WORKING PAPERS
RESEARCH DEPARTMENT

WORKING PAPER 16-27
IDENTITY THEFT AS A TEACHABLE MOMENT

Nathan Blascak
Payment Cards Center, Federal Reserve Bank of Philadelphia

Julia Cheney
Payment Cards Center, Federal Reserve Bank of Philadelphia

Robert M. Hunt
Payment Cards Center, Federal Reserve Bank of Philadelphia

Vyacheslav Mikhed
Payment Cards Center, Federal Reserve Bank of Philadelphia

Dubravka Ritter
Payment Cards Center, Federal Reserve Bank of Philadelphia

Michael Vogan
Moody's Analytics

October 2016

RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA
Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/research-and-data/

Identity Theft as a Teachable Moment

Nathan Blascak

Julia Cheney

Robert Hunt

Vyacheslav Mikhed

Dubravka Ritter

Payment Cards Center, Federal Reserve Bank of Philadelphia

Michael Vogan

Moody's Analytics

October 2016

ABSTRACT

This paper examines how a negative shock to the security of personal finances due to severe identity theft changes consumer credit behavior. Using a unique data set of linked consumer credit data and alerts indicating identity theft, we show that the immediate effects of fraud on consumers are typically negative, small, and transitory. After those immediate effects fade, identity theft victims experience persistent, positive changes in credit characteristics, including improved risk scores (indicating lower default risk). We argue that these changes are consistent with increased salience of credit file information to the consumer at the time of severe identity theft.

Keywords: inattention, salience, identity theft, extended fraud alert, risk score, consumer protection, credit report, Fair and Accurate Credit Transactions Act (FACTA)

JEL Codes: G02, D14, D18

---

* We wish to thank Dennis Carlson, Amy Crews Cutts, Bradley Dear, April Ferguson, and Henry Korytkowski of Equifax for assistance with the data. We thank Marieke Bos, Anat Bracha, Chris Carroll, Lynn Langton, Susan Herbst-Murphy, Blake Prichard, Peter Schnall, Victor Stango, Jialan Wang, Chet Wiermanski, and Stephanie Wilshusen for their helpful suggestions. We especially thank Loretta Mester for making this research possible. We thank seminar participants at the Federal Reserve Banks of Philadelphia and Cleveland, the Boulder Summer Conference on Consumer Financial Decision Making, the Federal Deposit Insurance Corporation Consumer Research Symposium, the Cherry Blossom Financial Education Spring Institute, and the Digital Information Policy Scholars Conference, the Workshop on the Economics of Information Security, and the Public Policy Conference on the Law & Economics of Privacy and Data Security for their comments.

I. Introduction

In 2012, the U.S. Department of Justice reported that more than 16 million U.S. consumers were victims of identity theft. Those victims experienced financial losses of about $11 billion (Harrell and Langton, 2013).[1] In the same year, the Federal Reserve estimates that unauthorized transactions initiated via check, automated clearinghouse (ACH), and credit and debit cards exceeded $6 billion (Gerdes and Liu, 2014).[2] Based on the anonymized credit records we analyze in this paper, we estimate that about 2 million consumers placed an alert of some sort in their credit bureau records in 2012 in response to identity theft. In 2013, criminals stole the payment data of 40 million consumers from the computer systems of the big box retailer Target. These hackers also had access to the names, home addresses, and e-mail addresses of 70 million Target customers (Yang and Jayakumar, 2014). These are but a few examples of the significance of identity theft for U.S. consumers and the payment system.

As noted previously, there are a number of measures on the prevalence and magnitude of data breaches as well as estimates of fraud that are partly attributable to compromised financial accounts. However, we know much less about the consequences of identity theft to consumers or how consumers respond to identity theft.[3] We contribute new information about both of these effects. Specifically, we investigate the immediate and longer-term effects of identity theft by assembling a new, unique data set combining anonymized consumer credit records contained in the credit reports with information about "extended fraud alerts" linked to these records. Extended fraud alerts require individuals to file an official report such as a police report with accompanying evidence of a crime. Therefore, we argue that consumers who file these alerts are very likely to be victims of severe identity theft.[4]

[1] This report is based on the Identity Theft Supplement to the *National Crime Victimization Survey*. The victim and loss data are for persons 16 years of age and older. Consumers are often reimbursed for their losses (for example, via the "zero liability" policies offered by the major debit and credit card networks in the U.S.). Harrell and Langton (2013) also note that double counting of losses may occur when account holders each report losses from a joint account.

[2] In Gerdes and Liu (2014), an unauthorized transaction is "a transaction made or attempted by an individual who is not authorized by the account holder or cardholder to use a payment instrument (e.g., ACH, check, credit card, or debit/ ATM card) to purchase goods and services, initiate funds transfers, or withdraw cash from an ATM." Unauthorized transactions are one form of existing account fraud committed by identity thieves.

[3] We provide an overview of the available research and statistics in Section II.

[4] For more detail, see Section III herein and Cheney et al. (2014).

To support this argument, we first document that the timing of these alerts is correlated with changes in variables (such as applications for credit and reverse address changes) that are indicative of criminals impersonating a consumer. In addition, we measure changes in indicators (such as risk score) that affect access to and the use of credit after fraud incidents have been resolved.[5] We document that, for most fraud victims, there is an immediate, negative effect of an identity theft event on credit bureau attributes. This increase in estimated credit risk is usually reversed within a few months of filing the alert. Thus, for many consumers, there is a moderate, albeit transient, negative effect of identity theft and fraud on their credit bureau attributes.

We also find evidence that a severe identity theft event can generate persistent and positive changes in the credit bureau files of victims who file extended alerts. For many of these consumers, risk scores increase in the quarter after fraud and by an amount significantly *greater* than the decline in scores that usually occurs at the time of fraud. Not only is this *net* increase statistically significant in absolute terms, it is also significant relative to the trend in risk scores among similar unaffected consumers in control groups.[6] We find that the average risk scores of fraud victims increases by at least 12 points and remains elevated for at least several years after the identity theft event. This observation coincides with persistent reductions in the share and number of bankcard accounts past due, delinquencies, third-party collections, and other derogatory events.

We find variation in our results by risk score when using a pooled conditional quantile regression model to examine heterogeneity across the risk score distribution over time. For example, the left tail of the risk score distribution (deeply subprime scores) does not dip before an extended fraud alert is filed; scores do dip before the event for the middle and right tails of the score distribution. Further, the left and middle quantiles in the score distribution exhibit increases in risk scores immediately following the event that are 10–20 points higher than scores one year before the event. These increases persist for several years in the middle of the score distribution and for up to two years for the left tail of the distribution.

---

[5] When we reference our data and analysis, we are using the Equifax Risk Score specifically. In other cases, we discuss risk scores generally.

[6] As we describe in Section IV, we identify the effect of fraud using the exogenous timing of fraud occurrence. In essence, this approach compares the outcomes of not yet victimized individuals with already victimized consumers. This method likely overcomes a possible selection bias induced by unobservable and potentially confounding factors determining fraud victim's choice to file an extended fraud alert.

Lastly, we also observe that the proportion of fraud victims with prime credit scores rises 5 percentage points (11 percent) after identity theft. This change has an important economic significance, as previous studies have shown that prime borrowers are more likely to be approved for credit and can receive better terms of credit (e.g., lower annual percentage rate (APR)). For example, by changing her FICO score from the 620–639 range to the 660–679 range, a borrower can reduce her APR from 4.7 percent to 3.7 percent on a 30-year fixed rate mortgage.[7] Bracha and Meier (2015) also show that moving from the 620–679 score range to the 680–739 range can decrease credit card APRs by 3.5 percentage points (from 19.1 percent to 15.5 percent, on average).

The patterns we find in the data are consistent with the conjecture that some consumers were less attentive to their credit records before experiencing a severe identity theft event. When consumers file an extended fraud alert, they gain access to their credit reports and have the opportunity to dispute and remove information about compromised or fraudulent accounts from their files. In addition, this event presents an opportunity for the consumer to dispute other erroneous or outdated information that has accumulated in the credit file.[8]

Both effects would explain the discrete improvement in the apparent credit risk of many consumers on impact. However, this in itself is insufficient to explain the *persistence* of this improvement for those consumers. The patterns we observe in our data suggest that a typical consumer with an extended fraud alert is, for some time at least, either more careful about using credit, more careful about monitoring their credit reports, or both. The observation may be particularly true for those fraud victims in the left tail and in the middle of the credit score distribution. The improvement in risk scores of victims in the middle of the score distribution is more persistent than the improvement in the scores in the left tail of the distribution. Thus, at least for some consumers, identity theft or fraud may serve as a *teachable moment*, increasing their attention to credit information.

---

[7] This example is based on the national average mortgage interest rates provided by FICO on August 8, 2016. See, www.myfico.com/CreditEducation/Calculators/loanrates.aspx.

[8] According to Harrell and Langton (2013), 41 percent of identity theft victims who contacted a credit bureau requested corrections to their credit reports. In addition, a 2012 FTC study found that 26 percent of 1,001 randomly selected consumers detected at least one potentially material error (including potential evidence of identity theft) in at least one of their three credit reports (FTC, 2012). About 9 percent of the consumers in the sample successfully disputed the alleged material errors, and, as a result, their risk scores increased by 10 points or more. See also, Smith et al. (2013).

This paper contributes to a few streams of studies. First, we add to the literature that examines the consequences of identity theft on the consumer. However, unlike previous studies that focused on consumer confidence in payment systems (e.g., Sullivan, 2010) and form of payment choice (Cheney et al., 2012; Kahn and Liñares-Zegarra, 2016; Stavins, 2013; Kosse, 2013), this paper examines how identity theft can affect consumers' credit performance and credit variables. This study is also related to papers considering the trade-off between information security and data privacy (Acquisti, 2004; Anderson and Moore, 2007); and incentives for consumers to prevent identity theft (Federal Trade Commission (FTC), 2003; Cheney, 2003).

Second, we also contribute to a large and growing literature showing that individuals in a wide variety of contexts pay limited attention to and do not process information completely when making important decisions. For example, a series of studies have demonstrated that investors react less than optimally to information readily available to them at no cost.[9] Lacetera, Pope, and Sydnor (2012) demonstrate that individual car buyers exhibit left-digit bias and fail to read all digits of car odometers correctly when purchasing cars. Stango and Zinman (2014) argue that surveying consumers about overdraft fees may increase customers' attention to these fees and help consumers to avoid them. Bracha and Meier (2015) conclude that text reminders can help low-credit score consumers to improve their risk scores by paying bills on time and reducing debt and credit use. We add to this literature by showing that identity theft can serve as a strong reminder to check credit reports and be more careful and attentive to their credit portfolios.

Our research can offer important policy implications. First, our research suggests that the Fair and Accurate Credit Transactions Act (FACTA) succeeded by providing consumers with valuable options to protect their identity and resolve fraud. Second, we argue that identity theft may be one of many examples of teachable or attentive moments for consumers (others may include unexpected fees, rejected credit applications, or errors in credit reports). Providing consumers with relevant information and, perhaps, financial education at these moments of attention may be a good strategy to improve financial behavior of these borrowers and credit related outcomes for both consumers and lenders.

[9] See, for example, Barber and Odean (2008), DellaVigna and Pollet (2009), Hirshleifer, Lim, and Teoh (2009), and Hirshleifer, Lim, and Teoh (2011).

II. Institutional Details and Data

A. Consumer Credit Reports and the Fair and Accurate Credit Transactions Act of 2003

A consumer credit report is an organized record of an individual's interaction with the credit market. Typically, 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. The three largest credit reporting agencies with national scope are Equifax, Experian, and TransUnion.

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 identity theft. FACTA permits consumers to obtain free copies of their credit reports from each of the three major bureaus once each year. FACTA also required federal regulators to develop "red flag" indicators of identity theft to aid in detecting identity theft. It also required credit reporting agencies to block information that results from identity theft and to implement a set of indicators or credit file flags that inform creditors that a consumer was, or may have been, a victim of identity theft. The credit file flags include initial and extended fraud alerts that we use in this paper.[10, 11] When the consumer files an alert with one national credit bureau, this information is communicated to the other two.

B. Extended Fraud Alerts

This paper uses extended fraud alerts to identify victims of severe identity theft and fraud. The extended fraud alert has characteristics implying that practically all filers of these alerts have been victimized.[12] In particular, extended fraud alert filers must submit a police

[10] An initial fraud alert may be placed in a credit file for 90 days (and may be renewed for multiple and consecutive 90-day periods) by consumers who make a good faith assertion of identity theft.

[11] Credit freezes are another option available to consumers who want to protect their credit file information. A credit freeze is typically a fee-based service offered by credit bureaus that prevents third parties from accessing a consumer's credit report until the consumer lifts the freeze. Although there is much variation across states, many states permit victims of identity theft to place a credit freeze in their credit bureau file free of charge and often do not charge fees to lift a freeze temporarily or to remove it permanently.

[12] In our companion paper (Cheney et al., 2014) we provide evidence that suggests most consumers who file initial alerts, credit freezes, or credit monitoring are often acting out of precaution rather than being actual victims. We choose to focus on extended fraud alerts in this study to be conservative about potential false positives, but we recognize that this implies additional false negatives.

report or an Identity Theft Report to place the alert in their credit bureau files. These reports require accompanying evidence of identity theft or fraud. Such evidence requires time and effort. In addition, consumers face criminal penalties for falsifying information in these reports.[13] Thus, filers of these alerts are unlikely to place alerts in their credit bureau files based simply on worries or as a precaution.

Under FACTA, when a credit reporting agency places an extended fraud alert in a consumer's credit file, it remains in the file for seven years unless the consumer chooses to remove it beforehand. In addition, an extended fraud alert removes the consumer's credit file from lists of prescreened credit and insurance offers for five years.

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 above applications.[14]

As mentioned earlier, an important element of the rights established in FACTA (and some state laws) is the opportunity for the consumer to obtain — at no cost — copies of his or her credit reports when placing a fraud alert. This 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 fraud or inaccuracies in a consumer's credit report are verified, the credit bureaus are required to remove this information and to prevent it from reappearing in subsequent reports. This has important implications for measurement and interpretation in this paper.

[13] 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."

[14] In comparison, credit report users may apply reasonable policies and procedures to confirm the identities of initial fraud alert filers when the filer, at his or her discretion, chooses not to provide a phone number for verification purposes as part of the alert information.

