

Retailer Bankruptcy and Its Spillover to Consumer Credit

March 2026

Abstract

Using detailed credit bureau data on U.S. consumer credit cards, we examine how retailer financial distress spills over to household balance sheets through retail credit cards. We show that retailers undergoing bankruptcy close retail credit cards en masse, resulting in sharp declines in affected consumers' total credit limits. We document significant heterogeneity in these bankruptcy spillover effects. Credit-constrained consumers (e.g., lower credit-score and higher credit utilization consumers) are unable to hedge against their retailer's bankruptcy risk by accessing alternative credit cards and experience stronger and more persistent reductions in total credit card borrowing capacity. Consequently, these consumers reduce their aggregate credit card borrowing, highlighting the real consequences and negative externalities imposed by retailer bankruptcies. Overall, we identify a novel retail credit card channel through which retailer bankruptcies impose negative externalities on households. Our findings reveal who bears the costs of retailer bankruptcy risk and demonstrate the persistence of these effects.

Keywords: Credit cards; Retail credit cards; Retail finance; Bankruptcy; Credit supply shock

JEL Classification: G33; G51; E21

Traditional corporate bankruptcy cost studies have primarily focused on direct costs (e.g., legal and administrative fees) and indirect costs (e.g., lost sales, reduced operational efficiency) to the filing firm and its immediate stakeholders (e.g., employees).¹ However, bankruptcies can impose significant negative externalities, and the effect of these externalities, especially on consumers, is less well understood. Using unique credit bureau data, this paper identifies a novel credit channel through which bankrupt firms impose negative externalities on consumers. Specifically, we study the bankruptcies of large retail stores and the spillover effects of these bankruptcies on consumers through retail store credit cards.

Retail store cards (e.g., Macy’s store credit card) are credit cards issued by a retail store in partnership with a financial institution. These cards are closed-loop credit cards which can only be used at the issuing store to make purchases. In contrast, general-purpose credit cards (e.g., Chase Sapphire credit card) are open-loop credit cards which can be used at multiple merchants to make purchases. Retail cards are widespread in the U.S., facilitating consumption smoothing of households. As of 2024 there were about 300 million retail cards in circulation. More than half of the top 100 retailers offer retail cards spanning a variety of merchant categories such as department stores, furniture retailers, home improvement chains, jewelry stores, electronics vendors, and grocery outlets.

Despite the ubiquity of retail cards, the economic rationale for this form of credit intermediation and the consequences of its disruption remain understudied. A key contractual feature of a retail card is its closed-loop structure, which restricts card usage to the issuing retailer’s stores. This restriction mitigates information frictions by limiting transactions to a narrow set of goods and merchants, enabling lenders to better predict spending patterns and reduce exposure to unexpected or surprise transactions. Store-specific card rewards also foster

¹Among others, see (Weiss, 1990; Bris et al., 2006) for legal and administrative fees, (Altman, 1984; Shleifer and Vishny, 1992; Opler and Titman, 1994; Andrade and Kaplan, 1998) for lost sales, reduced operational efficiency, and damaged stakeholder relationships, (Gilson, 1997; Hertz et al., 2008; Baghai et al., 2021; Graham et al., 2023a) on losses borne by firm’s immediate stakeholders.

customer loyalty and incentivize timely repayment. Further, the restriction of transactions to the issuing retailer’s store facilitates risk-sharing contracts between banks and retailers, thereby expanding credit supply.

However, when retailers experience financial distress or declare bankruptcy, they can significantly disrupt credit intermediation through the retail credit card channel. More than 150 retailers filed for bankruptcy between 2015 and 2023, a period that is often referred to as the “Retail Apocalypse”.² Bankrupt retailers may close existing retail cards en masse, lower credit limits, and reduce the issuance of new cards. However, the impact of this disruption on consumers is not obvious. The impact may be minimal if consumers can costlessly switch to other credit cards or have adequate unused credit. Conversely, such disruptions can adversely affect consumers if credit market frictions constrain consumers from switching to other sources of credit. Our analysis serves as an implicit test for the presence of such frictions and identifies consumer segments for whom these frictions may be binding.

Studying this question empirically presents several challenges. First, one must precisely identify consumers who are affected by a given retailer’s bankruptcy. Second, it is important to isolate the bankrupt retailer’s impact on consumers’ credit supply from credit demand changes. For instance, economic downturns can cause both a retailer’s failure and a reduction in income affecting credit extension to consumers.

We tackle these challenges by using a novel credit bureau dataset obtained from one of the three major credit bureaus in the U.S. The dataset allows us to observe *all* the credit cards of a consumer along with their type (retail or general-purpose), issuer information (underwriting bank and retailer), and a time-series of key credit outcomes (limits, balances). The dataset also contains other consumer-level credit report information (e.g., credit scores, total debt), thus allowing us to track the evolution of both credit cards and the broader credit portfolio

²“The Retail Apocalypse: A Timeline,” CB Insights, accessed June 27, 2025, <https://www.cbinsights.com/research/retail-apocalypse-timeline-infographic/>.

of each consumer over time.

Next, we exploit mass closures of retail cards by bankrupt retailers as a sudden unanticipated shock that severs credit intermediation through the retailer and results in a sharp drop in credit card borrowing capacities of the affected consumers. Our empirical strategy employs a stacked difference-in-differences design around these mass card closure events to estimate the causal effect of retailer bankruptcies on consumer credit outcomes. We compare the credit outcomes of the bankruptcy-affected cardholders (the treatment group) to those of otherwise similar consumers who have retail credit cards but were not exposed to any bankruptcy during our entire sample period (the control group). In addition to including a host of control variables, we also include high-dimensional fixed effects in our analysis to ensure we compare the bankruptcy-affected consumers with control consumers within the same five-point credit score bin and five-digit ZIP code. We find both groups exhibit parallel trends in key credit variables prior to the bankruptcy event, supporting the causal interpretation of our empirical results.

We find that card closures following retailer bankruptcies cause a sharp and persistent contraction in household credit access. Affected consumers lose approximately \$2,000 of credit limits on average in the quarter of the card closure, which represents a 6% reduction in credit card borrowing capacity relative to the total credit limits before the mass card closure event.

Next, we show that this spillover of retailer bankruptcy on consumers has real consequences for aggregate credit card borrowing. By including all the credit cards issued to each consumer in our analysis, we can examine whether consumers can hedge against the transmitted bankruptcy shock by substituting to other credit cards. Examining effects over a three-year period, we find that while affected consumers' total credit limits begin recovering soon after retail card closures, this recovery is short-lived and plateaus rapidly after 1.5 years. Three years post-closure, affected consumers still have \$1,500 lower borrowing capacity (4% relative to pre-closure mean), implying a recovery of only \$500 over the three year period from

the initial \$2000 loss.

Examining the total number of credit cards, we find that affected consumers still have one less credit card relative to control consumers three years post-closure, indicating that the recovery of \$500 in credit limits occurred through limit increases on existing credit cards. However, a reduction in borrowing capacities of consumers may not matter if consumers have adequate liquidity buffers in terms of their unused credit limits. Thus, we next test whether the spillover of retailer bankruptcy through credit limits has an aggregate impact on credit card borrowing, and consequently on the consumption smoothing ability of consumers.

We find that the decline in credit card borrowing capacity of affected consumers is accompanied by a significant reduction in credit card borrowing. Compared to control consumers, credit card balances of affected consumers are lower by \$144 (2.5% of pre-closure mean) in the quarter of the card closure and \$235 (4.0% of pre-closure mean) three years post-closure. These results imply a marginal propensity to borrow out of credit limits ranging from 7 cents (144/2000) to 16 cents (235/1500), which is consistent with estimates documented in prior literature ([Gross and Souleles, 2002](#); [Agarwal et al., 2017](#); [Aydin, 2022](#); [Chava et al., 2023](#)).

We also uncover strong heterogeneity in outcomes across borrower types. Affected consumers with ex-ante higher credit-utilization rates (ratio of balances to limits) and lower credit scores experienced the most severe and persistent adverse effects of the retailer’s bankruptcy spillover. Consumers in the highest utilization decile lost 11% of their total credit limits and cut credit card borrowing by about \$1,000 (13% of the pre-closure mean). By contrast, even though low-utilization consumers experienced a reduction in credit limits, there was no detectable change in their credit card borrowing, suggesting that the spillover of retailer bankruptcy did not result in binding credit constraints for low-utilization consumers. Similarly, consumers at the bottom end of the credit score distribution in the 500–550 range experienced a \$1,830 (38% of pre-closure mean) cut in credit limits resulting in a \$1,300 (40% of pre-closure mean) reduction in credit card borrowing. Affected consumers at the top end

in the 800–850 range experienced about a 7% decline in limits, but no significant changes in credit borrowing. Together, our findings show that the spillover effects of retailer bankruptcies are borne by those who are least able to absorb the shock.

Finally, to understand the economic mechanisms behind our results, we analyze how the spillover of retailer bankruptcy affects credit card debt dynamics among consumers facing different credit constraints. The patterns we uncover are consistent with consumption–savings models featuring precautionary saving motives, where consumers build unused credit limits as liquidity buffers to insure against future income risk (Deaton, 1991; Carroll, 1997). A reduction in credit limits tightens these buffers and thus strengthens precautionary motives. Consumers respond by reducing borrowing to rebuild liquidity buffers, even when they were initially operating far below their credit limits. Our empirical results mirror these theoretical predictions. First, consumers exposed to retailer bankruptcy curtail their credit card borrowing relative to unaffected consumers even when operating far below their limits. Second, affected consumers who are closer to their limits dramatically delever and reduce credit utilization ratios to actively build credit buffers. Taken together, our results support precautionary saving motives induced by the contraction in credit limits that occur through involuntary card closures resulting from a retailer’s bankruptcy.

