

Prior Fraud Exposure and Precautionary Credit Market Behavior

Nathan Blascak and Ying Lei Toh

November 2022

RWP 22-14

<http://doi.org/10.18651/RWP2022-14>

FEDERAL RESERVE BANK *of* KANSAS CITY



Prior Fraud Exposure and Precautionary Credit Market Behavior*

Nathan Blascak[†] Ying Lei Toh[‡]

October 2022

Abstract

We study how past experiences with privacy shocks affect individuals' likelihood to take precautionary behavior when faced with a new privacy shock in the context of credit markets. We focus on experiences with identity theft and data breaches, two kinds of privacy shocks that either directly lead to fraud or put an individual at an elevated risk of experiencing fraud. We show that immediately after the announcement of the 2017 Equifax data breach, individuals with either kind of prior fraud exposure were more likely to freeze their credit report and close credit card accounts than individuals with no prior exposure. We also find that prior victims of identity theft, a more serious type of exposure, were more likely to take precautionary actions than individuals who were victims of a previous data breach.

Keywords: Equifax data breach, consumer credit, credit freeze

JEL Classification: D14, D18, G50

*The authors would like to thank Alessandro Acquisti, Julia Cheney, Fumiko Hayashi, Bob Hunt, Thanh Lu, the anonymous reviewers for WEIS, and seminar participants at the Boulder CFDM conference, the Consumer Finance Round Robin, WEAI Conference, the Federal Reserve System Applied Micro Conference, and the Federal Reserve Bank of Kansas City for their helpful comments. Ian McGroarty and Bryant Wright provided excellent research assistance.

[†]Consumer Finance Institute, Federal Reserve Bank of Philadelphia (email: nathan.blascak@phil.frb.org)

[‡]Federal Reserve Bank of Kansas City (email: yinglei.toh@kc.frb.org)

Disclaimer: This Philadelphia Fed working paper represents preliminary research that is being circulated for discussion purposes. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia, the Federal Reserve Bank of Kansas City, or the Federal Reserve System. Any errors or omissions are the responsibility of the authors. No statements here should be treated as legal advice. Philadelphia Fed working papers are free to download at <https://philadelphiafed.org/research-and-data/publications/working-papers>.

1 Introduction

Privacy shocks such as data breaches and identity (ID) theft are increasingly common occurrences that expose an individual’s personal identifying information (PII) to outside parties and increase the risk of fraud victimization and the misuse of PII for affected individuals.¹² However, previous studies show that despite the increased fraud risk, many individuals affected by these kinds of events do not take steps to mitigate their risk exposure, such as enrolling in credit monitoring or ID theft protection services, and those who do take action generally only do so after victimization has occurred ((Ablon et al., 2016; Ponemon Institute, 2014; Brodtkin, 2007; Zou et al., 2018; Romanosky et al., 2011)). These findings suggest that (1) prior to experiencing a privacy shock, individuals may perceive the net benefit of taking precautions to be negative because the cost of taking precautions is relatively high or the expected losses from subsequent fraud victimization are relatively low and (2) previous experience with fraud or an increase in the risk of fraud (collectively “fraud exposure”) may increase the net benefit of taking precautions in response to a new privacy shock.

Motivated by these facts, we investigate how individuals’ past experiences with fraud exposure affects their likelihood of taking precautionary action in response to a new negative shock to their PII. Based on the findings from the previous literature, we hypothesize that individuals with different types of prior fraud experiences will have different likelihoods of taking precautionary actions when their PII is exposed. To test this hypothesis, we leverage the 2017 Equifax data breach, which unexpectedly exposed the sensitive personal information of 147 million U.S. consumers, over 70% of the U.S. adult population. We consider the adoption of two fraud prevention measures in the credit market following the breach: closing credit card accounts and placing a credit report freeze. These measures, respectively, help prevent criminals from misusing breach victims’ existing financial accounts or fraudulently opening new accounts under victims’ names, which are the most common types of fraud (Harell, 2021).

To guide our empirical analysis, we first develop a descriptive model of individual precautionary behavior in the credit market. The model incorporates insights from the extant literature that suggest that learning and salience/availability are important factors that determine if and how an individual may respond to a new breach. We classify individuals

¹Identity theft has multiple statutory definitions in the U.S., which include transferring/using another individual’s data with the intent to commit fraud and actual fraud using another individual’s information without permission. For more information, see Cheney (2005).

²In 2020 alone, U.S. consumers filed over 2.3 million reports of identity (ID) theft or fraud with the Federal Trade Commission’s Consumer Sentinel Network, and 34% of those victims reported losses. For more information, see <https://public.tableau.com/app/profile/federal.trade.commission/viz/ConsumerSentinel/Infographic>.

into three groups, based on their past fraud experiences: prior identity theft (fraud) victims, prior data breach victims (who did not experience fraud), and non-victims (who have never been exposed to fraud or data breaches). We hypothesize that prior fraud victims are the ones most likely to take precautionary actions, as they likely have the lowest cost of taking action — they learned from dealing with fraud in the past (learning effect) — and the highest perceived fraud risks — they can easily imagine fraud happening to them again since it has occurred to them before (positive availability effect). For prior breach victims, we are not able to formulate a testable hypothesis because it is *a priori* unclear whether they would be more likely to adopt precautionary measures than non-victims. On the one hand, prior breach victims may have learned about precautionary measures available to them when their data were exposed, lowering the cost of taking precautions (i.e., a positive learning effect). On the other hand, prior breach victims may place a low likelihood of fraud happening to them after a new data breach if they did not experience fraud during the prior breach, lowering the expected benefit of taking action (i.e., a negative availability effect). Prior breach victims are more likely to take precautions than non-victims if the learning effect outweighs the availability effect.

To empirically examine if individuals with prior fraud exposure adopt precautionary credit market measures in response to the announcement of a new data breach, we utilize a large nationally representative data set of anonymized credit bureau records that also contains information on credit report freezes and fraud flags. We use the detailed nature of this data to empirically identify the three groups of individuals we classify in our descriptive model: individuals who were previous victims of identity theft, individuals who were likely to have been victims of a prior data breach (the 2015 Anthem data breach), and individuals who did not experience either event (non-victims).

With this data, we estimate a series of difference-in-differences (DID) regression models to compare how each of our prior victimization groups responded to the Equifax data breach announcement relative to the non-victims. Consistent with our hypothesis, we find that prior fraud victims are 0.1 to 0.25 percentage points (relative 25%) more likely to place a credit freeze than non-victims but do not close more card accounts than non-victims. For prior breach victims, we find that they are 0.1-0.3 percentage points (relative 30%) more likely to place a credit freeze than non-victims and actually close *fewer* card accounts than non-victims in the quarters after the Equifax breach announcement. For prior breach victims, this result is consistent with the learning effect for credit freezes from exposure to the Anthem breach outweighing the availability effect of the non-occurrence of fraud, while the availability effect dominates the learning effect for the closure of credit card accounts.

Finally, we compare the response of prior ID theft victims to prior data breach victims to

the Equifax breach announcement. We find that although past ID theft victims placed fewer new freezes than prior breach victims in the quarter of the breach announcement, prior ID theft victims placed more new freezes overall (from the second quarter following the breach announcement until the end of our sample). We also find that prior ID theft victims closed on average 7.5% more credit card accounts compared to Anthem breach victims, though the coefficients are noisily estimated.

Our work contributes to the literature on consumer responses to data breaches. Prior research has found that some individuals adopt precautionary measures after falling victim to a data breach. For instance, Turjeman & Feinberg (2019) found that the victims of a prominent dating site data breach took steps, such as removing photos from their profile pages, to protect their identities in the aftermath of the breach. Specific to the context of credit markets, Mikhed & Vogan (2018) examined individuals' precautionary behavior following the 2016 South Carolina Department of Revenue data breach and found that residents of South Carolina were briefly more likely to sign up for precautionary measures, such as initial fraud alerts and credit freezes, compared to residents of neighboring states.

