Working Papers

FEDERAL RESERVE BANK OF BOSTON

No. 25-10

The Credit Card Spending Channel of Monetary Policy: Micro Evidence from Account-level Data

Falk Bräuning and Joanna Stavins

Abstract:
Monetary policy impacts consumer spending via the effect of interest rate changes on credit card borrowing. Using supervisory account-level spending and balance data, we estimate that a 1 percentage point increase in the interest rate reduces credit card spending by nearly 9 percent and revolving balances by close to 4 percent. Aggregate results are primarily driven by revolving accounts, while we estimate small and statistically insignificant interest-rate elasticity for transaction accounts. Consistent with financial constraints, low-credit-score accounts tend to adjust spending, while high-credit-score accounts adjust balances.

JEL Classifications: D12, D14, E43, G21
Keywords: Credit cards, interest rates, consumer spending, borrowing

---

Falk Bräuning (falk.braeuning@bos.frb.org) is a vice president and economist in the Federal Reserve Bank of Boston Research Department. Joanna Stavins (joanna.stavins@bos.frb.org) is a principal economist and policy advisor in the Federal Reserve Bank of Boston Research Department.

The authors thank Liam Chentoufi and Julian Perry for providing excellent research assistance.

This version: September 2025 **https://doi.org/10.29412/res.wp.2025.10**

1 Introduction

Personal consumption expenditures are a major component of economic activity, accounting for about two-thirds of GDP in the United States. As such, academics and policymakers have long highlighted the role of consumer demand in business cycle dynamics (e.g., Blanchard, 1993; Mian, Rao and Sufi, 2013; Matthes and Schwartzman, 2025). For example, consumption weakness was a key driver of the 1990–1991 recession (Blanchard, 1993), while, more recently, the post–COVID-19 recovery has been driven by strong consumer spending. Thus, an important question for monetary policy arises: How sensitive is consumer spending to changes in interest rates? We argue that if interest-rate elasticity is high, monetary policy is more potent in steering aggregate demand.

In this paper, we highlight one specific channel, namely how interest rate changes impact personal consumption expenditures by affecting the cost of credit card borrowing. Understanding this transmission channel is critical, as credit card payments have increased over the past several years, both in their total dollar value and their share of consumer spending (Figure 1, Panel a). On average, credit card payments increased in value by more than 10 percent per year during the past decade and amounted to $5.83 trillion in 2022 (Federal Reserve Payments Study)[1]. They constitute about 20 percent of total consumer spending (Diary of Consumer Payments Choice)[2], and, importantly, about 30 percent of credit card expenditures by consumers who are debt-financed—that is, consumers with revolving account balances (Figure 1, Panel b). These revolving balances, as well as any new spending on revolving accounts, carry interest expenses. Thus, the interest-rate elasticity of credit card spending is a key statistic for policy and model calibration.

Because interest rates are endogenous to macroeconomic conditions, our key identification challenge is to isolate the effect of interest rates on credit card spending. For empirical identification of interest-rate elasticity, we leverage microdata at the individual account level. These data are collected for bank stress-testing purposes and provide a unique window into consumer spending behavior, covering nearly 80 percent of all US credit cards. For each account, for each month, we observe the following key variables: (1) spending (flow), (2) debt balance (stock), and (3) interest rate (annual percentage rate, or APR), as well as other rich contract and borrower characteristics. Importantly, given the panel structure of the data, we observe spending and balance on the same card over time at a monthly frequency.

Our identification scheme leverages the predominant contractual practice whereby nearly all accounts are variable-rate, with the APR equaling a constant margin over a variable

[1] See https://www.federalreserve.gov/paymentsystems/fr-payments-study.htm

[2] See https://www.atlantafed.org/banking-and-payments/consumer-payments/survey-and-diary-of-consumer-payment-choice

Figure 1: Importance of Credit Card Spending and Debt for Consumer Spending

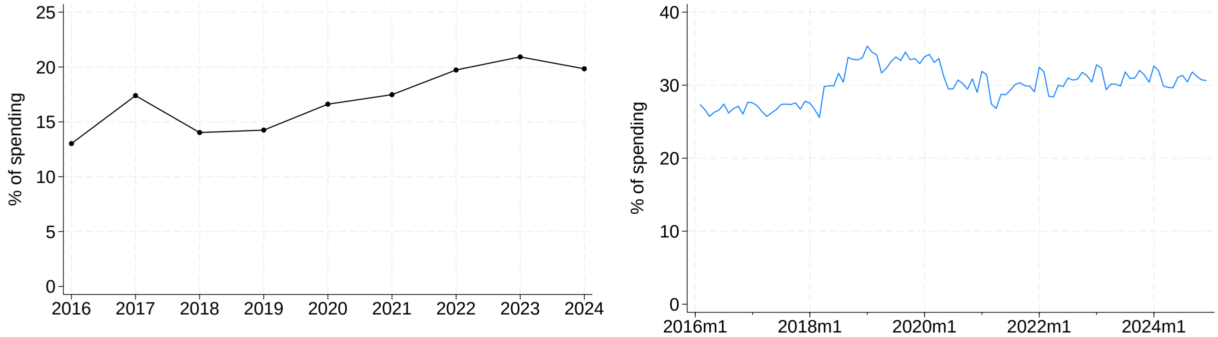

(a) Share of Credit Card Spending (b) Share of Spending by Revolvers

Notes: Panel (a) shows the percentage of purchases by dollar value made using credit cards, calculated using the Survey and Diary of Consumer Payment Choice. The sample consists of transactions made in October of each year, and observations are weighted to account for day-of-week effects. Panel (b) shows the percentage of credit card spending made by accounts with a positive revolving balance using Federal Reserve Y-14M data. In Panel (b), one bank was removed from the share calculation for all dates preceding March of 2018 due to a change in reporting methodology. See Appendix Figure A.1 for robustness. *Source*: SDCPC and Y-14M, authors' calculations.

index rate, typically the prime rate. Moreover, we leverage the prevalence of contractual APR ceilings, meaning that the APR is bound to not exceed a certain contractually defined maximum. These ceilings introduce a kink in the relationship between the APR and the index rate, allowing us to apply a regression kink design (RKD) for causal identification. The RKD identifies a local average treatment effect near the kink.

Our results reveal a sizable response of credit card spending to interest rate changes. Using spending-weighted regressions, we estimate an aggregate interest-rate elasticity of about –8.7. That is, for a 1 percentage point increase in the APR, credit card spending in the following month declines 8.7 percent. We identify these effects from changes in spending within the same credit card over time. Supporting the validity of our RKD for causal inference, our results are robust to a wide range of controls, including bank-by-month fixed effects, state-by-month fixed effects, and credit-score-by-month fixed effects.

An increase in the interest rate not only curtails credit card spending, but also leads to a deleveraging, as outstanding credit card balances decline about 3.7 percent in response to a 1 percentage point higher interest rate in the preceding month. While low-credit-score accounts decrease their spending in response to an interest rate increase, high-credit-score accounts pay down their existing revolving balances. That difference between the responses of low- and high-credit-score accounts is consistent with underlying financial frictions, as low-credit-score account holders are more likely to be liquidity-constrained and lack other sources of debt.

We also estimate heterogeneous elasticities. In particular, we show that aggregate results are driven by revolving accounts—those accounts with a positive carryover balance. Crucially, these accounts accrue interest not only on the existing balance, but also on any new spending. The estimated interest-rate elasticity for revolving accounts is –15, indicating that in response to a 1 percentage point increase in APR, credit card spending on revolving accounts declines 15 percent in the following month. We estimate the interest-rate elasticity of revolving credit card balances at –3.8. By contrast, for transaction accounts—defined as those without a positive balance at the beginning of the statement period—we find (in absolute value) small and statistically insignificant interest-rate elasticity. This is intuitive since, for transactors, spending is not subject to interest charges.

We then study the role of financial constraints. Separating credit card accounts by credit score shows that only spending on accounts held by low-credit-score consumers is affected and that the estimated interest-rate elasticity of spending is –18, while the estimated elasticity for high-credit-score consumers is small and statistically insignificant.[3]

This finding is consistent with the literature showing that low-credit-score households are much more financially constrained than high-credit-score households. For example, Agarwal et al. (2018) find that when credit limits expand, low-credit-score households increase their spending, whereas there is no effect on high-credit-score households.

However, high-credit-score accounts see a significant deleveraging when interest rates increase, with balances declining 7 percent after a 1 percentage point increase in APR. We do not find such a response in balances for low-credit-score consumers, consistent with limited access to credit and liquidity that would allow them to pay down their revolving credit card balances while maintaining spending.

Although our RKD provides credible causal identification, the estimated elasticities are local average treatment effects obtained from a subset of accounts near the maximum APR. To understand the relationship between monetary policy rate changes and total credit card spending growth more broadly, we complement our RKD analysis with aggregate analysis. Specifically, we use local projections in conjunction with monetary policy shocks from Jarociński and Karadi (2020) as instruments for endogenous federal funds rate changes. Results confirm a sizable reduction in spending growth of about 9 percentage points in response to a 1 percentage point funds rate increase. Importantly, the peak effect is delayed by about two months after the funds rate increase, which is consistent with APR changes taking as long as two months to fully reflect funds rate changes as well as potential consumer inattention.

[3] A low credit score is defined as a credit score below the median score in a given month, while a high credit score is above the median.

Related Literature Our paper relates to several strands of the literature. First, it aligns with the large body of work studying the transmission of monetary policy to consumption. Early research focuses on the response of aggregate consumption to monetary policy changes using time-series analysis (e.g. Bernanke and Blinder, 1992; Christiano, Eichenbaum and Evans, 1996). One key transmission channel highlighted in the literature is the wealth effect, as discussed in Lettau, Ludvigson and Steindel (2002). More recent papers focus on heterogeneous responses across households. For example, Berg et al. (2021) show that older households are more responsive than young or middle-aged households, attributing the differences to stronger wealth effects. Our paper contributes to this literature by identifying one specific channel related to debt-financed credit card spending and highlighting heterogeneous and aggregate effects.

Second, our paper relates most closely to the limited literature on the effect of interest rates on credit card spending. Nelson (2025) assesses the effect of changes in credit card lenders' ability to raise interest rates in response to changes in borrowers' credit risk, not in response to prime rate changes. Grodzicki (2023) uses credit card account data for the 2016–2019 period to estimate the effect of APR changes on credit card spending and finds that spending and borrowing decline when the federal funds rate (FFR) rises. Grodzicki et al. (2023) use the same data for the 2009–2011 period and find that both spending and borrowing on credit cards decline when interest rates rise, but the response to interest rate changes is not uniform across cardholders and varies with the credit-score quantile. Grigoli and Sandri (2023) use credit card data from Germany and find that credit card spending responds more quickly to changes in short-term interest rates than to changes in long-term interest rates.