Our paper makes two key contributions. First, we document a novel retail credit card channel through which retailer financial distress transmits to consumers. These findings contribute to the bankruptcy costs literature by identifying previously understudied negative externalities of firm bankruptcies on consumers. Prior research on bankruptcy externalities has primarily examined how firm failures affect other firms or local labor markets. For example, Graham et al. (2023b) find that employee annual earnings fall by approximately 13% in the first full calendar year following firm bankruptcy. Baghai et al. (2021) show that distressed firms lose workers with the highest cognitive and noncognitive skills as they approach bankruptcy. Additionally, Bernstein et al. (2019) document that forced liquidations lead to significant job losses

among nearby firms and reduce growth and entry of new businesses in local economies. We extend this literature by demonstrating that household balance sheets can also be materially affected by firm financial distress via the propagation of a credit contraction.

Second, our study highlights the critical role of retailers and retail credit cards in intermediating credit from banks to consumers. When this intermediation is disrupted, credit-constrained consumers are unable to substitute to other credit cards, exacerbating their credit constraints. Our results show that retailer financial distress risk is disproportionately and persistently borne by consumers with greater financial fragility. These findings inform policy discussions about whether the leverage of retailers that extend significant consumer credit should be regulated much like banks and other financial intermediaries. This is particularly relevant given the growing trend of private equity takeovers of retailers through leveraged buyouts.

The remainder of the paper is organized as follows. Section 1 provides institutional background on retail credit cards. Section 2 describes the data. Section 3 outlines the empirical methodology and presents the main results on the aggregate effects of retail card closures. Section 4 examines heterogeneous impacts across borrower types based on pre-shock credit utilization and credit scores. Section 5 examines which lenders provide replacement credit after retailer bankruptcies. Section 6 analyzes credit dynamics across pre-shock utilization bins. Section 7 concludes.

1 Institutional Background on Retail Credit Cards

Retail credit cards, also known as store cards, are closed-loop consumer credit lines offered through partnerships between retailers (e.g., Macy’s) and financial institutions (e.g., Citigroup). Unlike general-purpose credit cards (e.g., a Chase Sapphire card), which can be used

across merchants, retail cards are limited to purchases at the issuing retailer.³ Retail credit accounts constitute a significant and understudied segment of the U.S. consumer credit market. Over half of the top 100 U.S. retailers offer a store-branded card.⁴ In 2023, private-label balances totaled \$64.9 billion across 107 million accounts held by 68.5 million consumers. Notably, more than 60% of these balances are held by nonprime borrowers,⁵ highlighting the importance of retail credit as a source of financing and consumption smoothing for lower credit-score consumers.

The closed-loop nature of store cards is a crucial contractual feature. Because closed-loop store cards are usable only at the issuing retailer, they can reduce lender risk through greater control over spending and risk-sharing partnerships with retailers. This structure helps mitigate traditional credit market information frictions and enables credit extension to relatively higher-risk individuals who might otherwise be excluded from the general-purpose credit card market. Consistent with these ideas, Figure 1 provides evidence that retail credit cards are more accessible than general-purpose credit cards. Figure 1 plots the credit score and age distribution of consumers when they first enter the credit card market. Crucially, the sample in Figure 1 is limited to banks that underwrite both closed-loop retail store credit cards and open-loop general-purpose cards. Panel A and Panel B show that the *same* set of banks is willing to extend credit to lower-credit-score and younger consumers in partnership with a retailer through retail store credit cards than when they extend credit directly to the consumer through general-purpose credit cards. Thus, Figure 1 is consistent with the idea that the closed-loop nature of retail credit cards helps mitigate information frictions between

³Co-branded credit cards, which carry a retailer’s brand but function as open-loop general-purpose cards, are excluded from our analysis. Throughout this paper, “retail card” or “store card” refers exclusively to private-label closed-loop cards.

⁴Consumer Financial Protection Bureau, “Issue Spotlight: The High Cost of Retail Credit Cards,” accessed June 27, 2025, <https://www.consumerfinance.gov/data-research/research-reports/issue-spotlight-the-high-cost-of-retail-credit-cards/>.

⁵Board of Governors of the Federal Reserve System, “Estimating Retail Credit in the U.S.,” *FEDS Notes*, June 21, 2024, accessed June 27, 2025, <https://www.federalreserve.gov/econres/notes/feds-notes/estimating-retail-credit-in-the-u-s-20240621.html>.

lenders and borrowers. This aligns with theoretical insights from [Diamond \(1991\)](#) who shows that monitored contracts (closed-loop retail cards) dominate for high-risk borrowers, while unmonitored contracts (open-loop general-purpose cards) dominate for low-risk borrowers.

Our account-level credit-bureau data confirm that retail cards disproportionately serve lower-credit-score borrowers and often function as a gateway to credit. The data source is described in detail in [Section 2.1](#). As shown in [Figure 1](#), consumers opening their first retail card have materially lower baseline median VantageScore than those entering the bankcard market for the first time (median score 640 vs. 670), and they tend to be younger.

[Insert [Figure 1](#) here]

[Figure 2](#) further demonstrates that reliance on retail cards declines monotonically with credit score. Adopting credit-bureau practice tiers for VantageScore 3.0,⁶ [Panel A](#) focuses on credit-limit access: deep-subprime borrowers (300–499) hold 62% of their total credit limits on store cards, whereas super-prime borrowers (781–850) hold only 26%. [Panel B](#) turns to real borrowing and shows a parallel gradient: deep-subprime borrowers carry 56% of their revolving balances on retail cards, compared with just 19% for super-prime borrowers, providing further evidence of the central role retail cards play in financially constrained households’ credit portfolios.

[Insert [Figure 2](#) here]

⁶Experian, “Credit Score FAQs,” *Ask Experian*, accessed July 13, 2025, <https://www.experian.com/blogs/ask-experian/credit-education/faqs/credit-score-faqs/>.

2 Data Sources and Sample Construction

2.1 Consumer Credit Bureau Data

We use proprietary credit card-level data from one of the three major credit bureaus in the United States covering the near universe of all credit cards in the U.S. The data are fully anonymized and provided exclusively for academic research. Our dataset is a 1 percent random sample of consumers in the credit bureau’s data from January 2014 through June 2023. For each consumer in our sample, we observe monthly panel data on every credit card they hold. These data include end-of-month snapshots of card-level limits and balances. Balances capture both revolving debt carried over from previous months and current-month spending that is not yet repaid. For each card, we also observe the issuing institution and product type (private-label versus general-purpose). One limitation of our dataset is that it does not report any information on annual percentage rates (APRs) or reward structures.

In addition to card-level limits and balances, our dataset also contains a quarterly panel of consumer-level credit report information. This information includes Vantage credit scores, total outstanding debt, number of other credit accounts (e.g., auto loans, mortgages), and balances and delinquency information on those accounts, as well as five-digit ZIP code. Together, these data allow us to track the evolution of both credit card usage and the broader credit portfolio of each consumer over time.

2.2 Retailer Bankruptcy and Mass Card Closure

We construct a brand-level panel of U.S. retailer bankruptcies for firms that issued private-label credit cards. The bankruptcy sample spans 2014–2023 and includes 60 retail brands associated with 37 parent companies. A bankruptcy event is defined as the parent company’s filing for Chapter 11 or Chapter 7 protection. Filing dates from 2014 to 2019 are sourced from

[Alanis et al. \(2022\)](#). We extend this dataset through 2023 by hand-collecting public filings and announcements. For each case, we verify that the retailer operated an active private-label card program at the time of filing by cross-checking issuer information from partner banks.

Within the universe of bankrupt retailers, we isolate mass card closure events – terminations of private-label card programs that occur during bankruptcy and are clearly observable in our credit-bureau records. We flag a mass card closure when the share of open accounts for a retailer’s card falls by at least 40 percentage points in a single quarter relative to the level observed four quarters before the bankruptcy event. Each candidate is then rigorously validated using court dockets, press releases, and lender notices to confirm that the observed collapse aligns with a publicly announced program shutdown. We retain only those events occurring within three quarters of the bankruptcy filing. This screening yields 29 brands (19 parent firms) whose private-label card programs were terminated at or shortly before bankruptcy, producing an immediate account-closure spike in the bureau data. By contrast, bankruptcies that involve continued operations show no comparable decline, indicating that card programs either remained active or were transferred to successor issuers. Finally, to ensure adequate pre-trend and post-event windows in our quarterly panel, we further restrict the mass card closure sample to events between 2016 Q3 and 2021 Q1, leaving 27 brands in the final estimation sample.

2.3 Event-Study Sample

We hand-match retailer brand names to their card-program identifiers in the credit bureau data and merge this mapping with our consumer credit file. For each mass card closure event, we define the treated group as consumers who (i) held the retailer’s private-label card in the baseline quarter ($t = -4$) and (ii) experienced closure of the card account in the event quarter ($t = 0$). A threat to the identification is endogenous sorting: distressed retailers may entice

lower-quality borrowers with last-minute discounts, while higher-quality customers may preemptively cancel their cards. By restricting the sample to cardholders active a full year before the shutdown, we exclude these late entrants and early exiters, thereby greatly mitigating this concern. Pre-trends in Figure 7 to Figure 10 confirm that treated and control consumers follow parallel credit trajectories prior to the event. The never-treated control group comprises consumers who have held at least one private-label retail card, but never one facing a mass card closure event, ensuring comparable credit-card usage patterns.

Figure 3 provides a visual check on the sample construction. The treated series displays a discrete drop of roughly one retail card at the event quarter, consistent with the mechanical account closures underlying our design. The control series remains flat throughout, and the two groups move in close alignment during the pre-event period. The persistence of the post-event gap confirms that the sample captures the intended card loss and its aftermath rather than short-term fluctuations in cardholding.