Researchers have also identified several factors that may affect individuals' response to a data breach. These factors include the effectiveness with which the breached entity communicates the breach and available protective measures (Zou et al., 2019), individuals' knowledge of available protective measures (Zou et al., 2018), the presence of overlapping protective measures (Zou et al., 2018; Ablon et al., 2016), the cost of adopting these measures (Zou et al., 2018; Romanosky et al., 2011), and behavioral factors, such as the underestimation of the probability of falling victim to fraud (optimism bias), preference for remaining in the status quo, acting only after fraud has occurred (status quo bias), and the desensitization to data breaches due to breach or notification fatigue (Zou et al., 2018; Romanosky et al., 2011).

We extend the existing literature by exploring how individuals' past experiences with fraud affects their responses to a (future) data breach. We focus on two channels through which prior exposure to fraud or heightened fraud risks (specifically, following a data breach) may affect individuals' future precautionary responses. First, past exposure may raise individuals' awareness of the precautionary actions available to them, thereby reducing their cognitive costs of taking action. Individuals may have researched or received information (for instance, in a breach notification letter) about precautionary measures they could adopt during past events. Zou et al. (2018) found in their interview study that most individuals were either unaware of or misunderstand common protective actions such as placing fraud alerts and credit freezes, implying that individuals may face relatively high cognitive costs of taking protective actions. Moreover, the interviewees who were able to correctly describe

these measures have all been offered these services in a previous data breach. This finding suggests that individuals may learn about precautionary measures from past exposure, which lowers their costs of taking action in response to a future event, making them more likely to adopt protective actions.³

Given these results, we argue that prior fraud victims are likely to have lower costs of adopting precautionary measures than prior breach victims. Because fraud is more salient to fraud victims than breach victims, past fraud victims are likely to not only have a greater awareness of protective measures available, but also practical experience with taking these measures. Their experience with adopting precautionary measures in the past further lowers past fraud victims' costs of taking precautions in response to future fraud or breach events.

Second, past exposure to fraud or heightened fraud risks may affect individuals' perceived fraud risks and thus their expected losses from a future exposure to a data breach. The average individual is unlikely to know the true risks of experiencing fraud, conditional on falling victim to a data breach. Instead, they may rely on the availability heuristic to help them determine fraud risks, assigning higher risks to an event that is more salient or that they can more easily recall or imagine happening to them and vice versa (Tversky & Kahneman, 1974). Prior to any exposure to an event, individuals are likely to assign low probability to the event, as it is not highly available in their minds. In other words, their initial or baseline perceived risk is small. Exposure to an event in the past is likely to increase the salience and availability of the event for individuals, implying that they are likely to assign a higher probability to the event. Consequently, we expect individuals to assign different probabilities to an event, depending on their past exposures.

Past exposure to fraud and past exposure to a data breach may have only opposite effects on individuals' perceived breach-related fraud risk. Past exposure to fraud increases the availability of fraud, thereby raising the probability that fraud victims assign to the occurrence of fraud (both generally and conditional on exposure to a data breach). Past exposure to a data breach, in contrast, raises victims' perceived probability of data breaches but may lower the perceived probability of fraud occurring due to a data breach. These effects imply that past fraud victims are likely to assign the highest fraud risk to a data breach, while past breach victims may assign lower fraud risk than non-victims (individuals who have not been exposed to fraud and data breaches in the past).

The remainder of this paper is structured as follows. Section 2 provides a brief background of the Equifax data breach and initial evidence that individuals considered adopting

³Studies in the preventative health-care literature have also found similarly effects of learning from past exposure — past exposure to an illness personally or via a first-degree relative increases individuals' awareness of preventative health measures available (Baer et al., 2010; Mouchawar et al., 1999).

precautionary credit market measures following the breach. Section 3 presents our theoretical framework for examining consumer precautionary actions. In Section 4, we describe our data and sample construction, and in Section 5, we present our empirical specifications and our findings. Section 6 discusses other types of precautionary measures that individuals may adopt, and Section 7 concludes.

2 The Equifax Data Breach

2.1 Background

Equifax is one of the three major credit bureaus (also known as credit reporting agencies) in the United States. Credit bureaus collect information about individuals' credit accounts (for example, account balances, payment histories, credit limits, debt collections, and bankruptcies) from various third-parties, including banks, credit card companies, telecommunications and utilities companies, and landlords. Each credit bureau then compiles and maintains the information it has collected on an individual in a credit bureau file (or a credit report), which creditors and lenders can access and use for evaluating the individual's creditworthiness. Any individual who has recently used a traditional credit product (for example, a credit card, student loan, auto loan, or mortgage) almost certainly has a credit report at one or more of the three major credit bureaus. According to a 2015 study by the Consumer Financial Protection Bureau, the vast majority of adults (approximately 89%) in the United States has a credit report.

On September 7, 2017, Equifax publicly announced that it had suffered a data breach, potentially impacting 147 million individuals — the vast majority of adults — in the United States. Equifax revealed that hackers had gained unauthorized access to some of its data from mid-May through July 2017. The company detected the intrusion on July 29, 2017, and was able to identify and patch the vulnerability that the hackers had exploited to gain access to its system.⁴ The information the hackers accessed included names, Social Security Numbers (SSNs), birth dates, home addresses, and, in some cases, driver's license numbers. Additionally, the hackers also obtained the credit card numbers of about 209,000 individuals and dispute documents with PII of about 182,000 individuals.

⁴Equifax describes their actions since 2017 in the Equifax 2020 Security Annual Report. The authors of this paper have not verified the accuracy of this report.

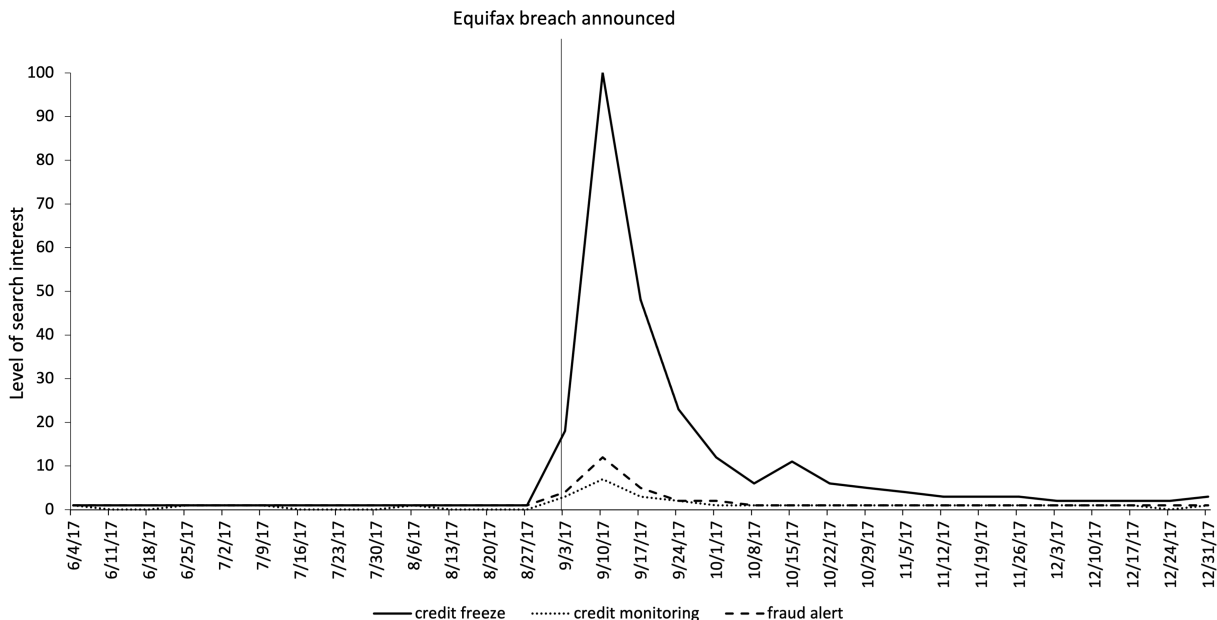
2.2 Individuals' Precautionary Behavior Following the Breach

The Equifax data breach placed affected individuals at heightened risks of ID theft and fraud. For instance, criminals may use the data stolen from Equifax to open new accounts (new account fraud) or take out additional credit on an existing account under an individual's name (existing account fraud). Following the announcement of the Equifax breach, many consumer protection groups, state attorneys generals, and security experts recommended affected individuals adopt precautionary measures to protect themselves against increased fraud risks.

