FEDERAL RESERVE BANK
PHILADELPHIA

Working Papers
RESEARCH DEPARTMENT

Financial Fraud Through the Lens of Extended Fraud Alerts

Nathan Blascak
Federal Reserve Bank of Philadelphia
Consumer Finance Institute

Julia Cheney
Federal Reserve Bank of Philadelphia
Consumer Finance Institute

Robert M. Hunt
Federal Reserve Bank of Philadelphia
Consumer Finance Institute

Vyacheslav Mikhed
Federal Reserve Bank of Philadelphia
Consumer Finance Institute

Dubravka Ritter
Federal Reserve Bank of Philadelphia
Consumer Finance Institute

Michael Vogan
Upstart

WP 25-29

PUBLISHED
October 2025

FRB CFI
PHILADELPHIA

ISSN: 1962-5361

Disclaimer: This Philadelphia Fed working paper represents preliminary research that is being circulated for discussion purposes. The views expressed in these papers are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors. Philadelphia Fed working papers are free to download at: https://www.philadelphiafed.org/search-results/all-work?searchtype=working-papers.

DOI: https://doi.org/10.21799/frbp.wp.2025.29

Financial Fraud Through the Lens of Extended Fraud Alerts*

Nathan Blascak Julia Cheney Robert Hunt Vyacheslav Mikhed
Dubravka Ritter Michael Vogan†

October 8, 2025

Abstract

We use extended fraud alerts in anonymized credit reports to examine how identity theft, and subsequent clean-up, affects consumers' credit outcomes. The immediate effects of fraud for these consumers are negative, relatively small, and transitory. After placing an alert, these consumers experience persistent declines in delinquencies and a 12-point increase in credit scores, and 11 percent of filers become prime consumers. Many of these consumers take advantage of their improved creditworthiness and obtain additional credit. Although alert filers have larger balances, their performance on loans is as good as or better than before fraud, suggestive of a change in behavior following fraud.

Keywords: identity theft, fraud alert, consumer credit, credit performance
JEL Codes: G51, D14, D18

*We wish to thank Dennis Carlson, Amy Crews Cutts, Bradley Dear, April Ferguson, and Henry Korytkowski for their assistance with the data. We thank Marieke Bos, Anat Bracha, Chris Carroll, Claire Célérier, Ronel Elul, Dan Grodzicki, Andrew Hertzberg, Lynn Langton, Wenli Li, Brianna Middlewood, Susan Herbst-Murphy, Blake Prichard, Peter Schnall, Victor Stango, Jialan Wang, Chet Wiermanski, Stephanie Wilshusen, and anonymous referees for their helpful suggestions and Bryant Wright for his excellent research assistance. We especially thank Loretta Mester for making this research possible. We thank seminar participants at the Federal Reserve Bank of Philadelphia, the Federal Reserve Bank of Cleveland, Charles River Associates, the Boulder Summer Conference on Consumer Financial Decision Making, the FDIC Consumer Research Symposium, the Cherry Blossom Financial Education Spring Institute, the Digital Information Policy Scholars Conference, the Workshop on the Economics of Information Security, the RAND Behavioral Finance Forum, the Second Quadrant Behavioural Finance Conference, and the Public Policy Conference on the Law & Economics of Privacy and Data Security for their comments.

†Blascak, Cheney, Hunt, Mikhed, Ritter: Consumer Finance Institute, Federal Reserve Bank of Philadelphia; Vogan: Upstart.

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 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

According to a U.S. Department of Justice (DOJ) report, in 2021 approximately 24 million persons age 16 or older (about 9 percent of the U.S population) reported exposure to identity theft over the preceding twelve months. These victims reported $16.4 billion in losses, with 60 percent of victims experiencing a loss of more than $1 and the average loss being $1,160 (Harrell and Thompson, 2023). Moreover, 1 in 5 respondents reported being a victim of identity theft at some point in their lifetime. Along with the incidence and scope of out-of-pocket losses, the number of reported identity theft cases is continuing to increase. The Federal Trade Commission hosts a complaint system – The Consumer Sentinel Network – where consumers submit reports about problems including identity theft. The number of identity theft complaints submitted to this network increased by a factor of five from 2011 to 2021.[1]

In addition to out-of-pocket losses, identity theft and fraud can have serious consequences for consumers' credit profiles. Identity theft victims may have accounts opened in their names without their knowledge or experience an increase in balances, missed payments, and other derogatory information in their credit reports. If not addressed by the consumer, this kind of negative and erroneous information in credit report records can have detrimental effects on consumers' finances and the access to and prices in credit, rental, and insurance markets.

In this paper, we analyze one of the tools available to consumers to mitigate the effects of fraud in their credit reports: extended fraud alerts. Extended fraud alerts are user-initiated flags that are placed in individual's credit report when they have been a victim of fraud, including identity theft. When consumers place extended fraud alerts, they (1) must submit detailed information regarding the fraudulent account(s) on their credit report and (2) receive a copy of their credit report, allowing them to dispute any additional fraudulent information or errors in their reports. Credit bureaus are required to block fraudulent information and fix any errors, while creditors must take additional steps to verify a consumer's identity when opening new accounts. Since extended alert filers must provide evidence of identity theft, this alert cannot be filed purely out of precaution. While this tool has the potential for substantial benefits for consumers, little research has been done to study the effect or efficacy of these alerts.

What are the experiences of consumers who are victims of identity theft and subsequently file an extended fraud alert? To document the changes in credit outcomes prior to and following

[1]For more details, see: `https://www.ftc.gov/system/files/ftc_gov/pdf/csn-annual-data-book-2024.pdf`.

fraud and extended alert filing, we use an anonymized 5 percent random sample of quarterly credit records of U.S. consumers from the credit bureau Equifax, supplemented with a novel data set containing detailed information on consumers' fraud alert placements. Our credit bureau panel tracks consumers in the form of data snapshots in a moment in time (otherwise known as an "as was" data view). The quarter-by-quarter nature of the data delivery enables us to observe potentially fraudulent credit activity (including new accounts or balance increases) because any such fraudulent activity is not removed or edited from our data. Our sample consists of the approximately 50,000 consumers who filed an extended fraud alert between Q1:2008 and Q3:2013 and we observe their credit history for several years before and after the fraud incident and alert filing. Leveraging the plausibly exogenous variation in the timing of extended alert filing, we estimate a standard event study model based on the Sun and Abraham (2021) methodology and document the evolution of consumer credit characteristics around the time of alert filing and for several years afterward.

With our data and empirical strategy, we provide evidence on three broad research questions. First, we document changes in consumer characteristics that are consistent with fraud that occurred prior to or around the time of extended alert filing. Second, we examine consumer credit outcomes after the immediate effects of fraud are removed. Third, we provide evidence of long-term consequences for extended alert filers in terms of their creditworthiness, borrowing, repayment, and loan default.

For the first question, we document that just preceding and at the time of extended alert filing, consumers see additional credit inquiries (applications for credit), new credit cards opened in their name, and a spike in the rate of address reversals (a change to a temporary address followed by a change back to the original address). These changes are consistent with criminals applying for credit with stolen consumer information, being approved for new credit using the stolen information to impersonate actual consumers, and deceiving lenders into sending credit products (e.g., new credit cards) to new addresses where criminals can collect them. As a result, the credit standing of victimized consumers is negatively affected. Equifax Risk Scores (Risk Scores)[2] of extended alert filers decline by 2.5 points, on average, in the quarter prior to alert filing. For some of these consumers, the effects are larger and as a result, their Risk Scores become subprime following the fraud.

Second, among extended fraud alert filers, we find that many of the negative consequences

[2]We use the term *Risk Score* or *Score* to refer to a proprietary credit score provided to us by Equifax. Subprime borrowers are defined as those with Scores less than or equal to 660, while prime borrowers have Scores above 660.

of fraud quickly disappear from credit bureau records following alert filing. In particular, credit inquiries, the number of new credit cards, and address reversals all decrease significantly and quickly. Further, we observe positive changes to consumer credit attributes initially unaffected by fraud. For example, we observe significant and persistent reductions in the incidence of third-party collections and major derogatory events, while the share of card balances in good standing increases. The removal of fraudulent information from these individuals' credit bureau files is insufficient, in itself, to explain these incremental positive changes that persist for many years following fraud.

Consistent with the improvements in individual credit characteristics, we find that the Risk Scores of extended alert filers increase by an average of 11 points after the immediate effects of fraud fade. For many of these consumers, this increase (11 points) is larger than the decrease (2.5 points) in the Risk Score due to fraud. We also observe that the proportion of extended alert filers with prime Scores increases by about 5 percentage points (a relative increase of about 11 percent) after filing an extended fraud alert. As with the improvements in the incidence of derogatory events and loan performance metrics, the increase in consumers' Risk Scores persists for up to 20 quarters after alert filing.

We also document that many extended alert filers appear to use their improved creditworthiness after fraud to apply for additional credit. In particular, we find that these consumers apply for new auto loans and increase auto loan balances. Similarly, some of these consumers apply for new mortgages and hold more mortgage debt. Although these consumers have additional mortgages and auto loans and increased balances on them, their loan performance is as good as or better than before fraud. The finding that some fraud victims use their improved creditworthiness to apply for new credit is consistent with prior studies showing that consumers apply for credit after positive shocks to their credit scores.[3]

The persistent improvements in Risk Scores, major derogatory events, collection accounts, and card balances in good standing suggest that many extended fraud alert filers were not actively scrutinizing their credit reports prior to the fraud. Many of these attributes improved more than they deteriorated due to fraud, and remained persistently better long after fraud alert filing. This suggests that consumers were not aware of preexisting errors in their reports and corrected those shortly after being victimized. Taken together, these results are suggestive that, among extended

[3]For example, Gross, Notowidigdo, and Wang (2020) show that borrowing on cards, auto loans, and mortgages rises after credit score increases following bankruptcy flag removals. Similar effects are documented by Musto (2004); Bos, Breza, and Liberman (2018); Herkenhoff, Phillips, and Cohen-Cole (2021); and Dobbie et al. (2020).

alert filers, an experience of identity theft may represent a shock to their attention.[4] Our results – and this interpretation – may not be generalizeable to all identity theft victims as the salience of the fraud-related information may not be as pronounced for milder fraud experiences, but helps to fill in the knowledge gap that currently exists in the literature on consumer experiences with financial fraud.

2 Contributions to the Literature

This paper contributes to several existing literatures. First, our paper relates to a large and growing literature documenting the effects of the removal of negative credit information and consumers' subsequent credit behavior. For example, Gross, Notowidigdo, and Wang (2020) show that borrowing on cards, auto loans, and mortgages rises after credit score increases following bankruptcy flag removals. Similar effects are documented by Musto (2004); Bos and Nakamura (2014); Bos, Breza, and Liberman (2018); Herkenhoff, Phillips, and Cohen-Cole (2021); and Dobbie et al. (2020).[5] We also document that the credit expansion does not lead to additional defaults in the long run for this population of consumers. This finding stands in contrast to some prior studies of credit expansions under different circumstances (e.g., Musto, 2004, finds that borrowers are more likely to default on credit obtained after a bankruptcy flag removal). Overall, our findings imply that identity theft is not likely to lead to long-term credit damage or consumer withdrawal from credit markets for extended alert filers. On the contrary, identity theft events among fraud alert filers may serve as a *teachable moment* for those consumers who were previously not knowledgeable about credit bureau records or consumer credit markets more broadly.

Second, our paper contributes to the literature on fraud in financial markets. While much of this literature focuses on the parties that commit fraud, such as financial advisors (Dimmock, Gerken, and Graham, 2018; Dimmock and Gerken, 2012; Egan, Matvos, and Seru, 2019; Qureshi and Sokobin, 2015), CEOs (Khanna, Kim, and Lu, 2015; Agrawal, Jaffe, and Karpoff, 1999), and firms (Piskorski, Seru, and Witkin, 2015; Povel, Singh, and Winton, 2007; Dyck, Morse, and Zingales, 2010, 2024), we examine the effects of fraud on victims' credit outcomes. Our study complements

[4]It is also consistent with other research documenting that many individuals do not pay regular attention to their credit reports, which in turn can lead to errors in credit files persisting (e.g., Federal Trade Commission, 2012, finds that 26 percent of consumers have material errors in their credit reports and 13 percent experienced credit score changes after the errors were corrected).

[5]In a different context, Guttman-Kenney (2025) shows that suppressing negative credit bureau information using disaster flags available in credit reports increases credit scores.

research on the effects of fraud on investment decisions by individuals and households (Gurun, Stoffman, and Yonker, 2018; Giannetti and Yang, 2016) and on the prevalence and cost of financial fraud victimization (Hamdi et al., 2024; Knüpfer et al., 2024) by showing evidence of and the fallout from financial fraud among consumers who file extended fraud alerts.

