WORKING PAPER NO. 14-28
IDENTITY THEFT AS A TEACHABLE MOMENT

Julia Cheney
Federal Reserve Bank of Philadelphia

Robert Hunt
Federal Reserve Bank of Philadelphia

Vyacheslav Mikhed
Federal Reserve Bank of Philadelphia

Dubravka Ritter
Federal Reserve Bank of Philadelphia

Michael Vogan
Federal Reserve Bank of Philadelphia

September 2014

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

Julia Cheney*

Robert Hunt

Vyacheslav Mikhed

Dubravka Ritter

Michael Vogan

Payment Cards Center, Federal Reserve Bank of Philadelphia

September 2014

ABSTRACT

This paper examines how instances of identity theft that are sufficiently severe to induce consumers to place an extended fraud alert in their credit reports affect their risk scores, delinquencies, and other credit bureau variables on impact and thereafter. We show that for many consumers these effects are relatively small and transitory. However, for a significant number of consumers, especially those with lower risk scores prior to the event, there are more persistent and generally positive effects on credit bureau variables, including risk scores. We argue that these positive changes for subprime consumers are consistent with the effect of increased salience of credit file information to the consumer at the time of the identity theft.

Keywords: Inattention, identity theft, fraud alert, consumer protection, credit report, Fair and Accurate Credit Transactions Act (FACTA), propensity score matching

JEL Codes: D14, D18, G02

* Contact: Julia Cheney, Payment Cards Center, Federal Reserve Bank of Philadelphia, Ten Independence Mall, Philadelphia, PA 19106; e-mail: julia.cheney@phil.frb.org. 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 Susan Herbst-Murphy, Blake Prichard, Peter Schnall, Chet Wiermanski**,** and Stephanie Wilshusen for helpful suggestions. We especially thank Loretta Mester for making this research possible. This paper is available free of charge at http://www.philadelphiafed.org/consumer-credit-and-payments/payment-cards-center/publications/.

I. Introduction

According to a U.S. Department of Justice report, more than 16 million U.S. consumers were victims of identity theft in 2012. These victims bore out-of-pocket costs of about $11 billion (Harrell and Langton, 2013).[1] The Federal Reserve estimates unauthorized transactions initiated via check, automated clearinghouse (ACH), and credit and debit cards exceeded $6 billion in 2012 (Gerdes and Liu, 2014).[2] Based on the anonymized credit records we analyze in this paper, we estimate that about 2.0 million consumers placed an alert of some sort in their credit bureau records in 2012.[3] 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 above, 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.[4] We contribute new information about both of these effects. Specifically, we investigate the immediate and longer-term effects of identity theft as measured by data contained in the credit reports of consumers who file an "extended fraud alert"[5] with a credit bureau.[6] First, we document that the timing of these alerts is correlated with

[1] This report is based on an Identity Theft Supplement to the National Crime Victimization Survey. The victim and loss data are for persons 16 years of age and older. The authors note that double counting of losses may occur when account holders each report losses from an account jointly held.

[2] In Gerdes and Liu (2014), an unauthorized transaction is defined as "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 a form of existing account fraud committed by identity thieves.

[3] In the United States, a credit bureau record represents a record of an individual's borrowing and repayment activity, including information on applications for credit (inquiries), loan balances, and delinquencies. These records are also referred to as "credit reports" or "credit bureau files." For more information on credit reporting in the United States, see Hunt (2005).

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

[5] Extended fraud alerts require individuals to file an official report — such as a police report — with the credit reporting agency. Therefore, we believe that consumers who file these alerts are more likely to be victims of identity theft fraud than are consumers who file other forms of alerts that appear more precautionary in nature. For more detail, see Section III herein and Cheney et al. (2014).

changes in variables (such as applications for credit and address changes) that are indicative of criminals impersonating a consumer. Second, we measure changes in indicators (such as risk scores) that affect access to and the use of credit. For many consumers, we document that there is an immediate, negative effect of an identity theft event on credit bureau attributes. This effect is usually reversed within a few months. This reversal is likely due to the removal of information from the credit bureau files about compromised accounts or accounts opened fraudulently by criminals and to disputes of any delinquencies associated with those accounts.[7] Thus, for many consumers, especially ones with prime risk scores a year or more prior to filing an extended fraud alert, there is a moderate, albeit transient, negative effect on their credit bureau attributes.

Third, and somewhat surprisingly, we find evidence that an identity theft event generates *persistent* changes in the credit bureau files of some victims, especially among consumers who had subprime risk scores a year or more prior to the event. For many of these consumers, risk scores increase in the quarter after the alert is filed and by an amount significantly greater than the decline in scores that usually occurs at the time the alert is filed. Not only is this increase statistically significant in absolute terms, it is also significant relative to the trend in risk scores among similar subprime consumers (selected using propensity score matching) who do not file any alert or credit freeze.[8] Using this approach, we find that the average risk scores of the "treated" and matched "control" groups are similar prior to the event, but scores for the treated group are at least 10 points higher than for the control group for two or more years after the event. This difference in risk scores coincides with persistent reductions in the share and number of credit cards past due, delinquencies, and third-party collections. Thus, for many subprime consumers with extended fraud alerts, an initial negative effect is followed by a persistent and larger improvement to their credit bureau attributes.

What explains this divergence in average outcomes for prime and subprime consumers exposed to identity theft? We argue that, on average, subprime consumers may initially be less

---

[6] The focus of this paper is on the effects of identity theft on consumer outcomes; the effects on merchants, financial institutions, and other entities are outside the scope of this study.

[7] The timing and extent of this "cleaning" process influences what we can and cannot measure in our data. In Section III, we describe how this process influences the way we analyze the data.

[8] As we describe in Section III, we implement our treatment–control approach using cohorts of consumers who likely experienced an identity theft event in a specific calendar quarter. We do this to allow for differences in selection as well as the magnitude of fraud at different points in time and to control for business cycle effects.

attentive to their credit records prior to experiencing an identity theft event. As discussed previously, information about fraudulent accounts is typically removed at the time the consumer files the extended fraud alert. But it is also likely the consumer will dispute other erroneous or outdated information that has accumulated in his or her credit file.[9] Thereafter, the patterns we observe in our data suggest that a typical subprime consumer with an extended fraud alert is, for some time at least, either more careful about using credit or more careful about checking credit reports and disputing erroneous information, or both. Among prime consumers who we might expect were paying more attention to their credit files and their credit more generally, the effect of this *teachable moment* is likely much smaller.

We organize the remainder of the paper as follows. Section II provides a brief overview of the literature. Section III describes the institutional detail behind identity theft and fraud alerts, our data, and methodology. We review our main results in Section IV. Sections V and VI discuss a number of extensions and robustness checks. We conclude in Section VII.

II. Existing Literature

This paper contributes to the existing literature on identity theft and fraud. Several papers consider the consequences of identity theft on consumer confidence in payment systems to move money between parties to a transaction efficiently and securely (Sullivan, 2010). Part of this discussion has been recognition that a lack of confidence may induce consumers to switch to less efficient forms of payments, such as paper payments (Cheney et al., 2012). There is some evidence that identity theft and assessments of payment security affect consumer payment choice. In particular, one paper has been able to measure the effect of an identity theft incident on consumer adoption and use of particular types of payment (Kahn and Liñares-Zegarra, 2013). Another examines how consumers' assessments of payment method security relate to consumers' actual payment behavior (Stavins, 2013). A third study argues that news about card fraud reduces consumer willingness to use debit cards in regular transactions (Kosse, 2013).

[9] 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 Federal Trade Commission (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) on 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).

Other papers have considered the trade-off between information security and data privacy that may have implications for identity theft prevention (Acquisti, 2004; Anderson and Moore, 2007). Other work shows that incentives for consumers to prevent some forms of identity theft (such as existing account fraud) are limited by regulatory liability limits and business practices, whereas consumer incentives to prevent new account identity theft are stronger (FTC, 2003; Cheney, 2003).

The frequency of extended fraud alerts is examined in a U.S. Government Accountability Office (GAO) report (2002), which suggests that extended fraud alerts may reliably proxy for incidence of identity theft. The popular press argues that these alerts do not affect risk scores (Saranow Schultz, 2010). Our study contributes to the literature on identity theft by examining how extended fraud alerts, which represent likely instances of identity theft, affect consumers' cost of credit, access to credit, and repayment behavior.

There is a large and growing literature showing that individuals in a wide variety of contexts pay limited attention 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.[10] 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. We contribute to the literature that discusses limited attention by showing that a consumer's behavior after a likely instance of identity theft is consistent with an increase in the salience of consumer personal financial information and increased consumer attention to credit information after the identity theft incident.

III. Institutional Details, Data, and Methodology

Identity theft and fraud may manifest in many forms. For example, criminals may use stolen credit or debit card account information to acquire goods and services fraudulently. If they succeed in stealing the personal identification number (PIN) associated with a debit card, they may be able to withdraw money from a victim's checking account. If criminals obtain additional

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

information (such as addresses, names, dates of birth, Social Security numbers (SSNs), or financial account passwords), they may be able to open fraudulent new credit accounts or take over consumers' existing accounts. Criminals use a variety of complex strategies to obtain this information. In some instances, they deceive their victims by sending "spear phishing" e-mails that include links to download viruses or by installing keystroke logging software onto a consumer's personal computer. Criminals also engage in database intrusion (data breaches), mail interception, false address changes, and dumpster diving.[11]

A. Extended Fraud Alerts

Although a variety of mechanisms have evolved to prevent identity theft and to protect victims from the consequences of ensuing fraudulent activity, this paper focuses on the study of extended fraud alerts.[12] Under the Fair and Accurate Credit Transactions Act (FACTA) of 2003, consumers have a right to place an extended fraud alert on their file at each of the three national credit reporting agencies and to receive a free copy of their credit report from each agency.[13] An extended fraud alert remains in the credit file for seven years unless the consumer chooses to remove it earlier. 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

[11] Cheney (2005) describes some of the tactics criminals use to perpetrate different forms of identity fraud.

[12] These mechanisms include an initial fraud alert, an extended fraud alert, and a credit freeze. 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. 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 remove it permanently. For more information about these other protection mechanisms, see Cheney et al. (2014).

[13] There are three major credit bureaus in the U.S.: Equifax, Experian, and TransUnion. When the consumer files an alert with one credit bureau, this information is communicated to the other two. Although FACTA provides all consumers with the right to receive a free credit report annually from each of the national credit reporting agencies, it further provides rights to request additional free reports as part of filing an alert.

method specified in the alert to verify the consumer's identity in the case of any of the above applications.[14]

Particularly relevant to our analysis, an extended fraud alert has characteristics that imply that most filers of these alerts have been actual victims of fraud or identity theft.[15] For example, extended alert filers must submit an Identity Theft Report in order to place the alert in their credit bureau files. Consumers face criminal penalties for falsifying information in these reports.[16]

To summarize, an extended fraud alert is a complex mechanism — initiated with considerable effort by the consumer — with several implications: 1) the release of credit file information to the consumer, 2) exclusion from prescreened credit and insurance offers for five years, 3) additional verification of identity by lenders at the time of the credit extension for seven years, and 4) the psychological shock from the fraud itself. All these activities increase a consumer's attention to personal finances and credit file information. It is for this reason that we treat an extended fraud alert as a shock to credit file salience and to the consumer's attention in order to study how creditworthiness and credit activity are affected by inattention on the part of the consumer.

B. Data

In order to explore the effect of extended fraud alerts 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

[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.

[15] In comparison, we surmise that a larger share of consumers who file an initial fraud alert could be doing so out of precaution — perhaps because they learned of a compromise at a business they frequent or they received a letter informing them their records may have been accessed — but they have not yet necessarily experienced an event of fraud, account takeover, etc. We choose to be conservative about potential false positives, but we recognize that this implies additional false negatives. As mentioned previously, Cheney et al. (2014) provide additional evidence that extended alerts are more likely to capture realized fraud incidents.

[16] 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."

anonymized 5 percent random sample of variables contained in the credit bureau records of U.S. consumers.[17] 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 SSN as the method of randomly selecting the sample.[18]

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 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 was 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.[19] 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.[20] We use the quarter in which an extended fraud alert first appears in the consumer's file to define "cohorts" of consumers likely affected by identity theft. In much of the following analysis, we examine changes in variables in *event time* — so many quarters before or after an extended fraud alert first appears.

We typically measure average effects for each cohort of consumers who file a first extended fraud alert rather than taking an average of all treated consumers. We do this for several reasons. First, conducting our analysis on a cohort-by-cohort basis allows us to compare individuals who experience the same seasonal and business cycle effects that may also influence the outcome variables we study. Second, this approach allows us to observe the patterns that are

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

[18] 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.

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

[20] 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.

common across the cohorts while also observing variations in timing, amplitude, and persistence across the cohorts. An important source of the fraud that consumers experience comes from data breaches, but we do not wish to assume that all data breaches have the same consequences.[21] Over time, more or less sensitive data are stolen from different organizations serving different populations of consumers. Thus, we want to allow for some variation in treatment and selection, especially since the specific circumstances of fraud change over time.[22]

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 dispute fraudulent accounts or delinquencies on compromised accounts as well as any other errors in their credit reports. This has important implications for measurement and interpretation in this paper.

The process of "cleaning" a credit report means that evidence of fraud is often removed from the data around the date, or shortly thereafter, at which a fraud alert is originally filed. We do not see which variables are cleaned or for what reasons. However, the manner in which each quarter of the CCP data is assembled implies that much of the evidence of any fraud in the *preceding* quarters 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. It is important to emphasize that this characteristic of our data does not necessarily apply to the actual credit report information consumers and creditors access every day. When an error is discovered in information contained in those credit bureau files, the erroneous information is removed from the entire history contained in those files even if the error was discovered long after it first appeared.[23]

There are two variables in our data that are especially noteworthy for our analysis. The first variable is inquiries — applications for credit or insurance made by the consumer — and the

[21] See, for example, the discussion of the South Carolina Department of Revenue data breach in Section III of Cheney et al. (2014).

[22] By selection, we mean primarily the decisions of the criminals stealing and exploiting sensitive data.

[23] This distinction between the dynamics of credit bureau data used by consumers and market participants as opposed to the construction of the CCP arises primarily because the panel was initially constructed from data archives, with new quarters of data added sequentially thereafter.

second is information about the consumer's address.[24] We also study the behavior of such variables as risk score, the number of bankcards with positive balances, the percent of credit cards in good standing, and the number of third-party collections.[25]

Table 1 presents the descriptive statistics for our data set. The table's first three columns provide summary statistics for the data set, including the entire population, the prime population, and the subprime population.[26] The last three columns summarize statistics for extended fraud alert filers at the time of the filing, including statistics for all filers, prime filers, and subprime filers. 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 696, while it is 654 for the entire extended alert population. The number of inquiries in the past three months is 0.55 for the entire population compared with 1.55 for the extended alert population. 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.

Another observation from Table 1 is that there are more subprime extended alert filers than prime filers. In comparison, there are about 1.8 prime consumers for every subprime consumer in our general population. In part, this observation led us to investigate outcomes for prime consumers separately from subprime consumers (see Section IV).

C. Methodology: Propensity Score Matching

In order to account for possible differences between identity theft victims and the general credit bureau population, as well as to control for possible selection into the treatment group, we use propensity score matching. We follow the standard model of propensity score matching, as described in Heckman, Ichimura, and Todd (1997, 1998) and many other studies. A recent study

[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] The risk score contained in the CCP is the Equifax Risk Score. Accounts in good standing are defined as those that are paid as agreed, without any delinquency. Third-party collections consist of accounts in collection and derogatory public records (e.g., judgments).

[26] Prime consumers are defined as those with risk scores > 660 four quarters before the alert, while subprime borrowers are those with risk scores ≤ 660 four quarters before the alert. We use risk score lagged four quarters to avoid information possibly contaminated by fraud in our definition of the prime and subprime groups.

in consumer finance that relies on propensity score matching to construct matched control and treatment groups is Agarwal and Qian (2013). As the first step in the propensity score matching model, we estimate the following model:

$$\Pr(D = 1) = F(X\beta), \quad (1)$$

where D is the indicator variable for a first extended fraud alert ($D = 1$), X is a vector of consumer characteristics that may influence selection into the treatment group, and β is a vector of coefficients.

We estimate equation (1) on a cohort-by-cohort basis using a Probit regression model. The vector of selection variables includes four-quarter lags of the following variables: an indicator for the presence of a first mortgage; number of inquiries within three months; number of inquiries within 12 months; borrower age and age squared; age of newest account; number of accounts with positive balances; risk score; number of occurrences of 30 days past due on credit cards within 24 months; age of the oldest credit card; indicators for credit card utilization between 25 and 50 percent, between 50 and 75 percent, between 75 and 100 percent, and over 100 percent; and an indicator for a change in address on file. Cheney et al. (2014) also provide evidence that these variables help to explain the prevalence of extended fraud alerts. We use four-period lags of these variables in order to avoid the possibility that the data are already affected by fraud (e.g., inquiries in the quarter preceding the alert) in modeling selection into the treatment group.[27] Four-period lags also allow us to avoid both using the same covariates to model selection into the treatment group and including these variables later among our outcome variables of interest.

Estimated coefficients from equation (1) are used to predict the probability of treatment (having a first extended alert) for each individual in a given cohort, including treated individuals and individuals in the control population (consumers without any alerts during our sample period). After that, individuals from the treatment and control populations are matched based on their predicted probability of treatment, also referred to as a propensity score. We use one-to-one nearest neighbor matching without replacement. In simple terms, for each consumer with a first

[27] Our choice of lag length is motivated by patterns in inquiries and address changes that we illustrate in Figures 1 and 3. These are discussed in Section IV.A. In Section VI, we verify that our results are robust to using even longer lags of the variables.

extended fraud alert, we identify a consumer in the control population who appears most similar in terms of the observable characteristics in the treatment selection model in equation (1).

In order to reduce the computational burden, we take a 10 percent random sample of all individuals in our control population and use them in the propensity score matching. Even with this reduced sample, we have about 1,000 untreated individuals for every fraud alert filer before matching.[28] We select individuals from the control population only once within each cohort (i.e., without replacement), so if they are matched to someone from the treatment group, we do not match them to anyone else in the same cohort. As can be seen from Table 2, which compares observable consumer attributes before and after the matching for a representative cohort (June 2009 extended fraud alert filers), our propensity score matching technique is effective in finding individuals in the population without any alerts who are comparable (in terms of observables) with the extended fraud alert group. In particular, at the time of the matching, there is no statistically significant difference between the treated and matched control groups in terms of risk score, number of inquiries, and other variables used in the matching.

