Existing Commitments, New Credit, and Monetary Policy

Xinhao Dong*

November 2026

Abstract

Credit-card lines are commitments: banks authorize them before households use them. We ask where monetary tightening adjusts card credit—on lines already granted or on the terms of new accounts. In quarterly data on the U.S. large-bank card market, a one-standard-deviation contractionary high-frequency surprise lowers the median initial limit on newly opened accounts by 5.2 log points relative to the median limit on open accounts, while total commitments and the number of open accounts do not decline. Because new authorizations are about 2.3 percent of committed credit each quarter, this margin moves the stock of card credit slowly and by little. In a household financing framework built from payment, balance-sheet and expenditure surveys, a dollar of authorization removed at origination reduces spending 2.35 times as much as a dollar withdrawn from existing lines, because new credit reaches households with little or no unused capacity.

JEL classification: E52, E21, D14, G21, G51.

Keywords: monetary policy; credit commitments; credit cards; credit limits; household liquidity.

*Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, BC V5A 1S6, Canada. Corresponding author. Email: chinasdqddxh@gmail.com.

1 Introduction

A credit-card line is a commitment to lend, made before the household decides when to borrow. This timing gives a lender two ways to change the quantity of card credit after a monetary contraction. It can revise lines it has already granted, or it can change the terms it offers on accounts opened after the shock. The two adjustments reach different households. A household that holds an unused line can keep drawing on an authorization obtained in the past; a household seeking a new card faces the terms available today. Measures of credit outstanding, or of limits across all open accounts, mix the two adjustments.

This paper asks which of the two margins monetary policy moves, and what the answer implies for the scale and incidence of the credit-card quantity channel. We use the public quarterly series derived from the Federal Reserve's FR Y-14M collection for the U.S. large-bank card market. The series report, for the same market and quarter, the distribution of initial limits on accounts opened during the quarter and the distribution of limits on all accounts open in that quarter, together with gross originations, total commitments and the incidence of line increases and decreases on open accounts. That combination lets us estimate the response of the origination margin and of the committed stock to the same monetary surprise, in the same regression design.

The two margins respond differently. Following a one-standard-deviation contractionary surprise in the orthogonalized high-frequency series of Bauer and Swanson (2023), the median initial limit on new accounts falls by 4.6 log points on impact. The directly estimated difference between the new-account median and the open-account median is -5.2 log points, with a 95 percent confidence interval of $[-7.5, -2.9]$. The contrast survives excluding the pandemic year and controlling for the change in the median credit score of new accounts. It describes the terms of credit rather than its volume: the mean initial limit, the number of new accounts and newly authorized dollars all have imprecisely estimated responses.

The committed stock does not contract. Under the same surprise, total authorized commitments respond by 0.01 log points [−0.16, 0.18], the number of open accounts by 0.01 [−0.19, 0.21], and the share of open accounts receiving a line decrease does not rise. The open-account median does not fall, and the upper quantiles of the open-account distribution move little. Associations with realized policy-rate changes point the same way, as does a bank-holding-company panel that spans monetary cycles back to the early 1990s. The price of revolving credit, by contrast, adjusts: in realized-rate data the assessed-interest APR rises by 1.08 percentage points for each 100-basis-point increase in the federal funds rate.

The origination result rests on one surprise series, and we treat this as the main threat to it. The new-minus-open contrast is negative under all four surprise series we examine (five measures, since we use the Bauer–Swanson series in both its orthogonalized and its published form), but with quarter-of-year effects only the two Bauer–Swanson measures yield intervals that exclude zero. Bauer and Swanson orthogonalize their surprises with respect to macroeconomic and financial information available before each announcement, and that step matters here. Using their unorthogonalized surprise shrinks the per-basis-point contrast to 0.61 of its orthogonalized value. Splitting the unorthogonalized surprise into its orthogonal and predictable parts shows why: the orthogonal part lowers new limits relative to open limits by 0.63 log points per basis point, while the predictable part has an opposite-signed but imprecisely estimated coefficient (0.48, [−0.11, 1.07]); the two coefficients differ ($p = 0.003$). This pattern is consistent with the weaker contrasts we find under surprise series that retain predictable news.

The two facts together indicate the scale of the quantity channel that runs through card credit. New authorizations average 2.3 percent of committed credit per quarter. Even if the estimated decline in median initial limits applied to every new dollar authorized over the following five quarters, and every such account survived, commitments would be lower by about 0.38 percent (0.74 percent excluding the pandemic year). The pattern parallels

the evidence of Morgan (1998), who shows that bank loans made under commitment do not contract after monetary tightening while loans without commitments do. In revolving consumer credit, the commitment is the line itself, and almost all of the stock is committed.

A small quantity channel can still be unevenly distributed. The second part of the paper asks whose spending the origination margin reaches. We combine the 2024 Diary of Consumer Payment Choice, the Survey of Consumer Finances and the Consumer Expenditure Survey into a weighted synthetic population in which each household has income, liquid assets, existing card limits and balances, and a quarterly expenditure requirement. A one-quarter financing rule maps available resources into the spending a household can finance. The framework is deliberately narrower than a structural model of credit demand: it takes balances and expenditure needs as given and asks how a specified change in authorized credit changes financeable spending.

Within that framework we remove the same amount of authorization in two ways: by imposing the estimated tightening on the initial limits of a quarter's new accounts, and by withdrawing an equal number of dollars uniformly from existing lines. In this framework, a dollar of authorization removed at origination costs 2.3 times as much spending as a dollar removed from existing lines, or 1.8 cents more. The comparison favors the origination margin in every one of the 50 cross-survey pairings, in all 57 alternative populations we construct by varying cross-survey dependence, balance-sheet donors and the allocation of new accounts (with ratios between 1.41 and 3.87), and under all five definitions of the expenditure requirement we consider. The reason is where each margin's dollars sit. Existing authorization is concentrated in large lines held by non-revolvers with substantial unused capacity, so a proportional withdrawal mostly removes slack. New authorization is spread across accounts and reaches households with little or no unused capacity, including households that hold no card at all. Those initially cardless households receive 9.9 percent of new authorization but bear 29.4 percent of the loss; per dollar of reduced authorization, their loss is 2.5 times that of revolvers and 6.5 times that of non-revolvers. Much of

the comparison runs through these households: among cardholders alone, the origination margin's loss per dollar is 1.8 times that of existing lines in the baseline, but it falls below that of existing lines (0.91) in the populations where cardholders' expenditure requirements rise most steeply with their unused capacity. The aggregate one-quarter effect is small, 0.0048 percent of reference spending, consistent with the small flow.

These per-dollar contrasts come from a rule in which spending responds to credit only when a household's financing constraint binds. Quasi-experimental estimates of responses to credit-line expansions, such as Aydin (2022) and Gross and Souleles (2002), find responses also among borrowers far from their limits, and they are larger than the model's per-dollar values. Adding, for households whose constraint is slack, a uniform response equal to the first-quarter estimate of Aydin (2022) raises losses on both margins and narrows their ratio to 1.14 (no less than 1.04 across the alternative populations); the cardless-to-non-revolver ratio narrows to 1.6. Taken one at a time, three departures would reverse the ordering of the two margins: responses away from the constraint that are larger for cuts in existing lines than for smaller new lines by about 1.9 cents per dollar, receipt of new accounts by households with a shortfall below 43 percent of their group's rate, or outside financing of 62 percent of every measured shortfall (Section 4.5).

Two features of the evidence shape its interpretation. The public card series describe changing populations of accounts rather than fixed incumbent lines, so they do not rule out cuts concentrated on particular borrowers. The household calculations combine surveys that do not jointly observe financing resources, expenditure requirements and new-card receipt, and they hold the price of borrowing and households' willingness to borrow fixed; they measure exposure to a change in authorized credit, not the causal effect of monetary policy on consumption or welfare.

The paper makes three contributions. The first is to separate, within one market and against one identified shock, the response of credit terms at origination from the response of committed credit. Work on the credit channel has long recognized that commitments

change how lending responds to policy: Sofianos et al. (1990) and Morgan (1998) document that commitment lending is insulated from monetary tightening, and Duca and VanHoose (1990) and Woodford (1996) study the policy implications of advance commitments. Drawdowns on firms' existing credit lines shape how bank lending responds to shocks (Greenwald et al., 2025), and consumer lending contracts after monetary tightening (Ludvigson, 1998; den Haan et al., 2007). Recent work examines how monetary policy reaches card borrowing costs and spending (Grodzicki, 2023; Bräuning and Stavins, 2025; Grigoli and Sandri, 2026; Littlejohn, 2025), new household lending by banks and nonbanks (Bosshardt et al., 2024; Argudo, 2025; Cucic and Gorea, 2026), and the supply and pricing of card credit (Han et al., 2018; Drechsler et al., 2025); lending-standard surveys measure credit supply more broadly (Bassett et al., 2014). Our evidence shows that, for card credit, monetary tightening operates through the terms of new commitments rather than through the stock. That contrasts with bank funding shocks, under which issuers do cut existing lines (Chava et al., 2023).

The second contribution concerns identification. The origination response is sensitive to how the surprise is measured, and we show that the sensitivity has an interpretable source: the component of raw surprises predictable from pre-announcement information moves new-account terms in the opposite direction from the orthogonal component, as the argument of Bauer and Swanson (2023) implies (see also Nakamura and Steinsson, 2018; Miranda-Agrippino and Ricco, 2021). We report all series side by side, including the event-study surprises of Acosta et al. (2025) and the information-purged series of Jarociński and Karadi (2020).

The third contribution is to show how the adjustment margin shapes the incidence of credit quantities across households. Prior work estimates the consumption value of credit limits and its heterogeneity (Gross and Souleles, 2002; Agarwal et al., 2018; Gross et al., 2020; Aydin, 2022), documents how variable and actively managed limits are (Fulford, 2015; Santucci, 2015; Bord et al., 2026), and models the joint use of cards for payments and

revolving credit (Fulford and Schuh, 2024). Quantitative models of the card market show how lending technology, consumers' search among offers and the revolving structure of lines shape who obtains credit (Drozd and Serrano-Padial, 2017; Galenianos and Gavazza, 2022; Raveendranathan et al., 2026). Heterogeneous-agent models of monetary transmission put household liquidity at the center (Kaplan et al., 2018; Auclert, 2019), as does evidence from household microdata (Cloyne et al., 2020; Holm et al., 2021; Crawley and Kuchler, 2023; Ampudia et al., 2024), and a tightening of borrowing limits is a standard way to model a credit crunch (Guerrieri and Lorenzoni, 2017). Our results are consistent with a more specific structure for card credit under monetary tightening: limits on existing accounts adjust little, while the terms on new accounts tighten.

The rest of the paper is organized as follows. Section 2 presents the card-market evidence. Section 3 sets out the household financing framework and the two ways of removing authorization. Section 4 calibrates the origination margin and compares its incidence with that of an equal withdrawal from existing lines. Section 5 discusses what the results imply for monetary transmission, and Section 6 concludes.

2 Origination terms and the committed stock

This section estimates how the terms of new card credit and the committed stock of card credit respond to monetary surprises, and how far the first can move the second.

2.1 Data and design

Our card-market data are the public Large Bank Credit Card and Mortgage Data of the Federal Reserve Bank of Philadelphia, constructed from the FR Y-14M collection and available quarterly from 2012Q3 to 2025Q4. The series cover the credit-card portfolios of the largest U.S. bank holding companies and are aggregated across institutions, with no account identifiers. Three kinds of object matter for us. The first is the distribution of

authorized limits on all accounts open in a quarter, which describes the committed stock. The second is the distribution of initial limits on accounts opened during the quarter, which describes the terms of new commitments. The third is a set of quantities: gross new accounts, gross newly authorized dollars, total commitments, the number of open accounts, and the shares of open accounts that receive a line increase or decrease. The assessed-interest APR on revolving balances comes from the Federal Reserve's G.19 release, which covers commercial banks generally. Appendix Table A1 lists the series.

The open-account distribution is a cross-section rather than a cohort. It includes accounts opened in the current quarter as well as accounts opened earlier, and its composition changes as accounts open and close. A comparison of the two medians therefore asks whether the terms of new commitments move differently from the prevailing distribution of commitments, not whether a given account's line changes.

For an outcome Z_t we estimate local projections (Jordà, 2005) in quarterly differences,

$$100\,(\log Z_{t+h} - \log Z_{t-1}) = \alpha_h + \beta_h S_t + \Gamma_h' W_{t-1} + \tau_{q(t)} + \varepsilon_{t+h}, \tag{1}$$

where S_t is the monetary surprise, W_{t-1} contains lagged unemployment, the VIX, card delinquency and utilization and a linear trend, and $\tau_{q(t)}$ are quarter-of-year effects, which matter because new-account counts and mean initial limits are strongly seasonal. Dollar outcomes are deflated by the CPI. Standard errors are Newey–West, and intervals use Student-t critical values with the residual degrees of freedom. Following the high-frequency approach of Gürkaynak et al. (2005) and Gertler and Karadi (2015), our identifying measure is the orthogonalized surprise of Bauer and Swanson (2023), summed within the quarter and scaled so that one unit is one standard deviation, 7.89 basis points. The shock series ends in 2023Q4, so the impact regressions use 45 quarters from 2012Q4.