[Insert Figure 3 here]

Our difference-in-differences sample is restricted to 27 retailer brand bankruptcies between 2016 Q3 and 2021 Q1, ensuring at least four quarters of pre-event data and sufficient post-event follow-up. These 27 mass card closure events fall into 12 distinct calendar quarters, so we define each quarter as an event cohort. For every cohort we build a relative-time panel that follows treated consumers and their never-treated counterparts from $t = -4$ through $t = +12$, where $t = 0$ is the closure quarter and negative (positive) integers denote pre- (post-) event periods. We drop 9.7% of consumers who would be treated in more than one cohort and additionally exclude those with missing or insufficient quarterly attribute data, resulting in 72,535 treated and 712,494 control individuals. Finally, we stack the 12 cohort-specific panels to construct the event-time dataset used in estimation.

2.4 Descriptive Statistics

We begin by examining the distribution and timing of card-program closure events in our sample and demonstrating their close link to retailer bankruptcy filings. Figure 4 presents the calendar-time distribution of retailer bankruptcies and their associated card-program closure events in our sample. Retailer bankruptcies occur throughout our six-year analysis period (2016–2021), but these events are primarily concentrated in 2018 and 2020.

Panel B plots the calendar-time distribution of mass card closure events resulting from the termination of entire retail store credit card programs. These mass card closure events constitute the primary variation exploited in our empirical analysis. Our sample contains 12 calendar quarters with 27 unique mass card closure events. Even though at least one mass card closure event occurs in every year of our analysis period between 2016–2021, these events are clustered around two quarters, namely, 2018-Q2 and 2020-Q4. To ensure that our results are not driven solely by mass card closure events occurring in any particular quarter or period (e.g., the COVID-19 period), we conduct robustness tests by sequentially excluding each calendar quarter from our analysis and re-estimating our baseline specification using mass card closure events from the remaining 11 quarters. These results are presented in Online Appendix Figure A.1 and show that our findings are not driven by events in any single period.

[Insert Figure 4 here]

Next, Figure 5 plots the timing of mass card closures relative to retailer bankruptcy filings in our sample. Close to 75% of our mass card closure events occur in the first two quarters after the bankruptcy filing quarter.

[Insert Figure 5 here]

Finally, Figure 6 plots the fraction of retail cards closed and opened within the ± 1 year event window surrounding mass card closure events. Approximately 80% of retail cards are closed during the mass card closure quarter relative to the stock of retail cards one year prior to the event. New card openings decline steadily in the year preceding the mass card closure quarter, and no new cards are opened following the mass card closure event. Collectively, Figure 6 is consistent with the complete wind-down of retail store credit card programs, providing empirical validation for the information we collected from news articles documenting the closure of retail store credit card programs following retailer bankruptcy filings.

Overall, these figures demonstrate that mass card closure events resulting from retailer bankruptcies are abrupt, large-scale contractions in retail credit supply. These characteristics facilitate causal interpretation in our empirical analysis of the spillover effects of retail store bankruptcies on consumers through the credit card channel.

[Insert Figure 6 here]

We now turn to a comparison of baseline characteristics between treated and control consumers, reported in Table 1. The table presents credit attributes measured one year prior to the mass card closure event.

The first three rows highlight a key distinction: treated consumers exhibit significantly greater reliance on retail credit. Retail cards account for 56.8% of their total credit card holdings, compared to 33.9% among controls. Similarly, 26.5% of their total credit card balances are on retail cards, versus 16.7% for controls. Crucially, this higher retail exposure does not signal weaker credit quality. On the contrary, treated consumers have higher average VantageScores (744 vs. 723) and lower utilization rates (21.7% vs. 29.3%). They also have marginally higher total credit limits and balances, consistent with ample borrowing capacity. Their longer credit histories (324 vs. 261 months) reflect our sample restriction to consumers holding the affected retail card for at least four quarters – excluding late entrants and early

exiters by construction. Collectively, these patterns help mitigate concerns that distressed retailers disproportionately attract riskier borrowers.

The control group consists of consumers who have held at least one private-label retail card but were never exposed to a mass card closure event. We restrict the control sample to consumers residing in the same five-digit ZIP codes as the treated group. This restriction does not affect identifying variation, since consumers outside these ZIP codes would be excluded once ZIP-code fixed effects are included. To further address observable differences, all regressions control for baseline characteristics including VantageScore, age, credit history length, retail-credit reliance, and delinquency status. We also include fixed effects for five-point VantageScore bins and five-digit ZIP-code groups, each interacted with event cohort. Despite level differences in observables, treated and control consumers exhibit flat and parallel pre-event trends, as shown in the dynamic DiD plots in Section 3.

Although treated consumers tend to appear more sophisticated in terms of credit behavior, this reflects the design of our sample, which includes only those who held the relevant retail card for at least four quarters prior to the event. While this sampling strategy may introduce relatively more credit-sophisticated consumers into the treated group, it serves an important purpose that it reduces the concern that financially weaker consumers disproportionately sort into relationships with struggling retailers prior to bankruptcy. Furthermore, in Section 4, we examine heterogeneous effects by borrower characteristics and show that the primary responses in total borrowing, credit scores, and utilization are concentrated among subprime and high-utilization consumers, despite their smaller share of the treated sample. These responses align closely with predictions from buffer-stock saving models, reinforcing the interpretation that the shock has disproportionately large and persistent consequences for constrained borrowers.

[Insert Table 1 here]

Table 2 tracks changes in credit outcomes from the quarter before ($t = -1$) to the quarter

after ($t = +1$) each mass card closure event. The first row confirms the mechanical effect of the program shutdown: treated borrowers lose nearly one retail card relative to controls (-0.9). Retail credit limits fall by about \$1,903 more than in the control group, reducing total limits by \$2,325. There is no offsetting increase in bankcard limits (-289), indicating limited substitution into general-purpose credit. Relative to controls, treated borrowers also reduce balances by about \$200 and outstanding debt by about \$205, while utilization rises by 1.3 percentage points. Together, these simple treated–control differences confirm that the mass card closure delivers an immediate liquidity shock that reduces borrowing capacity and credit quality, even for consumers who otherwise appear seasoned in the market.

[Insert Table 2 here]

3 Empirical Framework

To trace the quarterly path of each credit outcome around a mass card closure, we estimate the stacked difference-in-differences model

$$\Delta y_{i,c,q} = \beta_q \text{Treated}_i + \gamma' X_{i,c,-4} + \psi_{\text{score} \times c} + \theta_{\text{zip} \times c} + \varepsilon_{i,c,q}, \quad q = -4, \dots, 12, \quad (1)$$

In Equation (1), $\Delta y_{i,c,q} = y_{i,c,q} - y_{i,c,-1}$ denotes the change in outcome y for consumer i in cohort c between relative quarter q and the omitted quarter ($q = -1$). The indicator Treated_i equals one for borrowers whose retail card was closed in a mass card closure at $q = 0$ and zero for the never-treated comparison group, so the coefficients β_q trace quarter-by-quarter treatment effects. We follow borrowers from $q = -4$ to $q = +12$, allowing us to observe both the immediate disruption and the long-term adjustment to the credit shock. The control vector $X_{i,c,-4}$, measured at $q = -4$, comprises the logs of age, credit-history length, and imputed monthly income, a 90-day delinquency flag, and three ratios that capture ex-ante

reliance on retail credit (shares of cards, limits, and balances). To soak up residual cross-sectional heterogeneity, we include two high-dimensional grids of fixed effects: (i) five-point VantageScore bins interacted with event cohort ($\psi_{\text{score} \times c}$) and (ii) truncated five-digit ZIP codes interacted with cohort ($\theta_{\text{zip} \times c}$). Standard errors are clustered by retailer, using each consumer’s *largest-limit* card to assign them to a brand when the control group holds multiple store cards. Under the null of parallel trends, the pre-event coefficients $\beta_{-4}, \beta_{-3}, \beta_{-2}$ should be statistically indistinguishable from zero; we verify this in each of the dynamic DiD results.

The specification lets us include rich fixed effects without collinearity and offers several econometric advantages. First, we do not rely on propensity-score matching to construct a control group. Instead, we follow [Callaway and Sant’Anna \(2021\)](#) and use the large pool of “never-treated” retail-card holders, tightening comparability with high-dimensional fixed effects and baseline covariates. This design choice avoids the well-documented pitfalls of regression to the mean and spurious trend matching that arise when matching on pre-treatment outcomes or propensity scores ([Daw and Hatfield, 2018](#); [King and Nielsen, 2019](#)). Second, we stack the 12 event cohorts and interact all fixed effects with cohort indicators, so that each cohort has its own intercept and set of controls. This cohort-specific saturation follows the guidance in [Sun and Abraham \(2021\)](#) and eliminates the negative-weighting problem documented by [Goodman-Bacon \(2021\)](#) in staggered DiD designs. Third, we confirm that our main results are robust to the doubly robust estimator of [Callaway and Sant’Anna \(2021\)](#), which combines inverse-propensity weighting with outcome regression. These results are reported in the Online Appendix Figure [A.2](#).

3.1 Credit-Supply Shock: Cards and Limits

Table [3](#) presents an overview of the stacked DiD estimates derived from Equation [\(1\)](#) during the event quarter ($t = 0$; Panel A) and three years post-event ($t = 12$; Panel B), benchmarked

against pre-treatment means. The estimates document a severe, immediate contraction in credit supply triggered by retail card closures. On average, affected borrowers lose approximately one retail card and one total card, coinciding with a \$1.9 thousand decline in aggregate credit limits, representing a roughly 6% reduction from the baseline. Contemporaneously, credit card balances contract by 2.5% relative to the baseline. These short-run dynamics reveal that the mass closure shock directly restricts credit availability, and borrowers exhibit limited capacity to immediately substitute toward alternative borrowing channels.

The long-run effects, observed three years post-closure, are substantial both statistically and economically. Affected borrowers persistently hold one fewer credit card, and their aggregate borrowing capacity remains depressed by approximately \$1.5 thousand (4% below the pre-shock baseline). Furthermore, the contraction in borrowing deepens over time: total balances fall by \$235 (4% of baseline), and open debt obligations contract by \$710 (nearly 3% of baseline). Collectively, these magnitudes emphasize that the liquidity shock originating from retailer distress is not a transitory friction but a persistent structural shift that fundamentally reshapes household balance sheets.