As is standard in propensity score matching, after we obtain the treatment and matched control groups, we compare the average outcomes for the two sets of consumers. The effect of the extended fraud alert on consumers is calculated as the difference between the average outcomes of the individuals with extended fraud alerts and the average outcomes of the matched group of individuals without a fraud alert at any time during our study period. We calculate differences between the averages at the time of the treatment, as well as before and after the alert. We examine risk scores, number of inquiries within the past three months, the percent of bankcards in good standing within three months, the number of cards with positive balances, and the number of third-party collections within 12 months as our outcome measures. We calculate the effects for each cohort as far back as four quarters before the first extended alert, at the time of the alert (denoted time 0), and in as many quarters after the alert as our data allow. We focus especially on risk score because this variable captures the general financial health of consumers. We examine changes in the other variables because they may help to explain the drivers of

[28] To verify that sampling potential controls does not affect our results, we reran our tests using three different samples from the whole control population. All our results and conclusions are not sensitive to such changes in sampling.

changes in the risk score and they are good potential indicators of fraud (e.g., criminals' actions leading to address changes, inquiries, accounts past due, and changes in the number of cards).

IV. Results

In this section, we present evidence consistent with identity theft occurring just before the placement of an extended alert. Next, we discuss the effect of an extended fraud alert on such outcomes as risk score, inquiries, percent of cards in good standing, the number of cards with positive balances, and the number of third-party collections. These effects are examined using the propensity score matching technique described in the previous section.

A. Evidence of Fraud

Figure 1 plots the average number of inquiries for each cohort in event time, with time 0 equal to the quarter of the first extended fraud alert. As with all of our figures, there are many patterns that are common across cohorts, with some variations in levels or the amplitude of changes. The decline in inquiries observed for very long lags of the most recent cohorts (the leftmost portion of the figure) is an artifact of the financial crisis and recession, when both supply (prescreened solicitations) and demand (mortgage applications) fell. These patterns illustrate the importance of separating seasonal and business cycle effects from the effects we seek to measure.

The most important observation about 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, or the quarter just before. This is consistent with 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 also coincides with a transitory decline in risk score documented in Figure 2. Note that the average increase in score that follows is typically larger than the transitory decline (we revisit this point in the next section). As in Figure 1, a number of patterns are common across the cohorts, but it is clear there is considerable variation in the average risk scores of consumers in each cohort. It is also clear from Figure 2 that there is a long-

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

run increasing trend in average risk scores that is likely due to the business cycle. Both of these observations motivate our use of propensity score matching in the analyses presented in Section IV.B. and thereafter. That approach accounts for the variation in the characteristics of consumers filing extended fraud alerts at different times as well as the general trends in credit bureau variables over time evident from the simple averages such as the ones depicted in Figure 2.

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. But we can compare the pattern of address changes (or mobility) at the time an extended fraud alert is filed with patterns prior to and after the event.[31] Figure 3 plots the fraction of fraud alert filers who change their addresses over the quarters in our sample. As we observed with inquiries, there is a sharp increase in mobility just before the fraud alert is filed, followed by a decline to prealert levels in subsequent quarters. The share of persons with an address change coinciding with the extended fraud alert doubles from an average level of around 8 percent to around 16 percent. This pattern is observable for all cohorts of extended fraud alert filers and is not observed in the control population. The increases in inquiries and address changes 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 likely to have occurred at the time of the fraud alert or just before it.

As previously described, subprime consumers are overrepresented among extended alert filers in comparison with the general population. Further, Table 1 reveals that the average risk score of individuals with extended fraud alerts tends to be lower than the average score of the general credit bureau population. Similarly, the average risk score of prime extended alert filers is lower than the average risk score of all prime consumers. In contrast, we found that the average risk score for subprime consumers with extended fraud alerts is higher than the average risk score for the population of all subprime consumers. For these reasons, in the subsequent analysis, we examine prime and subprime consumers separately.

[30] Criminals may do this when taking over existing accounts, or they may apply for new accounts in the name of a consumer but use another address.

[31] Recall that while consumer address changes may be reversed after the discovery of fraud, the history of address changes in the Consumer Credit Panel is not updated and, therefore, is not affected by cleaning.

Figure 4 plots the average risk score before and after the extended fraud alert for prime borrowers. The average risk score of prime borrowers declines by around 10 points in the period preceding the extended fraud alert but recovers to the prealert level at the time of the alert and remains at the same level in subsequent quarters. This trend is evident for all cohorts. The decline in the risk score before the placement of the extended fraud alert may be the result of identity theft — fraudulent shopping for credit and new fraudulent accounts — not immediately discovered by the consumer.[32] The recovery in risk score after the alert is likely to be the result of the removal of fraudulent accounts or transactions after consumers dispute them.[33]

Figure 5 plots the average risk score for subprime borrowers who file an extended fraud alert. Average risk scores increase by 15 to 40 points in the quarter after the alert is filed. Not all of this increase appears to be attributable to the removal of fraudulent accounts, since any decrease in risk score just prior to the alert is very small. In addition, the increase in scores appears to be persistent. This difference in the patterns for prime and subprime consumers suggests somewhat different mechanisms for these two populations, thereby motivating our analysis in the following sections.

B. Propensity Score Analysis: Effects for Prime and Subprime Consumers

As we explained in the previous section, we use propensity score matching to account for possible selection into the treatment group and to find individuals in the control population who are similar to extended fraud alert filers. In this and the following subsections, we present the results from our analysis using propensity score matching. The dynamic of the difference in risk scores for each cohort of prime borrowers with extended fraud alerts and the relevant group of matched control individuals is summarized in Figure 6. In addition to plotting differences in risk scores as solid lines, we denote the statistical significance of our estimates at the 5 percent confidence level using dots.[34] The average risk score of prime borrowers with extended fraud alerts falls precipitously before the alert and reaches its trough one quarter prior to the alert. The

[32] Also, see Cheney et al. (2014).

[33] The steeper increase in risk scores up to four quarters prior to the placement of the extended fraud alert is largely mechanical because we apply the score cutoff at that point. A similar pattern occurs if we apply shorter or longer lags of the cutoff.

[34] Actual coefficients and *p*-values for these and all other results for propensity score matching are available from the authors.

average risk score recovers in the quarter of the extended fraud alert and, for most cohorts, largely remains statistically indistinguishable from the average risk score of the matched control group after the alert.

Figure 7 shows, for each cohort, the difference between average risk scores of subprime consumers in the treatment group and the subprime consumers in the matched control group, before and after extended fraud alerts are placed. As can be seen in Figure 7, the statistical significance of the estimates varies over time and across cohorts.

Several important results are apparent in Figure 7. First, for all cohorts, the differences between treatment and matched control groups are not statistically significant four quarters prior to the fraud alert. Propensity score matching therefore allowed us to identify a group of control individuals that appears similar to the group of treated individuals in terms of observable characteristics such as risk scores. Second, risk scores increase following the extended fraud alert for all cohorts considered. The increase in average score is between 9 and 33 points and is statistically different from zero. Third, these gains in risk score persist over time: Most cohorts maintain scores that are at least 10 points higher, on average, compared with matched persons who have never had an extended fraud alert. For some cohorts, the average increase in score is statistically different from zero even two to three years after the fraud alert. All these findings are consistent with our earlier descriptive results.

The results from propensity score matching indicate that the effects we observe are not likely to be due to advantageous selection or simple mean reversion, which may affect subprime borrowers. The persistence in effects is consistent with the hypothesis of behavioral changes on the part of the consumers who dedicate more attention to their finances and credit records after filing an extended fraud alert.

At the same time, Figure 7 also shows that there is no significant difference in the risk scores of subprime consumers with extended fraud alerts relative to matched control individuals immediately before the alert was filed. By construction, there should not be a difference four quarters prior to the alert, but it is somewhat surprising that we do not see a difference in the quarter immediately preceding the alert. An unlikely explanation is that the actual fraud occurred long before. This is unlikely given there is no clear declining trend in the average risk score of subprime consumers even eight or 12 quarters before the extended fraud alert (see Figure 5).

While some consumers may not notice fraud for a long time, it seems unlikely the majority of consumers would fail to notice it for three or more years, especially if the result of the fraud involved collections activity (see Section IV.F.).

As a robustness check, we also examine the average risk scores of subprime consumers who have no inquiries two quarters before the alert but place an extended fraud alert at some point in our sample period. These consumers may be regarded as less attentive to their credit record because they do not appear to be actively shopping for credit. We use propensity score matching to find consumers in the control sample who are comparable with the above-described consumers. Similar to Figure 7, there is no statistical difference between the average risk scores of the treated and matched control groups before the extended fraud alert for most of the cohorts considered. However, the average risk score of fraud alert filers substantially increases at the time of the fraud alert and remains elevated in the following quarters. By comparing coefficients for the potentially less attentive sample with the entire group of subprime extended fraud alert filers (those in Figure 7), we can conclude that the effect of the extended fraud alert for the less attentive consumers is somewhat larger. This is consistent with the drivers behind our results being inattention on the part of consumers before identity theft and increased salience of credit file information after fraud alerts.

C. Inquiries

Figure 8 shows, for each cohort, the differences in inquiries between subprime extended fraud alert filers and a matched group of control individuals, before and after the alert. The pattern is similar to our earlier results. First, the number of inquiries within three months tends to increase in the quarters preceding the extended fraud alert for all cohorts considered. Second, there is a peak in inquiries in the quarter of the fraud alert. This is consistent with attempts to open new fraudulent accounts in the consumer's name, at least for some of the defrauded individuals. Third, the number of inquiries remains elevated for extended fraud alert filers even after the fraud alert. Part of this increase in inquiries may represent increased shopping for credit by these consumers following the extended fraud alert and the cleanup of their credit files. This

may be the consumer's response to the improvement in risk scores illustrated in Figure 5.[35] Better scores may allow these consumers to obtain more credit, and on better terms, than they could before. Thus, at least some additional inquiries after the filing of the extended fraud alert may simply reflect legitimate shopping behavior by extended alert filers, though we cannot rule out additional fraudulent inquiries.

As revealed in Figure 9, the treatment and matched control groups of prime consumers follow a pattern similar to subprime borrowers prior to the filing of the extended fraud alert but not after. As with subprime consumers, inquiries of prime extended fraud alert filers begin to increase two to three quarters before the fraud alert, and they peak at the time of alert filing. By a quarter after the alert, inquiries of treated prime consumers decline to a level similar to the control group of prime borrowers. By contrast, the inquiries of treated subprime consumers remain elevated relative to their matched control group for as long as four years after a fraud alert. These results confirm our earlier findings that the effect of extended fraud alerts on prime consumers is short lived and that it primarily consists of a decline in risk score and an increase in inquiries before the alert, which quickly revert to their prior levels after the alert. The effects for subprime consumers are clearly different.

D. Share of Bankcards in Good Standing

Figure 10 demonstrates that, among subprime consumers, the share of bankcards in good standing within the preceding three months compared with all bankcards within the preceding three months increases substantially at the time of extended fraud alert filing. For all cohorts but one, there is no statistically significant difference between fraud alert filers and matched control individuals at the time of the matching (four quarters before alert filing), but the percent of cards in good standing increases by 5 to 16 percent for alert filers at the time of the alert. This increase

[35] Depending on the algorithm for a particular risk score, more credit applications may reduce one's score. Given that the treated consumers appear to be applying for more credit, this may attenuate somewhat the increases in risk score we observe in Figures 5 and 7. However, the effect of additional credit applications on risk score may be relatively small. For example, Avery et al. (2004) examine the effect of combining multiple inquiries into a single inquiry during credit file cleanup. They find that this action tends to change risk scores in their data by less than 2 points.

is consistent with the closure of fraudulent and unused delinquent accounts by fraud victims.[36] Figure 10 also shows that fraud alert filers maintain a higher proportion of credit card accounts in good standing, compared with their matched peers, for the duration of our sample (15 and more quarters, in some cases). This result may suggest that the increased salience of credit file information to subprime consumers at the time they file an extended fraud alert may lead to persistent improvements in their repayment behavior.

E. Cards with Positive Balances

As shown in Figure 11, the number of cards with positive balances in the preceding three months is not statistically different between the subprime extended fraud alert filers and matched control individuals four quarters before the alert. However, at the time of the alert filing, likely fraud victims reduce the number of cards with balances by around 0.20 to 0.35 more relative to the sample of matched controls. This difference is statistically significant for most of the cohorts, but the effect is less persistent than we observed for other variables, essentially disappearing after five quarters. This pattern is consistent with our earlier findings that alert filers close unused or fraudulent card accounts at the time of alert filing but later use their improved risk scores to access additional credit.

F. Third-Party Collections

Finally, as can be seen in Figure 12, the number of third-party collections also declines at the time of the extended fraud alert for the subprime treatment group compared with the matched control group. This is consistent with the closure of fraudulent accounts in collections concurrent with the alert filing. For the majority of the cohorts we study, there is no statistical difference between the treatment and matched control groups in the number of third-party collections soon after the placement of an extended fraud alert.[37]

[36] In addition to the proportion of bankcards in good standing, we examined the number of cards past due and found that the number of cards past due declined at the time of fraud alert filing.

[37] The number of third-party collections is not used in the propensity score matching. Hence, this variable is not aligned in the treatment and control groups four quarters before extended fraud alert placement.

G. Summary of Results

Taken together, these findings suggest that subprime consumers dispute errors in their accounts, including removing fraudulent bankcard accounts after the placement of an extended fraud alert. It is unlikely that all accounts removed at the time of the alert are due to fraud. A portion of cards for extended fraud alert filers are placed in third-party collections prior to the fraud alert, so consumers are probably aware of at least some of these trade lines even before the fraud. Our charts therefore imply that some of the trade lines that were disputed, removed, or closed by consumers after the alert are accounts that may not have been monitored or repaid carefully by borrowers before the identity theft incident. Thus, the extended fraud alert may allow subprime consumers to become better informed and take more responsibility over their personal finances.

For subprime consumers, the dynamics of risk scores, the share of bankcards in good standing, and the number of accounts in collections after the extended fraud alert are consistent with the hypothesis that these borrowers improve their repayment behavior and credit management after the alert. Extended fraud alert filers increase the share of bankcards in good standing and reduce the number of third-party collections. These customers maintain a similar number of cards with positive balances when compared with the matched control group. Better credit management increases creditworthiness, as reflected in improved risk scores, which persist at the higher level even several years after the extended fraud alert. Therefore, the information shock from the identity theft and subsequent extended fraud alert appears to help subprime consumers to focus attention on their credit records and personal finances and manage them more efficiently.

V. Extensions

A. The Expiration of Exclusion from Prescreened Credit and Insurance Offers

As described previously, the placement of an extended alert in an individual's credit file excludes the consumer from prescreened credit and insurance offers for five years. However, the alert itself remains in the credit file for seven years, unless removed by the consumer. Therefore, we can separate the effect of prescreened offers from the effect of extended fraud alert expirations by looking at extended fraud alert filers with alerts that have aged five years. In this

subsection, we explore whether the expiration of the no-prescreened credit and insurance offers affects consumer borrowing and repayment. To put it another way, we examine whether extended alert filers' exclusion from prescreened credit and insurance offers contributes to the risk score improvements that we document in the previous section.

Figure 13 illustrates average risk scores for 23 cohorts of consumers who have had an extended fraud alert for at least five years. We group consumers into cohorts based on the age of their extended alerts. Each cohort includes individuals with exactly five years with an extended fraud alert. The resulting cohorts are displayed in such a way that the expiry of the exclusion from prescreened credit offers corresponds to quarter 20 on the graph (five years after placement of an extended fraud alert).[38]

Figure 13 shows little change in average risk scores once prescreened credit and insurance offers are permitted to reach extended fraud alert filers. Risk scores continue to slowly increase over time for these consumers. Their risk scores are relatively high at the time of the expiration and, on average, are above 660 for all cohorts.[39] It appears that the exclusion from prescreened credit and insurance offers does not, by itself, affect the risk scores of consumers with extended fraud alerts. If it does, the effect is a permanent one and does not disappear after this restriction is removed.

B. The Expiration of the Extended Fraud Alert

In this subsection, we explore how the expiration of an extended fraud alert changes consumer credit outcomes. As we discussed in Section II, an extended fraud alert requires lenders to take extra steps to verify the identity of a consumer before extending credit. The protection was designed to prevent criminals from using the consumer's personal information to open fraudulent accounts. In theory, this mechanism, in itself, may also impede the procurement

[38] In Figure 13, we focus on the expiration of the exclusion from credit and insurance offers (henceforth, "the expiration"), which is 20 quarters after an extended fraud alert is placed in a credit bureau file. We identify cohorts by the quarter in which extended alerts are placed (t=0) and follow these cohorts around the expiration (t=20) to the extent that our data set allows (our observation period is Q1:2008–Q3:2013). So, for example, the Dec-03 cohort is defined as individuals with extended alerts placed in Dec-03 with credit and insurance offers expiring 20 quarters later in Dec-08. For this cohort, our data set allows us to observe credit files for the three quarters preceding the expiration (Q1:2008–Q3:2008), the time of the expiration (Q4:2008), and 19 quarters after that (Q1:2009–Q3:2013).

[39] We also calculated changes in inquiries, 30 days past due occurrences, third-party collections, the number of cards with past due amounts > 0, and cards with positive balances after expiration for extended alert filers. There is no apparent response in any of these variables to the renewed exposure to prescreened offers.

of legitimate credit because of the extra effort required for application. The expiration of an extended fraud alert after seven years disables this protection; in this subsection, we examine whether consumers or fraudsters (to the extent that we can observe in our data) react to this event. In other words, we attempt to separate the effect of more difficult access to credit induced by extended fraud alerts from the effect of credit information release and correction of the credit bureau files analyzed in the previous section.

Our findings are depicted in Figure 14. As with the previous graphs, in this graph we split all individuals into cohorts based on the quarter in which their extended fraud alert is placed. For example, the Sep-01 cohort includes all consumers whose extended alerts are filed in Q3:2001 and expire in Q3:2008. Figure 14 demonstrates that there is little response to the expiration of the alert. Risk scores remain at the same level or increase gradually after expiration. This result suggests that the changes in consumer behavior that we observed after alert placement are not simply due to the presence of the extended fraud alert in the credit record and/or due to more difficult access to credit.[40] Combined, these findings suggest that the changes in consumer outcomes after extended fraud alert placement are persistent and can be explained, at least for some consumers, by an increase in attention to credit management on the part of the consumer.

