# House of Stolen Cards: Does Payment Security Improve Credit Outcomes for Households?*

**Abstract**

Despite widespread adoption of new payment security technologies—chip-enabled cards, tokenized payment systems, and biometric authentication—the implications for household credit access remain poorly understood. We exploit the 2013–2014 rollout of chip-enabled cards as an exogenous shock to payment security. Because the rollout altered which fraud types succeed, fraud victims before and after 2014 face different fraud environments, making simple pre-post comparisons invalid. We instead adopt a stacked cohort event-study design, grouping consumers by their first fraud year and comparing each cohort's credit trajectory to never-defrauded controls around the cohort's own fraud event. We find that lenders expand credit access by $1940 (5.8% of the mean credit limit) to fraud-exposed households following the chip-card rollout. Additionally, we show fraud exposure is associated with higher account closure rates, elevated delinquency, and reduced credit inquiries — patterns that persist after the rollout and are not mitigated by the technology shock. These results point to an asymmetric response to payment security innovation, one that restores lender confidence without fully rebuilding consumer trust, underscoring the need for policy interventions.

Keywords: Payment Security, Financial Literacy, Consumer Credit, Fraud Risk, Credit Card, Lending, Credit Information, Household Finance.
JEL Classification: G21, G28, G50, G51, G53.

*The authors received no specific financial support for this research and have no relevant or material financial interests, relationships, or affiliations that relate to the research described in this paper.

# 1 Introduction

Over the past decade, lenders and payment networks have rolled out a wave of new payment security technologies aimed at curbing financial fraud, such as chip-enabled (EMV) cards, tokenized "numberless" credit cards, one-time passcodes, biometric authentication, and other multi-factor verification tools. These technologies are costly to deploy and are typically justified on the premise that more secure payment instruments reduce fraud losses and restore market participants' confidence in the financial system. Yet, despite the scale of this technological push, we know remarkably little about whether these payment security innovations have any real economic consequences for households.

In this paper, we examine what the causal effect is of improvement in payment security technology on household credit outcomes. We explore this question by focusing on the U.S. credit card market and the introduction of chip-enabled (EMV) credit cards between 2013 and 2014 — the largest nationwide shift in consumer payment security technology to date. Our identification comes from this technology shock: we exploit the rollout of chip-enabled cards as a source of exogenous variation in payment security to examine whether, and how, improved payment technology affects household credit. Specifically, we look at credit cards for two reasons; first credit cards are the most widely used financial instrument in the U.S. with about 81% of adults using at least one card (Fed, 2022). Second, credit card fraud is the most common type of identity theft in the U.S. with costs exceeding $11.64 billion (FTC, 2023; Nilson, 2023).

To study the effect of the chip-card rollout on household credit outcomes, we use 2004–2020 micro-data on consumer credit records from a major U.S. credit bureau via the Gies Consumer and Small Business Credit Panel (GCCP). The central identification challenge is that we cannot simply compare consumers who first experienced credit card fraud before

2014 to consumers who first experienced it afterward: the chip-card rollout itself changed which kinds of fraud succeed (e.g., shifting fraud away from card-present channels toward card-not-present channels), so pre- and post-2014 fraud victims likely differ systematically in the type of fraud they suffered and in who they are. A direct comparison across these groups would therefore confound compositional differences in who gets defrauded, and how, with the effect of the technology shock itself.

We address this challenge with a stacked cohort event-study design. We group consumers by the calendar year of their *first* credit card fraud (cohorts 2010–2013) and, for each cohort, build a stacked sub-sample consisting of that cohort's fraud-exposed consumers together with all never-defrauded consumers, observed in an event-time window around the cohort's fraud year. Because the 2014 chip-card shock falls at a different point in each cohort's event-time window, tracing each cohort's own credit outcome trajectory — and asking whether it bends precisely when the chip-card shock arrives for that cohort — lets us isolate the effect of the technology shock from the effect of the fraud event itself, without ever needing to compare different vintages of fraud victims to one another. The never-defrauded consumers in each cohort's sub-sample are not a conventional difference-in-differences control group. Because selection into fraud is non-random, we neither treat the gap between fraud-exposed and never-defrauded consumers as a causal effect of fraud nor require the two groups to follow parallel trends. Instead, the never-defrauded consumers serve as a market benchmark that absorbs economy-wide movements in credit supply, so that identification rests solely on whether each cohort's credit path, net of the market, bends precisely when 2014 enters its event-time window. Our specifications absorb cohort-by-consumer and cohort-by-year fixed effects and cluster standard errors at the cohort-by-consumer level.

We obtain three main sets of empirical results. Firstly, we show that adoption of chip-enabled credit cards in the U.S. led to a dramatic reduction in instances of stolen credit cards, the number of credit card frauds fell by 65%. In our descriptive event study, we also show that fraud experience is correlated with decline in credit limit, new inquiries and an

increase in default, and card closure.

Second, we find that the chip-card shock triggered a dramatic reversal in lenders' treatment of fraud-exposed consumers. For each pre-chip cohort (first fraud in 2010–2013), tracing credit limit trajectories in event time reveals a consistent pattern: before the chip-card shock arrives in each cohort's window, fraud-exposed consumers see declining credit limits relative to their never-defrauded counterparts, consistent with lenders managing risk exposure by restricting credit supply. Once calendar year 2014 enters each cohort's event-time window, this trajectory reverses sharply. Pooling across cohorts, we estimate an average increase of $1,942 (approximately 5.8% relative to the mean) in credit limits for fraud-exposed consumers after the shock. The consistency of this within-cohort reversal across all pre-chip cohorts — each experiencing the credit-limit jump at exactly the point their event-time window crosses 2014 — suggests that enhanced payment security reduced lenders' perceived value at risk, making them willing to extend credit that had previously been withheld.

Third, we examine the mechanism behind the credit limit recovery. We show that for pre-chip cohorts (first fraud in 2010–2013), the likelihood of future fraud falls sharply, by approximately 80 percent, precisely when the calendar year reaches 2014 for each cohort. That is, conditional on experiencing a first fraud, the likelihood of subsequent fraud drops significantly after the introduction of chip-enabled credit cards. This within-cohort decline in repeat-fraud risk coincides exactly with the credit-limit jump documented in our main results: chip technology closed the vulnerability that made fraud victims persistently high-risk, and lenders responded by extending more credit rather than less.

Overall, we conclude that the 2014 chip-card technology shock was able to reduce instances of credit card fraud, and lenders responded by restoring credit to fraud-affected consumers. However, consumer confidence in the marketplace did not improve as consumers still face declines new inquiries and an increase in default and card closure after the technology shock.

In this paper, we are the first to investigate how improvement in payment technology

can affect credit outcomes for consumers. To the best of our knowledge, no prior work has examined how payment security is associated with credit outcomes for households. We are also the first to study the widespread adoption of chip-enabled cards in the U.S. credit market.

The closest strand to our paper is the literature on payment security and fraud incidence. Sullivan (2013) examine the chip-card rollout in other countries and document fraud displacement — card-present fraud declines while card-not-present fraud rises — and Hayashi (2019) confirm a parallel shift in U.S. fraud patterns following the 2014 rollout. These papers establish that the technology shock changed which kinds of fraud succeed, but they focus on fraud incidence as the outcome of interest. We go a step further by asking how the resulting change in fraud risk alters lender behavior and household credit access — a consequence that aggregate fraud statistics do not reveal. Kahn et al. (2017) show that perceived payment security spills across consumers through peer networks. On the consumer side, Blascak and Toh (2022) document that prior fraud experience makes consumers more likely to close credit accounts upon a subsequent fraud event, and Hamdi et al. (2024) find that victimization is associated with elevated bankruptcy, foreclosure, and delinquency, particularly for low-income and younger borrowers. These papers establish that fraud exposure matters for consumer financial health we add to the literature by examining how a payment security technology shock can impact fraud exposure and credit access.

Our paper also contributes to the literature on lender behavior in consumer credit markets. Chava et al. (2023) show that lenders transmit funding shocks to consumers by reducing credit limits, establishing that lender balance-sheet constraints are a key driver of credit supply variation. We document a complementary channel: exogenous shifts in fraud technology change lenders' perceived risk of extending credit to a given consumer, generating large movements in credit supply even absent any change in lender funding costs. Agarwal et al. (2010) and Agarwal et al. (2018) document how reward programs and adverse selection shape credit card spending, debt, and default.

Finally, our paper speaks to the literature on how negative credit reporting signals affect consumer credit access. Dobbie et al. (2020) and Gross et al. (2021) exploit the removal of bankruptcy flags from credit reports and find large downstream effects on credit limits and interest rates. Guttman-Kenney (2023) studies the natural disaster flag and finds, by contrast, that it does not improve credit access. Fraud-exposed consumers in our setting carry an implicit negative signal — a "lost or stolen" account closure that appears in their credit report — whose net effect on credit extension is ex ante ambiguous, since it reflects both elevated risk and involuntary victimization. We show that the credit market consequences of this signal depend critically on the payment security environment: the same fraud event that triggered credit restriction before 2014 led to credit expansion afterward. This finding implies that the effect of a negative reporting flag on lender behavior is not fixed but is shaped by the risk environment in which lenders operate, a dimension that the existing literature has not examined.

The remainder of the paper is as follows: Section 2 discusses credit card transactions, fraud liability, and the EMV chip technology; Section 3 describe a simple one period model to lay foundation of the hypothesis, Section 4 describes the data; Section 5 presents the empirical strategy and results, Section 6 discusses the mechanism analysis; and Section 7 concludes.

# 2 Background

## 2.1 Credit Card Transactions & Fraud Liability

In a typical credit card transaction, the payment process involves four primary parties (see Figure A.1). The process begins with the consumer using their credit card to acquire a good or service. This action is supported by the consumer's lender[1], who finances the

[1]The six major credit card lenders in the US are Chase, American Express, Citibank, Capital One, Bank of America and Discover.

transaction. The lender is responsible for underwriting the consumer's credit. Technology for this exchange is provided by the processor[2], which maintains the infrastructure necessary for the transaction to propagate through the financial network. Finally, there's a merchant who has the point-of-sale (PoS) terminal provided by the processor. The merchant does not need to have any formal relationship with the consumer's lender, instead they interact with the processor. The processor generates revenue from charging a fee, called interchange fee, to the merchant which is typically around 2% to 2.5% of the transaction cost. The processor shares part of the interchange fee with the lender, who also generates revenue by charging a high interest rate, late fees, and an annual renewal fee to the consumer.

There are two broad kinds of credit card fraud: card-present fraud, where the transaction involves the card being swiped, inserted, or tapped at the PoS terminal; and card-not-present fraud, which occurs without the physical card — online, via mobile apps, over the phone, etc (see Figures A.2 and A.3). If the consumer suspects a transaction to be fraudulent or unauthorized, the Fair Credit Billing Act (FCBA) of 1974 enables them to raise a dispute with their lender. FCBA requires the lender to investigate and settle the dispute within 90 days. Sophisticated financial institutions typically maintain a specialized fraud department responsible for overseeing suspicious transactions, while less sophisticated lenders may choose to outsource fraud detection to external service providers. When a consumer raises a dispute, the lender closes the affected credit card to prevent further misuse. This fraudulent card is marked as closed due to *lost* or *stolen* in the consumer's credit report. The lender also issues a new replacement card with a new unique tradeline number; however, the credit history of the consumer remains the same in their credit report.

During the dispute investigation period, the lender cannot legally collect any amount for the disputed transaction from the consumer. More importantly, FCBA limits the liability of a cardholder for a fraudulent transaction to $50. Further, many lenders offer a zero liability policy which reduces the consumer's liability to $0 in case of fraud. FCBA also states that

[2]Four major US processors are Visa, MasterCard, American Express and Discover, with Visa and Mastercard together controlling 87% of market share by volume for credit card transactions.

a fraudulent transaction dispute cannot impact the consumer's credit history. In case of a card-present fraud, the merchant is not liable [3], if it follows the best practices such as using tap or chip processing instead of swiping the card to execute the transaction. In card-present fraudulent transactions, the cost of fraud is either borne by the lender or in some cases by the merchant. Whereas, the processor is responsible only for providing the technology for transaction execution, not for fraud.

Lenders also claim that the consumers' credit standing will not get impacted by the occurrence of the credit card fraud. For example; *"When you report your card as lost or stolen to the [lender], they will arrange for a new card to be sent to you. Replacing a credit card won't affect your credit score....the card account and the age of that account remain the same for the purpose of credit reporting." (Chase, 2024)*

Basic economic theory tells us that a fraudulent event should alter the lender's risk assessment of the affected consumer: either the consumer did not follow best practices to keep their credentials safe, or they may intentionally be disputing legitimate transactions to avoid paying for goods or services. In both scenarios, the lender's operational and capital costs for extending credit rise significantly.

## 2.2 EMV Chip Technology: How It Works

Historically, credit card transactions in the United States relied on a magnetic stripe — a 1960s-era technology that stored static, unencrypted information such as the card number, expiration date, security code, and cardholder name on a band of magnetic particles on the back of the card. Because this information never changed, it could be intercepted and replicated with relatively simple equipment. A widespread fraud technique known as "skimming" exploited this vulnerability: criminals attached illegal card readers to ATMs, gas pumps, and point-of-sale terminals to silently capture the magnetic-stripe data of any card that was swiped. The captured data could then be encoded onto a blank card to create

[3]For card-not-present fraud, which are more common in e-commerce, the liability is always on the merchant.

a near-perfect counterfeit, indistinguishable from the original at the point of sale. The ease and low cost of this attack meant that compromised card data could be harvested in bulk, used repeatedly, and sold on secondary markets before cardholders or lenders detected the fraud.

Chip-enabled credit cards — standardized under the Europay, Mastercard, and Visa (EMV) specification — replaced this static model with dynamic cryptography. Rather than storing fixed data, an EMV chip is a microprocessor that, at the moment of each transaction, generates a unique one-time transaction code (a cryptogram) using a secret key shared only between the chip and the card-issuing lender. The terminal transmits this cryptogram to the lender's authorization system, which verifies it against its own record of the key. Because the cryptogram is computed fresh for every purchase and encodes transaction-specific details such as the amount and timestamp, it cannot be replayed or used to construct a counterfeit card: an attacker who intercepts the cryptogram for one transaction gains nothing that is useful for any subsequent transaction. This property — sometimes called "dynamic data authentication" — directly closes the skimming attack vector that magnetic stripes left open.

