Mental Accounting for Credit: Evidence from a Large-Scale Credit-Limit Experiment*

Sean Higgins† Raymond Kluender‡

November 2026

Abstract

We develop a model of mental accounting for credit in which consumers budget according to non-fungible spending categories that govern their borrowing behavior. We test our model in a randomized controlled trial with Cashea, Venezuela's largest consumer credit provider. Cashea provides two distinct credit lines to all of its borrowers: one for grocery stores, pharmacies, and restaurants, and one for all other retail merchants. We randomized 0%, 50%, or 100% credit-limit increases on one or both credit lines for 209,996 active users. For both credit lines, we estimate large same-category monthly marginal propensities to consume (MPCs) of 8–10 cents per $1 of additional credit, with near-zero cross-category MPCs, consistent with mental accounting. We present five additional supporting results: (1) the results persist in every post-treatment month; (2) the same-category MPCs represent an increase in total consumption; (3) the results hold separately for both constrained and unconstrained borrowers; (4) we fail to identify any subgroup that treats the relaxation of the budget constraint fungibly; and (5) the results hold separately for each sub-category of consumption within each credit line. We conclude that consumers form mental accounts for credit, with economically large effects on borrowing and the allocation of spending across categories. This has important implications for the regulation of financial technology and the effects of monetary policy.

*We are grateful to Gonzalo Aguilera, Malena Amendola, Betsy Hernández Hernández, Jaime Toribio Arrieta, Cynthia Aiello, Arnoldo Gabaldón, Stefany Garcia, Paula Mogollon Mejia, Adrian Mota, Gal Oliveira, Pedro Vallenilla, and many others at Cashea for their collaboration on this research. For helpful comments and feedback, we thank John Beshears, Dean Karlan, David Laibson, Matthew Rabin, Stephan Siegel, and seminar participants at the Advances with Field Experiments Conference, CEAR-RSI Household Finance Workshop, Columbia New Empirical Finance Conference, Dartmouth College and Tuck School of Business, Harvard Behavioral and Experimental Economics Seminar, Ohio State University, Philadelphia Federal Reserve, Shanghai University of Finance and Economics, University of Michigan, Western Michigan University, Tsinghua University, and WEFIDEV. We thank Yixiang Gao, Eleanor Jenke, and Noah Jenkins for research assistance, and Irving Birkner for support on the data sharing agreement. IRB: Harvard IRB25-0237; Northwestern STU00224570. Pre-registration: AEARCTR-0016361.

†Northwestern University, Kellogg School of Management and NBER, sean.higgins@kellogg.northwestern.edu

‡Harvard Business School and NBER, rkluender@hbs.edu

1 Introduction

Consumer credit finances a large and growing share of global consumption. Consumer borrowing has important macroeconomic consequences: excessive borrowing can cause recessions (Mian, Sufi, and Verner, 2017) and financial crises (Müller and Verner, 2024), while the marginal propensity to borrow out of credit-limit increases mediates the effects of expansionary monetary policy (Agarwal, Chomsisengphet, Mahoney, and Stroebel, 2018).

Yet we know little about how consumers decide how much to borrow and which purchases to finance with debt. Mental accounting—or the idea that consumers budget according to non-fungible spending categories—has long been recognized as a driving force behind consumption and savings decisions (Thaler, 1985, 1990, 1999), but there is limited work on whether mental accounting shapes borrowing.

The ideal experiment to test for mental accounting for credit requires three elements. First, it requires observing which spending categories comprise the mental account associated with a credit line. Studies of mental accounting for income rely on restricted-use income such as food vouchers (Chorniy, Finkelstein, and Notowidigdo, 2025; Hastings and Shapiro, 2018); with general-purpose credit, the corresponding accounts are difficult to observe and potentially heterogeneous across consumers. Second, even when mental accounts for credit are observable through category-specific credit lines, testing fungibility requires exogenous variation in category-specific credit limits. Without mental accounting, a credit-limit increase can raise consumption through liquidity constraints (Gross and Souleles, 2002; Zeldes, 1989) or precautionary motives (Aydin and Kim, 2024). But unless consumers have exhausted the credit line whose limit is increased and have no cash to reallocate, they should respond to an increased credit limit—even for category-specific credit—by increasing consumption in all categories to equate the marginal utility of consumption across categories. Mental accounting instead predicts a consumption increase only for the spending categories associated with that line. Third, the ideal experiment requires measures of spending both on and off the credit line to distinguish changes in consumption from payment-method substitution.

We run this experiment by randomizing category-specific credit limits for 209,996 consumers in partnership with Cashea, the largest consumer lender in Venezuela. Cashea has more than 10 million users and an estimated 8% of Venezuela's GDP runs through their platform. Every Cashea user has access to two independent, dollar-indexed credit lines: one accepted only at grocery stores, pharmacies, and restaurants ("grocery line"), and the other for clothing, home goods, technology, and other retail ("retail line"). We randomly assigned credit-limit increases of 0%,

50%, or 100% on the grocery line, the retail line, or both lines to 209,996 active users. We estimate category-specific marginal propensities to consume out of these credit-limit increases using 12 months of administrative data supplemented with baseline and follow-up surveys which measure total spending in each category on and off the credit lines. We received 28,511 responses to the follow-up surveys.

We develop a model of mental accounting for credit that generates testable predictions for the effects of these category-specific credit-limit increases with and without mental accounting. We specify a two-category consumption-savings model with separate credit lines for each category (indexed by ℓ), mirroring our institutional setting. Future income is subject to exogenous shocks, so consumers' credit limits may bind in the future if they experience a negative shock; we refer to this as a *dynamic constraint*. A credit-limit increase on line ℓ relaxes this future constraint and can raise current consumption even for consumers who are not *statically constrained* (i.e., consumers who could already spend more on category ℓ through unused credit on their category-ℓ credit line or by reallocating cash).