Given the severity of the data breach and the recommendations for individuals to take precautionary actions, we may expect many individuals to have adopted measures to mitigate the fraud risks arising from the breach. Google search trend data indicate an increase in search interest in precautionary measures that can help reduce fraud risks following the breach announcement, suggesting that many individuals were considering adopting these measures. Figure 1 shows the Google search trends for three commonly recommended precautionary measures — *credit freeze*, *credit monitoring*, and *fraud alert* — in the U.S. between June and December 2017. Search interest in the above fraud protection measures increased sharply in the week following Equifax's breach announcement and began falling the week after. The timing of the rise and fall in levels of search interest for these terms strongly suggest that the spikes in interest for the above fraud protection measures were induced by the breach announcement. The term *credit freeze* saw the biggest and most persistent increase in search interest (from the pre-breach announcement period) after the announcement of the breach, which appears to be consistent with the strong emphasis that consumer protection agencies, state attorneys generals, security experts, and the media placed on credit freezes as a precautionary measure.

Individuals' decision to adopt precautionary measures may be influenced by a number of factors. In this paper, we are interested in individuals' past experiences of fraud and data breaches as a factor that may affect individuals' decisions. Past exposure to fraud and data breaches may influence how individuals perceive the risks and losses from these events and may also affect their cost of adopting precautionary measures. To better understand the effects of past exposure to these events on individuals' precautionary behavior following the Equifax breach, we develop a simple theoretical framework, which we present in the next section. Using our model, we formulate hypotheses on individuals' adoption of these precautionary measures. We then test our hypotheses empirically using credit bureau data, which we describe in Section 5.

Figure 1: Google Search Interest Over Time: Credit Freeze, Credit Monitoring, and Fraud Alert



Notes: Authors' calculations using data from Google trends from June 2017 to December 2017. Scale represents search interest relative to the highest point on the chart for the given time period.

3 A Descriptive Model of Consumer Precautionary Credit Market Action

3.1 Basic Setup

Consider a setting where we have a unit mass of individuals that are in the credit market for \bar{T} periods. At every period $t < \bar{T}$, individuals have to decide whether to adopt a precautionary measure j to reduce potential losses due to fraud of type $k \in [\text{New Account}, \text{Existing Account}]$ in the credit market, if they had not already done so.⁵ Suppose that at every period t , an individual i perceives the probability of fraud to be P_{it} and draw their potential fraud losses, L_i , from a continuous distribution G . Let τ^j denote the effective of measure j at reducing an individual i 's expected fraud losses and $C_i^j t$ denote the cost of adopting precautionary measure j in period t . Individual i would choose to adopt a measure j if the benefit of doing

⁵For precautionary measures that stay in place unless the individuals choose to remove them, we assume the cost of taking those precautionary measures to be zero in the periods following the initial placement of the measures.

so exceeds the cost; that is,

$$\tau^j P_{it}^k L_i > C_{it}^j.$$

Put differently, for a given set of parameter values, individuals would adopt precautionary measure j if their potential fraud losses are sufficiently large. The probability that individual i adopts is given by:

$$S_{it}^j = 1 - G\left(\frac{C_{it}^j}{\tau^j P_{it}^k}\right).$$

If individuals were informed about fraud risks (i.e., their perceived fraud risk matches the true fraud risk, \tilde{P}^k) and were fully rational, their decision on whether to adopt the precautionary measure would be optimal. In reality, however, individuals often do not know what the true probability of fraud is, and they are boundedly rational. Fraud may be considered a relative low-probability event. Existing research on consumer decision-making has found that individuals do not always seek information about the likelihood of low-probability events; instead, they tend to rely on heuristics to determine the probability of these events, and their decision processes are often subject to judgment biases (Camerer & Kunreuther, 1989). A well-known heuristic that individuals may use to assess the probability of an event is the *availability heuristic*.⁶ Using this heuristic, individuals assign higher probabilities to events that they can easily recall or imagine happening to them. Individuals' assessment of fraud risks may also be subject to an *optimism bias*; that is, they believe that fraud would not befall them.

Let P_0^k denote individuals' perceived probability of fraud at $t = 0$, before they were exposed to fraud or heightened fraud risks. Given that fraud can be considered a relatively rare event, we postulate that individuals are likely to underestimate this risk. Prior to any exposure to fraud, individuals are likely to find it difficult to imagine fraud occurring to them (low availability of fraud events in individuals' minds) and may be overly optimistic about their chances of not experiencing fraud (optimism bias). These behavioral factors tend to lead individuals to underestimate the true fraud risks (P_0^k is close to 0), implying that boundedly rational individuals are less likely to adopt a precautionary measure than fully rational individuals.

Individuals' lack of knowledge and awareness of the precautionary measures they can adopt further lowers the likelihood that individuals would take action prior to any exposure to fraud or fraud risks. Their lack of knowledge translates into a high (non-monetary) cost of

⁶Studies in the field of natural disaster risk mitigation have found that the use of the availability heuristic to assess the risk of a natural disaster helps to explain the pattern of individuals' adoption of precautionary measures (Kunreuther, 2006). Since natural disasters and credit market fraud are similar types of events (adverse events with low probability of occurrence but potentially high losses), we believe that individuals are also likely to rely on the availability heuristic to judge the risks of credit markets.

taking action; that is, their cost of taking action j at $t = 0$, C_0^j , is very high. For simplicity, we assume that P_0^k is sufficiently small (or C_0^j sufficiently high) such that individuals do not find it optimal to adopt any precautionary measure in the absence of a shock.

3.2 Past Exposure and Consumers' Response to a Data Breach

Suppose that at $t = T''$, all individuals are exposed to a data breach. Exposure to the data breach (objectively) increases the likelihood of fraud, which causes individuals' perceived fraud risk to jump discretely. The exposure also increases the salience of fraud events, which may further increase individuals' perceived fraud risks. Consequently, individuals are more likely to adopt precautionary measures than in the absence of a shock, independent of their past experience.

Individuals' past experience with ID theft or a data breach can affect how much more (or less) likely they are to adopt precautionary actions in response to the new data breach than unexposed consumers through two channels. The first channel is consumers' perceived fraud risk (via the availability effect). Suppose that at $t = T' < T''$, some individuals were exposed to ID theft and some others were exposed to a data breach (but did not experience fraud). Past ID theft (fraud) victims are likely to perceive the risk of fraud to be higher (in general) than those who have not been exposed to fraud because fraud events are more available to them.⁷ These individuals are therefore more likely to have the highest perceived fraud risk following the new breach at $t = T''$ compared to other individuals. Interestingly, whereas the occurrence of fraud is a highly available event for past ID theft victims, the reverse is true for past data breach victims who did not experience fraud as a result. These past data breach victims are likely to recall that no fraud occurred the last time their data was breached and assign a lower probability to fraud occurring following the new breach than past ID theft victims, as well as previously unexposed individuals (or non-victims). We thus conjecture the ordering of perceived fraud risks after the new breach to be $P_{T''}^{k,F} > P_{T''}^{U,k} = P_0^k > P_{T''}^{B,k}$.