C. Limited Consumer Attention to Credit Information

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 *score* in the previous 12 months, and 65 percent had not reviewed their credit reports.[15] Similarly, the U.S. Department of Justice found that only 38 percent of respondents reported that they had checked their credit reports in the past 12 months (Harrell and Langton, 2013).[16]

Consumer inattention may have implications for the accuracy of data contained in credit reports. A 2012 FTC study detected that 26 percent of 1,001 randomly selected consumers had at least one potentially material error (including potential evidence of identity theft) in at least one of their three credit reports (FTC, 2012). These errors may have significant consequences: 9 percent of the consumers in the FTC study sample successfully disputed the alleged material errors, increasing their risk scores by 10 points or more.[17] As mentioned in the introduction, depressed credit scores due to credit report inaccuracy or errors, including those stemming from identity theft, can have tangible economic effects for many consumers, especially those on the prime/subprime margin. Moreover, because the use of credit report information has become more widespread, credit report accuracy may also have adverse spillover effects. Credit reports are now frequently used in many areas unrelated to consumer credit: screening job applicants, underwriting insurance contracts, and for apartment rental applications.

Given the nature and potential economic effects of having one's identity stolen, it is likely that an identity theft incident may increase the salience of credit information and encourage increased monitoring of credit reports and/or scores. Evidence from the National Crime Victimization Survey (NCVS) shows that the number of victims who reported checking their credit report increased by up to 15 percentage points upon victimization, and the number of

[15] www.nfcc.org/NewsRoom/FinancialLiteracy/files2013/NFCC_2014FinancialLiteracySurvey_datasheet_and_key_findings_031314%20FINAL.pdf.

[16] See www.bjs.gov/index.cfm?ty=pbdetail&iid=4821.

[17] The FTC defined *material error* as potentially changing the credit score associated with the credit report.

victims that checked their bank or credit card statements increased by up to 26 percentage points.[18]

As shown in previous literature, limited attention can manifest itself in many forms and contexts. Grubb (2015) characterizes *consumer inattention* as consumers not keeping track of their past usage of goods and services. Stango and Zinman (2014) more broadly characterize it as "incomplete consideration of information that would inform choices." Sims (2003) describes inattention as consumers' limited information-processing capacity. DellaVigna (2009) models attention as a scarce resource that is a function of both salience and the number of other tasks at hand. Gabaix and Laibson (2000, 2001) and Gabaix et al. (2006) theoretically model inattention in multiple settings by having consumers minimize cognition costs. Although we do not propose our own definition of limited consumer attention, given the stylized facts described previously, we argue that many consumers pay limited attention to their credit information before an identity theft incident and that the incident may serve as a strong reminder for them to be more attentive to their credit reports.

D. 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 set (hereafter, CCP), combined with additional information on the timing (placement) and type of fraud alerts obtained from Equifax by the Payment Cards Center. This CCP data set consists of an anonymized 5 percent random sample of variables derived from the credit bureau records of U.S. consumers.[19] 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 Social Security number (SSN) as the method of randomly selecting the sample.[20]

The CCP is an unbalanced panel in which new individuals are included over time as they obtain or first report a SSN to a lender (e.g., after immigrating to the United States), open their first credit accounts, or gain their first public records. Similarly, consumers are dropped from the

[18] These statistics are based on our own calculations using the public use NCVS data set.

[19] We obtained data on fraud alerts for the period Q1:2008 to Q3:2013. The main CCP data set begins in 1999.

[20] Our data do not include actual SSNs. 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.

sample when they die, change their SSNs, 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 examine the credit files of individuals continuously present in the data set in all quarters of the Q1:2008 to Q3:2013 period so that we can trace the credit histories of these consumers.[21] Our sample consists of about 10.8 million consumers, of whom we observe that approximately 53,000 filed a first extended fraud alert in Q1:2008 or thereafter.[22] In much of the following analysis, we examine changes in variables in *event time* — the number of quarters before or after an extended fraud alert first appears.

In addition to risk score, there are two variables in our data that are especially noteworthy for our analysis.[23] The first variable is inquiries — applications for credit or insurance made by the consumer — and the second is information about the consumer's address.[24] We also study the behavior of such variables as the number of bankcards with positive balances, the proportion of bankcards in good standing, and the number of third-party collections.[25]

FACTA requires that credit reporting agencies block information resulting from identity theft four days after accepting a consumer's Identity Theft Report. The agencies must notify information furnishers 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

[21] Working with a balanced panel also mitigates concerns about "fragments" in the credit bureau files. See Wardrip and Hunt (2013).

[22] We call these *first* extended fraud alerts to distinguish between the quarter in which the alert is placed in the file and the subsequent quarters during which the alert is effective. In other words, we use the term to distinguish between the flow and stock of consumers with fraud alerts in our data.

[23] The risk score contained in the CCP is the Equifax Risk Score.

[24] The CCP contains information on the block or census tract corresponding to the address of the consumer. It also contains a "scrambled address" — a randomly generated set of characters derived from the consumer's address that can be used to detect a change of address.

[25] Accounts in good standing are defined as those that are paid as agreed, without any delinquency.

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 which 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 the 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, the CCP is similar to other real-time data sets used by researchers. It is important to emphasize that this property of the CCP 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 is blocked from the entire history contained in those files, even if the error was discovered long after it first appeared.

It is possible that the timing of the placement of extended fraud alerts may not coincide perfectly with changes in our 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 graph the changes in four credit variables of interest across event time by the month of extended alert filing (Appendix Figure A1). Results indicate that both timing of fraud and the effect of placement of fraud alerts does not systematically differ by the month of filing.

E. Summary Statistics

Table 1 presents the descriptive statistics for our data set. The table provides summary statistics for the data set, including consumers without any type of fraud alert, alert filers four quarters before filing, and alert filers at the time of filing. From this table, we can observe a number of differences between the entire population and the extended alert population for many of these variables. For example, the average risk score for the entire population is 695, while it is 655 for the extended alert population at event time zero. The number of inquiries in the past three months is 0.54 for the entire population compared with 1.5 for the extended alert population at time zero. What this means is that a typical filer of an extended fraud alert is not representative of the credit bureau population as a whole. For example, most consumers have prime credit scores, while the average extended alert filer has a subprime score.[26] These differences can potentially reflect both the selection of consumers into fraud alert protection as well as the effect of the fraud alert itself. We disentangle these two factors in the subsequent analysis.

III. Research Design

A. Selection into Treatment

We argue that filers of extended fraud alerts are very likely to be actual victims of identity theft or fraud. This is because these filers are required to submit an identity theft or police report with accompanying evidence of criminal activity and are subject to criminal penalties for falsifying information in these reports.

However, not every victim of identity theft or fraud files an extended fraud alert with a credit bureau. According to the data from the U.S. Department of Justice, only 8.1 percent of all identity theft victims file an extended fraud alert or a police report (Harrell and Langton, 2013). Among consumers suffering from more severe forms of identity theft — such as opening new accounts in the consumer's name — about 23 percent of victims contact the police and about 30 percent contact a credit bureau. It is therefore important to distinguish between actions that are exogenous to the consumer and actions that are endogenous. We argue that victimization itself, as well as its timing, is exogenous to the consumer but that the decision to file an extended alert is endogenous. This feature, in itself, introduces concerns about selection bias.

Selection into filing an extended fraud alert can occur because filers may be unobservably more motivated or attentive than victimized nonfilers and thus proceed and file an alert. This type of selection implies that fraud victims who decided to refrain from filing an alert may be "hidden" among individuals without any alert. If we compare extended fraud alert filers with all nonfilers or a selected group of nonfilers, we would not be able to separate the effect of fraud from the effect of unobservable motivation to file an alert.

[26] Prime consumers are defined as those with risk scores > 660, while subprime borrowers are those with risk scores ≤ 660.

B. Identification Strategy

To address this endogeneity issue of selection, we propose an identification strategy that relies on the timing of treatment (extended fraud alert filing) to identify the effect of fraud on credit bureau variables. In this strategy, we focus on the sample of fraud alert filers only. Since everyone in this sample files an alert at some point in time, we argue that all these individuals are motivated to file an alert once they have discovered evidence of fraud. However, because these individuals are victimized in different periods in our sample, those who are victimized later can serve as a control group for those victims who are defrauded earlier in the sample.[27] Thus, in this strategy, identification comes from the timing of victimization that is exogenous to the consumer. This feature addresses the selection issue discussed earlier that assumes self-selection among alert filers based on unobservable attributes.

C. Econometric Methodology

In our main analysis, we estimate the following distributed lag regression model:

$$Y_{it} = \beta_0 + \sum_{e=-8}^{22} \beta_{1e} T_e + \alpha_t + \delta_i + \varepsilon_{it,} \quad (1)$$

where Y is an outcome variable of interest, T is a set of event time dummy variables relative to the time of extended fraud alert filing. For example, T_2 is equal to 1 two quarters after alert filing by an individual and 0 otherwise. This specification measures the changes in the outcome variables up to eight quarters before fraud alert filing (to observe preexisting trends, if any), at the time of the filing, and up to 22 quarters after alert filing, all relative to the omitted period, which is more than eight quarters before the alert filing. This specification also includes calendar time fixed effects α_t and individual fixed effects δ_i.[28] Standard errors are clustered at the individual level.

[27] It is possible that the timing of alert relative to the timing of victimization may also be endogenous. For example, after a consumer experiences identity theft, she may be unobservably more/less motivated to report the theft immediately. Evidence from the NCVS directly refutes this hypothesis. Approximately 47 percent of all respondents reported that they discovered identity theft within 24 hours of misuse, and 95 percent said they found out within three months.

[28] Note that we cannot include cohort of fraud alert filing fixed effects into this specification because they would be perfectly collinear with individual fixed effects. As a robustness test, we included cohort fixed effects and zip code fixed effects, but omitted individual fixed effects, and obtained almost identical results.

The data used in these regressions only include extended fraud alert filers. Hence, the only source of variation exploited is the variation in the time of victimization and fraud alert filing. This specification is standard in the literature and is used by Gallagher (2014) and Gross, Notowidigdo, and Wang (2014), and Dobkin et al. (2016).

D. Reverse Causality

It is possible that some consumers may be trying to clean up their files in preparation for mortgage or other important credit applications. While there are limits to what consumers can do to clean up their credit file, it is possible that some consumers may "discover" something bad in their reports — such as fraud — simply because they were actively shopping for credit and paid more attention to their credit reports. This situation could be described as reversed causality in the sense that a future application for credit or good credit behavior may lead to fraud discovery. To the contrary, survey evidence from the NCVS indicates that only approximately 1 percent of respondents who were victims of identity theft said that they discovered misuse upon applying for credit, bank accounts, or loans.

We also use our data to test directly 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. This hypothesis implies that such consumers are equally likely to have fraud-related activity (e.g., address change reversals, new accounts that are closed immediately, increases in delinquent accounts) in their files at any time before they file an alert. If, instead, negative fraud-related activity is present in the credit reports of alert filers shortly before (one or two quarters) or in the quarter of fraud alert filing, the hypothesis of a simple credit file repair leading to identity theft discovery will not be supported by the data. Instead, the effects we observe are likely to be related to fraud or identity theft. As described in the next section, our results do not support the credit repair hypothesis.

IV. Main Results

In this section, we present evidence justifying our identifying assumptions. We also provide evidence consistent with identity theft or fraud occurring just before the placement of an extended alert. Next, we discuss the effects of an extended fraud alert on such outcomes as risk score, inquiries, proportion of cards in good standing, the number of cards with positive balances, and the incidence of third-party collections and other derogatory events.

A. Evidence of Fraud

Panel A of Figure 1 plots the average number of inquiries for all fraud victims in event time, with event time 0 equal to the quarter of the first extended fraud alert filing. In all panels of this figure, we remove business cycle and seasonal effects by including calendar time dummies. These plots display the residual average values of variables after we remove seasonal and business cycle effects. The most important observation about Panel A of Figure 1 is the very large and transitory increase in the number of credit applications that coincides with the quarter the extended alert is filed. The average number of inquiries increases from 1.2 before fraud to 1.8 (50 percent) at the time of fraud alert. This increase is consistent with consumers' personal information being stolen by criminals and used to shop for credit.[29] It is possible that consumers become aware of identity theft because this spike in applications triggers letters or phone calls from creditors.

The rapid buildup in inquiries before fraud alerts also coincides with a transitory decline in risk score documented in Panel B of Figure 1 (time −1). Note that the average increase in score that follows is typically larger than the transitory decline (we revisit this point in the next section).

Certain types 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 the credit bureau.[30] 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 reverse address changes at the time an extended fraud alert is filed with patterns prior to and after the event.[31] Panel C of Figure 1 plots the fraction of fraud alert filers who reverse address changes over event time. This figure shows a sharp increase in reverse address changes at the time the

[29] We provide formal statistical tests of these effects in the next section and in Cheney et al. (2014).

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

[31] 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 Consumer Credit Panel is not updated and, therefore, is not affected by the new information.

fraud alert is filed and one quarter after, consistent with consumer reversing address changes made by criminals.[32]

Finally, Panel D of Figure 1 plots the average number of new revolving accounts for fraud alert filers before and after they file an alert. This figure shows that new revolving accounts begin to increase sharply a few quarters preceding the fraud alert filing and peak one quarter before alert filing. This finding is consistent with criminals using consumers' stolen personal information to open new revolving accounts.

The increases in inquiries, reverse address changes, and the number of new accounts near the time of the fraud alert filing, as well as the decline in risk score shortly before the placement of the fraud alert, allow us to conclude that fraud is very likely to have occurred either during the quarter the fraud alert was filed or at most a quarter or two before that date. The patterns in these four credit report variables discount the hypothesis that alert filers simply attempt to repair their files due to another reason (e.g., preparing to apply for credit in the future) and discover fraud in their reports during this repair. This hypothesis implies that fraud-related activity would be equally likely to be present in all quarters before fraud alert filing. However, we find in our data that fraud related activity is tightly concentrated just before the alert filing or at the time of filing.

B. Regression Analysis: Evidence of Fraud and Its Effects over Time

Figure 2 and all subsequent figures report coefficients from the distributed lag regression model specified in equation (1). The four panels of Figure 2 show coefficients for four outcome variables: credit inquiries, risk score, reverse address changes, and new revolving accounts. The coefficients show the difference in the outcome variables between already victimized consumers (treatment group) and not yet victimized individuals (control group) over time before and after identity theft and fraud. In addition to point estimates displayed as dots, all these figures provide 95 percent confidence intervals as vertical bands.