[Insert Table 3 here]

Figure 7 and Figure 8 plot the event-time coefficients for four credit-supply margins alongside their 95% confidence intervals. The pre-event estimates are economically small and statistically indistinguishable from zero, validating the parallel trends assumption. Panel A of Figure 7 demonstrates an exact one-card drop in retail cards at event quarter 0, and Panel B confirms an identical one-card decline in total cards. This one-to-one reduction confirms the absence of contemporaneous substitution into general-purpose bankcards. Even after six quarters, the extensive margin recovers by less than 0.1 cards on average, underscoring a persistent shortfall in credit access.

[Insert Figure 7 here]

The erosion of borrowing capacity is both severe and highly persistent. As shown in Panel A of Figure 8, retail-card limits drop by roughly \$1.85 thousand at $t = 0$, which is a 20% collapse relative to the pre-shock average, and fail to recover through quarter 12. Total credit limits, presented in Panel B, contract by approximately \$1.94 thousand at the event quarter, equating to a 5.7% decline from the \$34.2 thousand pre-shock mean. Three years post-closure, this shortfall narrows only marginally to \$1.48 thousand, or 4.3% of the baseline. This protracted inability to restore credit capacity is consistent with recent findings that households only partially backfill lost credit following adverse shocks (Chava et al., 2023). Ultimately, our estimates capture a distinct contraction in credit supply that originates within the retail sector and is only partially offset by general-purpose credit card markets. The subsequent subsection explores how this persistent supply shock propagates into broader liquidity and credit-quality outcomes.

[Insert Figure 8 here]

3.2 Precautionary Saving as a Liquidity-Adjustment Channel

We next examine whether households respond to the sudden credit-supply shock by accumulating precautionary savings, a central mechanism emphasized by buffer-stock models of consumption. The pre-event coefficients across all specifications are economically small and statistically indistinguishable from zero, validating the parallel trends assumption.

The initial loss of retail credit limits induces a significant and persistent contraction in actual borrowing. As shown in Panel A of Figure 9, aggregate credit-card balances decline by \$144 in the event quarter and the gap widens to \$235 by quarter 12. Relative to the pre-shock mean of \$5.83 thousand, these reductions represent a 2.5% and 4.0% persistent contraction in revolving debt. Interpreting these balance reductions as suppressed borrowing (Barrot et al., 2022), the implied marginal propensity to borrow (MPB) out of liquidity is

approximately 7 cents per dollar of lost limit on impact (\$144 reduction against a \$1,943 limit drop), rising to 16 cents by quarter 12 (\$235 reduction against a \$1,484 limit shortfall). These MPB estimates align closely with, and slightly exceed, prior literature documenting MPBs of 10–14 cents ([Gross and Souleles, 2002](#)), 11 cents ([Aydin, 2022](#)), and 7 cents ([Chava et al., 2023](#)). While quarter-end snapshot reporting may introduce attenuation bias, the long-run persistence of this shortfall points to a structural reduction in household consumption rather than transient measurement noise.

Despite this behavioral reduction in absolute borrowing, borrowers experience a persistent deterioration in their available liquidity buffers. Panel B reveals that total credit utilization increases by 0.9 percentage points upon impact, which is a 4.0% rise relative to the 21.7% pre-shock baseline, and remains elevated by 1 percentage point through quarter 12. This upward drift in utilization highlights a binding constraint: the mechanical reduction in total credit limits outweighs the households’ active efforts to delever. The inability to stabilize utilization indicates that affected borrowers cannot seamlessly substitute the lost retail credit line with alternative sources of open-loop credit.

[Insert Figure 9 here]

The most compelling evidence of a broad precautionary saving response emerges from the dynamics of total open debt, which captures obligations across the entire household balance sheet. Figure 10 documents a modest initial decline of \$120 that accelerates to a \$710 reduction (2.8% of baseline) by quarter 12. Because the contraction in total debt substantially exceeds the decline in credit card balances alone, it is evident that households are not simply rotating debt to other credit facilities. Instead, they engage in comprehensive deleveraging to restore their target liquidity buffers. This magnitude of debt reduction strongly aligns with the theoretical predictions of [Aydin \(2022\)](#), wherein adverse shocks to borrowing capacity amplify precautionary motives and compress household consumption, particularly among borrowers

operating near their credit limits.

Collectively, the concurrent declines in credit-card balances and total debt, alongside elevated utilization rates, confirm the precautionary-saving channel. The discrete removal of retail credit shrinks available borrowing capacity, prompting households to curb aggregate expenditure rather than frictionlessly migrating their borrowing to alternative lenders. If this buffer-rebuilding mechanism drives the aggregate results, buffer-stock theory suggests these effects should concentrate among ex ante constrained consumers. We explore this heterogeneity in Section 4.

[Insert Figure 10 here]

4 Heterogeneous Effects Across Borrower Types

4.1 Results by Utilization Bins

We next examine whether the effects of mass card closures are concentrated among borrowers with tighter ex ante liquidity positions. We begin by sorting consumers into 10% bins of pre-event credit utilization. Within each bin, we estimate a collapsed first-difference specification that compares each consumer’s average post-event outcome (quarters 0 to 12) with the average pre-event outcome (quarters -4 to -1). This approach complements the dynamic event-study estimates in Section 3 by increasing precision in smaller subsamples and summarizing the average effect over the full three-year post-shock horizon.

$$(\bar{y}_{i,c,b}^{\text{post}} - \bar{y}_{i,c,b}^{\text{pre}}) = \beta_b \text{Treated}_i + \gamma' X_{i,-4} + \psi_{\text{score} \times c} + \theta_{\text{zip} \times c} + \varepsilon_{i,c,b}, \quad (2)$$

In Equation (2), the dependent variable is the change in the average outcome y for consumer i in cohort c and bin b , where $\bar{y}_{i,c,b}^{\text{pre}}$ and $\bar{y}_{i,c,b}^{\text{post}}$ denote the mean values of y over the pre- and post-event windows, respectively, and the regression is estimated separately within

each utilization bin b . Treated_i equals one for consumers whose private-label card is closed in a mass card closure at $q = 0$. The vector $X_{i,-4}$ includes baseline controls measured at $q = -4$: log age, log credit-history length, log imputed monthly income, a 90-day delinquency indicator, and three ratios capturing reliance on retail credit (shares of retail cards, limits, and balances). Following Equation (1), we include fixed effects for five-point credit score bins interacted with cohort ($\psi_{\text{score} \times c}$) and truncated five-digit ZIP codes interacted with cohort ($\theta_{\text{zip} \times c}$). Standard errors are clustered by retailer brand, assigned based on the largest-limit card in each consumer’s portfolio. The coefficient β_b therefore measures how the average effect of the mass closure shock varies across the utilization distribution. Because utilization summarizes how much unused credit a borrower has before the shock, heterogeneity in β_b is informative about which households can absorb the loss of retail credit and which households face binding liquidity pressure.

Heterogeneity across utilization bins reveals incomplete and uneven replacement of lost retail credit, with clear differences across the extensive and intensive margins. Figure 11 plots the estimated effects on total cards (Panel A) and total credit limits (Panel B) across 10% utilization bins.

On the extensive margin, no utilization group fully replaces the lost retail card. The net decline in total cards remains close to one throughout the distribution, indicating that extensive-margin substitution is limited even though the closure is mechanical. At the bottom of the utilization distribution, borrowers exhibit almost a one-for-one decline in total cards. Taken together with the negligible balance response below, this pattern suggests that the retail card was not central to liquidity management for these households, and they appear willing to let the relationship lapse rather than actively replace it. By contrast, consumers in the middle of the utilization distribution, especially in the 50–60% bin, offset somewhat more of the shock and therefore experience the smallest net decline in total cards. These borrowers appear to have both a motive to replace credit and sufficient access to do so. At the top of the

utilization distribution, however, replacement weakens again: highly utilized borrowers lose roughly one account on net despite having the strongest incentive to preserve access to credit. This U-shaped pattern is economically informative. Limited replacement at low utilization is consistent with low demand for substitute liquidity, whereas limited replacement at high utilization is more consistent with supply-side frictions in obtaining alternative credit.

On the intensive margin, total credit limits show a similar but sharper pattern. Low-utilization borrowers lose the most in dollar terms, about \$2,267, or 7.8% of pre-shock limits, because they begin with larger aggregate credit lines. High-utilization borrowers, by contrast, experience the largest proportional contraction: total limits fall by about \$1,113, or 11.3% of baseline capacity. The distinction between dollar and percentage losses matters. For borrowers with ample unused credit, the closure removes a card but leaves substantial remaining liquidity. For borrowers already operating near their borrowing limits, even a smaller dollar loss translates into a much tighter effective liquidity position. Thus, the economic incidence of the shock is steepest at the top of the utilization distribution, where replacement is weakest relative to need.

[Insert Figure 11 here]

Table 4 aggregates the same estimates into three broader utilization groups. The table confirms that the extensive-margin response is similar across groups: total cards fall by -1.187 for borrowers with utilization below 50%, -1.033 for those between 50% and 90%, and -1.014 for those above 90%. Expressed relative to group means, these are large declines of 30.6%, 25.3%, and 33.5%, respectively. On the intensive margin, however, the distributional burden becomes clearer. Total credit limits contract by \$2,007 (7.0%) for the 0–50% group, \$1,067 (6.5%) for the 50–90% group, and \$1,113 (11.3%) for the 90%+ group. The high-utilization group therefore does not lose the most credit in dollar terms, but it loses the most in economically relevant proportional terms.