A natural question is why chip cards continued to carry a magnetic stripe at all, given the stripe's well-known vulnerability. The answer is practical interoperability: at the time of the EMV rollout, a large share of point-of-sale terminals worldwide — particularly in smaller merchants, kiosks, and older infrastructure — did not yet support chip readers. Retaining the magnetic stripe ensured that cardholders could still complete transactions at non-chip-enabled terminals, preventing chip adoption from creating a two-tier payment system in which some consumers could not use their cards at certain merchants. More directly relevant to the liability framework discussed in Section 2.3, the magnetic stripe also serves as a technical fallback within chip-capable terminals themselves: if a chip fails to be read after repeated attempts, the card reader may prompt a swipe so that the transaction can still be completed. Removing the stripe entirely would mean that a malfunctioning chip renders the card unusable, an unacceptable reliability outcome for lenders. As chip

infrastructure has matured and contactless payments have grown, the case for the magnetic stripe has weakened.[4]

Between 2013 and 2014, chip-enabled credit cards achieved widespread adoption in the United States, driven by a combination of industry coordination and policy. Card networks — Visa, Mastercard, American Express, and Discover — had been working with lenders to migrate their card portfolios to EMV chips in the years prior, and by 2013–2014 the majority of newly issued U.S. credit cards carried chips. The 2014 BuySecure Initiative, introduced by the Obama administration, provided further political momentum and pushed federal agencies to adopt chip-and-PIN technology for government-issued cards, but the card-level rollout was already well underway through industry-led action (Sullivan, 2013; Hayashi, 2019). By 2016, chip-enabled terminals had become the standard across most retail channels. The consequence for card-present counterfeit fraud was immediate and dramatic: because a skimmed magnetic-stripe can no longer be used to produce a functional counterfeit once the merchant's terminal requires the chip, the frequency of card-present fraud fell sharply in the years following the rollout, as documented in Figure 1.

## 2.3 EMV Liability Shift

The technological upgrade to chip cards was accompanied by a fundamental reallocation of financial responsibility for fraud losses, known as the EMV liability shift. Prior to this shift, lenders absorbed the losses for the vast majority of card-present fraud regardless of how the transaction was processed. This arrangement reflected the practical reality that merchants had little ability to detect counterfeit magnetic-stripe cards, and the fraud losses ultimately stemmed from vulnerabilities in the card itself rather than in merchant behavior. Lenders therefore priced the expected cost of card-present fraud into their overall risk models, and fraud-exposed consumers were viewed as signals of elevated future loss.

Under the October 2015 EMV liability shift framework, the allocation of loss depends on

[4]Mastercard has announced a formal phase-out, under which no new Mastercard cards will carry a magnetic stripe from 2029 onward and all stripes will be retired by 2033.

whether the merchant has upgraded to a chip-compliant terminal. If a chip card is presented at a merchant that still only accepts a magnetic swipe — because it has not yet deployed a chip reader — and a fraudulent transaction occurs, liability for that transaction shifts to the merchant rather than the lender. The logic is straightforward: the chip would have prevented the fraud, and the merchant's failure to require it is the proximate cause of the loss. Conversely, if the merchant does use a chip reader and a fraud still occurs (e.g., a stolen physical card is used with a valid chip), liability remains with the lender, since the merchant followed protocol and the card's security was not at fault. The 2015 liability-shift deadline announced by the card networks gave merchants a strong financial incentive to upgrade their point-of-sale terminals to chip readers ahead of that date, accelerating terminal deployment on the merchant side.

A key feature of the EMV liability framework is that the fallback hierarchy is strictly governed by the merchant's contract with its payment processor. Merchants are not permitted to accept a magnetic-stripe swipe simply because it is more convenient or faster. The protocol requires that a consumer first attempt to insert or tap their chip card up to three times; only if all three chip-read attempts fail — indicating a possible chip malfunction — is the merchant contractually permitted to fall back to a magnetic-stripe swipe. Should three swipe attempts also fail, the merchant may then manually key in the card details. This escalation requirement is material for liability: if a merchant bypasses the chip-attempt requirement and accepts a swipe on the first try, it has violated its processing agreement and bears full liability for any resulting fraud, regardless of whether the chip was actually defective. The three-attempt rule therefore creates a paper trail that documents whether the fallback was genuinely necessitated by hardware failure, and payment processors can audit transaction logs to determine whether the merchant followed protocol before ruling on a fraud claim. In practice, this means that "convenience swipes" — merchants allowing customers to swipe to avoid the marginally longer chip-dip process — expose the merchant to liability that would otherwise fall on the lender.

Figure A.3 illustrates these three primary scenarios. Scenario 1 depicts card-present fraud with a counterfeit card: before 2014, lenders were liable; after the rollout, a merchant that accepts a swipe bears the loss while a merchant that requires the chip is protected. Scenario 2 depicts card-present fraud with a stolen physical card: lender liability is unchanged before and after 2014, since the chip was used correctly and the fraud reflects card theft rather than a technology failure. Scenario 3 depicts card-not-present fraud: merchant liability is unchanged throughout, as the EMV standard addresses only card-present channels.

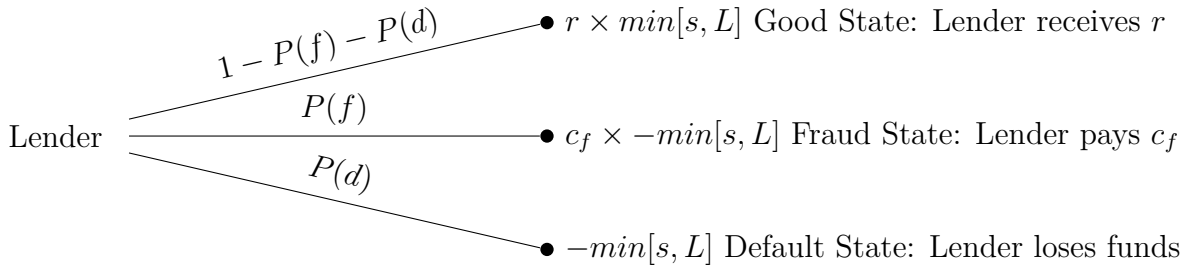
The liability shift changed the incentive structure facing lenders in a consequential way. Before 2014, every fraud-exposed consumer represented a prospective future loss for the lender, regardless of whether subsequent fraud would be card-present or card-not-present. After the shift, a lender whose consumer holds a chip card and transacts at chip-enabled terminals is largely insulated from card-present counterfeit losses — the most common and scalable form of fraud — because the merchant now bears that liability. This reduced the expected loss associated with serving fraud-exposed consumers, making them comparatively less risky to extend credit to, and providing the direct incentive channel through which we expect the technology shock to affect credit supply.

# 3 A Simple Model

To understand the lender's choice, we propose a simple one period model. The lender is tasked with the choice of extending a credit line of size $L$ to a consumer. The consumer will make a purchase of size $s$ and repay the loan immediately. The size of the purchase is unknown to the lender.

The lender faces three possible states of the world. In the default state, which occurs with probability $P(d)$, the consumer defaults on the loan and the lender loses the funds lent to the consumer, up to $L$. In the fraud state, which occurs with probability $P(f)$, the consumer experiences fraud and the lender must pay $c_f$, the cost of fraud. If the lender is

liable for the fraud, then $c_f = 1$ as they must cover the loss. If merchant is liable for fraud, then $c_f$ is some small amount to reflect that, while not liable, the lender must manage the process of making the consumer whole and collecting funds from the merchant. In the good state, which occurs with probability $1 - P(f) - P(d)$, the consumer repays the loan in full and the lender receives an interest rate of $r$ on the funds lent to the consumer. The lender's expected return can be summarized in the following decision tree:




It can be shown that the profit to the lender can be expressed as:

$$\Pi = (1 - P(f) - P(d))(r \times min[s, L]) - P(f)(c_f \times min[s, L]) - P(d)(min[s, L]) \tag{1}$$

The lender will choose to extend credit if the expected profit is positive, i.e. $\Pi > 0$. Rearranging this inequality and finding the FOC, we find that the lender follows this policy for $L^*$:

$$L^* = \begin{cases} L & \text{if } [(1 - P(f) - P(d)) \times r] > P(f)(c_f) + P(d) \\ 0 & \text{if}[(1 - P(f) - P(d)) \times r] \leq P(f)(c_f) + P(d) \end{cases} \tag{2}$$

From this model, we see that the lender's decision to extend credit depends on the relative size of the expected return in the good state, $[(1 - P(f) - P(d)) \times r]$, versus the expected costs in the fraud and default states, $P(f)(c_f) + P(d)$. If the expected return exceeds the expected costs, the lender will extend credit up to $L$. If not, the lender will choose not to extend any credit.

The purpose of this model is to illustrate the role of fraud in the lending decision. When

the probability of fraud, $P(f)$, or the cost of fraud, $c_f$, increases, the expected costs to the lender increase making it less likely that the lender will choose to extend credit. Conversely, if the probability of default, $P(d)$, decreases, the expected costs decrease, making it more likely that the lender will extend credit. This model leaves us with two testable predictions. First, when the probability of fraud increases, lenders will reduce the supply of credit and/or be less willing to extend new credit. Second, following the introduction of chip cards, lenders will increase the supply of credit as both $P(f)$ and $c_f$ decrease with the payment technology.

# 4 Data

We use the Gies Consumer and Small Business Credit Panel[5] (GCCP), an anonymized consumer credit reporting panel dataset[6] obtained from a major credit bureau. This dataset contains a one-percent random sample of individuals who have credit reports with the credit bureau. The dataset has annual snapshots of two to three million individual credit reports taken in the first quarter of each year from 2004 to 2025.

A key advantage of the GCCP is that it contains detailed annual information on credit balances, credit limits, payment history, past-due balances, collections, inquiries, and public records such as bankruptcies for each individual's reported credit line items ("tradelines") such as mortgage, auto loan, student loan, credit card etc. The dataset has Vantagescore, a credit scoring model with a range between 300 and 850. The dataset also has demographic variables such as age, gender, marital status and 5-digit zip code. We restrict our sample to credit cards for years 2004 to 2020. We drop records with missing age, credit score, or income, as well as those in which reported age increases by 10 years or more in a single year. For each credit card account we observe credit limit, account balance, account open date, charge-off amount, delinquency, collection amount along with a unique anonymized lender identification number.

[5]See Fonseca (2023) for a detailed discussion on GCCP.
[6]See Gibbs et al. (2023) for a detailed discussion of the consumer credit report databases.

## 4.1 Identifying Credit Cards with Fraud

For our main sample we keep tradelines where the account type is a credit card. To identify credit card fraud, we exploit a flag that indicates an account was closed with a status code of "*lost* and *stolen*". To account for the fact that such closed fraudulent accounts remain in the credit report data for an undefined period, we restrict our definition of fraud to be the first instance of a *lost* and *stolen* reporting flag for a credit card. Our definition of fraud avoids overstating the frequency of fraudulent credit card accounts.

Section 7 reports the summary statistics for the full sample; an average consumer has total credit limit of $32,970, total credit card balance of $4,710, a credit score of 696, holds about 4 cards and has a likelihood of 2.56% for experiencing a fraud. Whereas, the median consumer has total credit limit of $14,520, total credit card balance of $810, a credit score of 708, holds about 3 cards during the sample period.

Figure 1 plots the number of credit cards in our sample that experienced fraud each year from 2005 to 2020. The gray dashed line marks end of first quarter of 2014, the time of widespread chip-card adoption in the U.S. Following the widespread chip adoption, the number of *lost and stolen* credit cards in our data declines markedly. This time-series pattern provides the descriptive backdrop for our empirical strategy: the chip rollout constitutes a sharp, technology-driven reduction in card fraud that we exploit as a quasi-experimental shock.

A potential criticism for our fraud definition is that we cannot observe *lost* and *stolen* cards separately. We broadly define a lost card as a card that is replaced without any fraudulent transactions, and a stolen card is a card with fraudulent transactions. We do not believe that lenders would view these differently. A card that is lost, is simply one that has not yet been found by a criminal. While a consumer may have lost their card in their home, it could also be sitting in the back seat of a cab waiting for a fraudster to find and use it. One could argue that lenders are more receptive to lost cards as they did not observe the negative event of a fraudulent transaction. However, the lender should view these consumers

as riskier since they have revealed themselves as careless with control of their cards. Even if lenders are indifferent to lost cards, this would bias against our results. If lost cards are no concern then this increases noise in our observation of fraud. Hence, any bias in our results due to these lost cards should go against our hypothesis that extending additional credit to the fraudulent card consumers, increases the cost for the lenders.

Figure A.4 shows how total credit limits evolve around the fraud event for fraud-exposed consumers, estimated using a within-consumer event-study regression with no control group. Panel (a) restricts to consumers whose first fraud occurred between 2010 and 2013 (the pre-chip regime), while Panel (b) restricts to consumers whose first fraud occurred between 2014 and 2017 (the post-chip regime). In both panels the horizontal axis shows event time $j$ relative to each consumer's first-fraud year. Among pre-chip victims, credit limits follow an increasing path before the fraud event and then decline in the post-fraud years, consistent with lenders pulling back on credit extension once a fraudulent event is observed. Among post-chip victims, the same event-time window reveals a different pattern: credit limits continue to rise after the fraud event, with no visible pullback. Because this figure conditions only on fraud-exposed consumers. We take this raw pattern as motivating evidence for the causal design developed in Section 5.

# 5 Empirical Strategy and Results

## 5.1 Descriptive Event Study

We first report the pooled event-study associations between fraud exposure and credit outcomes such as total credit limit, credit card account closure, 90-day delinquency, and credit inquiries. We pool consumers into two groups: a pre-chip group comprising consumers whose first fraud occurred between 2010 and 2013, and a post-chip group comprising those whose first fraud occurred between 2014 and 2016. For each group, we estimate the following event-study regression at the consumer-year level:

$$Y_{i,t} = a + \sum_{j=-3}^{4} \beta_j \left( Fraud_i \times D_{i,t-j} \right) + \delta_t + \eta_i + \varepsilon_{i,t} \tag{3}$$

where $Y_{i,t}$ is the credit outcome of interest for consumer $i$ at calendar year $t$; $Fraud_i$ is a dummy equal to one for fraud-exposed consumers and zero for never-defrauded consumers; $D_{i,t-j}$ is an event-time dummy equal to one when $t$ is $j$ years from consumer $i$'s first-fraud year, for $j \in \{-3, \ldots, +4\}$; $\delta_t$ are calendar-year fixed effects; and $\eta_i$ are consumer fixed effects. Standard errors are clustered at the consumer level. The specification is estimated separately for the pre-chip sample (consumers with first fraud in 2010–2013) and the post-chip sample (consumers with first fraud in 2014–2017), and the $\beta_j$ coefficients trace the average credit outcome path across the event window relative to the fraud year for each regime. The resulting estimates reflect average associations between fraud exposure and outcomes within each regime rather than causal effects. We report them here as descriptive correlations; the cohort-level identification that isolates the causal effect of the chip-card liability shift is the focus of Section 5.2.1.

Figure 3 plots the event-study estimates for total credit limit. Among pre-chip cohorts (Figure 3 panel a), the credit limit is on an increasing trend before fraud event. At the fraud year ($j = 0$) credit limits for fraud-exposed consumers fall relative to the never-defrauded benchmark, remaining below it through the full post-event window. On average across pre-chip cohorts, fraud exposure is associated with a decline in total credit limits. Among post-chip cohorts (Figure 3 panel b), pre fraud event trends are again increasing, however at fraud year ($j = 0$) the credit limits for fraud-exposed consumers still rise relative to the never-defrauded benchmark and continue to increase through $j = +4$. The within fraud cohort design in Section 5 causally investigates the impact of payment security on consumer credit limit.