The object of interest is the difference between the responses of the two medians,

$$\Delta\beta_{50} = \beta_{new,50} - \beta_{open,50}, \tag{2}$$

which we estimate in a single regression with $\log A^{new}_{50,t} - \log A^{open}_{50,t}$ as the outcome, so that its standard error accounts for the covariance between the two medians.

Table 1: Monetary surprises, credit at origination, and the committed stock

	Estimate	95% CI	N
A. Credit at origination			
Median initial limit, new accounts	−4.63	[−6.98, −2.28]	45
Median limit, open accounts	0.58	[0.06, 1.11]	45
New minus open (estimated directly)	−5.22	[−7.54, −2.89]	45
Mean initial limit, new accounts	−1.40	[−3.08, 0.28]	45
Number of new accounts	0.08	[−2.26, 2.42]	45
Newly authorized dollars	−1.32	[−4.86, 2.22]	45
B. The committed stock			
75th percentile limit, open accounts	0.02	[−0.27, 0.30]	45
90th percentile limit, open accounts	−0.12	[−0.29, 0.06]	45
Total authorized commitments	0.01	[−0.16, 0.18]	45
Number of open accounts	0.01	[−0.19, 0.21]	45
Share of accounts with a line decrease (pp)	0.03	[−0.03, 0.10]	44
Share of accounts with a line increase (pp)	−0.05	[−0.20, 0.09]	44

Impact responses ($h = 0$) to a one-standard-deviation contractionary Bauer–Swanson surprise (7.89 basis points), in log points (100 times the change in log, CPI-deflated for dollar amounts) unless marked pp. Local projections in quarterly differences, 2012Q4–2023Q4, with lagged unemployment, VIX, delinquency, utilization, a linear trend, and quarter-of-year effects. Newey–West standard errors, Student-t intervals with residual degrees of freedom. The new-minus-open contrast is estimated as a single regression, so its interval accounts for the covariance of the two medians. Line-decrease and line-increase shares are percentage points of open accounts, with the same controls but no quarter-of-year effects (2013Q1–2023Q4); with quarter-of-year effects they are 0.03 [−0.03, 0.09] and −0.08 [−0.22, 0.07]. Open-account limits describe the cross-section of accounts open in each quarter, not a fixed cohort; the public series do not report quantiles below the median.

2.2 Terms at origination

Panel A of Table 1 reports the impact responses. A one-standard-deviation contractionary surprise lowers the median initial limit on new accounts by 4.6 log points. The median limit on open accounts does not fall, and the directly estimated contrast is −5.2 log points,

with an interval of [−7.5, −2.9]. The contrast does not depend on any single quarter, on the bandwidth of the variance estimator or on adding lagged terms: dropping one quarter at a time gives estimates between −6.0 and −4.4, and every interval in Appendix Table A3 excludes zero. Figure 1 traces the two medians over four quarters. The new-account median remains lower a quarter later, when the contrast is −7.3 log points, and the difference is no longer distinguishable from zero after two quarters. The evidence concerns the adjustment of origination terms on impact rather than a persistent divergence that this sample could estimate precisely. Appendix Table A2 reports every horizon, with and without the windows that touch 2020.

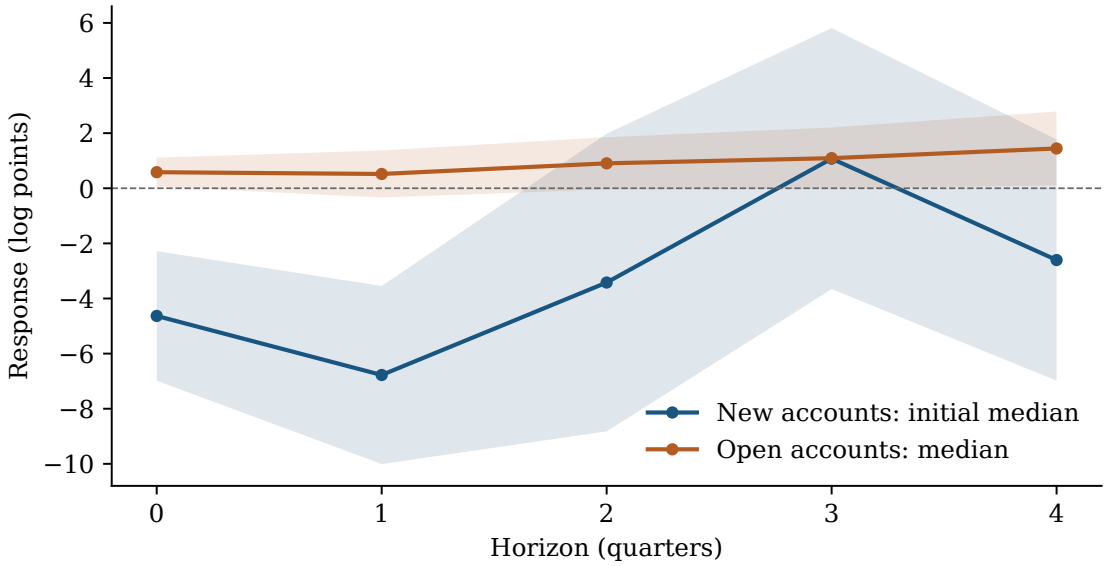


Figure 1: Median limits on new and open accounts after a contractionary surprise
Notes: Local projections of equation (1), in log points per one-standard-deviation Bauer–Swanson surprise (7.89 basis points), with quarter-of-year effects. Shading shows pointwise 95 percent intervals.

The point estimates are large relative to the size of the surprise. Rescaled linearly, the median response corresponds to −14.7 log points per 25 basis points. We avoid that rescaling. The quarterly series sum small surprises, the sample contains 45 quarters, and the largest surprise falls in the first quarter of 2020; a linear extrapolation well beyond the observed surprises would overstate what the data identify. The household calculations in

Section 4 therefore use the one-standard-deviation response itself.

The contrast concerns the terms of credit, not its measured volume. The mean initial limit, computed as newly authorized dollars divided by new accounts, falls by 1.4 log points with an interval of [−3.1, 0.3]; the number of new accounts and newly authorized dollars have responses of 0.08 and −1.3 log points, both with intervals that include zero. The evidence is therefore most precise about the middle of the distribution of initial limits. Whether total new authorization falls in proportion, and whether approval and issuance respond, the data leave open.

Changes in who receives new cards could move the median initial limit without any change in the terms offered to a given borrower. The public series allow one check on this. In the benchmark specification the median credit score of new accounts shows no significant change (−0.52 points, [−1.27, 0.24]), and controlling for its change leaves a contrast of −4.2 log points, [−6.4, −2.1] (Appendix Table A6). Within score bands, the median initial limit falls in all three bands in point estimate but significantly only for scores between 660 and 719 (−3.2 log points, [−5.2, −1.2]); without quarter-of-year effects the declines are larger and significant in the two upper bands. These checks argue against a shift toward lower-score borrowers as the main explanation; changes in unobserved applicant characteristics or card products remain possible.

2.3 The committed stock

Panel B of Table 1 turns to credit already committed. Total authorized commitments respond by 0.01 log points with an interval of [−0.16, 0.18], and the number of open accounts by 0.01. The 75th and 90th percentiles of the open-account limit distribution move by 0.02 and −0.12 log points. The share of open accounts receiving a line decrease does not rise, and the share receiving an increase changes little. The open-account median rises by 0.58 log points; since neither the number of accounts nor total commitments moves,

the rise reflects neither a smaller nor a larger stock, and the public series offer no way to attribute it further; we leave it uninterpreted.

Associations with realized policy-rate changes, over the slightly longer sample in which the card series are available, point in the same direction (Appendix Table A7). Per 100 basis points of realized federal funds increases, the open-account median, 75th and 90th percentiles change by 0.30, −0.19 and −0.27 percent, the line-decrease share by −0.02 percentage points, and the line-increase share by 0.38. Households use more of the capacity they hold: median utilization among active accounts rises by 1.07 percentage points, and the shares of accounts making only the minimum payment and of accounts thirty days past due rise as well. The price of revolving credit adjusts directly; in these realized-rate data the assessed-interest APR rises by 1.08 percentage points per 100 basis points, [0.83, 1.33]. Realized rate changes respond to economic conditions, so these are associations, not causal estimates. A longer panel of bank holding companies' unused card commitments in the FR Y-9C, which nets drawn balances and changes definition in 2010, shows no systematic response either before 2010 (−0.04 log points, standard error 0.45) or after.

Not every margin of line management is inert. Under the Jarociński–Karadi surprise, the 75th percentile of open-account limits falls by 0.78 log points, [−1.35, −0.21], and the 90th percentile falls under the SF Fed statement surprise with an interval that just reaches zero. When the windows touching 2020 are excluded, the share of accounts receiving a line increase declines for several quarters after a Bauer–Swanson surprise (−0.154 percentage points on impact, [−0.303, −0.005]). A placebo that regresses the line-decrease share on the following quarter's surprise also yields a positive coefficient (0.06, [0.01, 0.11]), a warning that the line-action series are noisy at this frequency. What the stock evidence supports is narrower than inertia: the center of the committed distribution, total commitments and the incidence of line cuts do not contract, while slower line growth and declines in the upper tail remain possible.

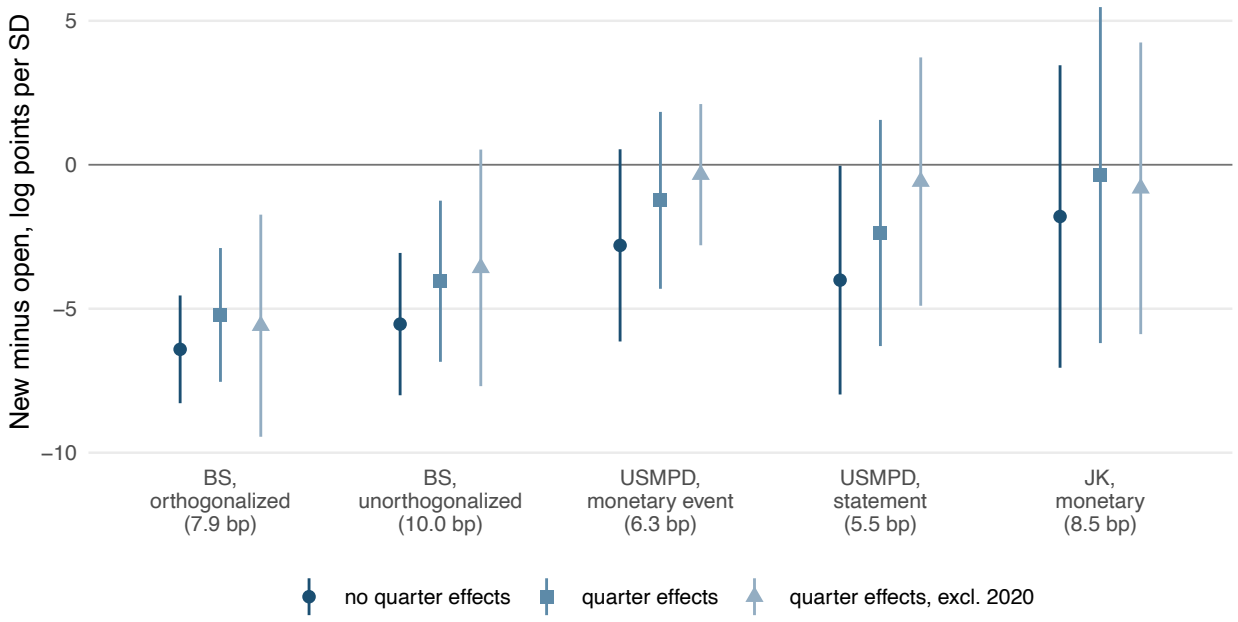


Figure 2: The new-minus-open contrast across surprise series

Notes: Impact estimates of equation (2) in log points per one standard deviation of each series (shown in basis points), with 95 percent intervals. BS: Bauer and Swanson (2023), orthogonalized and unorthogonalized. USMPD: SF Fed U.S. Monetary Policy Event-Study Database (Acosta et al., 2025), monetary-event and statement windows. JK: Jarociński and Karadi (2020), monetary component. The series differ in instruments, windows and units, which limits comparisons of magnitudes across series. Appendix Table A4 reports the estimates.

2.4 Evidence across surprise series

Figure 2 reports the new-minus-open contrast under five surprise measures and three specifications. Under all of them the point estimate is negative. With quarter-of-year effects, however, only the two Bauer–Swanson measures yield intervals that exclude zero; the unorthogonalized measure gives −4.0 log points, [−6.8, −1.2], per standard deviation. Without quarter-of-year effects the SF Fed USMPD statement surprise of Acosta et al. (2025) also does (−4.0, [−7.98, −0.04]), and when 2020 is excluded the contrasts under the USMPD and Jarociński–Karadi series lie between −0.8 and −0.3 log points, whereas the Bauer–Swanson contrast is −5.6 log points, [−9.4, −1.7]. The origination result therefore rests on the Bauer–Swanson family of measures, and the question is whether there is a reason, independent of this outcome, to prefer them.