Third, our paper relates to the literature studying the role of financial constraints, such as credit and liquidity constraints, on consumers. For example, early work by Zeldes (1989) shows that liquidity constraints play a significant role in consumption behavior. Gross and Souleles (2002) show that consumers use credit cards to smooth consumption during periods of income volatility, but this ability is constrained by credit limits and interest rates. They find that low-credit-score households increase spending when their credit limits expand. In addition, low-credit-score households are much more financially constrained than high-credit-score households: When credit limits expand, low-credit-score households spend more, while high-credit-score households do not increase their spending Agarwal et al. (2018). More recent work by Baugh et al. (2021) documents consumption smoothing behavior that challenges pure liquidity constraints or hand-to-mouth behavior. The findings are not limited to credit cards: Di Maggio et al. (2017) find that lower-credit-score households' consumption responds to changes in mortgage rates, suggesting financial frictions.

Our finding that, in response to interest rates, spending and debt changes vary by low- versus high-credit-score accounts also points to financial frictions.

The remainder of this paper is structured as follows. Section 2 describes the data. Section 3 provides institutional details on credit card interest rates, and Section 4 discusses our empirical approach. Section 5 presents our core results. Section 6 concludes.

2 Data Description

Our main data source is the Y-14M credit card account data. The Board of Governors of the Federal Reserve System collects Y-14M data monthly from bank holding companies (BHCs), intermediary holding companies (IHCs) of foreign banking organizations (FBOs), and covered savings and loan holding companies (SLHCs) with total consolidated assets of $100 billion or more; the Board of Governors uses the data in supervisory capital assessments and stress-test models. These institutions are required to report information on all the credit card accounts they have on file on the last business day of each calendar month. The accounts reported in the Y-14M data represent about three-quarters of the total bank card balances in the United States.

Given the large sample, we use a 1 percent random sample of active, non-fraudulent, non-corporate/business Y-14M accounts from January 2016 to December 2024, limiting our analysis to banks that reported their credit card portfolios for the entirety of this period. From these, we exclude accounts not indexed to the prime rate, accounts with a non-monthly APR reset frequency, and accounts covered by a promotional APR. Although cards with promotional APRs constitute a substantial share of the overall credit card market (Drozd and Kowalik, 2023), this paper examines changes in credit card use in response to changes in APR that take place when the FFR changes. Promotional interest rates typically increase dramatically when the promotional period ends, regardless of whether the FFR rises.

Notably, the Y-14M data do not provide account holder identification, so we cannot link multiple accounts by the same holder; that is, we cannot examine account-holder-level spending. Instead, our analysis focuses on account-level variation. Given our aggregate focus, however, we use weighted regressions to account for heterogeneous spending across accounts.

Among these observations, we further limit our analysis to accounts with a listed maximum APR, which account for more than two-thirds of the remaining observations. Although there is no legal limit to the APR that banks can provide to most consumers, most contracts have historically included language preventing the APR from rising above 30 percent.[4] We also

[4]See https://www.usatoday.com/story/money/2023/11/01/credit-card-interest-rates-climb-to-33-percent/71391622007/

Table 1: Descriptive Statistics

	mean	sd	p1	p25	p50	p75	p99	count
Panel A: Full Sample								
Spending (2016 $)	781.20	2,240.46	0.00	0.00	121.31	733.76	8,392.42	79,873,512
Balance (2016 $)	1,755.31	3,475.50	-44.91	2.47	401.97	1,949.85	16,927.36	79,872,558
Payment (2016 $)	802.50	2,369.72	0.00	1.85	146.42	668.33	9,129.04	79,873,512
Utilization	0.24	0.34	-0.01	0.00	0.06	0.38	1.03	79,872,486
APR	19.13	5.03	7.49	15.49	18.99	22.99	29.24	79,873,512
Dist. to Max APR	-11.02	5.40	-23.00	-14.75	-11.00	-7.00	-0.75	79,873,512
Max APR	30.17	1.93	29.99	29.99	29.99	29.99	36.00	79,873,512
APR Margin	13.84	4.48	3.15	10.99	12.74	17.49	23.99	79,873,512
Income at Orig. ($1000s)	97.84	1,479.71	0.00	37.99	63.00	100.00	499.99	1,583,305
Credit Score Decile	5.51	2.76	1.00	3.00	6.00	8.00	10.00	78,955,869

	mean	sd	p1	p25	p50	p75	p99	count
Panel B: RKD Sample								
Spending (2016 $)	856.18	1,856.86	3.79	76.16	294.97	956.31	7,389.78	2,108,158
Balance (2016 $)	2,058.16	3,039.69	4.69	268.23	946.19	2,633.33	14,429.86	2,108,158
Payment (2016 $)	811.47	1,955.80	0.00	71.84	228.43	786.31	7,947.19	2,108,158
Utilization	0.45	0.39	0.00	0.07	0.34	0.87	1.07	2,108,158
APR	28.59	0.82	27.74	28.24	28.49	28.99	29.99	2,108,158
Dist. to Max APR	-1.29	0.69	-2.00	-1.75	-1.50	-1.00	1.74	2,108,158
Max APR	29.94	0.69	29.99	29.99	29.99	29.99	29.99	2,108,158
APR Margin	20.71	1.62	19.49	19.74	20.49	20.74	26.99	2,108,158
Income at Orig. ($1000s)	84.96	125.94	2.00	38.40	60.00	98.75	450.00	44,457
Credit Score Decile	3.84	2.48	1.00	2.00	3.00	5.00	10.00	2,056,666

Note: The full sample is a 1 percent random sample of the population. Both samples are limited to accounts at relevant banks with non-promotional APRs and monthly reset frequencies. For each account, an observation is included for every month from January 2016 through December 2024 in which the account was open and active and had no known fraudulent activity. The RKD sample is further restricted to account-months with a positive balance, positive spending, and an APR no higher than 2 percentage points from the maximum. The utilization variable is defined as the balance/credit limit.

exclude all observations from one bank whose maximum APR variable had at times been erroneously reported, such that some of its accounts carried an APR higher than the reported maximum. Lastly, we remove a small number of accounts for which the reported maximum APR or the reported interest rate margin changed over the course of the months we observe.

Table 1, Panel A, displays descriptive statistics for the full sample, which includes nearly 80 million account-months from 13 different banks. Recall that, as mentioned earlier, our sample is a random 1 percent subset of the population. The number of observations with non-missing income is smaller than it is for the other variables, and this variable is not used in most of our analysis. Each variable associated with account activity—spending, balance, payment amount, and utilization—has a much higher mean than median, suggesting a right-tailed distribution of credit card activity, consistent with a skewed income and distribution of households. Average monthly spending per account equals $781 (constant 2016 dollars), and median spending is much smaller at $121. The average balance is $1,755, and the median is $402.[5]

Table 1, Panel B, shows the same descriptive statistics for the "RKD sample" that is used for our main elasticity estimates. This sample is limited to account-months with positive spending and positive balance (such that the log of both variables is non-missing) as well as accounts for which our running variable is within a 2 percentage point bandwidth around the maximum APR. The construction of this running variable is discussed further in Section 3. Observations in this bandwidth have higher spending, balance, payment amount, and utilization than the observations in the full sample shown in Panel A. Accounts in the RKD sample also tend to have a higher APR and APR margin than the full sample and a lower income and credit score. Credit score is measured in decile cohorts (as described later), and the median credit score decile for account holders in the RKD sample is 3, compared with 6 for the full sample.

In our regressions, we include fixed effects for the additional variables of account, month, state, bank, and credit score. Because different banks report accounts' credit scores using different score types, they cannot be easily converted to a single scale. Instead, we calculate percentiles within each month–score-type cell and assign observations to buckets 1 through 10 based on their decile within that cell.

3 Stylized Facts about APRs

What is the empirical relationship between the monetary policy interest rate and the interest rate on a credit card account? Empirically, Figure 2 shows—using the data underlying our analysis—that the average APR closely tracks changes in the FFR. This finding is consistent with Grodzicki (2023), who reports that interest rates on credit card accounts adjust fully to a change in the FFR within three months of a change in the FFR.

This close relationship between the policy rate and the APR is a direct consequence of the vast majority of credit card accounts being variable- rate contracts. The variable part is the base rate (or index rate), which is predominantly the prime rate, or FFR plus 300 basis points. The APR is then obtained by adding a margin (spread) to the prime rate. The prime rate is typically retrieved by the bank at the end of the month and applied to the billing

[5]Note that a few observations are actually negative, meaning that the account holder overpaid their balance the previous month, received a refund, earned statement credit, or was otherwise owed money by the bank.

Figure 2: Federal Funds Rate, Prime Rate, and Average APR

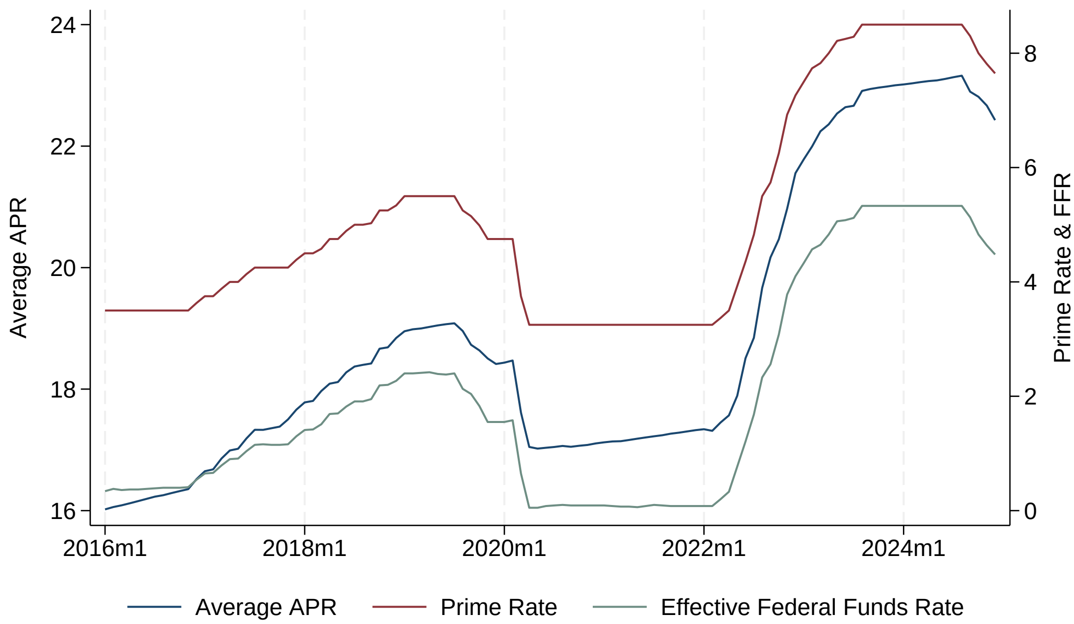

Note: Average APR is computed from the Y-14M data shown on the left-hand side scale. Federal funds rate and prime rate are shown on the right scale. Prime rate equals FFR plus 300 basis points. *Source*: Haver, Y-14M, authors' calculations.

cycle that includes that reference day.[6]

Given the pervasive use of the prime rate as the index rate, variation in APR is driven mostly by heterogeneity in margin.