Figures 4 to A.7 plot the event-study estimates for credit card account closure rate (%), 90-day delinquency rate (%), and number of new credit inquiries. Unlike credit limits, these consumer-side outcomes display qualitatively similar correlations in both the pre-chip

and post-chip periods: fraud exposure is associated with higher rates of account closure, elevated 90-day delinquency, and reduced credit inquiries in both regimes. The estimated magnitudes shift modestly across periods, but the sign of each association is stable. This pattern is consistent with fraud generating consumer financial distress — and a withdrawal from the credit market — regardless of whether chip technology was in place at the time of the fraud. The introduction of chip-cards did not insulate consumers from the disruption caused by the fraud event itself. Taken together, the technology shock did not materially alter the plausible consumer-side correlates of fraud exposure.

## 5.2 Fraud Cohort Difference-in-Difference Analysis

We examine the impact of the introduction of chip cards on consumer credit outcomes using a stacked cohort event-study design. We group fraud-exposed consumers into cohorts $c \in \{2010, \ldots, 2014\}$ according to the calendar year of their *first* credit card fraud. For each cohort $c$, we build a stacked sub-sample consisting of that cohort's fraud-exposed consumers together with never-defrauded consumers as controls, observed over an event-time window of $j \in \{-3, \ldots, 4\}$ years relative to cohort $c$'s fraud year. Stacking sub-samples in this way ensures that we never directly compare consumers whose first fraud occurred in different calendar years — and who therefore plausibly experienced different types of fraud — to one another. Our main specification estimates the following stacked regression at the consumer-by-cohort-year level:

$$Y_{i,c,t} = a + \sum_{j=-3}^{4} \beta_j \left(Fraud_{i,c} \times D_{c,t-j}\right) + \delta_{c,t} + \eta_{c,i} + \varepsilon_{i,c,t} \tag{4}$$

where $Y_{i,c,t}$ is the credit limit for consumer $i$, in cohort $c$'s stacked sub-sample, at calendar time $t$. $Fraud_{i,c}$ is a dummy equal to one for cohort $c$'s fraud-exposed consumers and zero for the never-defrauded consumers stacked alongside them. $D_{c,t-j}$ is an event-time dummy equal to one when calendar time $t$ is $j$ years from cohort $c$'s first-fraud year, for $j \in \{-3, \ldots, 4\}$.

Our coefficients of interest are the $\beta_j$, which trace out cohort $c$'s dynamic credit outcome path around its own fraud year, relative to the never-defrauded consumers stacked into that cohort's sub-sample. $\delta_{c,t}$ are cohort-by-calendar-year fixed effects, which absorb shocks common to all consumers within a cohort's stacked sub-sample in a given year. $\eta_{c,i}$ are cohort-by-consumer fixed effects, which remove time-invariant heterogeneity across consumers within each cohort's stacked sub-sample. Because a never-defrauded consumer can be stacked into the sub-samples of multiple cohorts, cohort-by-consumer (rather than consumer-only) fixed effects are required so that one cohort's variation is not attributed to another. We cluster standard errors at the cohort-by-consumer level, the level at which a consumer's treatment status is determined within each stacked sub-sample.[7]

It is important to be clear about the role of the never-defrauded consumers in each stacked sub-sample. They are not a conventional difference-in-differences control group, and we do not interpret the difference in credit outcomes between fraud-exposed and never-defrauded consumers as the causal effect of fraud. Because selection into fraud is non-random, fraud-exposed consumers differ from never-defrauded consumers along dimensions that lenders observe and act on, so the level gap between the two groups confounds the consequences of fraud with the characteristics that made these consumers likely to be defrauded in the first place. Nor does our identification rely on the two groups sharing parallel pre-trends. Rather, differenced through the cohort-by-year fixed effects, the never-defrauded consumers function as a market benchmark: they net out economy-wide movements in credit supply, such as the aggregate credit cycle, changes in lending standards, and the direct effect of calendar year 2014 on borrowers generally, that would otherwise contaminate the fraud-exposed trajectory. What survives this differencing is cohort $c$'s credit path net of prevailing market conditions, and the causal content of the design comes entirely from the timing of the bend in that net trajectory, as we explain next.

Crucially, the 2014 chip-card shock falls at a different point $j$ in each cohort's event-time

[7] Our results are robust to clustering at the consumer level alone.

window: for the 2014 cohort it coincides with the fraud event itself ($j = 0$), while for the 2010 cohort it does not arrive until $j = 4$. Tracing whether each cohort's $\beta_j$ path bends precisely when the chip-card shock reaches that cohort's cohort-specific event-time horizon — rather than at a common calendar date — lets us isolate the effect of the technology shock from the effect of the fraud event itself, without ever needing to compare different vintages of fraud victims to one another directly.

### 5.2.1 Results

Figure 5 presents our main cohort-level event-study results for credit limits, with the underlying regression coefficients reported in Table 4. For consumers whose first fraud occurred between 2010 and 2013, fraud exposure is *correlated* with a decline in total credit limits. Prior to the fraud event, these consumers see an increasing trend on their credit limits, however, at the event year ($j = 0$), credit limits fall relative to the control group and continue to decline in the post-fraud years. This association is consistent with lenders viewing fraud-exposed consumers as higher risk when card-present fraud losses fall on them, but we do not interpret the post-fraud decline itself as causal — selection into fraud and lender credit decisions are jointly determined. The causal content of the design lies within these same cohorts, when the calendar year reaches 2014, that is, when the chip-card technology shock arrives at each cohort's event-time horizon, we observe a sharp upward jump in credit limits relative to the never-defrauded group. Because the 2014 shock is the only systemic change arriving at that precise calendar moment, and because the jump appears cohort-by-cohort as each pre-chip cohort's event-time window reaches 2014 (at $j = +3$ for the 2011 cohort, at $j = +2$ for the 2012 cohort, and $j = +1$ for the 2013 cohort), we attribute this within-cohort credit-limit recovery to the causal effect of the more secure payment technology.

For consumers whose first fraud occurred in or after 2014, the pattern is markedly different from the start. Rather than an initial decline, credit limits rise continuously relative to the never-defrauded control group even at the fraud event year and remain elevated through

the end of the event window. Importantly, this increasing trajectory is present for every fraud cohort from 2014 onward, the lender response to fraud is uniformly positive once chip technology is in place. The monotone increase across all post-2014 cohorts corroborates the interpretation that the technology shock, rather than time trends or selection, is driving the change in credit limit.

Taken together, the cohort evidence suggests that the chip-card technology shock causally improved lender willingness to extend credit to fraud-exposed consumers, as seen both in the within-cohort jump for pre-chip victims when 2014 arrives and in the persistently rising credit limits for every post-chip cohort from the moment of their fraud onward.

## 5.3 Pooled Shock Estimate

The cohort-level event study in Section 5.2 traces credit limit trajectories indexed to each cohort's own fraud year, which ensures that within-cohort variation cleanly isolates the chip shock from the fraud event. A complementary approach is to pool all pre-chip cohorts and estimate a single difference-in-differences coefficient that measures the average effect of the 2014 chip-card shock on credit limits for fraud-exposed consumers relative to never-defrauded consumers. We implement this using the same stacked dataset but re-index time to the calendar year of the shock rather than to each cohort's fraud year.

The estimation sample consists of consumers in the pre-chip cohorts ($c \in \{2010, \ldots, 2013\}$) and the never-defrauded consumers stacked alongside them. $Treated_i$ equals one for fraud-exposed consumers and zero for never-defrauded consumers. We first estimate a calendar-year event study by constructing dummies $S_{t,j}^{2014} = \mathbf{1}[t - 2014 = j]$ for $j \in \{-4, \ldots, 4\}$, yielding the following specification:

$$Y_{i,c,t} = \alpha + \sum_{j=-4}^{4} \gamma_j \left( Treated_{i,c} \times S_{t,j}^{2014} \right) + \delta_{c,t} + \eta_{c,i} + \varepsilon_{i,c,t} \tag{5}$$

where $Y_{i,c,t}$ is the credit limit for consumer $i$ in cohort $c$'s stacked sub-sample at calendar

year $t$. The reference period is calendar year 2013 ($j = -1$, zeroed out). All other notation follows Equation (4): $\delta_{c,t}$ are cohort-by-calendar-year fixed effects and $\eta_{c,i}$ are cohort-by-consumer fixed effects. Standard errors are clustered at the cohort-by-consumer level.

The coefficients $\gamma_j$ trace the average differential in credit limits between fraud-exposed and never-defrauded consumers around the 2014 shock calendar year, absorbing all aggregate year effects separately within each cohort sub-dataset and removing time-invariant consumer heterogeneity. Because time is now indexed to 2014 uniformly across all pre-chip cohorts, $\gamma_j$ reflects the pooled response of fraud-exposed consumers relative to controls at each calendar date around the chip shock, irrespective of where each cohort's own fraud year falls in that window. As in Section 5.2, the never-defrauded consumers serve as a market benchmark rather than a matched counterfactual: the cohort-by-year fixed effects difference out the common calendar-year response, so $\gamma_j$ captures the credit-limit path of fraud-exposed consumers net of the market.

We also estimate a 2×2 version that collapses the event study into a single summary coefficient:

$$Y_{i,c,t} = \alpha + \beta \left( Treated_{i,c} \times Post_t \right) + \delta_{c,t} + \eta_{c,i} + \varepsilon_{i,c,t} \tag{6}$$

where $Post_t = \mathbf{1}[t \geq 2014]$. The coefficient $\beta$ identifies the average change in credit limits for fraud-exposed consumers relative to never-defrauded consumers after 2014, holding fixed cohort-specific year effects and consumer fixed effects. It provides a single-number summary of the aggregate credit supply response to the chip-card technology shock across all pre-chip fraud cohorts.

### 5.3.1 Results

Figure 7 presents the calendar-year event-study coefficients $\hat{\gamma}_j$ from Equation (5). The horizontal axis indexes calendar years relative to the 2014 chip-card shock, with calendar year 2013 ($j = -1$) as the omitted reference period. Starting at $j = 0$ (calendar year

2014), the differential turns positive and remains elevated through the post-2014 window: fraud-exposed consumers experience a relative increase in total credit limits compared to the never-defrauded benchmark, coinciding precisely with the arrival of chip technology.

The 2×2 collapsed estimate from Equation (6) provides a single-number summary of this shift, with the underlying event-study coefficients reported in Table 7. The coefficient $\hat{\beta}$ reported in Table 8 indicates that fraud-exposed consumers experienced an average increase in total credit limits of approximately $1,942 relative to never-defrauded consumers after 2014.

# 6 Mechanism: Repeated Fraud Risk

A natural question is why lenders extended more credit to fraud-exposed consumers once the chip-card technology arrived. We argue that a key mechanism is the reduction in the likelihood of future fraud for previously victimized consumers. Under the less secure magnetic-stripe regime, the underlying vulnerability that enabled a consumer's first fraud — easily cloned card data — remained intact afterward, leaving them at persistently elevated risk of repeat victimization. A lender rationally anticipating further fraud losses would restrict credit supply. Once chip cards became widespread after 2014, that vulnerability was closed: even a consumer who had already experienced fraud now held a card whose technology made future card-present fraud substantially harder. The same consumers who were high-risk before the shock became low-risk after it, giving lenders reason to extend credit rather than withhold it.

To examine this mechanism, we apply the same stacked cohort event-study specification (4) but replace the credit-limit outcome with an indicator for experiencing a subsequent fraud event. We focus on pre-chip cohorts (first fraud in 2010–2013) and trace the $\beta_j$ path of future fraud likelihood around each cohort's first fraud year. The key question is whether the probability of repeat fraud falls when each cohort's event-time window reaches the cal-

endar year 2014 — the same moment at which we observe the credit-limit jump in our main results.

The cohort-wise estimates confirm the mechanism; Figure 6 plots the results and Table 6 reports the underlying regression coefficients. For pre-chip cohorts, the likelihood of future fraud is elevated in the post-fraud years relative to the never-defrauded benchmark. Crucially, when the calendar year reaches 2014 for each cohort (at $j = +4$ for the 2010 cohort, $j = +3$ for the 2011 cohort, and $j = +2$ for 2012), the likelihood of future fraud drops sharply within those cohorts. This decline in repeat-fraud risk coincides precisely with the credit-limit jump documented in Section 5.2.1, linking the two results through a single channel: chip technology closed the vulnerability that made these consumers more risky, and lenders responded by relaxing credit supply. The mechanism thus operates entirely within the pre-chip cohorts i.e. the fraud-exposed consumers suddenly become less costly to serve once chip cards reduce future fraud risk, and lenders respond accordingly by extending credit limit.

To quantify these magnitudes, Table 6 shows that in the post-fraud years preceding 2014, the probability of repeat fraud for treated consumers is elevated by roughly 13 to 17 percentage points relative to the never-defrauded benchmark (0.159, 0.137, and 0.150 at $t+1$ through $t+3$ for the 2010 cohort; 0.168 at $t+2$ for the 2011 cohort; and 0.173 at $t+1$ for the 2012 cohort). Once calendar year 2014 enters each cohort's event-time window, the coefficient collapses to around 3 percentage points: 0.027 at $t+4$ for the 2010 cohort, 0.030 at $t+3$ for the 2011 cohort, and 0.035 at $t+2$ for the 2012 cohort. This implies a reduction in repeat-fraud probability of approximately 80 percent. This The timing of this collapse aligns one-for-one with the credit-limit recovery in Table 4. The 2013 cohort experiences a increase of around 4 percentage points. This is consistent with fraud vicimization increasing the rate of future fraud, even following the technology shock. However, this is dramatically smaller compaired to the roughly 15 percent increase that prior cohorts experienced following their first fraud.

# 7 Conclusion

In conclusion, our study provides new evidence on the causal impact of payment security innovation on household credit outcomes in the face of financial fraud. We leverage the 2013-2014 widespread adoption of chip-enabled credit cards as a natural experiment which led to a dramatic reduction in the incidence and recurrence of credit card fraud. This technological shift fundamentally altered the risk landscape for lenders, resulting in a substantial increase in credit supply to fraud-exposed consumers. Our stacked cohort event-study design allows us to trace each cohort's credit outcomes around its own fraud year and isolate the effect of the technology shock from the effect of the fraud event itself.

Importantly, while lenders responded to reduced fraud risk by expanding credit access, our results reveal a persistent adverse effect of fraud exposure on consumer behavior and financial distress. Fraud-exposure is correlated with higher rate of account closure, increased delinquency rate, and a reduction in number of credit inquiries, even after the implementation of enhanced payment security. These findings suggest that technological solutions, while effective in restoring lender confidence do not fully mitigate the erosion of consumer trust following fraud events.

Our results have broad implications for the design of financial regulation and consumer protection policy. While innovations such as chip-enabled cards can successfully reduce systemic risk and improve credit market functioning, restoring consumer confidence requires complementary interventions beyond technological upgrades. Future research should explore the role of financial education, targeted support, and institutional transparency in rebuilding trust among vulnerable households.

# References


Agarwal, Sumit, Souphala Chomsisengphet, and Chunlin Liu, 2010, The importance of adverse selection in the credit card market: Evidence from randomized trials of credit card solicitations, *Journal of Money, Credit and Banking* 42, 743–754.