Third, our empirical findings add to the literature that examines the consequences of financial fraud on consumers and consumers' use of fraud alerts to mitigate the effects of identity theft. However, unlike previous studies that focused on consumer confidence in payment systems (e.g., Sullivan, 2010) and payment choice (Cheney et al., 2012; Kahn and Liñares-Zegarra, 2016; Stavins, 2013; Kosse, 2013), this paper examines how financial fraud and the use of fraud alerts to mitigate it can affect consumers' credit performance and credit outcomes. This study is also related to papers considering the tradeoff between information security and data privacy (Acquisti, 2004; Anderson and Moore, 2007) and incentives for consumers to prevent identity theft (Blascak and Toh, 2022; Federal Trade Commission, 2003; Cheney, 2003).

3 Consumer Credit Bureau Records and the Use of Extended Fraud Alerts in the United States

3.1 Information Contained in Consumer Credit Bureau Records

In the United States, a consumer credit report (also referred to as credit file) is an organized record of an individual's interaction with the credit market. Typically, such a report will include information on the number, size, age, composition, and repayment status of the consumer's loans or lines of credit. A credit report may also include information obtained from public records, such as bankruptcy filings. In the United States, the three largest credit reporting agencies with national scope are Equifax, Experian, and TransUnion. Anonymized credit bureau data, such as that used in this study, has been used extensively both in academic research and in policy analysis (see Gibbs et al., 2025 for details on these data sets).

3.2 Extended Fraud Alerts

In 2003, the Fair and Accurate Credit Transactions Act (FACTA) became law, amending the Fair Credit Reporting Act (FCRA) of 1970. One of the goals of FACTA was to improve protections for consumers affected by fraud, including identity theft. FACTA required federal regulators to

develop "red flag" indicators of identity theft to aid in detecting fraud. It also required credit reporting agencies to block information that results from fraud and to implement a set of indicators (credit file flags) that inform creditors that a consumer is or may be a victim of fraud, including identity theft. The credit file flags include the extended fraud alerts that we use in this paper.

An important element of the rights established in FACTA (and some state laws) is the opportunity for the consumer to obtain a copy of his or her credit report at no cost from each of the three credit bureaus when filing a fraud alert, including an extended fraud alert. Receiving these reports gives consumers a chance to detect and dispute fraudulent accounts or delinquencies on compromised accounts as well as any other errors in their credit reports. If the information in a consumer's credit report cannot be verified by the creditor, the credit bureaus are required to remove this information and to prevent it from reappearing in subsequent reports. It is important to note that requesting a credit report or filing a fraud alert by itself does not remove fraudulent charges from credit accounts with individual creditors and does not prevent data on already open but not-yet-disputed fraudulent accounts from being added to the credit report. Even after filing an alert, consumers need to identify fraudulent information and dispute this information.

For a consumer to place an extended fraud alert in their credit bureau file, the filer must submit a police report or an Identity Theft Report to a credit reporting agency.[6] An Identity Theft Report requires submission of detailed information on the accounts that were compromised and accompanying evidence of fraud. Extended fraud alerts require a creditor to take additional steps in verifying the consumer's identity when a request is made to open a new credit account, increase an existing credit line, or issue an additional card associated with an existing credit account. The consumer specifies a telephone number or other reasonable contact method as part of the alert documentation. All creditors must contact the consumer by the method specified in the alert to verify the consumer's identity in the case of any of the applicable scenarios. Once filed, an extended fraud alert remains in a consumer's credit file for seven years unless the consumer chooses to remove it earlier than that. In addition, an extended fraud alert removes the consumer's credit file from lists of prescreened credit and insurance offers for five years. Under FACTA, when a consumer files an alert with one national credit bureau, the information submitted by the consumer is communicated

[6]FACTA, §111, defines an *Identity Theft Report* as, at a minimum, "a report that alleges an identity theft; that is a copy of an official, valid report filed by a consumer with an appropriate Federal, State, or local law enforcement agency, including the United States Postal Inspection Service, or such other government agency deemed appropriate by the Federal Trade Commission; and the filing of which subjects the person filing the report to criminal penalties related to the filing of false information if, in fact, the information in the report is false."

to the other two major bureaus.

Providing the required evidence of fraud requires both time and effort by or on behalf of the consumer. In addition, consumers face criminal penalties for falsifying information in these reports. Thus, filers of extended fraud alerts are unlikely to place these alert flags in their credit bureau files simply because of worry, out of an abundance of caution, or for a related reason.[7] In other words, the institutional features ensure that extended alert filers are almost certainly the victims of identity theft.

4 Data Description

To explore the effect of identity theft on consumer credit, we use the Federal Reserve Bank of New York Consumer Credit Panel/Equifax Data (CCP), combined with a unique data set detailing the timing (placement) and type of fraud alerts obtained from Equifax. The CCP contains credit characteristics for an anonymized 5 percent random sample of credit bureau records of U.S. consumers.[8] The CCP is an unbalanced panel in which new individuals are included over time as they obtain or first report an SSN to a lender (e.g., after immigrating to the United States), open their first credit accounts, or establish their first public record. Similarly, consumers are dropped from the sample when they die or "age off" following a prolonged period of inactivity and no new items of public record. The sample is designed to produce a panel with entry and exit behavior similar to the population that uses credit or has a credit history (Lee and van der Klaauw, 2010).

We begin with a sample of credit files of about 10.8 million individuals continuously present in the data set in all quarters from Q1:2006 to Q3:2013 so that we can trace the credit histories of these consumers and mitigate concerns about "fragments" in our data (Wardrip and Hunt, 2013). We then restrict our sample to 50,103 consumers who filed an extended fraud alert between Q1:2008 and Q3:2013 and observe their credit history for several years before and after an alert filing.[9] The

[7] In Section 4.2, we provide additional information about fraud alert filers.

[8] The sample is constructed by selecting consumers with at least one public record or one credit account currently reported and with one of five numbers in the last two digits of their SSNs as the method of randomly selecting the sample. Equifax uses SSNs to assemble the data set, but the actual SSNs are not shared with researchers. In addition, the data set does not include any names, actual addresses, demographics (other than age), or other codes that could identify specific consumers or creditors. Our data on fraud alerts span Q1:2008 to Q3:2013.

[9] It is important to emphasize that our analysis sample consists of consumers who *filed* their extended fraud alert during our sample period. Once an extended alert is filed, it is present in the credit file for many years, so we need to distinguish between the quarter in which the alert is placed in the file and the subsequent quarters during which the alert is active. To be precise, we use the date of filing contained in the alert. As discussed later, our sample is designed to capture consumers around the time they experience identity theft, so we focus on consumers who *filed*

distribution of the initial date of extended alert flag placement across calendar time is displayed in Figure 1. In much of the following analysis, we examine changes in consumer credit attributes in *event time* — the number of quarters before or after an extended fraud alert first appears.

Within the CCP, we have access to rich consumer-level information on mortgage accounts, home equity revolving accounts, auto loans, bank card accounts, student loans, and other loan accounts as well as public record and collection agency data. The CCP contains limited personal background information, such as the consumer's age (based on year of birth) and geographic information in the form of a scrambled address, state, zip code, metropolitan statistical area, and U.S. Census tract and block. We also have a credit score (specifically, the Equifax Risk Score) and the number of inquiries (i.e., applications for credit or insurance). We examine the number of and balances on bank card accounts, the proportion of cards in good standing, total credit card limit, prevalence of major derogatory events, prevalence of third-party collections, and other consumer characteristics for several quarters before and after alert filing. Table 1 presents the descriptive statistics for the sample of extended alert filers in our data set.

4.1 Timing and Evidence of Fraud Removal

FACTA requires that U.S. credit reporting agencies block information resulting from identity theft four days after accepting a consumer's dispute identifying this information. The agencies must notify information furnishers (lenders, servicers, etc.) that the information they submitted will be blocked from the consumer's credit file. This notification triggers actions required by FACTA for furnishers of the information, including that the furnisher may not continue to report this information to any credit reporting agency. Another option available to all consumers, not just identity theft victims, through the FCRA is the right to dispute errors (inaccurate or incomplete information) in credit reports. When such a dispute is verified, it may result in a change to or deletion of information in a consumer's credit report.

We cannot directly observe what kind of information is blocked or for what reasons. However, the manner in which each quarter of the CCP data is assembled implies that any fraud existing in the quarters *preceding* the filing of an extended fraud alert remains in our data. That is because, generally speaking, when a new quarter of data is added to the CCP, the information contained in the previous quarters is not revised. In this sense, this data set is similar to other real-time

an extended fraud alert during our sample period, not the larger group of consumers who have an extended fraud alert in their credit file.

data sets used by researchers. It is important to emphasize that this property of our data does not necessarily apply to the actual credit report information that consumers and creditors access every day. When an error is discovered in information contained in those credit bureau files, the erroneous information no longer appears anywhere in the credit history that a consumer or a creditor can see.[10] We utilize this specific nature of the CCP to detect evidence of both identity theft and its resolution in consumer credit attributes, as discussed in the next sections.

4.2 Describing Fraud Alert Filers

As mentioned in the previous section, the CCP does not contain any demographic or socioeconomic information. To better understand the characteristics of extended alert filers, we use data from the Identity Theft Supplement (ITS) of the National Crime Victimization Survey (NCVS). The NCVS is a nationally representative survey of individuals' experience with crime in the U.S. conducted by the DOJ's Bureau of Justice Statistics. The ITS is conducted every two years to individuals age 16 or older who complete an in-person interview for the NCVS. Given the relatively high response rates to the ITS, each wave of the ITS contains approximately 60,000-90,000 observations. We use the 2012, 2014, and 2016 waves of the ITS in a combined data set to compare identity theft victims and alert filers.

In Table 2, we examine a number of demographic characteristics of identity theft victims and extended alert filers in the NCVS ITS data. The majority of identity theft victims who filed an extended fraud alert are White, have higher incomes, and are home-owners. Alert filers also tend to be more highly educated (having a bachelor's degree or higher). These characteristics are suggestive of having the capacity to navigate the complex process of filing an extended fraud alert. As can been seen in the bottom row of Table 2, most victims of identity theft or fraud do not file a fraud alert with a credit bureau.[11] However, victims of severe identity theft –defined as those who experienced an identity theft incident that involved the opening of new accounts or the misuse

[10] It is possible that the timing of the placement of extended fraud alerts may not coincide perfectly with changes in credit variables. For example, consumers who file their alerts at the end of the third month of a quarter may not have their credit file updated until the first month of the following quarter. We considered the changes in key credit variables across event time by the month of extended alert filing to address this concern. Our results (available in Appendix Figure A1) indicate that both the timing of fraud and the effect of placement of fraud alerts do not systematically differ by the filing month.

[11] According to Harrell (2019), the NCVS identifies about 8 percent of all identity theft victims as those contacting a credit bureau following identity theft. Among consumers suffering from more severe forms of identity theft, such as opening new accounts in the consumer's name, a much higher percentage (about 33 percent) of victims contact a credit bureau, and about a third of those provide a police report to the credit bureau.

of their personal information –are 9 times more likely to file an extended fraud alert than other identity theft victims. Although extended alert filers selected into filing an alert, based on these data, their characteristics are similar to the overall population of identity theft victims.

5 Empirical Strategy

To identify the effect of extended alert flag placement on consumer credit outcomes, we leverage plausibly exogenous variation in the timing of extended alert filing for our sample of alert filers. The staggered nature of alert placement by consumers allows us to specify a standard event study framework. Often, the event study design is estimated via a two-way fixed effects (TWFE) regression specification, with unit and time fixed effects (Miller, 2023). However, it is now well known that this canonical TWFE regression yields consistent results only when treatment effects are homogeneous; estimated treatment effects will be biased if treatment effects are heterogeneous (de Chaisemartin and D'Haultfoeuille, 2020; Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021). Particular relevant in our context, Sun and Abraham (2021) show that in the presence of staggered treatment timing *and* treatment effect heterogeneity by treatment cohort, TWFE event study estimates for any given event time period in an event study model will be contaminated by the effects from other event time periods that are being estimated. This issue may arise if the effects vary by the alert placement cohort. For example, if flag filers earlier in our sample are more attentive and/or informed in different ways than later filers because of increased awareness or education, this may lead to different size treatment effects for placing an alert in their credit bureau files.

To address this concern, we estimate our event study models using the interaction-weighted estimator proposed by Sun and Abraham (2021) (henceforth SA). This estimator has multiple advantages in our context. First, it is consistent in the presence of heterogeneous treatment effects, which may be present given that we have over 20 cohorts of extended alert filers. Second, the SA estimator accommodates our analysis sample, which has no never-treated group. Because of this feature of our sample, we use the last treated cohort of individuals (extended flag filers in Q3:2013) as the control group. One consequence of using the last treated cohort as our control group is that we have to further restrict our analysis sample and exclude any quarters of data where the last treated cohort receives treatment. This necessarily requires us to trim our analysis sample to be from Q1:2006–Q2:2013. We also restrict our panel data to be balanced in calendar time, as the SA estimator is valid for balanced panels. To study the immediate and long-run effects of extended

alert filing, and given the sample restrictions for the SA methodology, we include up to 20 event time quarters post flag filing and 22 event time quarters prior to flag filing.