A second channel through which past exposure may affect individuals' response to a data breach is the cost of adopting precautionary measures. Prior breach and fraud victims may face lower (non-monetary) costs of adopting precautionary measures following the new breach, because they may have learned about the measures that are available and how to sign up for them from their past exposure to fraud or fraud risks. We argue that past ID theft victims are likely to have lower costs of taking precautions at $t = T''$ than past breach victims, as ID theft victims are likely to have both greater awareness of the actions available and practical experience with taking these actions. Put differently, past ID theft

⁷Events that have been personally experienced by individuals tend to be more available to them (Kahneman, 2011).

victims are likely to have learned more from their prior exposure regarding the adoption of precautionary measures than past breach victims. We therefore postulate that $E(C)_{T''}^{j,F} < E(C)_{T''}^{j,B} < E(C)_{T''}^{j,U} = C_0^j$.

All else equal, the availability and learning effects from past exposure imply that past ID theft victims are most likely to take action in response to the new breach. ID theft victims are likely to have both the highest perceived fraud risk and the lowest cost of taking action. Whether data breach victims may be more or less likely to take action relative to non-victims is *a priori* ambiguous. Prior breach victims are likely to assign lower fraud risks to the new breach than non-victims, as their past exposure to a breach did not lead to fraud. However, they are also likely to have a lower cost of taking precautions than non-victims due to learning from their past exposure. Prior breach victims are more likely to adopt precautionary measures following the new breach if the learning effect dominates the availability effect, and vice versa.

Assuming for simplicity that no individuals have a precautionary measure in place at the start of $t = T''$, we therefore hypothesize that prior fraud victims are more likely to adopt a new precautionary measure following the new breach than prior breach victims and previously unexposed individuals. However, we are unable to formulate a hypothesis on prior breach victims' uptake of a new precautionary measure relative to unexposed individuals, as we do not have a strong prior on whether the learning effect or the availability effect would dominate. We leave the direction of the difference in the adoption rates of a new precautionary measure between the two groups open for empirical investigation.

4 Data and Sample Selection

4.1 Data Description

To empirically examine how individuals previously exposed to fraud react to the Equifax breach announcement, we use data from the Federal Reserve Bank of New York Consumer Credit Panel/Equifax (CCP). The CCP data set is an anonymized 5% random sample of all U.S. individuals with a credit bureau record. To be included in the data set, individuals must have a Social Security number and at least one public record or credit account.⁸ The CCP is an unbalanced panel that follows individuals at a quarterly frequency and is constructed so that (1) new individuals are included over time as they open their first credit account or gain their first public records and (2) are dropped from the sample when they die or experience a prolonged period of credit market inactivity. The sample is designed in this way to mirror

⁸Public records include adverse events such as bankruptcy declarations or tax liens.

the entry and exit dynamics in the general credit bureau data population (Lee & van der Klaauw, 2010).

The CCP data are then merged with an unique data set of anonymized fraud alert information obtained by the Consumer Finance Institute from Equifax. These data contain detailed information on the type of fraud alert filed, the status of the alert, and utilization of other services such credit freezes and opt outs from prescreened credit offers. Importantly, we also observe the month and year of the placement of both the fraud alerts and credit freezes, allowing us to more precisely measure when individuals take these measures.

We also use additional anonymized data on individual credit card accounts (tradelines) for individuals in the CCP. This data set contains detailed disaggregated information on up to 10 credit card accounts for every individual in the CCP, though the periodicity of the data has changed over time. Tradeline data are available from 2017 to 2018 at a quarterly frequency and are available at a semiannual frequency from 2014 to 2016. For each account, we can observe the date it was opened, its payment status, current balance and limit, and the date of the last activity on the account. We also observe contextual information for each account in the form of “narrative codes.” For example, narrative codes can indicate if an account is part of a bankruptcy proceeding or if the account has been closed. Since both the main CCP and fraud alert data (both anonymized) are at the individual-quarter level, we aggregate the tradeline data up to the individual level. We provide more details on the use of the credit card tradeline data in the Appendix.

4.2 Sample Construction and Identifying Individuals Exposed to Fraud

To form our analytical sample, we merge all three data sets together from the years 2014 to 2018. We drop any individuals under the age of 18 or were recorded as being deceased at any point during our sample period. We also drop any individual with fewer than four total observations across the sample period to mitigate any problems due to “fragment” files in the credit bureau data.⁹ After these restrictions, we are left with approximately 12 million unique individuals.

To test our hypotheses on how *prior* exposure to fraud affects individuals’ credit market behavior, we identify three sets of individuals in the CCP data. The first two groups are our treatment groups: (1) individuals who have previously experienced fraud in the form of severe ID thefts and (2) those individuals who have faced heightened risks of fraud due

⁹Fragment credit bureau records occur when new records are created and then subsequently merged with preexisting records when a credit bureau discovers that the two records belong to the same individual.

to their exposure to the 2015 Anthem data breach (which we will elaborate on shortly). The third group is our control group, which comprises individuals who were not exposed to either severe ID theft or the Anthem data breach. We describe in detail how we identify our treatment groups in the following subsections.

4.2.1 Prior Fraud Exposure: Severe ID Theft Victims

To identify past victims of fraud, we use the extended fraud alert flags in the CCP data supplement, which is a strong indicator of being a victim of severe ID theft. Under the Fair and Accurate Credit Transactions Act (FACTA), the presence of an extended fraud alert on an individual’s credit report requires potential creditors to perform stringent identification verification requirements before extending credit to that individual. When filing an extended fraud alert, the individual must specify a telephone number or other reasonable contact method as part of the alert documentation; all creditors must contact the individual by the method specified in the alert to verify the individual’s ID when receiving an application for credit. Once placed, the extended fraud alert remains on the individual’s credit reports for seven years (unless the individual chooses to remove it beforehand) instead of one.¹⁰

The placement of an extended fraud alert flag in an individual’s credit file is an excellent proxy for severe ID theft victimization. Placing an extended fraud alert in a credit bureau file is an elaborate filing process, as the alert filer is required to submit either a police report or an ID Theft Report (ITR) to place the alert in their credit file. An ITR requires detailed information on the accounts that were compromised and accompanying evidence of ID theft or fraud. Providing such evidence entails both time and effort, and individuals face criminal penalties for falsifying information in these reports. Because of these requirements, filers of extended alerts are unlikely to place alerts in their credit bureau files based simply on worries or as a precaution.¹¹

4.2.2 Prior Exposure to the Risk of Fraud: 2015 Anthem Data Breach Victims

Anthem (Anthem Inc.) is one of the largest health insurers in the United States. Anthem offers private insurance plans through the Blue Cross (BC) and Blue Cross Blue Shield (BCBS) networks, as well as managed care plans (Medicaid). As of the end of 2014, Anthem and its affiliates served over 71 million individuals from all across the country. It offered private insurance plans in 14 states, operating as Anthem Blue Cross, Anthem Blue Cross Blue

¹⁰Additionally, an extended fraud alert removes the individual’s credit file from lists of prescreened credit and insurance offers for five years.

¹¹Using the same CCP data, Blascak et al. (2021) show that credit market behavior consistent with fraud, such as the opening of new accounts and changes of addresses, increases right before or concurrently with the placement of these flags.