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel, 2018, Do banks pass through credit expansions to consumers who want to borrow?, *The Quarterly Journal of Economics* 133, 129–190.

Blascak, Nathan, and Ying Lei Toh, 2022, Prior fraud exposure and precautionary credit market behavior, *Federal Reserve Bank of Kansas City Working Paper* .

Chase, 2024, What to do if you lost your credit card, https://www.chase.com/personal/credit-cards/education/basics/what-to-do-if-you-cant-find-credit-card#:~:text=Replacing%20a%20credit%20card%20won,the%20purpose%20of%20credit%20reporting.

Chava, Sudheer, Rohan Ganduri, Nikhil Paradkar, and Linghang Zeng, 2023, Shocked by bank funding shocks: Evidence from consumer credit cards, *The Review of Financial Studies* 36, 3906–3952.

Dobbie, Will, Paul Goldsmith-Pinkham, Neale Mahoney, and Jae Song, 2020, Bad credit, no problem? credit and labor market consequences of bad credit reports, *The Journal of Finance* 75, 2377–2419.

Fed, 2022, Report on the economic well-being of u.s. households in 2021, https://www.federalreserve.gov/publications/2022-economic-well-being-of-us-households-in-2021-executive-summary.htm.

Fonseca, Julia, 2023, Less mainstream credit, more payday borrowing? evidence from debt collection restrictions, *The Journal of Finance* 78, 63–103.

FTC, 2023, Consumer sentinel network data book (2022), chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/https://www.ftc.gov/system/files/ftc_gov/pdf/CSN-Data-Book-2022.pdf.

Gibbs, Christa, Benedict Guttman-Kenney, Donghoon Lee, Scott Nelson, Wilbert van der Klaauw, and Jialan Wang, 2023, Jel proposal:"consumer credit reporting data" .

Gross, Tal, Raymond Kluender, Feng Liu, Matthew J Notowidigdo, and Jialan Wang, 2021, The economic consequences of bankruptcy reform, *American Economic Review* 111, 2309–2341.

Guttman-Kenney, Benedict, 2023, Disaster flags: Credit reporting relief from natural disasters, *Available at SSRN* .

Hamdi, Naser, Ankit Kalda, and David Sovich, 2024, The costs of financial fraud victimization, *Available at SSRN 4787119* .

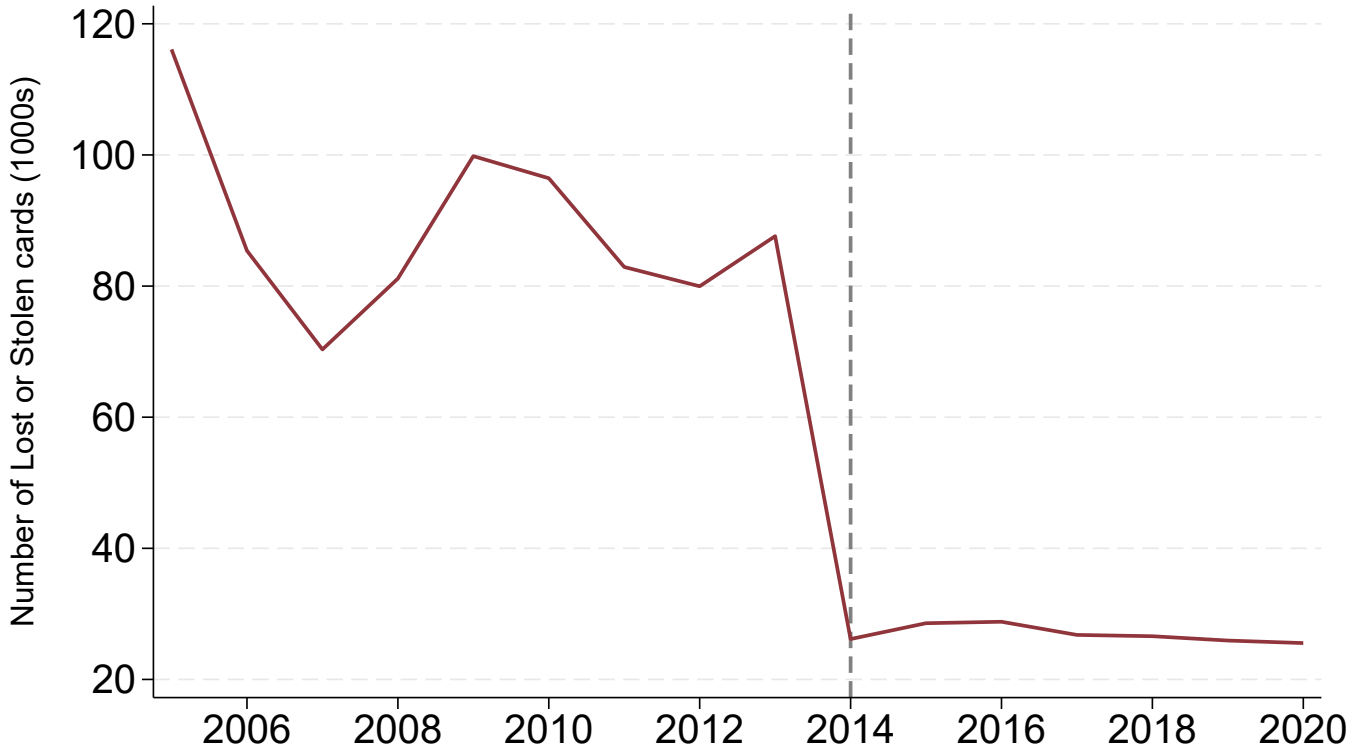
Hayashi, Fumiko, 2019, Payment card fraud rates in the united states relative to other countries after migrating to chip cards, *Economic Review* 104, 23–40.

Kahn, Charles M, José M Liñares-Zegarra, and Joanna Stavins, 2017, Are there social spillovers in consumers' security assessments of payment instruments?, *Journal of Financial Services Research* 52, 5–34.

Nilson, 2023, Nilson report card fraud worldwide 1254.

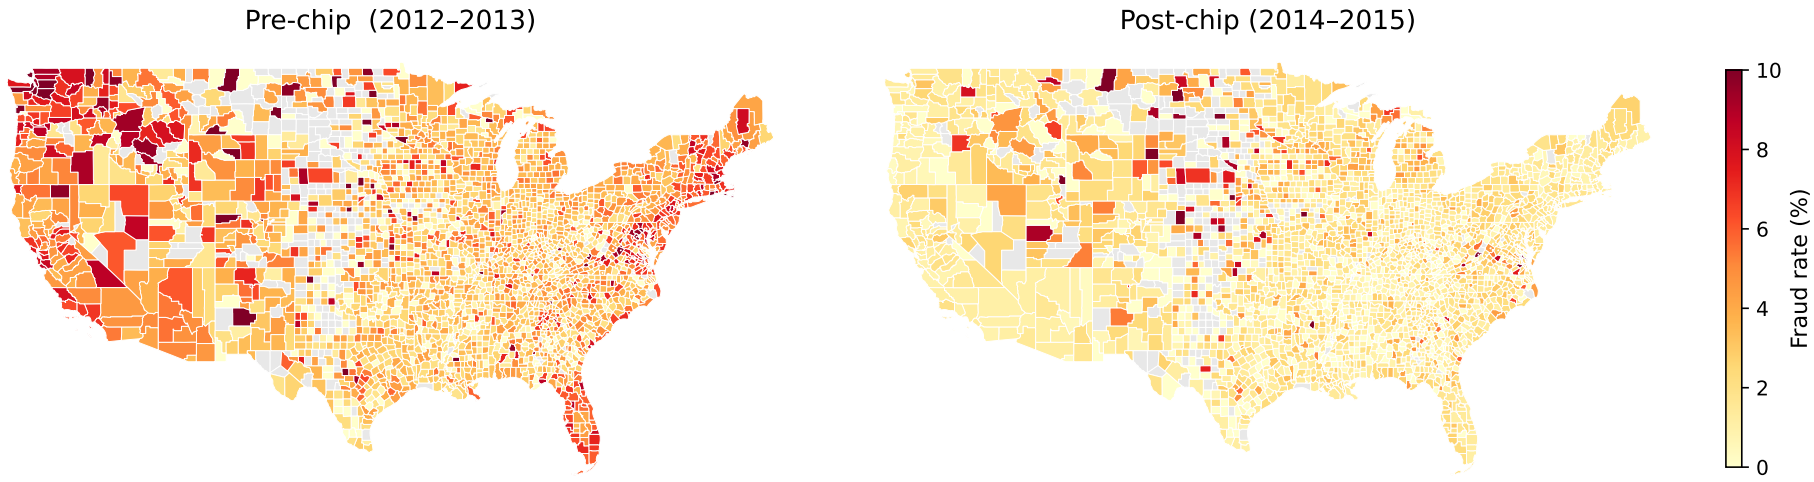
Sullivan, Richard J, 2013, The us adoption of computer-chip payment cards: Implications for payment fraud, *Economic Review-Federal Reserve Bank of Kansas City* 59.

**Figure 1:** Number of Lost or Stolen Credit Cards, 2005-2020




Notes: This figure shows the number of credit cards which were marked as closed due to lost or stolen in consumer credit reports by year, from 2005 to 2020. The gray dashed line represents widespread adoption of more secure chip enabled credit cards in the U.S.

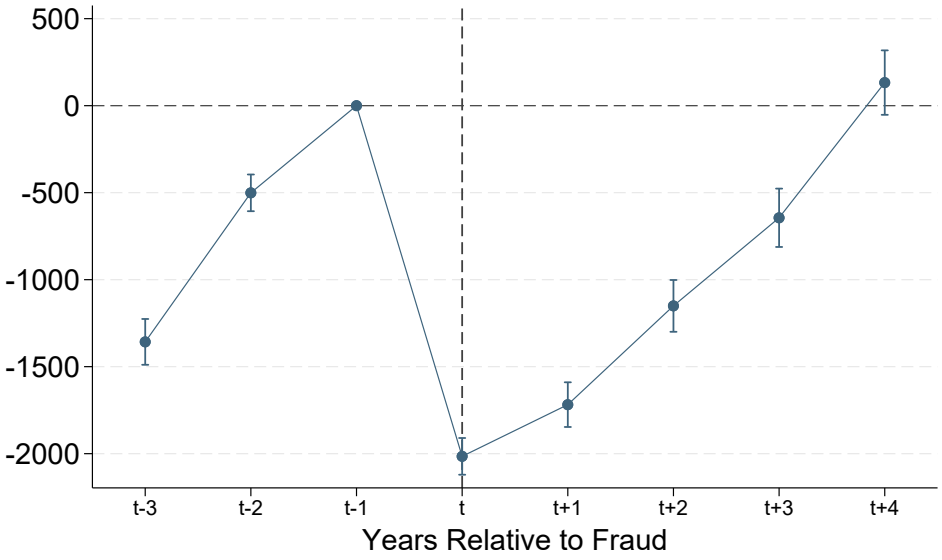
**Figure 2:** Credit Card Fraud Rate by County: Pre-Chip vs. Post-Chip



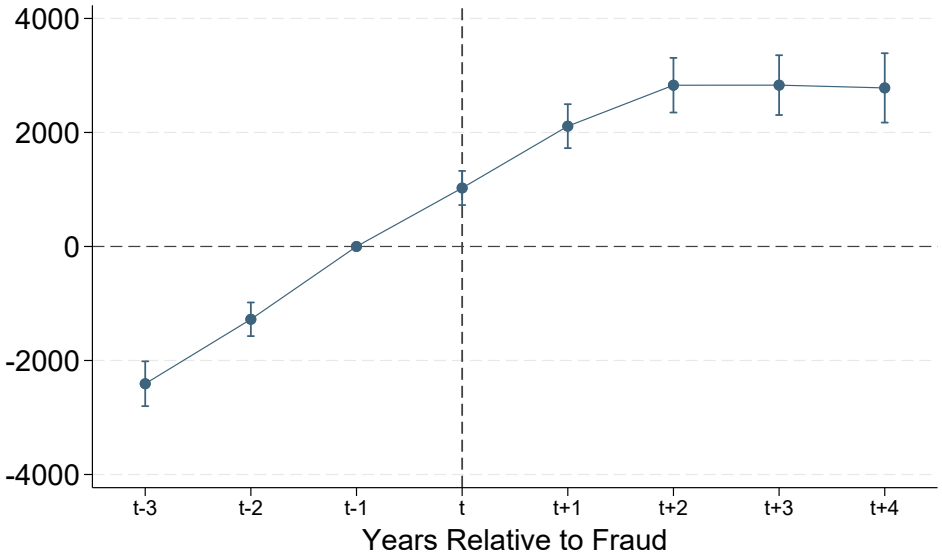

Notes: This figure plots, for each U.S. Counties, the share of sampled consumers who experience a credit card fraud event (in percent), comparing the pre-chip period (2012-2013, left panel) to the post-chip period (2014-2015, right panel).

**Figure 3:** Descriptive Event Study: Credit Limit ($)

**(a)** Pre-Chip Cohorts: Fraud Years 2010–2013



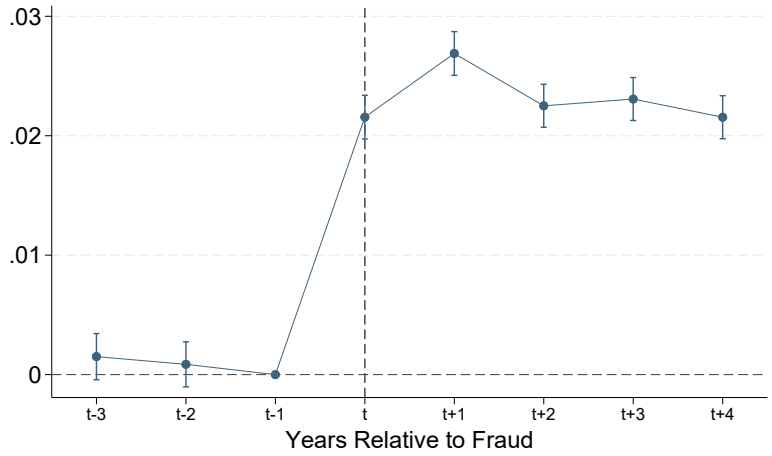

**(b)** Post-Chip Cohorts: Fraud Years 2014–2017



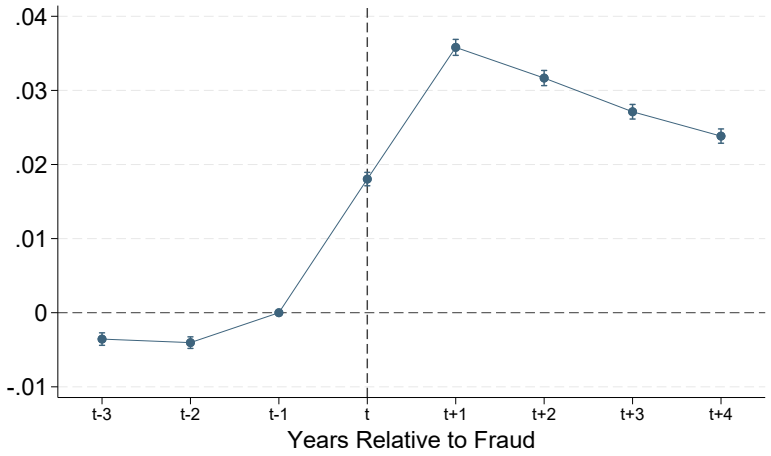

Notes: Each panel plots event-study coefficients ($\hat{\beta}_j$) and 95% confidence intervals from Equation (3), with total credit limit ($) as the outcome. Panel (a) pools consumers whose first fraud occurred between 2010 and 2013 (pre-chip); Panel (b) pools consumers whose first fraud occurred between 2014 and 2017 (post-chip). Standard errors are clustered at the consumer level.