Figure 3 shows the distribution of margin. There is substantial heterogeneity in margin. The distribution has a large range (the 1st percentile is 3 percent and the 99th percentile is 24 percent) and is multi-modal, with some clustering of observations above 10 and below and above 20. Because all credit cards index to the prime rate, accounts that have a higher APR, by construction, have higher margin.

Why the variation in margin? As one may expect, and as shown in Panel (b) of the preceding figure, margin correlates strongly with credit quality, as measured by our computed credit-score deciles. The lowest-credit-quality account holders, in bin 1, have a median margin of 19 percent compared with a margin of only 11 percent for the highest-credit- quality account holders (bin 10). The figure also reveals that, within the credit-score bins, there is still some variation in margin, due potentially to our coarse grouping of credit scores or other factors, such as bank-level heterogeneity.

[6] In our data, we find that billing cycles end and start days are relatively uniformly distributed across the month. See Appendix Figure A.4 for additional details.

Figure 3: Margin and Credit Score

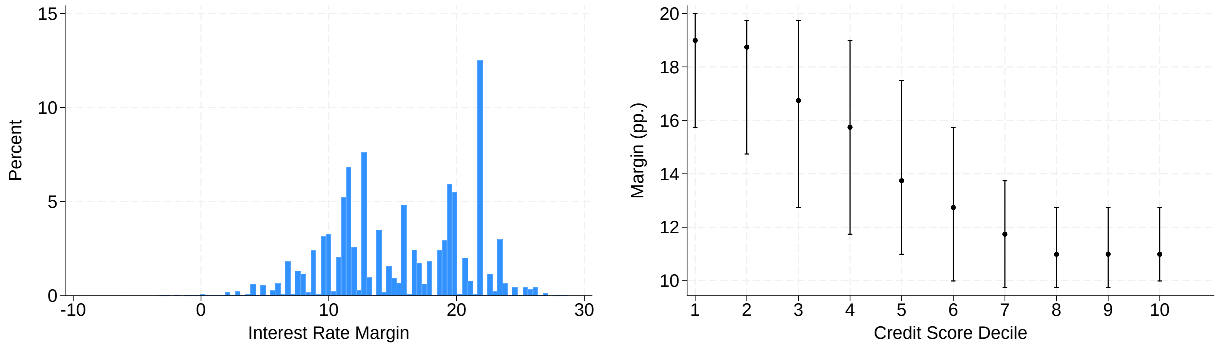

(a) Margin

(b) Margin and Credit Score

Note: Panel (a) shows the distribution of credit card interest rate margin, calculated using Y-14M credit card account data. Panel (b) shows the 25th, 50th, and 75th-percentiles of interest rate margin for observations within each credit score decile, calculated using Y-14M credit card account data. *Source*: Y-14M, authors' calculations.

Another key feature of credit card APRs—which is at the core of our empirical identification strategy—is that the APR is prohibited from increasing above a maximum; the "ceiling" is the maximum APR level specified in the credit agreement. Among accounts with a maximum APR, the vast majority report a maximum APR of 29.99 percent, consistent with the upper-bound industry norm of 30 percent. Given the APR ceiling, the benchmark (variable) interest rate of credit card account l at time t becomes

$$r_{l,t} = \min(\text{APR Ceiling}_l, \text{ Prime Rate}_t + \text{Margin}_l). \tag{1}$$

The minimum function in the APR on the credit card account introduces a kink in the relationship between the index rate and the interest rate on the card: There is a one-to-one relationship between the prime rate and the interest rate whenever the prime rate plus the margin is below the maximum APR, but when the prime rate plus the margin is above the maximum APR, the interest rate is flat and does not respond to prime rate changes.

An important variable in our empirical analysis is the distance to the kink, that is, the ceiling in our application. We define the so-called running variable as

$$d_{l,t} \equiv (\text{Prime Rate}_t + \text{Margin}_l) - \text{APR Ceiling}_l,$$

or the difference between the contractual APR ceiling and the base rate plus the margin.

The running variable is also called the forcing variable in terms of the RKD terminology.[7] According to our definition, a negative distance means that the APR is not bound by the contractual ceiling.

Figure 4: Kink in the APR

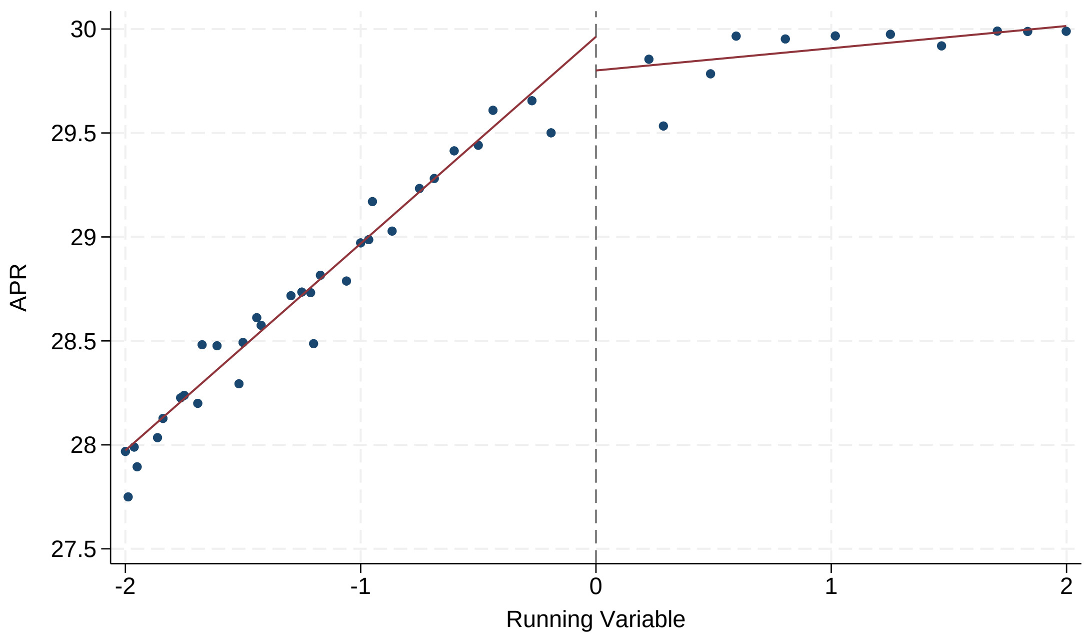

Note: The figure shows the average APR across values of the running variable among accounts with a revolving balance, calculated using Y-14M credit card account data. The running variable is equal to the prime rate plus interest rate margin minus maximum APR. The sample is limited to observations with values from −2 to 2, excluding those with a value of 0. Accounts whose interest rate margin and maximum APR fluctuate are also omitted. *Sources*: Y-14M, authors' calculations.

Figure 4 shows the relationship between the distance to the maximum APR (running variable) and the actual APR using a binned scatter plot, with a linear fit on both sides of the kink point. The relationship is very close—as one would expect, given how APRs are determined. With a negative distance, the actual APR increases one-for-one with the running variable, while for a positive distance. the relationship is completely flat, meaning that the actual APR does not respond to prime rate changes (because it is bound by the contractual ceiling).

Figure 5, Panel (a), shows the sample distribution of the distance to the maximum APR. For the large mass of account-months, the APR is below its contractual maximum. Note again that the maximum APR in this market is relatively common, with the vast majority of

[7] Note that we can rewrite the APR on the credit card account as a function of the distance: $r_{l,t} = \text{APR Ceiling}_l + \min(\mathrm{d}_{l,t}, 0)$.

Figure 5: Distance to Maximum APR and Share of Accounts at Maximum APR

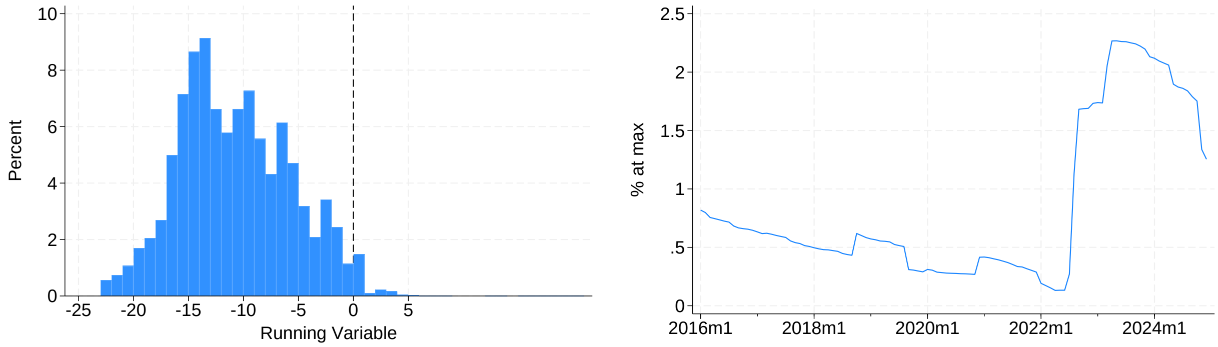

(a) Distance to Max APR

(b) Share Account at Max APR

Note: Panel (a) shows the distribution of the running variable, calculated using Y-14M credit card account data. The running variable is equal to the prime rate plus interest rate margin minus maximum APR. The bottom 1 percent of values is not shown. Panel (b) shows the percentage of accounts with the maximum possible APR, calculated using Y-14M credit card account data. The sample is restricted to accounts whose maximum APR remains constant over the January 2016–December 2024 period. *Sources*: Y-14M, authors' calculations.

accounts having a maximum APR close to or equal to 30 percent (typically 29.99 percent). Also, given that margins are typically constant within accounts over time, the share of observations that are at or near the maximum APR is a sole function of the level of the FFR and hence the prime rate. As the interest rate increases, more accounts will hit their maximum APR. Panel (b) illustrates this point by showing that, during the strong monetary tightening with large and fast funds rate hikes in 2022 and 2023, the share of accounts at the maximum APR increased from less than 0.5 percent to more than 2 percent.