Consumers form mental accounts around each category, which we model following Farhi and Gabaix (2019) as a category-specific consumption target and a psychological cost of deviating from the target. We develop a novel microfoundation for how these targets respond to changes in the category-specific credit limits. We treat a category-ℓ credit-limit increase as a salient cue about category-ℓ consumption (*à la* Bordalo, Gennaioli, and Shleifer, 2020). This cue prompts the consumer to re-optimize, but because re-optimizing the full consumption portfolio is difficult, the consumer solves a simpler single-good problem to set the new consumption target, holding fixed their target in the other category—in the spirit of similarity-based attention (Bordalo, Gennaioli, Lanzani, and Shleifer, 2026).

With mental accounting, consumers face a high psychological cost of deviating from their targets. The model therefore predicts that they either do not change consumption (if they are not dynamically constrained) or increase consumption only in the category whose credit limit increased, holding fixed consumption in the other category. Without mental accounting, deviating from the targets is costless and money is fungible: although credit on line ℓ can only be spent on category ℓ, it reduces the precautionary motive and frees up cash that can be spent on either category. Consumers therefore either do not change consumption (if they are not dynamically constrained) or increase consumption in both categories to equate marginal utilities across them.

We find overwhelming evidence of mental accounting for credit. For both categories, category-specific credit-limit increases cause large same-category consumption responses but near-zero spillover effects to other-category spending. In our central result, we estimate that a \$1 increase

in the grocery-line credit limit increases monthly grocery-line purchases by 9.8 cents, a $1 increase in the retail-line credit limit increases monthly retail-line purchases by 8.1 cents, and spillover effects to the other category are, in each case, zero or even slightly negative: -1.1 cents for the grocery-line-to-retail spillover and a not statistically significant point estimate of 0.2 cents for the retail-line-to-grocery spillover. Across treatment arms, spending increases roughly linearly with the size of the credit-limit increase. Increasing both credit limits raises spending on each line by roughly the same amount as increasing that line's limit alone. Because our experiment induced large increases in credit limits (50% or 100%), these MPCs imply meaningful changes in consumption. For example, doubling the retail-line limit raised monthly retail consumption by 5.9% of median monthly income.

We provide five additional pieces of evidence in support of mental accounting for credit.

First, both the same- and cross-category MPCs are remarkably persistent over time—consistent with consumers permanently updating their same-category consumption targets and neglecting to fully re-optimize due to mental accounting. The same-category MPCs in the first month of the experiment are close to the 12-month averages (0.079 vs. 0.098 for the grocery line and 0.064 vs. 0.081 for the retail line), which shows that the same-category effects are immediate and persistent. In the first month of the experiment, the point estimates of the grocery-line-to-retail spillover and retail-line-to-grocery spillover are 0.009 and 0.013, respectively. Both cross-category MPCs remain near zero and decline over time. Our central results are reproduced in every single one of the 12 post-treatment months.

Second, the same-category responses represent increases in total consumption rather than payment-method substitution. If consumers simply shifted purchases they would otherwise have made in cash onto the expanded credit line, spending on the line would rise with no change in total consumption. We test for payment-method substitution using follow-up surveys that measure spending in each category both on and off Cashea. Where we observe both administrative and survey outcomes, the survey estimates replicate the results from the administrative data: same-category on-Cashea MPCs are 0.065 for the grocery line and 0.050 for the retail line in the survey data, compared with 0.084 and 0.062 in the administrative data over the same period. Turning to spending outside Cashea, we find no evidence of payment-method substitution: for both the grocery and retail lines, we estimate statistically non-significant, near-zero effects of the credit-limit increases on both same- and cross-category spending. We therefore interpret our results throughout as increases in overall consumption.

Third, binding constraints cannot explain our results. Without mental accounting, a credit-limit increase can raise same-category spending with no cross-category spillover only if borrow-

ers are statically constrained (i.e., they have already exhausted the line whose limit is increased and have no cash to reallocate toward that category). Virtually none of our participants are statically constrained. Median baseline utilization is 16.8% on the grocery line and 12.0% on the retail line, 74.0% of users typically spend cash on both categories, and 79.6% have liquid savings. Nevertheless, given the possibility that static constraints could in theory explain a same-category increase with zero cross-category spillover, we prespecified heterogeneity tests by tercile of credit utilization at baseline. As expected, the same-category MPC is highest for the most-constrained tercile. But the cross-category spillover remains near zero in all three terciles across both lines, and we can rule out cross-category MPCs greater than 0.014. Consumers in the bottom two terciles are far from their credit limits, yet they show strong, statistically significant same-category responses and near-zero spillovers. This pattern, which holds in every tercile and is symmetric across categories despite different baseline limits, repayment terms, and merchant types, provides compelling evidence of mental accounting. In an additional prespecified test using self-reported credit constraints among baseline survey respondents, we again see large same-category and zero cross-category MPCs even among users who report they are unconstrained.

Fourth, we find no subgroups with meaningful cross-category spillovers in response to an increase in one line's credit limit, which contradicts standard models in which money is fungible. We use machine learning and all baseline characteristics available in the administrative data to test for heterogeneity in treatment effects separately for all four same- and cross-category responses (Chernozhukov, Demirer, Duflo, and Fernández-Val, 2025). We find substantial heterogeneity in same-category responses. However, we fail to identify any group with a meaningful positive spillover for either cross-category response: for eight out of 10 groups (five quintiles for each line), we fail to reject zero spillover, for one we find a negative spillover, and for one a positive spillover where the relative increase in cross-category consumption is less than 20% as large as the same-category response. We conclude that any subgroups that we can identify using observable baseline characteristics form mental accounts.