Our event study estimating equation takes the following form:

$$Y_{it} = \delta_t + \mu_i + \sum_{j=-8}^{j=20} \beta_j D_{ij} + \Omega X_{it} + \epsilon_{it}. \tag{1}$$

Y_{it} are our credit outcome variables of interest for consumer i in time t. D_{ij} contains our event time dummy variables, which are equal to one if consumer i has an extended fraud alert in their credit report j quarters from the initial extended alert flag filing at time $j = 0$. δ_t are calendar time fixed effects and μ_i are individual fixed effects. Respectively, these fixed effects allow us to control for fixed differences across individuals and general trends in credit characteristics for individuals in our sample. X_{it} includes an age polynomial, state fixed effects, and state by calendar time fixed effects. Our omitted time period are event quarters from -22 to -9; this implies that the coefficients for the event time indicators are the mean differences from the average value of the outcomes in these periods.[12] Standard errors are clustered at the individual level. Our coefficients of interest are the β_js, which are the coefficients for the event time dummy variables D_{ij}. The β_js estimate the effect of having an extended alert flag in individual i's credit report on outcome Y_{it} j quarters before or after the initial alert placement at $j = 0$.

6 Extended Fraud Alerts, Identity Theft, and Consumer Credit Outcomes

In this section, we use the econometric strategy described in Section 5 to examine three research questions. First, we examine key measures of credit activity for extended fraud alert filers around the quarter of alert filing. Second, we explore their credit performance after alert filing, including any cleanup of fraudulent information from credit records. Third, we document long-term trends in credit performance, borrowing and use of credit, loan default, and other consumer interactions with the credit markets up to five years following alert filing.

We summarize our results for these three lines of inquiry in Figures 2–5. The figures report event study coefficients for a number of outcomes, estimated using our previously described event study

[12] Also, as noted in Baker et al., (2023), in the case of no never-treated units in the sample, multiple relative time indicators need to be omitted to avoid multicollinearity.

models and the interaction-weighted SA estimator, as specified in Equation (1). The coefficients show the difference in each outcome between the treatment and control groups a certain number of quarters before or after the extended alert placement. In addition to point estimates displayed as dots, all figures provide 95 percent confidence intervals as vertical bands. In the following subsections, we discuss our results for all consumers and for some consumers with potentially more pronounced effects.

6.1 Evidence of Identity Theft in Credit Bureau Files

We document changes in four key credit bureau metrics preceding extended alert filing which are consistent with evidence of identity theft, in Figure 2. These potential indicators of fraud include new credit applications, new bank card accounts, reversal of address changes, and changes in credit score. We show the evolution in these measures around the time of extended alert filing in the four panels of Figure 2 and discuss them in detail in the following paragraphs.[13]

Panel A of Figure 2 displays a very large and transitory increase in the number of credit applications that coincides with the quarter the extended alert is filed. Relative to the base period (quarters 22 to 9 before extended alert filing), the average number of inquiries is increasing in the quarters leading up to alert filing, peaking at an additional 0.6 inquiries (relative to the sample average of 1.1) in the quarter of alert filing. This increase is consistent with consumers' personal information being stolen by criminals and used to apply for credit. Consumers often become aware of identity theft because this spike in applications triggers letters or phone calls from creditors. Data from the NCVS shows that almost 50 percent of identity theft victims discover identity theft through such communications. The number of inquiries decreases to pre-fraud levels by the third quarter after extended alert filing and remains on a downward trend in the long term.

Panel B of Figure 2 plots the average number of new bank card accounts before and after extended alert filing, which begins to increase sharply a few quarters before the extended alert filing and peaks one quarter before filing. Cumulatively, 1 out of 7 extended alert filers have an additional new bank card opened during the year prior to alert filing. This finding is consistent with criminals using consumers' stolen personal information to open new bank card accounts. The number of new bank card accounts declines quickly once the extended fraud alert is filed and

[13]We note that there is wide variation in the length of time consumers' information is typically used by fraudsters before being discovered, as displayed in Figure A2. For half of extended alert filers in the NCVS, their information was used for a week or less, but 30% of filers saw their data used for more than 1 month before being discovered, and 17% of filers more than 3 months.

continues to gradually decrease for up to 5 years after fraud. This decrease may reflect both behavioral changes on the part of consumers and the additional verification lenders must undertake in order to open new accounts for consumers with extended fraud alerts.

More serious forms of identity theft and subsequent fraud involve criminals changing the address on the consumer's financial accounts, which can trigger a change in the address that creditors report to credit bureaus.[14] In our data, we are unable to distinguish between fraudulent and genuine address changes. However, we can see if an address change is reversed to the original address in the subsequent quarter. Thus, we can compare the pattern of *address reversals* at the time an extended fraud alert is filed with patterns in this phenomenon prior to and after the event.[15] Panel C of Figure 2 shows the share of extended alert filers who revert their address to the prior quarter's address, relative to periods 20 to 9 before alert filing. The coefficients imply that around 1 percent of victims reverse address changes at the time of alert filing and an additional 1.5 percent of victims do the same one quarter after alert filing. Thus, we find evidence of a sharp increase in reverse address changes at the time the extended alert is filed and one quarter after, consistent with the consumer reversing address changes made by criminals.

Finally, panel D of Figure 2 shows a transitory decline in Risk Score of about 2.5 points shortly before the fraud alert filing and a subsequent recovery in the quarters after filing. However, the average increase in Scores that follows is considerably larger than the transitory decline. On average, Risk Scores increase by about 11 points relative to the omitted periods at the time of the alert filing; we will revisit this finding in the next subsection.

The increases in inquiries, address reversals, and number of new bank cards near the time of the extended alert filing, as well as the decline in Risk Scores shortly before the placement of the extended fraud alert, provide strong evidence that identity theft occurred within a couple of quarters before the extended fraud alert was filed, if not exactly in the same quarter in which the alert was filed.

[14] Criminals may change addresses when taking over existing accounts, or they may apply for new accounts using the consumer's name but a different address.

[15] Recall that consumer address changes may be reversed in the credit bureau file after the discovery of fraud, but the history of address changes in the CCP is not updated and, therefore, is not affected by the reversal in the "live" credit bureau file.

6.2 Persistent Changes in Credit Performance After Extended Alert Filing

In addition to the previously discussed indicators of identity theft in quarters just before and at the time of extended alert filing, panels A, B, and D of Figure 2 allow us to observe differences in the credit outcomes in the medium to long term, up to five years following extended alert filing. In particular, these panels show the number of credit inquiries, the number of new bank cards, and Risk Scores for up to 20 quarters after extended alert filing. While panel D of Figure 2 shows that, on average, victimized consumers' Risk Scores are 10 or more points higher even five years after extended alert filing, these consumers do not apply for credit as much as they did before fraud, and they have, on average, 0.2 *fewer* credit inquiries per quarter in this period (Figure 2, panel A). Consistent with this reduced number of inquiries, the number of new bank cards is about 0.04 (relative to the sample average of 0.1) lower each quarter four to five years out (panel B). These results also indicate that the persistent improvement in Risk Scores may be explained in part by reductions in the number of inquiries and the number of new bank cards after fraud, as reductions in both of these indicators of credit demand generally positively affect credit scores.

To examine the reasons behind the persistent improved creditworthiness of fraud alert filers, we consider measures of credit performance such as the share of card balances in good standing, the incidence of major derogatory events, and the incidence of third-party collections. Figure 3 provides evidence on the performance of extended alert filers with credit products several years after fraud. Panel A shows that consumers keep a higher proportion of their card balances in good standing (current) up to five years after extended alert filing. They also maintain a lower incidence of major derogatory events by about 4-7 percentage points (panel B) and a lower incidence of third-party collections by about 5 percentage points (panel C). These three measures of credit performance capture different margins of adjustment in debt repayment behavior. Card balances in good standing represent the strictest definition of performance (i.e., repaying debts on time without any delay). Major derogatory events capture more serious delinquency such as charge-offs, foreclosure, and internal collections. Third-party collections are delinquent accounts placed for collection with external firms that specialize in recovering at least a portion of an outstanding debt that a consumer owes.

The sharp declines in the incidence of derogatory events and third-party collections at the time of an extended alert filing likely result from consumers disputing fraudulent accounts and other incorrect information in their credit reports. However, the persistence of these positive effects

suggests that these consumers changed their repayment habits to keep more credit accounts in good standing and out of collections, as shown by the long-term effects of fraud on these variables.

To summarize our findings, we plot the share of the population with prime Risk Scores (Scores higher than 660) in our sample in panel D of Figure 3. The proportion of prime score consumers is a good proxy for how material credit score changes resulting from consumers' changed behavior are. Consumers shifting scores from subprime to prime are most likely to benefit from significantly improved financing options. Consistent with the decline in the average Risk Score before extended alert filing, shown in Figure 2, this figure shows that fraud activity at event time $e = -1$ lowers the share of prime consumers by about 2 percentage points. However, after fraud alert filing, the share of prime consumers increases by about 3 percentage points relative to the base period, which is a 5.1 percentage point increase relative to the quarter before fraud. This is an 11 percent increase over the sample average of 47 percent of prime consumers. The share of prime consumers continues to grow over time, and it is 6 percentage points higher after five years relative to the baseline period. This is a substantial change, in both a statistical and economic sense.

It is important to note that the mechanical effect of credit file "cleaning" (i.e., the removal of fraudulent information from credit bureau files), while likely responsible for the initial gain in Risk Score immediately after the identity theft incident, is insufficient to explain the persistent positive changes in consumer behavior on *existing* credit accounts.

6.3 Consumer Use of Improved Credit Standing After Alert Filing

Our previous results that show large increases in both the average Risk Score and the share of consumers with prime Risk Scores in the years following fraud alert filing have potentially far-reaching economic consequences, as they may allow borrowers to obtain more credit and at better terms. For example, during our sample period, the annual percentage rate (APR) on a 30-year, fixed-rate mortgage decreased from 5.809 percent to 4.833 percent when a borrower moved from the 620–639 FICO score range to the 660–679 range.[16] Bracha and Meier (2015) show that moving from the 620–679 score range to the 680–739 range can decrease credit card interest rates by 3.5 percentage points (from 19.1 percent to 15.5 percent, on average). Thus, positive changes in the Risk Score may allow borrowers to save on financing expenses and have more access to credit to smooth negative income or expense shocks.

In this section, we consider how consumers use credit after alert filing and whether lenders

[16]This example is based on the national average mortgage interest rates provided by FICO on December 31, 2013.

alter the supply of credit to these individuals. While some fraud alert filers appear to be more creditworthy after cleaning up their files (as we document in the previous sections), we examine whether these consumers apply and receive more credit from lenders, which particular types of credit they receive, and how they perform on this additional debt. We show results for all alert filers in Figures 4 - 5; these figures depict the changes in the use of credit cards, auto loans, and mortgage loans.

For credit card borrowing, the majority of changes to credit card accounts appears to occur in the first year following a fraud alert filing. In Figure 4, panel A shows that around 1-2 percent of consumers become new bank card holders in a few quarters preceding and following alert filing. We define a new bank card holder as someone who transitions from having no bank cards in the previous quarter to having at least one card in the current quarter. It is important to note that the probability of becoming a new cardholder can increase because of identity theft or legitimate consumer activity. We also observe a significant decline in the probability to become a new cardholder after alert filing (starting in year 2). Thus, consumers seem to be reducing their new credit card acquisitions.

Figure 4, panel B, shows that more consumers reduce their total number of bank cards after initial increases in cards at the time of identity theft. On average, consumers have 0.13 fewer cards in the year after alert filing compared with their pre-fraud levels, which is a 6 percent decrease over the sample average of two cards per consumer. Total card balances (panel C) slightly decline in the year following alert filing, then largely revert to pre-alert levels in subsequent years. Notably, total card limits (panel D) are persistently significantly lower even five years after alert filing, mirroring the pattern in the number of bank cards in panel B and signifying –in combination with similar total balances relative to pre-alert –that consumers are closing inactive card accounts or not opening new ones. This can be a good strategy to limit fraud on credit cards as having fewer cards would reduce criminals' opportunities to commit fraud with them. Additionally, the demand for unsecured credit stays the same as before fraud (as card balances return to pre-fraud levels shortly afterwards).

We present the changes in automobile-related borrowing (auto loans) after alert filing in Figure 5. Panel A shows that around 0.5-1 percent of consumers in our sample become new auto loan holders (defined as switching from zero loans to a positive number of loans) in each of the three quarters following alert filing. This is a significant effect since, on average, only 2.3 percent of consumers become new auto loan holders per quarter in our sample (see Table 1). The average auto loan balances (panel B) increase significantly by up to $500 (a 7 percent change relative to the

mean), indicating that some consumers are opening new auto loans and increasing their borrowing on these loans with improved creditworthiness after resolving identity theft.

These results for auto loans are significant for a number of reasons. First, auto loans are used to finance purchases of automobiles, which are durable consumption goods. Thus, we find evidence that some identity theft victims use their improved creditworthiness to finance durable consumption. Second, we do not find evidence that consumers withdraw from secured credit markets after identity theft. To the contrary, some consumers use their improved creditworthiness to borrow more and finance consumption. Third, despite additional debt, these consumers still perform well on their credit as shown in Figure 3.