4 Empirical Approach

Our goal is to measure the interest-rate elasticity of credit card spending and debt. The key identification challenge is that interest rates can be endogenous to the economy. For example, when the economy runs hot, consumer spending may be high and interest rates may increase as monetary policy tightens. We address this endogeneity of interest rates by using the regression kink design, following Card et al. (2015), Card et al. (2016) and its recent application by Indarte (2023). In our setting, the RKD is possible due to the prevalence of interest rate ceilings on the APR.

The basic idea of the RKD is to examine the change in the slope of the relationship between the outcome of interest (credit card spending and balance) and the running variable (applicable interest rate) at the exact location of the kink imposed by the rule, that is, the APR ceiling in our case. It is important to stress that, for the RKD to be valid, the market rule that governs the ceiling or movements in the index rate need not be exogenous. Provided that observations on either side of the kink threshold are similar—that is, they have a smooth density function at the threshold—any kink in the outcome can be attributed to the treatment effect of the policy variable (see Card et al. (2015) for technical details and additional standard regularity assumptions). Simply put, if we observe a kink in credit card spending when the APR ceiling becomes binding, we causally attribute this change to the change in APR.

We express credit card spending (our main outcome variable of interest) as a function of this distance, $c_{l,t} = c_{l,t}(d_{l,t})$, where we chose symbol c for consumption. Dropping subscripts to simplify notation, the local average treatment effect can be written as

$$\tau = \frac{\lim_{d_0 \to 0^+} \frac{c(d)}{d}\Big|_{d=d_0} - \lim_{d_0 \to 0^-} \frac{c(d)}{d}\Big|_{d=d_0}}{\lim_{d_0 \to 0^+} \frac{r^l}{d}\Big|_{d=d_0} - \lim_{d_0 \to 0^-} \frac{r^l}{d}\Big|_{d=d_0}}, \tag{2}$$

that is, the change in the slope of the outcome variable (numerator) scaled by the change in the slope of the first stage (denominator). Note that the denominator simplifies to 1 with equation (1), in which case the elasticity estimate is the change in the slope of credit card spending at the threshold. As we discuss later, we use a fuzzy design whereby the denominator can be different from 1.

The empirical implementation of the RKD involves estimating regression models for credit card spending (or any other outcome variable) and the interest rate (APR) for observations "close" to the kink using (local) polynomial regressions, similar to a regression discontinuity design. Instead of estimating a shift in the intercept, however, in the kink design, we are interested in estimating a slope change. We later discuss the exact empirical regression model used to estimate the change in slopes at the kink.

Our empirical estimation focuses on estimating the kink in the relationship between credit card spending and the distance to the APR ceiling. We do so by estimating the following regression model:

$$c_{l,t} = \begin{cases} \beta_1 d_{l,t} + f_1(d_{l,t}) + \theta_1 X_{l,t} + u_{l,t}, & \text{for } d_{l,t} < 0 \\ (\beta_1 + \beta_2) d_{l,t} + f_2(d_{l,t}) + \theta_2 X_{l,t} + u_{l,t}, & \text{for } d_{l,t} \geq 0, \end{cases} \tag{3}$$

where $c_{l,t}$ is the logarithm of spending on the card, and $d_{l,t}$ is the distance to the kink, as defined earlier. This number can be negative, indicating that the APR is below the ceiling. The piecewise regression equation is common in the RKD literature and means that we allow all parameters to vary freely on either side of the kink point; that is, we do not enforce pooled parameters.

The coefficient β_1 equals the derivative of credit card spending with respect to the distance as the distance approaches zero from above, and the coefficient β_2 measures the *change* in the derivative at $d_{l,t} = 0$, at the kink point. Flexible polynomial functions f_1 and f_2 can be included to account for a potential nonlinear relationship between the distance to the floor and credit card spending away from the kink point for the linear terms to accurately measure the derivative as the distance approaches zero. In our baseline specifications, we follow Card et al. (2015) and include polynomial functions of order 1, that is, linear terms in the distance. This choice is also consistent with the results of a statistical test for the optimal polynomial order, as outlined in Pei et al. (2022) (see Appendix Table A.1 for the test results).

The vector $X_{l,t}$ collects all control variables. We also allow the coefficients on these controls to change flexibly depending on whether the floor is binding. Throughout the analysis, we control for various sets of granular fixed effects. Importantly, we include account fixed effects and track spending and APR of the same credit card account over time so that the identification is driven by within-loan variation. This means that changes in APR come from variation in the base rate. As noted, changes in the base rate are likely to be endogenous, and to the degree that such changes are correlated with changes in economic activity, our RKD helps us overcome that endogeneity. Our most saturated models also include interactions of bank fixed effects, credit-score-group fixed effects, state fixed effects, and month fixed effects, thereby accounting for heterogeneity in narrowly defined cells of the data. This allows us to purge locational and credit-quality-related factors.

Because we focus on estimating the effects of interest rate changes on consumption from an aggregate viewpoint, we estimate weighted regression models using credit card spending in dollars as weights. Similarly, if the outcome variable is the balance on the account, we weight regressions with the dollar balance of the account.

The Y-14M data contain detailed contract information at the individual account level, including the APR, the base rate type, and the reset frequency. Yet, given the complexity of contractual interest rate schedules, the data may not fully capture mapping between the base rate and the applicable interest rate. Lack of such detail requires a fuzzy RKD in which the "first-stage" assignment rule—in our case the mapping between the index rate and the APR on the account—is also modeled as a regression. As with the credit card spending, we therefore estimate a similar piecewise model for the account's APR:

$$r_{l,t} = \begin{cases} b_1 d_{l,t} + g_1(d_{l,t}) + \gamma_1 X_{l,t} + e_{l,t}, & \text{for } d_{l,t} < 0 \\ (b_1 + b_2) d_{l,t} + g_2(d_{l,t}) + \gamma_2 X_{l,t} + e_{l,t}, & \text{for } d_{l,t} \geq 0. \end{cases} \tag{4}$$

The key parameter of interest is b_2, the change in the slope of the relationship between the (observed) APR and the distance at the kink point.

As is common practice in RKD applications (Card et al., 2016), we estimate both spending and APR models for observations within a given bandwidth around the interest rate floors, which allows for a narrow identification of the slope near the kink. The baseline results are estimated for a bandwidth of 2 percent around the floor, and we show the robustness of the effect for several bandwidth selections later. Moreover, we adopt a commonly used uniform kernel implying equal weighting of observations in our regressions.

The fuzzy RKD estimate of the interest-rate elasticity of credit card spending is then obtained as the ratio of the estimated changes in the slopes of the spending and the estimated change in the slope of the interest rate:

$$\hat{\tau} = \frac{\hat{\beta}_2}{\hat{b}_2}. \tag{5}$$

We compute standard errors using the delta method and based on multi-way clustered errors at the account and time levels.

While RKD allows us to identify the interest-rate elasticity, we should acknowledge that the RKD approach identifies a local average treatment effect (LATE) using only the interest rates of credit card accounts that are in the neighborhood of their respective APR ceiling. Given that there is heterogeneity in the applicable interest rate across accounts, the LATE estimate could be not be representative of the average treatment effect (ATE).[8] Specifically, when we look at the spending of accounts with APRs near their maximum, we focus disproportionately on a period of higher interest rates and on accounts with a higher margin, as discussed in Section 3. In other words, the RKD approach is subject to concerns of external validity, as elasticities may vary throughout the business cycle.[9] In addition, elasticities may vary across accounts, and we provide direct evidence of this point in the next section.

5 Results

Baseline Results Table 2 presents our baseline RKD estimates of the interest-rate elasticity of credit card spending. The elasticity estimates are reported in Panel A (top). The different columns correspond to different fixed effect models, as indicated at the bottom of the table. The results show relatively stable elasticity estimates from –6.5 to –8.7, supporting evidence

[8]See Angrist and Fernández-Val (2013) for a discussion.

[9]Given the short sample period and the fact that maximum APRs are less frequently binding during easing periods, we cannot estimate state-dependent models.

of the validity of the identification approach. Column (1) includes only account fixed effects to control for difference in spending levels across accounts. Using this specification, we estimate an elasticity of –6.44, but standard errors are still large, given confounding variation. Column (2) shows that, once we also include month FE, state FE, and credit-score decile FE, the standard errors decline substantially, with a small increase in the point estimate, which in turn becomes statistically significant.

The most saturated model in column (4)—which uses variation in credit card spending across accounts within the same bank-month-state-risk while also partialling out time-invariant account-specific effects—indicates an elasticity of –8.7. When the interest rate increases 1 percentage point, spending decreases a sizable 8.7 percent.[10] To put this estimate further into perspective, for an average cardholder in the RKD sample, an 8.66 percent elasticity means that a 1 percentage point increase in the interest rate leads to a $74 decrease in monthly spending (in 2016 dollars).

Panel B (middle) and Panel C (bottom) zoom in on the two different components that define the RKD estimand, that is, the change in the slope of the APR and the change in the slope of the spending at the kink point. Consistent with the unconditional figure, Panel B shows that the APR exhibits a tight relationship with the prime rate but not once the APR hits the ceiling, as indicated by a tightly estimated slope change of negative one at the kink point. Panel C reports a positive slope change for the relationship between the running variable and log spending. Again, in the least-saturated specification shown in column (1), estimates in Panel C are insignificant. Once we add time effects, the change in slope (interaction term) becomes highly significant, but we obviously cannot identify the distance-to-ceiling effect. Note that the elasticity estimate closely resembles the estimates of Panel C, given that the denominator of the RKD estimand is close to one.

While our analysis focuses mainly on credit card spending as an important component of consumption, we also study what interest rate changes mean for the account balance. We do so by estimating a similar RKD but replacing the outcome variable log(spending) with log(balance). Balance is measured in the Y-14M data as an end-of-reporting-month stock. While one may view the immediate margin of adjustment as new spending, account holders may also pay down existing debt in response to higher interest rates. In either case, relative to the counterfactual, credit card balances may be expected to decline. Again, the elasticity may be expected to vary across account types (revolvers versus transactors, and high versus low credit score). Table 3 reports the results of estimating the elasticity of credit

[10]Because our results are based on spending-weighted estimates, these numbers represent aggregate effects of interest changes. Our baseline results present elasticities for spending one month ahead; Appendix Figure A.5 shows different response horizons. Appendix Figure A.6 shows the robustness of our RKD estimates to different fixed effects and bandwidth choices.