Finally, we test the assumption that the two credit lines represent mental accounts in the spirit of Thaler (1985, 1999) and the two categories in our model. Each line bundles several different types of merchants. The grocery line includes grocery stores, pharmacies, and restaurants, while the retail line includes clothing, home goods, services, technology, tools, health, and sports. We estimate MPCs separately for each of these 10 merchant sub-categories for both same-category and cross-category credit-limit increases, and normalize the estimates by each merchant sub-category's baseline share of spending on its line. We estimate positive and statistically significant same-category MPCs for *every single merchant sub-category* out of an own-line credit-limit increase. The

cross-category MPCs are near zero for all 10 of the merchant sub-categories, with no normalized cross-category MPC exceeding 0.013. The same-category responses appear even for the most divisible sub-categories on each line, such as groceries and clothing, indicating that indivisible durable purchases are not the central driver of our results. These results provide evidence that users create mental accounts around the categories Cashea assigns to each credit line, consistent with evidence from Read, Loewenstein, and Rabin (1999) that consumers accept the brackets imposed upon them.

Our credit-limit increases induced large changes in the level and allocation of consumption. Doubling the retail-line credit limit increased retail consumption by $29 per month (5.9% of median monthly income) with essentially no change in grocery consumption, and doubling the grocery-line limit had a mirror-image effect. Given that these are short-maturity loans—consumers cannot delay repayment far into the future or pull consumption forward indefinitely—the money for the additional consumption must come from somewhere. In our second follow-up survey, we measured spending on categories not covered by either credit line, the stock and flow of savings, other debt, income, and peer transfers. None changed significantly for the treated groups, and we cannot isolate any universal source of funds for the additional spending. We can, however, rule out meaningful average reductions in savings or increases in financial distress. Effects on the probability of a user being at least 90 days delinquent are null or small, ranging from 0.5 (4.2%) to 1.7 percentage points (13.0%) across arms, and the credit-limit increases were profitable for Cashea in all arms. Just 13.1% of treated participants report wishing they had spent less through Cashea.

Mental accounting for credit has important policy implications. First, there is a heated debate over the rapid growth of technology-enabled, merchant-specific "buy now, pay later" (BNPL) loans, with particular concern over the provision of credit for basic necessities like groceries, rent, and utilities.[1] Our results suggest that expanding category-specific credit induces rebudgeting and higher spending only within the expanded categories. Further, this higher spending is driven by mental accounting rather than binding constraints. The welfare implications of expanding category-specific credit (e.g., BNPL) thus depend on the marginal utility of increased spending *only* on the categories of goods that can be purchased with that credit line—contrary to the logic of standard models, which overlook this non-fungibility.

Second, mental accounting can dampen the effect of monetary policy pass-through. Marginal propensities to borrow mediate effects of expansionary monetary policy (Agarwal, Chomsisengphet, Mahoney, and Stroebel, 2018), but shifts in the yield curve in response to monetary policy

[1]See, for example, The New York Times (August 2026), "'Buy Now, Pay Later' Lenders Pitch Loans for Needs Like Electricity and Rent."

are sometimes non-parallel, and changes to rates on different parts of the yield curve affect the supply of different credit products. Short-term rates pass through to credit-card limits (Agarwal, Chomsisengphet, Mahoney, and Stroebel, 2018; Chava, Ganduri, Paradkar, and Zeng, 2023), while longer-term rates pass through to mortgage loan amounts (Bosshardt, Di Maggio, Kakhbod, and Kermani, 2024). Without mental accounting, home buyers should respond to an increase in mortgage limits by increasing consumption of both housing and groceries. In contrast, with mental accounting, higher mortgage limits may only increase housing consumption. Furthermore, because utility is concave in the consumption of each good, the optimal increase in overall consumption in response to an increase in mortgage limits is smaller when mental accounting creates non-fungibility across categories.

Related Literature. Our finding—that mental accounting governs consumer borrowing decisions—brings new evidence to three literatures.

First, we contribute to the literature on borrowing and consumption responses to credit supply. Experimental and quasi-experimental studies consistently find that consumers—including those who are far from their limits and hold liquid assets—borrow significantly out of credit-limit increases (Agarwal, Chomsisengphet, Mahoney, and Stroebel, 2018; Aydin, 2022; Garber, Mian, Ponticelli, and Sufi, 2024; Gross and Souleles, 2002; Gross, Notowidigdo, and Wang, 2020). The large responses of unconstrained consumers are a puzzle for models in which credit matters only through liquidity constraints, and the literature has advanced two explanations: a precautionary motive, under which consumers spend out of expansions to their buffer of unused credit (Aydin, 2022; Aydin and Kim, 2024; D'Acunto, Rauter, Scheuch, and Weber, 2020), and signaling, under which consumers infer higher future income from the lender's decision to extend credit (Yin, 2026). These studies examine general-purpose credit and do not observe which goods cash is spent on. We randomize category-specific credit-limit increases and measure spending on and off the platform, which allows us to estimate how these expansions change the allocation of consumption across categories. Because models without mental accounting predict spillovers across categories while a model with mental accounting does not, the cross-category responses in our experiment directly test whether consumers create mental accounts around each of their credit lines.

