of personal information compromised, and toll-free numbers for consumers to call.

- **Timing of Notification**: How long an organization has to provide notice once it has discovered a breach. Common requirements are that notice must be provided within 30, 45, 60, or 90 days, though some states do not specify a timeframe at all.
- **Penalty** The civil penalty associated with failure to comply with the requirements of the law. Some states penalize by days over the time limit, while others penalize by number of individuals affected.
- **Cause of Action**: Whether consumers have the right to bring a cause of action following a breach notification.
- **Notice to State AG or CRAs**: Some states include provisions that require additional notice to the state Attorney General and/or consumer credit reporting agencies.

While states may differ on some of these specific requirements, much of the legislation shares common language and policies. These similarities may indicate some degree of policy diffusion across state jurisdictions. This diffusion is helpful primarily because states share a remarkably similar common baseline for breach notification requirements, and researchers can analyze the differences that arise from particular policy choices that are layered on top of this common baseline.

In 2016, California amended its existing breach notification standards to change the "content" aspect of its law. An example disclosure that conforms to this law can be seen in Figure 1. In particular, it requires that the notice has the following headings:

- Subject Line that says "NOTICE OF DATA BREACH"
- **What Happened**: A description of when and how the breach occurred.
- **What We Are Doing**: Actions that the organization is taking to mitigate potential harms.
- **What You Can Do**: Suggestions for how data subjects can safeguard their identity.
- **If You Have Questions**: Contact information and toll-free numbers for consumers to call with concerns.

These requirements essentially mandate that all disclosures provide certain content and are organized in a specific way - there is little room for an organization to change the style mandated by law. Thus, this amended law provides a good vehicle for exploring the effect of mandated disclosure on reported medical identity theft.

# 3 Literature Review

Data breach notification legislation has been explored in a few different pieces, but is largely an understudied area of law. In part, this may be because data breaches are a relatively new phenomenon. Furthermore, despite their costliness to victims, they are still somewhat rare events. Moreover, companies may not even be aware that they have been the victims of a cyberattack, and

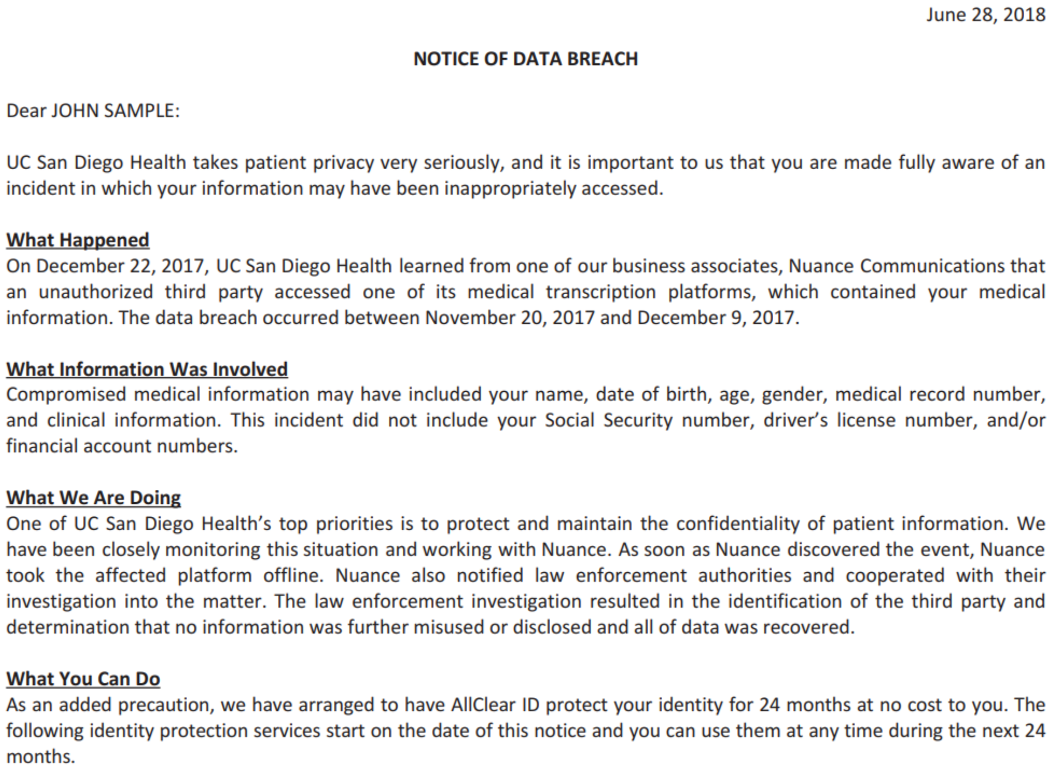

June 28, 2018

**NOTICE OF DATA BREACH**

Dear JOHN SAMPLE:

UC San Diego Health takes patient privacy very seriously, and it is important to us that you are made fully aware of an incident in which your information may have been inappropriately accessed.

**What Happened**

On December 22, 2017, UC San Diego Health learned from one of our business associates, Nuance Communications that an unauthorized third party accessed one of its medical transcription platforms, which contained your medical information. The data breach occurred between November 20, 2017 and December 9, 2017.

**What Information Was Involved**

Compromised medical information may have included your name, date of birth, age, gender, medical record number, and clinical information. This incident did not include your Social Security number, driver's license number, and/or financial account numbers.

**What We Are Doing**

One of UC San Diego Health's top priorities is to protect and maintain the confidentiality of patient information. We have been closely monitoring this situation and working with Nuance. As soon as Nuance discovered the event, Nuance took the affected platform offline. Nuance also notified law enforcement authorities and cooperated with their investigation into the matter. The law enforcement investigation resulted in the identification of the third party and determination that no information was further misused or disclosed and all of data was recovered.

**What You Can Do**

As an added precaution, we have arranged to have AllClear ID protect your identity for 24 months at no cost to you. The following identity protection services start on the date of this notice and you can use them at any time during the next 24 months.

Figure 1: Example of a Breach Notification Under the 2016 California Amendments from UC San Diego

the consumers who lost data may be unaware of identity theft for months of years after the breach. Finally, absent legal mandates, they are likely to be underreported, as victimized companies have strong incentives to not report or delay reporting serious incidents.

The major obstacle to empirically studying identity theft is lack of data. Chris Hoofnagle describes this problem as one of making known unknowns known (Hoofnagle, 2007). One issue with studying the extent of identity theft is that much of the collected data is drawn from surveys that are fraught with sampling errors. Without comprehensive data on identity theft, studying the effect of interventions becomes difficult. In another piece, Hoofnagle addresses this issue by using FTC Consumer Complaint Data (Hoofnagle, 2008). Here, he argues that identity theft victim data provides evidence that different financial institutions have different identity theft protection practices, and therefore different footprints. This evidence suggests that there is a plausible market for identity theft and fraud protection, and that government interventions can be targeted at a handful of firms with outsized identity theft footprints. Unfortunately, as discussed later, the full FTC Consumer Complaint Database is no longer available to researchers.

One such intervention is the "data breach notification law." States have adopted breach notification laws under the theory that breach notification will create incentives for companies to preemptively invest in cybersecurity measures. The idea is that a company that is forced to disclose a cybersecurity incident will face consumer backlash (Romanosky et al., 2011). Hoping to avoid such a consequence, companies will invest in cybersecurity beforehand so as to avoid the possibility of hurting their image. Additionally, should a breach actually be disclosed, and consumers are made aware of it, then consumers may pro-actively take steps to protect their identities, thus reducing the overall identity theft rate.

Sasha Romanosky and colleagues have offered the most extensive study of these data breach notification laws to date. In "Do Data Breach Disclosure Laws Reduce Identity Theft?," Romanosky et. al. directly address this issue. Using a panel data regression approach, they find that states that adopted breach notification laws experienced about a 2% decrease in per capita rates of identity theft. They included controls for state adoption of a breach notification law. They also checked for potential endogeneity issues (states adopted laws because they were experiencing a lot of cybercrime), but found this was a negligible factor with robustness checks (Romanosky et al., 2011). However, such checks are never able to provide absolute guarantees.

Another strand of the literature is concerned with how markets react to the announcement of a breach. This focus tests the notion that negative consequences stem from a breach announcement. The absence of an effect here would imply that firms do not have a strong incentive to invest in cybersecurty measures when faced with breach disclosures. Sanjay Goel and Hany Shawky, in two separate pieces, examine these market effects. In "Estimating the Market Impact of Security Breach Announcements on Firm Values," they use event study methodology to examine the effect a breach disclosure has on a firm's stock market value. In this piece, they find about a 1% decline in market value immediately following the disclosure of a cybersecurity incident (Goel and Shawky, 2009). In "The Impact of Federal and State Notification Laws on Security Breach Announcements," they use a similar event study method and conclude that after the passage of data breach notification legislation, the negative effect on stock prices is somewhat mitigated. They conclude that this implies that the legislation is effective in forcing firms to mitigate the potential damage from a cyberattack (though it could also be evidence that firms are mitigating the message sent to investors) (Goel and Shawky, 2014).