Bauer and Swanson (2023) show that raw high-frequency surprises are correlated with macroeconomic and financial news released before the announcement, and they construct their series by removing the component that this news predicts. That component plausibly moves credit terms in the opposite direction from a policy surprise: when the economy has been strong, both market rates and lenders' willingness to grant generous initial limits tend to be high. Replacing the orthogonalized series with the published unorthogonalized surprise leaves the contrast negative and significant, −4.0 log points per standard deviation (10.0 basis points), [−6.8, −1.2], but shrinks it per basis point to 0.61 of its orthogonalized value, and to between 0.51 and 0.68 across specifications. Regressing the contrast jointly on the orthogonal component and on the predictable component (the difference between the two series) separates the two. The orthogonal component lowers new limits relative to open limits by 0.63 log points per basis point, [−0.91, −0.36]; the predictable component enters with the opposite sign, 0.48 [−0.11, 1.07]; and equality of the two coefficients is rejected with $p = 0.003$ (Appendix Table A5).

The decomposition shows that purging predictable news matters for this outcome and in the direction that would attenuate unpurged series. Other differences between the series may also contribute to their weaker results. The USMPD surprises measure market moves in narrow windows around FOMC communications without that purge, while the Jarociński–Karadi series removes central-bank information effects identified from the co-movement of rates and stock prices, a different contamination; its estimates are also the least precise in Figure 2. We take the orthogonalized Bauer–Swanson contrast as the benchmark and keep the cross-series evidence in view when interpreting it.

2.5 How much can the origination margin move the stock?

The two sets of results are consistent with each other because the origination flow is small relative to the committed stock. Between 2012Q4 and 2023Q4, newly authorized dollars averaged 2.3 percent of total commitments per quarter, and new accounts were 3.4 percent

of open accounts in the median quarter. Suppose the estimated path of the median initial limit over the five quarters after a surprise applied to every newly authorized dollar, and every new account survived. Commitments would then be lower by 0.38 percent four quarters after the surprise, or 0.74 percent using the estimates that exclude 2020. The calculation carries the median response over to all new dollars, although the imprecisely estimated mean response is smaller on impact, and it ignores closures. Its figures can be set against the estimated response of commitments four quarters out, which is a different object. In the full sample that response is 0.66 percent of the lagged stock, [−0.19, 1.50], an interval that excludes a decline as large as this one; excluding 2020 it is −0.34, [−0.86, 0.17], which does not. The stability of the stock is therefore consistent with a small origination flow but offers no independent confirmation of the origination result, especially since closures and line changes on existing accounts move commitments by as much as new authorizations do.

The arithmetic has a substantive implication. Because nearly all card credit outstanding at any date was committed in earlier quarters, a monetary contraction that works through origination terms changes the quantity of card credit available to households slowly and by little. What it changes quickly is the credit available to households that need a new account. Sections 3 and 4 ask who those households are and how much their spending depends on it.

3 A household financing framework

The card-market evidence identifies where credit terms adjust. It leaves open whose spending depends on the adjustment. This section sets out a one-quarter financing framework that maps a change in authorized credit into a change in financeable spending, and defines two ways of removing the same amount of authorization: from existing lines, and from the initial limits of new accounts.

3.1 Households and available financing

The population is anchored to the 5,583 respondents of the 2024 Diary of Consumer Payment Choice, with the survey's weights. The Diary classifies each respondent at the start of the quarter into one of three groups: households with no credit card (G_0), cardholders who carry balances from month to month (revolvers, G_R), and cardholders who do not (non-revolvers, G_N), the distinction between revolving and convenience use (Fulford and Schuh, 2023). Membership in G_0 means currently without a card, not necessarily never having held one. Each respondent is matched to a complete Survey of Consumer Finances record, which supplies liquid assets X_i, authorized card limits A_i and card balances B_i, and to a complete Consumer Expenditure Survey record, which supplies quarterly income Y_i and an expenditure requirement D_i. Matching takes place within income and card-status cells and keeps each donor record intact, so relationships observed within a survey are preserved; within a cell we assign CEX records at random, which assumes that the SCF and CEX blocks are independent given the cell, and average results over 50 fixed assignments. We call each weighted Diary respondent, together with its matched SCF household record and CEX consumer-unit record, a synthetic household. Sections 3 and 4 report model calculations for this population; there, 'household' means a synthetic household, and dollar amounts are averages per household and quarter, in 2024 dollars. Appendix B describes the construction and its diagnostics.

A household's unused capacity on cards it already holds is

$$K_i^{old} = H_i \max\{A_i - B_i, 0\}, \tag{3}$$

where H_i indicates card ownership, so $K_i^{old} = 0$ in G_0. An authorized dollar can differ from an unused dollar, since part of a line may already be drawn. Let $M_i = Y_i + \lambda_X X_i$ be

current income plus usable liquid assets. Financeable spending in the quarter is

$$c_i = \min\{D_i,\ M_i + K_i\}, \qquad d_i = \left[D_i - M_i - K_i^{old}\right]_+, \tag{4}$$

where K_i is the household's unused card capacity and d_i is the shortfall that remains when only existing capacity is available. A smaller new line changes spending only for households with $d_i > 0$, and only until the shortfall is exhausted. A cut in an existing line can also bind a household with little slack, although such households account for about 0.1 percent of the existing-line loss below.

Because the expenditure and balance-sheet records come from different surveys, the reference state reproduces recorded spending for most, but not all, households. Before any policy change, 12.6 percent of households have $d_i > 0$ (between 11.4 and 13.7 percent across pairings). For these households the shortfall equals 39.5 percent of recorded expenditure (46 percent for those without a card); across all households it equals 5.4 percent of expenditure. The share of households with a shortfall is 36 percent among initially cardless households, 12 percent among revolvers and 4 percent among non-revolvers. Excluding housing outlays from the requirement lowers the overall share to 5.8 percent. These shortfalls combine financing pressure and matching: the CEX consumer unit did finance its recorded spending, by means the matched SCF record may not show. Since nearly all of the losses below arise in this region, we report how the results change with narrower definitions of D_i, under which the share falls to 3.5 percent, and across alternative dependence assumptions (Section 4.6).

The rule is an accounting device rather than a model of consumption choice. D_i is the benchmark expenditure requirement, the spending the rule finances when resources allow, close in spirit to the expenditure thresholds of Miranda-Pinto et al. (2025). We measure it by total quarterly expenditure in the matched CEX record, rather than consumption expenditure, because liquidity is fungible across outlays. The rule abstracts from borrowing

below capacity at a higher interest rate, from precautionary saving, and from intertemporal adjustment of the requirement. Its value is transparency: given the matched states, it states exactly which households a change in authorization can affect and by how much at most. Section 4.4 compares its per-dollar implications with quasi-experimental estimates.

3.2 Two ways to remove authorization

Existing lines. A proportional withdrawal κ from every existing cardholder's line replaces unused capacity with

$$K_i^{cut}(\kappa) = H_i \max\{(1-\kappa)A_i - B_i,\, 0\}, \tag{5}$$

holding balances fixed, and lowers spending by $\ell_i^{old} = \min\{D_i, M_i + K_i^{old}\} - \min\{D_i, M_i + K_i^{cut}(\kappa)\}$. Households in G_0 have no line to cut.

New accounts. A household in group g receives a new account during the quarter with probability p_g. The account carries an initial limit L drawn from a common distribution, of which a fraction ω is usable within the quarter. In the reference economy the limit is L; in the tighter-terms economy the same household receives the same account with limit $(1-\delta)L$. Existing lines are unchanged. Each new account adds to the household's existing capacity; the gross originations that calibrate it also include replacements of existing cards, which we do not model, and ω is the share of a new limit usable within the quarter rather than a replacement rate. The expected loss is

$$\ell_i^{new} = p_{g(i)}\, E_L\big[\min\{d_i, \omega L\} - \min\{d_i, \omega(1-\delta)L\}\big]. \tag{6}$$

Receipt is held fixed by construction, so the comparison isolates the terms of new credit; applications and approvals may well respond to policy as well.

A common amount. We set δ to the proportional equivalent of the estimated response of the median initial limit and compute the authorization it removes, $\mathcal{T} = \sum_i w_i p_{g(i)} E_L[\delta L]$,

where w_i are the normalized survey weights. We then choose κ so that the withdrawal from existing lines removes exactly the same amount, $\sum_i w_i \kappa A_i H_i = \mathcal{T}$. For any group g, let $\mathcal{L}_g = \sum_{i \in g} w_i \ell_i$ be its contribution to the expenditure loss and $\mathcal{A}_g$ its contribution to the authorization removed. We summarize incidence by

$$\eta_g = \mathcal{L}_g / \mathcal{A}_g, \tag{7}$$

the expenditure loss per dollar of authorization removed. It is a finite-change ratio under the rule in (4), distinct from an estimated marginal propensity to consume.

The two margins can differ for three reasons. They place the removed dollars on different households: a withdrawal from existing lines removes dollars in proportion to the limits households already hold, while tighter initial limits remove dollars in proportion to new accounts. They interact differently with the minimum in (4): a cut in an existing line removes capacity at the margin above a household's balance, while a new line adds capacity on top of whatever the household already has. And only the origination margin reaches households without a card. All calculations in the main text use the household state before any policy-induced change in income. Appendix C repeats them after a 100-basis-point cash-flow shock, which requires an external marginal propensity to consume and an income response, and reports the earlier one-percent withdrawal experiments on that scale.

4 The incidence of tighter origination terms

This section carries the estimated tightening of initial limits into the framework of Section 3, asks which households bear it, and compares it with an equal withdrawal from existing lines.

4.1 Calibrating new-card opportunities

Equation (6) requires the probability of receiving a new account in each group and the size distribution of initial limits. Table 2 lists every input and its status. Three data sources contribute, and they measure different units, so we connect them through explicit accounting restrictions rather than equating them.

For initially cardless households, adjacent annual waves of the Diary from 2021 to 2024 record 356 transitions from no card to card ownership among 1,663 valid transitions, a weighted annual entry rate of 21.3 percent. Assuming a constant hazard within the year gives a quarterly opportunity of 5.81 percent. This measures entry into current ownership, which differs from lifetime first adoption and from gross account openings.

The scale of new authorization comes from the bank data. In 2024, newly authorized dollars averaged 2.21 percent of the beginning-of-quarter authorized stock. Applying that ratio to the existing authorization of the population gives a reference budget of $T^{new} =$ \$607.06 per household-quarter. Initial limits follow a log-normal distribution matched to the 2024 mean and median initial limits, \$5,941 and \$3,663, which implies a dispersion of 0.98. The budget and the mean limit imply 10.2 new accounts per 100 households per quarter. Given the cardless entry rate, we solve for a common opportunity rate for revolvers and non-revolvers, 11.15 percent; the annual Diary data leave gross issuance between the two groups unidentified. An alternative route that starts from observed card counts and market account turnover implies 7.3 accounts per 100 households, and Section 4.6 reports the results under it.

The estimated response of the median initial limit implies a proportional reduction of $\delta = 4.53$ percent, which we apply to every simulated new limit. This shifts the whole distribution by the median's response and therefore removes more authorization than the imprecisely estimated mean response would; the sensitivity analysis matches both moments instead. The benchmark sets $\omega = 1$. The resulting authorization reduction, $\mathcal{T} =$ \$27.48

Table 2: Inputs to the household calculations

Object	Value	Source or status
Population shares W_0, W_R, W_N	17.4, 37.9, 44.8	2024 Diary, weighted (%)
Y_i, D_i	household-specific	matched CEX record
A_i, B_i, X_i	household-specific	matched SCF record
Usable share of liquid assets λ_X	1	assumption
Annual entry into card ownership $p_{0,y}$	21.3%	Diary panel, 2021–2024
Quarterly opportunity, initially cardless $p_{0,q}$	5.81%	constant-hazard conversion
Quarterly opportunity, cardholders $p_R = p_N$	11.15%	residual; equality assumed
New authorization relative to lagged stock	2.21%	Y-14M, 2024
New-authorization budget T^{new}	\$607.06	per synthetic household-quarter
Mean and median initial limit	\$5,941, \$3,663	Y-14M, 2024 dollars
Log-normal dispersion σ_L	0.98	implied by mean and median
Net usable fraction ω	1	assumption; lower values in text
Proportional cut in new limits δ	4.53%	from Table 1, median
Common authorization reduction $\mathcal{T}$	\$27.48	$\delta\, T^{new}$; used for both margins

All dollar amounts are 2024 dollars per synthetic household and quarter, averaged over the weighted 2024 Diary population; a synthetic household is one of the 5,583 Diary respondents with matched SCF household and CEX consumer-unit records. The receipt probabilities, the net usable fraction and the distribution-wide application of the median coefficient are calibration assumptions, not estimates. The 100-basis-point cash-flow experiment, which additionally requires an external marginal propensity to consume and an income response, is reported in Appendix C.

per household-quarter, is the common amount we also remove from existing lines.

4.2 Who bears tighter initial limits

Table 3: Tighter initial limits: incidence by initial card status

	Cardless (G_0)	Revolvers (G_R)	Non-revolvers (G_N)	All
Population share (%)	17.4	37.9	44.8	100.0
Quarterly receipt opportunity (%)	5.8	11.1	11.1	10.2
Share of new authorization (%)	9.9	41.3	48.8	100.0
Authorization reduction (\$)	2.71	11.35	13.42	27.48
Expenditure loss (\$)	0.26	0.43	0.20	0.88
Share of expenditure loss (%)	29.4	48.3	22.4	100.0
Loss per dollar of reduced authorization, η_g	0.095	0.038	0.015	0.032
Recipients whose spending falls (%)	18.4	6.5	2.4	

Model calculation. Every new initial limit is reduced by 4.53 percent, the proportional equivalent of the median response in Table 1; the same households receive new accounts in the reference and tighter-terms economies, and existing lines are unchanged. Dollar amounts are 2024 dollars per synthetic household-quarter; group dollar amounts are population contributions, already weighted by the receipt probability. Averages over 50 cross-survey pairings of the 5,583 synthetic households; shares are ratios of pairing means.