[Insert Table 4 here]

Figure 12 shows how these credit-supply losses translate into borrowing adjustment. Panel A plots effects on total credit card balances. The pattern is broadly monotonic: borrowers in low-utilization bins show little change, while balance reductions grow larger as pre-shock utilization rises. By the top decile, consumers reduce total credit card balances by about \$1,055, or roughly 13% of baseline balances. This is a large contraction in borrowing for a group that simultaneously faces the largest proportional decline in available credit. The joint pattern is more consistent with forced deleveraging than with a voluntary retreat from credit use.

Panel B turns to total open debt, which captures adjustment in the broader household balance sheet. Low-utilization consumers reduce debt only modestly, for example, by about \$293 (1.7% of baseline) in the 0–10% bin. By contrast, borrowers in higher-utilization bins exhibit much larger declines. The 60–70% bin, for example, shows a debt reduction of about \$2,068, or 5.5% of pre-shock debt. The fact that debt reduction extends beyond credit card balances suggests that affected households are not simply rotating borrowing across accounts. Instead, they appear to delever more broadly in response to the loss of liquidity. In other words, the retail credit shock spills over from a single card program into the overall household balance sheet.

[Insert Figure 12 here]

Credit score responses by utilization bin are reported in Online Appendix Figure A.3, Panel A. The larger balance and debt reductions in the upper utilization bins are accompanied by the largest subsequent declines in credit scores. Two mechanisms likely contribute. First, there is a mechanical utilization effect: removing a credit line raises utilization on the remaining accounts even if borrowers cut balances, and scoring models penalize that higher utilization.

Second, lender-initiated account closures may enter credit records in a way that is interpreted adversely by scoring models or future lenders, even when the closure is unrelated to borrower default. The important point is that high-utilization borrowers cannot fully protect their credit profiles by deleveraging. Once the account structure is disrupted, their credit quality can still deteriorate, creating a persistent form of path dependence in credit access. This interpretation is also consistent with evidence that seemingly neutral credit events can worsen perceived creditworthiness (Lieberman et al., 2021).

Table 5 summarizes the borrowing responses by broad utilization groups. Balances change little among low-utilization borrowers ($-\$162$, or -4.6%), decline more for mid-utilization borrowers ($-\$895$, or -9.1%), and fall most sharply for the most constrained group ($-\$1,055$, or -13.1%). Open debt is directionally similar, falling by $\$428$, $\$1,340$, and $\$1,788$ across the three groups, although the estimate for the 90%+ group is less precise. Taken together, the utilization-bin results show that retailer bankruptcy transmits a highly unequal liquidity shock. Households with the least spare credit capacity are least able to replace lost retail credit, delever the most aggressively, and nonetheless experience the greatest deterioration in effective credit access.

[Insert Table 5 here]

4.2 Results by Credit Score Bins

We next examine how the response to credit line removal varies with borrowers' creditworthiness. Figure 13 reports heterogeneity across baseline VantageScore bins: Panel A plots the change in total credit limits, and Panel B plots the corresponding change in total credit card balances.

Panel A shows that consumers in the highest score bins lose more dollars on average, reflecting their larger pre-shock credit portfolios, but these losses are modest relative to available

credit. For borrowers in the 800–850 bin, the decline is only 7.2% of the pre-shock mean, which is economically similar to losing a single retail card rather than experiencing a broad contraction in liquidity. By contrast, lower-score borrowers lose less in absolute dollars but face much steeper proportional contractions. Consumers in the lowest score bin (VantageScore 500–550) lose about \$1,832, which equals 37.6% of their pre-shock total credit limits. Thus, the same card closure represents a much larger liquidity shock for borrowers with weaker credit profiles.

Panel B reveals an even sharper gradient in actual borrowing responses. Consumers in the 500–550 bin reduce outstanding balances by about \$1,348, or roughly 40% of baseline balances, following the card closure. The response attenuates steadily as credit scores rise. Other subprime and near-prime borrowers (VantageScore below 660) also cut balances materially, with declines ranging from 8% to 20% of pre-shock balances. At the top of the score distribution, by contrast, balance responses are economically close to zero, and the point estimate in the highest bin is slightly positive. The contrast indicates that low-score borrowers cannot readily replace the lost retail credit line and instead respond by deleveraging, whereas high-score borrowers absorb the lost line within existing liquidity slack.

[Insert Figure 13 here]

Table 6 summarizes the same pattern using broader prime and non-prime groups. Total limits fall by \$1,214 (−9.9%) for non-prime consumers and by \$2,027 (−6.9%) for prime consumers. Yet the borrowing response is concentrated almost entirely among non-prime households: balances decline by \$807 (−11.4%) for non-prime borrowers, compared with only \$79 (−1.8%) for prime borrowers, an estimate that is economically small and statistically indistinguishable from zero. In other words, prime consumers lose more dollars of credit supply, but non-prime consumers bear the economically meaningful contraction in borrowing. What matters is therefore not the dollar size of the removed line alone, but whether the loss is large relative to the borrower’s remaining liquidity.

[Insert Table 6 here]

These results sharpen the liquidity channel documented earlier. The aggregate MPB estimated in Section 3 masks substantial heterogeneity across the score distribution: it is economically meaningful for subprime and near-prime borrowers, but close to zero for prime consumers. For weaker-score households, retailer bankruptcy operates as a genuine credit-supply shock that forces balance-sheet adjustment because outside options are limited. For prime households, by contrast, the closed retail card appears to have served mainly convenience, rewards, or store-specific demand rather than essential liquidity. This interpretation is also consistent with the utilization-bin results in Section 4.1: in both dimensions, the strongest responses are concentrated among borrowers most likely to face binding credit constraints.

VantageScore responses are reported in the Online Appendix Figure A.3 Panel B. Across the full sample, the mean change is statistically insignificant. Subprime borrowers experience an average decline of four points, while prime borrowers see a small increase. Prime consumers thus absorb the limit cut without harming their credit files, consistent with ample alternative liquidity. In contrast, non-prime borrowers preserve their scores only by sharply deleveraging. This path-dependent dynamic underscores how a uniform credit supply shock can lead to divergent financial trajectories, which is most damaging for those already least creditworthy.

5 Lender Composition of Replacement Credit

To examine which lenders help offset the loss of retail credit after the shock, we classify each lender–consumer relationship into three categories. First, *incumbent retail lenders* are banks that issued a retailer-branded card (e.g., Citi on a Macy’s card) to the consumer during the pre-shock period (quarters $t = -4$ to $t = -1$). Second, *incumbent non-retail lenders* are all other financial institutions with a pre-shock credit card relationship with the borrower. Third, *new lenders* are institutions with no pre-existing credit card exposure to the consumer prior to

the shock. For each treated and control consumer, we identify whether at least one lender in a given category extends new credit, either through a new account or a limit increase, within four quarters after the shock, and we then re-estimate Equation (1) within the corresponding subsamples. Because these groups are defined by post-shock credit origination, the results should be interpreted as descriptive conditional patterns rather than causal comparisons across lender types. Even so, they are informative about which lender relationships are most closely associated with replacement credit after the bankruptcy shock.

Figure 14 shows that meaningful replacement credit is associated primarily with new lenders rather than incumbent relationships. Panel A indicates that recovery in the number of cards is limited across all groups: even consumers who obtain new credit from new lenders remain roughly one card below their pre-shock level throughout the post-event window. Nevertheless, relative to borrowers whose post-shock credit expansion comes from incumbent lenders, they exhibit the greatest recovery on the extensive margin. Borrowers linked to incumbent retail lenders show the weakest recovery in card counts and, if anything, continue to drift downward over time. Panel B reveals an even sharper contrast for total credit limits. Consumers who obtain new credit from new lenders recover a substantial share of the initial limit loss over the subsequent quarters, whereas those receiving credit from incumbent lenders remain far below their pre-shock borrowing capacity. Among incumbents, retailer-affiliated lenders provide the least replacement credit, suggesting that the bank connected to the bankrupt retailer does little to offset the lost line once the retail-card program is shut down.

[Insert Figure 14 here]

Figure 15 shows that these differences in limit replacement are mirrored in borrowing behavior. Panel A indicates that consumers who obtain new credit from new lenders experience a noticeably smaller and less persistent decline in credit card balances than borrowers whose

replacement credit comes from incumbent lenders. By contrast, balance reductions remain large and persistent when post-shock credit is supplied by incumbent retail or incumbent non-retail lenders, consistent with weaker outside options and tighter effective liquidity. Panel B broadens the analysis to total open debt and shows that deleveraging occurs in all groups, implying that even successful card-level replacement does not fully undo the broader balance-sheet adjustment induced by the shock. That said, the decline in total open debt is largest and most persistent for borrowers obtaining new credit from incumbent retail lenders. Overall, the evidence suggests that outside lenders account for most economically meaningful replacement credit, while existing lender relationships, especially retailer-affiliated one, provide only limited insurance against the loss of a retail credit line.

[Insert Figure 15 here]

6 Credit Dynamics by Pre-Shock Utilization

We next examine how post-shock borrowing dynamics vary with consumers' ex ante liquidity position. We sort households into ten equal-width bins based on total credit card utilization in quarter $t = -4$ and, within each bin, plot the event-time paths of treated and control consumers. This descriptive exercise complements the regression-based heterogeneity estimates in Section 4.1 by showing when the treated-control gap opens and whether adjustment occurs through an immediate deleverage or through slower subsequent balance growth.

Figure 16 reports these trajectories. Panel A scales balances by pre-shock total credit limits, $Balance_t / Limit_0$, so the denominator is fixed at its pre-event level. This normalization is useful because any post-shock divergence cannot be driven mechanically by the reduction in current credit limits. Across almost the entire utilization distribution, the treated path drops or flattens at the time of the retail-card closure and then grows more slowly than the control path thereafter. The gap is largest in the upper utilization bins, where borrowers

closest to their borrowing constraints cut balances immediately relative to controls. In the lower utilization bins, the response is less abrupt, but treated borrowers still accumulate balances more slowly, indicating foregone borrowing growth even when the shock does not bind immediately.