More recently, Joshua Mitts and Eric Talley in an upcoming Harvard Business Law Review piece entitled "Informed Trading and Cybersecurity Breaches," examine whether there is evidence for insider trading prior to a breach announcement. Using matched sampling to compare breached and unbreached firms, they find systematic evidence that arbitrage occurs prior to a breach announcement. This implies that there are individuals who have prior knowledge of a breach before the market does. They argue that this raises normative concerns that go beyond run-of-the mill insider trading because the harms a hacker causes a company are endogenous to the company's cybersecurity practices, and thus creates an opportunity for sophisticated arbitrage based on knowledge of a company's security vulnerabilities. This is distinct from an informed trader using exogenous information because allowing insider trading on cyberattacks effectively subsidizes hacking activity, whereas traditional insider trading is generally an exercise in price discovery (Mitts and Talley, 2018).

In the legal literature, much of the attention toward data breach notification laws has been targeted at whether a federal standard is appropriate, and what the contours of mandatory notification should look like. In "Federal Security Breach Notifications: Policy and Approaches," Priscilla Regan gives an overview of congressional debates around adopting a federal breach notification law during the 2000s. She notes that procedurally, the U.S. Congress faces higher barriers to enacting legislation than many states do because of its extensive committee structure, and various veto points within both Houses. Substantively, the Democratic and Republican parties had bitter disagreements about the extent to which regulators or companies themselves should maintain discretion in when a notification is necessary. At the time the piece was written (2009), Regan expressed some optimism that a federal law would be enacted, but as of 2020, no such legislation has passed yet (Regan, 2009).

Sara Needles, in "The Data Game: Learning to Love the State-Based Approach to Data Breach Notification Law," explicitly argues against adopting a federal standard. In her view, the fact that the states have quite easily adopted their own breach notification standards while the federal government has struggled is strong evidence that the current state-by-state approach is working well. Moreover, she notes that the various different intricacies in state law (what types of data are covered, how a breach notification should be worded etc.) will be difficult to reconcile in a federal law, and the result may be unsatisfactory (Needles, 2009).

Overall, the literature reaches a few important conclusions. First, there is tentative evidence that breach notification laws do reduce identity theft. Second, one mechanism by which this occurs, namely that consumers and investors punish breached firms, seems somewhat plausible in that there are negative effects on stock valuations in the immediate aftermath of a breach disclosure. However, the lack of enduring effects on stock prices may indicate that this market mechanism does not work well, particularly if there is insider trading. Finally, the debates surrounding breach notification span back to at least the early 2000s, but are perhaps receiving renewed attention because of recent incidents.

## 4 Data

Collecting data that on incidents of data breaches before and after the passage of breach notification law suffers from endogeneity issues. Prior to the enactment of a breach notification law, firms presumably have little incentive to report data breaches. Indeed, Romanosky et. al. showed that the number of reported breaches increases quite rapidly immediately after a state breach notification law is adopted, suggesting that there are a large number of unreported breaches that previously occurred. Assessing the effectiveness of breach notification laws by looking at the number of reported breaches would be a fruitless endeavor as the measurement of the outcome is confounded with treatment.

That being said, a plausible way forward is to use victimization rates, in particular with regards to identity theft. Since the late 1990s, the U.S. government has tracked identity theft through various law enforcement agencies, both at the federal and state levels. The gold standard for this data is the Consumer Sentinel, which is a Federal Trade Commission database that contains over 20 million self-reported identity theft complaints collected from a variety of different agencies and non-profit organizations. Unfortunately, the Consumer Sentinel is only available to law enforcement agencies. Some organizations and scholars have used Freedom of Information Act (FOIA) requests to retrieve some of this data. However, a recent case in the District Court for the District of Columbia (Ayuda v. FTC (2014)) ruled against a general FOIA request for detailed records. Because these data are no longer available, I instead use the Consumer Financial Bureau's (CFPB's) consumer complaint database. Consumers who have disputes with companies can file complaints to the CFPB, and the named companies are obligated to respond to the complaints. One such category of complaints is related to "identity theft," and within that, there are a number of "medical debt" complaints. Although there will necessarily be a good deal of undercoverage, the CFPB is a major contributor to the Consumer Sentinel, and the data contain rich information about the complaints. The data go back to 2013, which precludes studying the effects of breach notification laws passed between 2003 and 2008.

The CFPB database offers the raw text of a complaint, its category, subcategory, and state that it took place. I specifically subset to medical identity theft by using the "identity theft," "debt collection," and "medical" filters. In total, there are approximately 12,000 records. Figure 2 gives a sample of what these data look like, with a full example available here.

| Date Received | Product | Sub-Product | Company | State |
|---|---|---|---|---|
| 10/26/2015 | Debt collection | Medical | Collection Information Bureau, Inc. | FL |
| 10/12/2016 | Debt collection | Medical | Diversified Consultants, Inc. | FL |
| 8/15/2018 | Debt collection | Medical debt | Phoenix Financial Services LLC | TX |
| 8/22/2018 | Debt collection | Medical debt | Credence Resource Management, LLC | VA |
| 9/15/2016 | Debt collection | Medical | Rash Curtis and Associates | CA |

Figure 2: Sample CFPB Medical Identity Theft Data

Within the CFPB, I focus on medical identity theft primarily because doing so avoids Stable Unit Treatment Value Assumption (SUTVA) violations. Organizations with a multi-state presence may not create fifty different disclosures to comply with each state's individual laws. Rather, they will tend to default to the strictest disclosure law because of the California Effect. Some exceptions to this general pattern may exist. For instance, Massachusetts forbids organizations from disclosing "How it Happened." Outside of these idiosyncrasies though, it is common for requirements adopted in one state to leak into the notices given in other states, thus violating SUTVA. Thus focusing on identity theft as a whole is unlikely to yield a credible causal estimate.

Medical identity theft differs from other identity theft in that the institutions responsible for handling data - hospitals, health insurers, etc. - typically either do not have an interstate presence or localize protected health information data storage. Indeed, policymakers and health data stakeholders are more concerned with facilitating health data transfers across state lines than preventing such transfers, because these sorts of transfers are difficult to make under the current regulatory patchwork. The Center for Medicare & Medicaid Services (CMS) has baseline rules for how stakeholders (clinics, hospitals, pharmacies, etc.) in states can exchange information across state lines and with the federal government. The federal government is piloting "State Health Information Exchanges" where local, regional, and state governments harmonize their health data protocols to allow for easier transfer between organizations and across state boundaries. These programs may make state lines less meaningful for health data in the future, but their existence suggests that transferring medical records across those lines is quite difficult. While the inability to easily move data across state lines is a problem for medical service providers, patients, insurers, and governments, it does come with the advantage that state health privacy laws can be more plausibly analyzed without violating SUTVA.

## 5 Exploratory Data Analysis

Data breaches and medical identity theft have both grown in number over the last decade. Using data drawn from the HIPAA data breach portal maintained by the U.S. Department of Health and Human Services (HHS), I examine incidents of medical data breaches over time. U.S. federal law requires that breaches of unencrypted medical data be reported to HHS. These data thus provide a useful baseline for examining the growth in data breaches over time.