Table 3 reports the incidence. In the model, tighter initial limits reduce spending by \$0.88 per household in the quarter, 0.0048 percent of reference spending; across the five

expenditure definitions of Appendix Table D5, it ranges from \$0.12 to \$0.88. The aggregate is small because the flow of new authorization is small, as Section 2.5 implies.

The distribution of the loss is uneven. Initially cardless households receive 9.9 percent of new authorization dollars, in line with their lower opportunity rate, but bear 29.4 percent of the loss. Revolvers receive 41.3 percent and bear 48.3 percent; non-revolvers receive 48.8 percent and bear 22.4 percent. Per dollar of reduced authorization, the loss is 0.095 for cardless households, 0.038 for revolvers and 0.015 for non-revolvers (Figure 3). The ordering reflects existing capacity. A new account received by a household without a card can relieve a shortfall that no existing line covers; one received by a revolver supplements a partly used line; one received by a non-revolver usually adds to a large unused buffer. Among recipients, the reduction in the initial limit lowers spending for 18.4 percent of cardless households, 6.5 percent of revolvers and 2.4 percent of non-revolvers.

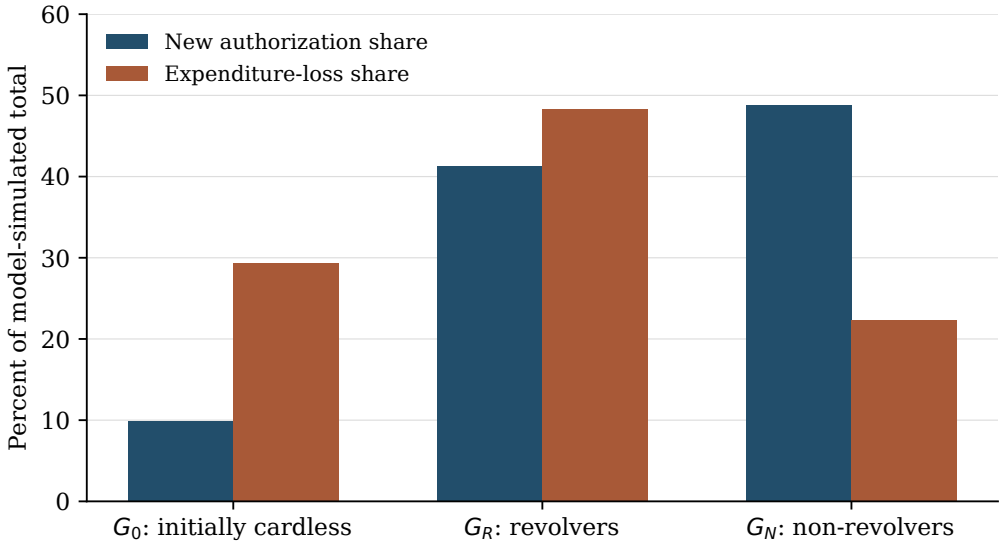


Figure 3: Shares of new authorization and of the expenditure loss, by initial card status
Notes: Model calculation of Table 3. Shares of the quarter's new authorization and of the expenditure loss from tighter initial limits, averaged over 50 pairings. These are model allocations rather than observed issuance shares.

Being without a card is an initial condition rather than a treatment. Cardless households are assigned zero existing card capacity by definition, but they keep their income, liquid assets and expenditure requirements from the matched records; the ordering comes from

how often and by how much those households face a shortfall.

4.3 The same authorization removed from existing lines

Table 4: The same authorization reduction on the two margins

	Cardless (G_0)	Revolvers (G_R)	Non-revolvers (G_N)	All
A. Tighter initial limits on new accounts				
Authorization reduction ($)	2.71	11.35	13.42	27.48
Expenditure loss ($)	0.259	0.426	0.197	0.882
Loss per dollar, η_g	0.0954	0.0375	0.0147	0.0321
B. Uniform withdrawal from existing lines				
Authorization reduction ($)	—	9.61	17.87	27.48
Expenditure loss ($)	—	0.275	0.101	0.376
Loss per dollar, η_g	—	0.0286	0.0057	0.0137
C. Withdrawal from existing lines of one group only				
Revolvers only: expenditure loss ($)	—	0.787	—	0.787
Revolvers only: loss per dollar, η	—	0.0286	—	0.0286
Non-revolvers only: expenditure loss ($)	—	—	0.155	0.155
Non-revolvers only: loss per dollar, η	—	—	0.0057	0.0057

Model calculation in the pre-shock household state. Each margin removes the same 27.48 dollars of authorization per synthetic household-quarter. Panel A reproduces Table 3. Panel B removes 0.10 percent of every existing cardholder's authorized limit, holding balances fixed; initially cardless households have no existing line. Panel C removes the full amount from one group, in proportion to its members' limits. The origination margin produces more loss per dollar in all 50 pairings; in the post-shock state of Appendix C the ratio of the two all-household values of η is 2.35. Averages over 50 pairings of the 5,583 synthetic households; dashes mark cells with no authorization removed.

Table 4 removes the same $27.48 of authorization per household from existing lines, which requires a uniform withdrawal of 0.10 percent of every cardholder's limit. The loss is $0.38, so the loss per dollar is 0.014, against 0.032 on the origination margin. A dollar of authorization removed at origination thus costs 2.35 times as much spending as a dollar removed from existing lines, or 1.8 cents more, and the ordering holds in all 50 pairings. In the household state after a 100-basis-point cash-flow shock the ratio is 2.35 (Appendix Table C1). We recompute both margins, with the existing-line withdrawal rate reset to remove the same amount, in each of 57 alternative populations: 35 cross-survey dependence structures, 20 balance-sheet donor assignments, and two allocations of new accounts tilted by liquid assets. The ratio ranges from 1.41 to 3.87 across populations and is above one in each of their 2,850 pairings, the lowest being 1.12 (Appendix Table D4).

In the baseline population the gap extends beyond the cardless households, whom existing lines miss entirely. Within revolvers, a dollar removed at origination costs 1.3 times as much spending as a dollar removed from existing lines; within non-revolvers, 2.6 times as much; among cardholders together, 1.8 times. In the 6 populations in which cardholders' expenditure requirements rise most steeply with their unused capacity, the cardholders-only ratio falls to 0.91, and the advantage of the origination margin comes entirely from cardless households. The difference lies in where each margin places the removed dollars. A proportional withdrawal from existing lines removes dollars in proportion to the limits households already hold, and those limits are concentrated among non-revolvers with large unused capacity, whose spending a small cut leaves unchanged. New authorization is spread across new accounts roughly independently of existing limits, so it reaches more households near their financing constraint. A withdrawal concentrated on revolvers comes closer: removing the full amount from revolvers' lines lowers spending by \$0.79, against \$0.88 on the origination margin, while removing it from non-revolvers' lines lowers spending by \$0.16. Under the uniform withdrawal, revolvers hold 35 percent of the authorization removed but bear 73 percent of the loss.

The comparison has a direct reading for transmission. In this card market, the quantity margin that monetary tightening moves is the one whose dollars matter more for current spending per dollar, but it is also the one that moves the fewest dollars.

4.4 Comparison with quasi-experimental responses

The per-dollar values in Tables 3 and 4 are lower than quasi-experimental estimates of responses to credit-line expansions (Table 5). Aydin (2022) finds that randomized line increases raise credit-line debt by about 0.11 per unit of additional limit in the first quarter and 0.28 after twelve quarters, and Gross and Souleles (2002) estimate that debt rises by 0.10 to 0.14 per dollar of limit over about a year. These estimates concern expansions for existing borrowers in other settings and differ in horizon and outcome, so they are

Table 5: Per-dollar responses: model and quasi-experimental benchmarks

Source	Response per dollar of credit limit	Object
This paper: new limits, η (all; G_0, G_R, G_N)	0.032; 0.095, 0.038, 0.015	one quarter; spending, hard financing cap
This paper: existing lines, η (all)	0.014	one quarter; spending, hard financing cap
Aydin (2022)	0.11 (first quarter); 0.28 (twelve quarters)	randomized line increases; credit-line debt per unit of limit
Gross and Souleles (2002)	0.10–0.14	about one year; change in debt per dollar of limit
This paper with $\phi = 0.11$: new limits; existing lines	0.139; 0.122	adds a uniform response away from the cap

The benchmarks are causal estimates for credit-line expansions to existing borrowers in other settings, not estimates of the objects computed here; they differ in horizon, outcome and population, and they concern line expansions, whose effects need not mirror those of reductions. ϕ is the response per dollar of usable capacity added to households whose financing constraint does not bind (Section 4.4); it is a sensitivity device, not an estimated parameter.

benchmarks rather than targets. One reason for the gap is structural. Under (4), a household with a slack financing constraint has no response to a change in its limit, whereas both studies find responses among borrowers far from their limits. In survey scenarios, by contrast, households report no spending out of the offer of a one-year interest-free loan (Fuster et al., 2021).

To gauge how much this matters for incidence, we add to each household's loss a response ϕ per dollar of usable capacity removed beyond the amount that binds, $\ell_i^\phi = \ell_i + \phi(\Delta K_i - \ell_i)$, with ϕ common to all households (Appendix Table D2). With $\phi = 0.05$, the ratio of the origination margin's loss per dollar to that of existing lines falls from 2.35 to 1.28, and the ratio between cardless households and non-revolvers falls from 6.5 to 2.2. With $\phi = 0.11$, the size of Aydin's first-quarter estimate, the two ratios fall to 1.14 and 1.6. The orderings survive in all pairings, partly by construction since ϕ is common across households; across the alternative populations the origination-to-existing-line ratio is at least 1.08 with $\phi = 0.05$ and 1.04 with $\phi = 0.11$. Absolute losses per dollar rise with ϕ on both margins; only their ratios fall. Under a common response, the hard-constraint ratios are thus the upper end of the relative contrasts. Responses that differ across groups, such as the larger borrowing responses of low-score borrowers documented by Agarwal et al.

(2018), could compress or widen the ratios.

4.5 What would reverse the comparison

Three departures from the baseline could reverse the ordering of the two margins, and each has a computable threshold (Appendix Table D3). First, because losses are linear in ϕ, the response away from the constraint would have to be larger for cuts in existing lines than for smaller new lines by 1.9 cents per dollar, close to the 1.8-cent difference between the margins under the hard constraint. Second, new accounts could go disproportionately to households without a shortfall. Only households with a shortfall lose from tighter initial limits, so scaling their probability of receipt by q, with each group's accounts and authorization held fixed, scales the origination loss by q; the ordering reverses at $q = 43$ percent, whereas the income-quintile allocations of Section 4.6 lower the origination loss to 95–99 percent of its baseline value and a two-to-one tilt toward households with high liquid assets lowers it to 77 percent. Third, we treat each measured shortfall as unmet need. Financing outside the framework would lower the origination loss but not the existing-line loss, since a constrained household loses an existing-line cut dollar for dollar; the ordering reverses once such financing covers 62 percent of every shortfall.

4.6 Sensitivity

Appendix Table D1 collects the remaining sensitivity analysis; Appendix D gives details. The definition of the expenditure requirement matters most for magnitudes. Replacing total expenditure with narrower definitions (excluding insurance and cash contributions, housing, or everything except nondurables and services) lowers the share of households constrained before any policy change to as little as 3.5 percent and the origination margin's loss to as little as \$0.12. The ratio of the two margins' losses per dollar stays between 2.1 and 2.7, and the ordering of cardless households, revolvers and non-revolvers holds in at

least 49 of 50 pairings for every definition (Appendix Table D5). The next largest source of uncertainty is the assumed dependence between expenditure requirements and financing resources across the two donor surveys. Varying the rank correlation of requirements with liquid assets among cardless households, and with unused existing capacity among cardholders, between -0.5 and 0.5 moves the aggregate loss between \$0.44–1.33 across 35 settings. The ordering of losses per dollar fails in 1 of 1,750 setting-by-pairing evaluations, and in none of 20 alternative balance-sheet donor assignments. Tilting receipt toward households with low or high liquid assets within each group moves the loss to \$1.08 or \$0.68. Reweighting receipt across income quintiles, using annual net card additions and observed card entry, lowers it to \$0.835; combined with the dependence setting in which that single reversal occurs, the ordering of cardless households and revolvers reverses in 6 of 50 pairings while the ordering of the averages is preserved.

The magnitude also depends on how the bank evidence is carried into the model. The account-count route for new-card volume lowers the loss to \$0.68. If cardless entrants receive a quarter of the average initial limit, with aggregate dollars held fixed, their share of the loss falls to 19 percent while their loss per dollar rises to 0.22. Matching both the median and the mean responses through a change in the shape of the limit distribution gives \$0.71, a proportional shift by the mean response alone gives \$0.27, and letting only half of each new limit be usable within the quarter gives \$0.68. These variations change the size of the aggregate loss considerably, from \$0.27 to \$0.71 among them, but not the conclusion that the origination margin reaches households with less existing capacity.