Next, we examine the use of mortgage credit after alert filing and summarize our results in Figure 5. Similar to the auto loan results, we find that a significant number of alert filers become new mortgage holders after identity theft (panel C). Around 0.5-1 percent of these consumers become new mortgage holders in each quarter post-fraud. Panel D shows that average mortgage loan balances increase as a result of new mortgage activity relative to the pre-theft levels by up to $12,000 (a 5 percent increase relative to the mean). This additional debt does not lead to increased financial distress as Figure 3 shows.

Overall, there is clear evidence of increased secured credit usage following the improvement in creditworthiness for fraud alert filers in the quarters and years following fraud. Despite holding additional auto and mortgage loans and increased balances for these credit products, consumers continue to perform no worse than before fraud on these accounts for at least five years following identity theft. Comparing across credit cards, auto loans and mortgages, we find a shift in the debt portfolio away from credit cards and into auto loans and mortgages, which may benefit consumers as they can increase their consumption of durable goods (cars and houses) and reduce their finance charges (interest rates). Note that consumers are applying for and obtaining this additional credit despite frictions introduced by alerts (identity verification, etc.).

In addition to the event study graphs, we present our results in a differences-in-differences (DID) specification in Tables 3 and 4. We estimate these effects using Equation 1, but replacing event study variables with two indicator variables for fraud time (event quarters -4 to -1) and post-fraud time (event quarters 0 to 20). Overall, the DID results are similar to our event study results in Figures 2 to 5. They show negative effects from identity theft in the quarters before alert filing (e.g., increases in address reversals and new cards) and positive effects in the quarters after filing an extended fraud alert, such as increases in auto and mortgage borrowing and declines in derogatory

events and collections.

7 Heterogeneity Analyses

7.1 Subprime-to-Prime Transition Group

While any increase in credit scores may improve credit access and terms of credit, crossing into a higher score category is especially beneficial. While there are several thresholds in scores affecting creditworthiness, the most significant is the subprime-prime distinction, as documented extensively in the previous literature. To understand whether changes in consumer credit outcomes are more pronounced for extended alert filers that cross into a higher score category, we apply our econometric strategy from Section 5 to the subsample of extended alert filers who crossed the prime Risk Score threshold after filing their alert. We define consumers as transitioning from the subprime to prime status if their Risk Scores are less than or equal to 660 in any of the four quarters before extended alert filing and are more than 660 at the time of alert filing or any of the four quarters after filing.

The changes in credit performance and borrowing of the subprime-to-prime transition subgroup of extended alert filers are presented in Appendix Figure A3. Comparing results in Appendix Figure A3 (subprime-to-prime transition group) to the results in Figures 2 and 5 (all borrowers), we can conclude that the subprime-to-prime transition group experiences improvements in Risk Scores that are three to four times the magnitude than those for all consumers who file extended alerts. For example, the initial increase in Risk Scores after alert filing was 11 points for all consumers, but this increase is around 40 points among the subprime-to-prime consumers. In general, the subprime-to-prime transition group experiences considerably larger improvements in Risk Scores, although the magnitude of the increase is more likely to diminish over time. Figure A3 also shows that the subprime-to-prime transition group is more likely than all alert filers to obtain new auto loans and new mortgage loans (panels C and D). However, the patterns of new credit card use (panel B) are similar in the subprime-to-prime transition group and the overall population of alert filers.

7.2 Consumers Without Credit Inquiries Before Alert Filing

To understand whether the effects of identity theft vary depending on the credit characteristics of the borrower, we study the effects of fraud on consumers without credit inquiries before and at

the time of extended alert filing and compare these effects with the results for all consumers. The presence of credit inquiries may capture two activities: (1) shopping for credit by consumers, and (2) shopping for credit by criminals using stolen consumer personal information. We hypothesize that consumers without inquiries may be (1) less attached to the credit market and less attentive to their credit information, or (2) subject to existing account fraud or other fraud that does not result in credit inquiries.

Results for the estimation of Equation (1) for consumers without inquiries in the three quarters before and the quarter of extended alert filing can be found in Appendix Figures A4–A5. The decline in credit inquiries in the four quarters before alert filing shown in panel A of Figure A4 is mechanical (we include only people without inquiries in this period), but the other results are not. Our results suggest that no-inquiry filers experience larger positive effects after identity theft on credit performance indicators (e.g., incidence of derogatory events and third-party collections, Risk Score). The results for no-inquiry victims in Figures A4–A5 are broadly similar to the results for the whole population of extended alert filers. These consumers reduce their credit card activity and keep more balances in good standing. These consumers also apply for credit after identity theft, as suggested by the number of inquiries in Figure A4. Overall, the findings for no-inquiry consumers suggest that these consumers are similar to other extended alert filers as they do not withdraw from credit markets and they use their improved creditworthiness to obtain new credit.

Studying consumers without credit inquires also allows us to test if our base results are affected by reverse causality, where, instead of consumers correcting credit reports in response to identity theft, some consumers may set out to clean their credit files in preparation for a mortgage or other major credit application. During this process, consumers may discover negative episodes in their reports — such as fraud — because they are actively applying for credit and paying more attention to their reports. This hypothesis implies that such consumers are likely to have indicators of fraudulent activity (e.g., address change reversals, new accounts that are closed immediately, and increases in delinquent accounts) in their files at *any* time before they file an extended fraud alert.

Our results summarized in Figure A4 do not support the hypothesis that consumers who file extended fraud alerts are simply engaged in credit file repair before a major credit application or some other event. In particular, panel C in Figure A4 shows address reversal (an activity indicative of identity theft) is tightly concentrated just before the alert filing or at the time of filing and not distributed across the quarters prior to alert filing. Similarly, we do not find evidence that fraud related activity (additional inquiries, new bank cards, or delinquent debt) are present in any

pre-alert-filing quarters. Analysis of the NCVS data supports our conclusion that individuals who experience identity theft were not typically in the process of shopping for credit, with approximately only 1 percent of respondents who were victims of identity theft stating that they discovered misuse upon applying for credit, bank accounts, or loans.

8 Discussion and Conclusion

This paper uses a unique data set of anonymized U.S. credit bureau records, including details on extended fraud alert filings, to examine the experiences of consumers who are victims of identity theft and subsequently file an extended fraud alert. The procedures necessary to file an extended fraud alert require fraud victims to file a police report or an Identity Theft Report (with accompanying evidence of identity theft and penalties for misrepresenting this information). Thus, these are consumers who are motivated by a sufficiently salient event to take action to protect themselves and to mitigate the damage experienced. The nature of the extended fraud alert filing process and our data—collected quarter by quarter on an "as was" basis, and thus reflecting potential fraud activity present in each quarter—provide us with unique advantages relative to other studies of consumer financial fraud.

We find suggestive evidence of identity theft shortly before extended alert filing, with declines in average Risk Scores of fraud victims and increases in new (likely fraudulent) card accounts, inquiries, and instances of address change reversals. The negative effects generally persist between one and two quarters. After these initial effects have passed, Risk Scores of extended alert filers rise by an average of 12 points, increasing the proportion of these consumers with prime Risk Scores by 5 percentage points, or 11 percent. For many extended alert filers, this effect is persistent over time and remains for as long as 20 quarters after fraud. We also find that, after filing an extended alert, these consumers have more card balances in good standing and a lower average incidence of derogatory events and third-party collections. The persistence of the reduction in the incidence of major derogatory and third-party collection events is particularly striking.

We also document that some extended alert filers take advantage of their improved financial standing to obtain additional credit. In particular, we find evidence that these consumers obtain additional auto loans and mortgages after alert filing and consolidate their use of credit cards while maintaining or improving their performance on existing card balances. We observe increases in the number of new loans and balances for these individuals. This finding is especially important because

car loans and mortgages are used to finance large durable consumption purchases or investment in real estate. Despite having additional loans and higher balances on those loans, these consumers manage their credit as well as or better than they did before financial fraud.

Our results are consistent with the idea that, prior to the event, many of these consumers were not focused on their credit reports. They were sufficiently motivated to take action and one result was a persistent improvement in their consumer credit attributes. This suggests those consumers were more attentive after the event than they were before. Several sources show that consumers do not pay close attention to their credit reports, credit scores, or other credit information. For example, according to a 2013 poll conducted by the National Foundation for Credit Counseling, 60 percent of adults 18 years or older had not checked their credit scores in the previous 12 months, and 65 percent had not reviewed their credit reports. Similarly, the DOJ found that only 42 percent of non-victims reported that they had checked their credit reports in the past 12 months; this number was considerably higher, at 62 percent, among consumers who had experienced identity theft (Harrell, 2019).[17]

It is possible that a fraud incident may increase the salience of credit information, temporarily increase the cost of acquiring/retaining credit, or encourage increased monitoring of credit reports and/or scores. Fraud alert filers may experience a number of negative feelings (e.g., shock, anger, anxiety) that may be action-inducing because of the seriousness of the event in a way that additional disclosures or reminders are not. Our calculations based on the NCVS data show that the number of identity theft victims who acknowledged checking their credit report increased by up to 15 percentage points upon victimization, and the number of victims who checked their bank or credit card statements increased by up to 26 percentage points. Although we cannot directly test the inattention hypothesis with our data, our empirical results are consistent with a change in behavior that occurred after the fraud incident, at least among extended alert filers. Definitively establishing this mechanism will require additional research.

[17]There is also a large and growing literature showing that individuals in a variety of contexts pay limited attention to and do not process information completely when making important decisions. Previous work has demonstrated that investors react less than optimally to information readily available to them at no cost (Barber and Odean, 2008; DellaVigna and Pollet, 2009; Hirshleifer, Lim, and Teoh, 2009, 2011), and that providing relevant information to consumers may increase attention and improve financial outcomes (Stango and Zinman, 2014; Bracha and Meier, 2015). Our work is also closely associated with the theoretical work on inattention and salience (DellaVigna, 2009; Gabaix and Laibson, 2000, 2001; Gabaix et al., 2006; Bordalo, Gennaioli, and Shleifer, 2013a, 2013b).

References

Acquisti, A. 2004. Privacy and Security of Personal Information. In Economics of Information Security, edited by L. Jean Camp and Stephen Lewis, 179–186. Boston: Kluwer Academic Publishers.

Agrawal, A., J. Jaffe, and J. Karpoff. 1999. Management Turnover and Governance Changes Following the Revelation of Fraud. Journal of Law and Economics 42, 309–342.

Anderson, R., and T. Moore. 2007. Information Security Economics — and Beyond. In Advances in Cryptology — CRYPTO 2007, edited by Alfred Menezes, 68–91. New York: Springer Berlin Heidelberg.

Baker, A., D Larcker, and C. Wang. 2022. How Much Should We Trust Staggered Difference-in-Differences Estimates? Journal of Financial Economics 144(2), 370-395.

Barber, B., and T. Odean. 2008. All That Glitters: The Effect of Attention on the Buying Behavior of Individual and Institutional Investors. Review of Financial Studies 21, 785–818.

Blascak, N. and Y. Toh. 2022. Prior Fraud Exposure and Precautionary Credit Market Behavior. Federal Reserve Bank of Philadelphia Working Paper 22-36.

Bordalo, P., N. Gennaioli, and A. Shleifer. 2013a. Salience and Asset Prices. American Economic Review: Papers and Proceedings 103, 623–628.

Bordalo, P., N. Gennaioli, and A. Shleifer. 2013b. Salience and Consumer Choice. Journal of Political Economy 121, 803–843.

Bos, M., E. Breza, and A. Liberman. 2018. The Labor Market Effects of Credit Market Information. Review of Financial Studies 31(6), 2005–2037.

Bos, M., and L. Nakamura. 2014. Should Defaults Be Forgotten? Evidence from Variation in Removal of Negative Consumer Credit Information. Federal Reserve Bank of Philadelphia Working Paper 14-21.

Bracha, A., and S. Meier. 2015. Nudging Credit Scores in the Field: The Effect of Text Reminders on Creditworthiness in the United States. Federal Reserve Bank of Boston Working Paper 15-2.

Callaway, B., and P. Sant'Ana. 2021. Difference-in-Differences with Multiple Time Periods. Journal of Econometrics 225(2), 200-230.

Cheney, J. 2003. Identity Theft: A Pernicious and Costly Fraud. Federal Reserve Bank of Philadelphia Payment Cards Center Discussion Paper 03-18.

Cheney, J., R. Hunt, K. Jacob, R. Porter, and B. Summers. 2012. The Efficiency and Integrity

of Payment Card Systems: Industry Views on the Risks Posed by Data Breaches. Federal Reserve Bank of Philadelphia Payment Cards Center Discussion Paper 12-04.

de Chaisemartin, C., and X. D'Haultfoeuille. 2020. Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. American Economic Review 110(9), 2964-2996.

DellaVigna, S. 2009. Psychology and Economics: Evidence from the Field. Journal of Economic Literature 47, 315–372.