**Figure 4:** Descriptive Event Study: Other Credit Outcomes

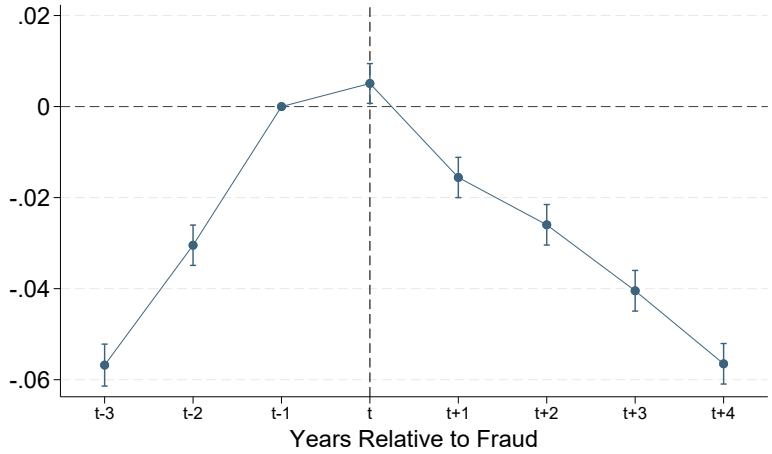
**(a)** Account Closure Rate (%)




**(b)** Default Rate (%)

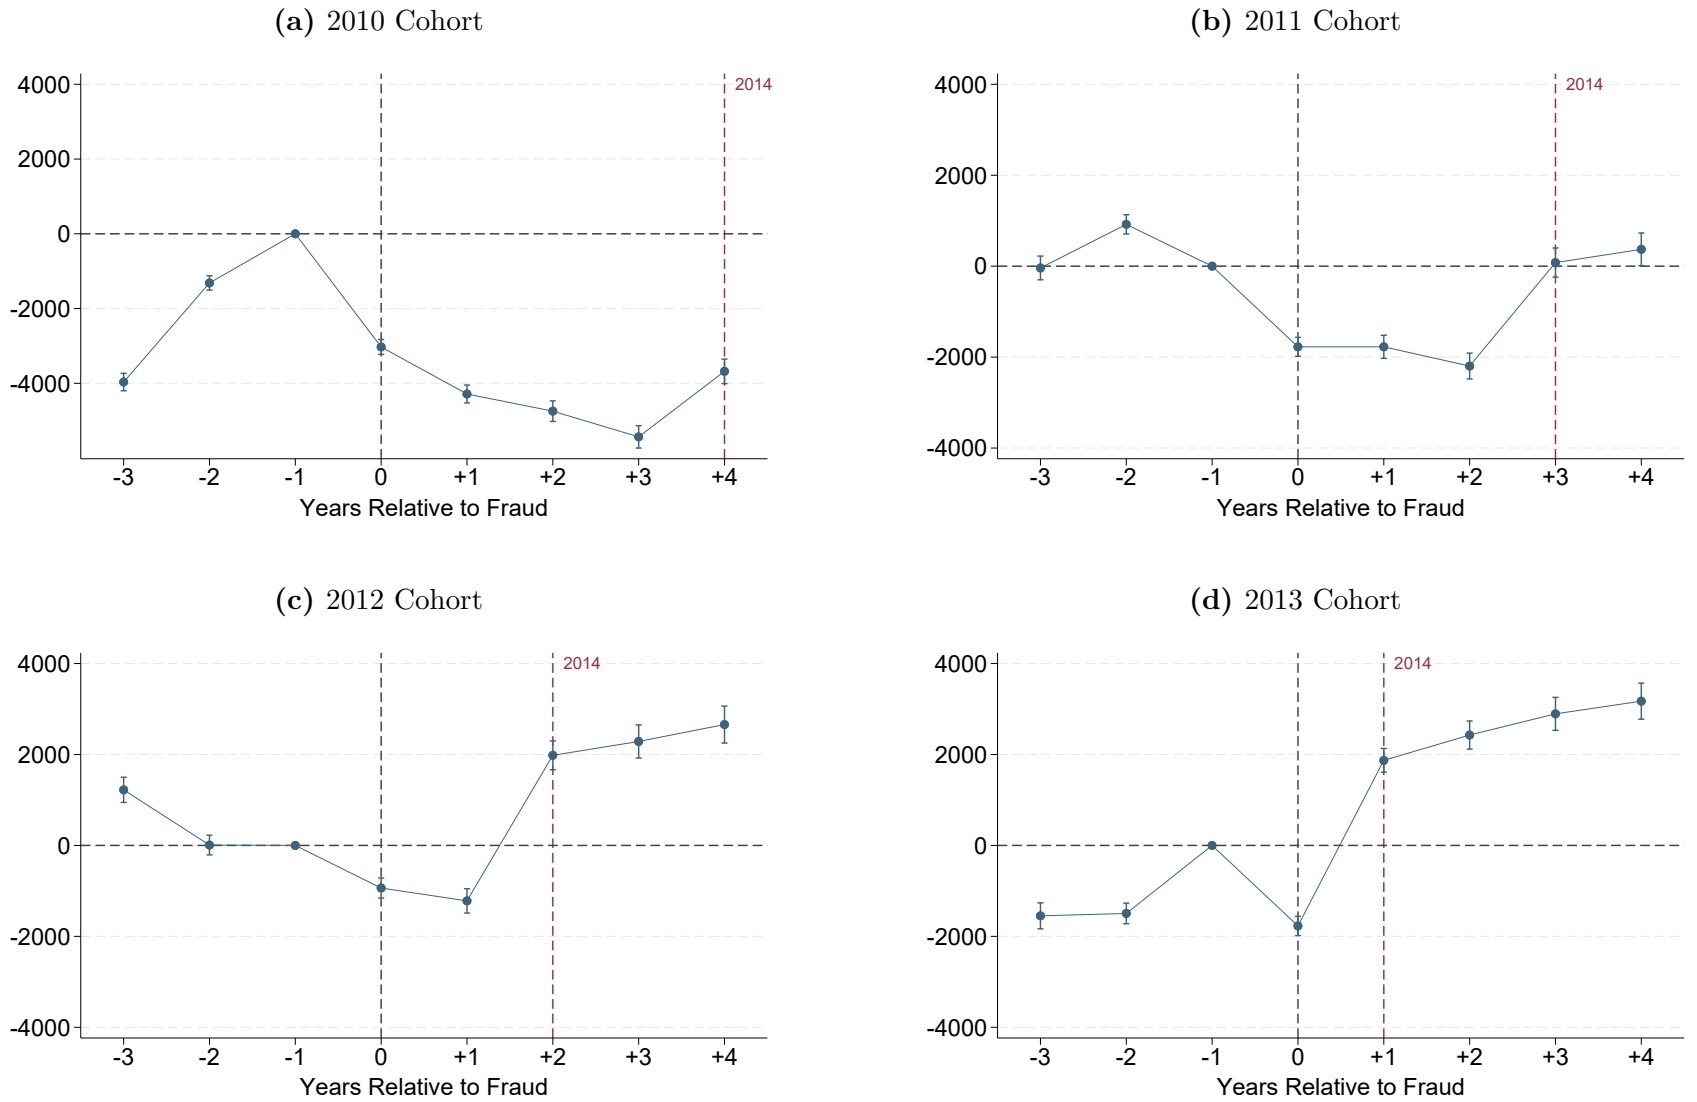



**(c)** Number of New Inquiries




Notes: Each panel plots event-study coefficients ($\hat{\beta}_j$) and 95% confidence intervals from Equation (3). Panel (a) shows account closure rate (%) as the outcome. Panel (b) shows 90 days past due default rate (%) as the outcome. Panel (b) shows number of new credit inquiries in past six months as the outcome. Standard errors are clustered at the consumer level.

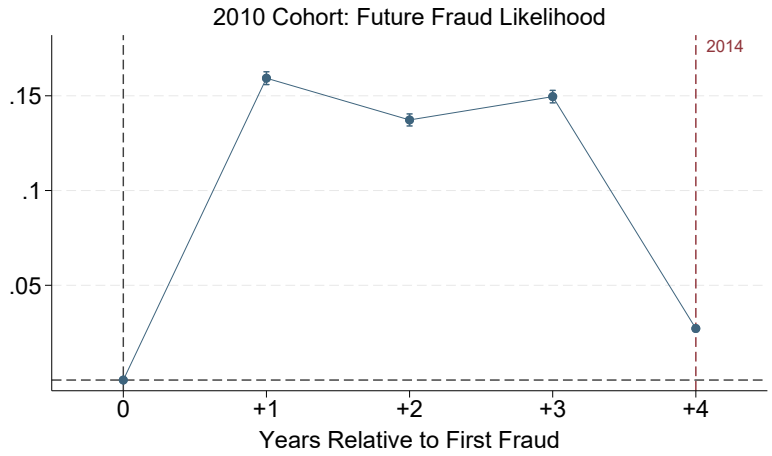
**Figure 5:** Stacked Cohort Event Study: Credit Limit ($) by Pre-Chip Fraud Cohort



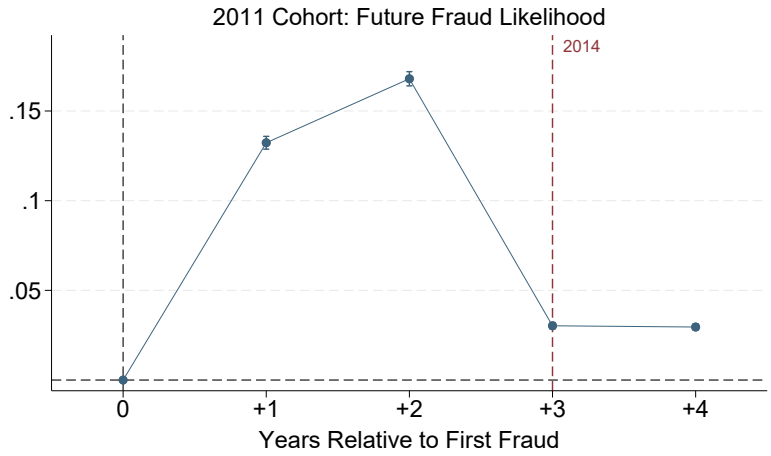

Notes: Each panel shows stacked cohort event-study coefficients ($\hat{\beta}_j$ from Equation (4)) and 95% confidence intervals for total credit limit ($), restricted to the indicated fraud cohort's sub-sample. The horizontal axis shows event time $j$ relative to the cohort's first-fraud year; $j = -1$ is the omitted reference period. The gray dashed vertical line marks fraud year. The maroon dashed vertical line marks calendar year 2014 in each cohort's event-time window. Standard errors are clustered at the cohort-by-consumer level.

**Figure 6:** Stacked Cohort Event Study: Future Fraud Likelihood by Pre-Chip Fraud Cohort

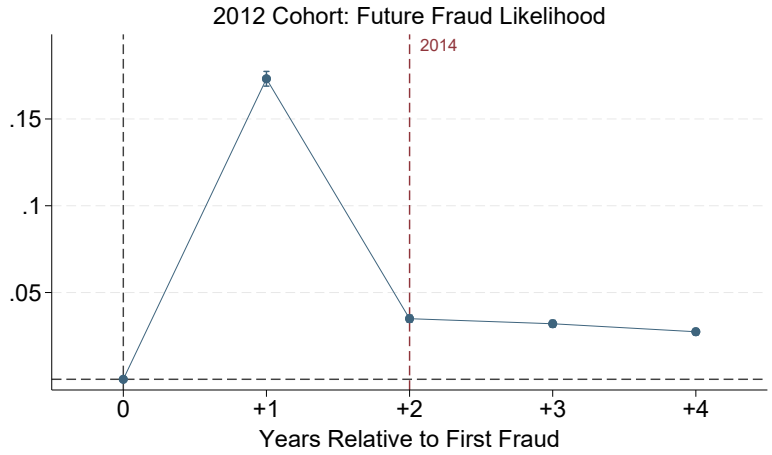
**(a)** 2010 Cohort



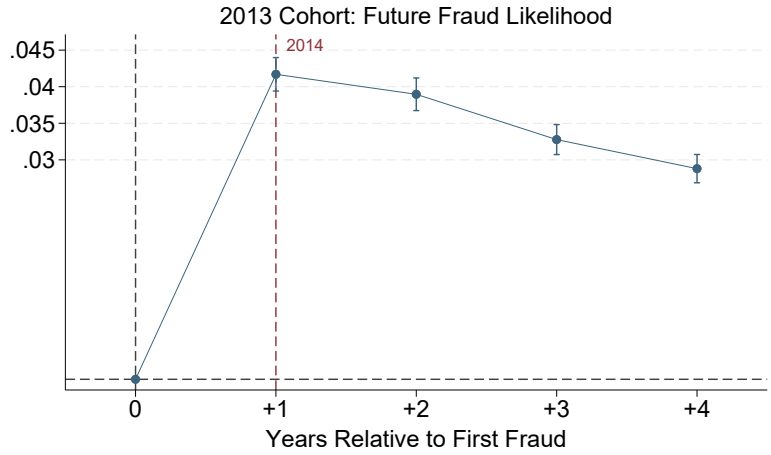

**(b)** 2011 Cohort




**(c)** 2012 Cohort

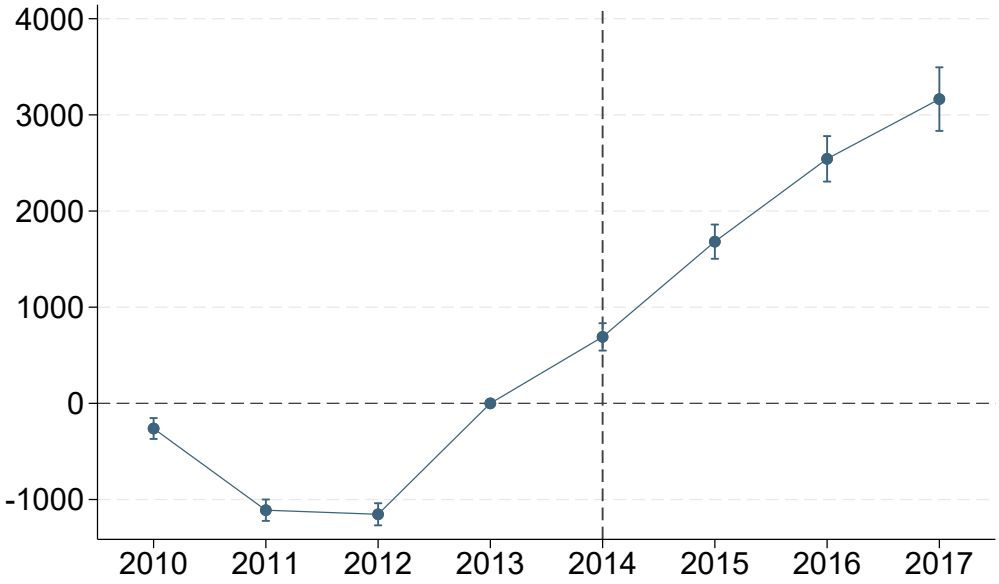



**(d)** 2013 Cohort




Notes: Each panel shows stacked cohort event-study coefficients and 95% confidence intervals for a future-fraud indicator, using the same specification as Equation (4) with the outcome replaced by an indicator equal to one if the consumer experiences a subsequent fraud at event time $j > 0$. The outcome is zero by construction at $j \leq 0$; the reference period is $j = 0$. The maroon dashed vertical line marks calendar year 2014 in each cohort's event-time window. Standard errors are clustered at the cohort-by-consumer level.