5 Implications for monetary transmission

The results describe how a commitment structure shapes the transmission of monetary policy through household credit. Three implications follow.

First, the evidence is consistent with commitments insulating the quantity of card credit

from monetary tightening, as loan commitments do for commercial lending (Sofianos et al., 1990; Morgan, 1998). The insulation depends on the shock. When banks face a wholesale funding shock, they do cut existing card lines, and constrained borrowers reduce their borrowing (Chava et al., 2023). Under the monetary surprises in our sample, the center of the committed distribution, total commitments and the incidence of line cuts do not contract, while the terms of new commitments tighten. Whether a shock reaches the committed stock is therefore an empirical property of the shock, not of revolving credit as such.

Second, the quantity of authorized card credit adjusts little and slowly. Total commitments show no measurable response to the surprise, and new authorizations are about 2.3 percent of commitments per quarter, so even the calculation of Section 2.5 changes the stock by less than one percent within a year. A persistent restriction can do more: in the model of Raveendranathan and Stefanidis (2025), the limit restrictions of the 2009 CARD Act account for more than half of the subsequent decline in revolving credit. The price of existing revolving balances moves within the quarter in the realized-rate data, and account-level evidence finds that card spending responds to such changes (Bräuning and Stavins, 2025; Grigoli and Sandri, 2026). The price and quantity results come from different designs and measure different objects, so decomposing transmission between them would require a common design.

Third, the results suggest a specific structure for borrowing limits in models of monetary transmission. Heterogeneous-agent models place household liquidity at the center of transmission (Kaplan et al., 2018; Auclert, 2019), and credit conditions typically enter through a borrowing limit common to all households or through a borrowing spread (Guerrieri and Lorenzoni, 2017; Faccini et al., 2026). The evidence here is consistent with limits that have a vintage structure: the committed stock does not contract measurably, while the terms on new accounts tighten. Mortgage models already give credit conditions such a vintage dimension, through payment-to-income limits that bind at origination (Greenwald,

2016) and through inherited contracts whose refinancing incentives and rate resets shape transmission (Di Maggio et al., 2017; Eichenbaum et al., 2022; Elias et al., 2026). Such a structure would imply a smaller and slower aggregate response of credit quantities than a common limit, and a concentrated response among households entering the card market or holding little unused capacity. The calculations of Section 4 indicate that this second group is where an additional dollar of authorization matters most for current spending. Since access to unsecured credit also shapes how households absorb job loss (Herkenhoff, 2019; Braxton et al., 2024), the distribution of that access across vintages may matter beyond the quarter we study.

Identifying the consumption effect of tighter origination terms would require account-level data on offers and originations linked to borrowers' existing lines and spending around identified shocks. The account-level FR Y-14M data used by Bord et al. (2026) and Bräuning and Stavins (2025) record limits and spending for individual accounts; linking new accounts to the same borrowers' existing lines and to their total spending would allow the incidence calculations here to be replaced by estimates.

6 Conclusion

Credit-card lines are authorized at one date and used at later ones. We find that monetary tightening adjusts the terms of new card credit, while the committed stock shows no measurable contraction. After a contractionary high-frequency surprise, the median initial limit on newly opened accounts falls relative to the limits on accounts already open, while total commitments, the number of open accounts and the incidence of line cuts do not move. The result is clearest under surprises purged of information that predicts them, and we show that the purged component matters for this outcome.

Because new authorizations are a small flow relative to the committed stock, this adjustment changes the quantity of card credit slowly and by little. It changes, instead, the

credit available to households that need a new account. In a household financing framework built from three surveys, a dollar of authorization removed at origination reduces spending more than a dollar removed from existing lines, and it reaches households with little or no unused capacity, including those without a card. The comparison holds across alternative survey matchings and expenditure definitions. The relative differences are largest under a rule in which only constrained households respond; with responses away from the constraint of the size found in quasi-experimental studies, absolute losses rise, the relative differences narrow, and their direction remains.

The analysis does not estimate the effect of monetary policy on household consumption through card credit, nor its welfare consequences. It shows that the margin on which credit adjusts and the financing position of the households it reaches belong together in any account of how monetary policy works through household credit.

References

Acosta, M., Ajello, A., Bauer, M. D., Loria, F., and Miranda-Agrippino, S. (2025). Financial market effects of FOMC communication: Evidence from a new event-study database. Working Paper 2025-30, Federal Reserve Bank of San Francisco.

Agarwal, S., Chomsisengphet, S., Mahoney, N., and Stroebel, J. (2018). Do banks pass through credit expansions to consumers who want to borrow? *Quarterly Journal of Economics*, 133(1):129–190.

Ampudia, M., Cooper, R., Le Blanc, J., and Zhu, G. (2024). MPC heterogeneity and the dynamic response of consumption to monetary policy. *American Economic Journal: Macroeconomics*, 16(3):343–388.

Argudo, E. P. (2025). Monetary policy transmission via nonbank lending: Evidence from peer-to-peer loans. *Journal of Financial Stability*, 80:101455.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–2367.

Aydin, D. (2022). Consumption response to credit expansions: Evidence from experimental assignment of 45,307 credit lines. *American Economic Review*, 112(1):1–40.

Baker, S. R., Farrokhnia, R. A., Meyer, S., Pagel, M., and Yannelis, C. (2023). Income, liquidity, and the consumption response to the 2020 economic stimulus payments. *Review of Finance*, 27(6):2271–2304.

Bassett, W. F., Chosak, M. B., Driscoll, J. C., and Zakrajšek, E. (2014). Changes in bank lending standards and the macroeconomy. *Journal of Monetary Economics*, 62:23–40.

Bauer, M. D. and Swanson, E. T. (2023). A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37:87–155.

Bord, V. M., Kovacs, A., and Moran, P. (2026). Automated credit limit increases and consumer welfare. *Journal of Monetary Economics*, 159:103901.

Bosshardt, J., Di Maggio, M., Kakhbod, A., and Kermani, A. (2024). The credit supply channel of monetary policy tightening and its distributional impacts. *Journal of Financial Economics*, 160:103914.

Bräuning, F. and Stavins, J. (2025). The credit card spending channel of monetary policy: Micro evidence from account-level data. Research Department Working Paper 25-10, Federal Reserve Bank of Boston.

Braxton, J. C., Herkenhoff, K. F., and Phillips, G. M. (2024). Can the unemployed borrow? Implications for public insurance. *Journal of Political Economy*, 132(9):3025–3076.

Chava, S., Ganduri, R., Paradkar, N., and Zeng, L. (2023). Shocked by bank funding shocks: Evidence from consumer credit cards. *Review of Financial Studies*, 36(10):3906–3952.

Cloyne, J., Ferreira, C., and Surico, P. (2020). Monetary policy when households have debt: New evidence on the transmission mechanism. *Review of Economic Studies*, 87(1):102–129.

Crawley, E. and Kuchler, A. (2023). Consumption heterogeneity: Micro drivers and macro implications. *American Economic Journal: Macroeconomics*, 15(1):314–341.

Cucic, D. and Gorea, D. (2026). Nonbank lending and the transmission of monetary policy. *Review of Financial Studies*, 39(4):966–1014.

den Haan, W. J., Sumner, S. W., and Yamashiro, G. M. (2007). Bank loan portfolios and the monetary transmission mechanism. *Journal of Monetary Economics*, 54(3):904–924.

Di Maggio, M., Kermani, A., Keys, B. J., Piskorski, T., Ramcharan, R., Seru, A., and

Yao, V. (2017). Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. *American Economic Review*, 107(11):3550–3588.

Drechsler, I., Jung, H., Peng, W., Supera, D., and Zhou, G. (2025). Credit card banking. Staff Report 1143, Federal Reserve Bank of New York.

Drozd, L. A. and Serrano-Padial, R. (2017). Modeling the revolving revolution: The debt collection channel. *American Economic Review*, 107(3):897–930.

Duca, J. V. and VanHoose, D. D. (1990). Loan commitments and optimal monetary policy. *Journal of Money, Credit and Banking*, 22(2):178–194.

Eichenbaum, M., Rebelo, S., and Wong, A. (2022). State-dependent effects of monetary policy: The refinancing channel. *American Economic Review*, 112(3):721–761.

Elias, M., Gillitzer, C., Kaplan, G., La Cava, G., and Prasad, N. (2026). Household liquidity and the adjustable-rate mortgage cash-flow channel of monetary policy. *Journal of Monetary Economics*, 163:103994.

Faccini, R., Lee, S., Luetticke, R., Ravn, M. O., and Renkin, T. (2026). Financial frictions: Micro vs. macro volatility. *American Economic Review*, 116(2):464–501.

Fulford, S. L. (2015). How important is variability in consumer credit limits? *Journal of Monetary Economics*, 72:42–63.

Fulford, S. L. and Schuh, S. D. (2023). Revolving versus convenience use of credit cards: Evidence from U.S. credit bureau data. *Journal of Money, Credit and Banking*, 55:1667–1701.

Fulford, S. L. and Schuh, S. D. (2024). Credit cards, credit utilization, and consumption. *Journal of Monetary Economics*, 148:103619.

Furceri, D., Mazzola, F., and Pizzuto, P. (2019). Asymmetric effects of monetary policy shocks across U.S. states. *Papers in Regional Science*, 98(5):1861–1891.

Fuster, A., Kaplan, G., and Zafar, B. (2021). What would you do with $500? Spending responses to gains, losses, news, and loans. *Review of Economic Studies*, 88(4):1760–1795.

Galenianos, M. and Gavazza, A. (2022). Regulatory interventions in consumer financial markets: The case of credit cards. *Journal of the European Economic Association*, 20(5):1897–1932.

Gertler, M. and Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76.

Greenwald, D. L. (2016). The mortgage credit channel of macroeconomic transmission. Working paper, MIT Sloan School of Management.

Greenwald, D. L., Krainer, J., and Paul, P. (2025). The credit line channel. *Journal of Finance*, 80(6):3137–3183.

Grigoli, F. and Sandri, D. (2026). Monetary policy and credit card spending. *European Economic Review*, 185:105294.

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

Gross, D. B. and Souleles, N. S. (2002). Do liquidity constraints and interest rates matter for consumer behavior? *Quarterly Journal of Economics*, 117(1):149–185.

Gross, T., Notowidigdo, M. J., and Wang, J. (2020). The marginal propensity to consume over the business cycle. *American Economic Journal: Macroeconomics*, 12(2):351–384.

Guerrieri, V. and Lorenzoni, G. (2017). Credit crises, precautionary savings, and the liquidity trap. *Quarterly Journal of Economics*, 132(3):1427–1467.

Gürkaynak, R. S., Sack, B., and Swanson, E. T. (2005). Do actions speak louder than words? the response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, 1(1):55–93.

Han, S., Keys, B. J., and Li, G. (2018). Unsecured credit supply, credit cycles, and regulation. *Review of Financial Studies*, 31(3):1184–1217.

Herkenhoff, K. F. (2019). The impact of consumer credit access on unemployment. *Review of Economic Studies*, 86(6):2605–2642.

Holm, M. B., Paul, P., and Tischbirek, A. (2021). The transmission of monetary policy under the microscope. *Journal of Political Economy*, 129(10):2861–2904.

Jarociński, M. and Karadi, P. (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43.

Johnson, D. S., Parker, J. A., and Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, 96(5):1589–1610.

Jordà, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to HANK. *American Economic Review*, 108(3):697–743.

Littlejohn, M. (2025). The credit card and small business lending channels of monetary policy. *Journal of Economic Dynamics and Control*, 171:105045.

Ludvigson, S. (1998). The channel of monetary transmission to demand: Evidence from the market for automobile credit. *Journal of Money, Credit and Banking*, 30(3):365–383.

Miranda-Agrippino, S. and Ricco, G. (2021). The transmission of monetary policy shocks. *American Economic Journal: Macroeconomics*, 13(3):74–107.

Miranda-Pinto, J., Murphy, D., Walsh, K. J., and Young, E. R. (2025). A model of expenditure shocks. *Journal of Monetary Economics*, 154:103807.

Morgan, D. P. (1998). The credit effects of monetary policy: Evidence using loan commitments. *Journal of Money, Credit and Banking*, 30(1):102–118.

Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics*, 133(3):1283–1330.

Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, 103(6):2530–2553.

Raveendranathan, G. and Stefanidis, G. (2025). The unprecedented fall in U.S. revolving credit. *International Economic Review*, 66(1):393–451.

Raveendranathan, G., Stefanidis, G., and Sublet, G. (2026). Rate caps on revolving credit lines. *Journal of Political Economy Macroeconomics*, 4(1):164–210.

Santucci, L. (2015). A tale of two vintages: Credit limit management before and after the CARD Act and Great Recession. Payment Cards Center Discussion Paper DP 15-01, Federal Reserve Bank of Philadelphia.

Sofianos, G., Melnik, A., and Wachtel, P. (1990). Loan commitments and monetary policy. *Journal of Banking and Finance*, 14(4):677–689.

Woodford, M. (1996). Loan commitments and optimal monetary policy. *Journal of Monetary Economics*, 37(3):573–605.

Online Appendix

A Card-market data and empirical specifications

A.1 Series

Table A1 lists the card-market series, their units and coverage. Initial and open-account medians describe changing populations of accounts. The mean initial limit is newly authorized dollars divided by new accounts. The flow-to-stock ratios used in the household calibration divide by the beginning-of-quarter stock.