C. Address Change as a Measure of Fraud Intensity

In this subsection, we explore whether the severity of identity theft, as measured by the intensity of extended fraud alert filers' address changes, affects risk scores, number of inquiries, and number of cards with past due amounts. Address change, or mobility, may capture both valid address changes as well as illegal activity consisting of fraudulent accounts opened in the consumer's name and direction of mail correspondence to an address other than the consumer's. As Figure 3 showed for all 19 cohorts, we observe an increase in mobility for extended fraud alert filers that coincides with the alert quarter. This increase in mobility at the time of the extended fraud alert is present for both subprime and prime consumers.

[40] Graphs for the number of cards past due, the number of third-party collections, and other outcome variables discussed previously are not included for space considerations. The only measure that does increase for a few cohorts after expiration is the number of inquiries within three months. No other variables change discontinuously after alert expiration.

Because address changes in the quarter of the extended fraud alert may capture the severity of fraud (i.e., whether criminals succeeded in opening fraudulent accounts and directing mail correspondence), we consider extended fraud alert filers with and without an address change separately. We hypothesize that immobile consumers are less likely to be subject to severe fraud, while mobile consumers might experience more serious fraud incidents. Risk scores of fraud victims without an address change tend to decline somewhat during the quarter preceding the alert and dramatically increase after the alert is placed. The decline in risk scores is considerably larger for fraud victims with an address change than for alert filers without an address change. This stronger decline in risk scores is then followed by a considerably larger increase in the alert quarter. Consistent with our earlier results, the number of inquiries within three months increases, and the number of credit cards past due declines in the quarter of the extended fraud alert both for mobile and for immobile consumers. Therefore, an address change in the quarter of the extended fraud alert does not appear to affect our results on inquiries or cards past due qualitatively but is associated with more dramatic declines and subsequent increases in average risk scores for consumers with both an extended fraud alert and a change in address on file.[41]

VI. Robustness Checks

A. Changing the Definition of Subprime Consumers

In order to examine the robustness of our findings to the definition of subprime consumers, we changed the risk score cutoff used to define this group. In our first exercise, we use a risk score cutoff of 620 four quarters before the extended fraud alert to differentiate subprime from prime consumers. This change in the cutoff has no bearing on our findings. It is also possible that the same risk score is associated with different default probabilities for different risk score distributions (e.g., in different time periods). For instance, while an individual may have a risk score of 660 in all time periods, lenders may adjust their underwriting thresholds during an economic recovery. To take into account such a scenario, we define subprime consumers as those with risk scores below the 25th percentile of the risk score distribution four

[41] These figures are available upon request.

quarters before the extended fraud alert. There is some variation in the 25th percentile — it ranges from 611 to 628 in our sample. Even when we use this variable cutoff based on the 25th percentile rather than a fixed value to define subprime consumers, the patterns in risk scores and other variables we found in our main analysis are robust.[42]

B. Matching in Different Time Periods

While our results presented so far indicate that fraud was unlikely to occur more than four quarters before the placement of an extended fraud alert for the majority of consumers in our data set, we explore the possibility of slow and gradual fraud activity in this section. If perpetrators were able to infiltrate credit files of identity theft victims more than four quarters before the extended fraud alert, our matching strategy would not be able to account for that. We conduct two exercises to rule out this possibility. First, we match individuals with and without extended fraud alerts based on their credit bureau characteristics eight quarters before the extended fraud alert. Eight quarters are likely far enough removed from the extended fraud alert for fraudulent credit accounts to go into collection, which would alert the consumer of any fraud. Overall, all our results are robust to matching on eight quarters before the alert.[43]

Second, we match individuals not on credit bureau characteristics four quarters earlier but on future credit bureau characteristics. It is possible to argue that matching on historical data can be biased because we do not know exactly when the identities of fraud victims were compromised. If their identities were compromised slowly over time and the placement of an extended fraud alert reveals only the consumer's discovery of long-term criminal activity, then our matching on past characteristics (no matter how far back) may be inappropriate. In order to account for this possibility, we match individuals on their credit characteristics in quarters after the placement of the extended fraud alert. Our assumption here is that the "true" credit characteristics of consumers are revealed after credit files are cleaned. Therefore, we perform propensity score matching for prime and subprime consumers eight quarters *after* fraud declaration (assuming, implicitly, that eight quarters are sufficient for any transient effects of an extended fraud alert to dissipate).

[42] These tables and figures are available upon request.

[43] These figures are available upon request.

Figure 15 shows the difference in average risk scores between prime consumers with and without extended fraud alerts, with propensity score matching performed eight quarters after the alert placement. While there is no difference between extended fraud alert filers and the matched controls in the matching quarter, there are large and persistent differences before alert placement. This finding indicates that matching on the future is likely not effective in accounting for selection on observable characteristics into the treatment group. However, the pattern of change in the average risk score over time is very similar to our earlier results. First, risk scores decline one quarter prior to extended fraud alert placement, then recover to higher levels in the alert quarter, and remain at a steady level afterward. Figure 16 shows differences in average risk scores between treated and matched control groups for subprime consumers matched on credit characteristics eight quarters after extended fraud alert placement; the patterns are comparable with those described for Figure 15.

VII. Conclusion

This paper uses a unique data set of anonymized credit bureau records of the U.S. population to examine the effects of extended fraud alerts on risk scores, access to credit, and credit portfolios. We isolate the most likely victims of identity theft or fraud as those individuals who place an extended fraud alert in their credit file. This type of fraud alert requires the filing of a police report or a report with a government agency, with accompanying evidence of identity theft. Thus, an extended fraud alert is more likely to be placed because of genuine identity theft, rather than other nonfraud-related factors.

We contribute to the existing literature on identity theft by examining longer-term effects of identity theft on consumer risk scores, inquiries, credit card holdings, balances, delinquencies, and address changes. Our results indicate that the average risk scores of consumers with extended fraud alerts increase after they place such alerts in their credit record. For subprime consumers, this effect is often persistent over time and remains for as long as 18 quarters after the extended fraud alert. We also find that the average number of cards past due and the average number of third-party collections all decline after the alert and remain at lowered levels for several years after the extended fraud alert. On the other hand, among prime consumers the effects of an extended fraud alert on their credit bureau attributes are transient. We do not find

evidence that prime and subprime consumers experience identity theft of different intensity, as measured by the likelihood to experience an address change in the quarter of the extended fraud alert. The improvement in risk scores and other credit file measures is consistent with limited attention to credit file information on the part of subprime consumers before the placement of an extended fraud alert and their increased attention to their credit portfolio after the alert.

Our findings suggest that existing fraud mitigation mechanisms such as extended fraud alerts may have unintended and potentially positive consequences of increasing consumer awareness of financial and credit information. This result may provide support for continuing efforts at improving the financial literacy and education of U.S. consumers. On the other hand, we do not observe a number of other, potentially relevant, outcomes of extended fraud alert filers. Given that the average person's attention is limited, it is possible that dedicating extra attention to credit file monitoring results in less attention to other important topics, which may or may not negatively affect these consumers.

Bibliography

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

Agarwal, Sumit, and Wenlan Qian. 2013. "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore," mimeo, National University of Singapore.

Anderson, Keith, Erik Durbin, and Michael Salinger. 2008. "Identity Theft," *Journal of Economic Perspectives*, Vol. 22, No. 2, pp.171–192.

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

Avery, Robert, Paul Calem, and Glenn Canner. 2004. "Credit Report Accuracy and Access to Credit." *Federal Reserve Bulletin*, Vol. 90, No. 3, pp. 297–322.

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*, Vol. 21, No. 2, pp. 785–818.

Cheney, Julia. 2003. "Identity Theft: A Pernicious and Costly Fraud," Federal Reserve Bank of Philadelphia Payment Cards Center Discussion Paper No. 03-18.

Cheney, Julia. 2005. "Identity Theft: Do Definitions Still Matter?" Federal Reserve Bank of Philadelphia Payment Cards Center Discussion Paper No. 05-10.

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 No. 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 No. 14-04.

DellaVigna, Stefano, and Joshua Pollet. 2009. "Investor Inattention and Friday Earnings Announcements," *Journal of Finance*, Vol. 64, No.2, pp. 709–749.

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.

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.

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

Heckman, James, Hidehiko Ichimura, and Petra Todd. 1997. "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme," *Review of Economic Studies*, Vol. 64, No. 4, pp. 605–654.

Heckman, James, Hidehiko Ichimura, and Petra Todd. 1998. "Matching as an Econometric Evaluation Estimator," *Review of Economic Studies*, Vol. 65, No. 2, pp. 261–294.

Hirshleifer, David, Sonya Lim, and Siew Hong Teoh. 2009. "Driven to Distraction: Extraneous Events and Underreaction to Earnings News," *Journal of Finance*, Vol. 64, No. 5, pp. 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*, Vol. 1, No. 1, pp. 35–73.

Hunt, Robert. 2005. "A Century of Credit Reporting in America," Federal Reserve Bank of Philadelphia Working Paper No. 05-13.

Kahn, Charles, and José Liñares-Zegarra. 2013. "Identity Theft and Consumer Payment Choice: Does Security Really Matter?" mimeo, University of Illinois.

Kosse, Anneke. 2013. "Do Newspaper Articles on Card Fraud Affect Debit Card Usage?" *Journal of Banking and Finance*, Vol. 37, No. 12, pp. 5382–5391.

Lacetera, Nicola, Devin Pope, and Justin Sydnor. 2012. "Heuristic Thinking and Limited Attention in the Car Market," *American Economic Review*, Vol. 102, No. 5, pp. 2206–2236.

Lee, Donghoon, and van der Klaauw, Wilbert. 2010. "An Introduction to the FRBNY Consumer Credit Panel," Federal Reserve Bank of New York Staff Report No. 479.

Saranow Schultz, Jennifer. 2010. "Fraud Alerts Don't Hurt Your Credit Score," Bucks: Making the Most of Your Money (blog). *New York Times*, May 25.

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*, Vol. 47, No. 3, pp. 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*, Vol. 27, No. 4, pp. 990–1030.

Stavins, Joanna. 2013. “Security of Retail Payments: The New Strategic Objective,” Federal Reserve Bank of Boston Public Policy Discussion Paper No. 13-9.

Sullivan, Richard J. 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).

U.S. Government Accountability Office and United States of America. 2002. “Identity Theft: Prevalence and Cost Appear to be Growing.”

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 No. 13-04.

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

Table 1: Descriptive Statistics

	Entire Data Set			Extended Alerts		
Variable	**All**	**Prime**	**Subprime**	**All**	**Prime**	**Subprime**
Risk Score						
Mean	695.71	765.37	570.50	654.40	748.08	582.23
Median	721.00	776.00	581.00	653.00	763.00	586.00
Standard Deviation	108.30	48.85	64.21	109.14	66.71	75.06
Number of Inquiries Past 3 Months						
Mean	0.55	0.43	0.71	1.55	1.22	1.66
Median	0.00	0.00	0.00	1.00	1.00	1.00
Standard Deviation	1.01	0.80	1.22	2.16	1.75	2.15
Number of Inquiries Past 12 Months						
Mean	1.91	1.48	2.49	4.12	2.88	4.55
Median	1.00	1.00	2.00	3.00	2.00	3.00
Standard Deviation	2.27	1.60	2.85	4.36	2.84	4.47
Age of Newest Account (Months)						
Mean	31.14	32.79	28.16	18.77	17.69	21.59
Median	16.00	14.00	19.00	10.00	9.00	12.00
Standard Deviation	46.00	53.13	28.80	24.79	26.47	24.94
Change in Number of Accounts						
Mean	-0.03	-0.05	0.00	-0.96	-0.56	-1.26
Median	0.00	0.00	0.00	-1.00	0.00	-1.00
Standard Deviation	1.02	0.96	1.13	2.57	1.92	2.85
Utilization Rate (Percent)						
Mean	0.62	0.36	1.52	0.40	0.26	0.60
Median	0.15	0.08	0.73	0.28	0.13	0.61
Standard Deviation	350.19	367.59	282.24	1.25	0.33	2.17
Age (Years)						
Mean	50.84	54.49	42.84	44.48	51.48	39.81
Median	49.00	54.00	41.00	42.00	51.00	37.00
Standard Deviation	17.79	17.64	14.06	14.68	15.19	12.16
Number of Bankcard Accounts w/ Update w/in 3 Months w/ Balance >$0						
Mean	1.61	1.47	2.02	1.74	1.74	1.62
Median	1.00	1.00	1.00	1.00	1.00	1.00
Standard Deviation	1.47	1.32	1.76	1.55	1.50	1.48
Number of Trades Currently 30 Days PD						
Mean	0.05	0.00	0.12	0.06	0.03	0.07
Median	0.00	0.00	0.00	0.00	0.00	0.00
Standard Deviation	0.27	0.07	0.43	0.30	0.21	0.32
Number of Bankcard Accounts with PD Amount >0						
Mean	0.19	0.01	0.60	0.23	0.04	0.42
Median	0.00	0.00	0.00	0.00	0.00	0.00
Standard Deviation	0.66	0.11	1.07	0.67	0.32	0.84
Number of Bankcards Currently 30 Days PD						
Mean	0.01	0.00	0.04	0.02	0.01	0.02
Median	0.00	0.00	0.00	0.00	0.00	0.00
Standard Deviation	0.14	0.03	0.23	0.16	0.12	0.16
Total Number of 30 Days PD Occurrences on Bankcards w/in 24 Months						
Mean	0.38	0.05	1.09	0.42	0.10	0.69
Median	0.00	0.00	0.00	0.00	0.00	0.00
Standard Deviation	1.41	0.36	2.31	1.28	0.49	1.65
Total Number of 120 Days PD Occurrences on Bankcards w/in 24 Months						
Mean	0.60	0.03	1.84	0.53	0.05	1.03
Median	0.00	0.00	0.00	0.00	0.00	0.00
Standard Deviation	3.23	0.71	5.48	2.77	0.57	3.99
Total Past Due Amount Bankcard Accounts w/ Update w/in 3 Months ($)						
Mean	340.96	20.11	1254.30	377.30	63.67	790.22
Median	0.00	0.00	0.00	0.00	0.00	0.00
Standard Deviation	2560.78	758.53	4739.17	2954.11	1443.69	4145.66
Total Number of Observations	245,263,237	147,662,502	82,149,486	52,649	17,198	22,009

Notes: Unit of observation is person-quarter. Statistics for the treatment groups are calculated at the time of extended alert placement. Observations with missing values are omitted from calculations. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

Table 2. Average Values of Matching Variables for Matched and Unmatched Fraud Alert Filers and Control Populations, Four Quarters Before Fraud Alert (Time of the Matching), Subprime Consumers, June 2009 Fraud Alert Filers

Variable name	All Treated	All Controls	Matched Treated	Matched Controls	Treated Minus Controls	P-value
Mortgage indicator	0.18	0.21	0.20	0.21	-0.01	0.73
Inquiries within 12 months	5.58	3.14	5.82	5.88	-0.06	0.82
Inquiries within 3 months	1.64	0.87	1.73	1.74	0.00	0.97
Person's age	38.51	40.98	38.93	38.63	0.31	0.54
Age squared	1646.54	1883.78	1679.37	1641.72	37.66	0.40
Age of newest account	15.10	21.90	12.03	12.43	-0.39	0.47
Number of accounts with positive balance	4.82	4.23	5.27	5.13	0.14	0.38
Risk score	544.25	564.64	541.93	540.90	1.03	0.75
Number of 30 days past due	1.24	1.33	1.31	1.30	0.00	0.99
Age of oldest card	91.12	98.24	91.21	87.50	3.71	0.16
Utilization 25 – 50%	0.10	0.07	0.12	0.12	0.00	0.81
Utilization 50 – 75%	0.11	0.11	0.13	0.14	-0.01	0.64
Utilization 75 – 100%	0.15	0.16	0.18	0.16	0.02	0.29
Utilization over 100%	0.49	0.53	0.41	0.42	-0.01	0.49
Address mobility	0.12	0.07	0.14	0.12	0.02	0.19

Notes: We compute average values of all matching variables before and after matching in treatment and control groups. While there are statistically significant differences between treatment and control groups before the matching, these differences become statistically insignificant after the matching (as shown in the last column of the table). We use propensity score matching to account for possible selection in the treatment group. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