Panel A of Figure 2 plots the difference in credit inquiries between treated and not yet treated extended fraud alert filers. Similar to Figure 1, this figure shows that, on average,

[32] Note that the removal of the business cycle effects results in reverse address changes being close to zero or negative in most events. These negative values can be interpreted as deviations from the pattern induced by the business cycle. Even though some values are negative, they are not statistically significant. We formally test this in the next section.

inquiries increase by 0.6 at the time of fraud, and this spike is highly statistically significant. The inquiries decrease after fraud and stay at a lower level than before fraud.

Consistent with criminal activity, Panel B of Figure 2 shows a statistically significant decline in risk score of about three score points one quarter before the fraud event. However, this decline is reversed and, on average, risk scores increase by about 12 points at the time of the fraud alert. Panel B of Figure 2 also shows that this improvement in risk scores is persistent and remains highly statistically significant for several years.

Panel C of Figure 2 plots the effect of fraud on reverse address changes. The only statistically significant coefficients are in the quarter before and a few quarters after extended fraud alert filing. The coefficients imply that around 1 percent of victims reverse address changes at the time of fraud and an additional 1.5 percent of victims do the same one quarter after fraud alert filing.

Similar to Panel D of Figure 1, Panel D of Figure 2 shows that the number of new revolving accounts peaks in the quarter preceding extended fraud alert filing. On average, 1 in 10 fraud victims have one new revolving account opened at that time. The number of new revolving accounts declines quickly after fraud alert filing. Fraud victims have, on average, between 0.05 and 0.1 fewer new revolving accounts in the quarter after identity theft or fraud is resolved.

An important feature in all panels of Figure 2 is that there are no pre-trends in credit inquiries, risk score, reverse address changes and the number of new accounts a year or more before alert filing. This finding is consistent with the identifying assumption that not-yet treated individuals are a reasonable control group for already victimized consumers. In addition, this result also suggests that fraud-related activity happens shortly before an extended fraud alert filing (up to one year before the alert filing).

In addition, the results in Figure 2 suggest that negative effects of fraud on risk score and increases in credit inquiries, reverse address changes, and new accounts (which are likely due to fraud) are robust and statistically significant. The results for the time after fraud also indicate that the persistent improvement in risk score after the resolution of identity theft or fraud may be explained in part by reductions in the number of inquiries and the number of new accounts after fraud. The reductions in both of these credit variables can positively affect risk score. We explore positive changes in other credit attributes after fraud in the next few figures.

Figure 3 presents the effects of fraud on such credit variables as the number of open bankcards, cards with positive balances, and age of the newest card. While the number of open cards increased just before fraud alert filing (Panel A of Figure 3), it drops sharply at the time of fraud alert filing and continues to decline thereafter. This may imply that consumers close fraudulent cards opened by criminals at time −1. After the fraud is resolved, consumers have fewer cards than before the fraud took place. This result suggests that, after identity theft or fraud, fraud victims close cards they do not actively use. Alternatively, they may dispute accounts reported as open when they had been closed for some time.

Panel B of Figure 3 shows that cards with positive balances (actively used cards) also follow a similar pattern, even though the decline after fraud alert filing is less pronounced. Finally, Panel C of this figure shows that the age of the newest card evolves naturally after the drop in this variable one quarter before fraud alert filing. It appears that victims of severe identity theft open bank cards less frequently while maintaining the cards they actively use.

Figure 4 provides evidence that, after fraud is resolved, consumers keep a higher proportion of their cards in good standing (Panel A). Fraud victims also reduce the incidence of major derogatory events on cards by about 4 percentage points (Panel B) and the incidence of third-party collections by 8 percentage points (Panel C). The sharp declines in the incidence of derogatory events and third party collections at the time of fraud alert filing might result from the consumer disputing fraudulent accounts as well as other incorrect information in their credit reports. However, the persistence of these effects over time suggests that consumers may have changed their payment habits to keep more cards in good standing and out of collections.

To summarize our findings, we plot the share of the population with prime scores (higher than 660) in our sample in Figure 5. This plot shows that fraud activity at event time −1 lowers the share of prime consumers by about 1.8 percentage points (4 percent decrease relative to the baseline share of prime consumers of 45 percent). However, after fraud resolution, the share of prime consumers increases by 5 percentage points (an 11 percent increase) and remains elevated for a few years. These changes have far-reaching economic consequences as they may allow borrowers to obtain more credit and at better terms. For example, on average, the interest rates (APR) on a 30-year fixed rate mortgage decrease from 4.7 percent to 3.7 percent when a borrower moves from the 620–639 FICO score range to the 660–679 range.[33] 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. It is also important to emphasize that lower risk borrowers can benefit lenders because these borrowers are less likely to default on loans.

V. Robustness Checks

A. Initial Fraud Alert Filers as a Control Group

In addition to filing extended fraud alerts, consumers can request a number of other fraud protection services from credit bureaus. One of these services is an initial fraud alert. The major difference between an extended fraud alert and an initial fraud alert is that initial alerts do not require police reports and evidence of fraud.[34] Therefore, initial fraud alerts can be filed out of precaution or suspicion about possible identity theft or fraud. However, in order to file an extended fraud alert, the consumer must present evidence of fraud and file a police report or an Identity Theft Report. Thus, we can argue that initial alert filers may be at least as motivated as extended fraud alert filers to request an alert but do not have the evidence of fraud that would allow them to file an extended fraud alert. This difference between the two types of alerts provides us with another potential mechanism to separate the effects of fraud from unobserved factors that motivate the filing of a fraud alert. In particular, this feature of the initial fraud alert allows us to use initial alert filers as a control group for the extended alert filers.

To measure the difference between these two types of alert filers, we estimate an alternate specification of the regression model in Equation (1):

$$Y_{it} = \beta_0 + \sum_{e=-8}^{22} \beta_{1e} T_e + \sum_{e=-8}^{22} \beta_{2e} T_e \times 1_{ext} + \beta_3 1_{ext} + \alpha_t + \delta_i + \varepsilon_{it}. \qquad (2)$$

[33] This example is based on the national average mortgage interest rates provided by FICO on August 8, 2016. See www.myfico.com/CreditEducation/Calculators/loanrates.aspx.

[34] In addition, unlike extended fraud alerts, which are active for seven years, initial fraud alerts are active for 90 days only. However, consumers can renew initial fraud alerts in multiple and consecutive quarters. An initial fraud alert does not remove the consumer from lists used to make pre-screened offers of credit, but it requires lenders to have additional policies and procedures in place to verify a consumer's identity when they receive a request to open a new account or other credit inquiries.

The only difference between this specification and the specification in Equation (1) is that we add a new indicator variable, 1_{ext}, which is equal to one for extended alert filers and to zero for initial alert filers. We also interact this variable with event time indicators T_s. Thus, β_{2s} coefficients will capture the differences between initial and extended alert filers before and after they file an alert.[35]

Figure 6 presents results from comparing extended alert filers to initial alert filers. All results in this figure are qualitatively very similar to our main results in Figure 2. In particular, even after using initial alert filers to control for unobserved motivation to file a fraud alert, we can see evidence of fraud among extended alert filers just before an alert such as increased inquiries (Panel A), decreased risk score (Panel B), a higher address reversal (Panel C), and a higher number of new revolving accounts (Panel D). The behavior of fraud victims after fraud resolution is also similar to the main results with persistent increases in risk score, fewer inquiries and new accounts opened by these consumers. Based on these results, we argue that our main findings are unlikely to be driven by unobservable motivation of some fraud victims to file an alert, but these results are more likely to be due to the effect of victimization on consumers.

B. Controlling for Long-Term Event Time Trends

As can be seen in Figure 1, some credit variables such as risk score may have long-term trends in event time. These long-term trends may be explained by mean reversion in risk score and other variables. For example, risk scores of a group of subprime individuals may rise over time simply because the effects of adverse past events, which decreased their scores in the first place, receive less weight in their current score.

To separate the effect of mean reversion in credit variables from the longer-term effects of fraud, we estimate the following parametric econometric model adopted from Dobkin et al. (2016):

$$Y_{it} = \beta_0 + \beta_1 e + \beta_2 e^2 + \beta_3 1_{e\geq 0} + \beta_4 1_{e\geq 0} \times e + \beta_5 1_{e\geq 0} \times e^2 + \beta_6 1_{-4\leq e\leq -1} + \beta_7 1_{-4\leq e\leq -1} \times e + \alpha_t + \delta_i + \varepsilon_{it}. \quad (3)$$

[35] We estimate the specification in Equation (1) on the sample of extended fraud alert filers only, while the results in this section include extended and initial alert filers.

In this specification, e denotes fraud event time (which varies from −22 to 22), $1_{e\geq 0}$ is an indicator variable equal 1 for non-negative event time, and $1_{-4\leq e\leq -1}$ is an indicator variable for time periods between −4 and −1. All other variables are as defined in Equation (1).

The specification in this model is motivated by the patterns in the data observed using nonparametric specification in Equation (1). In particular, Figures 2 through 4 show evidence of fraud shortly before alert filing and a discontinuous change in credit attributes at the time of alert filing. These two patterns motivate us to allow for discontinuous (intercept) shifts at the time of fraud (from time −4 to −1) and after fraud resolution (time larger than or equal to zero). We also allow for a quadratic trend in event time. However, because this trend may shift after fraud resolution, we interact the quadratic trend with the positive time indicator. Finally, we interact the linear component of the trend with the fraud time indicator.[36]

Table 2 summarizes results for this specification. The coefficients on the event time variable in this table show that there are (mostly) linear trends in many credit variables. For instance, on average, risk score increases by 0.9 points every quarter. The square of event time, however, is statistically insignificant and very small economically for the credit variables we consider. Similar to our previous results in Figures 2 to 4, we find generally negative effects of fraud on credit attributes. In particular, fraud ($-4 \leq \text{Time} \leq -1$) decreases risk score by 5 points, increases credit inquiries by 0.4, and increases address reversals, new revolving accounts, and derogatory events.

We also find generally positive changes in credit variables after fraud resolution. On impact ($\text{Time} \geq 0$), risk scores increase by 11 points, collections decrease by 8 percentage points, and the proportion of cards in good standing goes up 2.3 percentage points. Similar to our earlier results, inquiries and reverse mobility are elevated after fraud and cards with positive balances decrease. There is some attrition in these initial effects as indicated by the interactions of time trends with the after-fraud indicator variable. For example, the coefficients on the interactions indicate that about 5 points of the initial jump in the score is gone after 10 quarters since fraud event. Overall, these results are very similar to our main results obtained without controlling for long-term event time trends.

[36] Since there are only four periods for which the fraud time indicator is equal to 1, we do not to interact it with the square of event time to avoid multicollinearity.

C. Controlling for Individual Level Mean Reversion

The econometric model in Equation (3) assumes a common mean reversion for all individuals in both the pre- and post-alert filing time periods. It is possible there is substantial heterogeneity in mean reversion across individuals. Imposing a common mean reversion across individuals may mask the true effect of fraud on individuals. Because of the granular nature of our data, we have a long-time series for each individual in our sample, which can allow for panels to have their own individual time trends.

To distinguish the effect of mean reversion from that of fraud, we specify a model in the spirit of Musto (2004):

$$Y_{it} = \beta_0 + \delta_i + \delta_i \times t + \delta_i \times t^2 + \beta_1 D_{it} + \alpha_t + \varepsilon_{it}, \quad (4)$$

where δ_i is an individual fixed effect to be estimated and $\delta_i \times t + \delta_i \times t^2$ is an individual-level quadratic time trend.[37] The variable of interest in this specification is D_{it}, an indicator variable equal to 1 when individual i has an extended fraud alert filed at time t. This variable captures the difference in a variable of interest between before and after filing an extended fraud alert. By specifying an individual quadratic time trend for each consumer, we can more precisely separate the effect of mean reversion from the effect of the extended fraud alert. However, estimating individual fixed effects and individual quadratic time trends introduces computational restrictions. For each panel in Equation (4), we perform our analysis on a 6.7 percent random subsample of our data.[38]

We present results of this analysis from Equation (4) for risk score, proportion of bankcards in good standing, and new revolving accounts in Table 3. The estimates are quantitatively similar to those previously reported in Figures 2 and 4. After controlling for individual fixed effects and mean reversion, we find that having an active extended fraud alert increases risk scores, on average, by 13.6 points, increases the proportion of cards in good standing by 1.98 percentage points, and decreases the number of new revolving accounts by 0.085. Reported R^2 are high because the estimated individual effects, along with the individual quadratic time trends, account for a significant portion of the variation in these credit variables.

[37] As mentioned in the previous section, use of individual-level quadratic time trends is motivated by observed patterns in the data. Estimates using a linear time trend produce similar results.

[38] Since we estimate individual fixed effects and individual quadratic time trends, there is a minimum of 3*n* variables that are estimated, where *n* is the number of panels in the subsample.

VI. Heterogeneous Effects

In this section, we study the effects of fraud on consumers with various initial conditions or those subject to various forms of fraud or criminal activity. We focus on such variables as credit inquiries and risk score. To compare the effects of fraud on consumers with different values of these variables, we estimate the following model:

$$Y_{it} = \beta_0 + \sum_{e=-8}^{22} \beta_{1e} T_e + \sum_{e=-8}^{22} \beta_{2e} T_e \times 1_{sub} + \beta_3 1_{sub} + \alpha_t + \delta_i + \varepsilon_{it}, \qquad (5)$$

where $1_{sub}$is an indicator variable equal one for one of the subgroups we consider later in this section. These subgroups are defined based on the characteristics of consumers in either the two quarters preceding alert filing or in the quarter of alert filing (event time −2, −1, and 0).

A. Consumers with Credit Inquiries

Consumers with credit inquiries at the time of fraud may behave differently after fraud is resolved compared with consumers without credit inquiries. 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 cannot clearly distinguish between these two types of inquiries in our data. However, we can compare fraud victims without inquiries with fraud victims with inquiries. We hypothesize that consumers without inquiries may be 1) less attached to the credit market and less attentive to their credit information and 2) subject to existing account fraud or other fraud that does not result in credit inquiries.

Figure 7 shows changes in credit variables of fraud victims without credit inquiries at time −2, −1, or 0 compared with consumers with credit inquiries in that time span. The decline in credit inquiries at the time of fraud shown in Panel A of Figure 6 is, of course, mechanical, but the other results are not. Panel B of this figure suggests that no-inquiry fraud victims experience larger effects of fraud on risk scores in the quarters after the extended fraud alert was filed. This suggests that this subgroup of consumers exhibited more inattention before the fraud than the victim population as a whole. However, we cannot rule out the possibility that this subgroup experienced even more serious fraud than the entire population of extended alert filers.