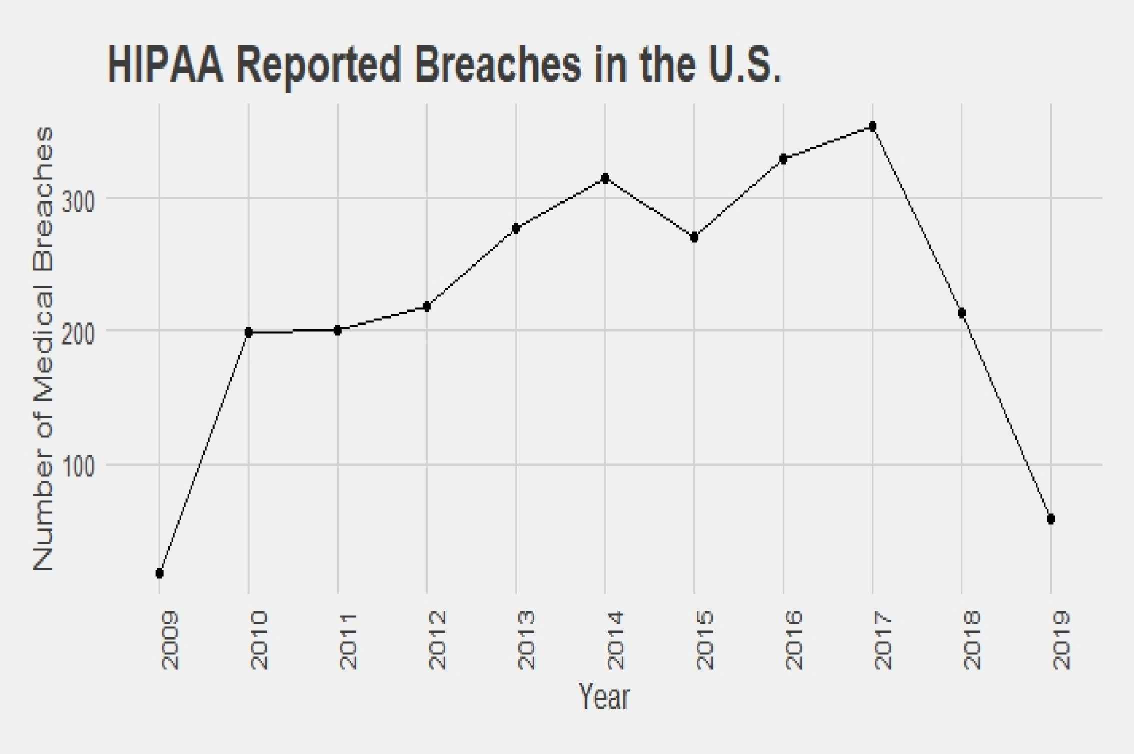

Figure 3: HIPAA Reported Breaches Over Time

In Figure 3, we see the growth in reported breaches over time. While the number does seem to be declining since 2018, these figures may be subject to lag as organizations may not discover breaches for months or years after the breach occurred. More organizations may also be encrypting their data, which would reduce the amount of unencrypted data that could be stolen. Broken down by state (see Figure 4), there are some clear patterns that emerge. California, Texas, and Florida lead the nation in reported data breaches, and several less populous states oftentimes do not report any breaches in some years.

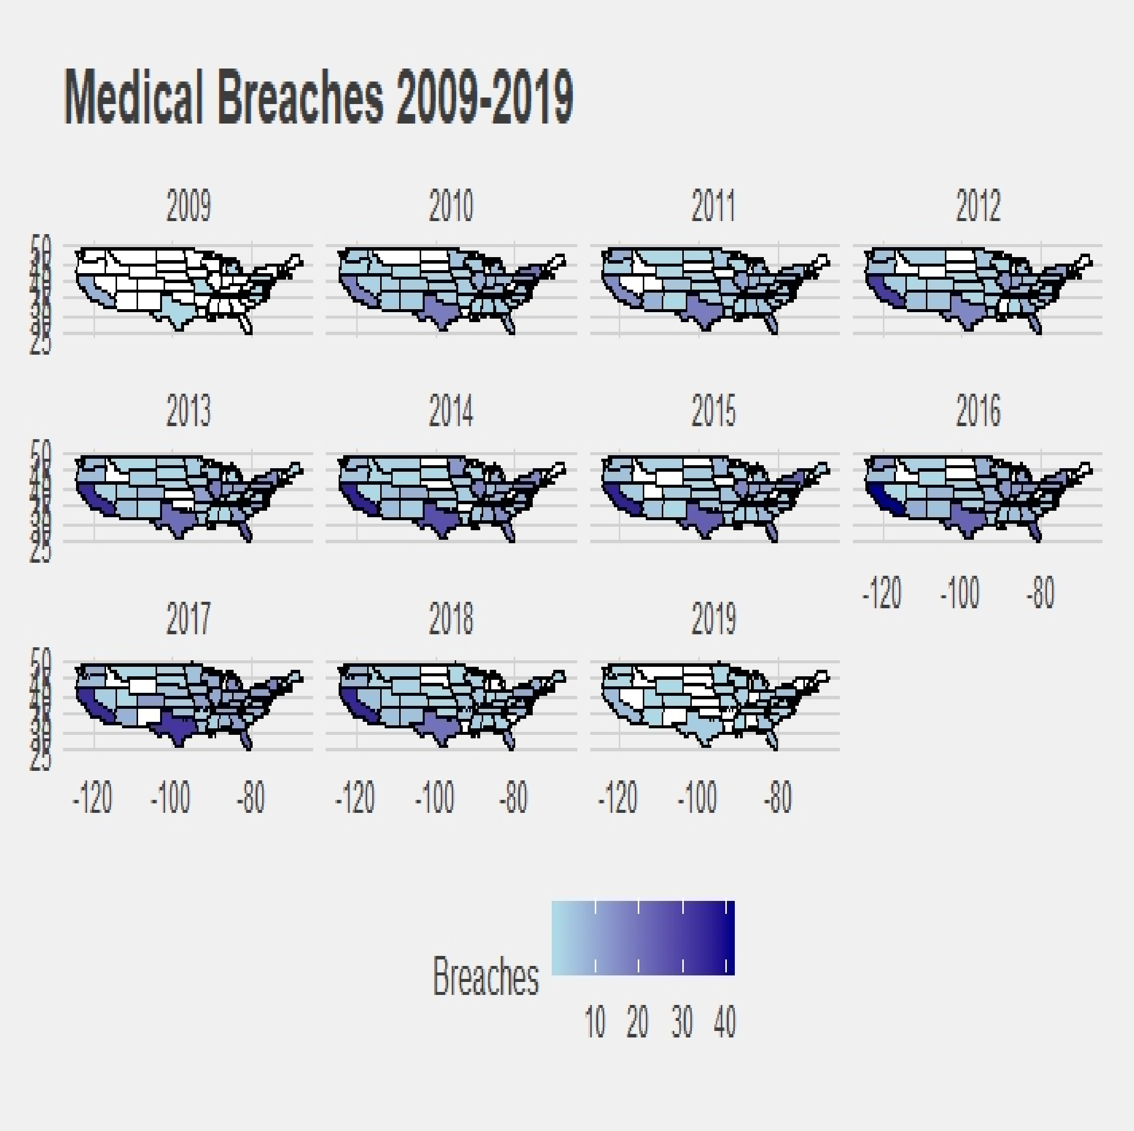

Figure 4: HIPAA Reported Breaches Over Time By State

Examining the sources of medical breaches, we can see that medical data breaches resemble data breaches more broadly. Hacking, theft, and accidental loss are the major categories contributing to data breaches. Interestingly, hacking seems to grow as a share of the overall number of breaches as time goes on. This may indicate that while incidents of employees stealing data or losing it continue in similar numbers, the number of incidents of malicious external attacks grow over time.

As breaches become more common, identity theft consequently rises as well. Figure 6 shows the number of medical identity theft reports to the CFPB over time. Again, reports to the CFPB represent a sample of the total extent of medical identity theft in the country. Victims of identity theft may never report this fact to the CFPB or law enforcement. They may report it to local law

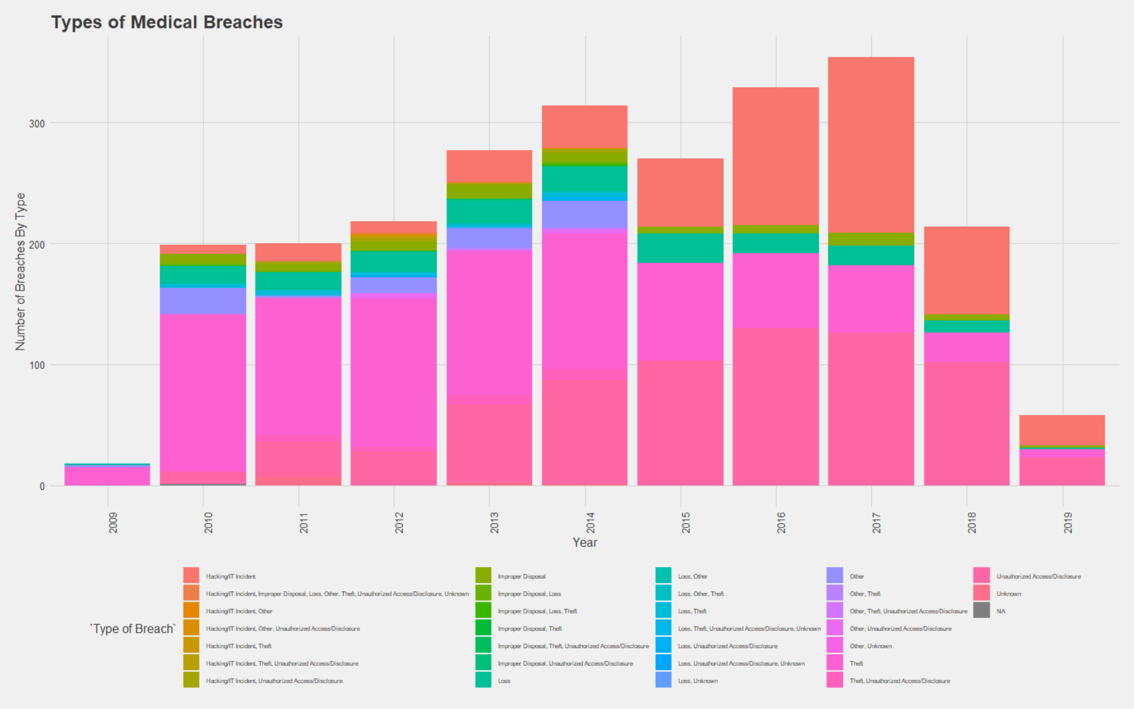

Figure 5: Types of HIPAA Reported Medical Breaches

enforcement, the Federal Trade Commission (FTC), state attorneys general, or credit reporting agencies instead of the CFPB. That being said, the amount of reported identity theft has grown each year, from fewer than 500 in 2013 to about 2500 in 2019.