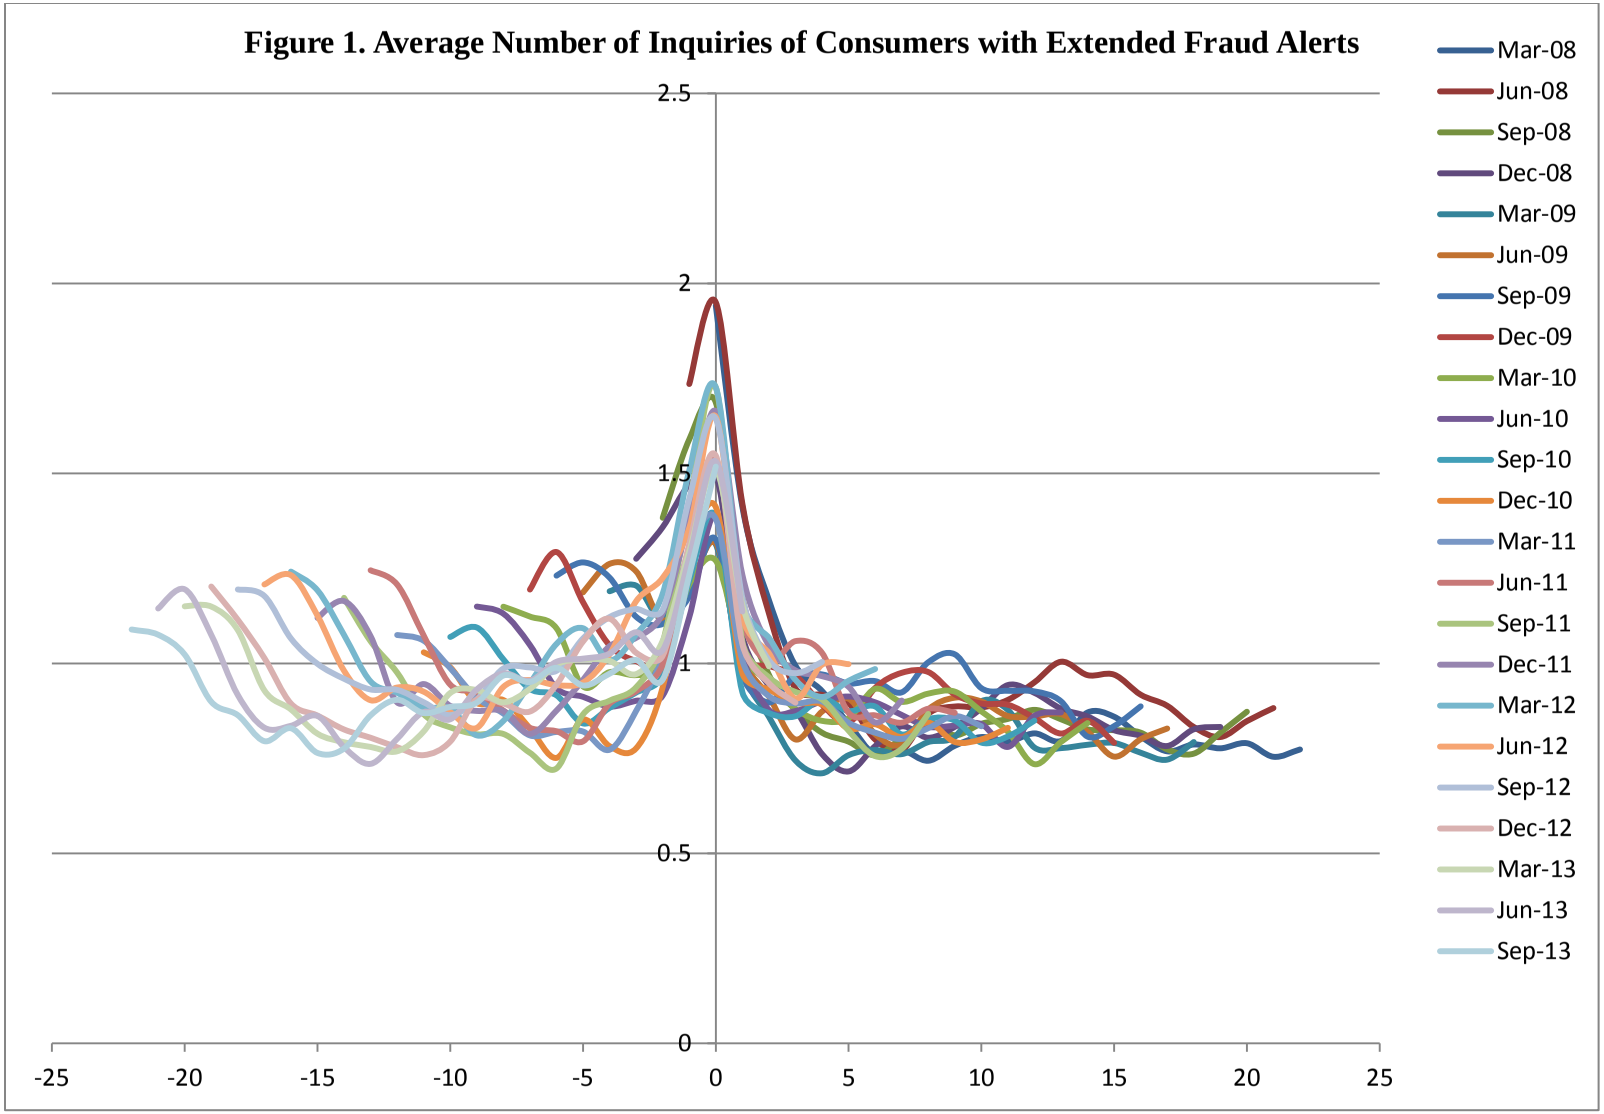

Notes: This figure depicts the average number of inquiries for consumers who filed an extended fraud alert at time 0. Cohorts are defined based on the quarter of fraud alert placement. The average number of inquiries increases at the time of fraud alert filing. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

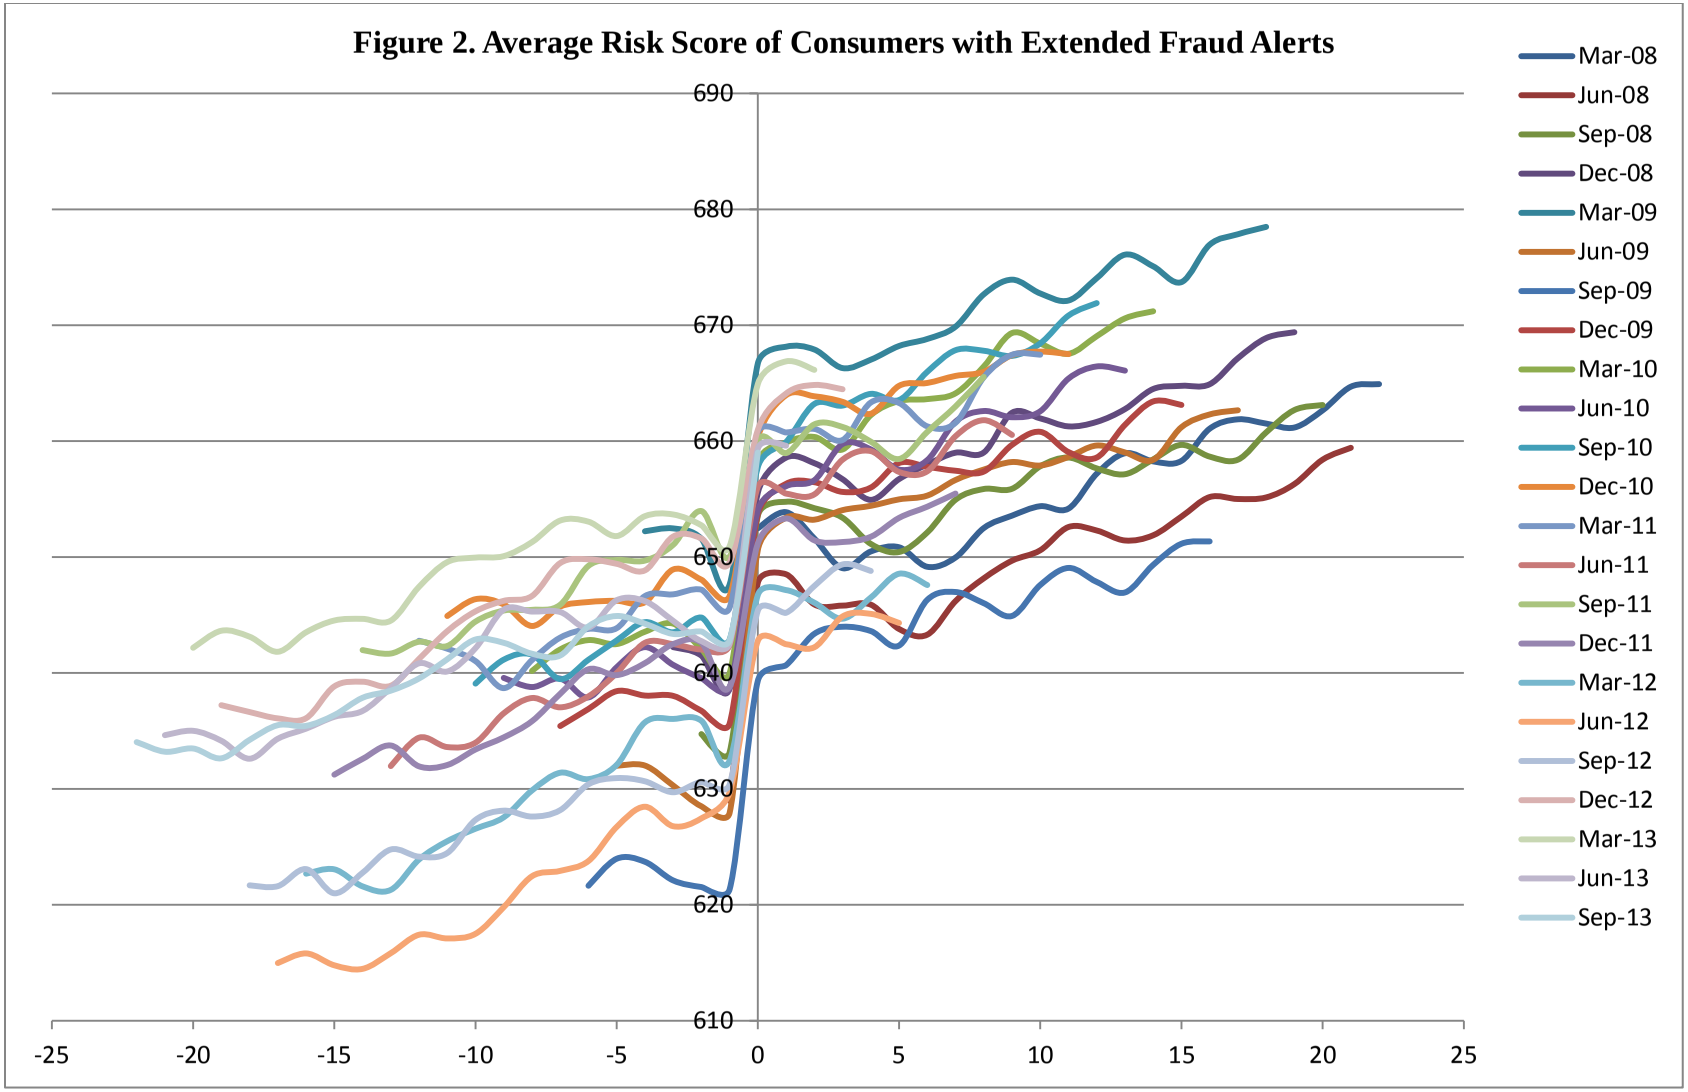

Notes: This figure depicts the average risk score for consumers who filed an extended fraud alert at time 0. Cohorts are defined based on the quarter of fraud alert placement. The average risk score increases at the time of fraud alert filing. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

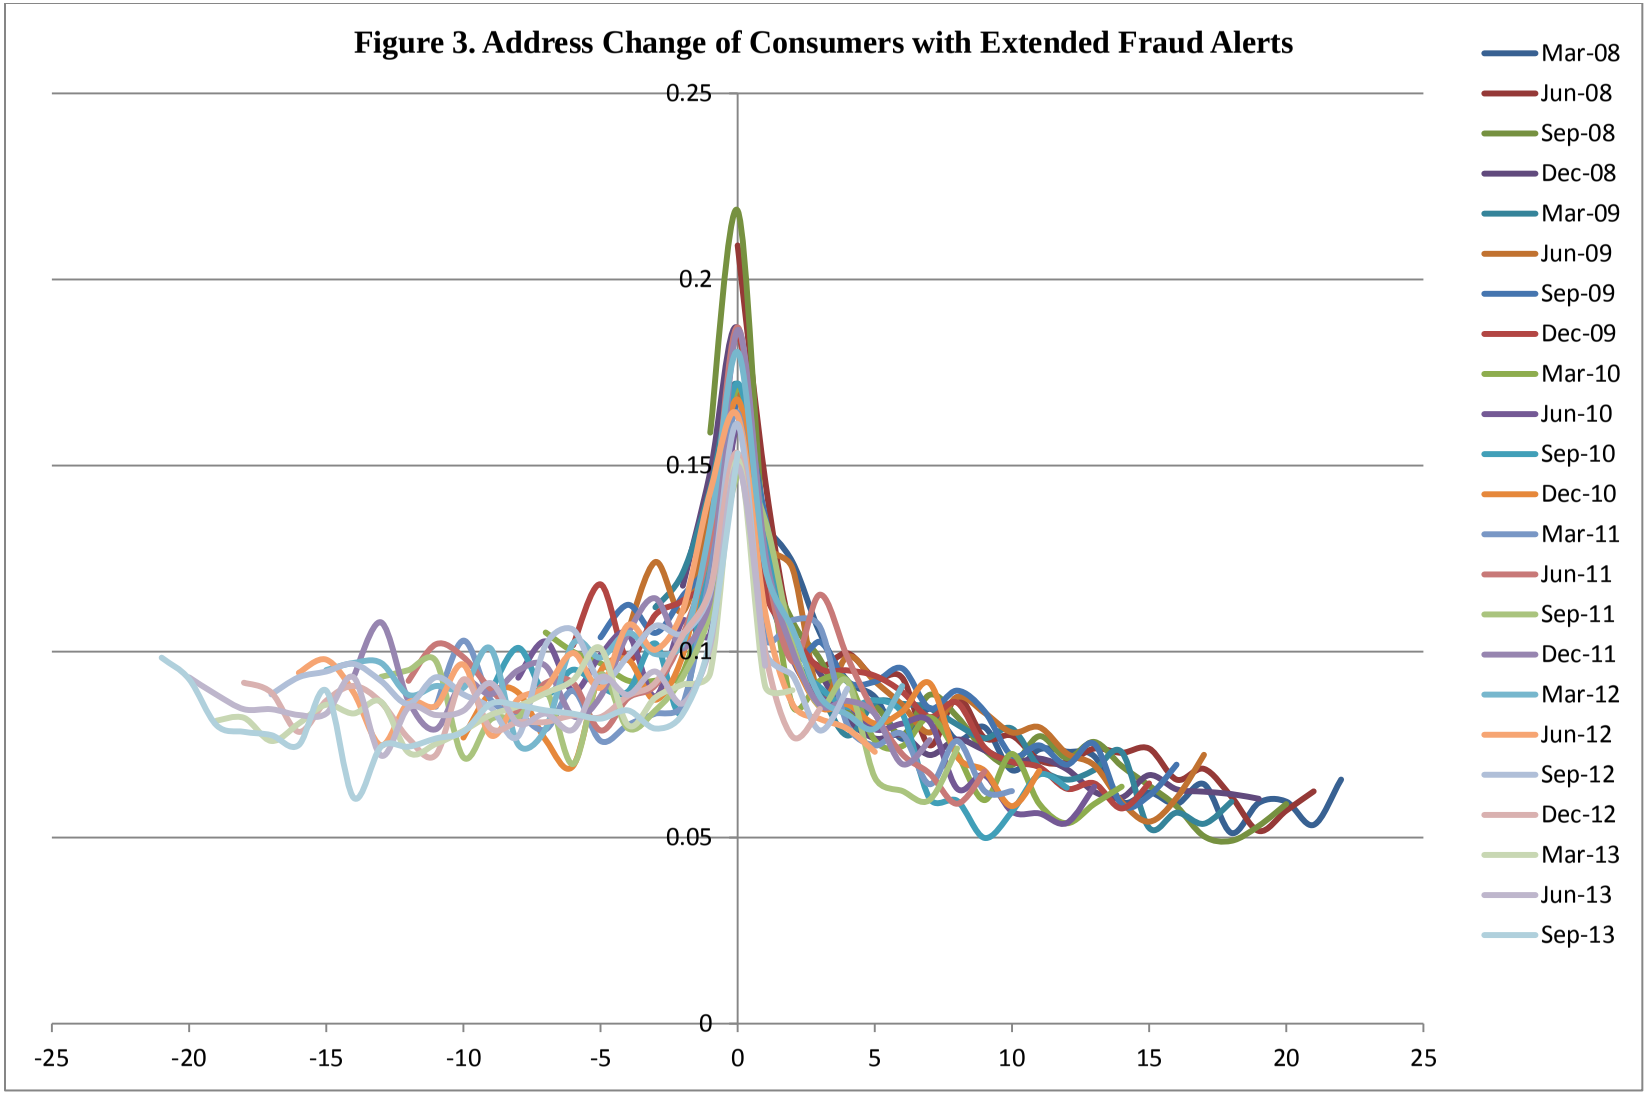

Notes: This figure depicts the average address change for consumers who filed an extended fraud alert at time 0. Cohorts are defined based on the quarter of fraud alert placement. The average number of address changes increases at the time of fraud alert filing. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