Panel C of Figure 7 suggests that there is no statistically significant difference in terms of reverse address changes between fraud victims with and without inquiries. However, Panel D of this figure shows that victims with inquiries are more likely to have new accounts opened at the time of fraud. This result is consistent with criminals being successful in applying for credit with stolen consumer information for a fraction of the population with credit inquiries or a valid consumer acquiring new credit.

B. Subprime and Prime Fraud Victims

It is possible that subprime consumers react to fraud differently than prime consumers, either because they exhibit more inattention or they experience more severe forms of fraud. We define subprime consumers as having a risk score less than or equal to 660 at event times −2, −1, or 0 and compare them with consumers with risk scores above 660 in that time period. We summarize the results of this exercise in Figure 8.

Panel A of this figure shows that, before and after fraud, subprime consumers have more credit inquiries than prime consumers, which is consistent with the prior literature suggesting that subprime customers shop for credit more than prime consumers. However, Panel A also shows that, at the time of fraud, both prime and subprime consumers have the same number of inquiries. This finding may imply that, at the time of fraud, the behavior of inquiries for both groups is driven by common factors such as criminals applying for credit using stolen consumer information. This result likely corroborates our argument that most of credit inquiries at the time of fraud are generated by criminals, not consumers.

Panel B of Figure 8 also shows that the effect of fraud on risk scores of subprime consumers is much larger than the comparable effect for prime consumers. In particular, immediately after fraud, the average subprime population's risk score jumps 18 points higher than the score of the prime population. There are three possible interpretations that are not mutually exclusive. First, this may be evidence that subprime consumers exhibit relatively more inattention prior to the fraud. Second, these consumers may have been exposed to more severe forms of fraud. Third, there may have been more errors unrelated to the fraud on the credit reports of these consumers. It is also worth noting that the effect of fraud resolution on the risk scores of the subprime population is persistent.

Panel C of Figure 8 suggests that there are few differences between prime and subprime fraud victims in reverse address changes. Finally, Panel D of this figure shows that subprime consumers have more new accounts opened as fraud occurs compared with prime victims. In addition, after identity theft is resolved, these consumers seem to continue opening a larger number of new accounts compared with prime borrowers. This last result may suggest that some subprime consumers use their higher risk scores to apply for additional credit cards.

C. Quantile Regressions

An alternative way to address heterogeneity is to study the distribution of our outcome variables of interest by specifying quantile regressions. This methodology is particularly attractive because many of our credit variables are not symmetrically distributed, which may mask heterogeneous effects in different parts of the distribution. Specifically, we estimate the distributed lag model from Equation (1) using a pooled conditional quantile regression model to analyze the effects of fraud on the entire risk score distribution. The equation to be estimated is:

$$Y_{it} = \beta_0 + \sum_{e=-8}^{22} \beta_{1e\rho} T_e + \alpha_t + C_t + S_s + \varepsilon_{it}, \tag{6}$$

where T is a set of event time dummy variables relative to the filing of the extended fraud alert, α is a vector of calendar time dummy variables, C is a vector of extended alert filer cohort dummy variables, and S is a vector of state dummy variables. The coefficient vector $\beta_{1e\rho}$ for the event time dummy variables will be estimated for each quantile ρ.[39] Each $\beta_{1e\rho}$ can be interpreted as the increase in risk score at a given percentile for the specified event time T relative to the omitted period (quarters −22 to −9).[40]

In estimating their quantile regressions, Dobkin et al. (2016) examine only the right tail of the distribution for their continuous credit variables of interest. While this approach is valid for variables whose distribution is zero inflated, such as collections or card balances, it ignores potential heterogeneity in variables such as risk score, for which the distribution is less truncated. To understand the heterogeneity across the entire population, we estimate Equation (6) for 19 different quantiles. We present the results of the conditional quantile regressions for the center and the tails of the risk score distribution for fraud victims graphically in Figure 9.

[39] Since the model is estimated with a constant, the omitted categories are event times e =−22 to e = −9, the state of Arkansas, the first quarter of 2008, and extended alert filing cohort 1.

[40] It is important to note that the model neither follows consumers across time nor does it fix consumers in any specific quantile at a specific point in time. Although individuals are allowed to move across quantiles across time, the nature of the event time dummy variables ensures that there is no double counting of individuals in any given event time period in estimating the event time coefficients.

Panel A of Figure 9 displays the effects of identity theft for the left tail of the risk score distribution. Risk scores in this part of the distribution are subprime, with scores at $e = -8$ ranging from 445 for the 5th percentile to 543 for the 25th percentile. Panel B displays risk score from the central part of the distribution. Scores here are near prime with a median risk score of 631. Panel C shows risk scores for the right tail of the distribution. Scores in this region are highly prime, with the 90th percentile risk score being 805.

In the pre-identity theft period from $e = -8$ to $e = -2$, the risk score distribution is relatively stable, with event time coefficients not statistically different than zero. This can be seen in all three panels of Figure 9. Similar to our results in Part B in Section IV, we observe a significant drop in risk score at $e = -1$, which, as we have previously argued, is indicative of fraud. However, unlike our previous results, the estimates from the quantile regressions show that the immediate effect of fraud is concentrated in the right tail of the distribution, with risk scores at the 75th percentile experiencing a statistically significant decrease of −12.7 points. This can be seen in Panel C. Panel A shows that risk scores in the left tail of the distribution experience almost no immediate effect from identity theft, with point estimates at $e = -1$ being positive and statistically insignificant. This indicates that the effects of identity theft are heterogeneous, with those individuals with higher risk scores being disproportionally affected by fraud.[41]

There is also evidence of heterogeneity after filing an extended fraud alert. As can be seen in all three panels, risk scores across the entire distribution experience an increase from $e = -1$ to $e = 0$. While scores for the median and the left tail of the risk score distribution increase substantially, scores for the right tail of the distribution experience smaller effects. Our results show that the immediate gains in risk scores after filing an extended fraud alert are concentrated in the center and left tail of the distribution.[42]

In the post-extended alert filing period, there is substantial heterogeneity in the evolution of risk scores through time across the risk scores distribution. For risk scores in the left tail, the gains from fraud resolution diminish within one to two years: scores at the 5th percentile, which are deeply subprime, lose all gains by $e = 5$. Scores at the 25th percentile lose all their gains to

[41] This is intuitive; criminals are likely to be more successful in opening new accounts using the information of consumers with higher risk scores.

[42] The conditional median increases by 22 points from $e = -1$ to $e = 0$.

their risk score by $e = 10$. By $e = 22$, consumers in the left tail of the distribution have risk scores that are lower than in the pre-theft period, although we cannot reject the null hypothesis that they are significantly different from zero. This can be clearly seen in Panel A.

Panel B shows that increases in the central part of the risk score distribution are highly persistent, with the median remaining stable across all event time periods. This increase across time for the median risk score is economically significant as a substantial mass of consumers move from being subprime to prime upon filing an extended fraud alert.[43] In Panel C, we see that for scores in the right tail of the distribution, there is a steady increase over time, with scores in the 90th percentile approximately 9.5 points higher at $e = 22$ than in the pre-alert filing period.

To examine the overall effect of identity theft on risk scores across time, Figure 10 presents a heatmap of the risk score distribution with coefficient values represented by colors. Note that both the intensity and persistence of the effect on risk score is greatest for the center of the distribution when compared with the tails.[44]

These results highlight the existence of substantial heterogeneity in how fraud affects risk scores and the reponse of risk scores to identity theft. First, the immediate effects of identity theft on risk score are not equally dispsered, with risk scores in the right tail of the distribution being disporportionally affected. Second, the risk score distribution skews right upon the filing of an extended fraud alert, with scores in the tails of the risk score distribution demonstrating heterogenous levels of response: Scores in the left tail of the distribution show a 16.2 point increase in risk score relative to the previous quarter, while scores in the 90th percentile of the risk score distribution receive only a modest 3.4 point increase upon filing.[45] Third, these tail effects are transitory: for the 10th percentile of the risk score distribution, risk scores decrease consistently over time, with the initial gains eliminated by $e = 10$. In contrast, for scores the 90th percentile of the risk score distribution, risk scores are increasing over time and are 10 points higher than their pre-theft average by $e = 22$.

[43] The unconditional median risk score increases from 630 one quarter prior to filing an extended fraud alert to 664 five years after the alert has been filed. The conditional median risk score increases from 640 to 665 during the same time period.

[44] 3-D graphs showing how the risk score distribution changes over time are available in the Appendix.

[45] Note that risk scores for the 90th percentile are still on average below their pre-theft level at this time.

VII. Conclusion

This paper uses a unique data set of anonymized U.S. credit bureau records, including details on fraud alert filings, to examine the effects of identity theft on risk scores, access to credit, and credit portfolios. We classify those individuals who place an extended fraud alert in their credit bureau files as the most likely fraud victims because this type of fraud alert requires the filing of a police report (with accompanying evidence of identity theft).

Our results show that, in the short-term, identity theft decreases the average risk score of victims and increases new (fraudulent) accounts, inquiries, and instances of reverse address changes. The effects of fraud generally persist between one to two quarters. The long-term effects of fraud on credit bureau attributes are often positive and larger than the negative short-term effects. For identity theft victims, risk scores increase by an average of 12 points after fraud-related activity is eliminated from their credit reports. For many consumers, this effect is persistent over time and remains for as long as 20 quarters after the fraud incident. We also find that the average proportion of cards in good standing increases and the average incidence of third-party collections and other derogatory events decrease after identity theft and remain at lowered levels for several years.

This paper contributes to the existing literature on identity theft by documenting the effects of fraud on important consumer credit attributes. We also argue that the improvement in risk scores and other credit outcomes is consistent with identity theft victims paying limited attention to credit file information before identity theft and then subsequently increasing their attention to their credit portfolios.

Thus, we also contribute to the recent literature on the importance of salient events for investors and consumers in various contexts. Specifically, our results suggest that experiencing a negative shock to one's personal finances due to fraud and identity theft may be a "teachable moment" for consumers, increasing their awareness of financial and credit information.

Our findings, therefore, may have direct policy implications. While we obviously do not suggest imposing adverse events (such as identity theft) on consumers, we argue that individuals may learn better when they are more attentive to a specific subject, even if this attention is due to a negative shock. In particular, the manner in which the details of salient events are communicated to consumers may amplify the educational effects of any "teachable moments." In addition, providing consumers with tools to address issues directly related to the salient event, such as correcting credit information problems due to identity theft, as in FACTA, can produce tangible, economically significant improvements in consumers' outcomes. In other words, providing the right information at the moment when consumers care about the information may be an effective policy in improving their financial literacy and welfare.

Bibliography

Acquisti, Alessandro. (2004). "Privacy and Security of Personal Information." In *Economics of Information Security*, edited by Jean Camp and Stephen Lewis, 179–186. Boston: Kluwer Academic Publishers.

Anderson, Ross, and Tyler Moore. (2007). "Information Security Economics — and Beyond." In *Advances in Cryptology — CRYPTO 2007*, edited by Alfred Menezes, 68–91. New York: Springer Berlin Heidelberg.

Barber, Brad, and Terrence Odean. (2008). "All That Glitters: The Effect of Attention on the Buying Behavior of Individual and Institutional Investors." *Review of Financial Studies,* 21(2), 785–818.

Bracha, Anat, and Stephan 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.

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

Cheney, Julia, Robert Hunt, Katy Jacob, Richard Porter, and Bruce 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.

Cheney, Julia, Robert Hunt, Vyacheslav Mikhed, Dubravka Ritter, and Michael Vogan. (2014). "Consumer Use of Fraud Alerts and Credit Freezes: An Empirical Analysis." Federal Reserve Bank of Philadelphia Payment Cards Center Discussion Paper 14-04.

DellaVigna, Stefano, and Joshua Pollet. (2009). "Investor Inattention and Friday Earnings Announcements." *Journal of Finance,* 64(2), 709–749.

DellaVigna, Stefano. (2009). "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature*, 47(2), 315–372.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew Notowidigdo. (2016). "The Economic Consequences of Hospital Admissions." National Bureau of Economic Research Working Paper 22288.

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, Xavier, and David Laibson. (2000). "A Boundedly Rational Decision Algorithm." *American Economic Review*, 90(2), 433–438.

Gabaix, Xavier, and David Laibson. (2002). "The 6D Bias and the Equity Premium Puzzle." *NBER Macroeconomics Annual 2001,* 16, 257–312.

Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg. (2006). "Costly Information Acquisition: Experimental Analysis of Boundedly Rational Model." *American Economic Review*, 96(4), 1043–1068.

Gallagher, Justin. (2014). "Learning About an Infrequent Event: Evidence from Flood Insurance Take-Up in the United States," *American Economic Journal: Applied Economics*, 6(3), 206–233.

Gerdes, Geoffrey, and May Liu. (2014). The 2013 Federal Reserve Payments Study, Recent and Long-Term Trends in the United States: 2000–2012, Detailed Report and Updated Data Release. Washington, D.C.: Federal Reserve System.

Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang. (2014). "Liquidity Constraints and Consumer Bankruptcy: Evidence from Tax Rebates." *Review of Economics and Statistics*, 96 (3): 431–443.

Grubb, Michael D. (2015). "Consumer Inattention and Bill-Shock Regulation." *Review of Economic Studies*, 82(1), 219–257.

Harrell, Erika, and Lynn Langton. (2013). "Victims of Identity Theft, 2012." U.S. Department of Justice, Bureau of Justice Statistics, Bulletin NCJ 243779.

Harrell, Erika. (2015). "Victims of Identity Theft, 2014." U.S. Department of Justice, Bureau of Justice Statistics, Bulletin NCJ 248991.

Hirshleifer, David, Sonya Lim, and Siew Hong Teoh. (2009). "Driven to Distraction: Extraneous Events and Underreaction to Earnings News." *Journal of Finance*, 64(5), 2289–2325.

Hirshleifer, David, Sonya Lim, and Siew Hong Teoh. (2011). "Limited Investor Attention and Stock Market Misreactions to Accounting Information." *Review of Asset Pricing Studies*, 1(1), 35–73.

Hunt, Robert. (2005). "A Century of Credit Reporting in America." Federal Reserve Bank of Philadelphia Working Paper 05-13.

Kahn, Charles, and José Liñares-Zegarra. (2016). "Identity Theft and Consumer Payment Choice: Does Security Really Matter?" *Journal of Financial Services Research*, 50(1), 121–159.

Kosse, Anneke. (2013). "Do Newspaper Articles on Card Fraud Affect Debit Card Usage?" *Journal of Banking and Finance*, 37(12), 5382–5391.