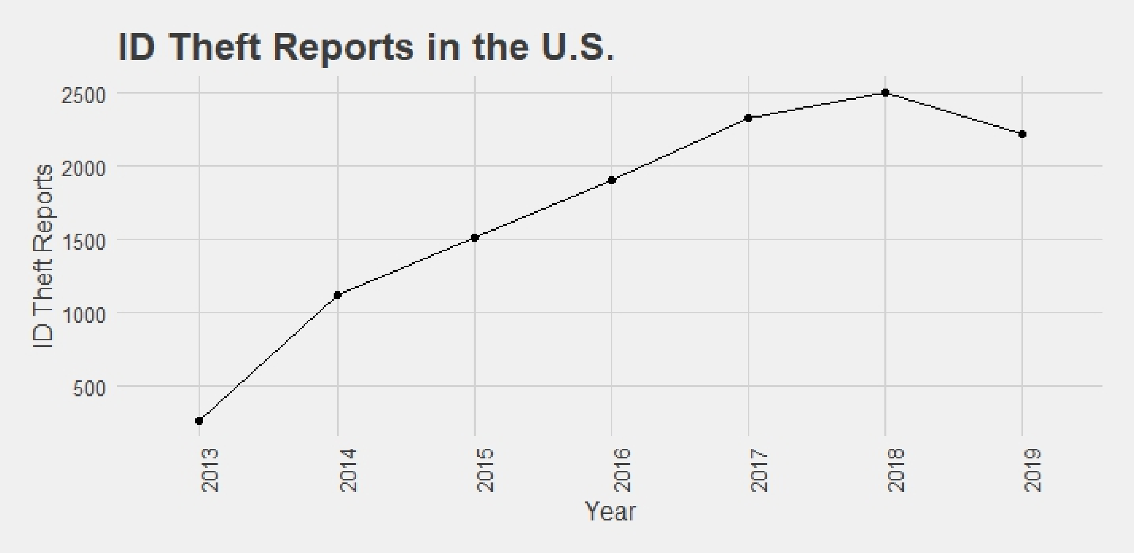

Figure 6: U.S. Identity Theft Over Time

Again, breaking down this information by state, some interesting trends emerge. Figure 7 shows the number of medical identity theft reports by state, scaled per 100,000 population (according to the 2010 U.S. Census). As with the HIPAA breach data, the most populous states unsurprisingly also have the most medical identity theft reports. Interestingly, this pattern holds even when scaling to identity theft theft per 100,000 people.

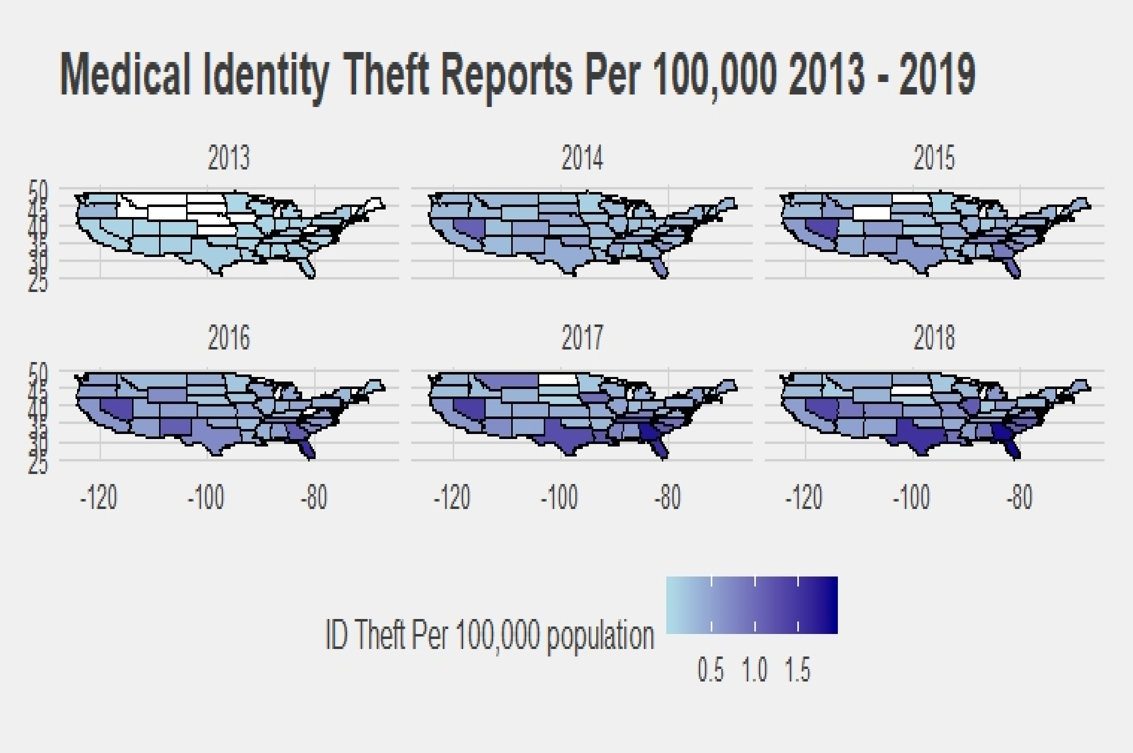
Medical Identity Theft Reports Per 100,000 2013 - 2019
2013
2014
2015
2016
2017
2018
50
45
40
35
30
25
-120
-100
-80
ID Theft Per 100,000 population
0.5
1.0
1.5

Figure 7: Medical Identity Theft by State Per 100,000 Population

Virtually every state has seen the number of reported thefts increase over time, but some regional patterns emerge as well. In particular, as seen in Figure 8, Florida and Texas each have around 400 reports in 2019, far more than the 200 or so that California, New York, and Illinois all have. Scaled by population (Figure 9), the South as a whole outpaces the rest of the country. Florida and Georgia each have victimization rates close to 2 reports per 100,000 people, compared to California, New York, and Illinois with closer to .5 reports each. One possible explanation for this trend could be Florida's large share of senior citizens, and therefore Medicare recipients, however this hypothesis is undercut by Georgia's similarly high victimization rates and low proportion of senior citizens. It is also possible that people in Southern states are more likely to report medical identity theft than people in other states, however the theoretical reason why this might be the case is unclear.

These regional patterns are important primarily because they suggest that there are genuine state-level differences. Some Southern states were relatively late adopters of data breach notification laws (for example, Alabama was the last state to enact one in 2018), however this does not tell the whole story as many were early adopters and their laws shared nearly the exact same provisions as laws in other states.