**Figure 7:** Pooled Chip-Shock Event Study: Credit Limit ($)




Notes: This figure plots calendar-year event-study coefficients ($\hat{\gamma}_j$) and 95% confidence intervals from Equation (5), with total credit limit ($) as the outcome. The horizontal axis indexes calendar years relative to the 2014 chip-card shock ($j = 0$); the omitted reference period is calendar year 2013 ($j = -1$). The sample consists of all pre-chip fraud cohorts (first fraud in 2010–2013) and never-defrauded consumers stacked alongside them. Cohort-by-calendar-year and cohort-by-consumer fixed effects are included. Standard errors are clustered at the cohort-by-consumer level.

**Table 1:** Summary Statistics, 2010-2017

| | (1) | (2) | (3) |
|---|---|---|---|
| | Mean | St. Dev. | Med. |
| **A. Demographics** | | | |
| Female (%) | 51.77 | 49.97 | 100.00 |
| Age (years) | 49.34 | 17.88 | 49.00 |
| Income ($1,000) | 48.19 | 26.88 | 41.00 |
| **B. Credit Supply, Utilization, and Demand** | | | |
| Credit Score | 696.40 | 101.79 | 707.00 |
| Total Credit Card Limit ($1,000) | 32.12 | 45.65 | 13.73 |
| Total Balance Amount ($1,000) | 4.32 | 9.59 | 0.68 |
| Balance Amount on Credit Cards 90+ Days Past Due | 0.01 | 0.25 | 0.00 |
| Credit Card Utilization (%) | 27.18 | 32.79 | 12.00 |
| Number of Credit Cards | 4.20 | 4.64 | 3.00 |
| Number of Credit Cards Opened in Past 6 Months | 3.13 | 3.54 | 2.00 |
| Total Number of Inquiries | 0.42 | 0.74 | 0.00 |
| **C. Default Risk and Accounts Churn Risk** | | | |
| Credit Card Closure (%) | 5.60 | 22.99 | 0.00 |
| Number of Credit Cards 90+ Days Past Due | 0.08 | 0.55 | 0.00 |
| Credit Card Default (%) | 3.61 | 18.66 | 0.00 |
| **D. Fraud Experience** | | | |
| Fraud-Affected Consumers (%) | 12.22 | | |
| Observations | 16,487,864 | | |

Notes: This table presents summary statistics from the Gies Consumer and Small Business Credit Panel, 2010–2017. The unit of observation is the consumer-year. Fraud-Affected Consumers reports the share of unique consumers who experienced at least one credit card fraud event during 2010–2017. The sample is limited to consumers with a non-missing credit score who had at least one credit card account in the sample period.

**Table 2:** Summary Statistics: Pre-Chip Fraud Cohorts vs. Never-Fraud Consumers

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | Treated | | Control | |
| | Mean | St. Dev. | Mean | St. Dev. |
| **A. Demographics** | | | | |
| Female (%) | 51.76 | 49.97 | 51.36 | 49.98 |
| Age (years) | 47.68 | 16.11 | 48.42 | 18.19 |
| Income ($1,000) | 57.78 | 30.14 | 44.96 | 24.37 |
| **B. Credit Supply, Utilization, and Demand** | | | | |
| Credit Score | 742.47 | 73.98 | 684.66 | 104.49 |
| Total Credit Card Limit ($1,000) | 56.82 | 56.32 | 25.92 | 39.99 |
| Total Balance Amount ($1,000) | 8.57 | 14.64 | 3.38 | 8.01 |
| Balance Amount on Credit Cards 90+ Days Past Due | 0.01 | 0.16 | 0.01 | 0.20 |
| Credit Card Utilization (%) | 24.96 | 29.26 | 28.15 | 33.92 |
| Number of Credit Cards | 6.83 | 5.39 | 3.49 | 4.06 |
| Number of Credit Cards Opened in Past 6 Months | 5.33 | 3.91 | 2.54 | 3.07 |
| Total Number of Inquiries | 0.48 | 0.77 | 0.42 | 0.74 |
| **C. Default Risk and Accounts Churn Risk** | | | | |
| Credit Card Closure (%) | 11.42 | 31.81 | 3.79 | 19.09 |
| Number of Credit Cards 90+ Days Past Due | 0.03 | 0.32 | 0.06 | 0.46 |
| Credit Card Default (%) | 1.54 | 12.32 | 3.27 | 17.79 |
| **D. Fraud Experience** | | | | |
| Ever Experienced Fraud (%) | 100.00 | 0.00 | 0.00 | 0.00 |
| Observations | 149,816 | | 1,452,604 | |

Notes: Treated consumers are those whose first credit card fraud occurred between 2010 and 2013 (pre-chip cohorts), observed one year before their fraud event. Control consumers never experienced credit card fraud, observed in 2013. The unit of observation is the consumer. Sample drawn from the Gies Consumer and Small Business Credit Panel (GCCP).

**Table 3:** Pre- vs. Post-Chip Fraud Cohort Comparison

| | Pre-Chip (2010–2013) | | Post-Chip (2014–2017) | | |
|---|---|---|---|---|---|
| | Mean | (SD) | Mean | (SD) | Diff |
| Female (%) | 51.76 | (49.97) | 55.89 | (49.65) | 4.14*** |
| Age (years) | 47.68 | (16.11) | 49.26 | (16.95) | 1.59*** |
| Income ($1,000) | 57.78 | (30.14) | 51.54 | (24.41) | –6.24*** |
| Credit Score | 742.47 | (73.98) | 739.98 | (74.59) | –2.50*** |
| Total Credit Card Limit ($1,000) | 56.82 | (56.32) | 44.37 | (48.65) | –12.45*** |
| Total Balance Amount ($1,000) | 8.57 | (14.64) | 6.73 | (10.63) | –1.84*** |
| Delinquent CC Balance ($1,000) | 0.01 | (0.16) | 0.00 | (0.07) | –0.00*** |
| Credit Card Utilization (%) | 24.96 | (29.26) | 25.90 | (27.68) | 0.94*** |
| Number of Credit Cards | 6.83 | (5.39) | 6.15 | (5.12) | –0.68*** |
| New Credit Cards (6m) | 5.33 | (3.91) | 5.55 | (4.26) | 0.22*** |
| Number of Inquiries (6m) | 0.48 | (0.77) | 0.56 | (0.84) | 0.08*** |
| Credit Card Closure (%) | 11.42 | (31.81) | 3.70 | (18.88) | –7.72*** |
| Revolving Balance 90+ DPD | 0.03 | (0.32) | 0.02 | (0.28) | –0.01*** |
| Credit Card Default (%) | 1.54 | (12.32) | 1.27 | (11.21) | –0.27*** |
| Observations | 149,816 | | 45,319 | | |

Notes: Sample consists of fraud-exposed consumers with first fraud in 2010–2013 (Pre-Chip) or 2014–2017 (Post-Chip), each observed one year before their fraud event (t = –1). Money variables rescaled to $1,000; binary variables to percent. Diff = Post−Pre. Standard errors clustered at the consumer level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

**Table 4:** Fraud Cohort D-i-D Analysis: Credit Limit ($) for Pre-Chip Cohorts (2010–2013)

| | (1) 2010 Cohort | (2) 2011 Cohort | (3) 2012 Cohort | (4) 2013 Cohort |
|---|---|---|---|---|
| t-3 | -3965.864*** | -38.937 | 1222.009*** | -1548.639*** |
| | (119.458) | (132.935) | (141.612) | (145.973) |
| t-2 | -1315.144*** | 920.050*** | 8.673 | -1495.768*** |
| | (97.646) | (107.838) | (110.626) | (115.521) |
| t | -3027.600*** | -1773.258*** | -936.803*** | -1770.345*** |
| | (101.459) | (105.941) | (112.020) | (108.587) |
| t+1 | -4285.249*** | -1773.595*** | -1219.285*** | 1870.791*** |
| | (121.863) | (129.426) | (136.454) | (133.217) |
| t+2 | -4745.218*** | -2196.462*** | 1982.019*** | 2427.664*** |
| | (140.357) | (145.229) | (161.582) | (157.930) |
| t+3 | -5433.070*** | 78.927 | 2286.623*** | 2894.070*** |
| | (152.727) | (163.557) | (185.600) | (184.986) |
| t+4 | -3680.567*** | 370.401** | 2658.489*** | 3172.711*** |
| | (167.573) | (183.453) | (207.419) | (202.303) |
| Observations | 10,922,702 | 11,027,428 | 11,199,717 | 11,384,415 |
| R-squared | 0.904 | 0.904 | 0.904 | 0.906 |

Notes: This table presents stacked DiD event study coefficients ($\hat{\beta}_j$ from Equation (4)) for pre-chip fraud cohorts (first fraud 2010–2013). Each column restricts the stacked dataset to the indicated cohort's sub-sample (treated + never-fraud controls). The omitted reference period is $t-1$. Standard errors (in parentheses) are clustered at the cohort-by-consumer level. */**/**** denotes significance at the 10%, 5%, and 1% levels.

**Table 5:** Fraud Cohort D-i-D Analysis: Credit Limit ($) for Post-Chip Cohorts (2014–2017)

| | (1) 2014 Cohort | (2) 2015 Cohort | (3) 2016 Cohort | (4) 2017 Cohort |
|---|---|---|---|---|
| t-3 | -2407.676*** | -2731.877*** | -3685.340*** | -4058.220*** |
| | (200.662) | (173.368) | (196.830) | (218.978) |
| t-2 | -1277.538*** | -1743.508*** | -2457.635*** | -2135.932*** |
| | (150.598) | (142.276) | (164.655) | (179.166) |
| t | 1025.850*** | 1872.789*** | 1666.220*** | 1292.091*** |
| | (152.783) | (156.947) | (169.687) | (171.915) |
| t+1 | 2110.623*** | 2817.589*** | 2629.265*** | 2064.422*** |
| | (196.649) | (205.946) | (205.786) | (197.415) |
| t+2 | 2828.915*** | 3236.228*** | 2830.832*** | 1829.347*** |
| | (244.422) | (236.292) | (242.641) | (254.051) |
| t+3 | 2830.667*** | 3246.650*** | 1936.033*** | 1775.599*** |
| | (267.552) | (272.506) | (279.852) | (292.339) |
| t+4 | 2781.906*** | 2482.982*** | 1560.292*** | 1102.399*** |
| | (310.300) | (304.200) | (323.250) | (303.496) |
| Observations | 11,427,434 | 11,656,785 | 11,861,274 | 12,059,764 |
| R-squared | 0.907 | 0.905 | 0.898 | 0.896 |

Notes: This table presents stacked DiD event study coefficients ($\hat{\beta}_j$ from Equation (4)) for post-chip fraud cohorts (first fraud 2014–2017). Each column restricts the stacked dataset to the indicated cohort's sub-sample (treated + never-fraud controls). The omitted reference period is $t-1$. Standard errors (in parentheses) are clustered at the cohort-by-consumer level. */**/**** denotes significance at the 10%, 5%, and 1% levels.

**Table 6:** Mechanism: Future Fraud Likelihood by Pre-Chip Cohort

| | (1) 2010 Cohort | (2) 2011 Cohort | (3) 2012 Cohort | (4) 2013 Cohort |
|---|---|---|---|---|
| t+1 | 0.159*** | 0.132*** | 0.173*** | 0.042*** |
| | (0.002) | (0.002) | (0.002) | (0.001) |
| t+2 | 0.137*** | 0.168*** | 0.035*** | 0.039*** |
| | (0.002) | (0.002) | (0.001) | (0.001) |
| t+3 | 0.150*** | 0.030*** | 0.032*** | 0.033*** |
| | (0.002) | (0.001) | (0.001) | (0.001) |
| t+4 | 0.027*** | 0.030*** | 0.027*** | 0.029*** |
| | (0.001) | (0.001) | (0.001) | (0.001) |
| Observations | 10,922,702 | 11,027,428 | 11,199,717 | 11,384,415 |
| R-squared | 0.270 | 0.256 | 0.251 | 0.210 |

Notes: Each column restricts the stacked dataset to the indicated pre-chip fraud cohort (treated units + never-defrauded controls). The outcome is an indicator equal to one if the consumer experiences any fraud at that event time, conditional on having experienced a first fraud at = 0. The reference period is = 0; the outcome is zero by construction at $\leq 0$ for all observations. Standard errors (in parentheses) are clustered at the cohort-by-person level. */**/*** denotes significance at the 10%, 5%, and 1% levels.

**Table 7:** Pooled Chip-Shock Event Study: Credit Limit

| | (1)<br>Total Credit Card Limit ($) |
|---|---|
| t-4 | -262*** |
| | (55) |
| t-3 | -1111*** |
| | (57) |
| t-2 | -1154*** |
| | (59) |
| t-1 | 0 |
| | (.) |
| t | 691*** |
| | (73) |
| t+1 | 1681*** |
| | (91) |
| t+2 | 2543*** |
| | (121) |
| t+3 | 3164*** |
| | (169) |
| Observations | 44,534,262 |
| R-squared | 0.904 |

Notes: Calendar-year event study indexed to the 2014 chip-card shock (Equation (5)). Treated = first fraud in 2010–2013; controls = never-fraud consumers (unmatched stacked). Reference period: calendar year 2013 (indicator omitted). Cohort-by-calendar-year and cohort-by-person fixed effects included. Standard errors in parentheses clustered at the person level. * p<0.10 ** p<0.05 *** p<0.01.