Table 2: RKD Estimates of Interest-rate Elasticity of Credit Card Spending

	(1)	(2)	(3)	(4)
Panel A: APR Elasticity				
Elasticity	-6.44	-7.43***	-6.50**	-8.66***
	(4.36)	(2.30)	(2.88)	(2.91)
Panel B: Dep. Var. is APR				
Dist to Ceiling	0.010***			
	(0.000)			
At Ceiling * Dist to Ceiling	-0.010***	-0.010***	-0.009***	-0.009***
	(0.000)	(0.000)	(0.000)	(0.000)
Panel C: Dep. Var. is ln(f.spending)				
Dist to Ceiling	-0.048			
	(0.037)			
At Ceiling * Dist to Ceiling	0.065	0.075***	0.061**	0.082***
	(0.044)	(0.023)	(0.027)	(0.027)
Account FE	Yes	Yes	Yes	Yes
Month, State, Credit Score FE	No	Yes	Yes	Yes
Bank*Month FE	No	No	Yes	Yes
State*Month FE	No	No	Yes	Yes
Credit Score*Month FE	No	No	Yes	Yes
Credit Score*Month FE	No	No	Yes	Yes
State*Bank*Credit Score*Month FE	No	No	No	Yes
N	1984112	1925069	1924856	1899757

Note: The RKD sample is defined subject to specifications outlined in the "Data Description" section using observations for which the running variable is within a 2 percentage point bandwidth (inclusive) around the maximum APR. Standard errors are clustered by account and month. Estimates are weighted by an account-month's share of total spending that occurred in the subsequent month, winsorized at the 97.5th percentile.

card balances with respect to interest rates. We find that a 1 percentage point increase in APR is associated with a 3.7 percent decline in credit card balances.

Revolvers versus Transactors We next zoom in on the mechanism. Credit card spending should be most sensitive to APR changes if the credit card account carries a positive balance, such that spending additional dollars accrues interest. An account that carries a positive unpaid balance from the previous month is referred to as revolving. Revolving accounts accrue interest on the unpaid balance as well as on any new transactions, and cardholders

Table 3: RKD Estimates of Interest-rate Elasticity of Credit Card Balance

	(1)	(2)	(3)	(4)
Panel A: APR Elasticity				
Elasticity	-1.94	-2.29	-3.70**	-3.71*
	(4.26)	(1.65)	(1.82)	(2.02)
Panel B: Dep. Var. is APR				
Dist to Ceiling	0.010***			
	(0.000)			
At Ceiling * Dist to Ceiling	-0.010***	-0.010***	-0.010***	-0.010***
	(0.000)	(0.000)	(0.000)	(0.000)
Panel C: Dep. Var. is ln(f.balance)				
Dist to Ceiling	0.002			
	(0.036)			
At Ceiling * Dist to Ceiling	0.020	0.023	0.035**	0.035*
	(0.043)	(0.016)	(0.017)	(0.019)
Account FE	Yes	Yes	Yes	Yes
Month, State, Credit Score FE	No	Yes	Yes	Yes
Bank*Month FE	No	No	Yes	Yes
State*Month FE	No	No	Yes	Yes
Credit Score*Month FE	No	No	Yes	Yes
Credit Score*Month FE	No	No	Yes	Yes
State*Bank*Credit Score*Month FE	No	No	No	Yes
N	1984112	1925069	1924856	1899757

Note: The RKD sample is defined subject to specifications outlined in the "Data Description" section using observations for which the running variable is within a 2 percentage point bandwidth (inclusive) around the maximum APR. Standard errors are clustered by account and month. Estimates are weighted by an account-month's share of total balances that occurred in the subsequent month, winsorized at the 97.5th percentile.

who carry unpaid balances are called revolvers. By contrast, cardholders who pay their balance in full at the end of each billing cycle are called convenience users or transactors. There is a strong serial correlation in revolving behavior, in that most account holders either always/usually or never/rarely revolve (Grodzicki and Koulayev, 2021). Because convenience users are not charged interest on their accounts, they are not directly affected by changes in APR. Approximately 40 to 50 percent of credit card holders revolve their balances every year, so the interest rate elasticity affects a substantial portion of the population.

Credit card accounts with revolving balances are typically associated with cardholders

Figure 6: Credit Card Debt and Credit Score

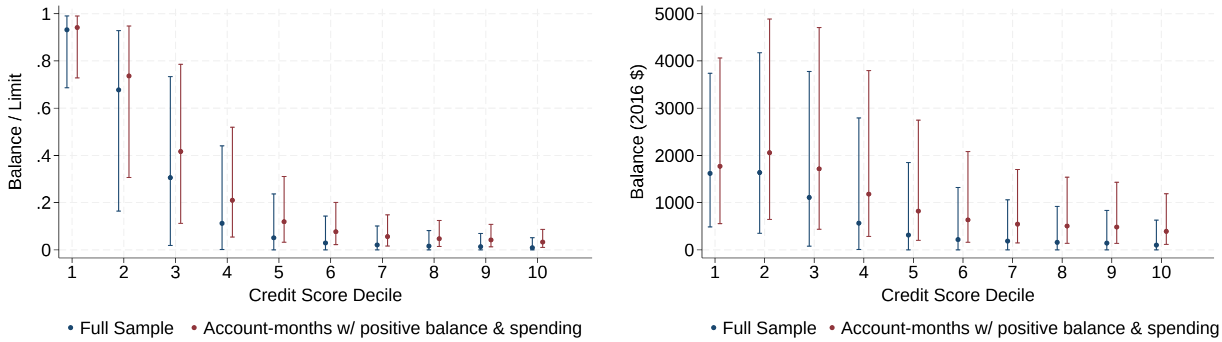

(a) Utilization

(b) Balance

Note: Panel (a) shows the 25th, 50th, and 75th percentiles of account utilization for observations within each credit-score decile, calculated using Y-14M credit card account data. Utilization is calculated by dividing account balance by credit limit. Panel (b) shows the 25th, 50th, and 75th percentiles of account balance, in 2016 dollars, for observations within each credit-score decile. *Sources*: Y-14M, authors' calculations.

with lower credit quality. Appendix Table A.2 shows that revolving is highly correlated with credit score. Figure 6 reports two measures of credit card balances against the credit-score decile. Panel (a) shows that credit-limit utilization, defined as the dollar balance over the limit, is higher for accounts associated with a lower credit score. The lowest-credit-score decile, in fact, has a median utilization that is close to one, meaning those cards tend to max out their limits. Credit utilization falls monotonically across credit-score deciles, that is, when moving toward accounts with higher credit scores. Panel (b) reports the dollar balance (2016 dollars) against the credit-score deciles, confirming that the balance of lower-credit-score accounts is higher not only in relative terms, but also in levels, even though lower-credit-score account holders tend to have lower income (Appendix Figure A.2)

Table 4 shows elasticity estimates for subgroups of accounts to zoom in on the role of revolvers. Consistent with the imposed channels, in column (1), revolvers have a high interest-rate elasticity of credit card spending. Our estimate indicates an elasticity of −15—almost twice as large as we found in the pooled sample. By contrast, the results for transactors shown in column (2) show an insignificant elasticity close to zero. In columns (3) and (4), we show the results for balances. Column (3) shows a significant elasticity estimate of −3.8 for revolvers, while column (4) shows ,for transactors, a small and insignificant positive elasticity estimate.

Heterogeneity by Credit Quality We also explore the role of financial constraints. Table 5 reports the interest-rate elasticity of credit card spending and balances by credit score. We

Table 4: RKD Estimates of Interest-Rate elasticity of Spending and Balance by Revolver Status

	Spending		Balance	
Account Type	Revolvers (1)	Transactors (2)	Revolvers (3)	Transactors (4)
Elasticity	-15.29*** (4.07)	-1.66 (3.88)	-3.84** (1.82)	4.59 (5.35)
Account FE	Yes	Yes	Yes	Yes
State*Bank*Credit Score*Month FE	Yes	Yes	Yes	Yes
N	1095134	697213	1095134	697213

Note: This table reports interest-rate elasticity of spending (columns 1 and 2) and balance (columns 3 and 4) for revolvers versus transactors. The RKD sample is defined subject to specifications outlined in the "Data Description" section using observations for which the running variable is within a 2 percentage point bandwidth (inclusive) around the the maximum APR. Standard errors are clustered by account and month. Estimates are weighted by an account-month's share of the total spending or balance across all accounts in the subsequent month, winsorized at the 97.5th percentile.

find that lower-credit-score accounts have a large elasticity of credit card spending, with a point estimate of –18, while spending on high-credit-score accounts does not significantly respond to interest rate changes (small estimate of –3.07). This finding aligns with the previously shown result that revolver status is highly correlated with credit score (see Appendix Table A.2).

For balances, we find a relatively large elasticity estimate of –7.12 for high-credit-score accounts and a small and insignificant elasticity estimate of –0.98 for low-credit-score accounts. Taken together, the results indicate that low-credit-score consumers respond to interest rate changes through adjustments in spending (consumption) rather than balance (leverage). The opposite is true for high-credit-score consumers.[11]

This finding is consistent with prevailing financial frictions. Low-credit-score consumers have limited access to credit and are likely to be liquidity-constrained; therefore, they may lack alternative sources that would allow them to pay down their revolving credit card balances while maintaining spending. By contrast, high-credit-score consumers may have access to savings or less costly credit sources, allowing them to pay down their balances when APR increases while maintaining spending. Indeed, this finding is consistent with financially

[11] Appendix Tables A.3 and A.4 report disaggregated results when splitting the sample by credit-score type and revolver status. Results depending on credit-score type are concentrated in revolving accounts and not present in transactor accounts, consistent with our previous finding.

Table 5: RKD Estimates of Interest-Rate elasticity of Spending and Balance by Credit Score

	Spending		Balance	
Credit Score	Low Credit (1)	High Credit (2)	Low Credit (3)	High Credit (4)
Elasticity	-17.89***	-3.07	-0.98	-7.12**
	(5.90)	(3.11)	(1.82)	(3.32)
Account FE	Yes	Yes	Yes	Yes
State*Bank*Credit Score*Month FE	Yes	Yes	Yes	Yes
N	943944	935278	943944	935278

Note: This table reports interest-rate elasticity of spending (columns 1 and 2) and balance (columns 3 and 4) for high- versus low-credit-score accounts. Low- and high-credit-score groups are determined by whether an account falls above or below the median credit score of accounts in the same month reporting the same credit-score type (within the RKD sample). The RKD sample is defined subject to specifications outlined in the "Data Description" section using observations for which the running variable is within a 2 percentage point bandwidth (inclusive) around the the maximum APR. Standard errors are clustered by account and month. Estimates are weighted by an account-month's share of the total spending or balance across all accounts in the subsequent month, winsorized at the 97.5th percentile.

unconstrained consumers tapping into credit in response to shocks (that is, interest rate changes) to smooth their consumption, while financially constrained consumers cannot do this.