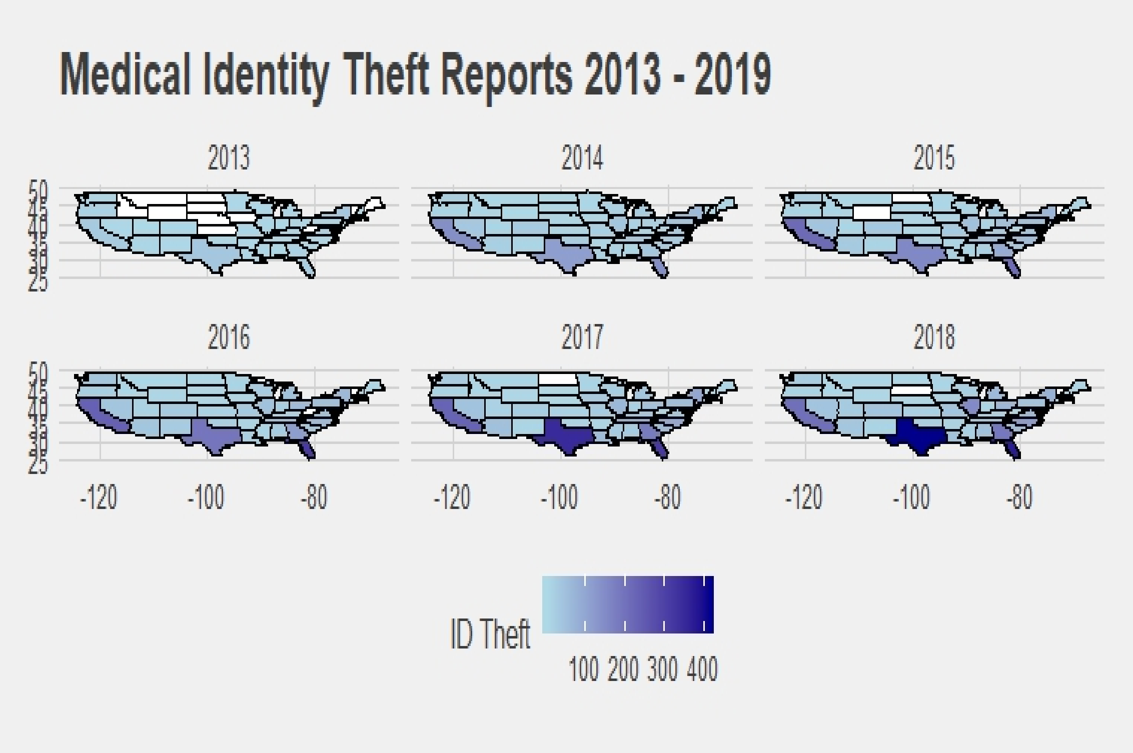

Figure 8: Medical Identity Theft by State

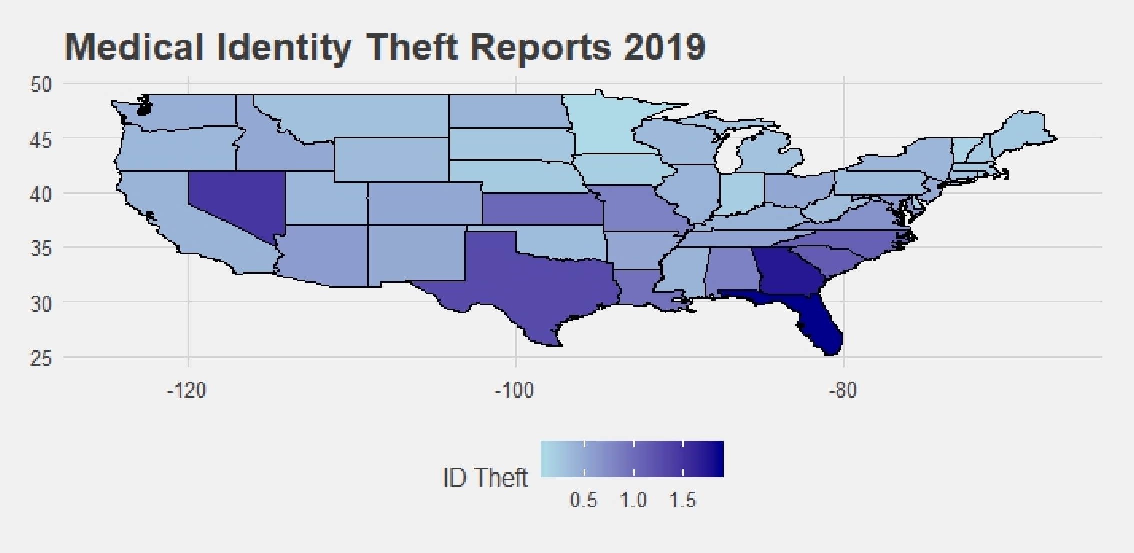

Figure 9: 2019 Medical Identity Theft Reports per 100,000 Population

## 6 Identification & Methodology

I specifically ask what is the effect of California's 2016 breach notification requirements in a data breach notification law on rates of medical identity theft. By 2008, 47 states already adopted some form of a breach notification law that covered businesses, thus making data generated after that time unhelpful in assessing a treatment effect. Also in 2008, California explicitly included health information under its definition of personal information. The U.S. Department of Health and Human Services (HHS) also has a long-standing regulation that requires processors of medical data to report breaches of unencrypted data. California's 2016 amendments expanded the notice requirements to require that they follow a particular format, be posted on an organization's website, and follow updated definitions of personal information. Thus, the mechanism that I am examining is whether a clear disclosure affects reported identity theft. Theoretically, data subjects who become aware of a breach following a disclosure may take precautions to safeguard their identities, and organizations may also be more careful about their cybersecurity practices if they know that data subjects will have such clear information about how incidents occurred.

To address this question, I employ an augmented synthetic control. Synthetic control was introduced by Abadie and Gardeazabal (Abadie and Gardeazabal, 2003) in 2003 where the authors studied the effect of terrorist conflict on the Basque region's GDP. It was further explored by Abadie, Diamond, and Hainmueller in 2010 where the authors examined the effect of California's cigarette tax on cigarette consumption (Abadie et al., 2010). The method is useful for comparative analyses, particularly when evaluating policies at the level of a state or country where there are relatively few units in the dataset (Athey and Imbens, 2017). The basic logic underlying synthetic control is that a real-world "treated" unit is compared to a synthetic control of itself. The synthetic control is created by borrowing covariates from other units. The key to creating a successful synthetic control is making sure that the synthetic unit matches the real world observed unit in the pre-treatment period. Once a successful synthetic unit is created, it is straightforward to calculate the treatment effect by subtracting the post-treatment control outcome from the post-treatment treatment outcome. It is essentially an extension of the difference-in-difference method that allows for the comparison of the treated unit against a hypothetical controlled version of itself, rather than an observed control unit.

Formally, the authors motivate the model as follows. Imagine that there $J + 1$ regions of interest. Allow $Y_{it}^N$ to be the outcome for region $i$ at time $t$ for units $i = [1 : J + 1]$ and time periods $t = [1 : T]$. Let $T_0$ be the number of preintervention periods, with $1 \leq T_0 < T$. Let $Y_i^I t$ be the outcome if unit $i$ at time $t$ was exposed to the intervention.

In this case, we wish to estimate the effect of the updated data breach notification law on California's medical identity theft rates. This relationship can be expressed as:

$$ATT = Y^I - Y^N$$

Where $ATT_t$ is the treatment effect[2], $Y^I$ is California's medical identity theft rates with the law, and $Y^N$ is California's medical idenitity theft rates without the law. This setup is analogous to the Neyman-Rubin potential outcomes framework. The fundamental problem of causal inference is that we can never observe the same unit under both treatment and control conditions (Rubin,

[2] ATT = Average Treatment Effect on the Treated. ATT is used in applications like difference-in-differences and synthetic control because the effect is being estimated for units that received treatment. This concept can be distinguished from Average Treatment Effect (ATE) which is an estimate of a treatment effect in a randomized control trial.

1974). Addressing this problem requires estimating the potential outcome for a unit under the counterfactual condition (i.e. estimating the value under treatment for a control unit, or estimating the value under control for a treated unit). The synthetic control method handles the fundamental problem of causal inference by creating an estimate of a counterfactual treated unit through a weighted combination of other untreated units.

One issue with the synthetic control method is that it is only valid when the synthetic unit matches the observed unit in the pre-treatment period. Ben-Michael, Feller, and Rothstein proposed an augmented synthetic control method that offers bias correction tools in situations where such pre-treatment matching is infeasible (Ben-Michael et al., 0). They propose using an outcome model to estimate bias in the pre-treatment fit of the synthetic control, and then debias the original synthetic control estimate. The authors specifically recommend using a ridge regression, and also provide random forest and matrix completion methods (Athey et al., 2017), among others [3]. Augmented synthetic control essentially has all of the transparency advantages of a standard synthetic control, but provides additional options in situations where perfect pre-treatment fits are not possible.

Extending the basic framework to this problem, imagine there are a set of states $i \epsilon S = 1:50$, and a set of time periods, $t \epsilon T$. The general problem of estimating the treatment effect for a unit, $i$ at time $t$ can be described as:

$$ATT = Y_{it}^{I} - Y_{it}^{N}$$

In this specific instance, assume that California is unit 1. Thus for each time period, the estimate is:

$$ATT = Y_{1t}^{I} - Y_{1t}^{N}$$

Where:

- $ATT$ = Average Treatment Effect on the Treated. Reduction in medical identity theft rates per 100,000 people.
- $Y_{1t}^{I}$ = Observed medical identity theft rate in California
- $Y_{1t}^{N}$ = Synthetic estimate of medical identity theft rate in California

Where $Y_{1t}^{N}$ is estimated by constructing a synthetic control. For an overall treatment effect, I average the ATTs across each post-treatment time period. I also fit various models, both with and without augmentations.

[3] Athey et. al. argue that "The [Neyman-Rubin] unconfoundedness approach estimates patterns over time that are assumed to be stable across units, and the synthetic control approach estimates patterns across units that are assumed to be stable over time." These different assumptions impose different restrictions on the missingness of the outcome. The authors suggest matrix completion methods that use regularization to estimate missing data, and relax the assumptions of either the uncoufnoundedness or synthetic control approaches that are popular in econometrics.