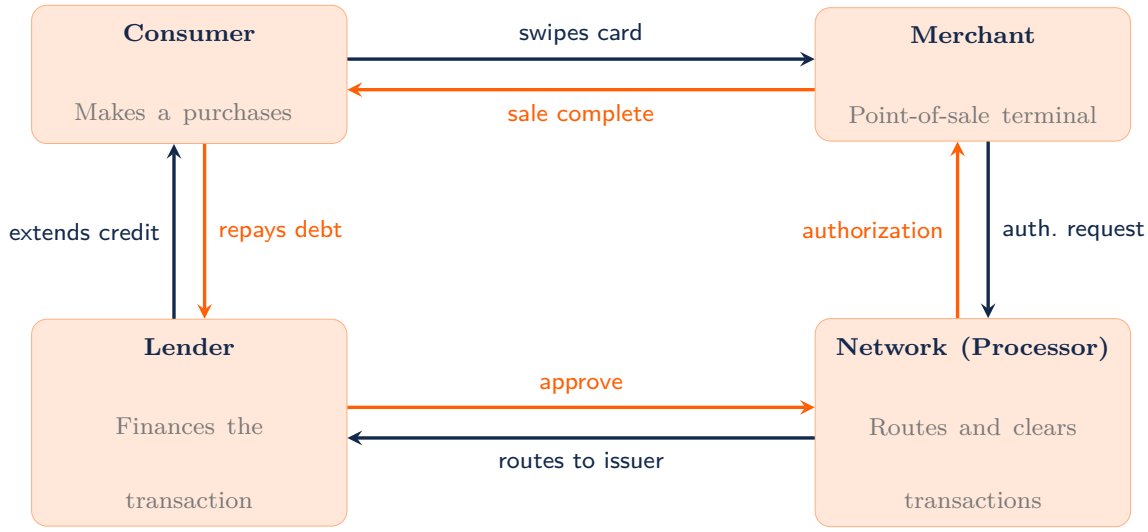
**Table 8:** Pooled Chip-Shock 2×2 DiD: Credit Limit

| | (1)<br>Total Credit Card Limit ($) |
|---|---|
| Treated × Post-Chip | 1942*** |
| | (70) |
| Observations | 44,534,262 |
| R-squared | 0.904 |

Notes: Collapsed 2×2 difference-in-differences (Equation (6)), where Post equals one for calendar years 2014 and later. Treated = first fraud in 2010–2013; controls = never-fraud consumers (unmatched stacked). Cohort-by-calendar-year and cohort-by-person fixed effects included. Standard errors in parentheses clustered at the cohort-by-person level. * p<0.10 ** p<0.05 *** p<0.01.

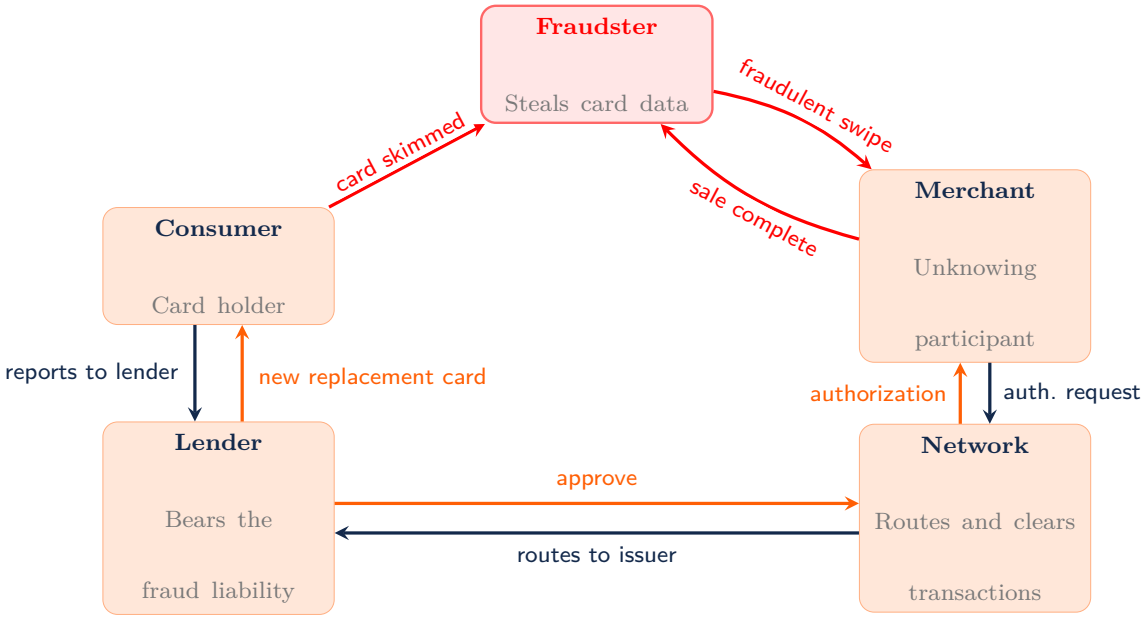
# A Additional Figures

**Figure A.1:** A Typical Credit Card Transaction



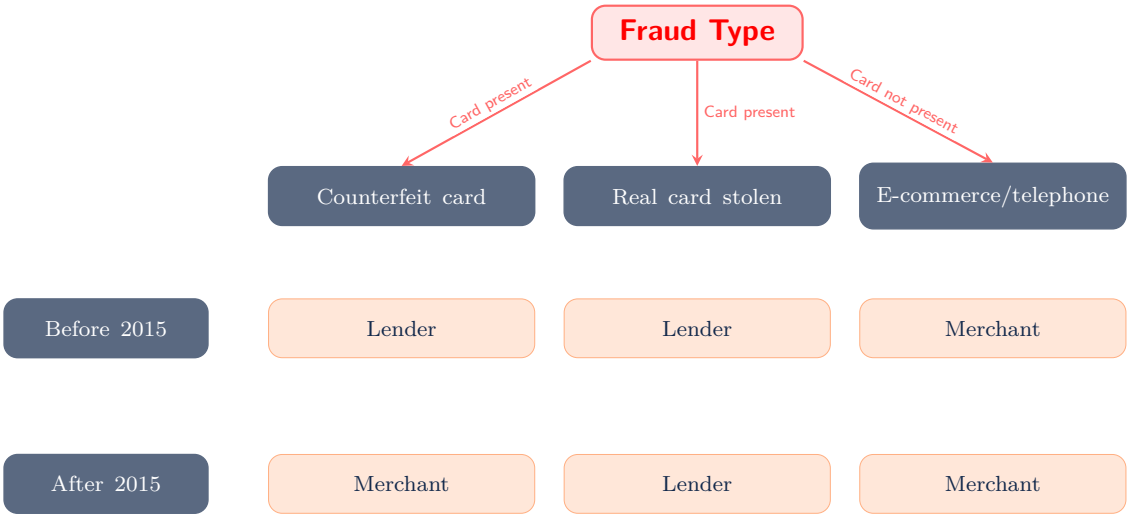

Notes: This figure illustrates the parties involved in a typical credit card transaction: the consumer, the merchant, the payment network (processor), and the lender. Blue arrows show the transaction flow; orange arrows show the return flows. The consumer swipes their card at the merchant's terminal; the merchant routes an authorization request through the network to the lender, which approves the transaction; the lender extends credit to the consumer, who subsequently repays the debt.

**Figure A.2:** What Happens During a Credit Card Fraud?




Notes: This figure illustrates how a card-present credit card fraud occurs. A fraudster skims the consumer's card data and uses it to make fraudulent purchases at a merchant. The transaction is authorized through the payment network and approved by the lender, who bears the cost of fraud. Upon reporting the fraud, the lender closes the compromised card — recording a "lost or stolen" flag in the consumer's credit report — and issues a replacement card.

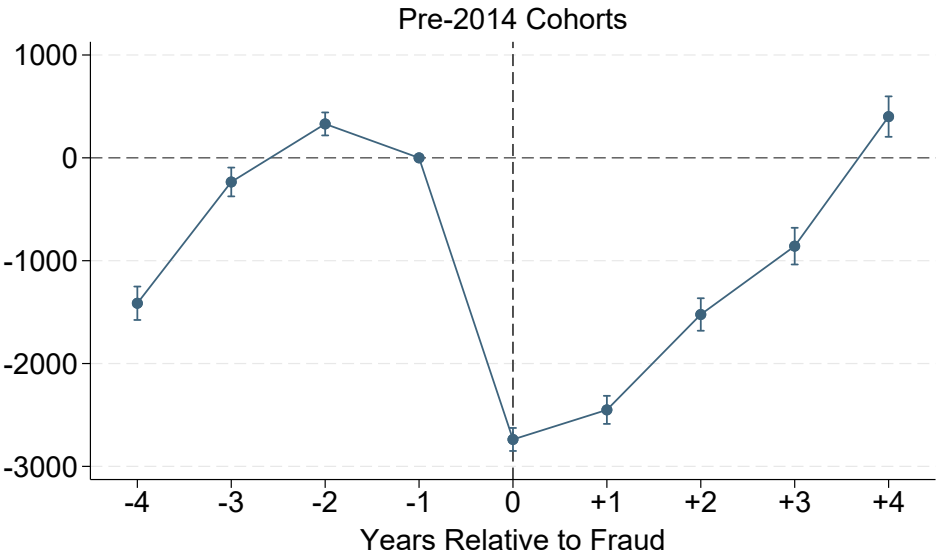
**Figure A.3:** Fraud Liability: Who Bears the Cost?



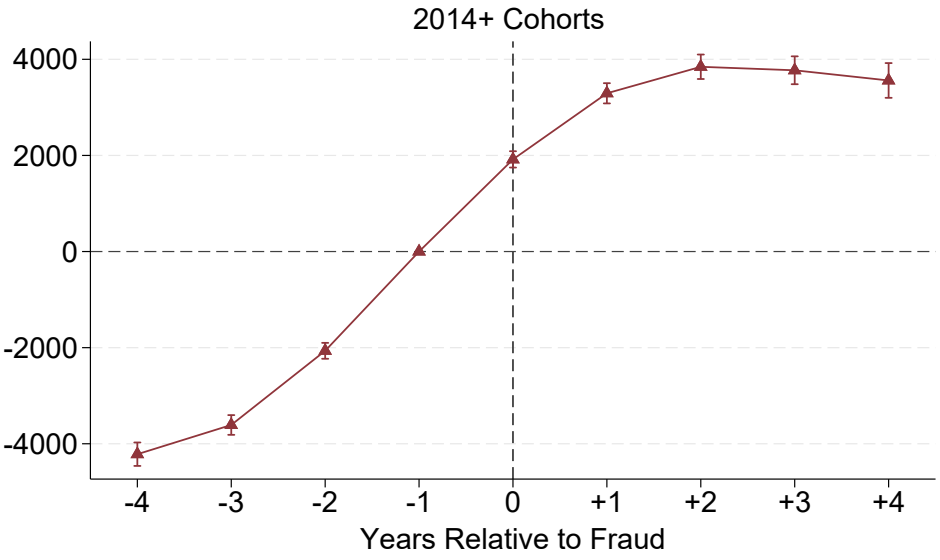

Notes: This figure shows who bears the cost of a fraudulent transaction for each combination of fraud type and time period. For card-present fraud with a counterfeit card, liability shifted from the lender to the merchant in October 2015: once chip-enabled terminals were in place, the chip would have prevented the counterfeit, so a merchant accepting a swipe bears the loss. For card-present fraud with a real stolen card, lender liability is unchanged, since the chip was used correctly and the fraud reflects physical card theft. For card-not-present fraud (e-commerce, telephone), merchant liability applies in both periods.

**Figure A.4:** Credit Limit ($) of Consumers with Fraud Experience

**(a)** Fraud Years 2010-2013



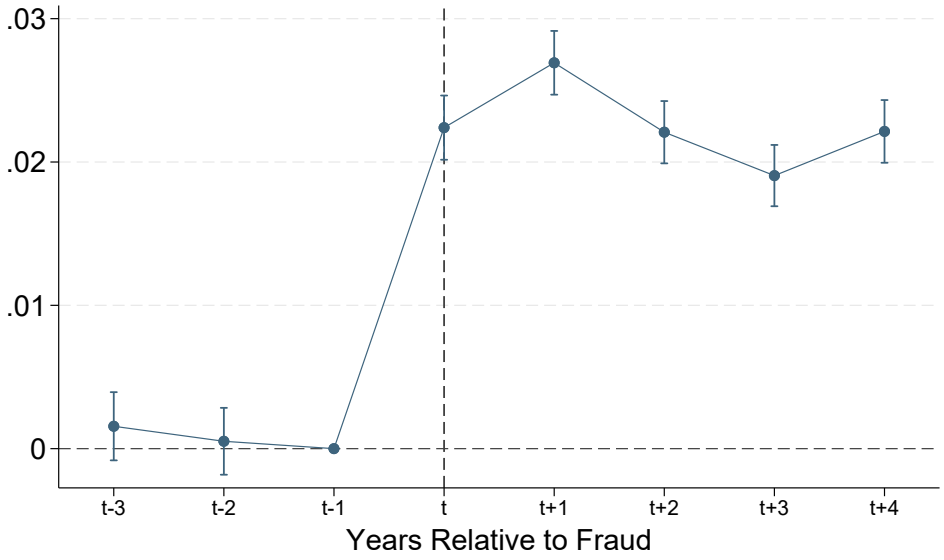

**(b)** Fraud Years 2014-2016



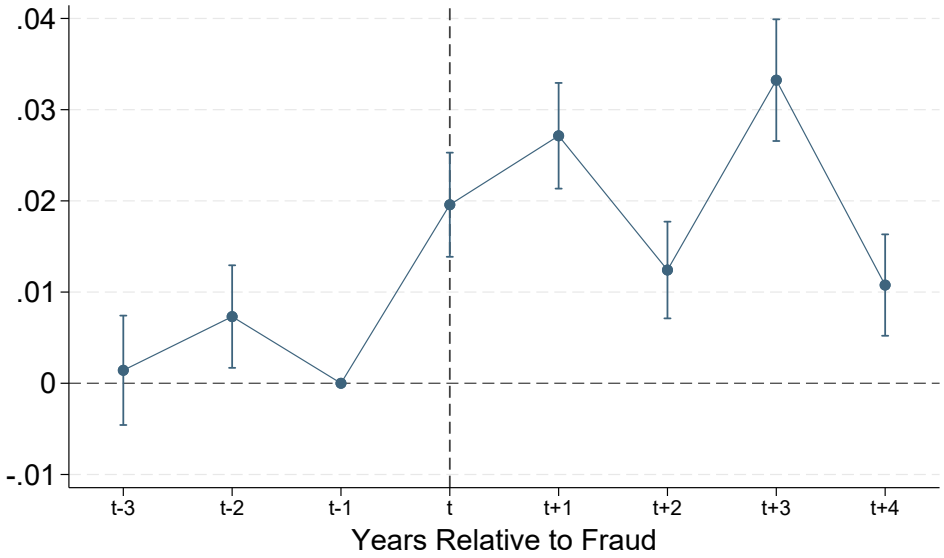

Notes: This figure plots raw event-study coefficients (and 95% confidence intervals) from consumer fixed-effects regressions of total credit limit, where $t = 0$ denotes the calendar year of a consumer's first fraud. The sample is restricted to fraud-exposed consumers only. Panel (a) shows cohorts whose first fraud occurred between 2010 and 2013 (pre-chip); Panel (b) shows cohorts whose first fraud occurred between 2014 and 2017 (post-chip).