The significant lowering of carried balances by high-credit-score individuals is also consistent with their higher levels of financial literacy. There is evidence that higher credit scores are correlated with higher levels of financial literacy and better financial outcomes (Lusardi and Tufano, 2015; Urban et al., 2020; Kaiser et al., 2022). High-credit-score individuals are more likely to be financially savvy and have the means to reduce their debt, while low-credit-score individuals reduce their spending instead.

Macro Evidence We next provide aggregate evidence of the impact of the monetary policy rate on credit card spending. This exercise serves several purposes. First, our RKD analysis pins down the APR elasticity of spending, and we try to directly gauge the pass-through of monetary policy rate changes (that is, FFR changes) into aggregate credit card spending. We focus on the size and timing of the aggregate effect. Second, our aggregate analysis is valuable given that the RKD identifies a local average treatment effect from observations in a small subsample of the data in which the APR is near or at its contractual maximum. Therefore, the LATE nature of our elasticity estimates may raise potential concerns about the validity of our RKD estimates in a broader context, specifically because we have shown that our RKD sample contains, disproportionately, observations from the high-interest period as well as from accounts with a high margin. Looking at aggregate evidence helps us identify these concerns.

For our aggregate analysis, we also rely on the Y-14M spending data, but we focus on aggregate (that is, spending-weighted) spending growth in a given month. We then estimate how changes in the FFR affect changes in spending growth at different response horizons. To estimate these effects, we apply standard local projections using the following linear regression model:

$$\Delta c_{t+h} = \beta^h \Delta r_t + \sum_{i=1}^{3} \phi_i^h \Delta c_{t-i} + u_t^h, \tag{6}$$

where Δc_{t+h} is the log difference in spending between period $t+h$ and period $t+h-1$, and Δr_t is the FFR change from month t over $t-1$. The object of interest is the parameter β^h, which captures the effect of a monetary policy rate change on spending growth h month ahead.

Figure 7 shows the estimated response coefficients of equation (6) for key response horizons. The blue line reports the response of total spending, while the red line reports the response of spending by revolving accounts. The results indicate a significant reduction in spending growth two months after the policy rate change. (Note that we estimate the model with differenced variables, hence a one-time negative downward shift in growth rate implies a permanent shift in the level.) The negative spending growth rates then gradually increase to zero over the next two months, with the effect on revolver spending growth being also significantly negative in $h = 3$.

Table 6 reports the peak response coefficient for revolvers (column 1) and total spending (column 3). The estimates indicate an FFR elasticity of spending growth of –0.08 for revolvers, while the total effect is slightly below –0.06. These effects are estimated using least squares. Because interest rate changes are endogenous, ordinary least squares (OLS) estimates may be biased, and we therefore also present results in which we estimate equation 6 using instrumental variables (IVs). Specifically, we instrument the (endogenous) prime rate change with monetary policy shocks from Jarociński and Karadi (2020), who use information inherent in high-frequency movement of interest rates and equity prices around monetary policy announcements. (Remember that the prime rate closely tracks the FFR.)

Columns (2) and (4) show that IV estimates are somewhat larger then their OLS counterparts, with estimates of –0.14 (revolvers) and –0.095 (total). First-stage effective F-statistics are reported in the bottom panel of the table. We also report Anderson–Rubin 90 percent

Figure 7: Response of Spending Growth to a Federal Funds Rate Increase

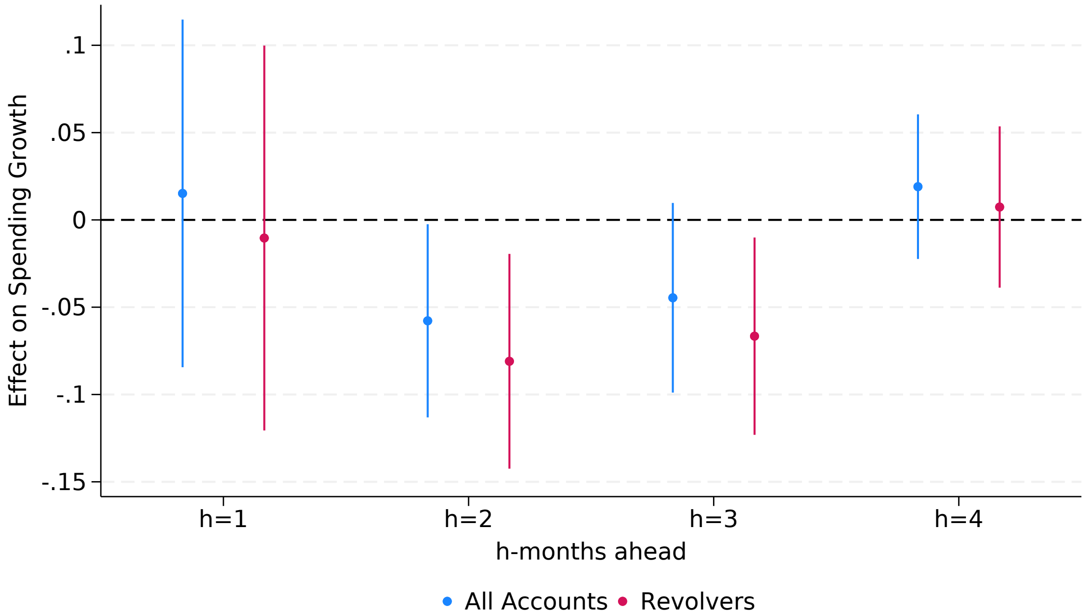

Note: This figure shows the response of spending growth (log difference) to an FFR increase of 1 percentage point. The dots refer to the estimates for β_h in equation (6), and the bars represent 90 percent robust confidence intervals. *Sources*: Haver, Y-14M, authors' calculations.

Table 6: Response of Spending Growth at h = 2

	Δc^{rev} (1)	Δc^{rev} (2)	Δc (3)	Δc (4)
Δr_{FF}	-0.0810**	-0.141***	-0.0578*	-0.0949**
	(-2.17)	(-2.90)	(-1.72)	(-1.98)
Eff F Stat		5.175		4.629
Lower CI		-0.241		-0.188
Upper CI		-0.030		0.044
Cumulative at h = 4	-0.145	-0.309	-0.069	-0.167
p-val	0.064	0.044	0.366	0.288

Note: This table provides more details on the response of spending growth to an FFR increase. The top panel reports the peak effect at $h = 2$ months after the funds rate change. Columns (1) and (2) focus on spending by revolving accounts; columns (3) and (4) focus on total spending. OLS and IV estimates are both presented. For IVs, we also report Anderson–Rubin confidence intervals to account for potentially weak instruments.

confidence intervals to account for potentially weak instruments. For revolver effects, the robust confidence intervals do not include the zero, confirming the central finding that spending growth drops significantly two months after an FFR increase.

Why the delayed response of spending growth to policy rate changes? That is, why does spending not adjust immediately the month after the policy rate change? There are several non-mutually exclusive reasons. As discussed, it typically takes two months for policy rate changes to transmit fully into APR, due to the way the credit agreements are written. Moreover, in the Y-14M data, reporting of spending in a given month refers to spending of all accounts with the last billing cycle day in that month.[12] Further, consumers are typically not notified about APR changes related to movements in the base rate (which is closely linked to the FFR). As a result, lack of information and/or inattention may also contribute to the delayed effect.

As shown earlier, the drop in spending growth happens two months later, and growth remains below counterfactual through Quarter 4, although at diminishing rates. For context, the bottom of the table also reports the cumulative spending effects through Quarter 4. Point estimates indicate a drop in cumulative consumption over the four months after a 1 percentage point interest rate increase of about 16.7 percent (total spending) and of about 30.9 percent for revolvers. Note, however, that the effects are statistically significant only for revolvers. Generally, we find the standard errors of the cumulative estimate to be larger, potentially given large standard errors, especially for $h = 1$ effects.

6 Conclusion

We use a large data set of the majority of credit card accounts issued by US banks from 2016 to 2024 to assess the effect of monetary policy on consumer spending through the effect of interest rate changes on credit card borrowing. Taking advantage of a kink in the relationship between prime rate and APR generated by a contractual rate ceiling, we estimate the elasticity of credit card spending with respect to interest rates. We find that credit card spending declines significantly when interest rates increase and that this result is driven by credit card revolvers and by low-credit-score cardholders. By contrast, credit card spending by convenience users and by high-credit-score cardholders does not change significantly when interest rates rise.

Our tightly identified RKD estimates are based on a subset of accounts that overrepresents revolvers and low-credit-score cardholders, which drive our results. In particular, more than 60 percent of accounts in the RKD sample have a revolving balance compared with about 30 percent of the full sample. Therefore, we complement our RKD analysis using aggregate analysis. Specifically, we verify a strong response of aggregate credit card spending to federal

[12]The Y14 data for 2015:m5 contains spending on all credit cards with a billing cycle ending in m5.

funds rate changes using local projections with monetary shocks as instrumental variables.

Our findings are relevant for understanding the transmission of monetary policy to aggregate consumption. In particular, our findings highlight an economically sizable credit card spending channel. This channel likely has grown in importance over time, a development that can be attributed to the increasing prevalence of credit card use and the rising proportion of purchases financed through credit card debt.

References

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(1): 129–190.

Angrist, Joshua D., and Iván Fernández-Val. 2013. "ExtrapoLATE-ing: External Validity and Overidentification in the LATE Framework." *Advances in Economics and Econometrics: Tenth World Congress*, , ed. Daron Acemoglu, Manuel Arellano and EddieEditors Dekel *Econometric Society Monographs*, 401–434. Cambridge University Press.

Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A. Parker. 2021. "Asymmetric Consumption Smoothing." *American Economic Review*, 111(1): 192–230.

Berg, Kimberly A., Chadwick C. Curtis, Steven Lugauer, and Nelson C. Mark. 2021. "Demographics and Monetary Policy Shocks." *Journal of Money, Credit and Banking*, 53(6): 1229–1266.

Bernanke, Ben S., and Alan S. Blinder. 1992. "The Federal Funds Rate and the Channels of Monetary Transmission." *The American Economic Review*, 82(4): 901–921.

Blanchard, Olivier. 1993. "Consumption and the Recession of 1990-1991." *The American Economic Review*, 83(2): 270–274.