Second, we contribute to the literature on mental accounting (Kőszegi and Matějka, 2020; Lian, 2021; Rabin and Weizsäcker, 2009; Thaler, 1985, 1990, 1999; Tversky and Kahneman, 1981). Empirical evidence shows that consumers budget by spending category (Zhang, Sussman, Wang-Ly, and Lyu, 2022) and violate fungibility of consumption across categories in response to restricted-

use income, price changes, and small windfalls (Chorniy, Finkelstein, and Notowidigdo, 2025; Hastings and Shapiro, 2013, 2018; Milkman and Beshears, 2009). Mental accounting also affects consumption responses to tax refunds versus tax payments due (Baugh, Ben-David, Park, and Parker, 2021) and receipt of a new payment card (Boehm, Fize, and Jaravel, 2025; Gelman and Roussanov, 2024), as well as portfolio choice (Choi, Laibson, and Madrian, 2009). In credit markets, mental accounting has been shown to affect debt repayment (Katz, 2023; Katz, Russel, and Shi, 2026; Ponce, Seira, and Zamarripa, 2017), the choice of loan terms (Argyle, Nadauld, and Palmer, 2020), and borrowing while coholding savings (Sussman and O'Brien, 2016). We build on this work with an experimental test of whether mental accounting for credit affects the allocation of consumption and borrowing across categories: we randomize the credit limits on category-specific credit lines and find that consumers treat each line as a budget for its category, declining to re-optimize across categories even when statically unconstrained.

Third, we contribute to a growing literature on fintech lending (Berg, Fuster, and Puri, 2022) and BNPL credit in particular. For consumers, BNPL access increases total spending and shifts its composition toward retail (Bian, Cong, and Ji, 2023; Di Maggio, Katz, and Williams, 2022). For merchants, offering BNPL increases sales (Alok, Ghosh, Kalda, and Mitra, 2026; Berg, Burg, Keil, and Puri, 2025). Recently, BNPL has drawn regulatory scrutiny over concerns about overborrowing (deHaan, Kim, Lourie, and Zhu, 2024; Guttman-Kenney, Firth, and Gathergood, 2023), though BNPL repayment histories can also improve lenders' screening (Laudenbach, Molin, Roszbach, and Sondershaus, 2025). This evidence is mostly quasi-experimental and comes from markets where BNPL remains a small share of spending. We conduct a large-scale randomized evaluation of BNPL credit limits with the dominant consumer lender in a market where BNPL intermediates an estimated 8% of the country's entire GDP. We show that the defining feature of BNPL credit—its restriction to particular merchants and categories—shapes the allocation of consumption, not because consumers are credit-constrained but because they form mental accounts around their credit lines.

2 Cashea and Consumer Credit in Venezuela

We partner with Cashea, the largest consumer lender in Venezuela. Historically, Venezuela had a healthy consumer lending sector anchored by bank-issued credit cards. About 45% of Venezuelan adults held a credit card in 2011, card receivables totaled $7.8 billion at the official exchange rate in early 2012, and 90-day delinquency rates at the largest issuers were about 1%.[2] Over the next

[2]Electronic Payments International (April 2012), "Venezuela: held back by uncertainty."

decade, hyperinflation, reserve requirements, interest-rate caps, and requirements that consumer loans be extended in rapidly depreciating *bolívares* effectively ended consumer lending. By late 2022, total card credit outstanding had plummeted to roughly $16 million nationwide—just 0.2% of its value a decade earlier.[3]

Cashea entered this void in 2022 by launching dollar-indexed, interest-free installment loans through merchant partners at the point of sale. As of July 2026, Cashea had more than 10 million users and 40,000 partner stores spanning grocery stores, pharmacies, retail outlets, and service providers.[4] The Cashea app has been installed on more than half of smartphones in the country, processes transactions equivalent to roughly 8% of Venezuela's GDP, and is growing rapidly. Cashea's consumer lending balance sheet is larger than that of the entire traditional banking system. Consumers largely lack access to any credit outside of Cashea: in our survey data, only 26.4% of users had recently used any other form of credit, including personal loans, credit cards, store credit, other installment purchases, or loans from a private lender, family member, or friend. Cashea's position as the predominant consumer lender in Venezuela allows us to observe an unusually complete picture of consumer borrowing and repayment behavior.

Cashea provides a standard installment-loan contract typical of large international buy now, pay later (BNPL) lenders such as Klarna, Affirm, and Afterpay. To enroll, a consumer downloads the Cashea app and creates an account. Cashea's proprietary credit-scoring model uses alternative data commonly used by fintech lenders (Berg, Burg, Gombović, and Puri, 2020) to assign an initial limit on each of two credit lines, with only 2% of applicants rejected. Credit limits and other loan terms change when a user moves up or down a "level" in Cashea's levels system, where all users start at level one and move through the six levels based on repayment history and total purchase amounts. To make a purchase in store or online at a partner merchant, the consumer enters the purchase through the Cashea app to use their available financing and makes the down payment directly to the merchant at the point of sale.

The merchant finances the loan and Cashea underwrites the risk, guaranteeing the merchant against nonpayment.[5] After making the down payment at purchase, the consumer pays the remaining balance in installments at 14-day intervals, typically over 14 to 42 days. Merchants pay Cashea a small fee as a percentage of each transaction (similar to debit and credit cards in the US); 46.1% of Cashea users we surveyed report that merchants pass the fee through to the purchase price. Late payments trigger a $4 late fee, freeze both credit lines, and reduce future credit ac-

[3]Reuters (February 2023), "Venezuelans say credit cards that were once lifeline now 'useless'"; Rodner, Martínez & Asociados, "Venezuela: Lending & Secured Finance Laws and Regulations 2026," *International Comparative Legal Guide*.
[4]PR Newswire (July 2026), "Venezuela's Cashea Raises $100 Million From Global Investors".
[5]For a subset of merchants, Cashea also provides factoring, paying the merchant the full purchase amount up front and collecting the installments directly; for these transactions Cashea, rather than the merchant, finances the loan.

cess even after the borrower becomes current, by lowering the user's level and associated limit. Conversely, as users accumulate purchases and repay on time, their level increases and they earn higher credit limits, lower required down payments, and—at participating merchants—longer maturities.