Lacetera, Nicola, Devin Pope, and Justin Sydnor. (2012). "Heuristic Thinking and Limited Attention in the Car Market." *American Economic Review*, 102(5), 2206–2236.

Lee, Donghoon, and Wilbert van der Klaauw. (2010). "An Introduction to the FRBNY Consumer Credit Panel." Federal Reserve Bank of New York Staff Report 479.

Musto, David K. (2004). "What Happens When Information Leaves a Market? Evidence from Postbankruptcy Consumers." *Journal of Business*, 77(4), 725–748.

Sims, Christopher A. (2003). "Implications of Rational Inattention." *Journal of Monetary Economics*, 50(3), 665–690.

Smith, Douglas, Michael Staten, Thomas Eyssell, Maureen Karig, Beth Freeborn, and Andrea Golden. (2013). "Accuracy of Information Maintained by U.S. Credit Bureaus: Frequency of Errors and Effects on Consumers' Credit Scores." *Journal of Consumer Affairs*, 47(3), 588–601.

Stango, Victor, and Jonathan Zinman. (2014). "Limited and Varying Consumer Attention: Evidence from Shocks to the Salience of Bank Overdraft Fees." *Review of Financial Studies*, 27(4), 990–1030.

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

Sullivan, Richard. (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).

Wardrip, Keith, and Robert 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.

Yang, Jia Lynn, and Amrita Jayakumar. (2014). "Data on 70 Million Taken in Growing Target Breach." *Washington Post*, January 11.

Table 1. Summary Statistics

Variable	Nonfilers Mean	Nonfilers S.D.	Filers At $t = -4$: Mean	Filers At $t = -4$: S.D.	Filers At $t = 0$: Mean	Filers At $t = 0$: S.D.
Number of inquiries, past 3 months	0.541	0.989	1.028	1.646	1.480	2.027
Number of inquiries, past 12 months	1.873	2.217	3.410	3.896	3.861	3.995
Age of newest account	31.877	46.977	19.052	23.115	19.893	25.786
Change in number of accounts opened, past 6 months	−0.007	0.680	0.006	0.864	−0.048	1.006
Change in number of revolving accounts, past 6 months	−0.008	0.511	0.003	0.638	−0.072	0.797
Utilization rate (fraction)	0.309	0.367	0.393	0.411	0.390	0.397
Age (years)	51.004	17.919	43.943	14.742	44.869	14.744
Number of bankcard accounts w/ update w/in 3 months w/ balance>0	1.600	1.459	1.811	1.569	1.680	1.491
Number of trades currently 30 days past due	0.045	0.267	0.064	0.307	0.051	0.277
Number of bankcard accounts with past due amount>0	0.194	0.665	0.372	0.896	0.235	0.663
Total number of 30 days past due occurrences on bankcards w/in 24 months	0.382	1.417	0.596	1.702	0.394	1.255
Total number of 120 days past due occurrences on bankcards w/in 24 months	0.601	3.236	0.989	3.959	0.542	2.908
Total past due amount bankcard accounts w/ update w/in 3 months	344	2638.2	553.4	3571.3	403.1	3042.5
Risk score	695.11	108.18	641.34	119	655.1	109
Number of Observations	21,517,164		39,367		39,404	

Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center. Note: Risk score is the Equifax Risk Score.

Table 2. The Effect of Fraud on Credit Variables After Controlling for Event Time Trends

Variables	(1) Risk Score	(2) Inquiries	(3) Reverse Mobility	(4) New Revolving Accounts	(5) % Cards in Good Standing	(6) Accounts with Positive Balances	(7) Open Cards	(8) Collections	(9) Derogatory Events
Time	0.882***	−0.006	0.0003	−0.003	0.001	0.015**	−0.012***	0.002**	−0.007***
	(0.183)	(0.004)	(0.0002)	(0.002)	(0.001)	(0.007)	(0.005)	(0.001)	(0.001)
$Time^2$	0.005	0.000	0.0000	0.000	0.000***	0.000	0.000**	0.000	−0.000***
	(0.008)	(0.000)	(0.0000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
1(Time≥0)	11.26***	0.302***	0.0113***	0.004	0.023***	−0.271***	−0.084***	−0.082***	−0.005
	(0.933)	(0.023)	(0.0013)	(0.012)	(0.005)	(0.035)	(0.022)	(0.006)	(0.006)
1(Time≥0)×Time	−0.491**	−0.093***	−0.0019***	−0.021***	−0.004***	0.007	−0.019***	0.002*	0.006***
	(0.195)	(0.004)	(0.0002)	(0.002)	(0.001)	(0.008)	(0.005)	(0.001)	(0.001)
1(Time≥0)×$Time^2$	0.010	0.003***	0.0001***	0.000***	−0.000***	−0.002***	−0.001***	−0.000***	0.000***
	(0.008)	(0.000)	(0.0000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
1(−4≤Time≤−1)	−5.28***	0.388***	0.0048***	0.105***	−0.003	0.097***	0.044**	−0.019***	0.035***
	(0.922)	(0.025)	(0.0013)	(0.012)	(0.005)	(0.035)	(0.022)	(0.006)	(0.006)
1(−4≤Time≤−1)×Time	−1.38***	0.102***	0.0010***	0.026***	−0.001	0.025***	0.011**	−0.004***	0.008***
	(0.185)	(0.005)	(0.0003)	(0.003)	(0.001)	(0.007)	(0.004)	(0.001)	(0.001)
Constant	643.45***	1.174***	0.0018	0.441***	0.837***	4.789***	2.482***	0.279***	0.104***
	(0.924)	(0.022)	(0.0012)	(0.011)	(0.005)	(0.036)	(0.022)	(0.006)	(0.006)
Observations	1,190,342	1,039,903	1,210,927	1,109,464	838,848	1,136,894	1,053,260	1,210,927	1,210,927
R-squared	0.035	0.021	0.0028	0.017	0.008	0.005	0.070	0.005	0.011

Notes: This table shows the effect of fraud on credit outcomes indicated in the first row modeled using Equation (3). 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 resolution (time≥0). This model includes individual fixed effect and calendar time fixed effects. The results indicate that there are negative changes in credit variables at fraud time and positive changes after fraud resolution. Standard errors are clustered at the individual level and reported in parentheses. ***, **, and * indicate statistical significance at 1%, 5%, and 10% levels. Risk score is the Equifax Risk Score.

Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center

Table 3. The Effect of Fraud on Credit Variables, Controlling for Individual Mean Reversion

	(1) Risk Score	(2) Cards in Good Standing	(3) New Revolving Accounts
After extended alert filed	13.606*** (1.238)	0.018*** (0.006)	−0.085*** (0.019)
Number of panels	3520	3234	3438
Total observations	79715	55995	74179
Within R^2	0.553	0.565	0.317
Overall R^2	0.929	0.844	0.446

Notes: All specifications include individual fixed effects, individual quadratic time trends, and time fixed effects. Standard errors are clustered at the individual level. Risk score is the Equifax Risk Score.

Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center

Figure 1. Indicators of Potential Identity Theft in Event Time

Panel A. Credit Inquiries

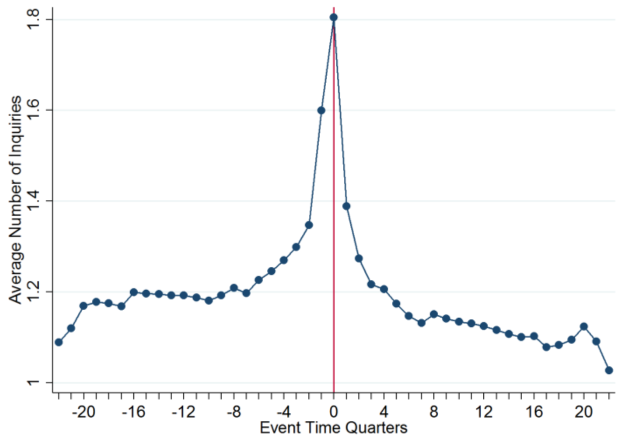

Panel B. Risk Score

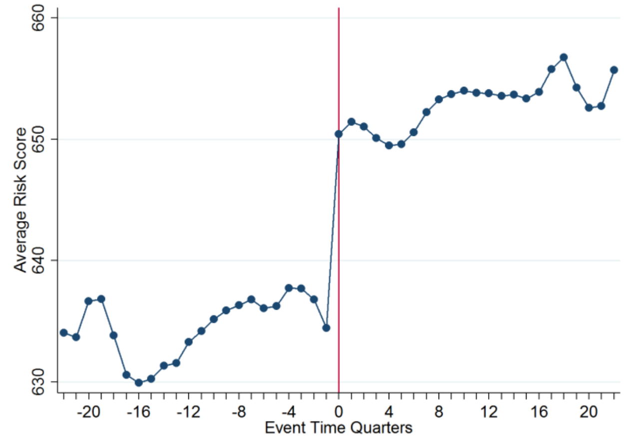

Panel C. Address Reversals

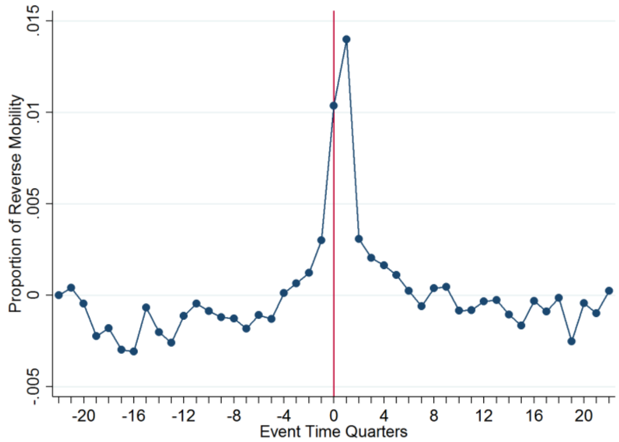

Panel D. New Accounts

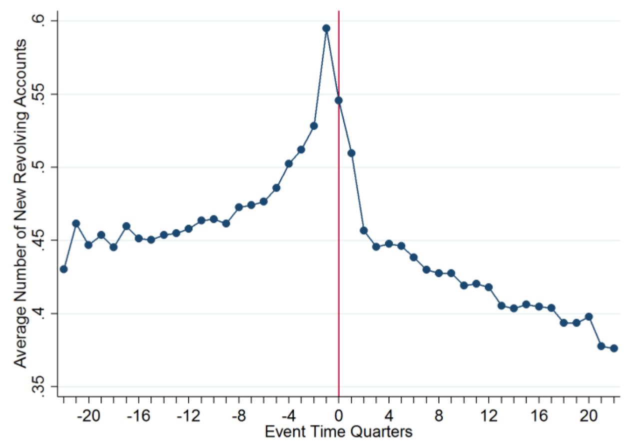

Notes: This figure depicts average values of credit bureau characteristics of fraud victims before and after fraud activity. Time 0 denotes the quarter of extended fraud alert filing with negative time being quarters before this event and positive time being quarters after the event. The data include only extended fraud alert filers in Q1:2008–Q3:2013. The effect of the business cycle is removed using calendar time dummies. Risk score is the Equifax Risk Score.

Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center

Figure 2. Indicators of Potential Identity Theft — Treatment versus Control

Panel A. Credit Inquiries

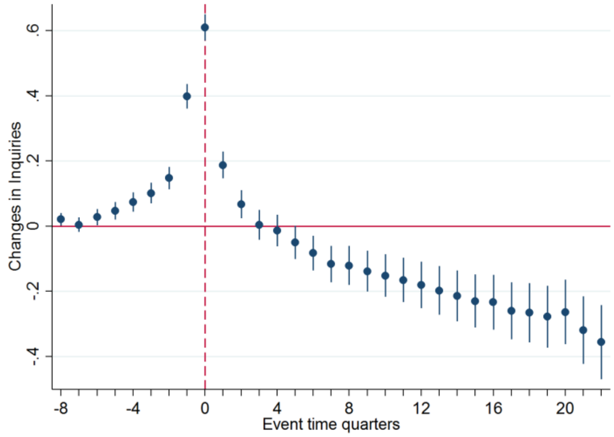

Panel B. Risk Score

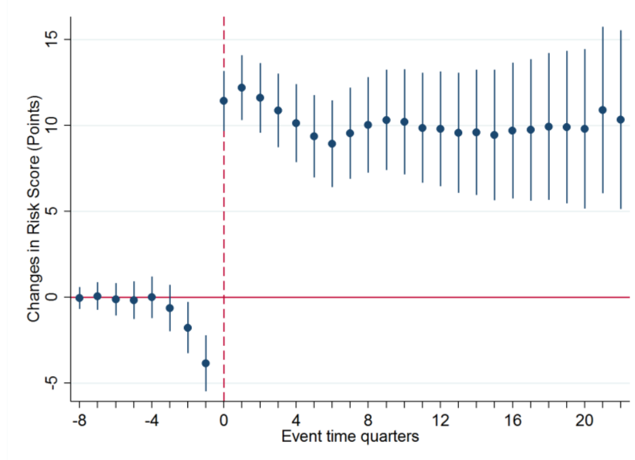

Panel C. Address Reversals

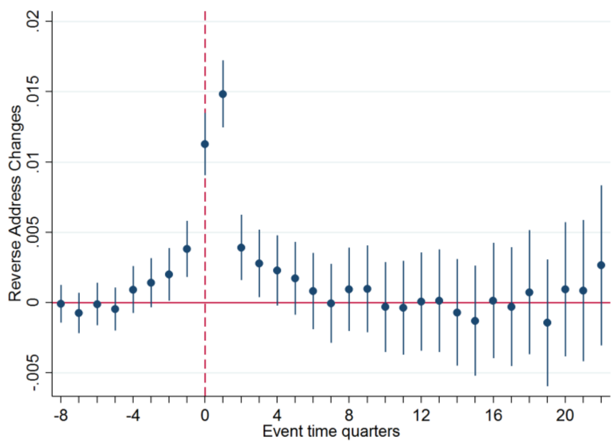

Panel D. New Accounts

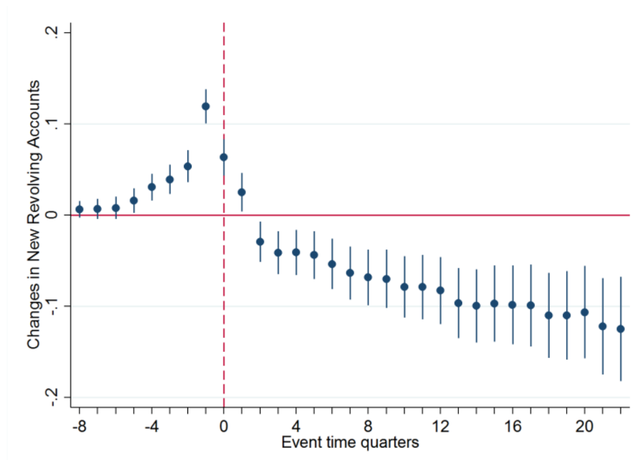

Notes: This figure depicts changes in credit bureau characteristics of fraud victims before and after fraud activity. These changes are estimated based on a distributed lag specification with event time dummies. Time 0 denotes the quarter of fraud alert filing, with negative time being quarters before this event and positive time being quarters after the event. All quarter dummies prior to quarter −8 are omitted. The dots represent point estimates and bands show 95% confidence intervals. The data include only fraud alert filers in Q1:2008–Q3:2013. The idenification comes from the exogenous timing of fraud activity. Risk score is the Equifax Risk Score.