Card, David, David S. Lee, Zhuan Pei, and Andrea Weber. 2015. "Inference on Causal Effects in a Generalized Regression Kink Design." *Econometrica*, 83(6): 2453–2483.

Card, David, David S. Lee, Zhuan Pei, and Andrea Weber. 2016. "Regression Kink Design: Theory and Practice." National Bureau of Economic Research Working Paper 22781.

Christiano, Lawrence J., Martin Eichenbaum, and Charles Evans. 1996. "The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds." *The Review of Economics and Statistics*, 78(1): 16–34.

Di Maggio, Marco, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao. 2017. "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging." *American Economic Review*, 107(11): 3550–88.

Drozd, Lukasz A., and Michal Kowalik. 2023. "The Conundrum of Zero APR: An Analytical Framework." Federal Reserve Bank of Philadelphia Working Papers 23-06.

Grigoli, Francesco, and Damiano Sandri. 2023. “Monetary policy and credit card spending.” Bank for International Settlements BIS Working Papers 1064.

Grodzicki, Daniel. 2023. “Credit cards and the floating rate channel of monetary policy.” *Applied Economics Letters*, 30(3): 249–254.

Grodzicki, Daniel, Alexei Alexandrov, Özlem Bedre-Defolie, and Sergei Koulayev. 2023. “Consumer Demand for Credit Card Services.” *Journal of Financial Services Research*, 63(3): 273–311.

Grodzicki, Daniel, and Sergei Koulayev. 2021. “Sustained credit card borrowing.” *Journal of Consumer Affairs*, 55(2): 622–653.

Gross, David B., and Nicholas S. Souleles. 2002. “Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data*.” *The Quarterly Journal of Economics*, 117(1): 149–185.

Indarte, Sasha. 2023. “Moral Hazard versus Liquidity in Household Bankruptcy.” *The Journal of Finance*, 78(5): 2421–2464.

Jarociński, Marek, and Peter Karadi. 2020. “Deconstructing Monetary Policy Surprises - The Role of Information Shocks.” *American Economic Journal: Macroeconomics*, 12(2): 1–43.

Kaiser, Tim, Annamaria Lusardi, Lukas Menkhoff, and Carly J. Urban. 2022. “Financial education affects financial knowledge and downstream behaviors.” *Journal of Financial Economics*, 145(2): 255–272.

Lettau, Martin, Sydney C. Ludvigson, and Charles Steindel. 2002. “Monetary policy transmission through the consumption-wealth channel.” *Economic Policy Review*, 8(May): 117–133.

Lusardi, Annamaria, and Peter Tufano. 2015. “Debt literacy, financial experiences, and overindebtedness.” *Journal of Pension Economics and Finance*, 14(4): 332–368.

Matthes, Christian, and Felipe Schwartzman. 2025. “The Consumption Origins of Business Cycles: Lessons from Sectoral Dynamics.” *American Economic Journal: Macroeconomics*, forthcoming.

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. “Household Balance Sheets, Consumption, and the Economic Slump.” *The Quarterly Journal of Economics*, 128(4): 1687–1726.

Nelson, Scott T. 2025. "Private Information and Price Regulation in the US Credit Card Market." *Econometrica*, 93(4): 1371–1410.

Pei, Zhuan, David S. Lee, David Card, and Andrea Weber. 2022. "Local Polynomial Order in Regression Discontinuity Designs." *Journal of Business & Economic Statistics*, 40(3): 1259–1267.

Urban, Carly, Maximilian Schmeiser, J. Michael Collins, and Alexandra Brown. 2020. "The effects of high school personal financial education policies on financial behavior." *Economics of Education Review*, 78: 101786.

Zeldes, Stephen P. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy*, 97(2): 305–346.

A Additional Figures and Tables

Figure A.1: Impact of Individual Bank on Aggregate Share of Revolver Spending

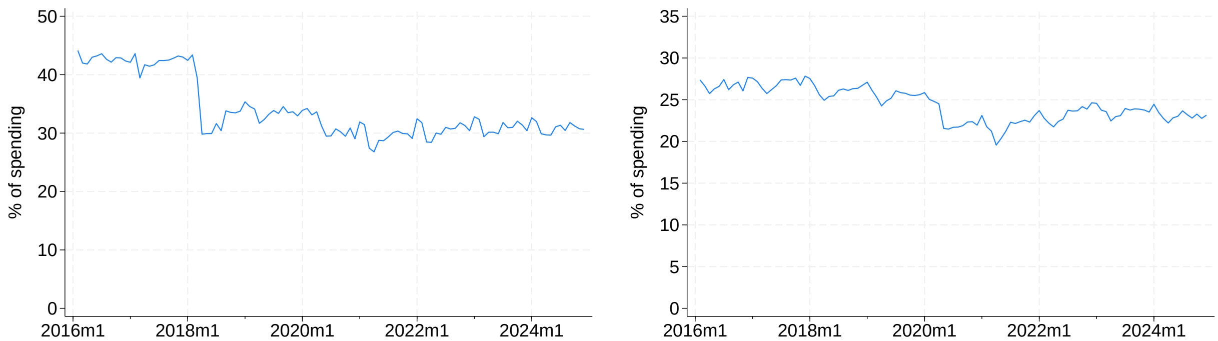

(a) Share of Spending by Revolvers (b) Share of Spending by Revolvers minus Bank

Note: Panel (a) shows the percentage of credit card spending made by accounts with a positive revolving balance, calculated using all Y-14M credit card account data. Panel (b) shows the percentage of credit card spending made by accounts with a positive revolving balance, excluding a singular Y-14M participating bank from the entire sample due to a bank-specific change in reporting methodology generating a discontinuity in the early months of 2018. *Sources*: Y-14M, authors' calculations.

Figure A.2: Income Distribution by Credit-Score Decile

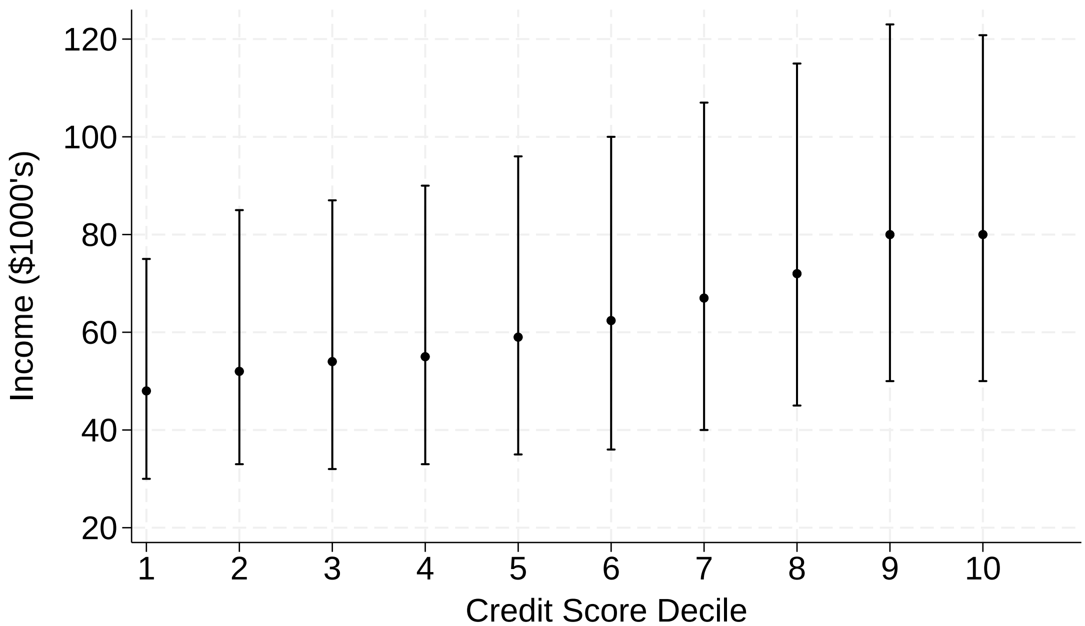

Note: The figure shows the 25th, 50th, and 75th percentiles of interest-rate margin for observations within each credit-score decile, calculated using Y-14M credit card account data. The figure reports only one observation for each account using the credit score from the earliest observed month of the account's existence. The credit-score decile is calculated according to specifications outlined in the "Data Description" section. Income is reported at origination. *Sources*: Y-14M, authors' calculations.

Figure A.3: Credit Card Spending and Credit Score

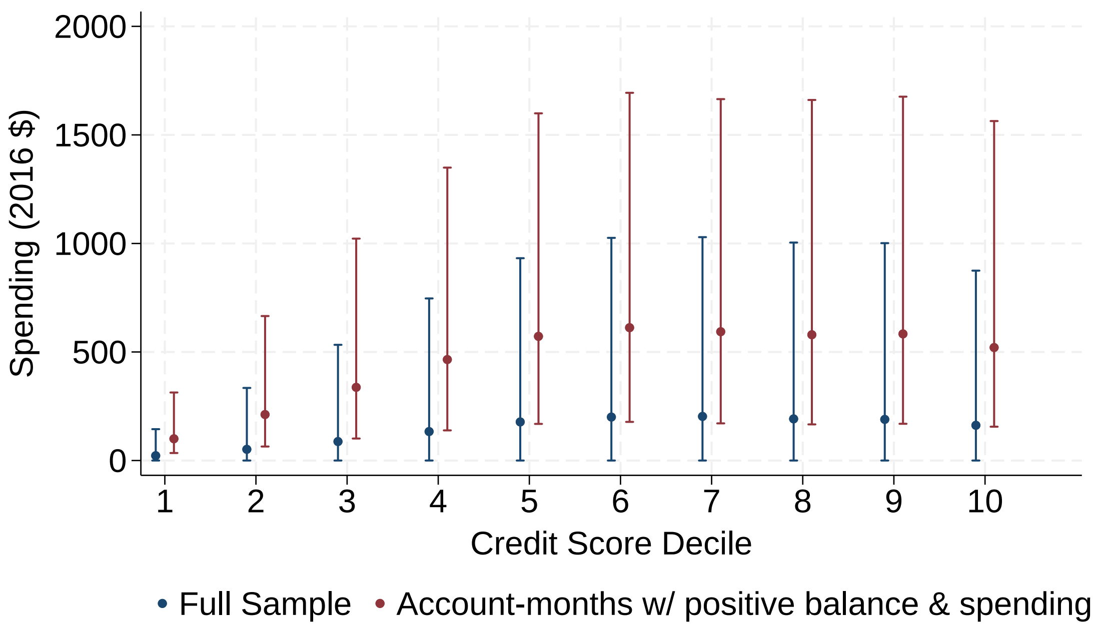

Note: The figure shows the 25th, 50th, and 75th percentiles of account spending, in 2016 dollars, for observations within each credit-score decile, calculated using Y-14M credit card account data. The credit-score decile is calculated according to specifications outlined in the "Data Description" section. *Sources*: Y-14M, authors' calculations.