Venezuela is functionally dollarized: the dollar has broadly become the *de facto* unit of account for wages, invoices, and most goods.[6] Cashea's loan amounts, installment payments, and credit limits are dollar-indexed. As a result, inflation in the local currency does not mechanically erode the loan's principal or the borrower's access to credit. Store prices are also typically posted in dollars (or in dollar-indexed local currency), which limits the influence of local-currency inflation on intertemporal consumption decisions.[7] Finally, wages are typically dollar-indexed: in our survey of Cashea users, 87.6% reported they were paid in dollar-indexed salaries. This prevents erosion in the borrower's ability to repay over the repayment horizon.

Cashea provides the ideal setting to study mental accounting for credit because each user receives two independent and non-fungible credit lines. The *Línea Cotidiana* (directly translated as "Everyday Line") can be spent at grocery stores, pharmacies, and restaurants, while the *Línea Principal* (directly translated as "Main Line") covers home goods, clothing, services, technology, tools, health, and sports. We refer to these two credit lines as the "grocery line" and "retail line" to streamline exposition. Appendix Table A1 reports baseline spending shares (for number of transactions and dollars) by store sub-category within each line: 75% of grocery-line purchases are at grocery stores, 14% at pharmacies, and 9% at restaurants, while 46% of retail-line purchases are on clothing, 25% on home goods, 9% on services, and the rest dispersed across other retail categories.

The grocery and retail lines offer different loan terms to mitigate risk and because grocery stores and retail stores have different margins. Purchases on the grocery line are paid with one down payment and a single installment payment 14 days later, with two installments over 28 days available at select merchants for higher-level users. The retail line follows the "pay in four" structure typical of BNPL loans (**guttman-kenney˙2026**): a down payment followed by three installments at 14-day intervals, with up to 12 installments available at select merchants for higher-level users. Grocery-line loans have shorter maturities than retail-line loans, in part because merchants finance the credit and grocery stores' thinner margins leave them less able to absorb delayed cash flows. Credit limits on the grocery line are also lower than on the retail line. The median baseline limits are $146 for the grocery line and $366 for the retail line, relative to a median monthly income

[6] Embassy of Switzerland in Venezuela (2025), "Economic Report 2024/2025 Venezuela."
[7] Fortune (August 2026), "Venezuela abandoning the bolivar and adopting the U.S. dollar would be the biggest currency switch since the advent of the euro, Hanke says."

of $480 based on our baseline survey of Cashea users. Available credit on the grocery line cannot be used at retail-line merchants, and vice versa. This restriction is enforced by the platform at the point of sale based on the type of merchant. In our model, we posit that consumers may form mental accounts around these two spending categories, which we then test in our experiment.

3 Model

We present a stylized model to generate predictions that distinguish mental accounting from theories in which money is fungible. We first set up a two-good consumption–savings problem with category-specific credit lines, reflecting the institutional features of our setting. We then introduce mental accounting following Farhi and Gabaix (2019), in which consumers face a psychological cost of deviating from category-specific consumption targets. Finally, we microfound how consumers update their consumption targets: a credit-limit increase is a salient cue that leads consumers to re-optimize their consumption target only within the mental account that they associate with that credit line, because re-optimizing their full consumption portfolio is costly.

3.1 Model Setup

We start with a three-period model with two consumption categories $\ell \in \{G, R\}$, where G represents consumption at grocery stores, pharmacies, and restaurants, and R represents consumption at other retail merchants. The agent has separate credit lines for each category with credit limits L_ℓ. The income process y_t is designed to capture two features: income uncertainty, which we assign to period 2, and a general increase in income over the lifecycle, which we capture by assuming $y_3 > y_1$. Specifically, in period 2, there is probability p of a good state in which there is no income shock and $y_2 = y_1$, and probability $1 - p$ of a bad state with a negative income shock of size δ such that $y_2 = y_1 - \delta$.

Each period, the agent chooses how much to consume of each category ℓ as a function of the state s (which we omit for notational convenience), both on and off their credit lines. We denote total consumption on category ℓ at time t as $c_{\ell t}$, and on-credit-line spending as $a_{\ell t}$; thus cash spending in category ℓ at time t is $c_{\ell t} - a_{\ell t}$. We assume that borrowing in period t must be repaid with interest rate r in period $t + 1$. Any resources not consumed in period t are carried over to the next period as a savings balance x_{t+1}, and we assume the agent starts the first period with no savings, $x_1 = 0$. For convenience, we define m_t as the net resources available in period t, which is comprised of current-period income and the savings balance carried over from the prior period, net of any repayments that must be made for debt incurred in the prior period:

$m_t \equiv y_t + x_t - (1+r)\left(a_{G,t-1} + a_{R,t-1}\right)$. In period 1, there was no prior period in which to save or borrow and this collapses to $m_1 = y_1$. Thus, the law of motion for savings carried to periods 2 and 3 is $x_{t+1} = m_t - [(c_{Gt} - a_{Gt}) + (c_{Rt} - a_{Rt})]$.