Sources: Authors’ calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center

Figure 3. The Effect of Fraud on the Number of Open Cards, Cards with Balances, and Age of Newest Card

Panel A. Open Bankcards

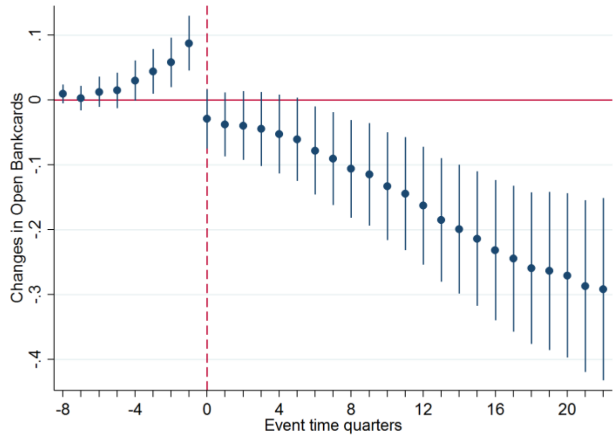

Panel B. Cards with Positive Balances

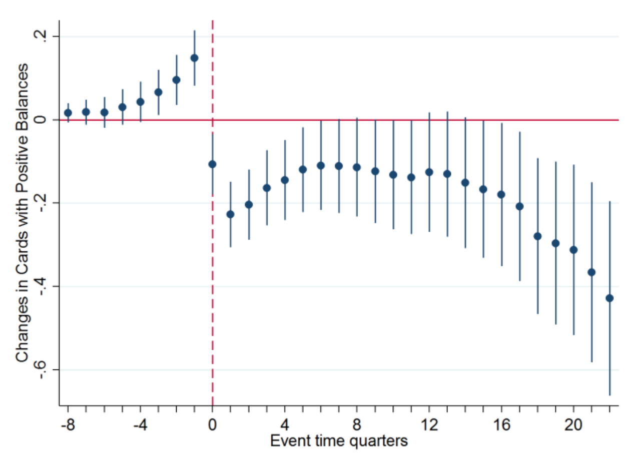

Panel C. Age of Newest Bankcard

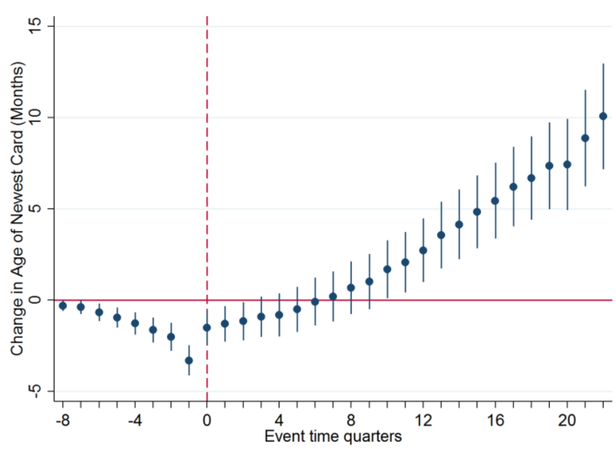

Notes: This figure depicts changes in credit bureau characteristics of fraud victims before and after fraud activity. These changes are estimated based on a distributed lag specification with event time dummies. Time 0 denotes the quarter of fraud alert filing, with negative time being quarters before this event and positive time being quarters after the event. All quarter dummies prior to quarter −8 are omitted. The dots represent point estimates and bands show 95% confidence intervals. The data include only fraud alert filers in Q1:2008–Q3:2013. The identification comes from the exogenous timing of fraud activity.

Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center

Figure 4. The Effect of Fraud on the Proportion of Cards in Good Standing, the Incidence of Major Derogatory Events, and Third-Party Collections

Panel A. Cards in Good Standing

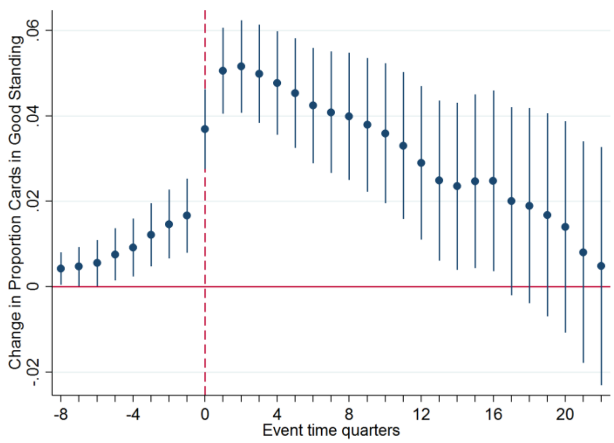

Panel B. Major Derogatory

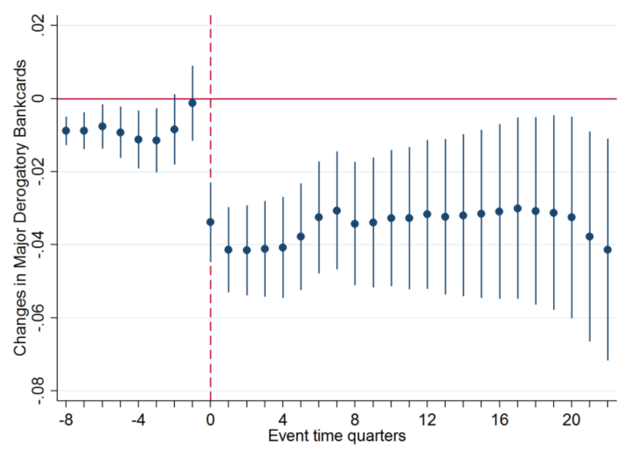

Panel C. Third-Party Collections

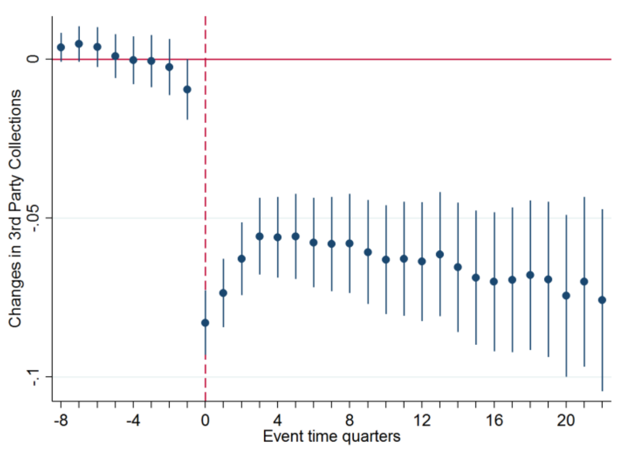

Notes: This figure depicts changes in credit bureau characteristics of fraud victims before and after fraud activity. These changes are estimated based on a distributed lag specification with event time dummies. Time 0 denotes the quarter of fraud alert filing, with negative time being quarters before this event and positive time being quarters after the event. All quarter dummies prior to quarter −8 are omitted. The dots represent point estimates and bands show 95% confidence intervals. The data include only fraud alert filers in Q1:2008–Q3:2013. The identification comes from the exogenous timing of fraud activity.

Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center

Figure 5. Effect of Fraud on the Share of Consumers with Prime (> 660) Risk Score

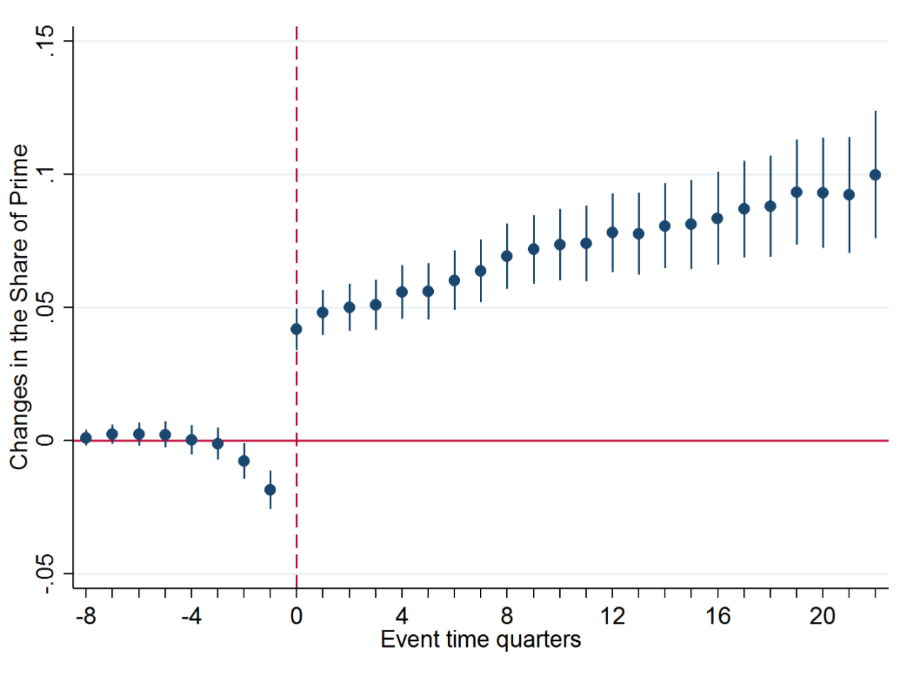

Notes: This figure depicts changes in credit bureau characteristics of fraud victims before and after fraud activity. These changes are estimated based on a distributed lag specification with event time dummies. Time 0 denotes the quarter of fraud alert filing, with negative time being quarters before this event and positive time being quarters after the event. All quarter dummies prior to quarter −8 are omitted. The dots represent point estimates and bands show 95% confidence intervals. The data include only fraud alert filers in Q1:2008–Q3:2013. The identification comes from the exogenous timing of fraud activity. Risk score is the Equifax Risk Score.

Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center

Figure 6. The Effect of Fraud on Credit Bureau Variables Using Initial Alert Filers as Controls

Panel A. Credit Inquiries

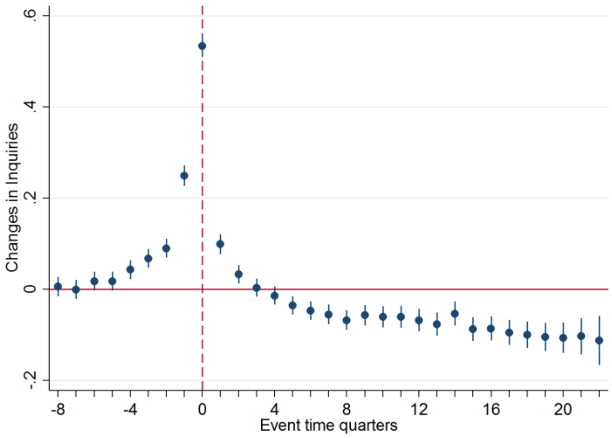

Panel B. Risk Score

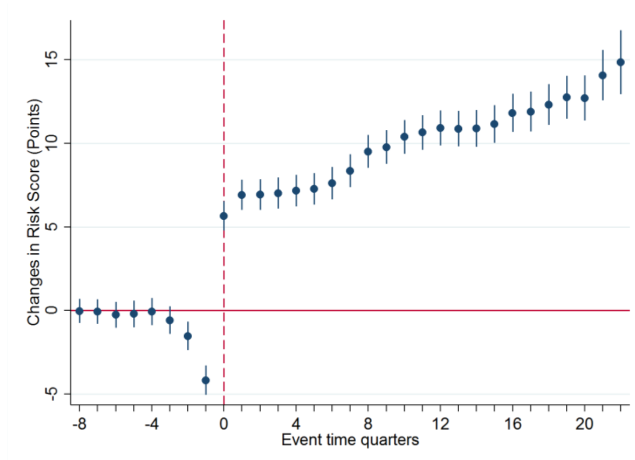

Panel C. Address Reversals

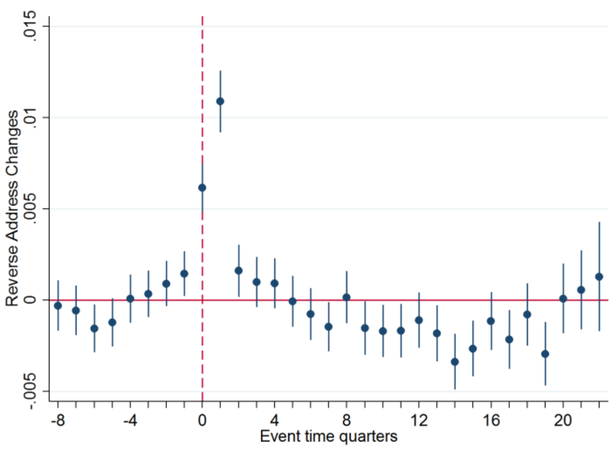

Panel D. New Accounts

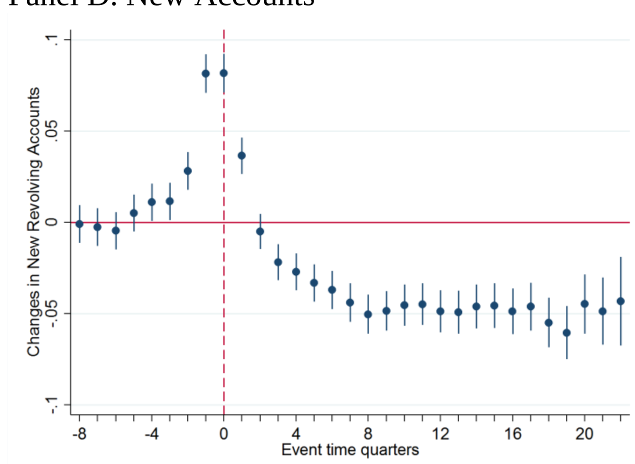

Notes: This figure depicts changes in credit bureau characteristics of fraud victims before and after fraud activity. These changes are estimated based on a distributed lag specification with event time dummies. Time 0 denotes the quarter of fraud alert filing, with negative time being quarters before this event and positive time being quarters after the event. All quarter dummies prior to quarter −8 are omitted. The dots represent point estimates and bands show 95% confidence intervals. The treatement group includes extended fraud alert filers. The control group consists of intial fraud alert filers. Risk score is the Equifax Risk Score.

Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center

Figure 7. The Effect of Fraud on Consumers Without Credit Inquiries Compared with Consumers with Credit Inquiries

Panel A. Credit Inquiries

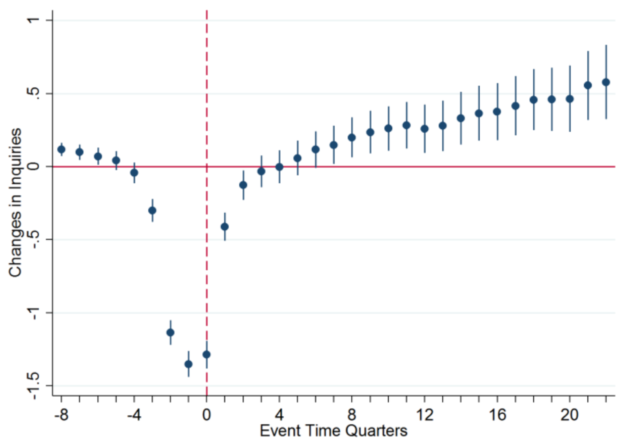

Panel B. Risk Score

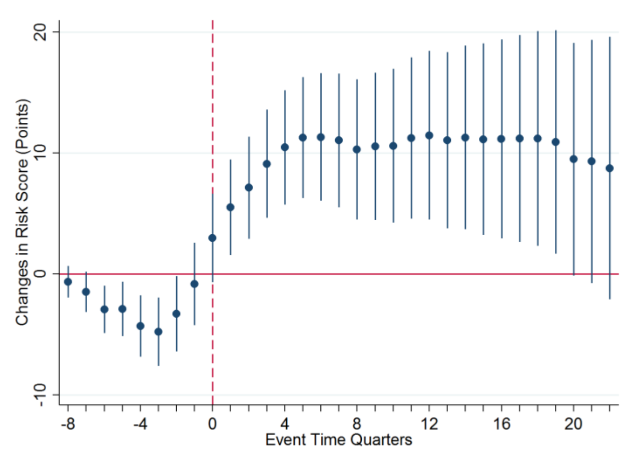

Panel C. Address Reversals

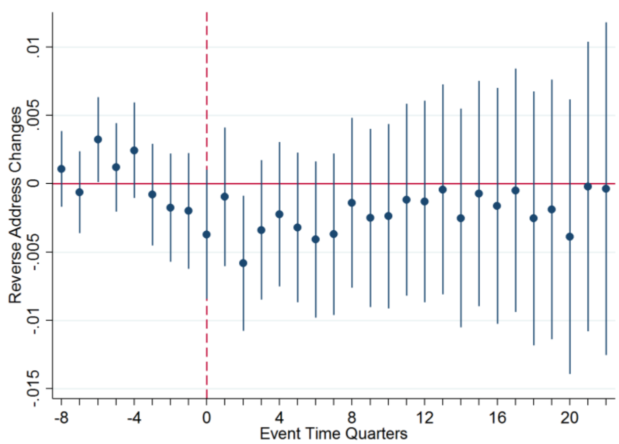

Panel D. New Accounts

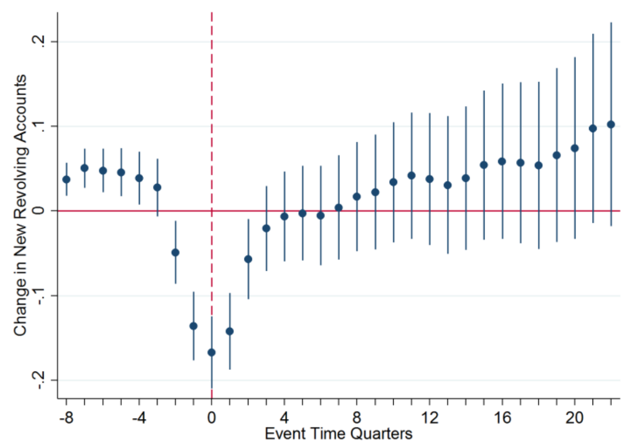

Notes: This figure depicts changes in credit bureau characteristics of fraud victims without credit inquiries at the time of fraud alert and the two quarters before that relative to credit bureau variables of fraud victims with credit inquiries in the same period. These changes are estimated based on a distributed lag specification with event time dummies interacted with the no inquiry/ inquiry indicator. 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. These figures imply that less attentive consumers who were not shopping for credit before fraud (or who had no new account fraud) have larger increase in risk score than more attentive consumer (who were shopping for credit). All quarter dummies prior to quarter −8 are omitted. The dots represent point estimates and bands show 95% confidence intervals. Risk score is the Equifax Risk Score.

Sources: Authors’ calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center

Figure 8. The Effect of Fraud on Consumers with Subprime Risk Scores Compared with Consumers with Prime Risk Scores

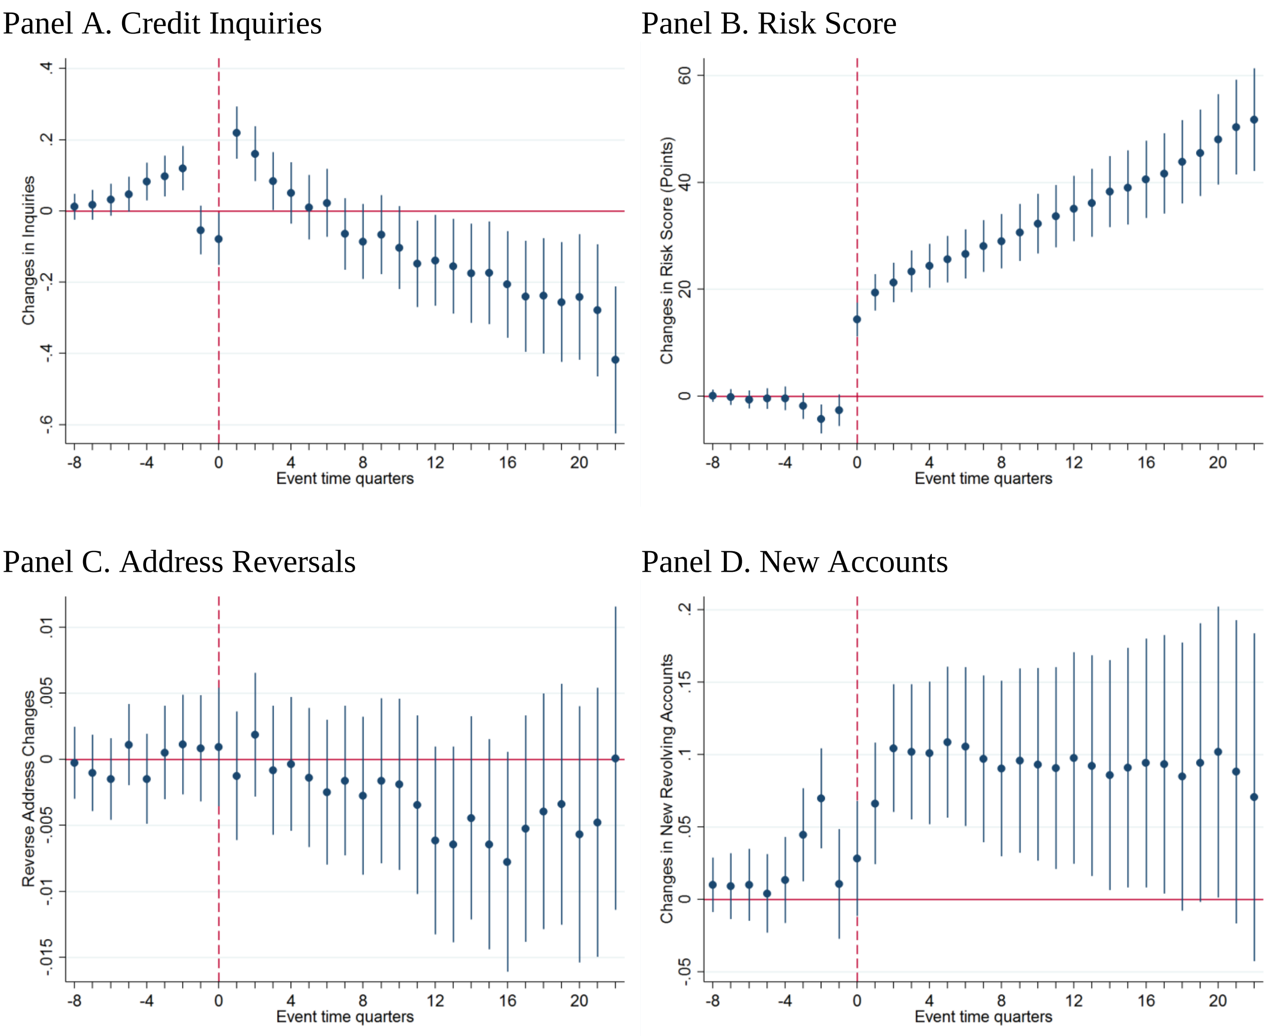

Notes: This figure depicts changes in credit bureau characteristics of subrime (less than or equal to 660) and prime (more than 660) fraud victims as of the time of fraud alert and the two quarters before that. These changes are estimated based on a distributed lag specification with event time dummies interacted with the subprime/ prime indicator. 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 quarter dummies prior to quarter −8 are omitted. The dots represent point estimates and bands show 95% confidence intervals. Risk score is the Equifax Risk Score.

Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center

Figure 9. Heterogeneity in the Effect of Fraud on Risk Score, by Percentile

Panel A. Left Tail of the Risk Score Distribution

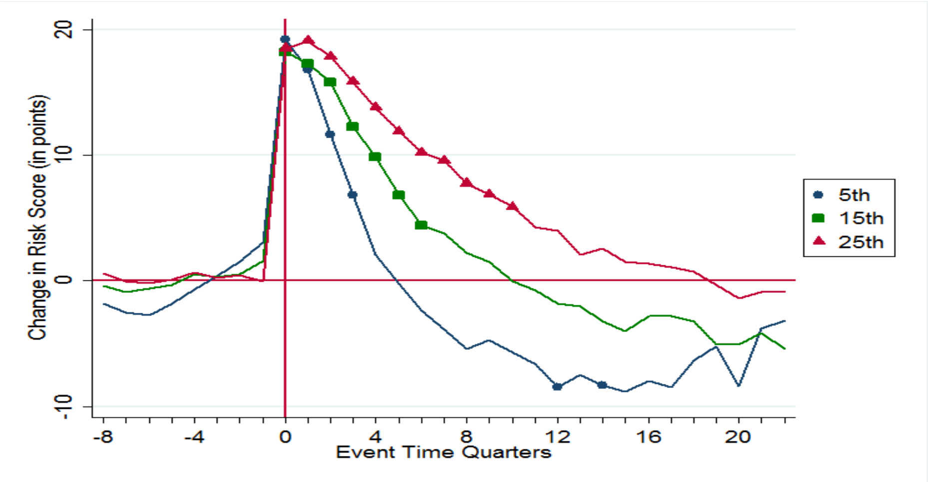

Panel B. Center of the Risk Score Distribution

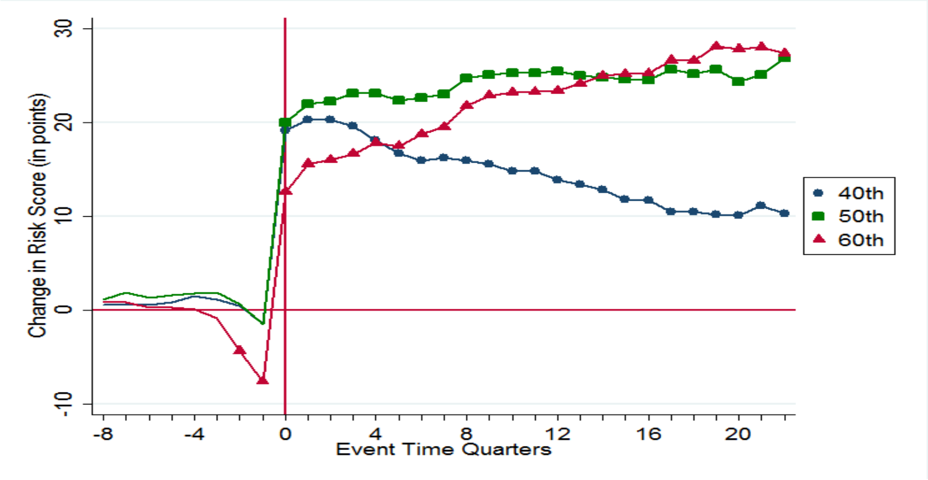

Panel C. Right Tail of the Risk Score Distribution

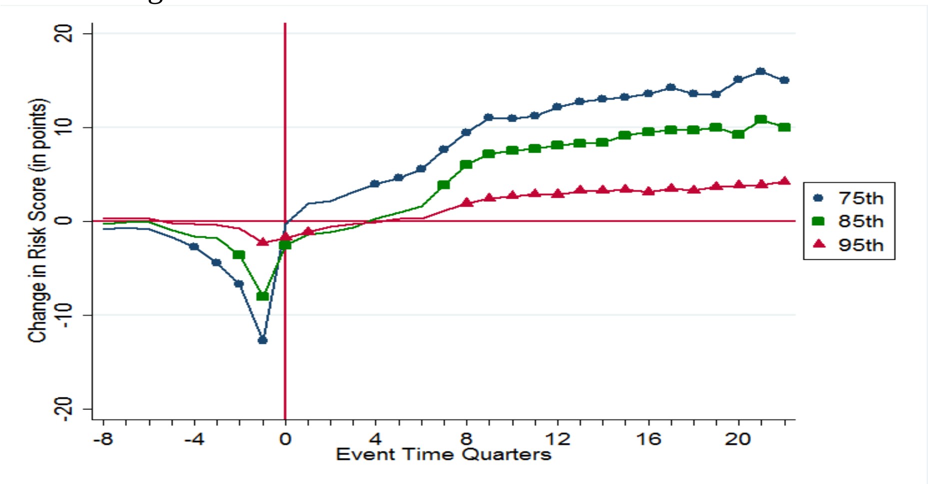

Notes: Only statistically significant coefficients are shown. Robust standard errors are calculated. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center. Risk score is the Equifax Risk Score.