Table A1: Credit-market series and measurement

Series	Units / transformation	Coverage	Source
Assessed-interest APR	pp, nominal annual rate	2012Q3–2025Q4	G.19/FRED
Open limit, median	real dollars; log outcome	2012Q3–2025Q4	FR Y-14M/FRED card panel
Initial limit, median	real dollars; log outcome	2012Q3–2025Q4	FR Y-14M public aggregates
Gross new authorization	billions of dollars; deflated in LP	2012Q3–2025Q4	FR Y-14M public aggregates
Gross new accounts	millions; no CPI deflation	2012Q3–2025Q4	FR Y-14M public aggregates
Open accounts	millions; stock	2012Q3–2025Q4	FR Y-14M public aggregates
Open authorized stock	billions of dollars	2012Q3–2025Q4	FR Y-14M public aggregates
Share with line increase	percentage points	2012Q3–2025Q4	FR Y-14M/FRED card panel
Share with line decrease	percentage points	2012Q4–2025Q4	FR Y-14M/FRED card panel
Active utilization, median	percentage points	2012Q3–2025Q4	FR Y-14M/FRED card panel
Minimum-payment share	percentage points	2012Q3–2025Q4	FR Y-14M/FRED card panel
Past-due account share	percentage points	2012Q3–2025Q4	FR Y-14M/FRED card panel

Notes: Coverage before regression-specific restrictions. Initial means are constructed as gross authorization dollars divided by gross new accounts; initial and open medians describe changing account populations. Dollar outcomes in the local projections are deflated by the quarterly CPI. The flow/stock transport uses lagged stock, not contemporaneous stock.

A.2 Horizons

Table A2 reports the responses of the two medians and of their directly estimated contrast at horizons zero to four. The first quarter of 2020 carries the largest Bauer–Swanson surprise in the sample and coincides with the pandemic collapse in originations, so the right-hand columns drop every observation whose window touches 2020. Excluding those windows makes the contrast larger at every horizon.

Table A2: Responses by horizon

			Full sample			Excluding windows touching 2020		
h	New $p50$	Open $p50$	Contrast	95% CI	N	Contrast	95% CI	N
0	−4.63	0.58	−5.22	[−7.54, −2.89]	45	−5.59	[−9.45, −1.73]	40
1	−6.78	0.52	−7.30	[−10.82, −3.77]	45	−8.70	[−13.06, −4.33]	39
2	−3.42	0.90	−4.32	[−9.89, 1.24]	45	−8.26	[−14.71, −1.80]	38
3	1.08	1.09	−0.02	[−4.81, 4.78]	45	−3.79	[−9.21, 1.62]	37
4	−2.60	1.44	−4.05	[−8.66, 0.56]	45	−8.36	[−15.70, −1.03]	36

Responses $100[\log Z_{t+h} - \log Z_{t-1}]$ per one-standard-deviation Bauer–Swanson surprise (7.89 basis points), with quarter-of-year effects. Newey–West lag $\min\{\max(1, h+1), 8\}$. The right-hand columns drop every observation whose window $[t-1, t+h]$ touches 2020; the first quarter of 2020 carries the largest surprise in the sample.

Table A3 varies the impact specification of Table 1 by dropping each quarter in turn, changing the Newey–West bandwidth and adding lagged terms.

Table A3: Sensitivity of the impact contrast to single quarters, bandwidth and dynamics

Variation	Specifications	Estimate	95% CI
Baseline (Table 1)	1	−5.22	[−7.54, −2.89]
Dropping one quarter at a time	45	−5.95 to −4.41	upper endpoint at most −1.83
Newey–West lags 0, 1, 2, 4, 8	5	−5.22	upper endpoint at most −2.65
Adding lagged outcome and shock terms	3	−4.85 to −4.50	upper endpoint at most −1.97

New-minus-open median contrast at $h = 0$ per one-standard-deviation Bauer–Swanson surprise, with the controls and quarter-of-year effects of Table 1. The lag specifications add, in turn, the lagged change in the outcome; that change and the lagged surprise; and the lagged level of the outcome and the lagged surprise. Every interval excludes zero.

A.3 Surprise series and the predictable component

Table A4 reports the estimates plotted in Figure 2. Comparing the published orthogonalized and unorthogonalized Bauer–Swanson surprises on a common basis-point scale, the

ratio of their per-basis-point contrasts is 0.61 with quarter-of-year effects and between 0.51 and 0.68 across specifications. Table A5 reports the joint regression on the orthogonal and predictable components.

Table A4: Origination and open-account medians across surprise series

Series (one SD)	Specification	New $p50$	Open $p50$	Contrast	95% CI	N
Bauer–Swanson, orthogonalized (7.9 bp)	no quarter effects	−5.71	0.70	−6.41	[−8.28, −4.54]	45
	quarter effects	−4.63	0.58	−5.22	[−7.54, −2.89]	45
	quarter effects, excl. 2020	−4.78	0.81	−5.59	[−9.45, −1.73]	40
Bauer–Swanson, unorthogonalized (10.0 bp)	no quarter effects	−5.13	0.41	−5.53	[−8.01, −3.06]	45
	quarter effects	−3.79	0.25	−4.05	[−6.84, −1.25]	45
	quarter effects, excl. 2020	−3.24	0.34	−3.58	[−7.69, 0.53]	40
SF Fed USMPD, monetary event (6.3 bp)	no quarter effects	−2.47	0.33	−2.80	[−6.14, 0.54]	53
	quarter effects	−0.98	0.26	−1.24	[−4.31, 1.84]	53
	quarter effects, excl. 2020	−0.08	0.27	−0.34	[−2.80, 2.11]	48
SF Fed USMPD, statement (5.5 bp)	no quarter effects	−3.75	0.26	−4.01	[−7.98, −0.04]	53
	quarter effects	−2.19	0.18	−2.37	[−6.30, 1.56]	53
	quarter effects, excl. 2020	−0.34	0.24	−0.58	[−4.90, 3.73]	48
Jarociński–Karadi, monetary (8.5 bp)	no quarter effects	−1.53	0.26	−1.80	[−7.05, 3.46]	46
	quarter effects	−0.21	0.15	−0.36	[−6.19, 5.48]	46
	quarter effects, excl. 2020	−0.58	0.24	−0.82	[−5.88, 4.25]	41

Impact responses per one standard deviation of each series (shown in basis points); the series differ in event windows, instruments and units, so per-standard-deviation magnitudes are not directly comparable across rows. Each row uses the common sample available for that series and specification. The unorthogonalized Bauer–Swanson series is the published surprise before removal of the component predictable from pre-announcement data.

Table A5: Orthogonal and predictable components of the Bauer–Swanson surprise

Specification	Orthogonal	95% CI	Predictable	95% CI	p, equal	N
no quarter effects	−0.782	[−0.989, −0.576]	0.489	[−0.079, 1.056]	0.0002	45
quarter effects	−0.632	[−0.909, −0.356]	0.481	[−0.111, 1.072]	0.0033	45
quarter effects, excl. 2020	−0.706	[−1.195, −0.218]	0.547	[−0.213, 1.307]	0.0114	40

Regression of the new-minus-open median contrast at $h = 0$ jointly on the orthogonalized surprise and on the predictable component (unorthogonalized minus orthogonalized surprise), both in basis points; coefficients are log points per basis point. In the estimation sample the two surprise series have correlation 0.86 and the predictable component has a standard deviation of 4.3 basis points. The last column tests equality of the two coefficients with the Newey–West covariance.

A.4 Credit-score composition

Table A6: Credit-score composition of new accounts

	Quarter-of-year effects		No quarter-of-year effects		
	Estimate	95% CI	Estimate	95% CI	N
Median credit score, new accounts (points)	−0.52	[−1.27, 0.24]	−0.29	[−0.82, 0.24]	45
25th percentile score, new accounts (points)	−0.61	[−1.39, 0.18]	−0.42	[−1.01, 0.18]	45
Median credit score, open accounts (points)	0.20	[−0.02, 0.42]	−0.02	[−0.28, 0.25]	45
Share of new accounts with score below 660 (pp)	0.21	[−0.12, 0.53]	0.16	[−0.12, 0.44]	45
Median initial limit, score below 660	−3.11	[−9.15, 2.93]	−4.10	[−9.14, 0.94]	45
Median initial limit, score 660–719	−3.21	[−5.24, −1.18]	−4.47	[−6.62, −2.32]	45
Median initial limit, score 720 or above	−0.83	[−2.38, 0.71]	−3.08	[−5.12, −1.03]	45
New minus open, all new accounts	−5.22	[−7.54, −2.89]	−6.41	[−8.28, −4.54]	45
New minus open, controlling for the change in median score	−4.25	[−6.42, −2.07]	−5.95	[−7.57, −4.34]	45

Impact responses per one-standard-deviation Bauer–Swanson surprise. The left columns use the benchmark specification of Table 1; the right columns omit quarter-of-year effects. Limit rows are log points of the median initial limit within the stated score band. Credit scores are those reported in the public large-bank series. The controlled contrast adds the change in the median score of new accounts as a regressor.

A.5 Standing-credit margins

Table A7 collects the associations with realized policy-rate changes, the high-frequency results that cut against the main reading of the stock evidence, and the bank-holding-company panel.

Table A7: Standing-credit margins: realized-rate associations, fragile results, and the long bank panel

	Estimate	95% CI	N
A. Realized federal funds rate, per 100 basis points (associational)			
Assessed-interest APR (pp)	1.08	[0.83, 1.33]	53
Open-account limit, median (%)	0.30	[−0.84, 1.44]	53
Open-account limit, 75th percentile (%)	−0.19	[−0.91, 0.53]	53
Open-account limit, 90th percentile (%)	−0.27	[−0.87, 0.34]	53
Line-decrease share (pp)	−0.02	[−0.10, 0.07]	52
Line-increase share (pp)	0.38	[0.00, 0.75]	53
Net account exit rate (pp)	−0.02	[−0.88, 0.85]	52
Active utilization, median (pp)	1.07	[0.37, 1.78]	53
Minimum-payment share (pp)	0.60	[0.34, 0.85]	53
Thirty-day past-due share (pp)	0.13	[0.06, 0.21]	53
B. Unfavorable or fragile high-frequency results disclosed in the text			
75th percentile, open accounts; Jarociński–Karadi	−0.78	[−1.35, −0.21]	46
90th percentile, open accounts; USMPD statement	−0.31	[−0.62, 0.01]	53
Line-decrease share, one quarter ahead of surprise (placebo)	0.06	[0.01, 0.11]	42
Line-increase share, excl. 2020, $h = 0$	−0.154	[−0.303, −0.005]	39
Line-increase share, excl. 2020, $h = 1$	−0.199	[−0.377, −0.020]	38
Line-increase share, excl. 2020, $h = 2$	−0.156	[−0.306, −0.006]	37
C. Bank holding companies, unused card commitments, per 100 basis points (log points, s.e.)			
Before 2010 (77 quarters)	−0.04	(0.45)	58554
After 2010 (33 quarters)	−1.01	(4.40)	8982

Panel A: impact responses to a 100-basis-point quarterly change in the federal funds rate, 2012Q3–2025Q4; associational, since realized changes respond to economic conditions. Limit responses are exact percentage changes of real limits. Panel B: impact responses per one standard deviation of the stated surprise series (Bauer–Swanson unless noted); the placebo regresses the line-decrease share on the next quarter's surprise. Panel C: FR Y-9C, institution fixed effects and lagged bank controls, standard errors clustered by institution and quarter; the reporting item nets drawn balances and changes definition in 2010, so the two regimes are estimated separately; N is bank-quarters.

B Household data construction

B.1 Anchors, donors and units

The population consists of the 5,583 respondents of the 2024 Diary of Consumer Payment Choice with positive weights, normalized to sum to one, so that dollar amounts are population averages per synthetic household. Card ownership and revolving status come from the Diary: a cardholder is a revolver if the respondent reports carrying a balance during the preceding year. Each respondent is assigned an intact record from the pooled 2016, 2019, 2022 Surveys of Consumer Finances within cells defined by income and revolving status. The SCF record supplies liquid assets, total card limits and card balances; SCF revolving status is based on repayment habits, which is related to, but not identical with, the Diary measure. Initially cardless respondents keep zero card limits and balances by construction and may receive a liquid-asset donor. Each respondent also receives an intact Consumer Expenditure Survey Interview record within income cells, which supplies quarterly income and total quarterly expenditure. All dollar amounts are in 2024 prices and quarterly.

Keeping each donor record intact preserves the joint distribution observed within each survey, for example between a household's card limit and its balance, or between its income and spending. The dependence between SCF and CEX variables within a matching cell is not observed. The baseline assigns CEX records at random within cells; we average over 50 such assignments and vary the dependence explicitly in Appendix D. Variation across assignments describes sensitivity to the matching procedure, not sampling uncertainty.

B.2 Entry into card ownership

Table B1 reports annual entry into card ownership among initially cardless Diary respondents, from adjacent waves between 2021 and 2024, and the constant-hazard conversion to a quarterly rate. Table B2 reports annual net additions to cards held by baseline income

quintile, which we use only to reallocate receipt across income cells in the sensitivity analysis. Net additions are not gross issuance: closures and replacements between waves can leave the number of cards unchanged.