Figure 4. Average Risk Score of Prime Consumers with Extended Alerts

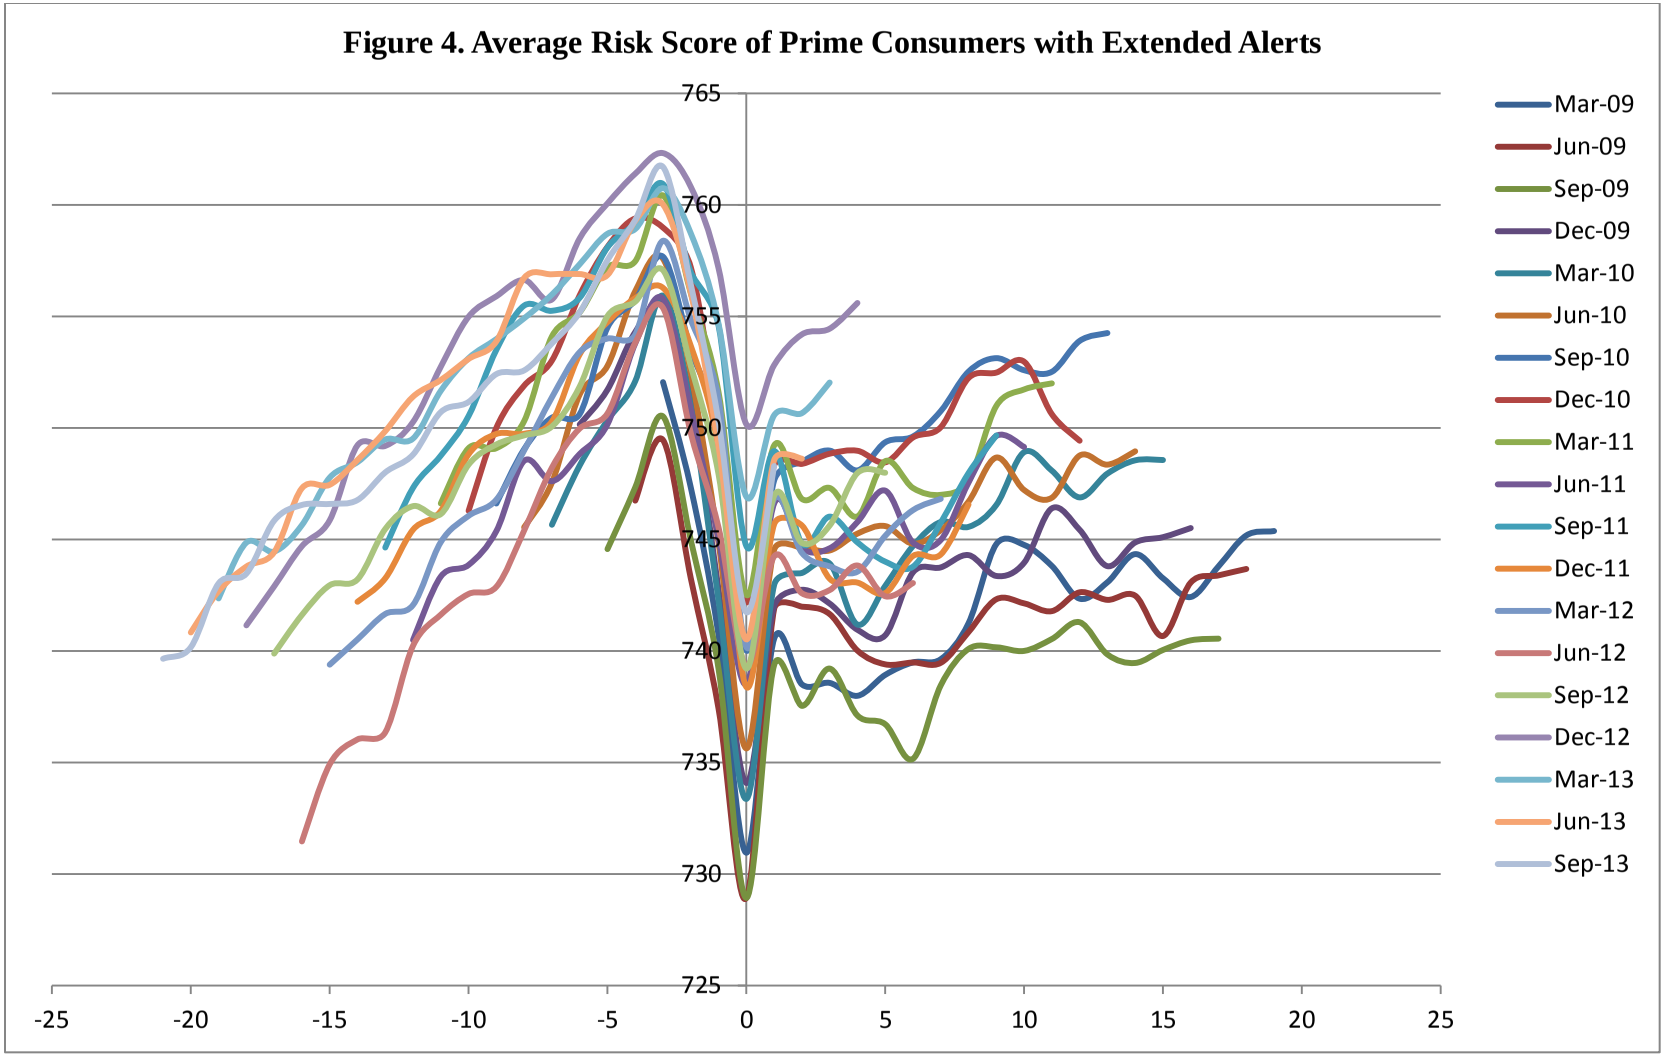

Notes: This figure depicts the average risk score for prime consumers who filed an extended fraud alert at time 0. Cohorts are defined based on the quarter of fraud alert placement. The average risk score increases after fraud alert filing. Prime consumers are defined as those with a risk score >660 four quarters before the alert. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

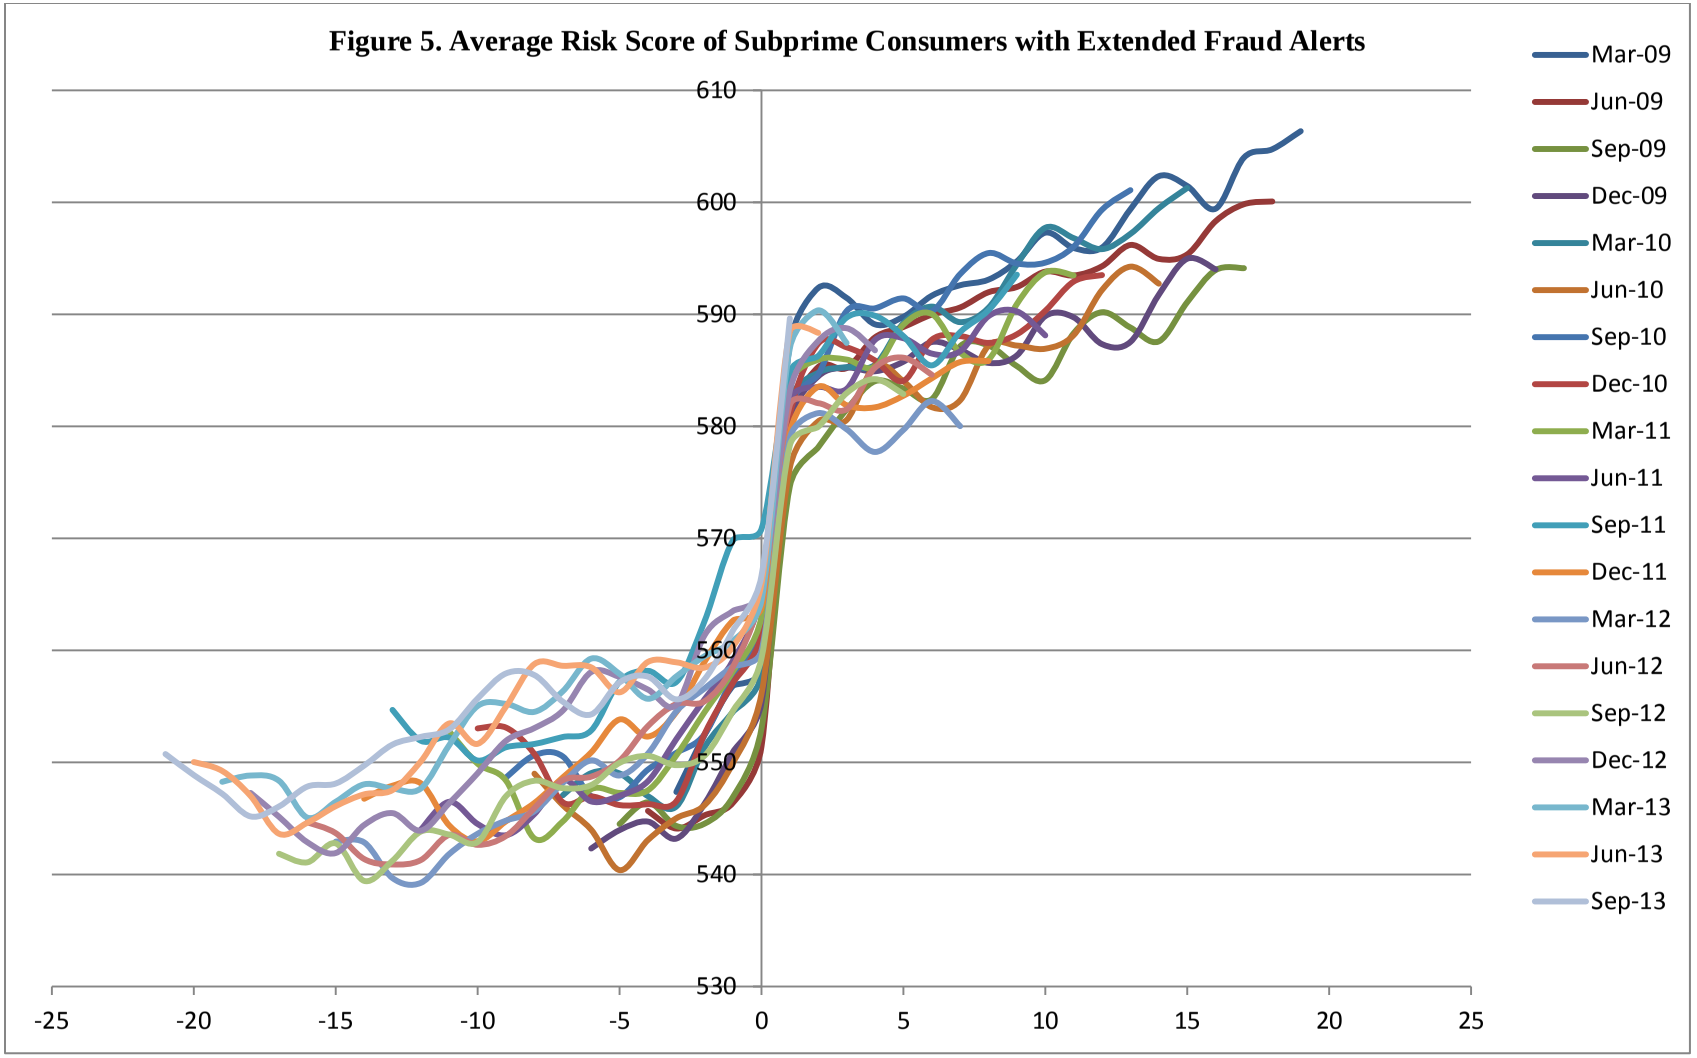

Notes: This figure depicts the average risk score for subprime consumers who filed an extended fraud alert at time 0. Cohorts are defined based on the quarter of fraud alert placement. The average risk score increases at the time of fraud alert filing. Subprime consumers are defined as those with risk scores ≤660 four quarters before the alert. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

Figure 6. Risk Score, Difference Between Propensity Score Matched Individuals with and Without Fraud Alert, Prime

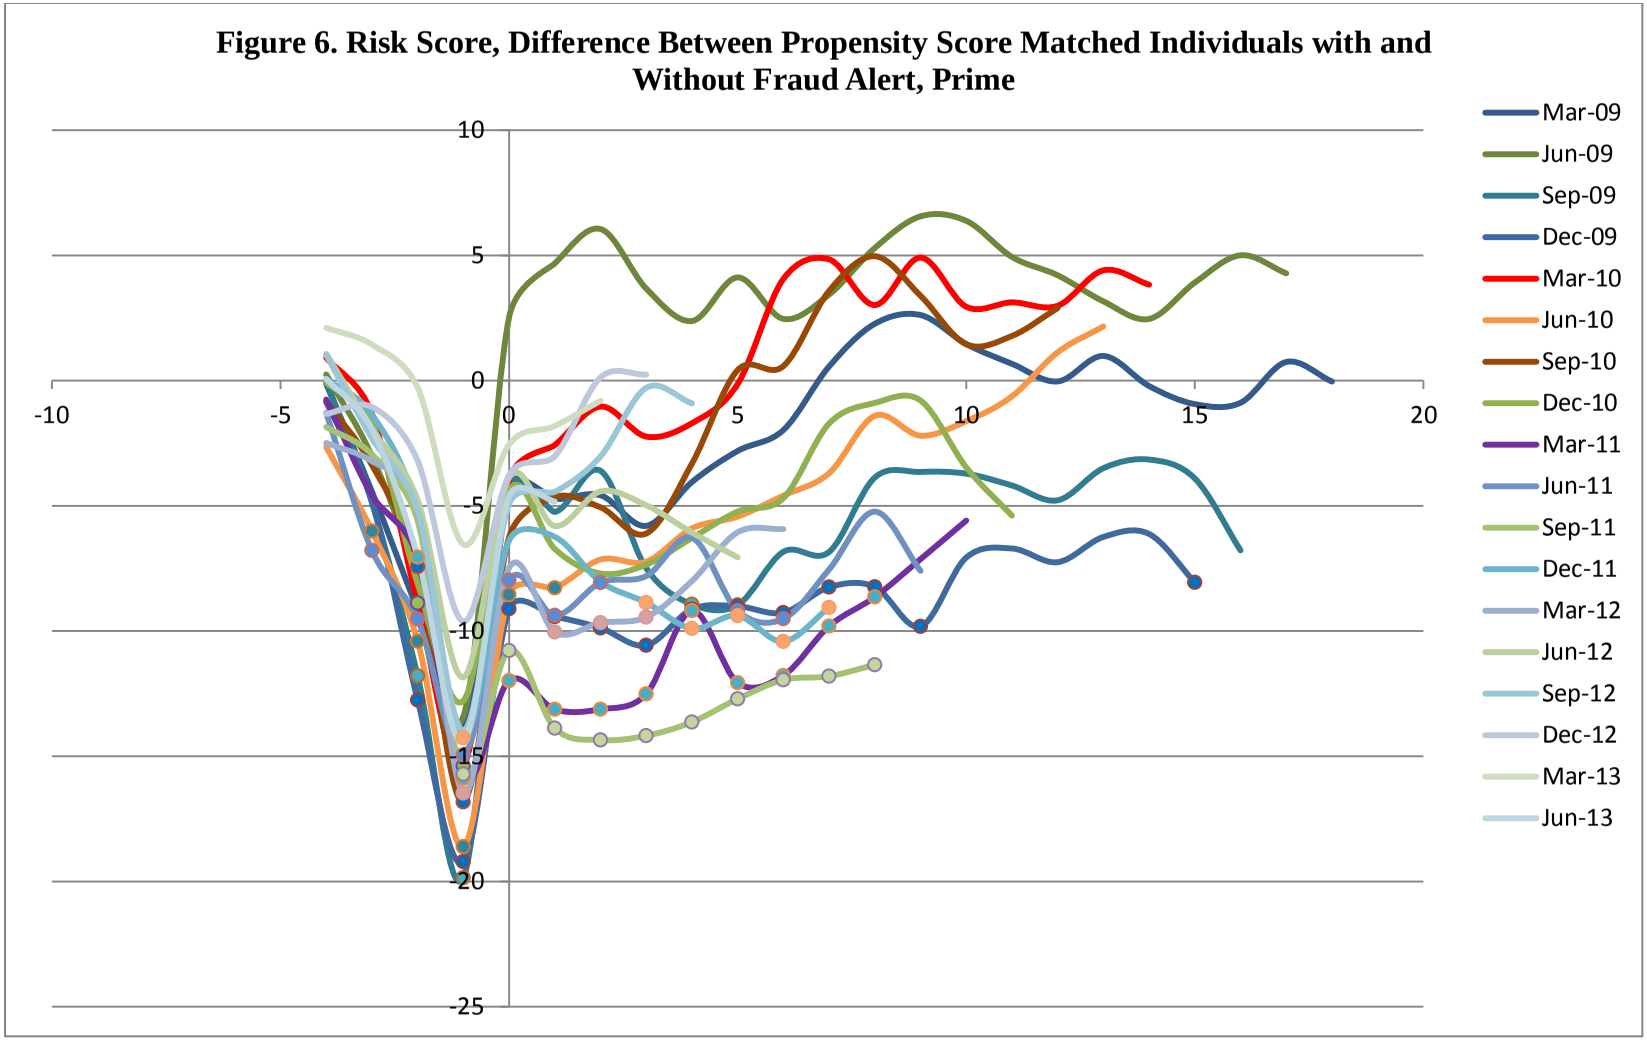

Notes: This figure depicts the difference in risk scores of prime consumers who filed an extended fraud alert at time 0 and risk scores of a propensity score matched control group. See the text of the paper for more details on the propensity score matching. Cohorts are defined based on the quarter of fraud alert placement. The average risk score increases at the time of fraud alert filing. Prime consumers are defined as those with risk scores >660 four quarters before the matching. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

Figure 7. Risk Score, Difference Between Propensity Score Matched Individuals with and Without Fraud Alert, Subprime

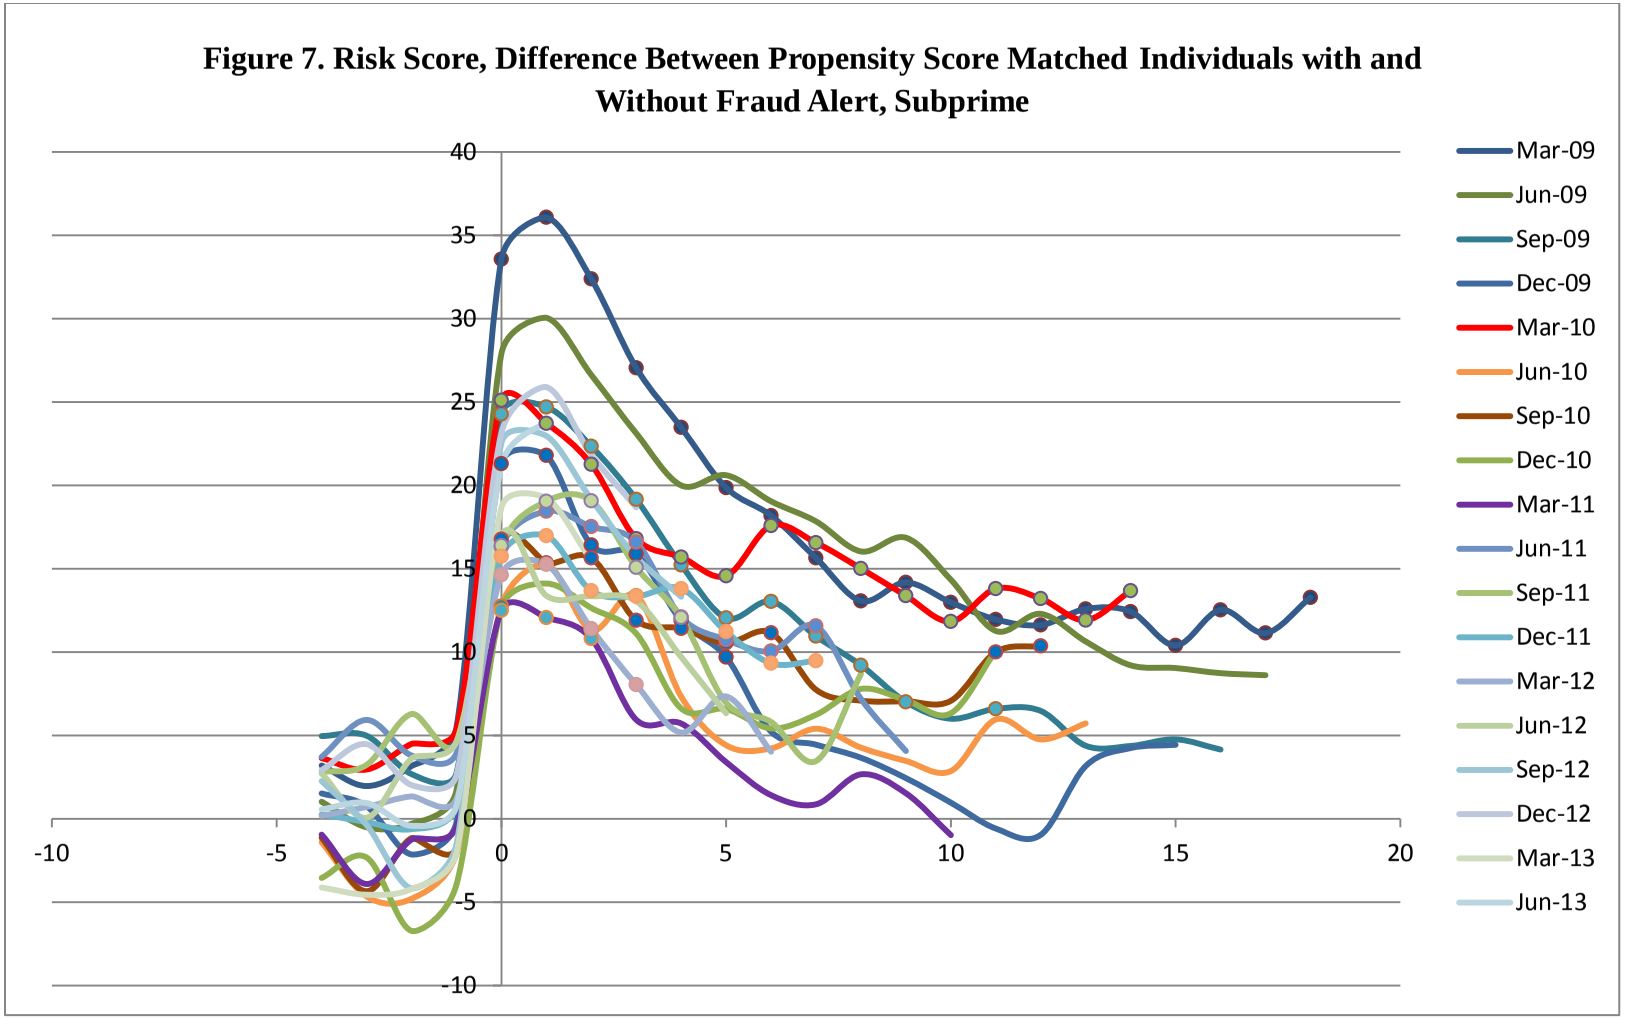

Notes: This figure depicts the difference in risk scores of subprime consumers who filed an extended fraud alert at time 0 and risk scores of a propensity score matched control group. See the text of the paper for more details on the propensity score matching. Cohorts are defined based on the quarter of fraud alert placement. The average risk score increases at the time of fraud alert filing. Subprime consumers are defined as those with risk scores ≤660 four quarters before the matching. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