Figure 10. Heatmap for the Effect of Fraud on Risk Score Across Time

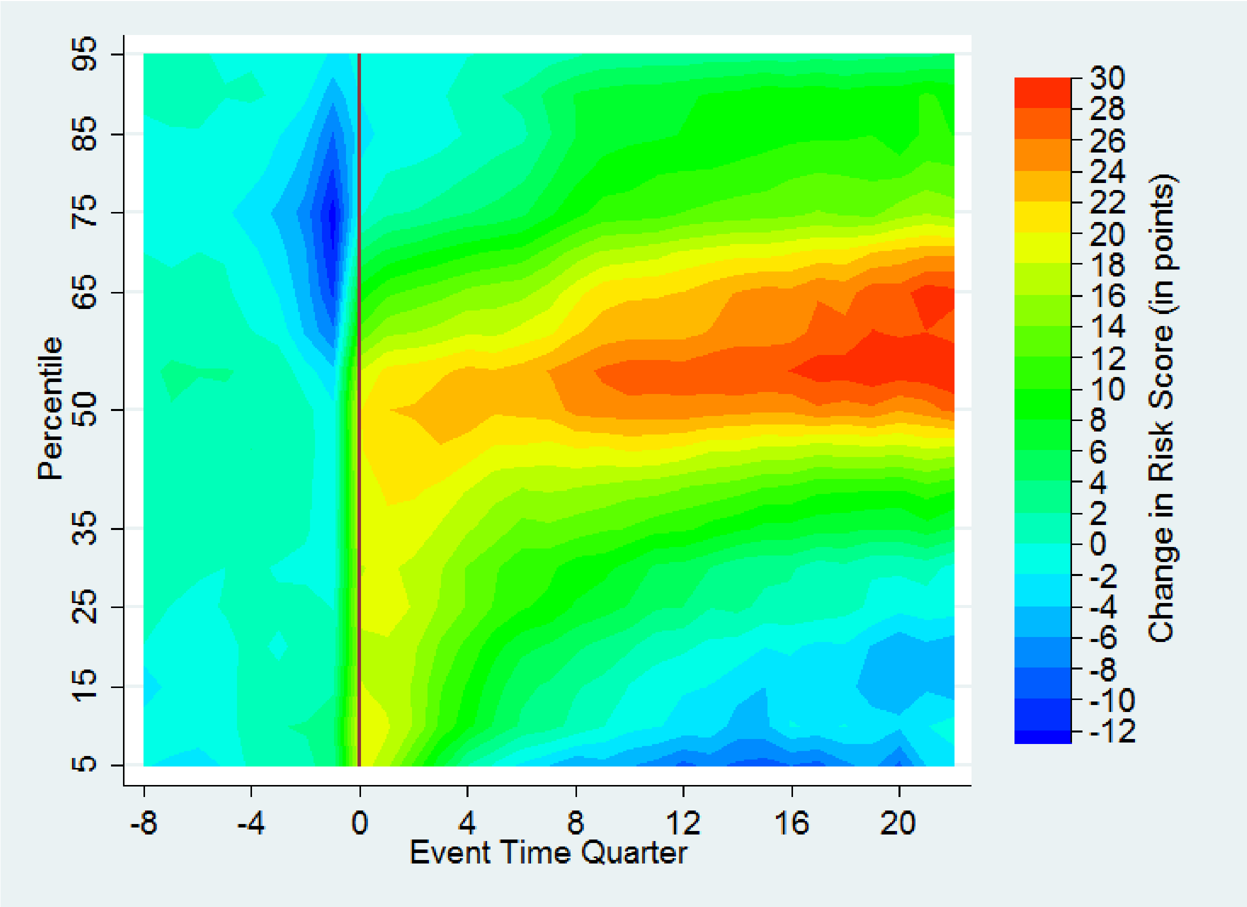

Note: Both statistically significant and insignificant coefficients are shown. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center. Risk score is the Equifax Risk Score.

Appendix

Figure A1. The Effect of Fraud on Credit Bureau Variables by Month of Alert Filing

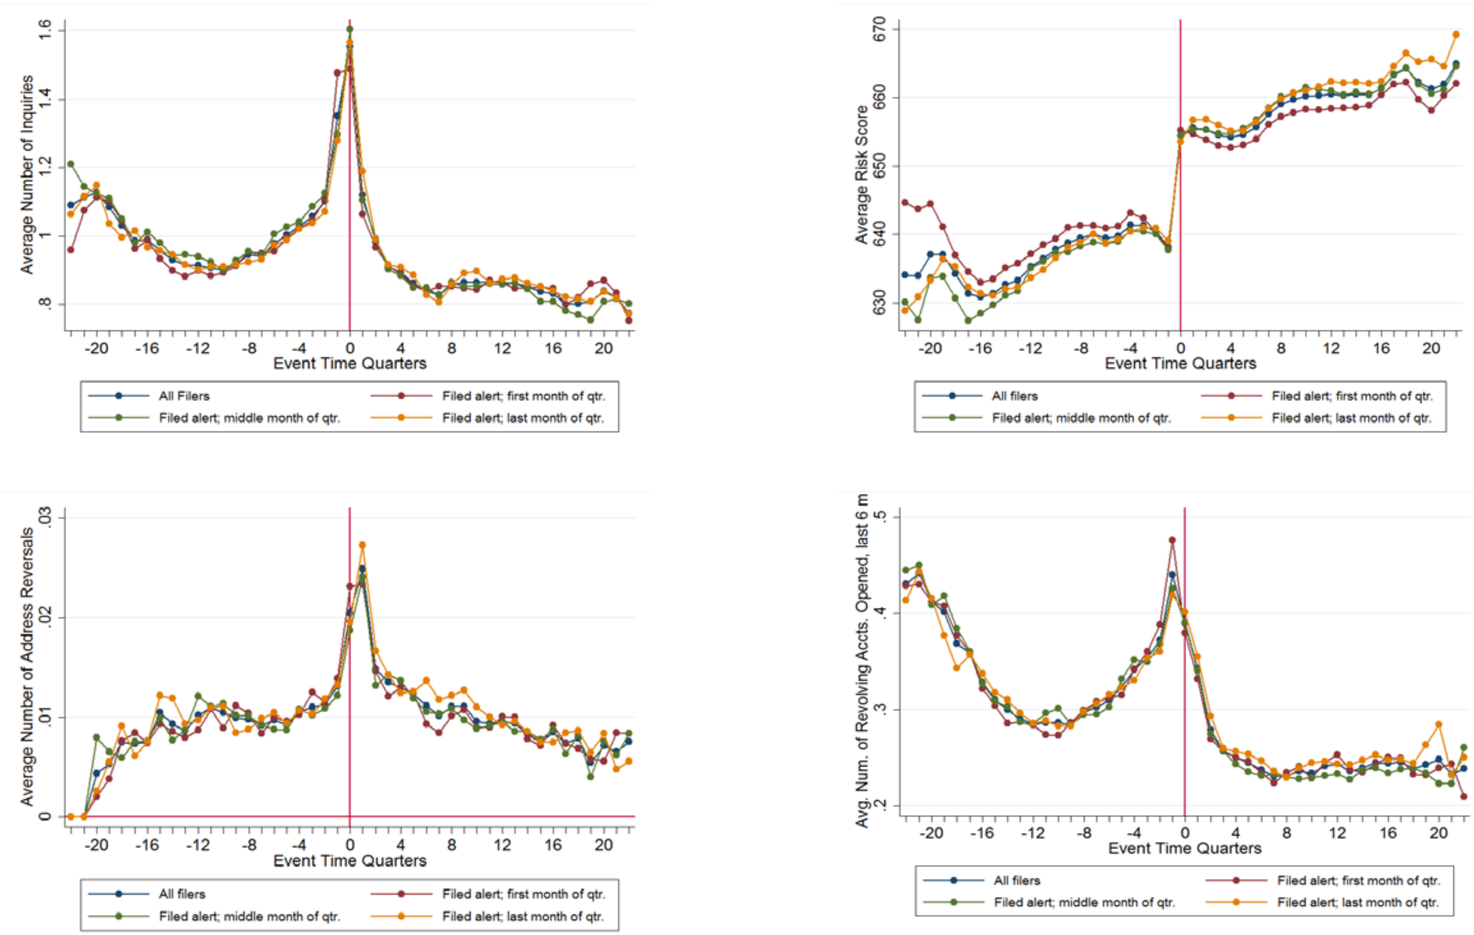

Notes: This figure depicts average values of credit bureau characteristics of fraud victims before and after fraud activity by the month of extended alert filing. Time 0 denotes the quarter of extended fraud alert filing with negative time being quarters before this event and positive time being quarters after the event. Risk score is the Equifax Risk Score. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center

Figure A2. The Effect of Fraud on the Risk Score Distribution

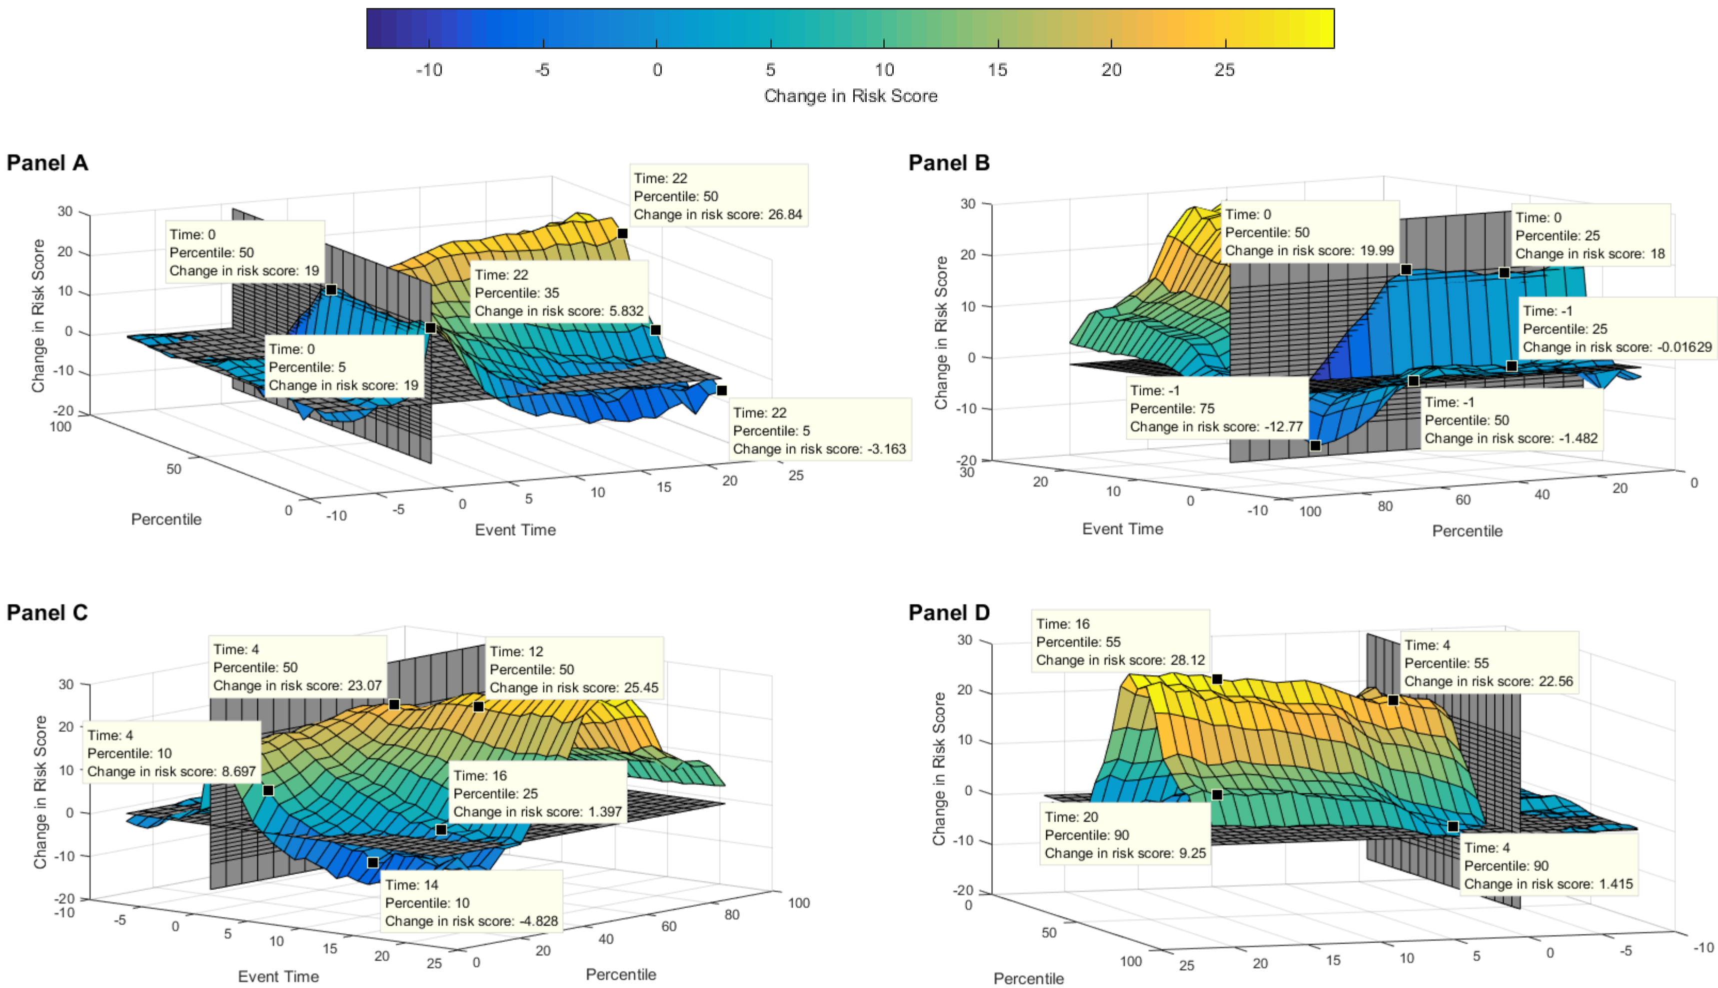

Notes: This figure depicts average values of credit bureau characteristics of fraud victims before and after fraud activity. The vertical gray plane at Time 0 denotes the quarter of extended fraud alert filing with negative time being quarters before this event and positive time being quarters after the event. The horitontal gray plane denotes zero change in the risk score. Risk score is the Equifax Risk Score.

Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center