The agent solves the following optimization problem:

$$\max_{\substack{\{c_{Gt}\},\,\{c_{Rt}\},\\ \{a_{Gt}\},\,\{a_{Rt}\}}} \quad \mathbb{E}_{t=1}\left[\sum_{t=1}^{3} \beta^{t-1} U(c_{Gt}, c_{Rt})\right] \tag{1}$$

$$\text{subject to} \quad (c_{Gt} - a_{Gt}) + (c_{Rt} - a_{Rt}) \leq m_t \;\; \forall t, \tag{2}$$

$$0 \leq a_{\ell t} \leq \min\left\{L_\ell,\, c_{\ell t}\right\} \text{ for } t \in \{1,2\} \text{ and } \ell \in \{G,R\}, \tag{3}$$

$$a_{\ell 3} = 0 \text{ for } \ell \in \{G,R\}, \tag{4}$$

where $U(c_{Gt}, c_{Rt})$ is the instantaneous utility in period t and β is the discount factor. Constraint (2) says that cash consumption cannot exceed the available resources. Given the definition of m_t, this constraint precludes access to any other form of credit, and implies that $x_{t+1} \geq 0$. Constraint (3) caps borrowing on each line at that line's credit limit, and establishes that credit on one line cannot finance consumption of the other good. Finally, constraint (4) prohibits borrowing in the model's terminal period, since there is no future period in which to repay.

3.2 Mental Accounting

We model mental accounting following Farhi and Gabaix (2019), which is a formulation also used in applications by Hastings and Shapiro (2018) and Chorniy, Finkelstein, and Notowidigdo (2025). The consumer has a consumption target $c^*_{\ell t}$ for each category and incurs a psychological cost $\gamma_\ell \geq 0$ of deviating from it. Instantaneous utility is

$$U(c_{Gt}, c_{Rt}) = u(c_{Gt}, c_{Rt}) - \gamma_G \left|c_{Gt} - c^*_{Gt}\right| - \gamma_R \left|c_{Rt} - c^*_{Rt}\right|. \tag{5}$$

We assume that $u(\cdot)$ is homothetic, increasing in consumption of both categories, and concave in consumption of both categories. When $\gamma_G = \gamma_R = 0$, the model collapses to a model without mental accounting; when γ_ℓ is large, the targets effectively determine the allocation of consumption across categories, as deviating from these targets is costly.

As in Farhi and Gabaix (2019), the formulation in equation (5) is agnostic about the source of the category-specific consumption targets. How these targets evolve is critical in our setting, because the experiment changes category-specific credit limits and the model's predictions depend on how the targets respond. Thus, we microfound the evolution of the $c^*_{\ell t}$ targets as follows.

A credit-limit increase on the category-ℓ line is a salient cue about category-ℓ consumption (*à la* Bordalo, Gennaioli, and Shleifer, 2020). The cue prompts the consumer to re-optimize, but re-optimizing the full consumption portfolio in equation (1) subject to constraints (2)–(4) is costly. Consumers instead take a shortcut to set their new consumption targets: they solve the simple "local" single-good problem in the category associated with the cue, holding fixed their target in the other category, in the spirit of similarity-based attention (Bordalo, Gennaioli, Lanzani, and Shleifer, 2026).

Once consumers have determined their new consumption target for the category whose credit limit increased, their consumption response depends on how strongly they form mental accounts (i.e., the size of γ_G and γ_R). For high-γ consumers (those with strong mental accounts), the new consumption targets determine consumption, since deviating from them is costly. For low-γ consumers, deviating is not costly, so they re-optimize consumption across both categories even though only the category-ℓ target changed. Thus, our microfoundation also provides an interpretation of γ as the cost of full rebudgeting—i.e., the mental cost of re-optimizing the entire consumption portfolio rather than only the account made salient by the category-specific cue. Appendix B formalizes the thresholds $\bar{\gamma}_G$ and $\bar{\gamma}_R$ that separate the two groups.

3.3 Model Predictions

We now turn to our model's predictions for how consumption across the two categories responds to an increase in one category's credit limit. In the predictions, "current consumption" refers to consumption in period 1, which can be thought of as the period during which our experiment is conducted, while period 2 refers to a future year with a potential income shock and period 3 refers to later on in the lifecycle.

Prediction 1. ***Without*** *mental accounting ($\gamma_G = \gamma_R = 0$), an increase in the category-ℓ credit limit will either not affect current consumption in* ***either*** *category or increase current consumption in* ***both*** *categories.*

Whether overall period-1 consumption will increase in response to a credit-line increase depends on whether the consumer is dynamically constrained in the bad state. If, in the counterfactual world with a lower credit limit on line ℓ, they would optimally consume more in period 2 than they are able to when they experience the income shock δ, the credit-limit increase will ease this constraint, reduce the precautionary motive, and lead them to increase period-1 consumption compared to the counterfactual. Thus, our model features standard dynamic liquidity constraints due to uninsured income risk—as in Aiyagari (1994) and consistent with empirical evidence in D'Acunto, Rauter, Scheuch, and Weber (2020), Aydin (2022), and Aydin and Kim (2024). Survey

evidence also supports the relevance of dynamic constraints in our setting: 51.1% of borrowers said that the reason they did not fully use one of their credit lines was that they wanted to leave some available credit limit in case of a shock.

Importantly, without mental accounting, a borrower equates marginal utilities of consumption across categories in period 1, even when dynamic constraints bind in the bad state in period 2. As a result, a credit-limit increase does not raise period-1 consumption in only one category. The one exception is a borrower who is *statically* constrained. Being statically constrained requires two conditions to hold in period 1 in the absence of a credit-limit increase: first, the borrower would fully use the category-ℓ credit line, and second, they would not have any cash available to reallocate towards category-ℓ consumption. The second requirement means that they would neither save nor spend cash on other categories. For example, an agent who uses all of their grocery-line credit but also saves or spends cash on clothing is not statically constrained: their marginal utilities of consumption across categories must be equal—otherwise they would reallocate some of their savings or cash spending on clothing to grocery consumption. This static constraints case is the only scenario in which marginal utilities of consumption would differ across categories in the absence of the credit-line increase. Without mental accounting, this difference in marginal utilities is necessary for a consumer to respond to a credit-limit increase by raising consumption in only one category.