Panel B turns to contemporaneous utilization, $Balance_t/Limit_t$. Here the picture is different but complementary. For most of the utilization distribution, treated borrowers remain at higher utilization than controls after the shock, even though their balance growth is weaker in Panel A. This pattern reflects the fact that the decline in available credit mechanically raises utilization and offsets part of the balance adjustment. Among high-utilization households, utilization falls in levels after the shock, consistent with active deleveraging, but it remains elevated relative to controls because limits contract at the same time. Among low-utilization households, utilization stays low in absolute terms, so the main response is not forced deleveraging but rather a flatter borrowing path that preserves future liquidity.

Taken together, the two panels clarify the mechanism behind the heterogeneity results in Section 4.1. Households already close to their borrowing limits respond to retailer-card closures by reducing borrowing, yet they still end up with tighter effective liquidity because limits fall faster than balances. Households with more initial slack are not pushed against a hard constraint, but they nonetheless borrow less than comparable controls, consistent with a precautionary desire to preserve unused credit.

These patterns are consistent with prior evidence on household responses to adverse credit-supply shocks. [Chava et al. \(2023\)](#) show that consumers exposed to bank funding shocks deleverage most sharply when initial utilization is high, while less constrained borrowers primarily slow the growth of borrowing. Our results point to a similar mechanism in the retail-credit setting. Exogenous reductions in available credit induce households to rebuild or preserve liquidity buffers, in line with buffer-stock models of precautionary saving ([Aydin, 2022](#)). The broader implication is that retailer bankruptcies tighten effective household liq-

uidity throughout the utilization distribution, but the adjustment is steepest for borrowers who enter the shock already close to their credit limits.

[Insert Figure 16 here]

7 Conclusion

Corporate bankruptcy is typically viewed through the lens of firms and their immediate stakeholders. This paper shifts the focus to households by showing how retailer failures transmit credit-supply shocks to households by terminating retail credit card programs. Specifically, we study the impact of 27 retailer bankruptcies between 2016 and 2021 that resulted in retail credit card closures, documenting a previously underexamined effect of corporate bankruptcies on household balance sheets.

Using U.S. microdata and a stacked difference-in-differences design, we find that affected borrowers lose approximately six percent of total credit limits in the quarter of retail card closures, with a four percent loss persisting three years later. This persistent reduction in credit card borrowing capacity translates to decreased credit card borrowing of 7-16 cents per dollar of lost credit limits.

These effects exhibit significant heterogeneity across consumers. High-utilization and sub-prime consumers experience the largest contractions in credit and borrowing, losing up to one-third of available credit limits and reducing borrowing by over one-third. By contrast, prime and low-utilization borrowers show minimal effects on credit card borrowing, underscoring the disproportionate costs borne by financially fragile households and amplifying the broader financial and social costs of retailer failures.

Overall, our results demonstrate that corporate distress can reverberate well beyond the firm, tightening household credit constraints in persistent and unequal ways that disproportionately harm those least able to absorb the shock. These findings have important implica-

tions for the development of consumer credit policy. They inform discussions about whether retailers extending significant consumer credit should face regulatory oversight on their leverage similar to banks and other financial intermediaries. This is particularly relevant given the growing trend of private equity takeovers of retailers through leveraged buyouts.


References

- Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2017, 07). Do banks pass through credit expansions to consumers who want to borrow?*. *The Quarterly Journal of Economics* 133(1), 129–190.
- Alanis, E., S. Chava, and A. Shah (2022). Benchmarking machine learning models to predict corporate bankruptcy.
- Altman, E. I. (1984). A further empirical investigation of the bankruptcy cost question. *The Journal of Finance* 39(4), 1067–1089.
- Andrade, G. and S. N. Kaplan (1998). How costly is financial (not economic) distress? evidence from highly leveraged transactions that became distressed. *The Journal of Finance* 53(5), 1443–1493.
- Aydin, D. (2022, January). Consumption Response to Credit Expansions: Evidence from Experimental Assignment of 45,307 Credit Lines. *American Economic Review* 112(1), 1–40.
- Baghai, R. P., R. C. Silva, V. Thell, and V. Vig (2021). Talent in distressed firms: Investigating the labor costs of financial distress. *The Journal of Finance* 76(6), 2907–2961.
- Barrot, J.-N., E. Loualiche, M. Plosser, and J. Sauvagnat (2022). Import competition and household debt. *The Journal of Finance* 77(6), 3037–3091.
- Bernstein, S., E. Colonnelli, X. Giroud, and B. Iverson (2019). Bankruptcy spillovers. *Journal of Financial Economics* 133(3), 608–633. JFE Special Issue on Labor and Finance.
- Bris, A., I. Welch, and N. Zhu (2006). The costs of bankruptcy: Chapter 7 liquidation versus chapter 11 reorganization. *The Journal of Finance* 61(3), 1253–1303.
- Callaway, B. and P. H. Sant’Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics* 225(2), 200–230. Themed Issue: Treatment Effect 1.
- Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics* 112(1), 1–55.
- Chava, S., R. Ganduri, N. Paradkar, and L. Zeng (2023, 05). Shocked by bank funding shocks: Evidence from consumer credit cards. *The Review of Financial Studies* 36(10), 3906–3952.
- Daw, J. R. and L. A. Hatfield (2018). Matching and regression to the mean in difference-in-differences analysis. *Health Serv Res* 53(6), 4138–4156. 1475-6773 Daw, Jamie R Orcid: 0000-0001-6380-8523 Hatfield, Laura A P01 AG032952/AG/NIA NIH HHS/United States CIHR/Canada Journal Article Research Support, N.I.H., Extramural Research Support, Non-U.S. Gov’t United States 2018/06/30 Health Serv Res. 2018 Dec;53(6):4138-4156. doi: 10.1111/1475-6773.12993. Epub 2018 Jun 29.


- Deaton, A. (1991). Saving and liquidity constraints. *Econometrica* 59(5), 1221–1248.
- Diamond, D. W. (1991). Monitoring and reputation: The choice between bank loans and directly placed debt. *Journal of political Economy* 99(4), 689–721.
- Gilson, S. C. (1997). Transactions costs and capital structure choice: Evidence from financially distressed firms. *The Journal of Finance* 52(1), 161–196.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics* 225(2), 254–277.
- Graham, J. R., H. Kim, S. Li, and J. Qiu (2023a). Employee costs of corporate bankruptcy. *The Journal of Finance* 78(4), 2087–2137.
- Graham, J. R., H. Kim, S. Li, and J. Qiu (2023b). Employee costs of corporate bankruptcy. *The Journal of Finance* 78(4), 2087–2137.
- Gross, D. B. and N. S. Souleles (2002, 02). Do liquidity constraints and interest rates matter for consumer behavior? evidence from credit card data*. *The Quarterly Journal of Economics* 117(1), 149–185.
- Hertzel, M. G., Z. Li, M. S. Officer, and K. J. Rodgers (2008). Inter-firm linkages and the wealth effects of financial distress along the supply chain. *Journal of Financial Economics* 87(2), 374–387.
- King, G. and R. Nielsen (2019, 2019). Why propensity scores should not be used for matching. *Political Analysis* 27(4), 435–454.
- Lieberman, A., D. Paravisini, and V. Pathania (2021). High-cost debt and perceived creditworthiness: Evidence from the uk. *Journal of Financial Economics* 142(2), 719–736.
- Opler, T. C. and S. Titman (1994). Financial distress and corporate performance. *The Journal of Finance* 49(3), 1015–1040.
- Shleifer, A. and R. W. Vishny (1992). Liquidation values and debt capacity: A market equilibrium approach. *The Journal of Finance* 47(4), 1343–1366.
- Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* 225(2), 175–199.
- Weiss, L. A. (1990). Bankruptcy resolution: Direct costs and violation of priority of claims. *Journal of Financial Economics* 27(2), 285–314.

Figure 1: Entry Attributes of Retail vs. Bankcard Borrowers

Panel A: Credit Score Distribution of Entrants




Panel B: Age Distribution of Entrants




This figure compares entry-quarter characteristics of consumers opening their first retail card with those opening their first bankcard. Data are drawn from a 1% random sample of U.S. credit-card accounts. Panel A plots the distribution of VantageScores at account opening. Panel B plots the distribution of borrower age in the same quarter. The sample is limited to the ten largest issuers by total card volume, each issuing at least 20% of accounts in both retail and general-purpose segments.

Figure 2: Reliance on Retail Credit by Credit Score Bin

Panel A: Retail Card Share of Total Credit Card Limit




Panel B: Retail Card Share of Revolving Balances




This figure plots average reliance on retail credit across credit-score bins. Credit-score bins follow the bureau practice for VantageScore 3.0: deep-subprime (300–499), subprime (500–600), near-prime (601–660), prime (661–780), and super-prime (781–850). Panel A shows the mean share of total credit limits on private-label cards. Panel B shows the mean share of revolving balances on those cards. Estimates are based on a 1% random sample of the U.S. credit-card population.

Figure 3: Retail Card Loss Around Program Termination




This figure shows the mean number of retail (private-label) credit cards by event quarter, expressed relative to $q = -6$. Values are computed as simple averages for treated and control groups, with no matching, controls, or fixed effects applied. Quarter $t = 0$ marks the mass card closure, and treated and control groups are defined as in Section 2.3.

Figure 4: Calendar-Time Distribution of Retailer Bankruptcies and Mass Card Closure Events




This figure displays the calendar-quarter timing of retailer bankruptcies and mass card closures. The sample is restricted to the 27 mass card closure events that satisfy our identification criteria and form the final regression sample (2016 Q3 - 2021 Q1). Panel A plots the quarterly count of the 27 retailer bankruptcy filings. Panel B plots the quarterly count of the 27 mass card closure events.