Figure 8. Inquiries, Difference Between Propensity Score Matched Individuals with and Without Fraud Alert, Subprime

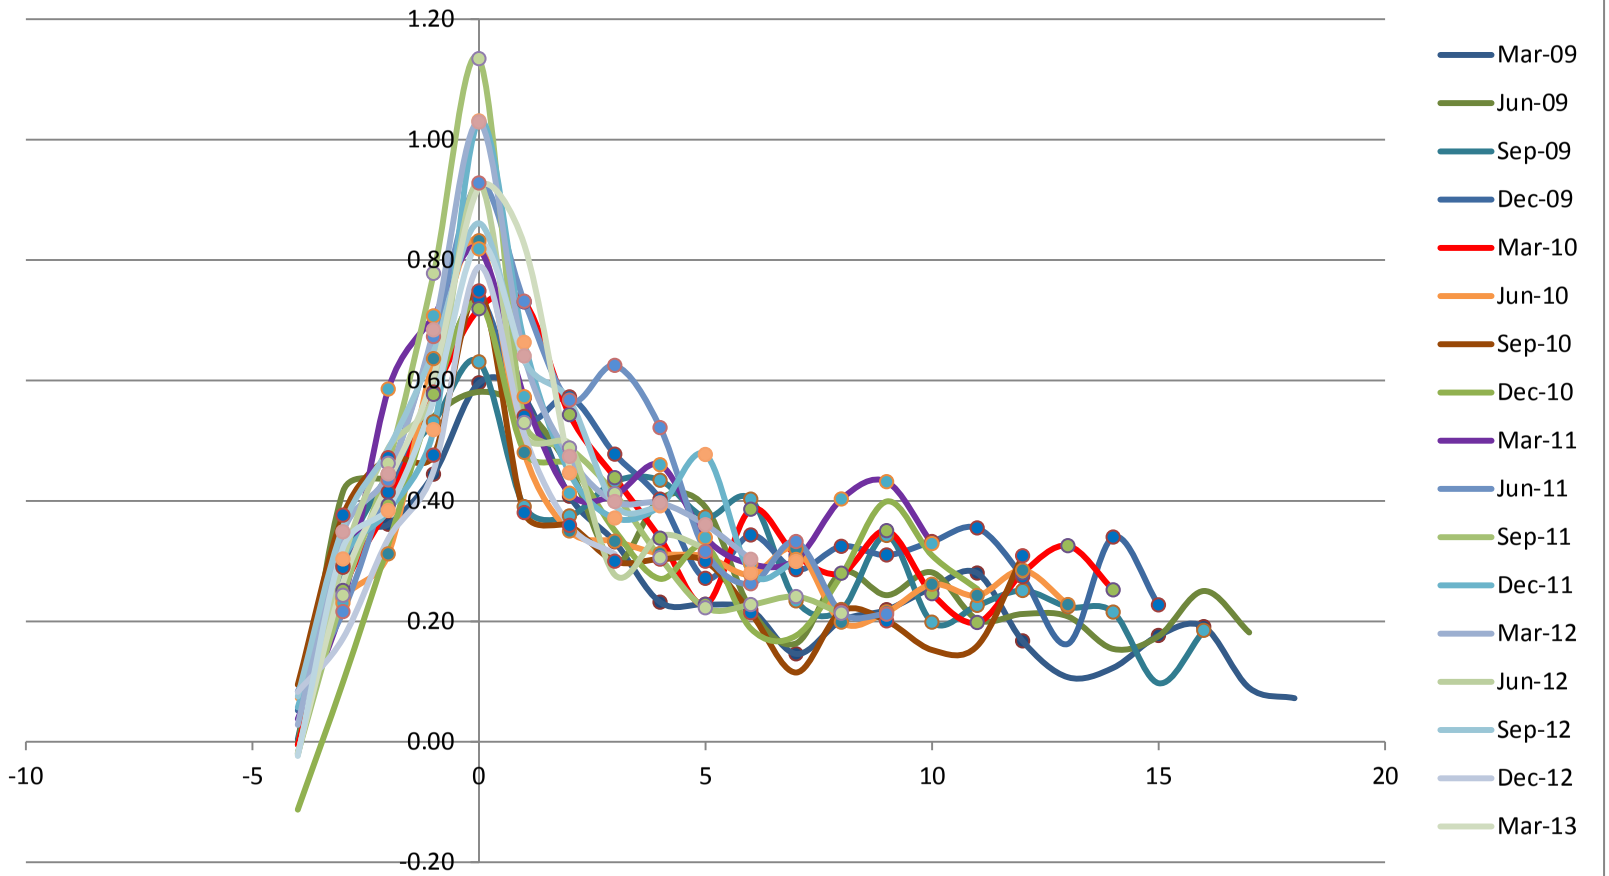

Notes: This figure depicts the difference in inquiries of subprime consumers who filed an extended fraud alert at time 0 and inquiries of a propensity score matched control group. See the text of the paper for more details on the propensity score matching. Cohorts are defined based on the quarter of fraud alert placement. The average number of inquiries increases at the time of fraud alert filing. Subprime consumers are defined as those with risk scores ≤660 four quarters before the matching. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

Figure 9. Inquiries, Difference Between Propensity Score Matched Individuals with and Without Fraud Alert, Prime

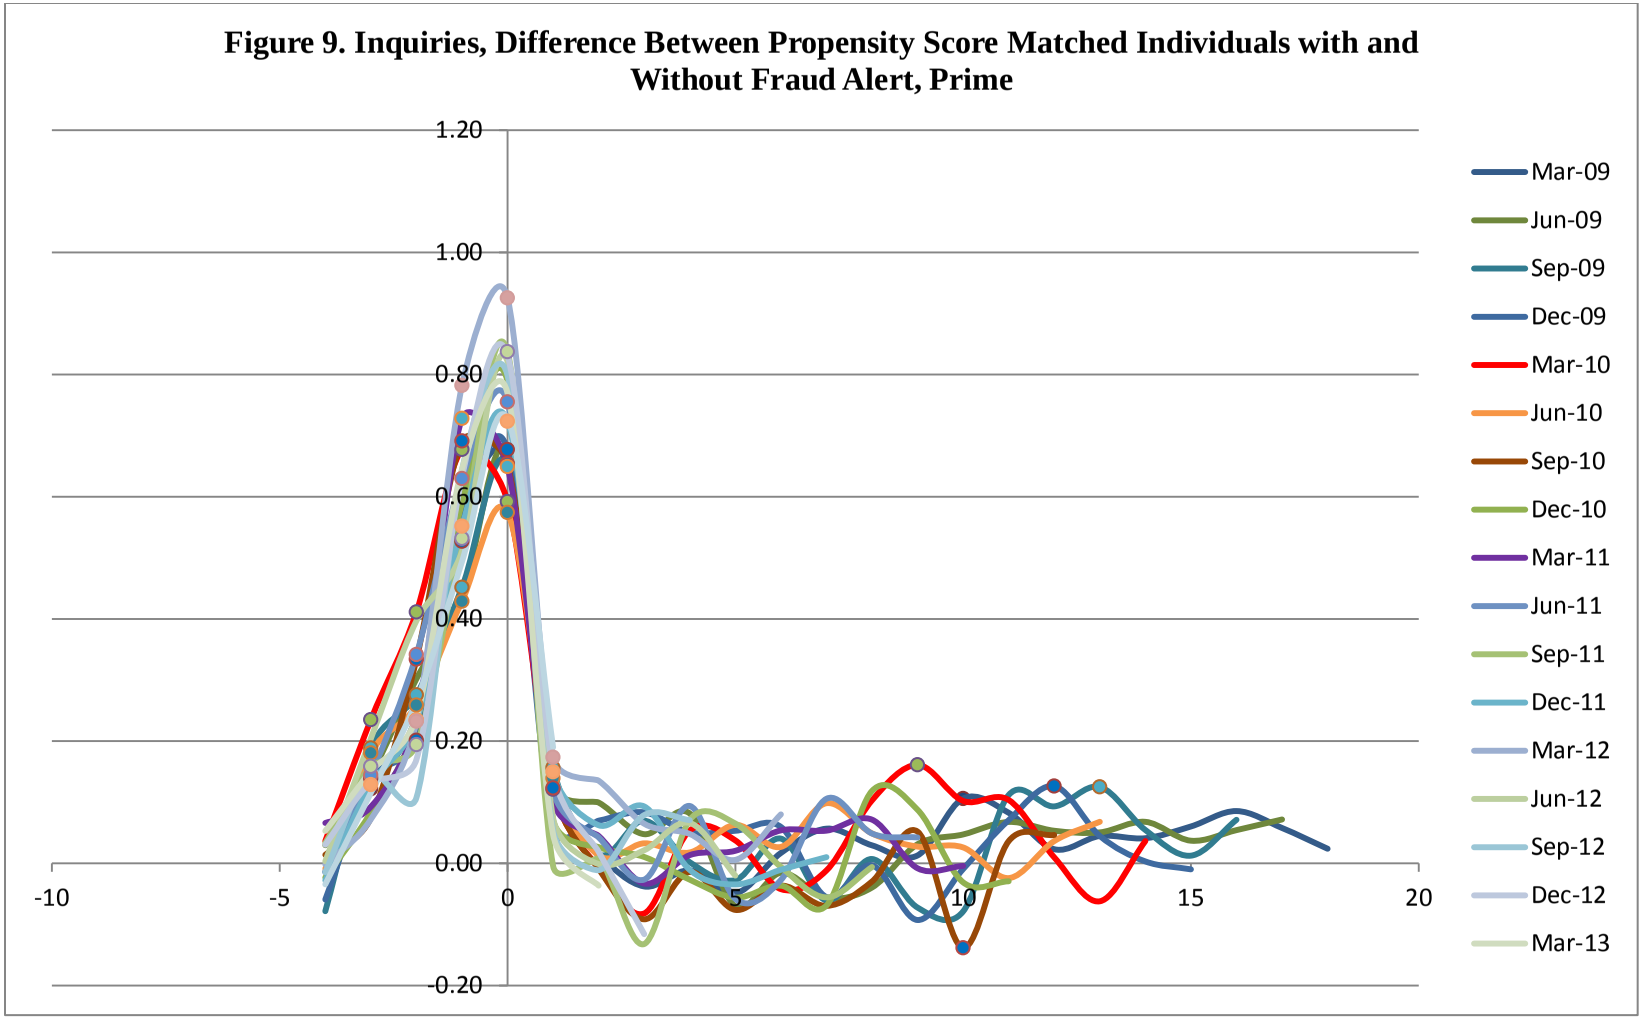

Notes: This figure depicts the difference in inquiries of prime consumers who filed an extended fraud alert at time 0 and inquiries of a propensity score matched control group. See the text of the paper for more details on the propensity score matching. Cohorts are defined based on the quarter of fraud alert placement. The average number of inquiries increases at the time of fraud alert filing. Prime consumers are defined as those with risk scores >660 four quarters before the matching. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

Figure 10. Percent of Bankcards Satisfactory, Difference Between Propensity Score Matched Individuals with and Without Fraud Alert, Subprime

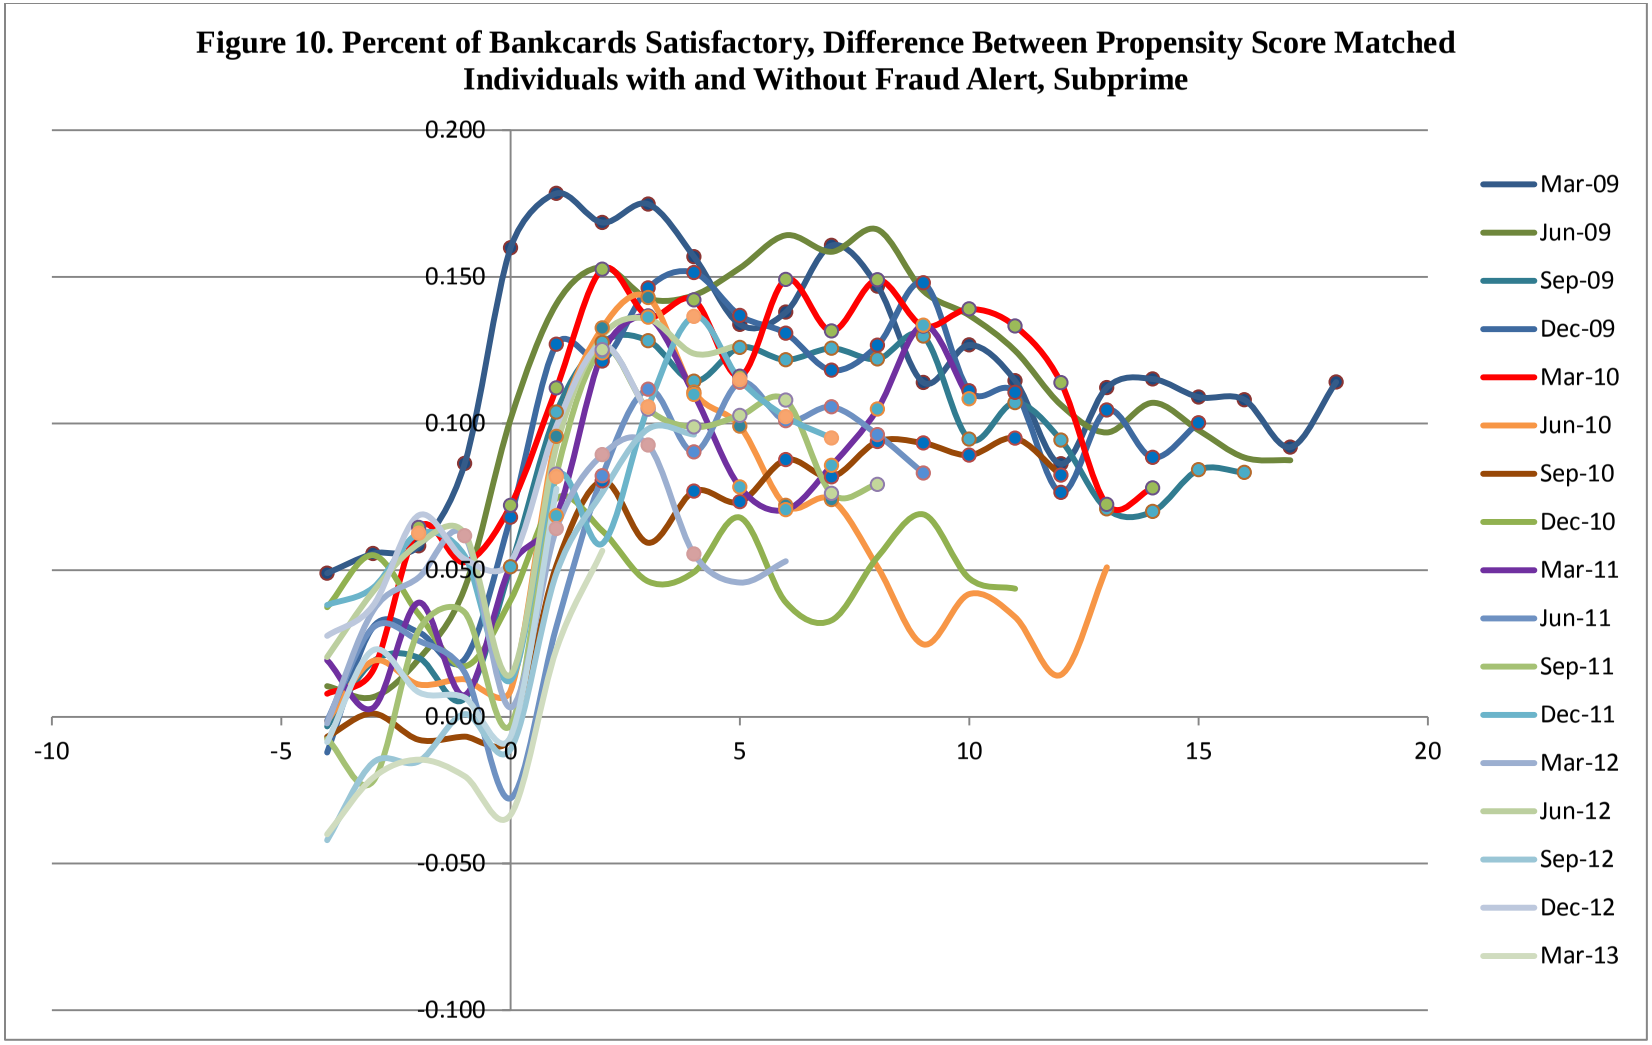

Notes: This figure depicts the difference in the share of credit cards in good standing of subprime consumers who filed an extended fraud alert at time 0 and the share of cards in good standing of a propensity score matched control group. See the text of the paper for more details on the propensity score matching. Cohorts are defined based on the quarter of fraud alert placement. The share of cards in good standing increases at the time of fraud alert filing. Subprime consumers are defined as those with risk scores ≤660 four quarters before the matching. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