**Figure A.5:** Descriptive Event Study: Credit Card Account Closure (%)

**(a)** Pre-Chip Cohorts: Fraud Years 2010–2013



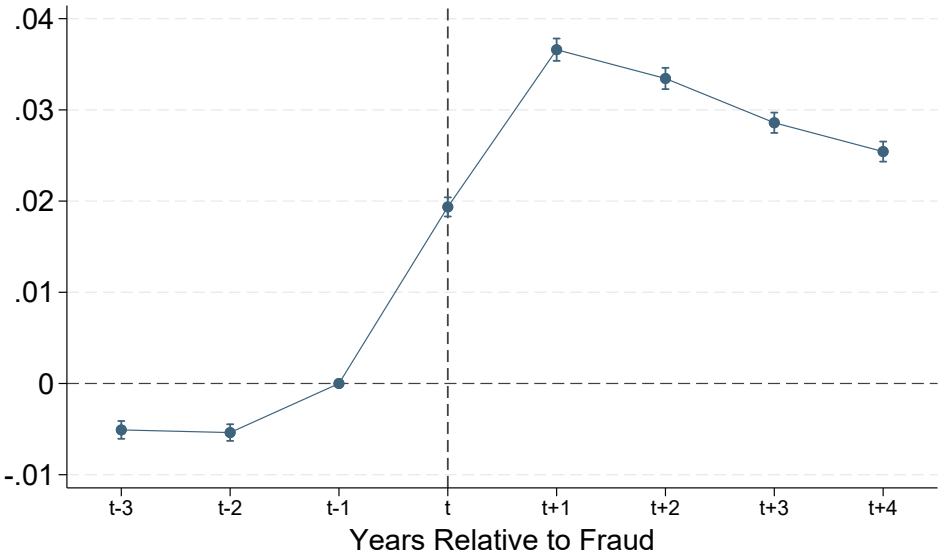

**(b)** Post-Chip Cohorts: Fraud Years 2014–2017



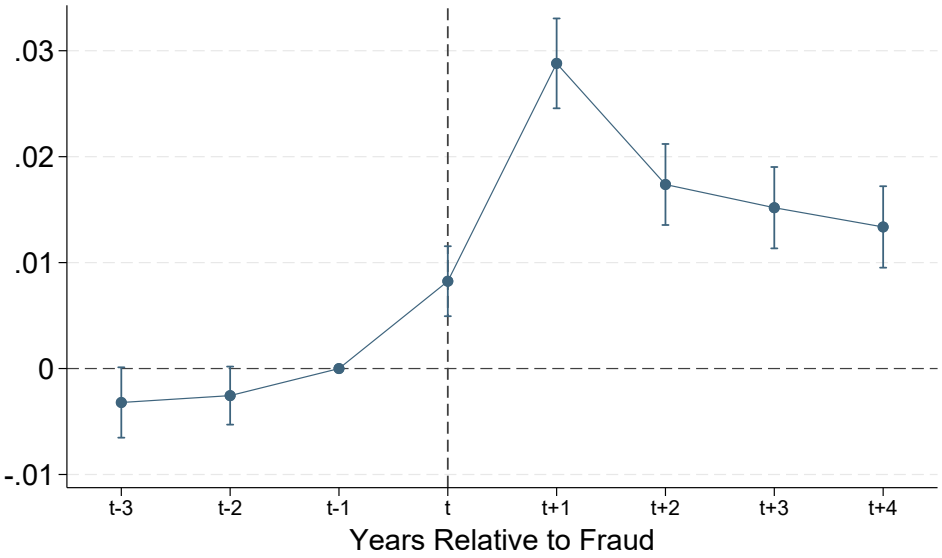

Notes: Each panel plots event-study coefficients ($\hat{\beta}_j$) and 95% confidence intervals from Equation (3), with account closure as the outcome. Panel (a) pools consumers whose first fraud occurred between 2010 and 2013 (pre-chip); Panel (b) pools consumers whose first fraud occurred between 2014 and 2017 (post-chip). Standard errors are clustered at the consumer level.

**Figure A.6:** Descriptive Event Study: 90-Day Delinquency Rate (%)

**(a)** Pre-Chip Cohorts: Fraud Years 2010–2013



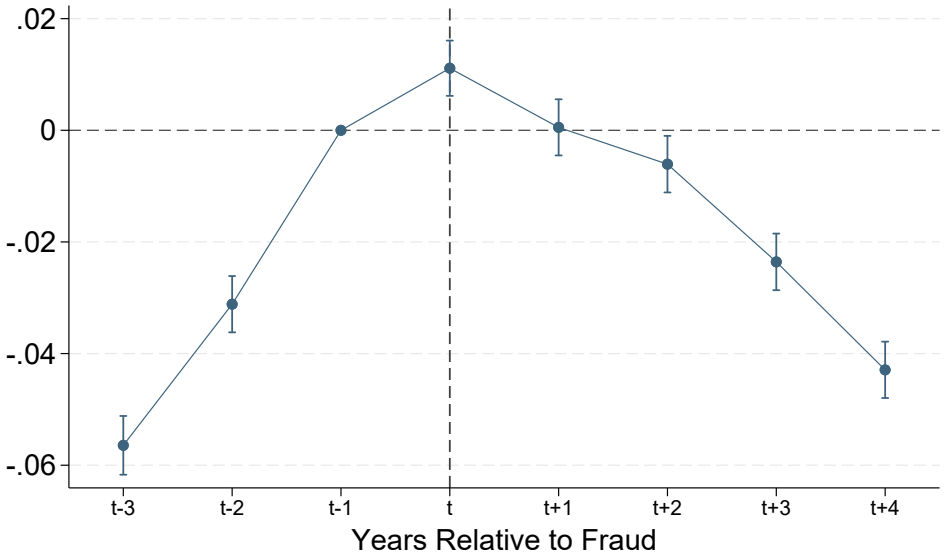

**(b)** Post-Chip Cohorts: Fraud Years 2014–2017



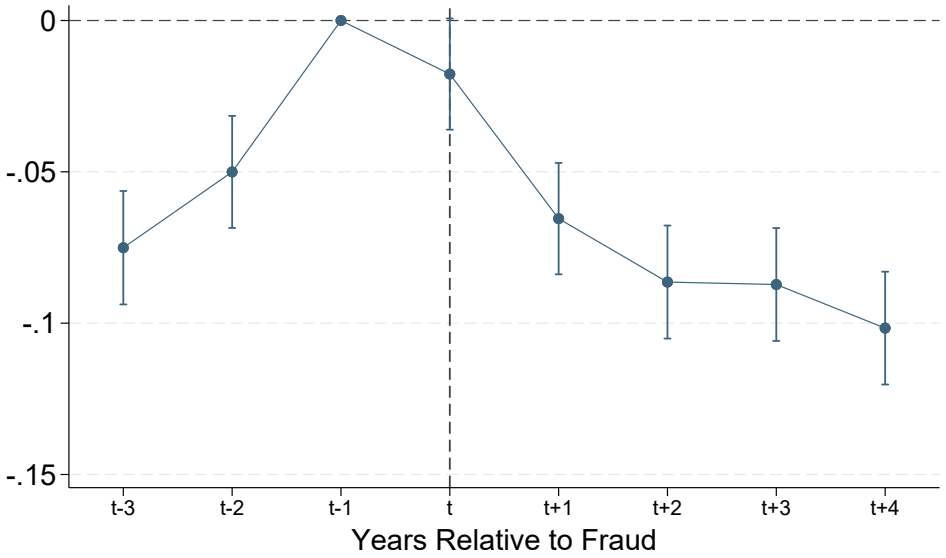

Notes: Each panel plots event-study coefficients ($\hat{\beta}_j$) and 95% confidence intervals from Equation (3), with an indicator for 90-day delinquency as the outcome. Panel (a) pools consumers whose first fraud occurred between 2010 and 2013 (pre-chip); Panel (b) pools consumers whose first fraud occurred between 2014 and 2017 (post-chip). Standard errors are clustered at the consumer level.

**Figure A.7:** Descriptive Event Study: New Credit Inquiries

**(a)** Pre-Chip Cohorts: Fraud Years 2010–2013

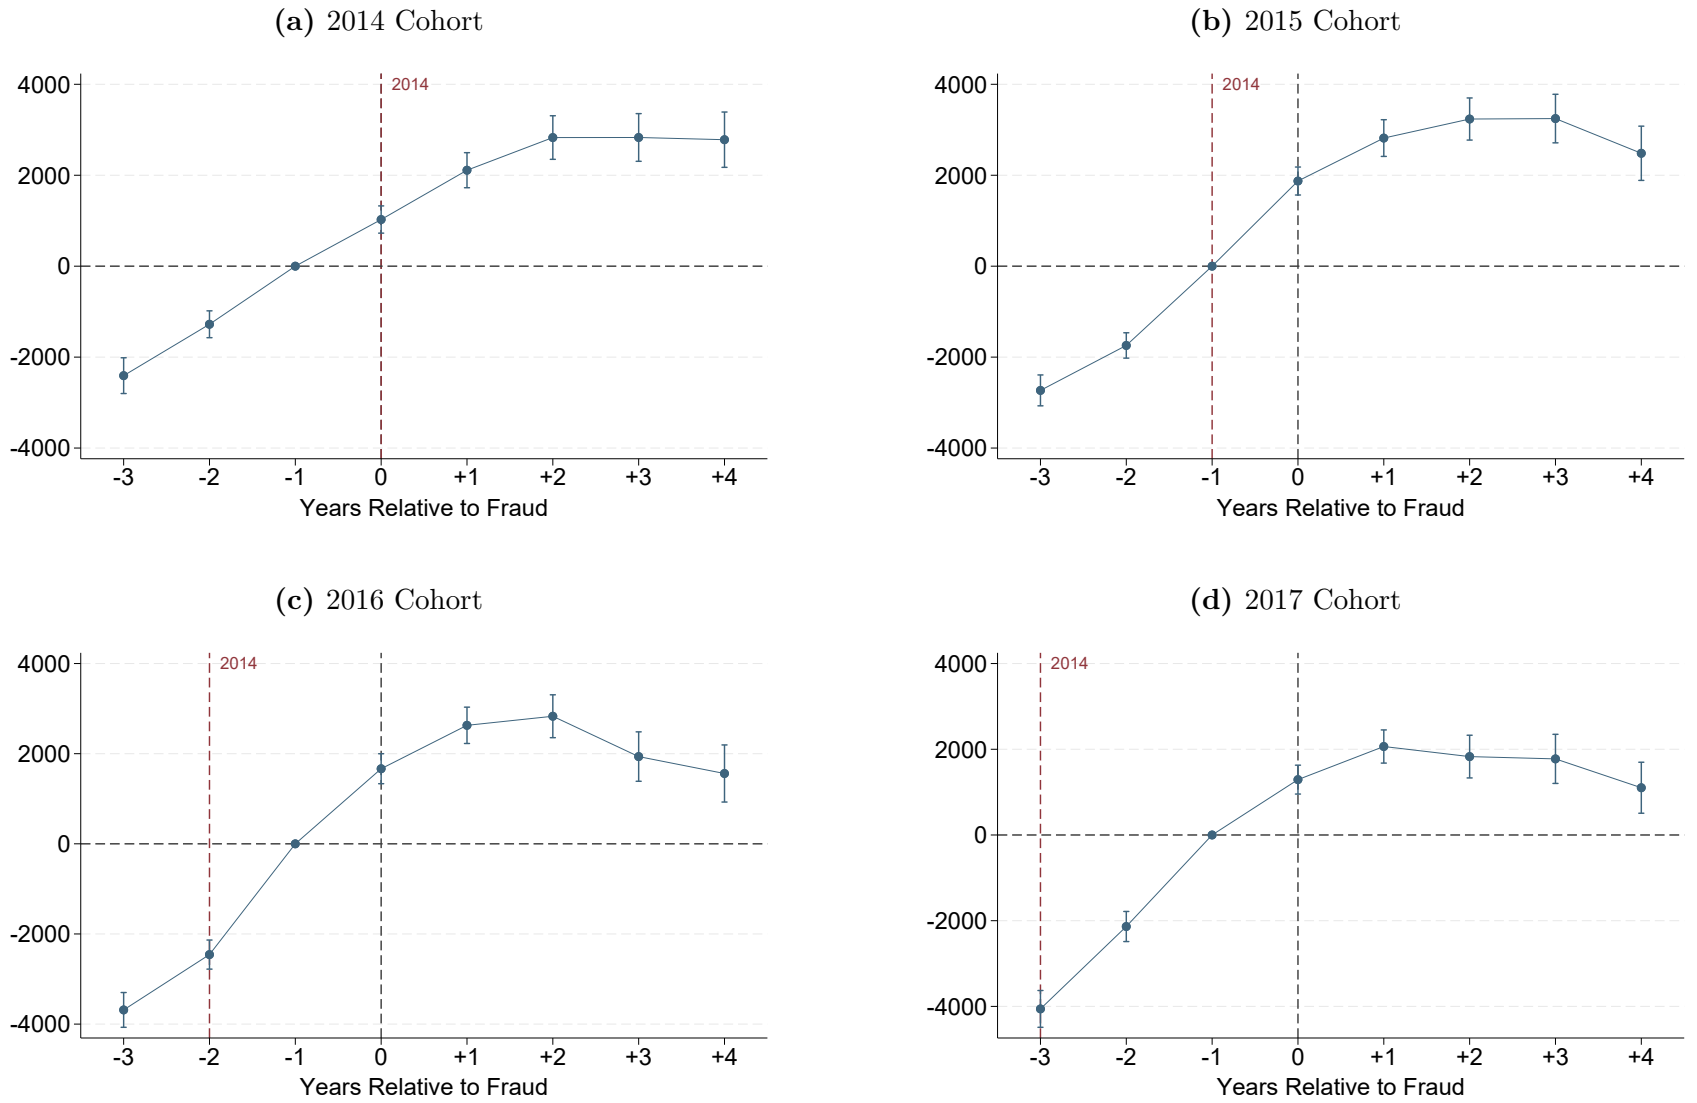



**(b)** Post-Chip Cohorts: Fraud Years 2014–2017




Notes: Each panel plots event-study coefficients ($\hat{\beta}_j$) and 95% confidence intervals from Equation (3), with the number of credit inquiries in a 6-month window as the outcome. Panel (a) pools consumers whose first fraud occurred between 2010 and 2013 (pre-chip); Panel (b) pools consumers whose first fraud occurred between 2014 and 2017 (post-chip). Standard errors are clustered at the consumer level.

**Figure A.8:** Stacked Cohort Event Study: Credit Limit ($) by Post-Chip Fraud Cohort




Notes: Each panel shows stacked cohort event-study coefficients ($\hat{\beta}_j$ from Equation (4)) and 95% confidence intervals for total credit limit ($), restricted to the indicated fraud cohort's sub-sample. The horizontal axis shows event time $j$ relative to the cohort's first-fraud year; $j = -1$ is the omitted reference period. The gray dashed vertical line marks the fraud year. Standard errors are clustered at the cohort-by-consumer level.