Figure 5: Proximity of Mass Card Closures to Bankruptcy Filings




This figure plots the timing differential, in calendar quarters, between each mass card closure and the retailer’s bankruptcy filing ($t = 0$). The plot shows the number of affected brands by differential quarter. Data are drawn from a 1% random sample of U.S. credit-card accounts and matched to retailer brands by card product name.

Figure 6: Card-Account Turnover Around Mass Card Closure Events




This figure shows the number of newly opened (light bars) and newly closed (dark bars) retail card accounts each quarter, scaled to the stock of open accounts in $t = -4$. Quarter $t = 0$ marks the mass card closure for each retailer. Data come from a 1% random sample of U.S. credit-card accounts, matched to retailer brands by card product name.

Figure 7: Dynamic Difference-in-Differences Estimates: Card Counts




This figure plots stacked difference-in-differences coefficients from Equation (1) for consumer credit card counts. Quarter $t = 0$ is the mass card closure, and quarter $t = -1$ is the omitted period for first differences. “Retail” refers to private-label cards usable only at the issuing retailer, while “Total” includes all card types (retail, co-branded, and general-purpose). Vertical capped lines represent 95% confidence intervals, with standard errors clustered by retailer brand.

Figure 8: Dynamic Difference-in-Differences Estimates: Credit Limits




This figure plots stacked difference-in-differences coefficients from Equation (1) for consumer credit limits. Quarter $t = 0$ is the mass card closure, and quarter $t = -1$ is the omitted period for first differences. “Retail” refers to private-label cards usable only at the issuing retailer, while “Total” includes all card types. Vertical capped lines represent 95% confidence intervals, with standard errors clustered by retailer brand.

Figure 9: Dynamic Difference-in-Differences Estimates: Balances and Utilization




This figure plots stacked difference-in-differences coefficients from Equation (1) for two consumer-level outcomes. Panel A shows the total balance carried across all credit cards. Panel B shows the total credit card utilization rate, defined as aggregate balances divided by aggregate limits across all credit cards. Quarter $t = 0$ marks the mass card closure, and quarter $t = -1$ is the omitted period for first differences. Vertical capped lines denote 95% confidence intervals with standard errors clustered by retailer brand.

Figure 10: Dynamic Difference-in-Differences Estimates: Total Open Debt




This figure plots stacked difference-in-differences coefficients from Equation (1) for total open debt outstanding. This outcome includes all active debt obligations reported to the credit bureau and is used as the numerator in standard debt-to-income (DTI) calculations. Quarter $t = 0$ marks the mass card closure, and quarter $t = -1$ is the omitted period for first differences. Vertical capped lines denote 95% confidence intervals with standard errors clustered by retailer brand.

Figure 11: Heterogeneous Effects by Borrower Utilization: Credit Margins




This figure plots average treatment effect estimates from Equation (2) across 10% bins of baseline utilization. Panel A reports the change in the number of credit cards. Panel B reports the change in total credit limit across all cards. Blue lines (left axis) plot the level change in outcome; red lines (right axis) plot the change as a percentage of the pre-treatment mean. Vertical capped lines indicate 95% confidence intervals, with standard errors clustered by retailer brand.

Figure 12: Heterogeneous Effects by Borrower Utilization: Credit Balances




This figure plots average treatment effect estimates from Equation (2) across 10% bins of baseline utilization. Panel A shows the change in total credit card balances, and Panel B shows the change in total open debt outstanding. For each bin, the blue line (left axis) reports the level change, while the red line (right axis) expresses the change as a percentage of the pre-treatment mean. Vertical capped lines denote 95% confidence intervals, with standard errors clustered by retailer brand.

Figure 13: Heterogeneous Effects by Borrower Score




This figure plots average treatment effect estimates from Equation (2) across 50-point VantageScore bins. Panel A displays the change in total credit limit across all cards. Panel B displays the change in total credit-card balances. For each score bin, blue lines (left axis) give the level change, and red lines (right axis) express the change as a percentage of the pre-treatment mean. Vertical capped lines represent 95 % confidence intervals, with standard errors clustered by retailer brand.

Figure 14: Dynamics of Credit for Borrowers Obtaining New Credit by Lender Type




This figure plots stacked difference-in-differences coefficients from Equation (1) for consumers who obtain new credit within four quarters of the event, separately by lender type (incumbent retail, incumbent non-retail, and new lenders). Panel A shows the dynamics of the number of credit cards, and Panel B shows the dynamics of total credit limits. Quarter $t = 0$ marks the mass closure of retailer cards associated with the bankruptcy event. Vertical capped lines denote 95% confidence intervals, with standard errors clustered by retailer brand.

Figure 15: Dynamics of Debt for Borrowers Obtaining New Credit by Lender Type



This figure plots stacked difference-in-differences coefficients from Equation (1) for consumers who obtain new credit within four quarters of the event, separately by lender type (incumbent retail, incumbent non-retail, and new lenders). Panel A shows the dynamics of total credit card balances, and Panel B shows the dynamics of total open debt. Quarter $t = 0$ marks the mass closure of retailer cards associated with the bankruptcy event. Vertical capped lines denote 95% confidence intervals, with standard errors clustered by retailer brand.

Figure 16: Balance and Utilization Dynamics by Treatment Status



This figure compares the long-run evolution of credit card balances and utilization for treated versus control consumers. Panel A reports balances scaled by pre-shock total credit card limits, $\text{Balance}_t / \text{Limit}_0$, and Panel B reports the contemporaneous utilization ratio, $\text{Balance}_t / \text{Limit}_t$, where t denotes event time and $t = 0$ is the quarter of mass closure of retailer cards. Treated consumers hold cards issued by the bankrupt retailer, whereas control consumers are matched peers without such exposure. Outcomes are averaged within each group.

Table 1: Descriptive Statistics on Consumer Characteristics

Variable	Treated		Control	
	Mean	SD	Mean	SD
<i>Retail card dependence</i>				
Number of Retail Cards	3.7	1.7	1.2	1.2
Retail Card Share (%)	56.8	19.9	33.9	30.8
Retail Limit Share (%)	38.3	27.5	25.0	31.6
Retail Balance Share (%)	26.5	34.1	16.7	30.9
<i>Credit characteristics</i>				
Credit Score	744.3	80.4	723.5	82.6
Total Credit Limit (\$)	34,207.5	28,826.2	24,210.4	26,038.4
Total Credit Utilization Ratio (%)	21.7	26.0	29.3	31.5
Total Outstanding Debt (\$)	25,586.5	35,120.7	26,185.1	36,363.5
Credit Experience (months)	324.0	125.7	261.1	142.1
Modelled Monthly Income (\$)	4,014.6	1,537.8	4,038.7	1,993.1
Number of Consumers	72,535		712,494	
Unique 5-digit ZIPs	16,640		16,640	

This table presents descriptive statistics for treated and control consumers in the baseline pre-shock quarter ($t = -4$), defined as one year prior to mass card closure events. Treated consumers are those whose retail credit cards were abruptly closed during mass card closure events resulting from retailer bankruptcies. Control consumers are those with retail credit cards issued by non-bankrupt retailers.

Table 2: Two-Quarter Changes in Credit Variables Around Mass Card Closure Events

Variable	Treated			Control			Difference	
	$t = -1$	$t = +1$	Δ^T	$t = -1$	$t = +1$	Δ^C	$\Delta^T - \Delta^C$	t -stat
Number of Retail Cards	3.6	2.7	-0.9	1.2	1.2	0.0	-0.9	-503.17***
Number of Bank Cards	3.1	3.1	0.0	2.5	2.5	0.0	0.0	-20.08***
Total Bankcard Limit (\$)	25,031	25,147	67	21,455	21,811	356	-289	-14.81***
Total Retail Card Limit (\$)	9,076	7,202	-1,890	3,579	3,592	13	-1,903	-249.39***
Total Credit Limit (\$)	34,163	32,203	-1,960	24,801	25,166	365	-2,325	-112.67***
Total Credit Balance (\$)	5,675	5,402	-273	4,866	4,791	-75	-199	-13.69***
Total Outstanding Open Debt (\$)	25,276	25,269	-107	26,174	26,245	98	-205	-3.76***
Total Credit Utilization (%)	20.8	21.0	0.2	27.6	26.5	-1.1	1.3	19.84***
Retail Card Share (%)	58.8	48.0	-10.8	33.3	32.9	-0.4	-10.4	-229.59***
Retail Limit Share (%)	37.8	31.7	-6.1	24.6	24.3	-0.3	-5.8	-125.27***
Retail Balance Share (%)	21.8	19.1	-2.7	13.9	13.5	-0.4	-2.3	-25.84***
Vantage Score	748.7	749.9	1.2	729.8	733.2	3.1	-1.9	-14.39***

This table presents two-quarter changes in credit outcomes from $t = -1$ (the quarter before the mass card closure) to $t = +1$ (the quarter after). Δ^T (Δ^C) is the within-consumer change for treated (control) borrowers. The last two columns show the simple difference-in-differences estimate ($\Delta^T - \Delta^C$) and its t -statistic. Standard errors are clustered at the consumer level and *, **, and *** indicate statistical significance at the 10%, 5%, and 1%, respectively.