Shield, Blue Cross Blue Shield of Georgia, or Empire Blue Cross Blue Shield, and provided managed care plans in 19 states and the District of Columbia through its subsidiaries, such as Amerigroup and UniCare.

In late January 2015, Anthem discovered that hackers had gained unauthorized access to one of its databases, which contained personal information, including names, dates of birth, Social Security numbers, home addresses, email addresses, and employment information of approximately 80 million customers and employees, dating back to 2004. On February 4, 2015, Anthem announced the data breach and began mailing out notification letters to individuals affected by the breach. In response to the breach, the company offered two-year free credit monitoring, child ID protection, and ID theft repair services, which affected individuals could sign up for beginning on February 13, 2015.¹²

The likelihood of exposure to the Anthem breach varies by state, depending largely on Anthem’s share of the insurance market in a given state. To identify how each state was affected by the breach, we gather data on the states where Anthem operated in 2015 (especially where it offered both BC or BCBS plans and Medicaid plans), and data on the number of individuals affected by the Anthem breach in each state.¹³ Using this information, we identify a group of states that were most affected by the breach (defined as having more than 25% of the state’s population having been a victim) and a group of states that were *least* affected by the breach (defined as having less than 5% of the state’s population reported being a victims). Table 1 provides a breakdown of these states. We define individuals living in the most exposed states each as individuals affected by a prior data breach (“prior data breach victims”).

In Appendix Section C, we provide event study results similar to those reported in Section 5.1 to demonstrate that individuals living in states that were most exposed to the Anthem data breach were more likely to respond to the announcement of the breach when compared to individuals living in states that were the least exposed.

4.3 Summary Statistics

As mentioned in Section 2.2, we are interested in two specific measures of precautionary behavior in credit markets: the likelihood of placing a new credit freeze and the number of

¹²We chose to focus on the Anthem breach for a few reasons. First, it was one of the largest data breaches in pre-Equifax breach period. Second, the types of data compromised in the Anthem breach are similar to those exposed in the Equifax breach. Thus, the knowledge and practical experience that individuals may gain from dealing with exposure to the Anthem breach are likely to be relevant in the case of Equifax’s breach.

¹³We collect data from news reports and press releases (whenever available). For states in which the victim count were not publicly reported, we contacted the state’s attorney general’s office or insurance department to request for the information.

Table 1: Anthem Breach: Most and Least Affected States

Most Affected States			
State	Number of Victims	State Population	% Affected
California	13,500,000	38,918,045	34.69%
Connecticut	1,700,000	3,587,122	47.39%
Georgia	3,700,000	10,178,447	36.35%
Indiana	4,500,000	6,608,422	68.09%
Maine	531,000	1,328,262	39.98%
Missouri	2,000,000	6,071,732	32.94%
New Hampshire	668,000	1,336,350	49.99%
New York	5,023,000	19,654,666	25.56%
Virginia	3,770,000	8,361,808	45.09%
Least Affected States			
State	Number of Victims	State Population	% Affected
Alaska	34,000	737,498	4.61%
Colorado	19,700	5,450,623	0.36%
Hawaii	18,000	1,422,052	1.27%
Illinois	215,000	12,858,913	1.67%
Montana	48,000	1,030,475	4.66%
New Jersey	209,000	8,867,949	2.36%
New Mexico	11,600	2,089,291	0.56%
North Dakota	27,000	754,066	3.58%
Oklahoma	100	3,909,500	0%
Utah	10,956	2,981,835	0.37%
West Virginia	220	1,842,050	0.02%

Note: Authors' calculations based on statistics from state websites, press articles, and correspondence from states' attorneys generals offices as of March 2020. Population data as of 2015.

closed credit card accounts. Table 2 provides the summary statistics of these precautionary credit market measures for prior fraud (ID theft) victims, prior data breach victims, and for the entire sample (including non-victims). The prevalence of having a credit freeze prior to the Equifax breach is very low, with less 1% of all individuals having an active freeze. This percentage increases by a factor of 4 (from 0.7% to 2.8%) after the breach. We also observe very few credit cards being closed on average (0.039 accounts in the pre-breach period), though the percent of individuals having closed a card is between 10% and 15%. Unsurprisingly, we observe relatively large differences in the means of our outcome variables of interest between our two groups of prior exposure victims. These gaps are unsurprising given the differences in severity of the prior victimization.

To explore whether individuals took precautionary actions following the breach, we examine how our two precautionary credit market variables evolved over time. To do so, we

Table 2: Summary Statistics

	Pre-Equifax Breach Average (S.D.)	Post-Equifax Breach Average (S.D.)
<u>Panel A: Full Sample</u>		
Number of bankcards	2.22 (2.32)	2.32 (2.39)
Number of closed bankcards	0.039 (0.21)	0.041 (0.22)
Number of inquiries	0.554 (1.06)	0.500 (0.97)
% with a freeze	0.7%	2.8%
% with a closed credit card account	11.0%	11.8%
Number of individuals	12,062,423	
<u>Panel B: Prior Fraud Victims</u>		
Number of bankcards	2.14 (2.39)	2.33 (2.53)
Number of closed bankcards	0.048 (0.24)	0.052 (0.25)
Number of inquiries	0.988 (1.68)	0.849 (1.47)
% with a freeze	3.8%	6.4%
% with a closed credit card account	14.7%	15.1%
Number of individuals	37,886	
<u>Panel C: Prior Breach Victims</u>		
Number of bankcards	2.46 (2.43)	2.61 (2.52)
Number of closed bankcards	0.040 (0.22)	0.042 (0.22)
Number of inquiries	0.500 (0.97)	0.456 (0.89)
% with a freeze	1.0%	3.5%
% with a closed credit card account	11.7%	12.3%
Number of individuals	3,329,793	

Notes: Authors' calculations using data from Federal Reserve Bank of New York Consumer Credit Panel/Equifax. Sample is from Q1:2016 to Q4:2018. Prior fraud victims are individuals who had filed an extended fraud alert any time between Q1:2010 to Q2:2017. Prior data breach victims are individuals living in states where at least 25% of the total population was affected by the 2015 Anthem data breach. We exclude individuals who fall in both victimization categories.


use a standard event study methodology to examine how individuals' adoption of precautionary measures changed in each time period before and after the breach announcement. Specifically, we estimate the following equation from the first quarter of 2016 to the fourth quarter of 2018:

$$y_{it} = \Pi_t + \mathbf{X}_{it}\boldsymbol{\Omega} + \delta_i + \gamma_s + \epsilon_{it}, \quad (1)$$

where y_{it} is the precautionary credit market measure of interest and Π_t is a vector of time dummy variables extending from six quarters before to six quarters after the Equifax breach announcement. Our omitted period is the 2nd quarter of 2017, which means that the estimated effects of the time dummies are relative to the quarter before the breach was announced. We include the county unemployment rate, county share of non-White, county total population, and age bin dummy variables as controls in the vector X_{it} . We also include state fixed effects, γ_s , and individual fixed effects, δ_i . Standard errors are clustered at the individual level.

Figure 2: Consumer Response to the Equifax Breach

Panel A: Probability of Having a New Credit Freeze Panel B: Number of Closed Credit Card Accounts



Notes: Authors’ calculations using data from Federal Reserve Bank of New York Consumer Credit Panel/Equifax. Analysis performed on the “full sample” outlined in Panel A of Table 2.