In our setting, virutally no borrowers are statically constrained: the median baseline credit-line utilization ratios are 16.8% on the grocery line, 12.0% on the retail line, and 16.1% combined. Most Cashea users (74.0%) also typically spend cash on both categories, and 79.6% have liquid savings.

Prediction 2. ***With*** *sufficient mental accounting ($\gamma_G > \bar{\gamma}_G > 0$ and $\gamma_R > \bar{\gamma}_R > 0$), an increase in the category-ℓ credit limit will either not affect current consumption in* ***either*** *category or lead to an increase in current consumption* ***only*** *in category ℓ with no spillover to the other category.*

As in the no-mental-accounting case, consumption will not respond in either category if the agent is not dynamically constrained, meaning that even in the bad state their credit limits would not bind in period 2 in the counterfactual without a credit-limit increase. The contrast occurs when an agent is dynamically constrained: in this case, a high γ would lead the agent to increase consumption *only* on the category with a credit-limit increase. This occurs because the credit-limit increase on line ℓ is a salient cue about category-ℓ consumption that leads the agent to re-optimize their consumption target on category ℓ while keeping their other consumption target fixed, since this is much easier than solving the full re-optimization problem. The large γ indicative of mental accounting then makes it costly for agents to deviate from their targets, so they do not.

Our setting is close to the ideal environment to test for mental accounting. Because our borrowers are not statically constrained, the cross-category response to a category-specific credit-limit increase distinguishes the two predictions. Without mental accounting, any consumption response should spill over to the other category. With mental accounting, it should not. We next describe the experiment we designed to measure these responses.

4 Experimental Design and Data

4.1 Experimental Design

The experiment sample includes 209,996 active users. We defined active users as those who made at least one purchase on each credit line in the prior 40 days and whose accounts were at least 62 days old as of August 1, 2025 (the date we drew the sample). The activity requirement ensures that participants were aware of both of their credit lines and were willing to borrow for both consumption categories. We preregistered the experimental design, primary and secondary outcomes, and heterogeneity analyses in a pre-analysis plan filed with the AEA RCT Registry on September 30, 2025 (AEARCTR-0016361). The experiment launched and Cashea implemented the credit-limit changes on October 3, 2025.

We randomly assigned users to one of six different treatment arms or the control group with equal probability: specifically, we randomized 50% or 100% increases on one or both credit lines, while limits for users assigned to the control group were unchanged.[8] Appendix Table A2 reports summary statistics by treatment arm and tests for balance. Joint tests do not reject the null of balance across all treatment arms for any variable, nor across all variables jointly for any individual treatment arm compared to the control group.

Figure 1 and Appendix Table A3 summarize the experiment's implementation: the figure plots average credit limits by credit line and treatment arm before and after launch, while the table also reports the number of users, baseline limits, and observed increases. Appendix Figure A1 shows that even 12 months after the experiment launched, despite some natural changes in limits due to users moving up or down levels, our experimental credit-limit changes swamp the naturally occurring credit-limit changes from users' changing levels.

To facilitate our prespecified heterogeneity analysis by baseline credit access and utilization, we stratified the randomization by terciles of (1) the total credit limit combining both lines, (2) the average utilization ratio (balance divided by limit) over the last 30 days for the grocery line,

[8] Our experiment did not directly affect users' levels (described in Section 2), and natural credit-limit changes due to changes in levels continued to occur according to Cashea's rules in both the treatment and control groups.

and (3) the average utilization ratio for the retail line, resulting in 27 strata.[9] Appendix Figure A2 shows the distribution of baseline utilization by line.[10]

Figure 2 shows the email messages users received about the credit-limit increases in a subset of the treatment arms for illustration. The messages announced the increase on the affected line(s) and displayed images with examples of goods users could purchase on the line(s). The control group received a comparable email promoting both Cashea lines, with examples of goods users could purchase on each. This design addresses the concern that the messages themselves, rather than the credit-limit increases, could drive our results through salience effects. Each message contained a graphics interchange format (GIF) animation that cycled through the images shown in Figure 2, along with the additional images shown in Appendix Figure A3. Users also received a push notification notifying them of the limit increase, with similar language for the control group.[11]

Our model assumes users are aware of the credit-limit changes and the changes generate the salient cue to re-optimize their budget on the affected line(s). Evidence from our follow-up survey suggests they were well-informed of their new credit limits: roughly half of treated respondents self-report a post-experiment credit limit that is within 2.5% of their actual observed limit, and most of the other half are close (Appendix Figure A4). The treated groups report their limits nearly as accurately as the control group, and knowledge of limits is high at baseline as well (Appendix Figure A5).

4.2 Administrative Data

We analyze administrative data from Cashea which includes each user's purchases, credit limits, installment schedules, and payments. We observe each user's credit limits on both credit lines at a daily frequency, which we use to construct the realized experimental variation and to track non-experimental limit changes. Purchase data are at the transaction level and include information on the purchase amount in USD, merchant, store category, date, and the associated credit line. For each installment of the repayment plan, we observe the scheduled due date and the payment date. Our analysis covers the 12 months from the launch of the experiment on October 3, 2025 through October 2, 2026.