DellaVigna, S., and J. Pollet. 2009. Investor Inattention and Friday Earnings Announcements. Journal of Finance 64, 709–749.

Dimmock, S., and W. Gerken. 2012. Predicting Fraud by Investment Managers. Journal of Financial Economics 105, 153–173.

Dimmock, S., W. Gerken, and N. Graham. 2018. Is Fraud Contagious? Coworker Influence on Misconduct by Financial Advisors. Journal of Finance 73(3), 1417–1450.

Dobbie, W., P. Goldsmith-Pinkham, N. Mahoney, and J. Song. 2020. Bad Credit, No Problem? Credit and Labor Market Consequences of Bad Credit Reports. Journal of Finance 75(5), 2377–2419.

Dyck, A., A. Morse, and L. Zingales. 2010. Who Blows the Whistle on Corporate Fraud? Journal of Finance 65, 2213–2253.

Dyck, A., A. Morse, and L. Zingales. 2024. How Pervasive Is Corporate Fraud? Review of Accounting Studies 29, 736–769.

Egan, M., G. Matvos, and A. Seru. 2019. The Market for Financial Adviser Misconduct. Journal of Political Economy 127(1), 233–295.

Federal Trade Commission. 2003. Identity Theft Survey Report. Washington, D.C.: Federal Trade Commission.

Federal Trade Commission. 2012. Report to Congress Under Section 319 of the Fair and Accurate Credit Transactions Act of 2003. Washington, D.C.: Federal Trade Commission.

Gabaix, X., and D. Laibson. 2000. A Boundedly Rational Decision Algorithm. American Economic Review 90, 433–438.

Gabaix, X., and D. Laibson. 2001. The 6D Bias and the Equity Premium Puzzle. NBER Macroeconomics Annual 16, 257–312.

Gabaix, X., D. Laibson, G. Moloche, and S. Weinberg. 2006. Costly Information Acquisition: Experimental Analysis of Boundedly Rational Model. American Economic Review 96, 1043–1068.

Giannetti, M., and T. Yue Yang. 2016. Corporate Scandals and Household Stock Market

Participation. Journal of Finance 71, 2591–2636.

Gibbs, C., Guttman-Kenney, B., Lee, D., Nelson, S., van der Klaauw, W., and Wang, J. 2025. Consumer Credit Reporting Data. Journal of Economic Literature 63 (2): 598–636.

Goodman-Bacon, A. 2021. Difference-in-differences with variation in treatment timing. Journal of Econometrics 225(2), 254-277.

Gross, T., M. Notowidigdo, and J. Wang. 2020. The Marginal Propensity to Consume over the Business Cycle. American Economic Journal: Macroeconomics 12(2), 351–384.

Gurun, U., N. Stoffman, and S. Yonker. 2018. Trust Busting: The Effect of Fraud on Investor Behavior. Review of Financial Studies 31, 1341–1376.

Guttman-Kenney, B. 2025. Disaster Flags: Credit Reporting Relief from Natural Disasters. Working Paper.

Hamdi, N., A. Kalda, A, and D. Sovich. 2024. The Costs of Financial Fraud Victimization. Mimeo, available at http://dx.doi.org/10.2139/ssrn.4787119.

Harrell, E. 2019. Victims of Identity Theft, 2016. U.S. Department of Justice, Bureau of Justice Statistics.

Harrell, E. and A. Thompson. 2023. Victims of Identity Theft, 2021. U.S. Department of Justice, Bureau of Justice Statistics.

Herkenhoff, K., G. Phillips, and E. Cohen-Cole. 2021. The Impact of Consumer Credit Access on Self-Employment and Entrepreneurship. Journal of Financial Economics, 141(1), 345-371.

Hirshleifer, D., S. Seongyeon Lim, and S. Hong Teoh. 2009. Driven to Distraction: Extraneous Events and Underreaction to Earnings News. Journal of Finance 64, 2289–2325.

Hirshleifer, D., S. Seongyeon Lim, and S. Hong Teoh. 2011. Limited Investor Attention and Stock Market Misreactions to Accounting Information. Review of Asset Pricing Studies 1, 35–73.

Kahn, C., and J. Liñares-Zegarra. 2016. Identity Theft and Consumer Payment Choice: Does Security Really Matter? Journal of Financial Services Research 50, 121–159.

Khanna, V., E. Kim, and Y. Lu. 2015. CEO Connectedness and Corporate Fraud. Journal of Finance 70, 1203–1252.

Knüpfer, S., V. Rantala, E. Vihriälä, and P. Vokata. 2024. Household Responses to Phantom Riches. Working Paper.

Kosse, A. 2013. Do Newspaper Articles on Card Fraud Affect Debit Card Usage? Journal of Banking and Finance 37, 5382–5391.

Lee, D., and W. van der Klaauw. 2010. An Introduction to the FRBNY Consumer Credit

Panel. Federal Reserve Bank of New York Staff Report 479.

Miller, D. 2023. An Introductory Guide to Event Study Models. Journal of Economic Perspectives 37(2), 203-230.

Musto, D. 2004. What Happens When Information Leaves a Market? Evidence from Post-bankruptcy Consumers. Journal of Business 77, 725–748.

Piskorski, T., A. Seru, and J. Witkin. 2015. Asset Quality Misrepresentation by Financial Intermediaries: Evidence from the RMBS Market. Journal of Finance 70, 2635–2678.

Povel, P., R. Singh, and A. Winton. 2007. Booms, Busts, and Fraud. Review of Financial Studies 20, 1219–1254.

Qureshi, H., and J. Sokobin. 2015. Do Investors Have Valuable Information About Brokers? FINRA Office of the Chief Economist Working Paper.

Stango, V., and J. Zinman. 2014. Limited and Varying Consumer Attention: Evidence from Shocks to the Salience of Bank Overdraft Fees. Review of Financial Studies 27, 990–1030.

Stavins, J. 2013. Security of Retail Payments: The New Strategic Objective. Federal Reserve Bank of Boston Public Policy Discussion Paper 13-9.

Sullivan, R. 2010. The Changing Nature of U.S. Card Payment Fraud: Industry and Public Policy Options. Federal Reserve Bank of Kansas City Economic Review (Second Quarter 2010).

Sun, L., and S. Abraham. 2021. Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects. Journal of Econometrics 225(2), 175-199.

Wardrip, K., and R. Hunt. 2013. Residential Migration, Entry, and Exit as Seen Through the Lens of Credit Bureau Data. Federal Reserve Bank of Philadelphia Payment Cards Center Discussion Paper 13-04.

Table 1: Summary Statistics – Extended Alert Filer Sample

Variable	Obs.	Mean	Std. Dev.
Number of inquiries, past 3 months	1,273,519	1.1	1.7
Number of new bankcards	1,468,543	0.1	0.4
Number of address reversals in the last quarter	1,468,722	0.009	0.094
Risk Score	1,446,719	648.4	118.7
Percent of bankcard balances in good standing (%)	855,013	94.3	20.1
Incidence of major derogatory events (%)	1,468,722	19.2	39.4
Incidence of third-party collections (%)	1,454,164	24.8	43.2
Share of prime consumers (Risk Score $>$ 660) (%)	1,446,719	47.2	49.9
Share w/ new bankcard(s) (%)	1,468,722	2.1	14.4
Number of bankcards	1,468,722	2.1	2.3
Total bankcard balance ($)	1,067,673	6,266	15,477
Total bankcard credit limit ($)	1,067,673	21,672	32,097
Share w/ new auto loan(s) (%)	1,468,722	2.3	14.9
Total balance on auto loans ($)	1,468,722	7,061	13,869
Share w/ new mortgage(s) (%)	1,468,722	1.5	12.2
Total balance on mortgage loans ($)	476,649	239,657	278,711

Notes: Risk Score is the Equifax Risk Score. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax Data, augmented with variables obtained from Equifax. Sample is from Q1:2006 to Q3:2013.

Table 2: Characteristics of Identity Theft Victims and Extended Alert Filers (NCVS ITS Data)

	(1) ID Theft Victim	(2) Severe ID Theft Victim	(3) Filed Extended Fraud Alert
Share high school diploma or less (%)	23	33	19
Share bachelor's degree or higher (%)	46	34	46
Average age (years)	47	46	46
Share non-Hispanic white (%)	75	61	66
Share married (%)	59	48	56
Share female (%)	52	54	53
Share own home (%)	71	61	65
Share income <$50,000 (%)	27	42	32
N	19,427	1,746	315

Notes: Authors' calculations using pooled data from the 2012, 2014, and 2016 National Crime Victimization Survey's (NCVS) Identity Theft Supplements. All calculations made using NCVS ITS sample weights. Severe identity theft is defined as an individual experiencing new account fraud and/or use of personal information for other fraudulent uses.

Table 3: The Effect of Alert Filing on Credit Outcomes (DID)

	(1) Risk Score	(2) Inquiries	(3) Address Reversals	(4) New Bank Cards
1(-4 ≤ Time ≤ -1)	-0.266	0.226***	0.003***	0.023***
	(0.269)	(0.007)	(0.0003)	(0.001)
1(Time ≥ 0)	12.428***	0.151***	0.007***	-0.012***
	(0.487)	(0.01)	(0.0004)	(0.002)
R^2	0.826	0.275	0.072	0.086
Observations	1,404,172	1,237,290	1,425,174	1,424,995
	(5) Percent of Card Balances in Good Standing	(6) Derogatory Events (%)	(7) Collections (%)	(8) Prime Indicator (%)
1(-4 ≤ Time ≤ -1)	-0.0002	-0.005***	-0.005***	-0.012***
	(0.001)	(0.001)	(0.002)	(0.001)
1(Time ≥ 0)	0.024***	-0.07***	-0.057***	0.043***
	(0.002)	(0.003)	(0.002)	(0.002)
R^2	0.414	0.460	0.454	0.757
Observations	828,185	1,425,174	1,411,404	1,404,172

Notes: This table shows the effect of alert filing on credit outcomes indicated in the first row modeled using Equation 1. This model allows for a quadratic trend in event time (Time). This feature is designed to remove mean reversion in credit attributes. This model also allows for changes in trend at the time of fraud (-4 to -1) and after fraud (Time≥0). This model includes individual fixed effects and calendar time fixed effects. The results indicate that there are negative changes in credit variables before alert filing and positive changes after alert placement. Standard errors are clustered at the individual level and reported in parentheses. Risk Score is the Equifax Risk Score.

Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax Data, augmented with variables obtained by the Consumer Finance Institute from Equifax.

Table 4: The Effect of Alert Filing on Credit Outcomes (DID)

	(9) New Bank Card Holder	(10) Number of Cards	(11) Cards Balance	(12) Cards Limit
1(-4 ≤ Time ≤ -1)	0.004***	0.052***	70.547	-19.672
	(0.0005)	(0.006)	(64.162)	(91.607)
1(Time ≥ 0)	0.001	-0.131***	-218.56**	-1,497.175***
	(0.001)	(0.011)	(104.758)	(150.261)
R^2	0.056	0.776	0.620	0.800
Observations	1,425,174	1,425,174	1,037,174	1,037,174
	(13) New Auto Loan Holder	(14) Total Auto Balance	(15) New Mortgage Holder	(16) Total Mortgage Balance
1(-4 ≤ Time ≤ -1)	0.002***	15.241	0.001**	991.768
	(0.0005)	(54.032)	(0.0004)	(1130.776)
1(Time ≥ 0)	0.004***	141.098	0.004***	4,429.127**
	(0.001)	(95.043)	(0.001)	(2004.97)
R^2	0.034	0.493	0.049	0.859
Observations	1,425,174	1,425,174	1,425,174	462,819

Notes: This table shows the effect of alert filing on credit outcomes indicated in the first row modeled using Equation 1. This model allows for a quadratic trend in event time (Time). This feature is designed to remove mean reversion in credit attributes. This model also allows for changes in trend at the time of fraud (-4 to -1) and after fraud (Time≥0). This model includes individual fixed effects and calendar time fixed effects. The results indicate that there are negative changes in credit variables before alert filing and positive changes after alert placement. Standard errors are clustered at the individual level and reported in parentheses. Risk Score is the Equifax Risk Score.
Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax Data, augmented with variables obtained by the Consumer Finance Institute from Equifax.