Figure 2 presents the results of our simple event studies. Panel A of 2 shows that the probability an individual would place a new credit freeze on their Equifax credit report increased sharply during the quarter of the breach announcement by 1 percentage point (a relative 142% increase) and remained elevated, though at a lower rate, from the fourth quarter of 2017 through the end of 2018. In Panel B, we see that the number of closed bankcard accounts did not immediately respond during the quarter of the breach announcement. However, the average number of closed card accounts dips slightly one quarter after the announcement and then increases by 0.006 accounts (a relative 14% increase) starting in the first quarter of 2018, and remains elevated for the rest of the year. We do note that there is evidence of a significant pre-trend in the quarters prior to the Equifax breach, and it is plausible that the increase in the number of closed accounts is an increase in the preexisting trend.

Overall, these figures show that consumers took precautionary actions in response to

the Equifax breach. Interestingly, these general results differ from the prior literature in that we observe persistent, elevated effects for our precautionary effects after the breach announcement. For example, Mikhed and Vogan (2018) found that after a serious data breach in South Carolina in 2012, there was increased uptake of credit freezes, but this only lasted for two quarters after the breach announcement.

5 Past Fraud Exposure and Individuals’ Response to the Equifax Breach

Having established that individuals, on average, took precautionary actions following the Equifax breach announcement, we now examine how prior exposure to fraud or the risk of fraud (i.e., prior exposure to a data breach or prior ID theft victimization) affects individuals’ responses. Specifically, we aim to (1) test if our hypothesis that prior *fraud* victims are more likely to adopt both precautionary measures than prior breach victims and non-victims and (2) examine if prior data breach victims are more or less likely to take up precautionary measures than non-victims.

5.1 Estimating Changes in Precautionary Behavior by Prior Fraud Exposure

To estimate how different types of prior fraud exposure affect precautionary credit market behavior, we examine how each prior exposure group responded to the Equifax data breach announcement relative to our control group of non-victimized individuals in a difference-in-differences (DID) framework. In addition to examining how each prior fraud exposure group responds to the Equifax breach relative to non-victims, we also directly compare prior fraud victims to prior data breach victims. Based on our conceptual framework presented in Section 3, we may expect prior fraud victims to take up precautionary measures at a higher rate than prior breach victims. Our DID estimating equation takes the following form:

$$y_{it} = \alpha_0 + (\mathbf{\Pi}_t \times D_i)\mathbf{\Psi} + \alpha_1 D_i + \Pi_t + \mathbf{X}_{it}\mathbf{\Omega} + \delta_i + \gamma_c + \epsilon_{it}, \quad (2)$$

where $D_i = 1$ if individual i belongs to the prior fraud exposure group and $\mathbf{\Pi}_t$, \mathbf{X}_{it} , δ , and γ are as defined in Equation (1). The vector $\mathbf{\Psi}$ contains the coefficients of interest on the interaction terms of the treatment indicator variable D_i and the time dummy variables $\mathbf{\Pi}_t$. We cluster our standard errors at the individual level. The identifying assumption for our DID specification is that absent the Equifax data breach, our outcomes of interest


for the treatment and control individuals, conditional on our control variables, would have trended similarly over time. We provide evidence that our outcomes meet the parallel trends assumption in the form of event study plots in Figures 3 to 5. In Table 3, we also report estimates of DID coefficients where we pool the individual time periods into short-run and long-run dummy variables, with the short-run dummy variable equal to one for two quarters immediately after the breach and the long-run dummy variable equal to one for the year 2018, which is three to six quarters after the announcement.

Results from estimating Equation (2) for individuals with prior exposure to ID theft are presented in Figure 3. Panel A shows that compared to non-victims, previous ID theft victims have a 0.1 percentage point higher likelihood of placing a new credit freeze in their credit report in the quarter immediately after the Equifax breach, and this likelihood increases to over 0.25 percentage points by the third quarter after the breach. In percentage terms, prior fraud victims were 2.6% to 6.6% more likely to place a new credit freeze relative to their pre-Equifax breach mean in the quarters after the breach announcement. In Panel B of Figure 3, we see that there is relatively little change in the number of closed bankcard accounts between the two groups after the breach. Our DID results in Panel A of Table 3 are consistent with the event study results, with the likelihood of having a freeze increasing by a statistically significant 0.2 percentage points in the short-run period and by 0.15 percentage points in the long-run period, and no significant changes in the number of closed credit cards in either period.

Figure 3: Equifax Breach Results: Prior Fraud Victims vs Non-Victims

Panel A: Probability of Having a New Credit Freeze

Panel B: Number of Closed Credit Card Accounts



Notes: Authors' calculations using data from Federal Reserve Bank of New York Consumer Credit Panel/Equifax.


Figure 4 shows the results for individuals who had been previously been exposed to the

2015 Anthem breach relative to non-victims. We can see that there is a large, statistically significant increase in the probability of placing a new credit freeze (0.3 percentage points, or a 30% increase) in the quarter of the breach announcement, which is three times larger than the response of prior ID theft victims. This probability remains elevated for the following two quarters, but it returns to zero by the end of our sample period. For the number of closed accounts in Panel B, we see that prior breach victims close *fewer* accounts after the breach compared to non-victims, with a decline of 0.002 accounts (a relative 5% decline) in the quarter immediately after the breach. This downward trend continues through the remaining five post-breach quarters, which is in contrast to the behavior we observe from prior ID theft victims, who closed accounts at the same rate as non-victims in the post-breach period.


Our DID results in Table 3 Panel B show similar results, with the likelihood of having a credit freeze increasing by 0.2 percentage points in the short-run period and increasing by 0.01 percentage point in the long-run period. The number of closed credit card accounts *decreases* by 0.001 cards in the short-run period and by 0.002 cards in the long-run period. While prior breach victims increased their likelihood of having a new credit freeze (relative to non-victim) by the same fraction of a percentage point as prior identify theft victims (relative to non-victims) as shown in Panel A, the increase is of different economic magnitude for the two groups: The 0.2 percentage point increase represents an increase of an additional 130,000 new credit freezes for prior breach victims, whereas the 0.2 percentage point increase for prior fraud victims translates to approximately 1,500 new freezes.

Figure 4: Equifax Breach Results: Prior Breach Victims vs. Non-Victims

Panel A: Probability of Having a New Credit Freeze



Panel B: Number of Closed Credit Card Accounts



Notes: Authors' calculations using data from Federal Reserve Bank of New York Consumer Credit Panel/Equifax.

In Figure 5, we directly compare prior ID theft victims to prior breach victims to see if their responses to the Equifax breach are statistically different from each other. To do this, we redefine D_i in Equation (2) so that $D_i = 1$ if an individual was a prior ID theft victim and $D_i = 0$ if an individual was a prior breach victim. In Panel A, we see that before the Equifax breach, both groups of individuals placed new freezes on their credit reports at the same rate. In the quarter of the breach (Q3 2017), previous ID theft victims placed new freezes at a *lower* rate than previous breach victims by 0.2%. However, in the quarters after the breach, the relationship reverses and previous ID theft victims place freezes at a *higher* rate than previous breach victims by a similar magnitude, with the effect persisting until the end of our sample. This immediate reversal is consistent with the summary statistics, where prior fraud victims have a higher rate of freeze placement than prior breach victims.

Panel B of Figure 5 shows that after the breach, prior ID theft victims closed 0.003 more accounts (a 7.5% increase relative to the prior breach group’s pre-breach mean) on average than for prior breach victims in the quarters after the Equifax breach. However, we note that our event study coefficients are estimated noisily, and we cannot rule out a zero effect. Our short-run and long-run DID coefficients in Panel C of Table 3, where we pool all post-breach time periods into two separate post-period dummy variables, are more precisely estimated and indicate that prior fraud victims have a 0.2 percentage point higher likelihood of having a credit freeze in the long-run period and close 0.004 and 0.005 more accounts in both the short-run and long-run periods, respectively. While the number of card accounts closed is economically small, these short-run and long-run coefficient estimates represent a 9%-11% increase.