Table B1: Observed annual ownership entry and the quarterly timing bridge

Transition/sample	N	Events	Rate (%)	Cluster SE (pp)	95% CI (%)
Pooled	1,663	356	21.278	1.403	[18.528, 24.027]
2021–2022	578	123	21.759	2.202	[17.443, 26.075]
2022–2023	516	126	23.578	2.464	[18.748, 28.407]
2023–2024	569	107	18.573	2.171	[14.318, 22.829]
Common IPW sample	1,607	345	21.612	1.440	[18.790, 24.434]
Prior stabilized IPW	1,607	345	21.613	1.452	[18.768, 24.458]

Annual positive ownership at the next wave conditional on explicit zero ownership at baseline; original baseline weights, adjacent 2021–2024 waves. SEs cluster repeated transitions by respondent; normal 95% intervals are clipped to [0,1]. Topcoded positive counts remain valid ownership endpoints. Pooled distinct respondents: 985. Constant-hazard quarterly rate: 5.81%, transformed 95% interval [4.99, 6.64]%. The inverse-probability-weighted rows use retention weights fitted on the common complete-case sample. Ownership entry is not gross issuance or necessarily lifetime first adoption.

Table B2: Annual net additions to cards held by baseline-income quintile

Group	n	Kish n	Weight share (%)	Rate (%)	SE (pp)	95% CI (%)
Q1	1,676	889.1	20.81	19.33	1.39	[16.60, 22.05]
Q2	1,733	939.6	21.96	22.28	1.34	[19.66, 24.90]
Q3	1,707	918.2	20.76	22.36	1.37	[19.68, 25.04]
Q4	1,392	811.1	18.51	26.99	1.52	[24.00, 29.97]
Q5	1,330	814.4	17.96	25.14	1.56	[22.09, 28.20]

Year-specific equality tests (five groups)

Transition	Transitions	$\chi^2(4)$	p-value
Pooled	7,838	16.177	0.002791
2021–2022	2,600	6.038	0.196357
2022–2023	2,438	2.822	0.588029
2023–2024	2,800	12.774	0.012437

Notes: Descriptive survey estimates; no monetary-policy shock is imposed. Annual positive net additions are $1(cc_{y+1} > cc_y)$ among observed adjacent transitions with exact counts from zero through five at both endpoints, valid baseline income and age, and positive baseline weights. Category > 5 endpoints are excluded; net additions are not gross issuance. Income cutoffs use each full baseline wave's weighted quintiles, retain ties, and need not give equal sample shares. n counts transitions; Kish n describes weight dispersion and is not an independent-cluster count. Confidence intervals use respondent-clustered ratio-estimator standard errors. The pooled joint test includes cross-group covariance from respondents who change income group and uses a common cluster correction. It rejects equality of these annual net-addition rates, not equality of gross quarterly openings. The supplied survey design does not support a full strata/PSU variance calculation here. Attrition remains complete-case selection with original weights.
Source: 2021–2024 Survey and Diary of Consumer Payment Choice individual releases with baseline respondent weights; 3,937 distinct pooled respondents. These are observed descriptive panel estimates, not model-simulated new-card allocations.

B.3 From market flows to household opportunities

Table B3 reports the 2024 Y-14M accounting behind the new-authorization budget and the alternative account-count route. The two routes transport market accounts to survey households through different denominators, and they imply 10.2 and 7.3 new accounts per 100 households per quarter. We keep both rather than forcing them to agree.

Table B3: Gross-account accounting and the alternative count bridge

Quarter	New auth. (bn)	Lagged auth. (bn)	New accounts (mn)	Flow/stock (%)	Mean limit (USD)
2024Q1	103.872	4,596.23	16.80	2.2599	6,182.88
2024Q2	106.219	4,595.70	17.53	2.3113	6,059.27
2024Q3	99.436	4,625.49	16.93	2.1497	5,873.36
2024Q4	99.404	4,656.05	17.60	2.1349	5,647.98

Y-14M aggregate accounting; authorization columns are 2024 real dollars, account counts are millions. Each lagged stock uses its own CPI date. Quarterly flow is gross authorization, not net balance growth. The lower panel transports these account-based denominators to survey respondents, assigning at most one prospective card per receiving household; it is a diagnostic calibration, not an observed count.

Route	Group	Accounts/100	New dollars	Cut dollars	Loss
Main	G0	1.009	59.944	2.714	0.259
Main	GR	4.220	250.706	11.350	0.426
Main	GN	4.989	296.413	13.419	0.197
Main	ALL	10.218	607.062	27.483	0.882
Count route	G0	1.009	59.944	2.714	0.259
Count route	GR	2.876	170.839	7.734	0.290
Count route	GN	3.400	201.985	9.144	0.134
Count route	ALL	7.285	432.768	19.592	0.684

C The 100-basis-point cash-flow experiment

The main text evaluates both credit margins in the household state before any policy-induced change in income. This appendix places the calculations in a state that has already absorbed a cash-flow shock, and reports the withdrawal experiments of earlier versions of this paper on that scale.

C.1 The post-shock state

A 100-basis-point tightening lowers each household's quarterly income by $\Delta Y_i = \delta_Y Y_i$ with $\delta_Y = 0.54$ percent, a scale taken from the state-level personal-income responses of Furceri et al. (2019). The household's requirement falls by an ordinary spending response $m\,\Delta Y_i$ with $m = 0.25$, an external calibration consistent with short-run spending responses to transitory income (Johnson et al., 2006; Parker et al., 2013; Baker et al., 2023); responses to income losses need not mirror responses to gains (Fuster et al., 2021). With existing lines unchanged, post-shock spending is

$$c_i^{stable} = \min\{D_i - m\,\Delta Y_i,\ Y_i - \Delta Y_i + \lambda_X X_i + K_i^{old}\}. \tag{8}$$

Neither m nor δ_Y is estimated from the card or survey data used here, and neither enters the main text.

C.2 Same authorization in the post-shock state

Table C1 repeats the comparison of Table 4 in both states. The losses per dollar on each margin barely change, and the origination margin's loss per dollar remains 2.35 times that of existing lines.

C.3 One-percent withdrawals from existing lines

Earlier versions of this paper studied a hypothetical synchronized withdrawal of one percent of every existing line in the post-shock state. We report those calculations for completeness. A one-percent withdrawal from all cardholders (CF0) removes \$274.20 of authorization per synthetic household and lowers spending by an additional \$3.79; the modeled spending decline after the cash-flow shock grows from \$39.87 to \$43.66 per synthetic household, 9.49 percent larger on average across pairings. This is the consequence of a hypothetical with-

Table C1: Same authorization reduction: both household states

State	Margin	Loss ($)	Pairing range	η	η_0	η_R	η_N	Affected (%)
Pre-shock	New limits	0.882	[0.749, 1.021]	0.0321	0.0954	0.0375	0.0147	0.58
	Existing lines, uniform	0.376	[0.295, 0.504]	0.0137	—	0.0286	0.0057	5.62
	Existing lines, revolvers only	0.787	[0.580, 1.073]	0.0286	—	0.0286	—	3.77
	Existing lines, non-revolvers only	0.155	[0.089, 0.239]	0.0057	—	—	0.0057	1.85
Post-shock	New limits	0.887	[0.753, 1.026]	0.0323	0.0958	0.0377	0.0148	0.58
	Existing lines, uniform	0.378	[0.301, 0.505]	0.0137	—	0.0287	0.0057	5.65
	Existing lines, revolvers only	0.790	[0.583, 1.075]	0.0288	—	0.0288	—	3.79
	Existing lines, non-revolvers only	0.156	[0.089, 0.241]	0.0057	—	—	0.0057	1.86

Every row removes 27.48 dollars of authorization per synthetic household-quarter. The pre-shock state is the reference economy of Section 3; the post-shock state first applies the 100-basis-point cash-flow experiment of Appendix C. Pairing range is the minimum and maximum over the 50 pairings. Affected is the share of the population whose spending falls.

drawal of a stated size, not a bound on actual transmission: the public open-account data cannot establish that no line falls by more than one percent. Restricting the withdrawal to revolvers (CF1) or non-revolvers (CF2) gives $2.77 and $1.02; revolvers bear 73 percent of the CF0 loss while holding 35 percent of authorization. Removing the full CF0 dollar amount from revolvers only (CF3) or non-revolvers only (CF4) gives $7.98 and $1.57. Tables C2 to C6 report the details and the sensitivity of these calculations to the ordinary spending response, to the definition of the expenditure requirement, and to a borrowing-cost response that discourages drawing on unchanged lines.

Table C2: Existing authorization withdrawal and expenditure incidence

Quantity	All cardholders	Revolvers only	Non-revolvers only
Authorization withdrawn (dollars)	274.195	95.874	178.322
Additional expenditure loss (dollars)	3.787	2.769	1.018
Revolver component (dollars)	2.769	2.769	0.000
Non-revolver component (dollars)	1.018	0.000	1.018
Rate on revolvers (percent)	1.000	1.000	0.000
Rate on non-revolvers (percent)	1.000	0.000	1.000
Revolver population share (percent)	37.859	37.859	37.859
Non-revolver population share (percent)	44.761	44.761	44.761
Revolver old-authorization share (percent)	34.965	34.965	34.965
Non-revolver old-authorization share (percent)	65.035	65.035	65.035

Model-simulated outcomes, averaged over 50 cross-survey pairings. The experiment combines the 100-basis-point cash-flow shock of Appendix C with a hypothetical 1% withdrawal of authorized credit in the eligible group. Dollars are 2024 real dollars per synthetic household-quarter. Balances are fixed; initially cardless units have no old line to cut. No sampling confidence intervals are inferred from pairings.

Table C3: Existing-line accounting and pairing variation

Arm	Unused capacity removed	Mean loss	Pairing min	Pairing max	Affected population (%)
CF0	272.603	3.787	3.025	5.060	5.679
CF1	94.290	2.769	2.035	3.770	3.812
CF2	178.314	1.018	0.589	1.564	1.867

Model-simulated outcomes, averaged over 50 cross-survey pairings. The experiment combines the 100-basis-point cash-flow shock of Appendix C with a hypothetical 1% withdrawal of authorized credit in the eligible group. Dollars are 2024 real dollars per synthetic household-quarter. Balances are fixed; initially cardless units have no old line to cut. No sampling confidence intervals are inferred from pairings. Pairing extrema describe assignment sensitivity, not sampling intervals. Affected shares count households whose spending falls by more than 10^{-8} dollars.

Table C4: Equal-dollar authorization withdrawal

Arm	R cut (%)	N cut (%)	Authorization removed	Expenditure loss	Unused capacity removed
CF3	2.860	0.000	274.195	7.981	269.114
CF4	0.000	1.538	274.195	1.568	274.183

Model-simulated outcomes, averaged over 50 cross-survey pairings. The experiment combines the 100-basis-point cash-flow shock of Appendix C with a hypothetical 1% withdrawal of authorized credit in the eligible group. Dollars are 2024 real dollars per synthetic household-quarter. Balances are fixed; initially cardless units have no old line to cut. No sampling confidence intervals are inferred from pairings. CF3 and CF4 withdraw the CF0 gross dollar budget; their percentage cuts differ. These are conditional quantity allocations, not estimated bank policies or welfare changes.

Table C5: Ordinary expenditure response and requirement sensitivity

MPC case	Mean MPC	Stable response	Ordinary term	Financing residual
$m = 0.25$	0.250	−39.871	−35.168	−4.703
$m = 0.15$	0.150	−26.434	−21.101	−5.333
$m = 0.40$	0.400	−60.028	−56.268	−3.760
Liquidity-graded MPC	0.263	−40.219	−36.102	−4.117

External calibrations of the ordinary MPC in the 100-basis-point cash-flow experiment; baseline uses CEX total expenditure. Columns decompose the stable-old-line response into the ordinary term and the financing residual, preventing double counting. Sensitivities are scenarios, not MPC estimates from matched records. Requirement results below use the same experiment with a one-percent authorization cut.

Requirement	Stable response	Cut response	Amplification (%)
Total expenditure	−39.871	−43.658	9.492
Excluding pensions/transfers	−38.974	−42.065	7.926
Nonhousing requirement	−37.216	−39.098	5.052
Legacy nondurable requirement	−36.617	−37.977	3.711
Nondurables and services	−35.780	−36.230	1.257

Table C6: Borrowing-cost diagnostics for the omitted price channel

Borrowing decline (%)	External anchor	Capacity-only response	Cost-wedge response	Stock amplification (%)
0.00	Scenario	−39.871	−39.871	9.492
1.50	Scenario	−39.910	−46.927	8.060
4.00	Scenario	−39.986	−58.690	6.392
5.60	GS short run	−40.040	−66.216	5.628
8.40	GS long run	−40.116	−79.399	4.675
15.00	Scenario	−40.317	−110.446	3.105

Soft-borrowing diagnostics on the scale of the 100-basis-point experiment. GS denotes the external Gross–Souleles borrowing-response calibration, at its stated short and long horizons. The two response columns are alternative implementations, not additive channels; neither is an identified one-quarter consumption mediation effect. Dollars use the 2024 dollars per synthetic household-quarter. Amplification averages pairing-specific ratios.

D Origination-margin sensitivity

Table D1 summarizes how the origination-margin calculation changes under alternative assumptions. Each variation changes one element of the baseline and holds the rest fixed.