Figure 11. Number of Cards with Positive Balances, Difference Between Propensity Score Matched Individuals with and Without Fraud Alert, Subprime

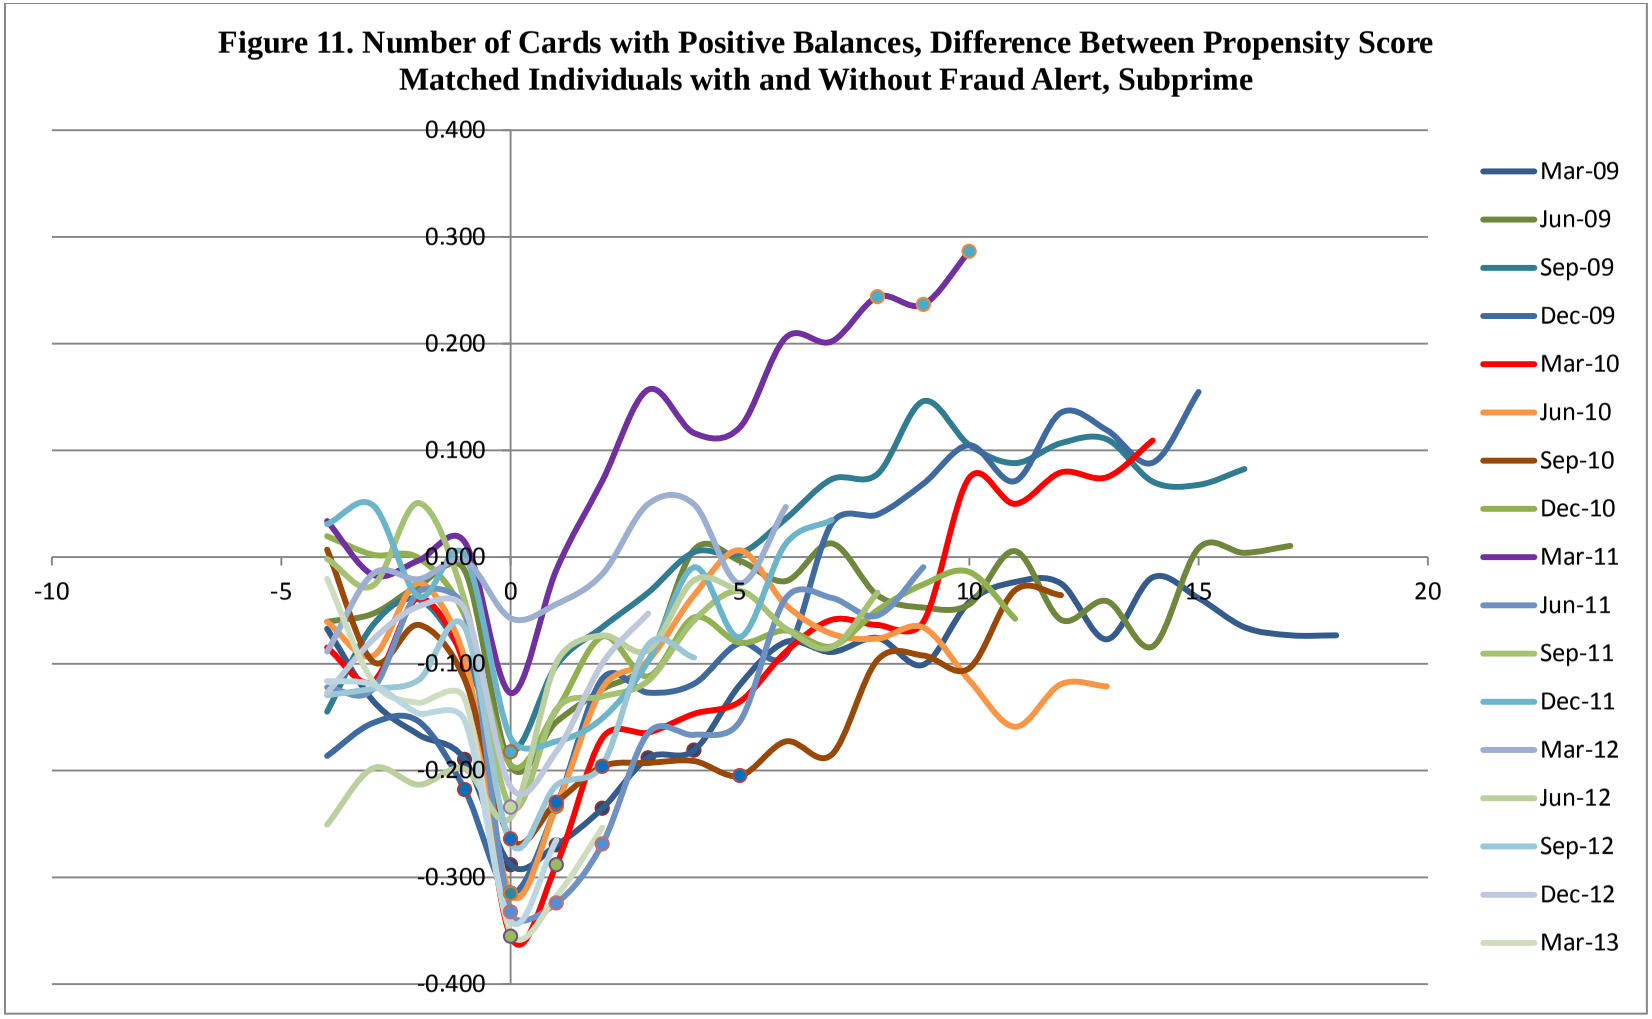

Notes: This figure depicts the difference in the number of credit cards with positive balances of subprime consumers who filed an extended fraud alert at time 0 and the number of cards with positive balances of a propensity score matched control group. See the text of the paper for more details on the propensity score matching. Cohorts are defined based on the quarter of fraud alert placement. The number of credit cards with positive balances stays constant after fraud alert filing. Subprime consumers are defined as those with risk scores ≤660 four quarters before the matching. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

Figure 12. Number of Third-Party Collections, Difference Between Propensity Score Matched Individuals with and Without Fraud Alert, Subprime

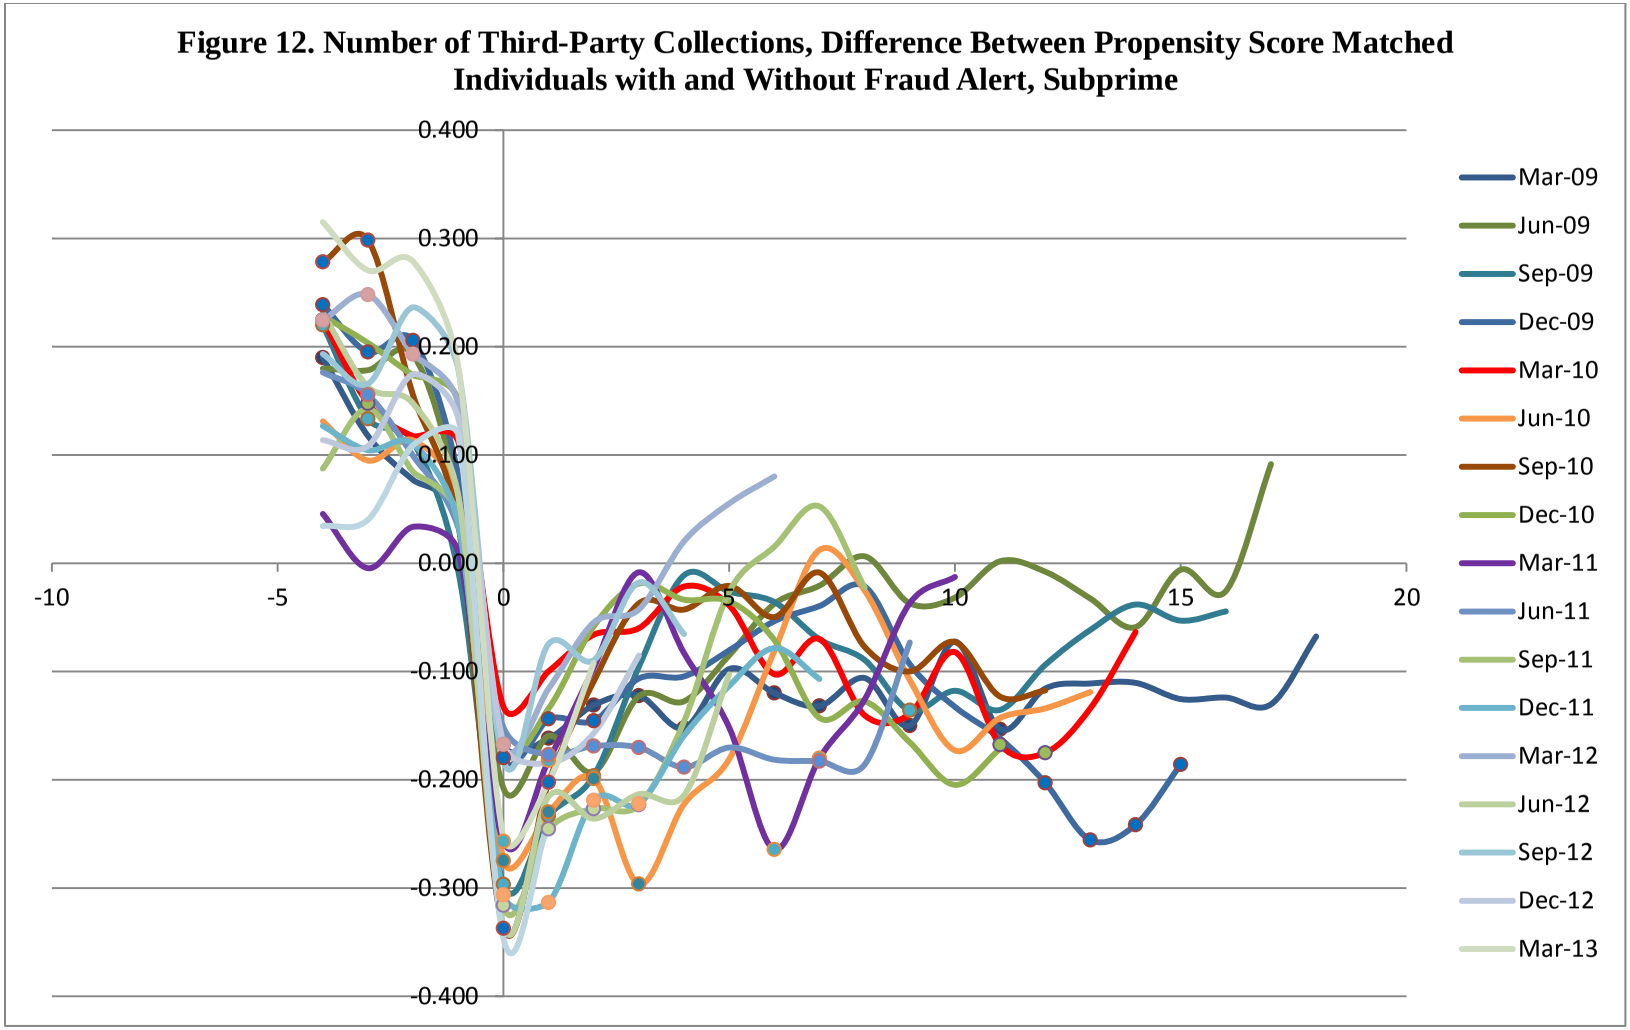

Notes: This figure depicts the difference in third-party collections of subprime consumers who filed an extended fraud alert at time 0 and third-party collections of a propensity score matched control group. See the text of the paper for more details on the propensity score matching. Cohorts are defined based on the quarter of fraud alert placement. The number of third-party collections declines at the time of fraud alert filing. Subprime consumers are defined as those with risk scores ≤660 four quarters before the matching. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

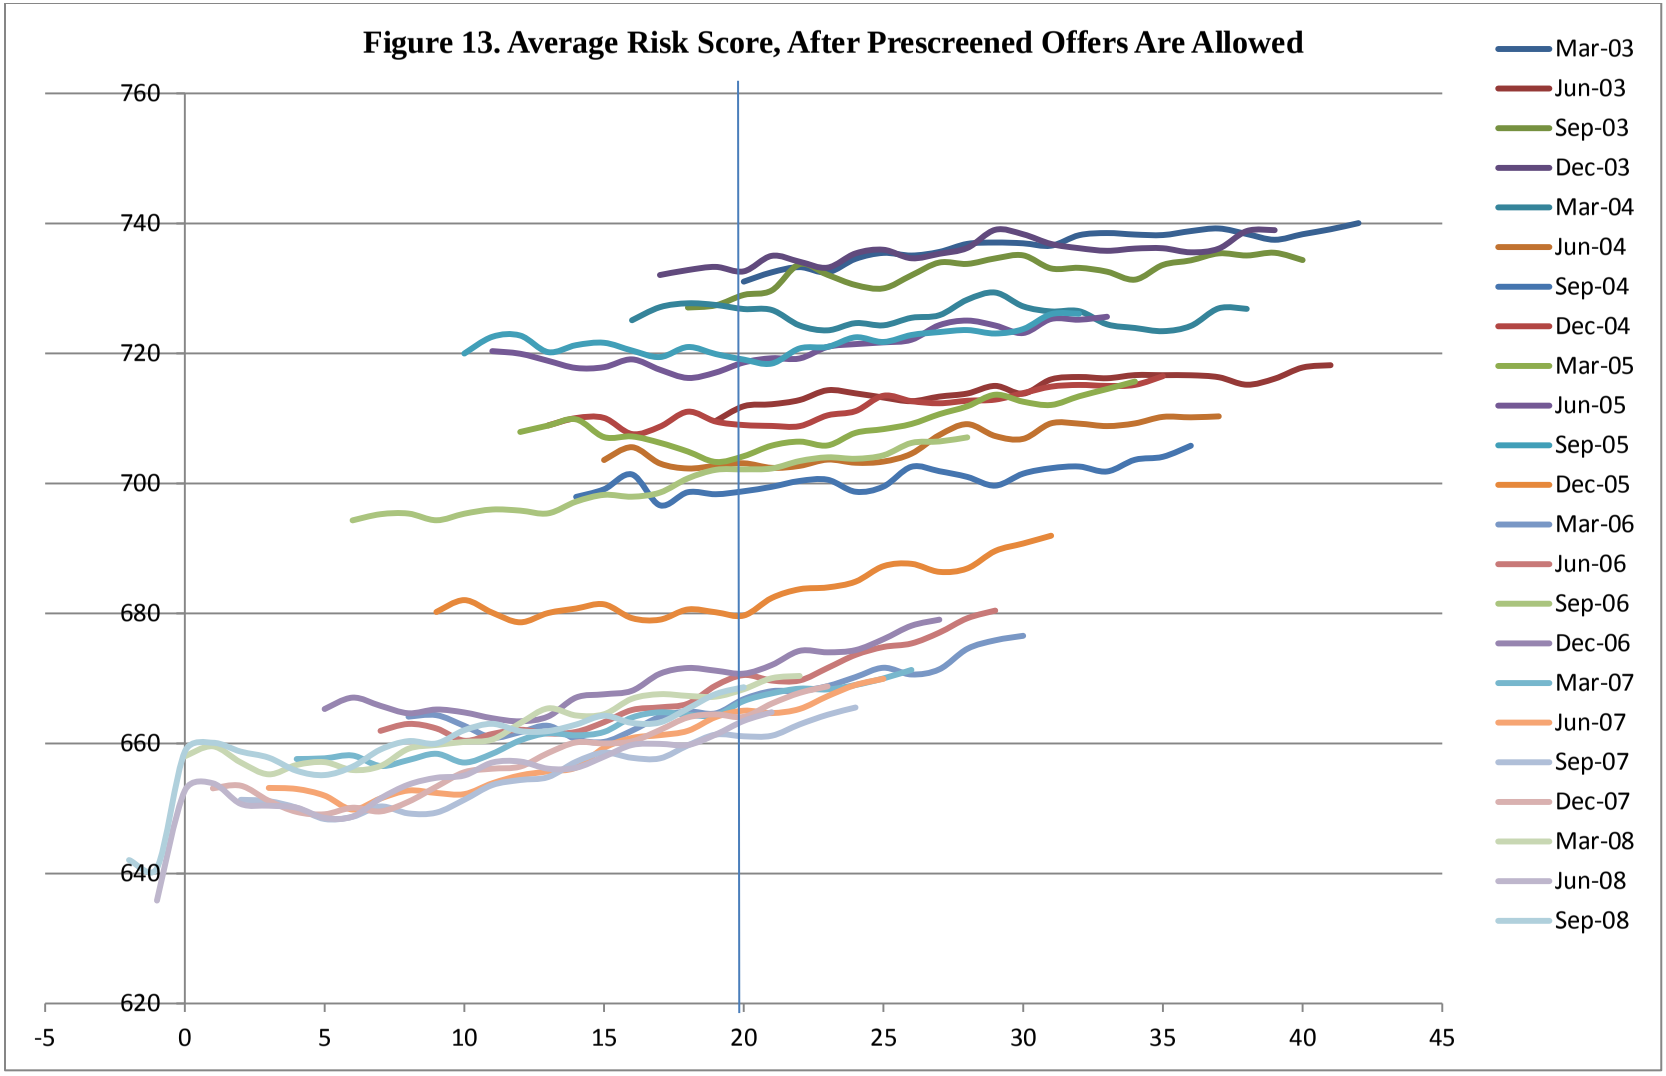

Notes: This figure depicts the average risk score of consumers who filed an extended fraud alert at time 0. Cohorts are defined based on the quarter of fraud alert placement. There is no change in risk score after prescreened credit and insurance offers are allowed to reach these consumers (quarter 20 denoted by a vertical line). Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