Table 3: Immediate and Long-run Treatment Effects

Panel A: Immediate response in event quarter							
	Δ Retail Cards (1)	Δ Total Cards (2)	Δ Retail Limit (3)	Δ Total Limit (4)	Δ Total Balance (5)	Δ Open Debt (6)	Δ Credit Card Utilization (7)
$I(\text{Treated})$	-1.004*** (0.031)	-1.027*** (0.030)	-1,846.141*** (186.597)	-1,943.455*** (197.918)	-143.749*** (31.257)	-120.325* (57.589)	0.009*** (0.001)
Pre-shock average	3.7	6.8	9,208.6	34,207.5	5,829.0	25,586.5	21.7
Relative to pre-shock avg. (%)	-27.1	-15.1	-20.0	-5.7	-2.5	-0.5	+4.0
VantageScore bin $\times$ cohort FE	✓	✓	✓	✓	✓	✓	✓
ZIP $\times$ cohort FE	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓
N	1,450,354	1,450,354	1,450,354	1,450,354	1,450,354	1,345,158	1,449,837
Adj. R^2	0.312	0.177	0.100	0.029	0.005	0.006	0.013

Panel B: Long-run response in three years							
	Δ Retail Cards (1)	Δ Total Cards (2)	Δ Retail Limit (3)	Δ Total Limit (4)	Δ Total Balance (5)	Δ Open Debt (6)	Δ Credit Card Utilization (7)
$I(\text{Treated})$	-0.992*** (0.031)	-1.062*** (0.035)	-1,873.676*** (218.479)	-1,484.475*** (234.962)	-235.084*** (61.348)	-710.227*** (137.233)	0.010*** (0.003)
Pre-shock average	3.7	6.8	9,208.6	34,207.5	5,829.0	25,586.5	21.7
Relative to pre-shock avg. (%)	-26.8	-15.6	-20.3	-4.3	-4.0	-2.8	+4.6
VantageScore bin $\times$ cohort FE	✓	✓	✓	✓	✓	✓	✓
ZIP (truncated) $\times$ cohort FE	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓
N	1,347,079	1,347,079	1,347,079	1,450,354	1,450,354	1,123,086	1,346,554
Adj. R^2	0.172	0.115	0.068	0.088	0.038	0.035	0.053

This table reports stacked difference-in-differences estimates from Equation (1) for seven consumer-level outcomes. Panel A presents immediate effects in the quarter of the mass card-program closure ($t = 0$), and Panel B presents long-run effects three years later ($t = 12$). Outcomes include changes in retail cards, total cards, retail and total credit limits, total credit card balances, total open debt, and overall credit card utilization. “Retail” refers to private-label cards usable only at the issuing retailer. “Total” includes all card types. The specification includes baseline controls (log age, log credit-history length, log monthly income, a 90-day delinquency flag, and retail shares of cards, limits, and balances) and VantageScore $\times$ cohort and ZIP $\times$ cohort fixed effects. Standard errors are clustered by retailer brand. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 4: Heterogeneous Effects by Utilization Group: Card Counts and Credit Limits

	Δ Total # Cards			Δ Total Limit		
	0–50% (1)	50–90% (2)	90%+ (3)	0–50% (4)	50–90% (5)	90%+ (6)
$I(\text{Treated})$	−1.188*** (0.043)	−1.033*** (0.060)	−1.015*** (0.088)	−2008.629*** (282.426)	−1068.547*** (239.409)	−1112.074** (453.801)
Pre-shock average	3.88	4.08	3.03	28,616	16,544	9,858
Relative to pre-shock avg. (%)	−30.57	−25.34	−33.48	−7.02	−6.46	−11.28
VantageScore bin $\times$ cohort FE	✓	✓	✓	✓	✓	✓
ZIP (truncated) $\times$ cohort FE	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓
N	1,086,331	265,829	72,699	1,086,331	265,829	72,699
Adj. R^2	0.139	0.072	0.053	0.091	0.036	0.040

This table reports average treatment effect estimates from Equation (2) across utilization groups. Borrowers are partitioned into three bins based on pre-shock credit utilization: low (0–50%), medium (50–90%), and high (90%+). Reported coefficients correspond to $I(\text{Treated})$, with standard errors in parentheses. Outcomes are the change in the total number of credit cards and the change in total credit limit. Specifications include baseline controls (log age, log credit-history length, log monthly income, 90-day delinquency flag, and retail shares of cards, limits, and balances) and VantageScore-bin $\times$ cohort and ZIP $\times$ cohort fixed effects. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 5: Heterogeneous Effects by Utilization Group: Balances and Open Debt

	Δ Total Balance			Δ Open Debt		
	0–50% (1)	50–90% (2)	90%+ (3)	0–50% (4)	50–90% (5)	90%+ (6)
$I(\text{Treated})$	−162.260** (63.801)	−894.623*** (139.967)	−1053.285*** (247.137)	−427.160** (149.766)	−1335.870** (457.331)	−1776.537 (1024.446)
Pre-shock average	3,564	9,842	8,048	22,788	37,440	33,918
Relative to pre-shock avg. (%)	−4.55	−9.09	−13.09	−1.87	−3.57	−5.24
VantageScore bin $\times$ cohort FE	✓	✓	✓	✓	✓	✓
ZIP (truncated) $\times$ cohort FE	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓
N	1,086,331	265,829	72,699	1,031,657	252,905	69,080
Adj. R^2	0.029	0.095	0.091	0.028	0.057	0.051

This table reports average treatment effect estimates from Equation (2) across utilization groups. Borrowers are partitioned into three bins based on pre-shock credit utilization: low (0–50%), medium (50–90%), and high (90%+). Reported coefficients correspond to $I(\text{Treated})$, with standard errors in parentheses. Outcomes are the change in total credit card balance and the change in total open debt outstanding. Specifications include baseline controls (log age, log credit-history length, log monthly income, 90-day delinquency flag, and retail shares of cards, limits, and balances) and VantageScore-bin $\times$ cohort and ZIP $\times$ cohort fixed effects. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 6: Heterogeneous Effects by Credit Score Group

	Δ Total Limit		Δ Total Balance	
	Non-prime (1)	Prime (2)	Non-prime (3)	Prime (4)
$I(\text{Treated})$	-1214.182*** (231.977)	-2026.555*** (271.003)	-806.736*** (128.837)	-78.995 (65.498)
Pre-shock average	12,239	29,249	7,067	4,351
Relative to pre-shock avg. (%)	-9.92	-6.93	-11.41	-1.82
VantageScore bin $\times$ cohort FE	✓	✓	✓	✓
ZIP $\times$ cohort FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
N	321,814	1,103,045	321,814	1,103,045
Adj. R^2	0.031	0.097	0.073	0.041

This table reports average treatment effect estimates from Equation (2), comparing non-prime (VantageScore 3.0 < 660) and prime (VantageScore 3.0 ≥ 660) borrowers. Reported coefficients correspond to $I(\text{Treated})$, with standard errors in parentheses. Outcomes are the change in total credit limit and the change in total credit card balance. Specifications include baseline controls (log age, log credit-history length, log monthly income, 90-day delinquency flag, and retail shares of cards, limits, and balances) and fixed effects for VantageScore group $\times$ cohort and truncated ZIP $\times$ cohort. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Appendix

Online Appendix: Robustness Checks

This appendix reports three sets of robustness results that complement the empirical framework in Section 3 and Section 4. Unless stated otherwise, specifications mirror the baseline design: outcomes are expressed relative to $q = -1$, the control vector and fixed effects are identical, and standard errors are clustered by retailer.

A. Excluding Event Cohorts

To verify that our findings are not driven by mass card closures occurring in a single quarter, we re-estimate the baseline specification while sequentially excluding one event cohort (calendar quarter) at a time. That is, we sequentially exclude one cohort (calendar quarter) of mass card closures and re-estimate the average treatment effect from Equation (2) using the remaining 11 cohorts to assess robustness.

The results are summarized in Figure A.1 Panel A and Panel B. In both outcomes, the coefficient estimates remain negative, statistically significant, and stable in magnitude regardless of which cohort is excluded. This confirms that the baseline effects are not driven by closures during any particular period, including the COVID-19 quarters, and instead reflect a consistent pattern across the full sample.

[Insert Figure A.1 here]

B. Doubly Robust Estimates

We implement the doubly robust estimator of Callaway and Sant’Anna (2021) as a robustness check. We estimate group-time average treatment effects $ATT_{g,t}$ using inverse-probability weighting with outcome regression, restricting attention to never-treated controls. We aggregate to event time as in the baseline and report the dynamic path over event quarter $q \in [-4, 12]$. Standard errors are clustered by retailer.


[Insert Figure A.2 here]

C. Heterogeneous Effects on Credit Score

Figure A.3 plots VantageScore responses to mass card closure shocks by 10%-utilization bins and 50-point credit-score bins, as discussed in Section 4.2.


[Insert Figure A.3 here]

Figure A.1: Robustness to Excluding Event Cohorts




This figure reports robustness results in which we exclude one cohort (calendar quarter) of mass card closures at a time and re-estimate average treatment effects from Equation (2). The pre- and post-event horizons, control variables, and fixed effects are identical to the baseline specification. Dots plot coefficient estimates of the average treatment effect, and vertical capped lines denote 95% confidence intervals with standard errors clustered by retailer brand.

Figure A.2: Doubly Robust Dynamic Difference-in-Differences Estimates



This figure plots dynamic treatment effect estimates from Equation (1) using the doubly robust estimator of [Callaway and Sant’Anna \(2021\)](#). Panels A–D report results for retail credit card count, total credit limit, total credit balance, and credit utilization, respectively. The outcome is defined as the change relative to the omitted quarter ($q = -1$), with the mass card closure quarter normalized to $q = 0$. The indicator $Treated_i$ equals one for consumers whose retail card was closed at $q = 0$ and zero for the never-treated comparison group, so the coefficients β_q trace quarter-by-quarter treatment effects. Vertical capped lines denote 95% confidence intervals, with standard errors clustered by retailer brand. Controls, fixed effects, and sample selection follow Section 3.

Figure A.3: VantageScore Responses by Credit Score Bin and Utilization Bin



This figure reports average treatment effect estimates from Equation (2) across two types of consumer bins. Panel A displays estimates by 10% utilization bins, and Panel B displays estimates by 50-point VantageScore bins. For each bin, the blue line (left axis) shows the level change in VantageScore, and the red line (right axis) expresses the change as a percentage of the pre-treatment mean. The mass card closure quarter is normalized to $q = 0$, with averages computed relative to the pre-treatment window ($q = -4$ to $q = -1$). Vertical capped lines represent 95% confidence intervals, with standard errors clustered by retailer brand.