Cross-survey dependence. We permute intact CEX records within matching cells to impose rank correlations between the expenditure requirement and liquid assets among initially cardless households, and between the requirement and unused existing capacity among cardholders, each at −0.5, −0.25, 0, 0.25 and 0.5, singly and jointly (35 settings). The ordering of losses per dollar holds for every setting's average; it fails in one individual pairing, at a positive correlation for cardless households combined with a negative correlation for cardholders.

Balance-sheet donors. We redraw the SCF donor for every household 20 times with the same matching algorithm, keeping group membership, CEX pairings and the new-authorization budget fixed.

Receipt selection. Within each initial group, we tilt the probability of receiving a new account toward households with low or high liquid assets, holding each group's expected accounts and authorization fixed.

Income allocation. We reallocate each group's fixed receipt across baseline income

quintiles using relative annual net card additions, or, for cardless households, observed entry into card ownership (Table D6). High-income cardless cells are thinly populated, and we do not estimate income-specific rates for every group. Table D7 reports the variants, alone and combined with the adverse dependence setting above.

Transport and shape. The account-count route replaces the dollar-based number of new accounts with the count implied by observed card ownership and market turnover. The group-amount variation lowers the mean initial limit of cardless recipients while preserving the market mean and median. The shape variations match both the median and mean responses of initial limits, or the mean response only. The availability variation sets $\omega = 0.5$.

Responses away from the constraint. Table D2 reports the calculation of Section 4.4.

Reversal thresholds. Table D3 reports the thresholds of Section 4.5.

The same-dollar comparison. Table D4 recomputes both margins of Table 4 in each alternative population above, resetting the existing-line withdrawal rate so that it removes the same authorization, with and without responses away from the constraint. Table D5 recomputes the baseline comparison under alternative definitions of the expenditure requirement, together with the share of synthetic households that cannot finance the requirement before any policy change.

Table D1: Sensitivity of the origination-margin calculation

Variation	Loss ($)	η_0	η_R	η_N	Order reversals
Baseline	0.882	0.095	0.038	0.015	0 of 50
Cross-survey dependence (35 settings)	0.442–1.326	0.068–0.119	0.018–0.055	0.004–0.028	1 of 1750
Balance-sheet donor reassignment (20 draws)	0.851–0.891	0.093–0.101	0.036–0.039	0.013–0.016	0 of 1000
Receipt tilted to low liquid assets	1.081	0.114	0.045	0.019	0 of 50
Receipt tilted to high liquid assets	0.683	0.077	0.030	0.010	0 of 50
Receipt reweighted by income (combined scores)	0.835	0.090	0.036	0.014	0 of 50
with adverse cross-survey dependence	1.127	0.061	0.053	0.027	6 of 50
Account-count route for new-card volume	0.684				
Cardless receive one-quarter the mean limit	0.800	0.223	0.036	0.014	
Median and mean responses matched jointly	0.711				
Proportional shift by the mean response only	0.268				
Half of each new limit usable within the quarter	0.678				

Expenditure loss in 2024 dollars per synthetic household-quarter from tighter initial limits, averaged over pairings; ranges are across settings or draws. Order reversals count individual pairings in which $\eta_0 > \eta_R > \eta_N$ fails. Cross-survey dependence varies the rank correlation of expenditure requirements with liquid assets (cardless households) and with unused existing capacity (cardholders) between -0.5 and 0.5. Receipt tilts and income reweighting hold each group's expected new accounts and authorization dollars fixed. Blank cells were not computed for that variation. Details are in Appendix D.

Table D2: Responses away from the financing constraint

ϕ	η^{new}	η_0^{new}	η_R^{new}	η_N^{new}	η^{old}	η_0/η_R	η_0/η_N	η^{new}/η^{old}	Order kept
0.00	0.0321	0.0954	0.0375	0.0147	0.0137	2.54	6.49	2.35	50
0.05	0.0805	0.1406	0.0857	0.0640	0.0627	1.64	2.20	1.28	50
0.11	0.1386	0.1949	0.1434	0.1231	0.1215	1.36	1.58	1.14	50

Pre-shock state, same authorization reduction as Table 4. Each household's expenditure loss becomes $\ell_i + \phi(\Delta K_i - \ell_i)$, where ℓ_i is the hard-cap loss and ΔK_i the reduction in usable capacity; $\phi = 0$ is the baseline. ϕ is common to all households, so the group ordering is largely preserved by construction; the informative object is how much the ratios shrink. Order kept counts pairings (of 50) with $\eta_0 > \eta_R > \eta_N$ on the new-limit margin. Because losses are linear in ϕ, the margins can also be given different responses: with slopes of 0.968 (new limits) and 0.981 (existing lines) per unit of ϕ, their losses per dollar are equal when the response to existing-line cuts exceeds that to new limits by 1.9 cents per dollar at a zero response to new limits and by 1.7 cents at 0.11.

Table D3: Conditions that would reverse the ordering of the two margins

Departure from the baseline	Reversal threshold	Range	Evidence
Responses away from the constraint larger for existing-line cuts than for smaller new lines (difference per dollar)	1.9 cents	0.5–3.6 cents (populations)	External estimates concern line expansions (Table 5)
New accounts go less often to households with a shortfall (their receipt relative to the group rate)	43%	35–57% (pairings); 26–71% (populations)	Receipt by income quintile: 95–99%; two-to-one tilt toward high liquid assets: 77%
Part of each measured shortfall is financed by means the surveys do not record (share financed)	62%	48–71% (pairings)	Not measured; loss ratio 1.88 at 25%, 1.32 at 50%, 0.63 at 75%

Pre-shock state; each row changes one assumption and reports the value at which the origination margin's loss per dollar of gross authorization equals that of the existing-line margin. Row 1 adds responses away from the constraint that differ across the margins (Table D2). Row 2 scales the receipt probability of households with a positive shortfall by q and raises that of other households in the same group so that each group's expected accounts and authorization are unchanged; since only households with a shortfall lose from tighter initial limits, the threshold equals the ratio of the existing-line loss to the origination loss, and the construction preserves quarterly receipt and authorization moments but not necessarily the annual entry process. The evidence column reports the origination loss under the income-quintile allocations of Table D7 and the liquid-asset tilts of Table D1 relative to its baseline value, which under this construction equals q. Row 3 adds a share θ of each household's measured shortfall to its resources; the existing-line loss is unchanged because constrained households lose an existing-line cut dollar for dollar. The loss ratio is the origination margin's loss per dollar divided by the existing-line margin's. Ranges are across the 50 pairings or across the alternative populations of Table D4, as indicated.

Table D4: The same-dollar comparison across alternative synthetic populations

		All households		Cardholders only	
ϕ	Variation (number of populations)	η^{new}/η^{old}	< 1	$\eta^{new}_{R,N}/\eta^{old}$	< 1
0.00	Baseline matching (1)	2.35	0	1.84	0
	Cross-survey dependence (35)	1.41–3.87	0	0.91–3.24	6
	Balance-sheet donor reassignment (20)	2.14–2.69	0	1.66–2.07	0
	Receipt tilted by liquid assets (2)	1.82–2.88	0	1.40–2.28	0
0.05	Baseline matching (1)	1.28	0	1.18	0
	Cross-survey dependence (35)	1.08–1.56	0	0.99–1.44	6
	Balance-sheet donor reassignment (20)	1.26–1.31	0	1.15–1.20	0
	Receipt tilted by liquid assets (2)	1.17–1.39	0	1.09–1.27	0
0.11	Baseline matching (1)	1.14	0	1.09	0
	Cross-survey dependence (35)	1.04–1.27	0	1.00–1.21	6
	Balance-sheet donor reassignment (20)	1.13–1.15	0	1.08–1.10	0
	Receipt tilted by liquid assets (2)	1.09–1.19	0	1.05–1.13	0

Pre-shock state; every population removes the same authorization from each margin, with the existing-line withdrawal rate recomputed for each population. Ratios are pairing-mean losses per dollar of gross authorization removed; ranges are across the populations in each row, and the columns headed < 1 count populations with a ratio below one. The cardholders-only ratio compares the origination margin's loss per dollar among revolvers and non-revolvers with the existing-line margin's loss per dollar; the two margins distribute their dollars differently within this group, so it is descriptive rather than a separate same-budget experiment. Across the 2,850 individual pairings of the alternative populations, the lowest all-household ratio is 1.12 with $\phi = 0$, 1.03 with $\phi = 0.05$ and 1.02 with $\phi = 0.11$. Cross-survey dependence varies rank correlations between -0.5 and 0.5 as in Table D1; ϕ is the response away from the financing constraint of Table D2.

Table D5: Alternative definitions of the expenditure requirement

Requirement D_i	Constrained (%) All	Cardless	Loss ($) New	Existing	η^{new}	η^{old}	Ratio	Order kept
Total expenditure (baseline)	12.6	36.2	0.882	0.376	0.0321	0.0137	2.35	50
Excluding insurance, pensions and cash contributions	11.1	33.1	0.747	0.305	0.0272	0.0111	2.45	50
Excluding housing	5.8	17.9	0.398	0.186	0.0145	0.0068	2.13	50
Nondurables, including transportation	4.7	15.0	0.293	0.135	0.0107	0.0049	2.17	50
Nondurables and services, excluding vehicles	3.5	13.3	0.119	0.044	0.0043	0.0016	2.69	49

Pre-shock state, 50 pairings. Constrained is the share of synthetic households whose requirement exceeds income, liquid assets and existing unused card capacity before any policy change. Losses are 2024 dollars per synthetic household-quarter from removing the same authorization on each margin; Ratio is η^{new}/η^{old}; Order kept counts pairings (of 50) with $\eta_0 > \eta_R > \eta_N$ on the new-limit margin. Requirement definitions follow the CEX Interview Survey categories; under the earlier one-percent experiment these definitions give the amplification figures of Table C5.

Table D6: Annual card accumulation and direct ownership entry by baseline income

Income quintile	Net add. (%)	G_0 entry (%)	Entry n	Events	Kish n	Support
Q1	19.33	16.17	798	125	432.2	Adequate
Q2	22.28	27.50	428	117	246.0	Adequate
Q3	22.36	19.42	238	53	150.8	Adequate
Q4	26.99	36.74	94	36	60.7	Low
Q5	25.14	27.78	49	14	32.5	Low

Notes: Pooled adjacent annual transitions, 2021–2024, using baseline respondent weights and baseline-wave weighted income quintiles. Net additions require valid, non-topcoded card counts at both endpoints; a reported count above five is topcoded and is excluded from that outcome. Ownership entry conditions on explicitly no card initially and positive ownership in the next wave; a topcoded positive endpoint counts as entry, and a missing subsequent card count does not exclude explicitly observed ownership. Missing ownership or attrition is not coded as no entry. n counts transitions, not distinct respondents; events count ownership entries. Kish n measures effective weighted support. Low support denotes fewer than 200 transitions, 30 events, or Kish $n < 100$. Tests use respondent-clustered covariance with a finite-cluster correction, conditional on published weights and fixed cutoffs. The joint entry-equality test gives $p = 0.0012$; within-status net-addition tests give $p = 0.0013, 0.7136, 0.0218$ for G_0, G_R, G_N. Holm-adjusted values for these four tests are 0.0046, 0.0046, 0.7136, and 0.0436, respectively. The Q4–Q1 and Q5–Q1 within-status Kitagawa contributions are 8.2743 and 4.8166 percentage points.

Table D7: Income allocation of fixed origination budgets

Scenario	Loss ($)	Change (%)	η_{G_0}	η_{G_R}	η_{G_N}	Reversals (of 50)
A. Baseline cross-survey matching						
Uniform receipt	0.882	0.00	0.095408	0.037531	0.014700	0
Net-addition scores	0.843	-4.45	0.092622	0.035676	0.013909	0
Cardless entry scores	0.867	-1.77	0.089655	0.037531	0.014700	0
Combined scores	0.835	-5.36	0.089655	0.035676	0.013909	0
Cardless entry, two bins	0.874	-0.96	0.092277	0.037531	0.014700	0
Combined, two bins	0.842	-4.55	0.092277	0.035676	0.013909	0
Net-addition scores, $\theta = 1.5$	0.824	-6.63	0.091154	0.034769	0.013534	0
B. Joint adverse cross-survey matching						
Uniform receipt	1.185	34.28	0.067507	0.054852	0.028229	1
Combined scores	1.127	27.79	0.061011	0.052856	0.026961	6

Notes: Quarterly modeled losses in dollars per synthetic household, averaged over all 50 cross-survey pairings. Relative changes in both panels use the uniform baseline in Panel A as the denominator. η_g is pairing-mean expenditure loss divided by pairing-mean reduced gross authorization dollars; it is not a group loss share. The last column counts individual pairings with $\eta_{G_0} < \eta_{G_R}$, out of 50. The two-bin variants pool Q1–Q2 and Q3–Q5 for cardless entry. The $\theta = 1.5$ variant raises the net-addition scores to that power before group normalization; missing income retains a neutral raw score. Panel B applies the adverse joint setting with G_0 expenditure–liquidity dependence +0.50 and cardholder expenditure–unused-credit dependence −0.50. A variant with income-specific entry for every group is not reported owing to insufficient cell support: high-income G_0 cells fail the prespecified effective-sample requirements. Expected new accounts, gross authorization dollars, and authorization cuts are fixed within each initial card-status group.