Figure A.4: Distribution of Cycle-End Days

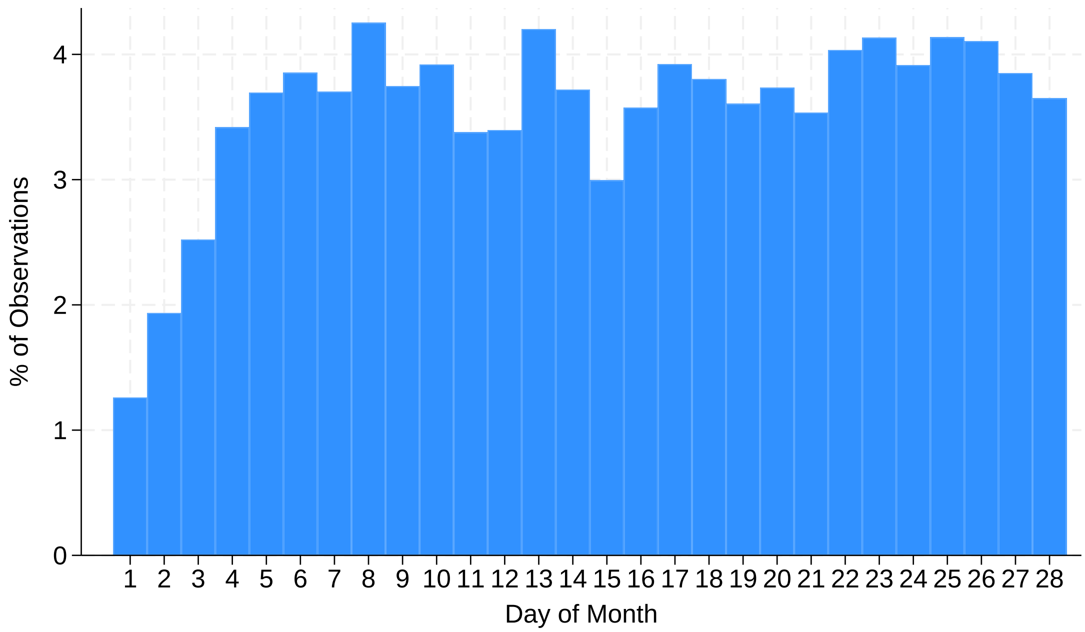

Note: The figure shows the distribution of cycle-end days, including only observations ending on day 1 through 28 of a given month, calculated using Y-14M credit card account data. *Sources*: Y-14M, authors' calculations.

Figure A.5: RKD Estimates of Interest-rate Elasticity by Horizon

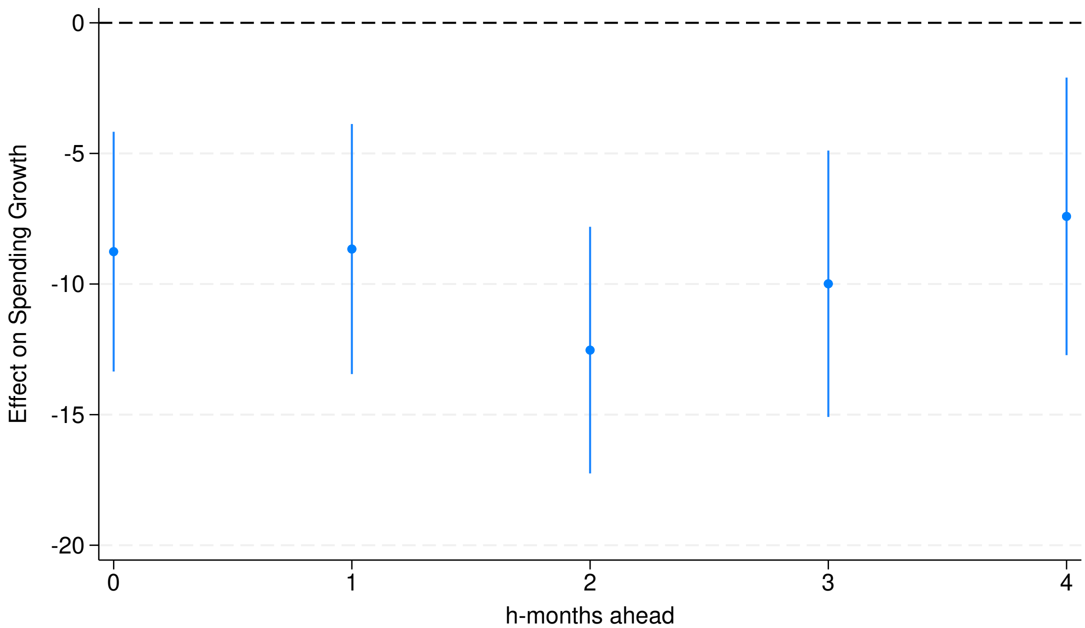

Note: The figure plots the estimated APR elasticity of credit card spending for varying time horizons. Horizon zero represents contemporaneous spending. Horizons one through four are monthly leads of account-level credit card spending. *Sources*: Y-14M, authors' calculations.

Figure A.6: Robustness to RKD Specification

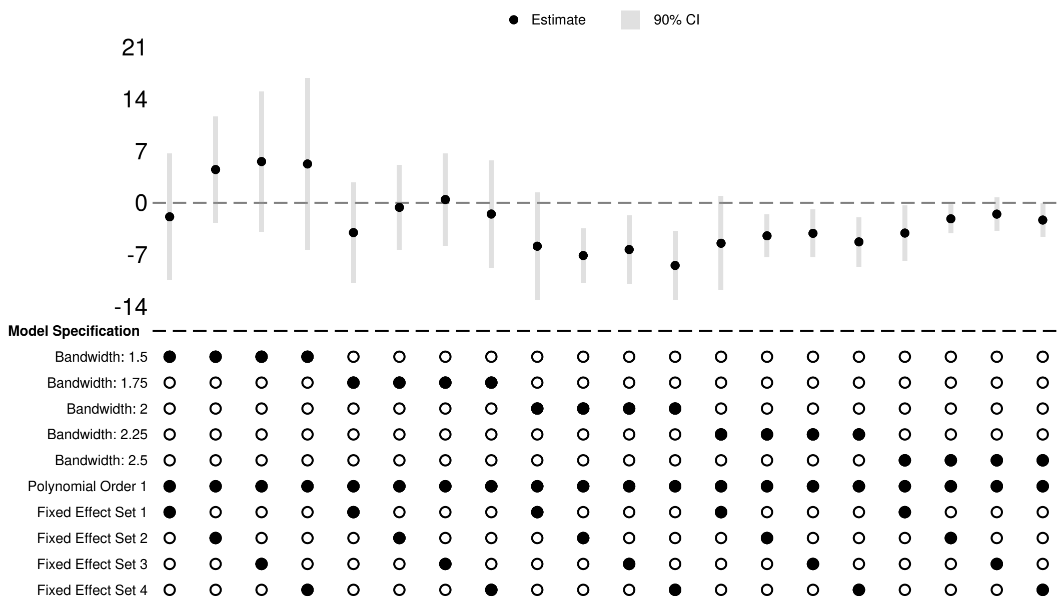

Note: This figure shows the robustness of our baseline elasticity estimates to changes in key parameter choices underlying the RKD. The parameters considered include sample bandwidth and fixed effects. Each specification uses a linear polynomial order. For each parameter combination indicated with black dots in the bottom panel of the figure, we estimate the interest-rate elasticity of credit card account spending. Point estimates shown in the top region of the figure are depicted as black dots, and 90 percent confidence intervals are shown as gray bars. Fixed Effect Set 1 corresponds to the fixed effects in column 1 of Table 2, Fixed Effect Set 2 to column 2, and so forth. *Sources*: Y-14M, authors' calculations.

Table A.1: Asymptotic Mean Squared Error by Polynomial Order

	Polynomial Order		
	Order 1	Order 2	Order 3
Estimated AMSE			
Conventional Sharp Estimator	0.000079	0.091931	1.076719
Bias-Corrected Sharp Estimator	0.467135	9.55413	0.579893

Note: The table shows the asymptotic mean squared error (AMSE) of the RKD estimator for varying polynomial orders, as outlined in Pei et al. (2022). *Sources*: Y-14M, authors' calculations.

Table A.2: Category Shares

	Credit Score		
	Low Credit	High Credit	Total
Account Type			
Transactor	9.58	30.09	39.67
Revolver	40.40	19.93	60.33
Total	49.98	50.02	100.00

Note: The table shows the percentage of "Transactor" and "Revolver" account types by credit score category. "Low Credit" is defined as having a credit score below the median. "High Credit" is defined as having a credit score above the median. *Sources*: Y-14M, authors' calculations.

Table A.3: Interest-rate Elasticity of Spending by Revolver and Credit Score Status

	Revolvers		Transactors	
	Low Credit (1)	High Credit (2)	Low Credit (3)	High Credit (4)
Elasticity	-20.06***	-5.76	-2.71	-1.34
	(6.17)	(5.39)	(11.58)	(4.47)
Account FE	Yes	Yes	Yes	Yes
State*Bank*Credit Score*Month FE	Yes	Yes	Yes	Yes
N	735810	342879	148969	531714

Note: The RKD sample is defined subject to specifications outlined in the "Data Description" section using observations for which the running variable is within a 2 percentage point bandwidth (inclusive) around the the maximum APR. Standard errors are clustered by account and month. Estimates are weighted by an account-month's share of total spending that occurred in the subsequent month, winsorized at the 97.5th percentile. *Sources*: Y-14M, authors' calculations.

Table A.4: Interest-rate Elasticity of Balance by Revolver and Credit Score Status

	Revolvers		Transactors	
	Low Credit (1)	High Credit (2)	Low Credit (3)	High Credit (4)
Elasticity	-1.48 (1.61)	-10.02*** (3.72)	20.69 (13.48)	1.64 (5.75)
Account FE	Yes	Yes	Yes	Yes
State*Bank*Credit Score*Month FE	Yes	Yes	Yes	Yes
N	735810	342879	148969	531714

Note: The RKD sample is defined subject to specifications outlined in the "Data Description" section using observations for which the running variable is within a 2 percentage point bandwidth (inclusive) around the the maximum APR. Standard errors are clustered by account and month. Estimates are weighted by an account-month's share of total spending that occurred in the subsequent month, winsorized at the 97.5th percentile. *Sources*: Y-14M, authors' calculations.