Figure 14. Average Risk Score, After Extended Alert's Expiration

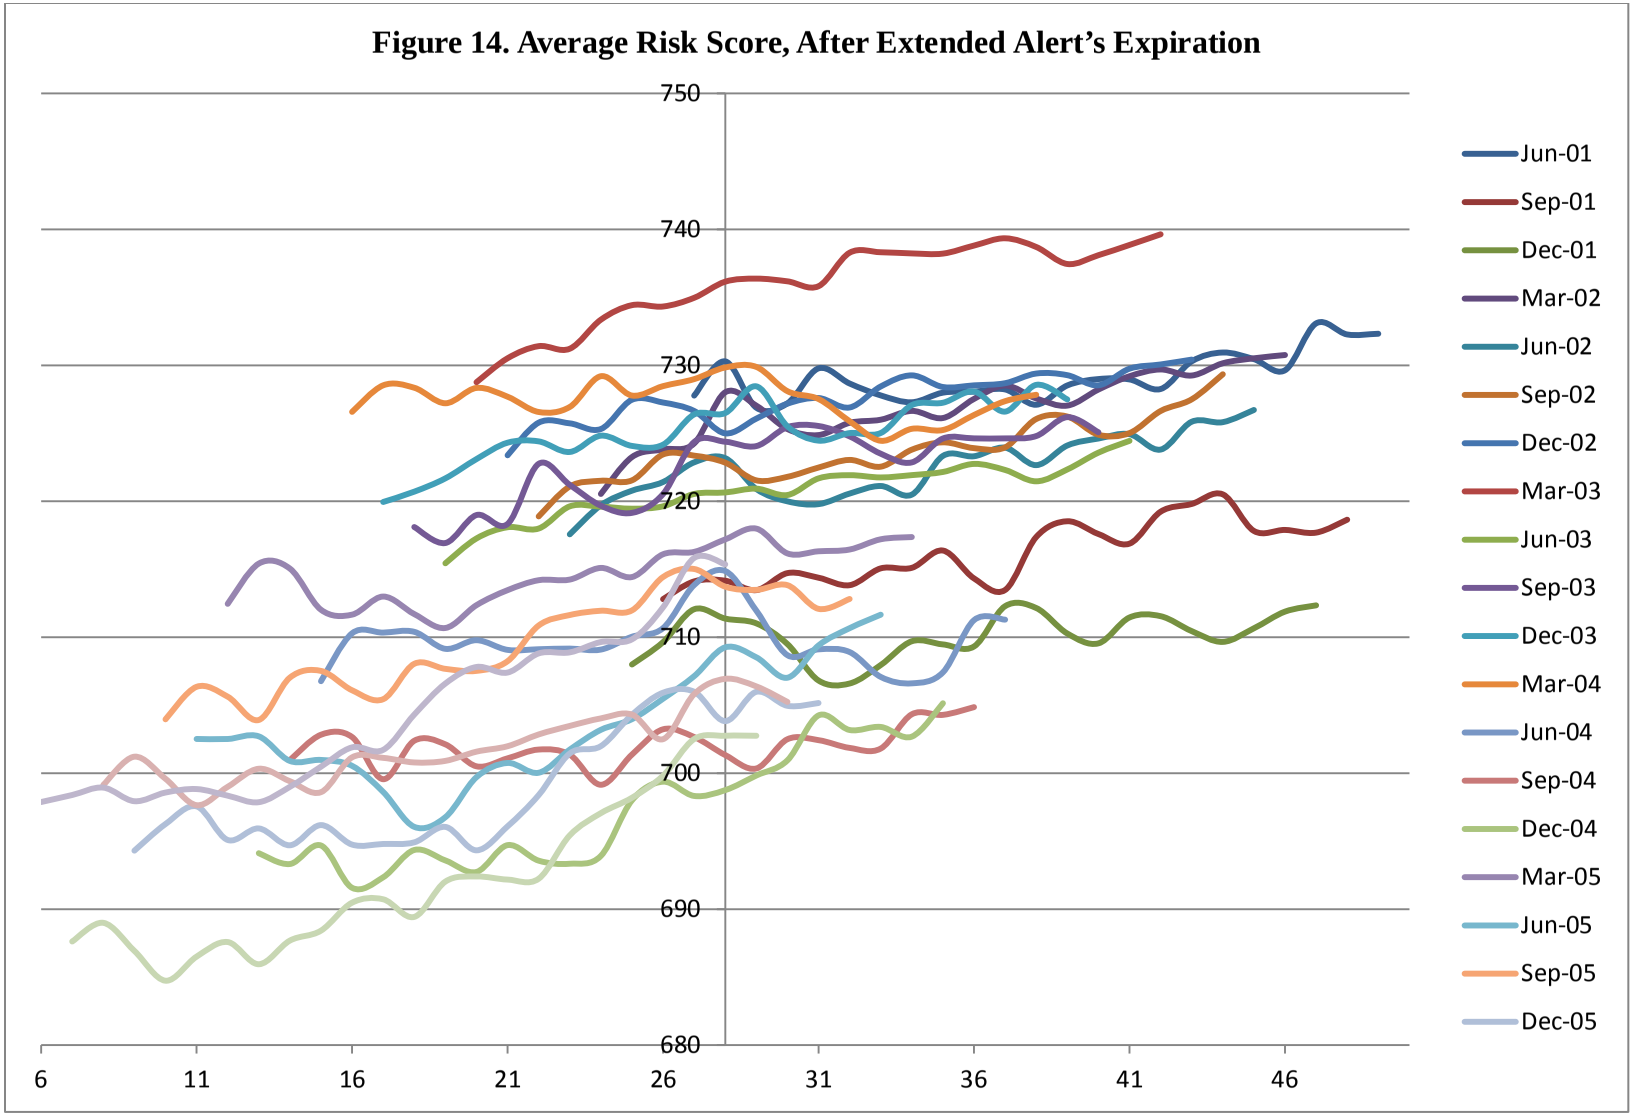

Notes: This figure depicts the average risk score of consumers who filed an extended fraud alert at time 0. Cohorts are defined based on the quarter of fraud alert placement. There is no change in risk score after the extended fraud alert expiration (quarter 28 denoted by a vertical line). Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

Figure 15. Risk Score, Difference Between Propensity Score Matched Individuals with and Without Fraud Alert, Prime Customers, Matched Eight Quarters After Extended Alert

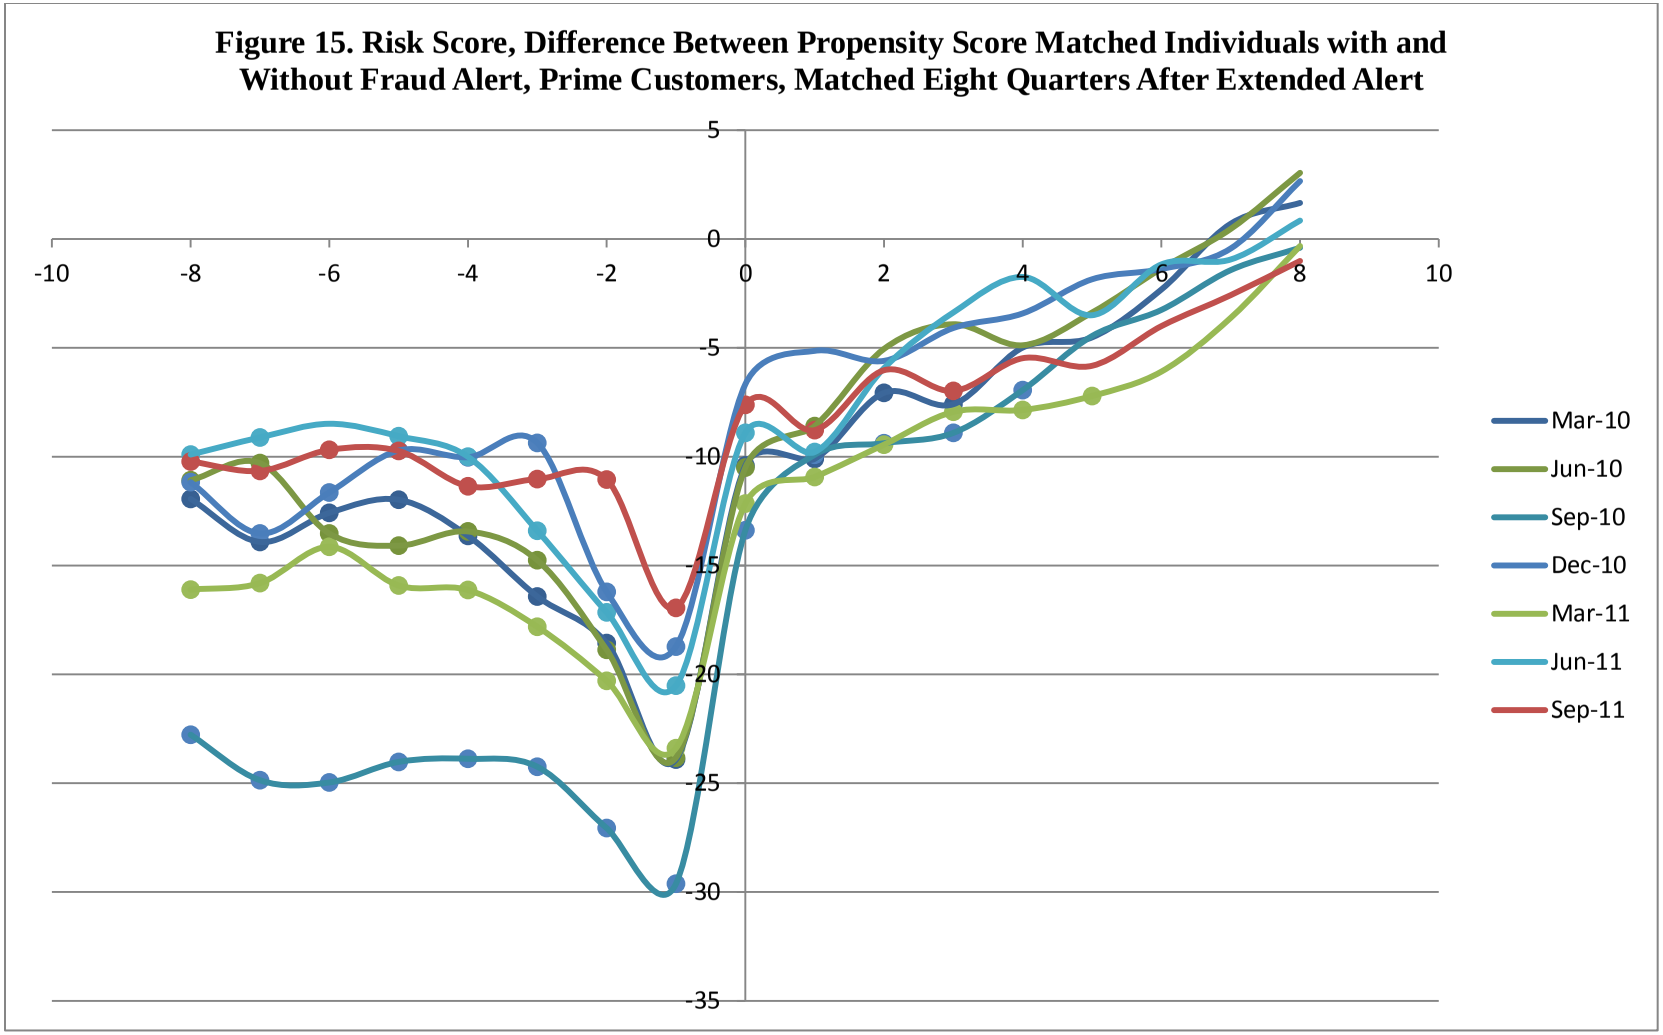

Notes: This figure depicts the difference in risk scores of prime consumers who filed an extended fraud alert at time 0 and risk scores of a propensity score matched control group. See the text of the paper for more details on the propensity score matching. Propensity score matching is done on the variables measured eight quarters after fraud alert filing. Cohorts are defined based on the quarter of fraud alert placement. The response of risk score is similar to earlier figures. Prime consumers are defined as those with risk scores >660 eight quarters after the matching. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.

Figure 16. Risk Score, Difference Between Propensity Score Matched Individuals with and Without Fraud Alert, Subprime Customers, Matched Eight Quarters After Extended Alert

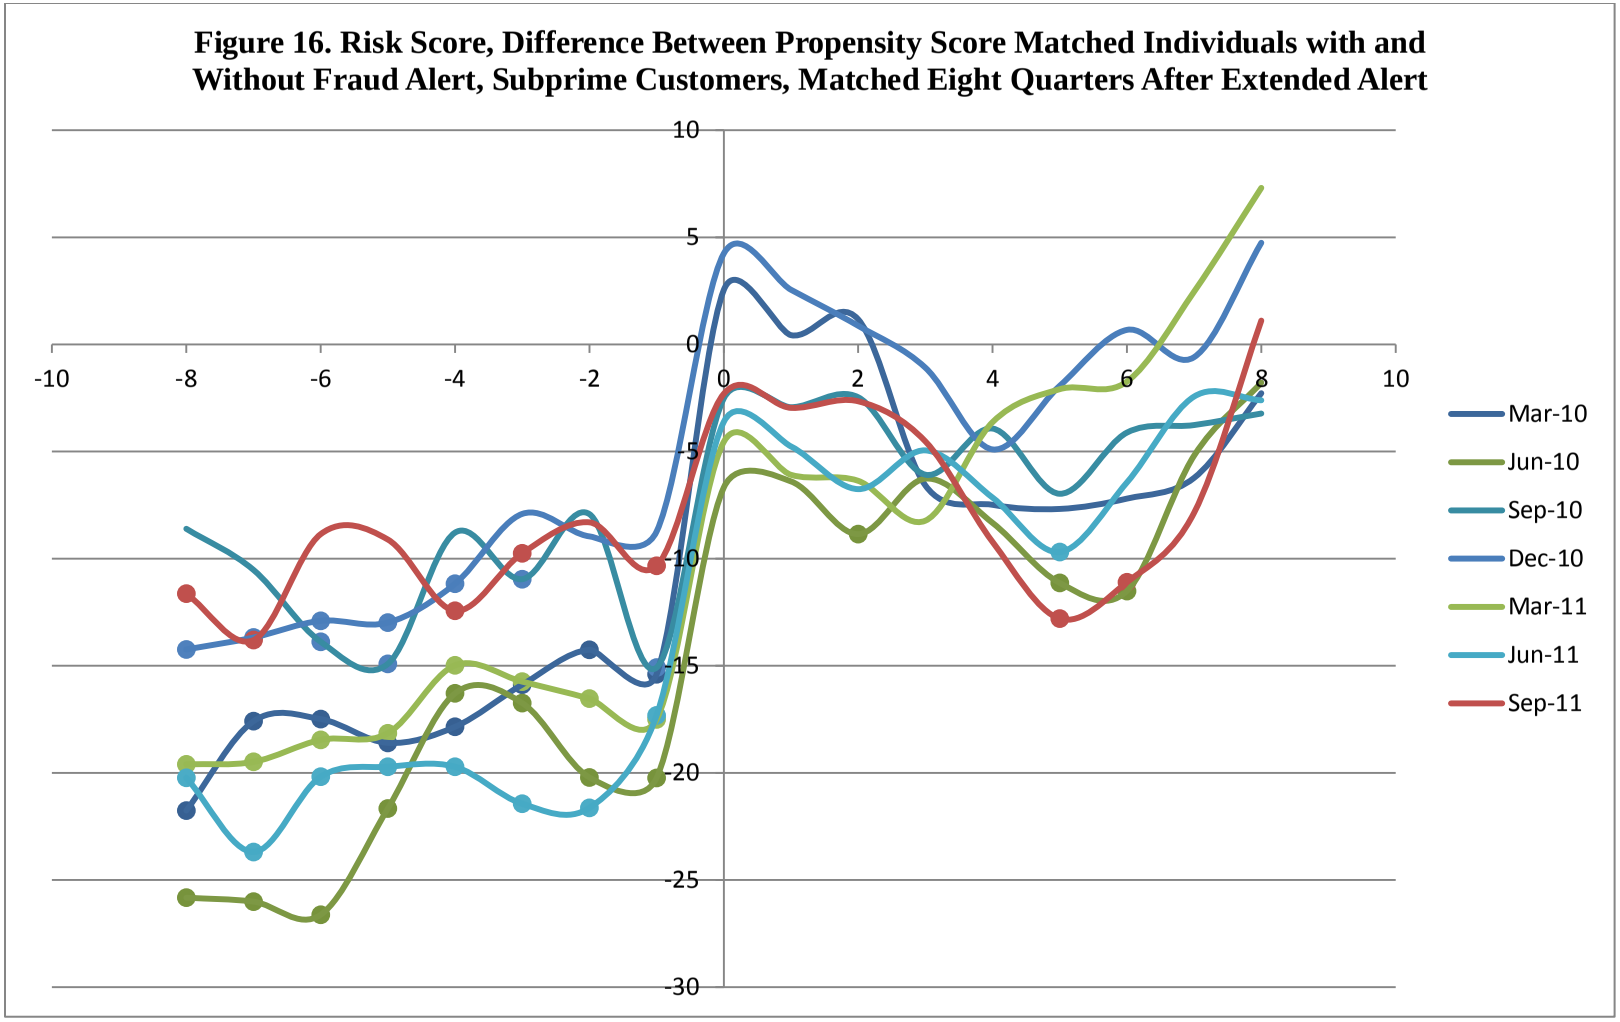

Notes: This figure depicts the difference in risk scores of subprime consumers who filed an extended fraud alert at time 0 and risk scores of a propensity score matched control group. See the text of the paper for more details on the propensity score matching. Propensity score matching is done on the variables measured eight quarters after fraud alert filing. Cohorts are defined based on the quarter of fraud alert placement. The response of the risk score is similar to earlier figures. Subprime consumers are defined as those with risk scores ≤660 eight quarters after the matching. Sources: Authors' calculations using data from FRBNY Consumer Credit Panel/Equifax, augmented with variables obtained by the Payment Cards Center.