Figure 1: Number of Extended Alert Filers over Time

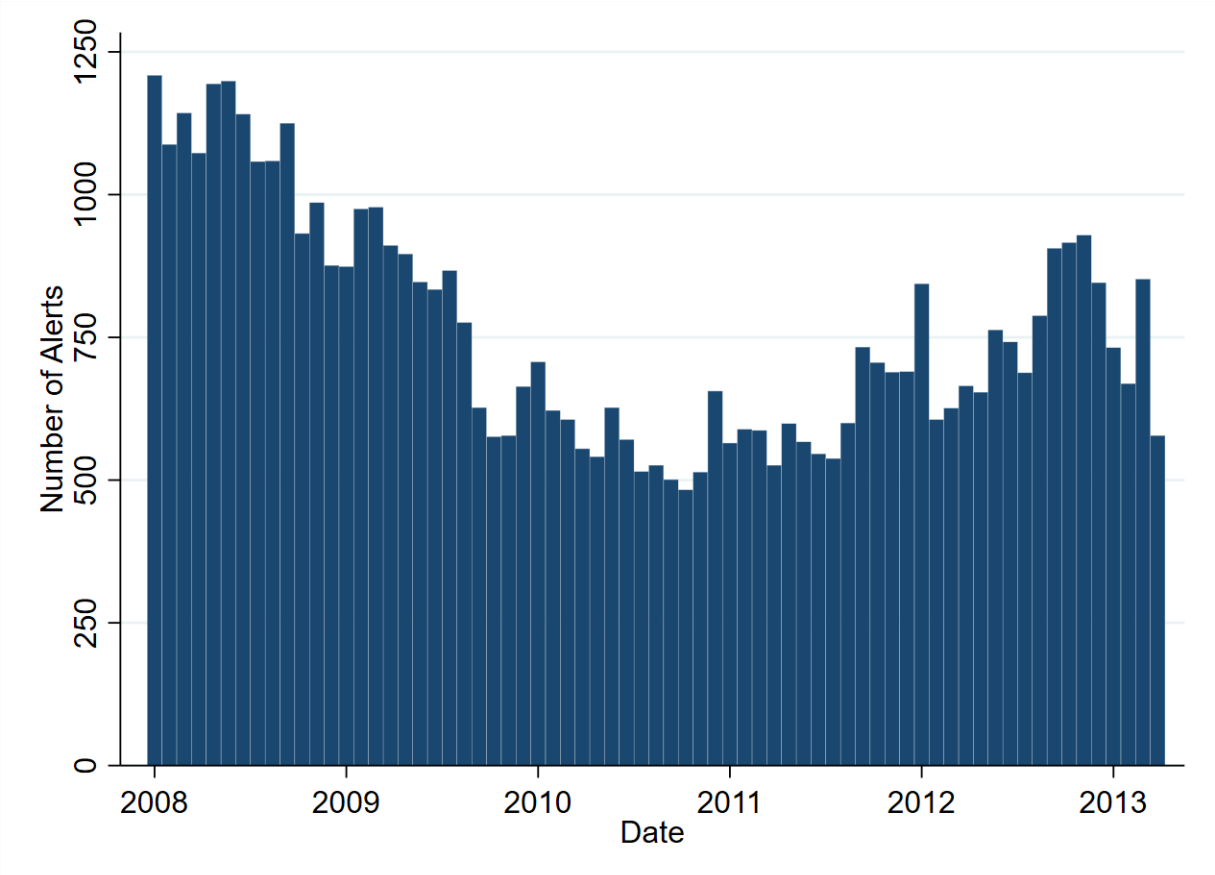

Notes: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax Data, augmented with variables obtained by the Consumer Finance Institute from Equifax. The data include only extended fraud alert filers in Q1:2008-Q3:2013.

Figure 2: Indicators of Potential Identity Theft

Panel A: Inquiries

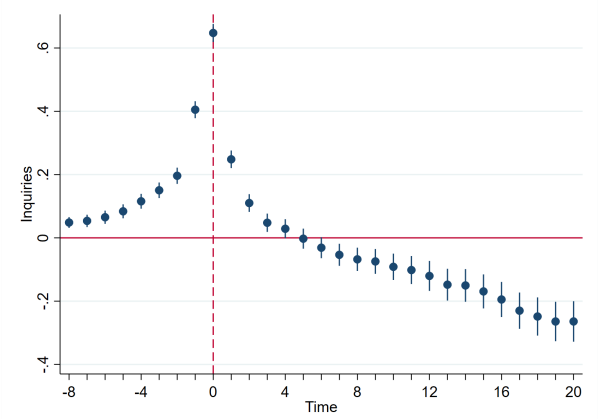

Panel B: New Bank Card Accounts

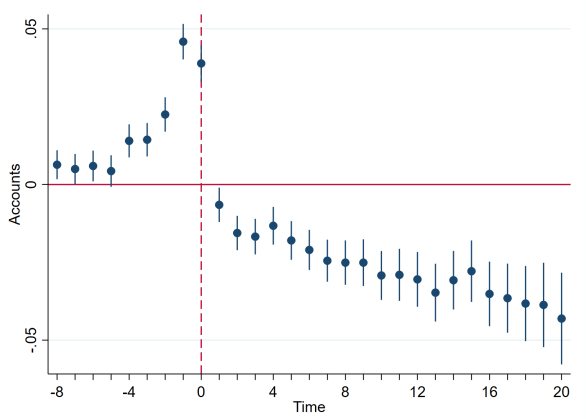

Panel C: Address Reversals

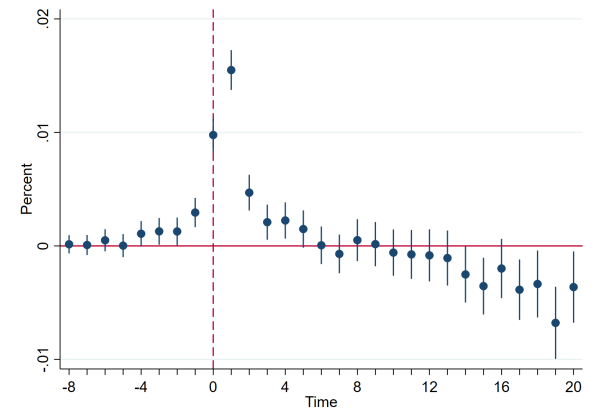

Panel D: Risk Score

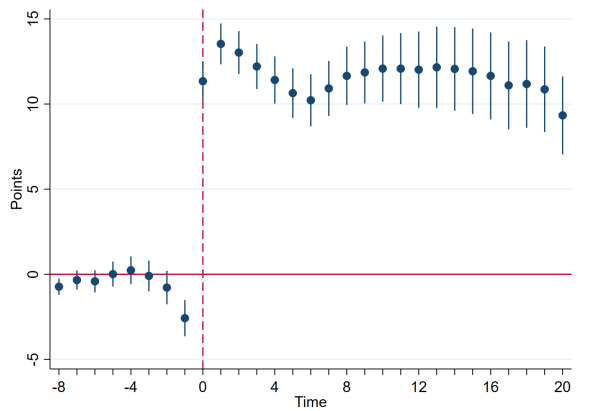

Notes: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax Data, augmented with variables obtained by the Consumer Finance Institute from Equifax. This figure depicts changes in the credit bureau characteristics of extended fraud alert filers before and after alert filing. These changes are estimated using Equation (1). Time 0 denotes the quarter of fraud alert filing, with negative time being quarters before this event and positive time being quarters after the event. All coefficients are estimated relative to the base category, which is quarters -22 to -9. The dots represent point estimates, and bands show 95 percent confidence intervals. The data include only extended fraud alert filers in Q1:2008-Q3:2013. Risk Score is the Equifax Risk Score.

Figure 3: Credit Performance Before and After Alert Filing

Panel A: Share of Card Balances Current

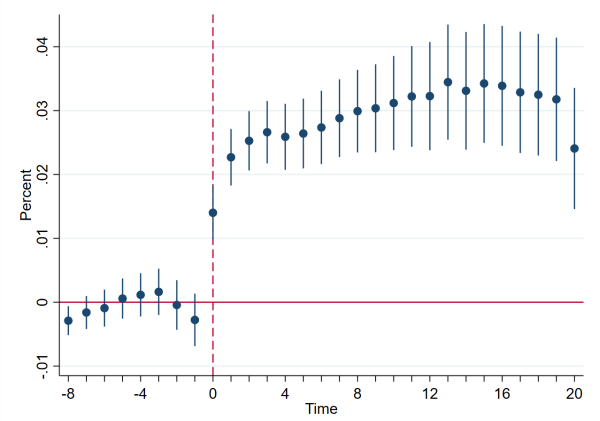

Panel B: Major Derogatory Events

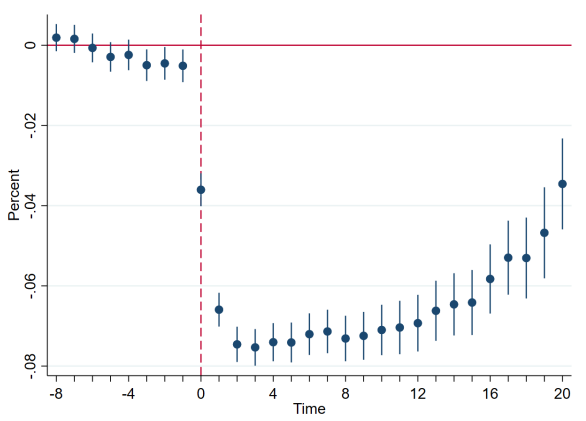

Panel C: Third-Party Collections

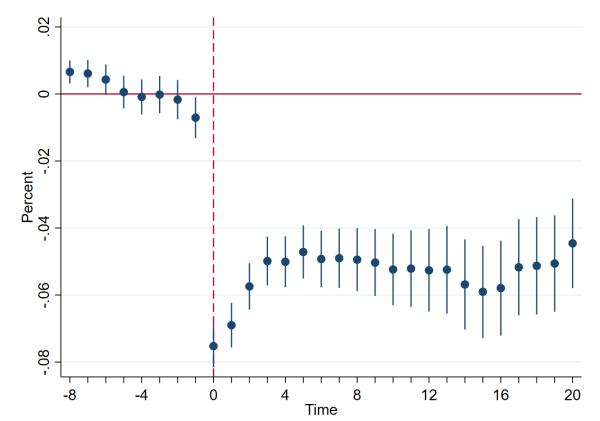

Panel D: Share of Prime Consumers

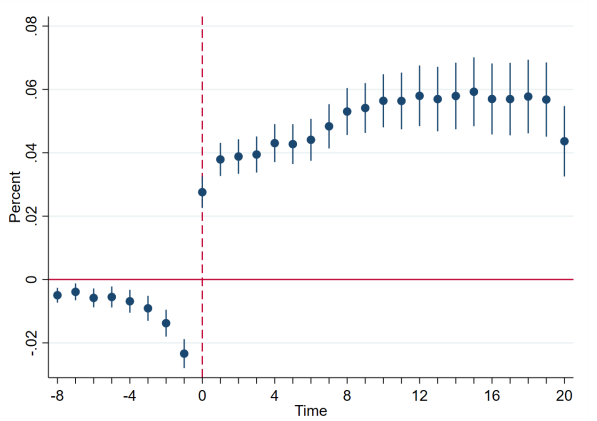

Notes: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax Data, augmented with variables obtained by the Consumer Finance Institute from Equifax. This figure depicts changes in the credit bureau characteristics of extended fraud alert filers before and after alert filing. These changes are estimated using Equation (1). Time 0 denotes the quarter of fraud alert filing, with negative time being quarters before this event and positive time being quarters after the event. All coefficients are estimated relative to the base category, which is quarters -22 to -9. The dots represent point estimates, and bands show 95 percent confidence intervals. The data include only extended fraud alert filers in Q1:2008-Q3:2013.