# 7 Results

## 7.1 Difference-in-Differences Baseline

I start with a demonstration of a simple difference-in-differences estimate that compares California to the U.S. average (minus California) for medical identity theft rates. Figure 10 shows the synthetic California estimate, the observed California, and the U.S. average (by state) for medical identity theft reports. The U.S. average does approximate California fairly well in the pre-treatment period, it tends to underestimate rates in California, and therefore is not an ideal comparison unit in the diff-in-diff framework. The synthetic control, which is a weighted sum of U.S. states, lessens this problem somewhat and tends to match California in the pre-treatment period, particularly in the periods immediately preceding treatment.

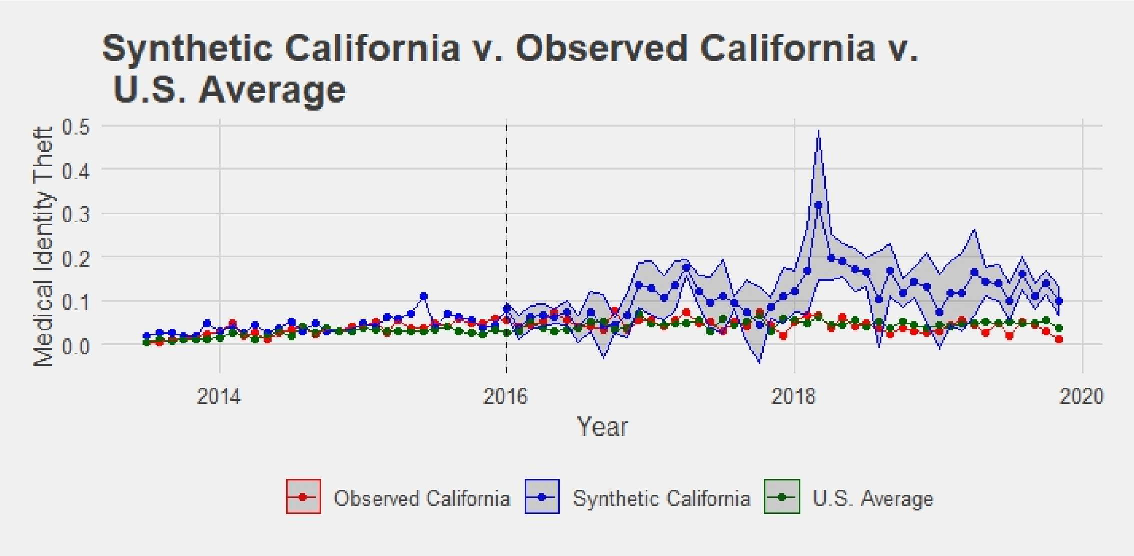

Figure 10: Synthetic California, Observed California, and the U.S. Average

Figure 11 illustrates the U.S. average differing from California rates in the pre-treatment period more clearly. Diff-in-diff relies on the "parallel trends" assumption that requires that the difference between the treated and control units is constant over time. The U.S. average tends to underestimate California, but sometimes exceeds it, and varies from time period to time period.

## 7.2 No Augmentation

Next, I show synthetic control results without augmentation. Augmenting a synthetic control is mainly useful in situations with poor pre-treatment fit. In this case, the fit is still fairly good without augmentation. Figure 12 shows the standard synthetic control estimate with no augmentations or additional covariates.

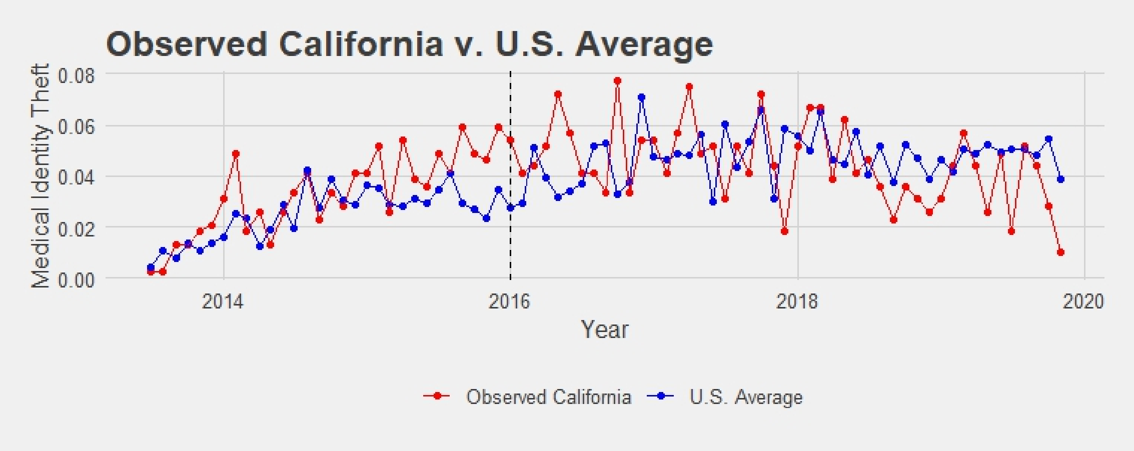

Figure 11: Difference in Differences Observed California v. U.S. Average

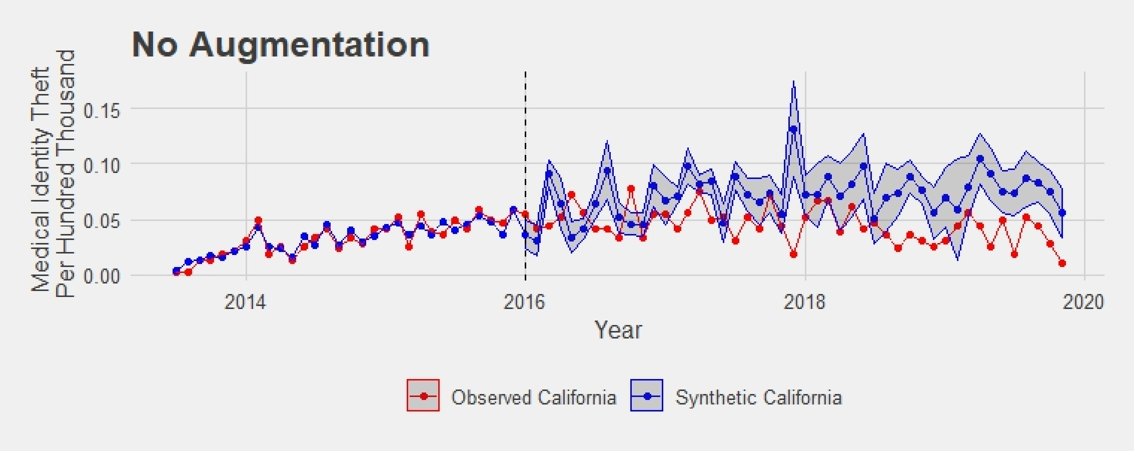

Figure 12: Standard Synthetic Control

## 7.3 Ridge Augmentation

Moving to the augmented estimates, I present results from the Ridge-augmented synthetic control. Preliminary results suggest that there is a small effect of these expanded notice standards on reported medical identity theft. I employ a Ridge-augmented synthetic control on reported medical identity theft. I provide estimates both per year and per month, and estimate total number of reports and reports per 100,000 people. The synthetic control also adds number of HIPAA reported breaches and the number of individuals affected by medical breaches in each state in a given year as additional covariates.

Figure 14 shows synthetic control estimates broken down by month. Table 1 shows the relative contributions of each state to the synthetic control, and Figure 13 visualizes these weights. The advantage of using synthetic control over the U.S. average is that the synthetic control attaches weights to each state so construct a more appropriate control unit. Scaled to victimization rates per 100,000 people, the average treatment effect on the treated is approximately .065 fewer reports of medical identity theft per 100,000 people (see Figure 14). The effect size grows over time, with the final estimate being close to .1 fewer reports per 100,000 people. This suggests that expanded disclosure requirements have a modest effect that potentially grows over time, though caution is advised against extrapolating too far into the future with relatively few pre-treatment periods.