Overall, our estimates from Figures 3 to 5 and Table 3 are consistent with our theoretical framework, which predicts that prior ID theft victims are more likely to take precautionary actions than prior breach victims, and that both types of prior fraud victims would take more action than non-victims. The magnitudes of our effect sizes, while statistically significant, are also consistent with the previous literature that, in aggregate, many individuals do not take precautionary actions in response to data breaches. In addition, the persistence of our results, when compared to the previous literature (e.g., Mikhed & Vogan, 2018), shows that individuals with prior exposure to fraud had a more persistent precautionary response to the Equifax data breach when compared to prior data breaches.


5.2 Robustness Checks

In Appendix Figures A3 to A5, we plot event study coefficients for two related credit market variables: the number of inquiries and the number of new accounts opened. While they are

Figure 5: Equifax Breach Results: Prior Fraud Victims vs. Prior Breach Victims

Panel A: Probability of Having a New Credit Freeze

Panel B: Number of Closed Credit Card Accounts



Notes: Authors' calculations using data from Federal Reserve Bank of New York Consumer Credit Panel/Equifax. Results are for prior ID theft victims relative to prior data breach victims (the control group).

not a direct indicator of precautionary behavior by themselves, when combined with our main measures, declines in either/both would provide some supporting evidence that individuals increased their precaution following the Equifax breach. More specifically, to reduce the likelihood that they experience fraud, individuals may take more passive actions like reducing their demand for credit, which can be measured by the number of applications sent or the number of new accounts opened. Appendix Figure A3 shows the results for prior ID theft victims vs. non-victims, Appendix Figure A4 shows the results for prior breach victims vs. non-victims, and Appendix Figure A5 shows the results comparing prior fraud victims and prior breach victims. Results for inquiries in all three figures have positive pre-trends. For the number of new accounts opened, we see that fraud victims open fewer accounts than non-victims, while prior breach victims open more new accounts. When comparing the two groups, we see that prior fraud victims open fewer new accounts than prior breach victims after the Equifax breach.¹⁴ We also report short- and long-run coefficients in Appendix Table A1.

¹⁴We also report results for these two measures for the Anthem breach in Appendix Figure A2.

Table 3: Difference-in-Difference Results

	Likelihood of Freeze	Number of Accounts Closed
Panel A: Prior Identity Theft vs. Non-Victims		
$Treat \times Short-run$	0.0019*** (0.0004)	0.0031 (0.0011)
$Treat \times Long-run$	0.0015*** (0.0002)	0.0018 (0.0009)
N	101,985,621	69,511,672
Panel B: Prior Breach vs. Non-Victims		
$Treat \times Short-run$	0.0020*** (0.0000)	-0.0014*** (0.0001)
$Treat \times Long-run$	0.0001*** (0.0000)	-0.0020*** (0.0001)
N	140,901,416	97,410,987
Panel C: Prior Identity Theft vs. Prior Breach		
$Treat \times Short-run$	-0.0001 (0.0004)	0.0046*** (0.001)
$Treat \times Long-run$	0.0015*** (0.0002)	0.0039*** (0.001)
N	39,820,443	28,532,171

Note: Authors' calculations using data from Federal Reserve Bank of New York Consumer Credit Panel/Equifax. $Short-run = 1$ for time periods Q3:2017 to Q4:2017 and $Long-run = 1$ for time periods Q1:2018 to Q4:2018.

6 Discussion

6.1 Other Types of Precautionary Measures

In addition to the two measures examined in this paper, individuals may also adopt a number of other precautionary measures, including purchasing credit monitoring services and signing up for initial fraud alerts. Unlike placing a credit freeze or closing a credit card account, which are proactive measures that can prevent fraud from occurring, these measures are passive and do not prevent fraud from occurring. Instead, these measures help alert individuals to fraud that has already arisen, enabling individuals to act to stop current fraudulent activity and prevent further fraud from occurring.

Credit monitoring services are commercial services that help individuals watch their credit reports for changes and activities, such as new account openings and credit inquiries, noti-

ying individuals whenever any change or activity occurs. These notifications may serve as early warnings of potential fraudulent activities — specifically, when individuals are alerted to changes or activities that they did not initiate — thereby enabling individuals to act to prevent the fraud or limit their losses. That said, if individuals take no action in response to the alerts, credit monitoring services will not help to prevent fraud, as they do not restrict creditors’ access to individuals’ credit reports. Individuals may purchase credit monitoring services from any of the three major credit bureaus, as well as third-party ID theft protection companies, such as LifeLock or ID Guard. Breached entities often offer affected individuals one to two years of free credit monitoring services. In Equifax’s case, the company offered all affected individuals one year of its credit monitoring services for free.


An initial fraud alert indicates to creditors that an individual is a possible victim of fraud.¹⁵ Under FACTA, when creditors observe an initial fraud alert in an individual’s credit file, they are required to take extra steps to verify the individual’s ID before granting any request to open a new credit account, increase an existing credit line, or issue an additional card associated with an existing credit account under the individual’s name. Because of the more stringent ID verification requirements, fraud alerts may help to lower the chances that an individual experiences new account fraud. However, unlike credit freezes, fraud alerts do not limit creditors’ access to the individual’s credit files; a criminal may thus still be able to open fraudulent new accounts under the individual’s name. An individual may place an initial fraud alert by contacting any one of the credit bureaus; the credit bureau that was contacted will notify the other bureaus of the alert. Currently, each initial fraud alert lasts for a year; individuals may place new alerts when their existing ones expire.¹⁶ Filing a fraud alert is free under the FACTA.

Although we are unable to analyze the adoption of credit monitoring services and initial fraud alerts in our paper, we find evidence that consumers may have adopted these measures using Google trends data. Figure 6 shows a spike in the volume of Google searches for the terms “credit monitoring” and “fraud alert” in the week following Equifax’s data breach announcement. These search trends suggest that consumers were seeking information on these precautionary measures after the breach announcement and as a result may have implemented some of them.

¹⁵Individuals who have been actual victims of fraud — specifically, severe ID theft — can file for an extended fraud alert. Extended fraud alerts are similar to initial fraud alerts, but they impose stricter filing requirements and last for a longer period time.

¹⁶Prior to the passage of a new federal law on September 21, 2018, each initial fraud alert lasts 90 days.

Figure 6: Google Search Interest Over Time: Credit Monitoring and Fraud Alert



Notes: Authors' calculations using data from Google trends from June 2017 to December 2017. Scale represents search interest relative to the highest point on the chart for the given time period.

7 Conclusion

This paper examines the effect of prior fraud and data breach exposure on future precautionary behavior in credit markets by exploiting the 2017 Equifax data breach, which exposed sensitive information for over 70% of the U.S. adult population. To guide our analysis, we first develop a descriptive model of credit market precautionary behavior, which enables us to formulate theoretically founded hypotheses on how past fraud or data breach exposure may affect individuals' precautionary responses to a future data breach. We test our hypotheses using a DID framework and large anonymized data set of consumer credit records. We find that, consistent with our hypotheses, past exposure to fraud or a data breach raises an individual's probability of adopting a precautionary measure following the announcement of the Equifax breach relative to previously unexposed individuals. Further, between prior fraud and prior breach victims, the former is more likely to have taken precautionary measures than the latter.

Although prior fraud and data breach victims were more likely to adopt precautionary measures than unexposed individuals, the vast majority of these individuals did not do so. One possible explanation is the cost reduction resulting from learning about precautionary measures from prior exposure is relatively small. Another potential explanation is that individuals' expected fraud losses are very small. The low expected fraud losses may be due to individuals' belief that they have nothing to lose, the liability protection that individuals

have, particularly against credit card fraud, or an optimism bias — individuals do not think fraud will occur to them. Thus, even if they were to experience a substantial reduction in the cost of taking precautionary actions, their incentive to do so would be relatively weak.