Figure 4: Credit Card Use Before and After Alert Filing

Panel A: New Bank Card Holder

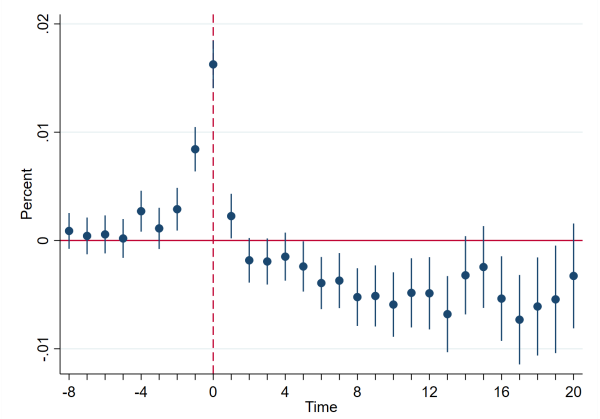

Panel B: Number of Bank Cards

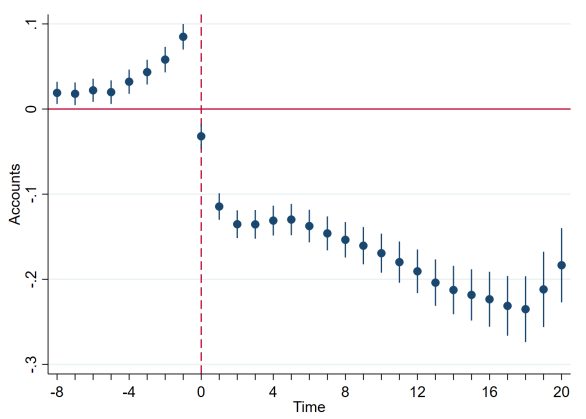

Panel C: Bank Card Balance

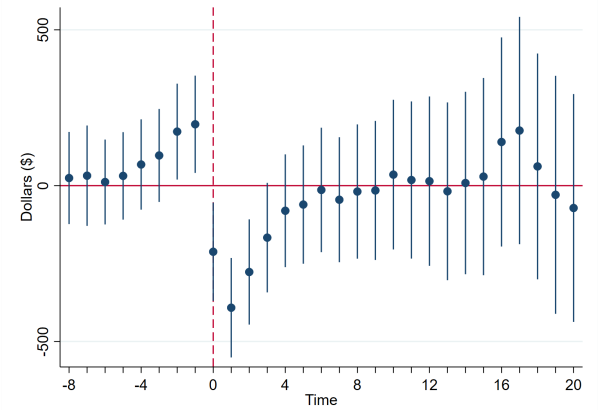

Panel D: Bank Card Credit Limit

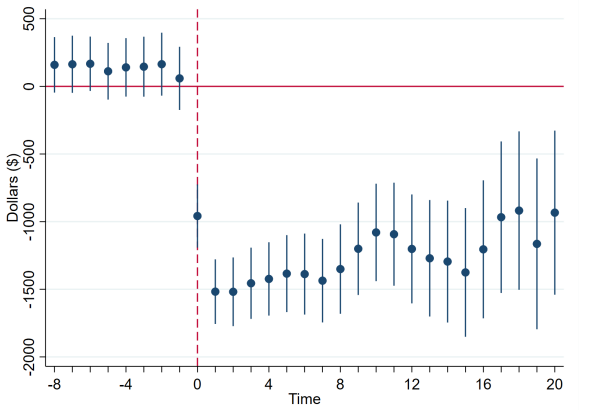

Notes: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax Data, augmented with variables obtained by the Consumer Finance Institute from Equifax. This figure depicts changes in the credit bureau characteristics of extended fraud alert filers before and after alert filing. These changes are estimated using Equation (1). Time 0 denotes the quarter of fraud alert filing, with negative time being quarters before this event and positive time being quarters after the event. All coefficients are estimated relative to the base category, which is quarters -22 to -9. The dots represent point estimates, and bands show 95 percent confidence intervals. The data include only extended fraud alert filers in Q1:2008-Q3:2013.

Figure 5: Auto Loan and Mortgage Use After Identity Theft Alert Filing

Panel A: New Auto Loan Holder

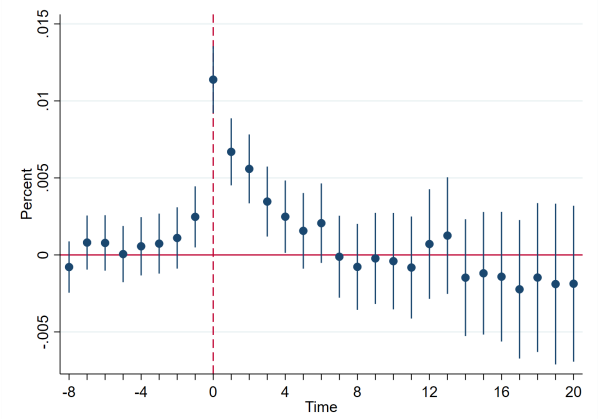

Panel B: Auto Loan Balances

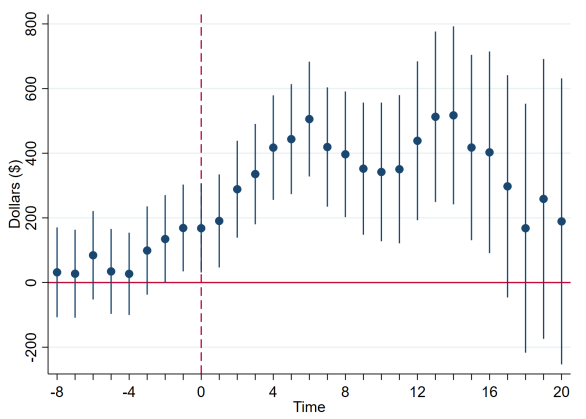

Panel C: New Mortgage Holder

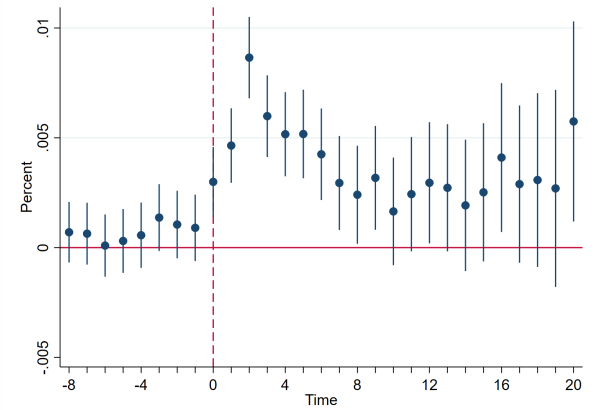

Panel D: Mortgage Balances

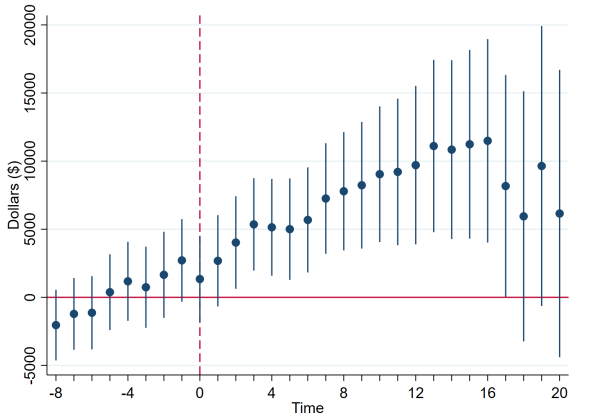

Notes: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax Data, augmented with variables obtained by the Consumer Finance Institute from Equifax. This figure depicts changes in the credit bureau characteristics of extended fraud alert filers before and after alert filing. These changes are estimated using Equation (1). Time 0 denotes the quarter of fraud alert filing, with negative time being quarters before this event and positive time being quarters after the event. All coefficients are estimated relative to the base category, which is quarters -22 to -9. The dots represent point estimates, and bands show 95 percent confidence intervals. The data include only extended fraud alert filers in Q1:2008-Q3:2013.

ONLINE APPENDIX

Figure A1: Changes in Credit Outcomes by Month of Alert Filing

Panel A: Credit Inquiries

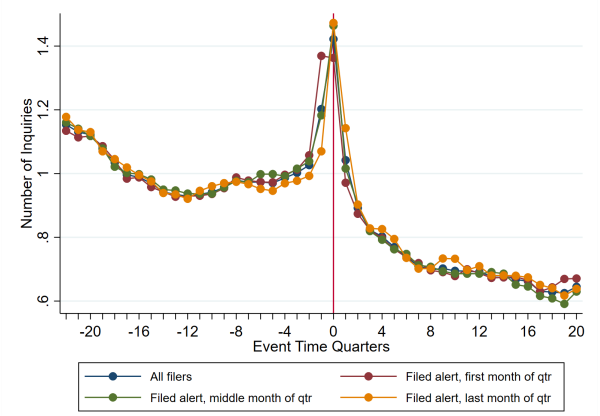

Panel B: Risk Score

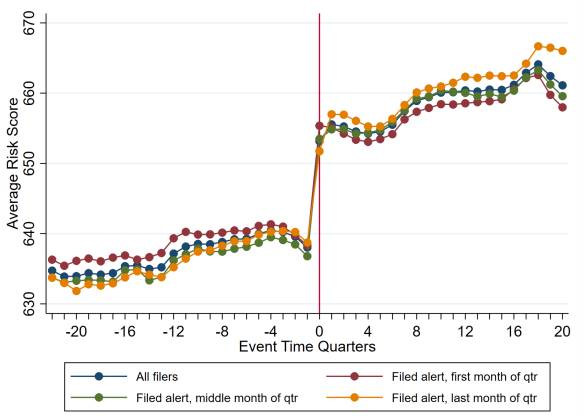

Panel C: Address Reversals

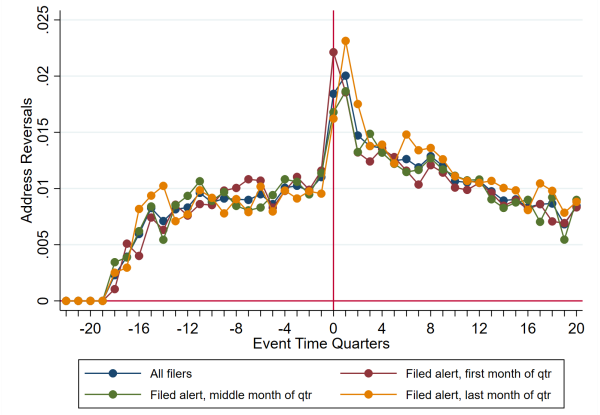

Panel D: New Bankcard Accounts

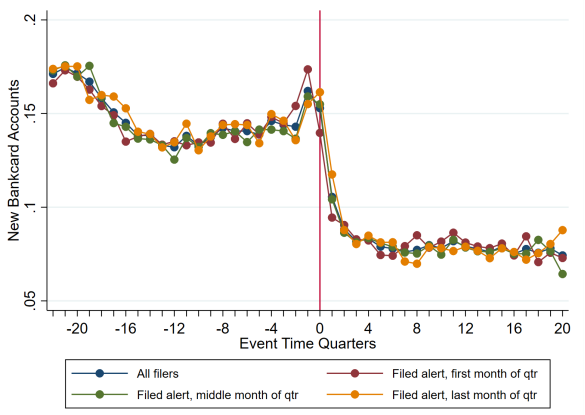

Notes: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax Data, augmented with variables obtained by the Consumer Finance Institute from Equifax. This figure depicts the average values of the credit bureau characteristics of extended alert filers by the month of the extended alert filing. Time 0 denotes the quarter of extended fraud alert filing, with negative time being quarters before this event and positive time being quarters after the event. Risk Score is the Equifax Risk Score. The data include only extended fraud alert filers in Q1:2008–Q3:2013.

Figure A2: Speed of Discovery and Clearing Up Identity Theft

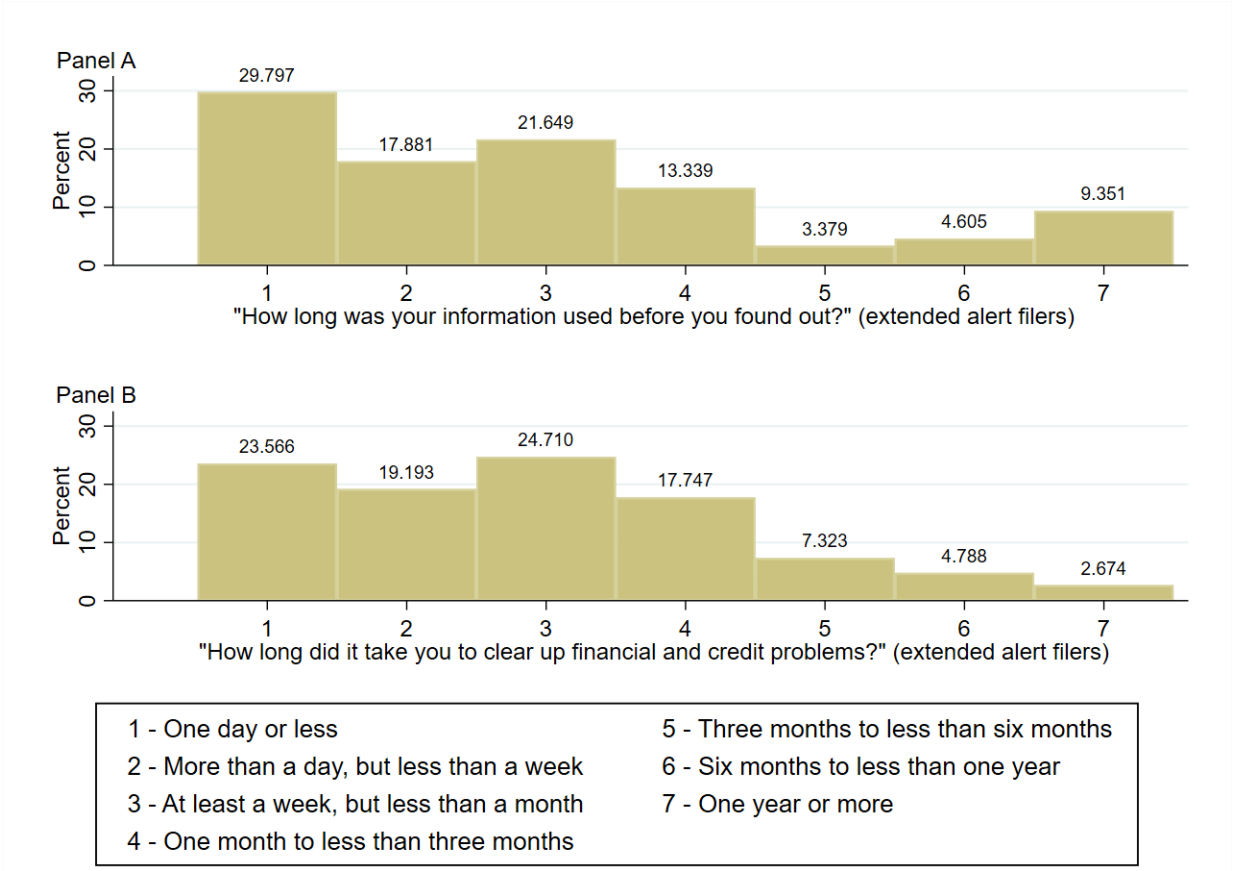

Notes: Survey responses from the Bureau of Justice Statistics' National Crime Victimization Survey 2012, 2014, and 2016 Identity Theft Supplements. All calculations made using NCVS ITS sample weights. Panel A: Percent of extended alert filers who discovered fraud less than three months after its occurrence is 82.7 percent. Panel B: Percent of extended alert filers reporting that clearing up financial and credit problems took less than three months is 85.2 percent.