| | State | Weight |
|---|---|---|
| 1 | AL | -0.03 |
| 2 | AR | -0.05 |
| 3 | AZ | -0.00 |
| 4 | CO | 0.00 |
| 5 | CT | -0.05 |
| 6 | DC | -0.01 |
| 7 | DE | -0.08 |
| 8 | FL | 0.20 |
| 9 | GA | 0.08 |
| 10 | IA | -0.05 |
| 11 | ID | -0.02 |
| 12 | IL | 0.66 |
| 13 | IN | -0.02 |
| 14 | KS | -0.04 |
| 15 | KY | 0.01 |
| 16 | LA | -0.03 |
| 17 | MA | 0.01 |
| 18 | MD | 0.01 |
| 19 | ME | -0.04 |
| 20 | MI | 0.01 |
| 21 | MN | 0.00 |
| 22 | MO | 0.00 |
| 23 | MS | -0.06 |
| 24 | MT | -0.00 |
| 25 | NC | 0.04 |
| 26 | ND | -0.05 |
| 27 | NE | -0.05 |
| 28 | NH | -0.05 |
| 29 | NJ | -0.00 |
| 30 | NM | -0.04 |
| 31 | NV | -0.06 |
| 32 | NY | 0.21 |
| 33 | OH | 0.06 |
| 34 | OK | -0.03 |
| 35 | OR | -0.03 |
| 36 | PA | 0.09 |

| | | |
|---|---|---|
| 37 | RI | -0.04 |
| 38 | SC | -0.03 |
| 39 | SD | -0.08 |
| 40 | TN | 0.03 |
| 41 | TX | 0.64 |
| 42 | UT | -0.04 |
| 43 | VA | -0.00 |
| 44 | VT | -0.06 |
| 45 | WA | -0.01 |
| 46 | WI | -0.05 |
| 47 | WV | -0.05 |
| 48 | WY | 0.06 |

Table 1: Weights Generated by Synthetic Control

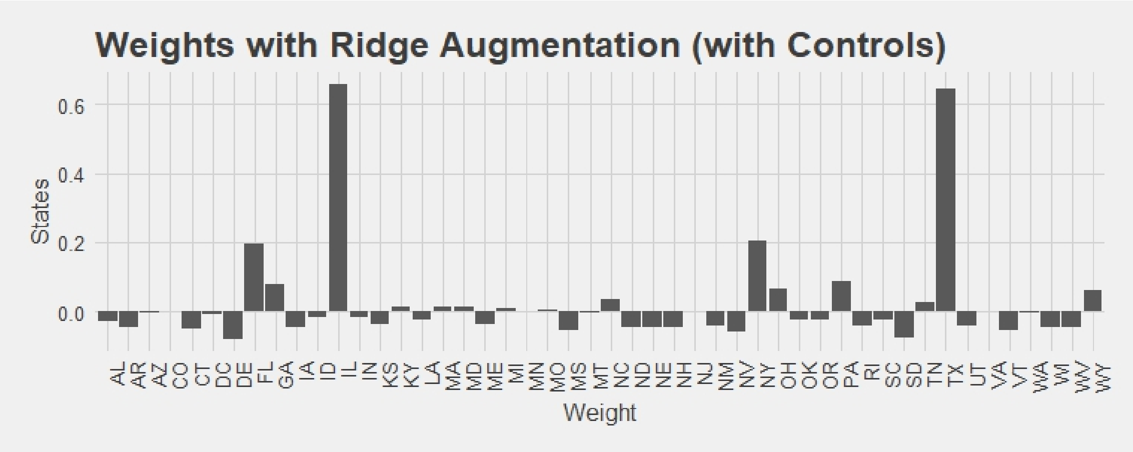

Figure 13: Barplot of Ridge Augmented Weights

The L2 imbalance (square root of the sum of the sum of squared vector values) is .11, with an average estimated bias of .05, and an average ATT estimate of -.069. The estimated treatment effect for November 2019 is about .1 fewer reports per 100,000 people. However, because there are so few reports in any given month, the estimates and observed values can be quite noisy. Thus, we should be careful about interpreting a large treatment effect from these estimates. That being said, pre-treatment fit is generally good regardless of the chosen model. Figure 15 shows synthetic control estimates across no augmentation, ridge, matrix completion, and gsynth (linear factor model). In each case, the overall pre-treatment fit is similar, as are the estimates.

The expanded breach notification requirements does imply a modest effect on medical identity theft reports. Both the synthetic estimates and the observed values of medical identity theft reports represent lower bounds. The CFPB medical identity theft reports are (likely unrepresentative) a sample of all reported identity theft reports, and not all identity theft is reported at all. More data

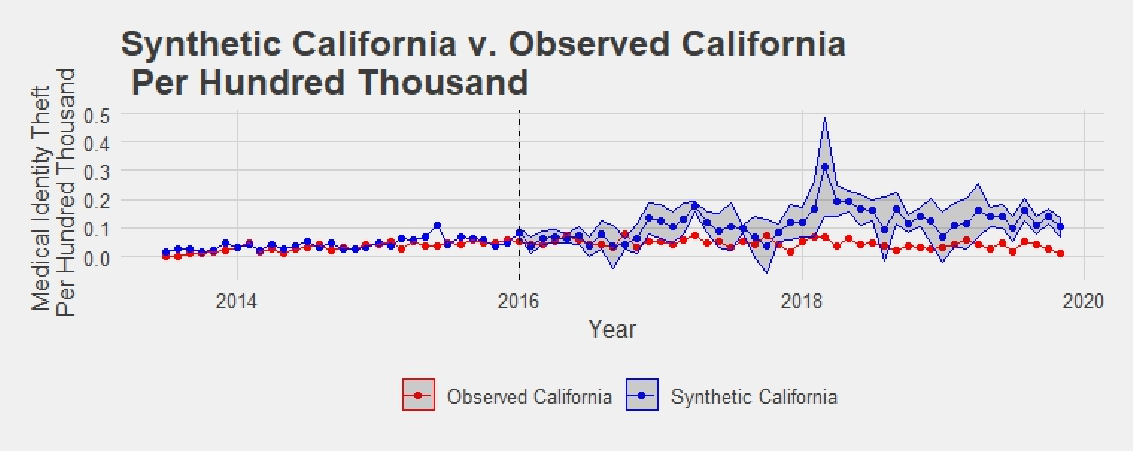

Figure 14: Estimates by Month

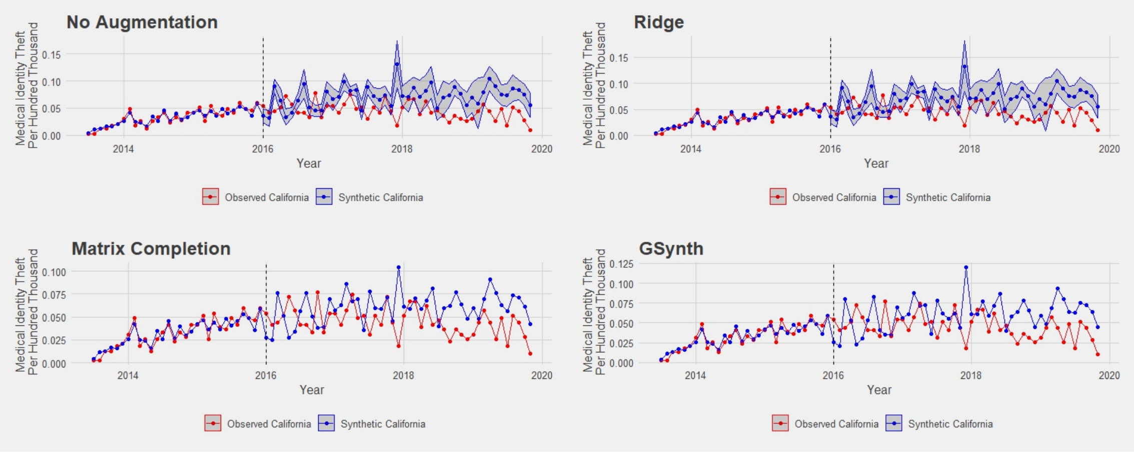

Figure 15: Augmented Synthetic Control Across Outcome Models

on identity theft reports would be helpful for reducing the noise in the estimates, particularly at the monthly level.

One main takeaway from these results is that the effect of the law grows over time. Figure 16 shows the estimated treatment effect on medical identity theft reports per 100,000 people per month. In the months immediately following the law going into effect, there is little change from a null effect, and matches the pre-treatment outcomes. Between January 2016 and November 2019 however, there is downward movement. Again, this pattern holds regardless of the specifications chosen, and suggests that the law has a modest but real effect on reported medical identity theft. Table 2 shows the results from various specifications, both with controls for state medical infrastructure and without.