While a low level of adoption of precautionary credit market actions does not necessarily imply that individuals are behaving suboptimally, we suspect that it may be in the case with their responses to the Equifax breach and to data breaches more generally. More specifically, we believe that individuals may be underestimating their potential losses because of the low-availability of fraud events, overoptimism regarding their likelihood of experiencing fraud, or mistaken beliefs that liability protection like that offered by credit card companies is available to them for all types of losses. Consumer protection agencies (CPAs) may thus have a role to play in improving individuals' understanding of the fraud risks and losses they could face from a breach of their personal data.

References


- Ablon, L., Heaton, P., Lavery, D. C., & Romanosky, S. (2016). *Consumer attitudes toward data breach notifications and loss of personal information*. RAND Corporation.
- Baer, H. J., Brawarsky, P., Murray, M. F., & Haas, J. S. (2010). Familial risk of cancer and knowledge and use of genetic testing. *Journal of general internal medicine*, *25*(7), 717–724.
- Blascak, N., Cheney, J., Hunt, R., Mikhed, V., Ritter, D., & Vogan, M. (2021). *Financial consequences of severe identity theft in the U.S.* (Working Paper No. 21-41). Federal Reserve Bank of Philadelphia.
- Brodkin, J. (2007, Apr). *Victims of choicepoint data breach didn't take advantage of free offers*. Network World. Retrieved June 27, 2022, from <https://www.networkworld.com/article/2297654/victims-of-choicepoint-data-breach-didn-t-take-advantage-of-free-offers.html>
- Camerer, C. F., & Kunreuther, H. (1989). Decision processes for low probability events: Policy implications. *Journal of Policy Analysis and Management*, *8*(4), 565–592.
- Cheney, J. S. (2005). *Identity theft: Do definitions still matter?* (Payment Cards Center Discussion Paper). Federal Reserve Bank of Philadelphia.
- Harell, E. (2021, Apr). *Victims of identity theft, 2018*. Bureau of Justice Statistic. Retrieved June 27, 2022, from <https://bjs.ojp.gov/content/pub/pdf/vit18.pdf>
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Kunreuther, H. (2006). Disaster mitigation and insurance: Learning from Katrina. *The Annals of the American Academy of Political and Social Science*, *604*(1), 208–227.
- Lee, D., & van der Klaauw, W. (2010). An introduction to the New York Fed consumer credit panel. *Federal Reserve Bank of New York Staff Reports*(479).
- Mikhed, V., & Vogan, M. (2018). How data breaches affect consumer credit. *Journal of Banking & Finance*, *88*, 192–207.
- Mouchawar, J., Byers, T., Cutter, G., Dignan, M., & Michael, S. (1999). A study of the relationship between family history of breast cancer and knowledge of breast cancer genetic testing prerequisites. *Cancer Detection and Prevention*, *23*(1), 22–30.

- Ponemon Institute. (2014). *The aftermath of a data breach: Consumer sentiment* (Research Report). Retrieved June 27,2022, from <https://www.ponemon.org/research/ponemon-library/security/the-aftermath-of-a-data-breach-consumer-sentiment.html>
- Romanosky, S., Telang, R., & Acquisti, A. (2011). Do data breach disclosure laws reduce identity theft? *Journal of Policy Analysis and Management*, *30*(2), 256–286.
- Turjeman, D., & Feinberg, F. M. (2019). *When the data are out: Measuring behavioral changes following a data breach*. Working Paper, SSRN 3427254. Retrieved June 27, 2022, from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3427254
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131.
- Zou, Y., Danino, S., Sun, K., & Schaub, F. (2019). You might be affected: An empirical analysis of readability and usability issues in data breach notifications. In *Proceedings of the 2019 chi conference on human factors in computing systems* (pp. 1–14).
- Zou, Y., Mhaidli, A. H., McCall, A., & Schaub, F. (2018). “ I’ve got nothing to lose”: Consumers’ risk perceptions and protective actions after the Equifax data breach. In *Fourteenth symposium on usable privacy and security (SOUPS 2018)* (pp. 197–216).

APPENDIX

A Anthem Data Breach

Figure A1: Google Search Trends for “Anthem Breach”; Nov. 1, 2014 - May 31, 2015




Notes: Authors’ calculations using data from Google trends from November 2014 to May 2015 for the states of Indiana, New Hampshire, Oklahoma, and West Virginia. Indiana and New Hampshire are “high” exposure states in that they had the highest percentages of their population affected by the Anthem data breach. Oklahoma and West Virginia are “low” exposure states as they had the lowest percentages of their population affected by the breach.


As a basic soundness check for our identification strategy, we compared the Google search interest in the term *Anthem breach* in two high-exposure states (Indiana and New Hampshire) relative to two low-exposure states (Oklahoma and West Virginia). Figure A1 shows that prior to the Anthem breach announcement, the level of search interest in the term was zero in any of the four states. In the week of the breach announcement (the week of Feb. 1, 2015), the search interest in *Anthem breach* increased substantially in the high-exposure states relative to the low-exposure states. The level of Google search interest for *Anthem breach* rose slightly in Oklahoma and not at all in West Virginia. These search trends suggest that residents in high- and low-exposure states are likely to have responded differently to the breach, lending support to our identification strategy.

Figure A2: Individuals' Response to the Anthem Breach

Panel A: Number of Inquiries




Panel B: Number of New Accounts




Notes: Authors' calculations using data from Federal Reserve Bank of New York Consumer Credit Panel/Equifax. These figures show difference-in-differences results comparing individuals who lived in states that were most impacted by the Anthem breach to individuals who lived in states that were the least impacted.

Figure A3: Equifax Breach Difference-in-Difference Results: Prior Fraud Victims vs. Non-Victims



Notes: Authors' calculations using data from Federal Reserve Bank of New York Consumer Credit Panel/Equifax.

Figure A4: Equifax Breach Difference-in-Difference Results: Prior Breach Victims vs. Non-Victims



Notes: Authors' calculations using data from Federal Reserve Bank of New York Consumer Credit Panel/Equifax.

Figure A5: Equifax Breach DID Results: Prior Breach Victims vs. Prior Fraud Victims



Notes: Authors' calculations using data from Federal Reserve Bank of New York Consumer Credit Panel/Equifax.

Table A1: Difference-in-Difference Results

	Inquiries	Number of New Accounts Opened
Panel A: Prior Identity Theft vs. Non-Victims		
<i>Treat</i> × <i>Short-run</i>	-0.054*** (0.008)	-0.0063*** (0.0019)
<i>Treat</i> × <i>Long-run</i>	-0.076*** (0.007)	0.0001 (0.0016)
<i>N</i>	63,420,083	69,511,672
Panel B: Prior Breach vs. Non-Victims		
<i>Treat</i> × <i>Short-run</i>	0.0060*** (0.0007)	0.0015*** (0.0002)
<i>Treat</i> × <i>Long-run</i>	0.0223*** (0.0007)	0.0067*** (0.0002)
Panel C: Prior Identity Theft vs. Prior Breach		
<i>Treat</i> × <i>Short-run</i>	-0.057*** (0.008)	-0.007*** (0.002)
<i>Treat</i> × <i>Long-run</i>	-0.092*** (0.007)	-0.005** (0.002)
<i>N</i>	23,848,944	28,532,171

Note: Authors' calculations using data from Federal Reserve Bank of New York Consumer Credit Panel/Equifax. *Short-run* = 1 for time periods Q3:2017 to Q4:2017 and *Long-run* = 1 for time periods Q1:2018 to Q4:2018.