Figure A3: Credit Activity Before and After Alert Filing —Subprime-to-Prime Transition Group

Panel A: Risk Score

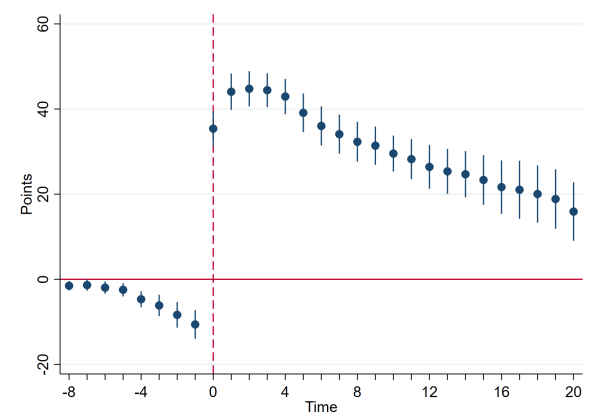

Panel B: New Card Holder

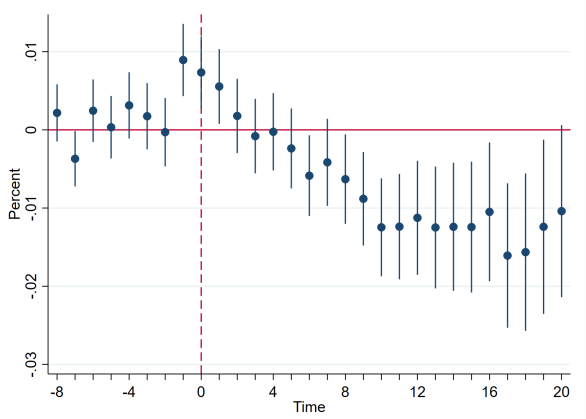

Panel C: New Auto Loan Holder

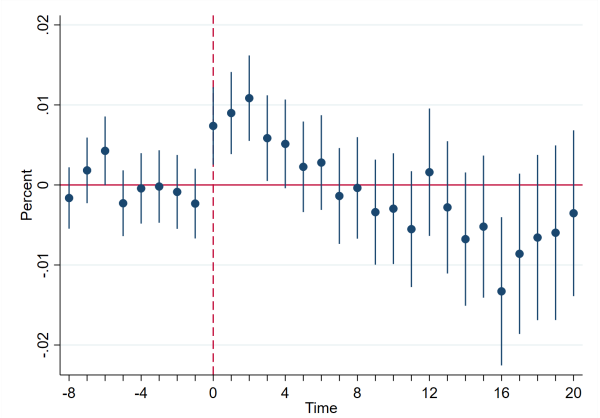

Panel D: New Mortgage Loan Holder

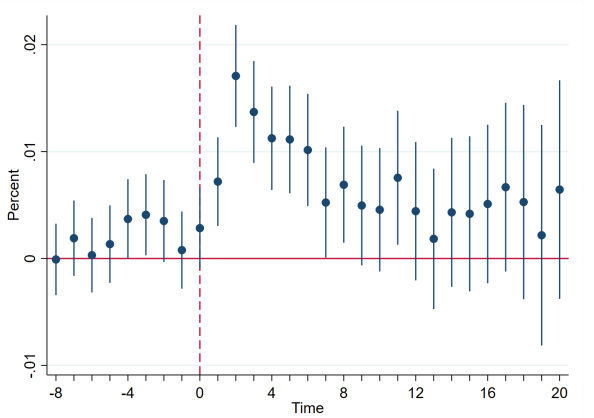

Notes: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax Data, augmented with variables obtained by the Consumer Finance Institute from Equifax. This figure depicts changes in the credit bureau characteristics of extended fraud alert filers before and after alert filing. These changes are estimated using Equation (1). Time 0 denotes the quarter of fraud alert filing, with negative time being quarters before this event and positive time being quarters after the event. All coefficients are estimated relative to the base category, which is quarters -22 to -9. The dots represent point estimates, and bands show 95 percent confidence intervals. The data include only extended fraud alert filers in Q1:2008-Q3:2013. Risk Score is the Equifax Risk Score.

Figure A4: Indicators of Potential Identity Theft —Filers Without Credit Inquiries

Panel A: Inquiries

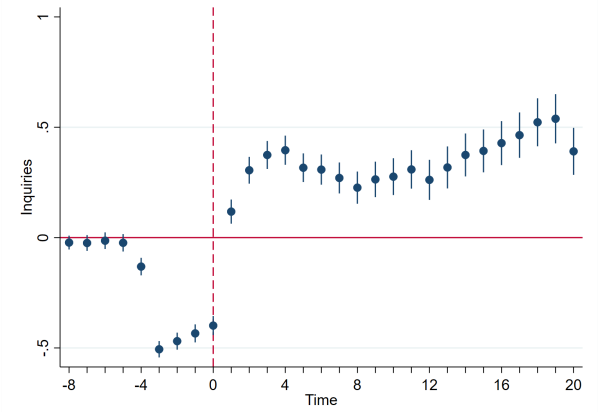

Panel B: New Bank Cards

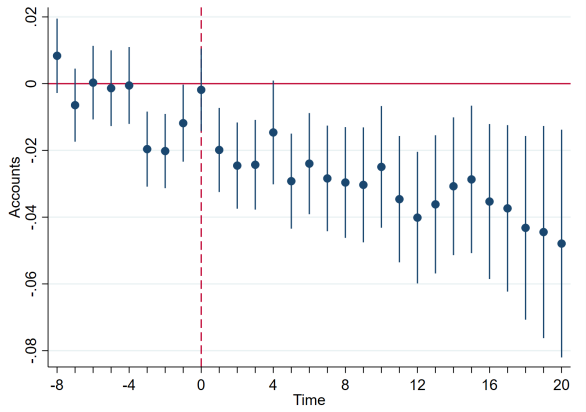

Panel C: Address Reversals

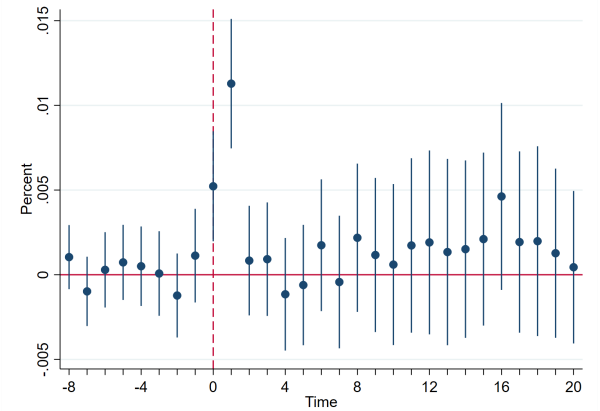

Panel D: Risk Score

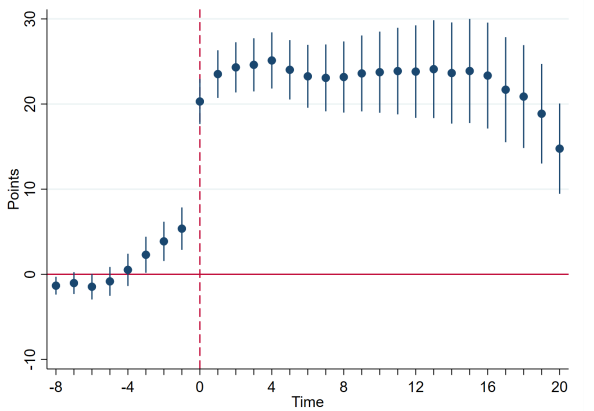

Notes: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax Data, augmented with variables obtained by the Consumer Finance Institute from Equifax. This figure depicts changes in the credit bureau characteristics of extended fraud alert filers before and after alert filing. These changes are estimated using Equation (1). Time 0 denotes the quarter of fraud alert filing, with negative time being quarters before this event and positive time being quarters after the event. All coefficients are estimated relative to the base category, which is quarters -22 to -9. The dots represent point estimates, and bands show 95 percent confidence intervals. The data include only extended fraud alert filers in Q1:2008-Q3:2013. Risk Score is the Equifax Risk Score.

Figure A5: Credit Performance Before and After Alert Filing —Filers Without Credit Inquiries

Panel A: Share of Card Balances Current

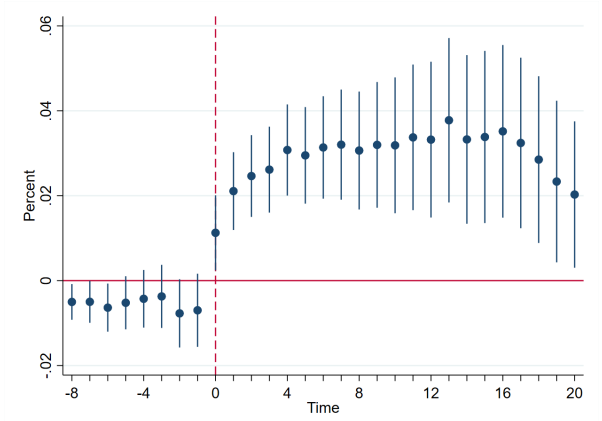

Panel B: Major Derogatory Events

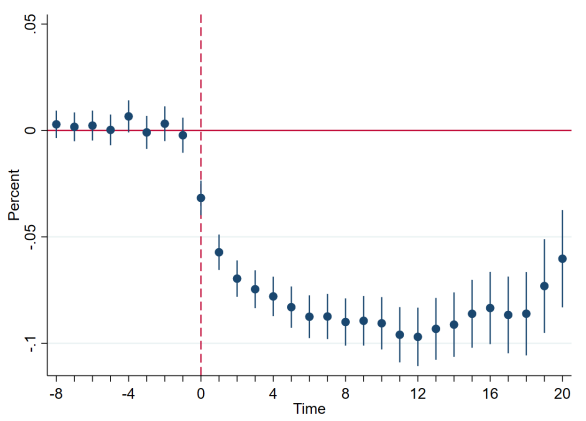

Panel C: Third-Party Collections

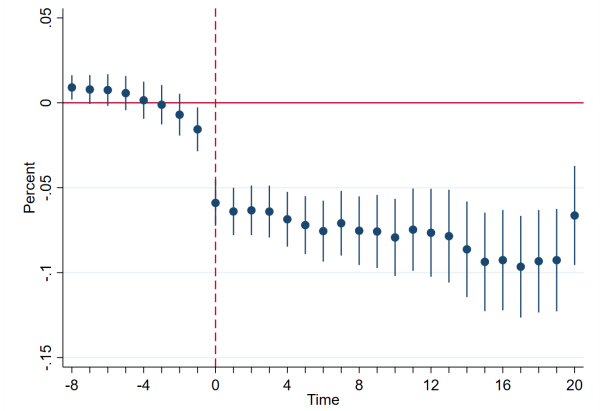

Panel D: Share of Prime Consumers

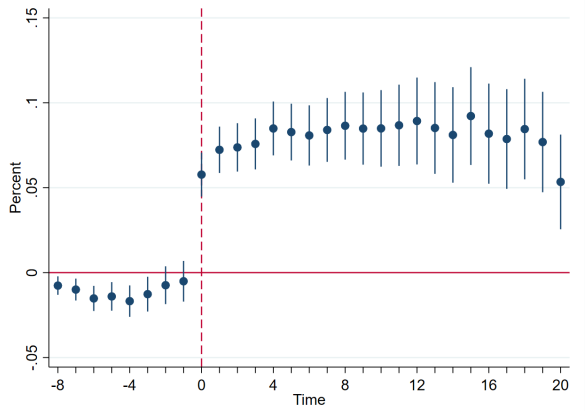

Notes: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax Data, augmented with variables obtained by the Consumer Finance Institute from Equifax. This figure depicts changes in the credit bureau characteristics of extended fraud alert filers before and after alert filing. These changes are estimated using Equation (1). Time 0 denotes the quarter of fraud alert filing, with negative time being quarters before this event and positive time being quarters after the event. All coefficients are estimated relative to the base category, which is quarters -22 to -9. The dots represent point estimates, and bands show 95 percent confidence intervals. The data include only extended fraud alert filers in Q1:2008-Q3:2013.