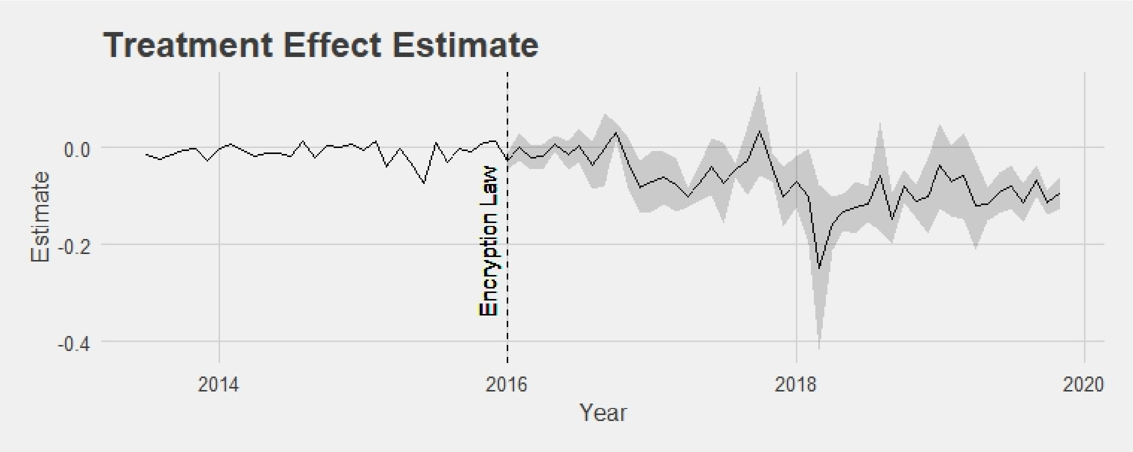

Figure 16: Average Treatment Effect on The Treated Over Time

| | Outcome Model | L2 Imbalance | Average ATT |
|---|---|---|---|
| 1 | None | 0.03 | -0.03 |
| 2 | None With Controls | 0.03 | -0.02 |
| 3 | Ridge | 0.03 | -0.03 |
| 4 | Ridge With Controls | 0.11 | -0.07 |
| 5 | Matrix Completion | 0.03 | -0.01 |
| 6 | GSynth | 0.03 | -0.02 |

Table 2: Outcome Models with L2 Imbalances and Average ATT

Otherwise, a question worth probing further based on these results is whether previous estimates of the effects of data breach notification laws potentially understated the magnitude of these effects. Previous work that estimated these effects using a panel regression found an effect of about 2% on reported identity theft (Romanosky et al., 2011). This work evaluated laws from the mid-2000s, and looked at identity theft rates across all kinds of identity theft. These estimates potentially understate the effect of state laws on state identity theft rates because of leakage of notifications across state lines for commercial breaches, underinclusiveness in the definition of "personal information" in the mid-2000s wave of laws, and omitted variable bias in the regression specification. In contrast, the synthetic controls imply an average treatment effect on the treated that corresponds to between a 2/100,000 a 7/100,000 decline in reported identity theft in California. In the final periods, the estimated percentage effect is closer to 40% [4]. Again, there are many potential sources of bias in the synthetic estimates as well, and smaller and null treatment effects are within one standard error of the point estimates in the periods immediately following adoption of treatment. That being said, the synthetic control estimate suggests that data breach notification may be more effective than previously thought.

## 8 Policy Discussion & Future Work

Policymakers are increasingly paying attention to privacy and cybercrime issues, and the data breach notification law remains the most popular and widespread tool used by U.S. states, federal agencies, and the European Union (EU). Despite its prevalence, there is little evidence about its efficacy, especially in recent years. As states rapidly and frequently adopt and update their data breach notification laws, understanding their effects will be of paramount importance.

One ongoing debate in privacy law is whether disclosure is an effective regulatory mechanism. Beyond data breach notification, governments are actively creating various notification requirements pertaining to individuals' privacy. The EU's GDPR contains several provisions that require companies to make their privacy agreements visually appealing and intuitive. The California Consumer Protection Act (CCPA) requires that companies that collect consumer information disclose what information is being collected about them, and whether it is sold. The theory underlying all of these regulations is that disclosures will help consumers control their information, and make rational decisions about their market participation. Both the GDPR and CCPA are recent developments, but indicate that policymakers continue to look to disclosure as the best regulatory option in this space.

This study provides an empirical examination of whether disclosure works. In the 15 years since California implemented the first data breach notification law in the U.S., every state has adopted some version of one. Previous attempts to study these laws were hampered by lack of access to data about the number of breaches and identity theft reports. By using CFPB data, this study overcomes some of those previous challenges. That being said, more comprehensive and publicly available data on identity theft reports would enhance researchers' ability to empirically answer important questions about privacy law and policy.

With regards to implications for privacy law, these results tentatively suggest that California's data breach notification updates in 2016 had an impact on the reported medical identity theft in the years after its adoption. There are a few important pieces to note here before generalizing to

[4] For the most part, I avoid expressing the effect in these percentage terms because of the low baseline rates of reported identity theft. Even a reduction of reports on the order of dozens or hundreds can have a sizable percentage effect, even in a large state like California. To effectively compare to previous literature I provide the percentage effect here, but with the understanding that it is highly susceptible to swing because of low base rates.

all data breach notification laws in all states across time. First, California already had a fairly strong data breach notification law in place, with a fairly expansive definition of "personal information," requirements that breached organizations notify consumers and the Attorney General, and penalties for failure to comply. The 2016 amendments required that notices use a particular format, provide clear information, and be labeled clearly. Thus, the 2016 amendments were more focused on the style and substance of the disclosure, rather than changing the types of disclosures that needed to be made. These results therefore point to the effect of mandating that disclosures look a particular way, not the effect of a generic data breach notification law. Moreover, the results may not generalized well beyond California; the exploratory data analysis showed that there are clear state and regional patterns in medical breaches and identity theft, so another state that adopts California's requirements may not enjoy the same benefits. Medical identity theft may also be different from other kinds of identity theft, and the law could be especially good or bad at deterring that category of cybercrime and not others. Keeping this caveats in mind though, the results suggest that disclosure does matter, and more importantly, that clear, well-designed disclosures matter.

Data breach notification laws continue to evolve, and these changes should provide researchers with ample opportunities to study the effects of various aspects of disclosure regulation. For example, one sources of variation between state laws is the presence or absence of private causes of action following a disclosure. Some states allow individuals to sue organizations following a breach notification, while others only allow the Attorney General to make that determination. Various states have different rules regarding who must be notified (consumers, attorneys general, and/or credit reporting agencies). Exceptions to breach notification requirements when data is encrypted are being reexamined. Differences in requirements for "likelihood of harm" analysis may also produce divergent outcomes. While breach notification laws share many similarities, these differences could provide a rich set of questions for more empirical work.

## 9 Conclusion

Data breach notification will likely continue to be a popular tool for policymakers regulating cybercrime, thus making evidence of how well the current regime works important for future policy decisions. Legislators and regulators are actively debating whether disclosure is an effective mechanism for protecting consumers without implementing heavy-handed market interventions. Quantifying the harms stemming from privacy invasions is a notoriously difficult problem, making it difficult for policymakers to know which policy levers to pull. Estimating the potential effect of disclosure on identity theft is a first step in understanding whether data breach notification laws are the most effective tools for protecting privacy. Using an augmented synthetic control approach, I estimate the effect of California's 2016 amendments to its breach notification law. These results tentatively suggest that breach notification does reduce identity theft, but more work is needed to get a complete picture.

# References

Abadie, A., A. Diamond, and J. Hainmueller (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American Statistical Association.*

Abadie, A. and J. Gardeazabal (2003). The Economic Costs of Conflict: A Case Study of the Basque Country. *American Economic Review 93(1)*, 113–132.

Athey, S., M. Bayati, N. Doudchenko, G. Imbens, and K. Khosravi (2017). Matrix Completion Methods for Causal Panel Data. *arXiv1710.10251v2.*

Athey, S. and G. W. Imbens (2017). The State of Applied Econometrics: Causality and Policy Evaluation. *The Journal of Economic Perspectives 31(2)*, 3–32.

Ben-Michael, E., A. Feller, and J. Rothstein (0). The Augmented Synthetic Control Method. *arXiv:1811.04170.*

Bradford, A. (2020). *The Brussels Effect.*

Cardenas, A., S. Radosavac, J. Grossklags, J. Chuang, and C. J. Hoofnagle (2010, September). An Economic Map of Cybercrime.

Cooter, R. and T. Ulen (2016). *Law and Economics, 6th Edition.* Berkeley Law Books.

Goel, S. and H. A. Shawky (2009, October). Estimating the Market Impact of Security Breach Announcements on Firm Values. *Information & Management*, 404–410.

Goel, S. and H. A. Shawky (2014, January). The Impact of Federal and State Notification Laws on Security Breach Announcements . *Communications of the Association for Information Systems 34.*

Hoofnagle, C. J. (2007). Identity Theft: Making the Known Unknowns Known. *Harvard Journal of Law & Technology 21*(1), 98–122.

Hoofnagle, C. J. (2008). Towards a Market for Bank Safety. *Loy. Consumer L. Rev. 21.*

Mitts, J. and E. Talley (2018). Informed Trading and Cybersecurity Breaches. *Harvard Business Review.*

Needles, S. A. (2009). The Data Game: Learning to Love the State-Based Approach to Data Breach Notification Law,. *North Carolina Law Review.*

Regan, P. M. (2009). Federal Security Breach Notifications: Politics and Approaches.

Romanosky, S., R. Telang, and A. Acquisti (2011). Do Data Breach Disclosure Laws Reduce Identity Theft? *Journal of Policy Analysis and Management 30*(2), 256–286.

Rubin, D. (1974). Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies. *Journal of Educational Psychology* (75(371)), 591.