[9] Misfits within a stratum were randomly assigned, which generated some small differences in the number of users allocated to each treatment arm, as in **mckenzie2011stratified**.

[10] We do not show the distribution of credit limits since these are proprietary and business sensitive.

[11] For example, for the Both +100% treatment group, the translated message sent was "We increased your Grocery and Retail Lines by 100%, check out your new available credit." For the control group, it read "Shop now and pay later, interest-free. Open the app to view your purchase limits."

4.3 Survey Data

We supplement the administrative data with three surveys to capture additional information about Cashea users and their financial situations, along with their economic activity that is not observable through the administrative data. Where applicable, the surveys requested answers in USD.

We fielded a baseline survey the week before the experiment launched. The baseline survey elicited spending in each category on and off Cashea, reasons for making purchases on and off Cashea, knowledge of credit limits, basic demographics, and self-reported credit constraints. Cashea sent invitations via email on September 22, 2025 to all users included in the experiment.[12] The baseline survey received responses from 8,728 users for a response rate of 4.2%.

We fielded two follow-up surveys. The first was fielded between one and two months after the experiment was launched, and the second was fielded between seven and eight months after launch. For the first follow-up survey, users received invitations to take the survey on November 6, 2025 via email. Those who had not completed the survey received a second invitation on December 4, 2025 via push notification on their phones. For the second follow-up survey, users received email invitations on May 12, 2026 and push-notification invitations on May 14, 2026. We received 16,341 responses to the first follow-up survey and 12,170 responses to the second, for response rates of 7.8% and 5.8%, respectively. In the main analysis, we pool data from both waves, where we received a total of 28,511 responses from 24,704 unique respondents, for an overall response rate of 11.8%.[13]

Both follow-up surveys elicited spending on each category on and off Cashea, along with knowledge of credit limits, use of other credit, income, and other demographics. The second follow-up survey additionally asked about spending on categories not covered by either credit line, transfers to others, their stock of savings, their flow of savings over the last 30 days, and whether their stock and flow of savings had increased or decreased since September 2025.[14]

Survey respondents tend to be more engaged users with slightly older accounts, higher limits, higher purchase amounts, and higher balances, but baseline covariates are balanced across

[12] Cashea sent a pilot version of the survey to 10,000 users on August 26, 2025, a revised pilot version to 10,000 more users on September 9, 2025, and the final version of the survey to the remaining 189,996 users on September 22, 2025. We include all survey responses across these three versions of the survey. The full questionnaire for the final version is available online at https://tinyurl.com/ccmfz226 in Spanish (original) and at https://tinyurl.com/2m25hwwr in English.

[13] As with the baseline survey, we sent a pilot version of the first follow-up survey to 10,000 users on November 6, 2025 and two pilot versions for the second follow-up survey on May 5, 2026 and May 11, 2026.

[14] The full questionnaire for the first follow-up is available online at https://tinyurl.com/57tcwb8n in Spanish (original) and at https://tinyurl.com/mphdwwd5 in English. Similarly, the full questionnaire for the second follow-up is available online at https://tinyurl.com/36djue7m in Spanish (original) and at https://tinyurl.com/4uusuu2a in English.

treatment arms within survey respondents. Appendix Table A4 compares covariates for survey respondents and non-respondents and Appendix Table A5 tests for balance across treatment arms within the follow-up survey respondents. Appendix Table A6 reports demographics for the baseline survey respondents.

5 Empirical Strategy

5.1 Marginal Propensities to Consume

We use the randomized credit-limit increases to estimate same-category and cross-category MPCs for each credit line. The cross-category MPC is the key test of our model. Without mental accounting, consumers who are not statically constrained respond to a credit-limit increase either not at all or by increasing spending in both categories, so a positive same-category MPC implies a positive cross-category MPC. With mental accounting, the cross-category MPC is zero even when the same-category MPC is positive.

Our experiment induced random variation in one or both credit limits, but credit limits also evolve for users over time as their repayment history increases (or decreases) their limits. Post-treatment credit-limit changes thus reflect both the experimentally induced variation and the ordinary progression that also occurs in the control group. Although natural changes are small relative to the experimental increases, they matter for estimating MPCs. Because the experimental increases are proportional, users with larger baseline limits receive larger dollar increases, and these users also experience larger natural growth in their limits. Regressing spending on post-treatment limits would therefore conflate the experimental variation with natural limit growth, which depends on users' own purchases and repayment.

We isolate the experimentally induced variation in credit limits with an instrumental variables (IV) strategy that instruments for post-treatment credit limits with the treatment assignment indicators. The first-stage equations are:

$$L_i^{\ell,\text{post}} = \sum_{j=1}^{6} \pi_j^{\ell} T_i^j + \theta_{s(i)} + \varepsilon_i^{\ell} \qquad \text{for } \ell \in \{G, R\}, \tag{6}$$

where $L_i^{G,\text{post}}$ and $L_i^{R,\text{post}}$ are user i's post-treatment credit limits on the grocery line and the retail line, T_i^j is an indicator for assignment to treatment arm $j \in \{1, \ldots, 6\}$ (control omitted), and $\theta_{s(i)}$ are randomization strata fixed effects. The fitted values measure the dollars of credit created by assignment.

The second stage regresses spending in one category on the instrumented category-specific

credit limits from the first stage regression:

$$y_i^\ell = \alpha_G \hat{L}_i^{G,\text{post}} + \alpha_R \hat{L}_i^{R,\text{post}} + \theta_{s(i)} + \varepsilon_i, \tag{7}$$

where y_i^ℓ denotes average monthly purchases in USD on line $\ell \in \{G, R\}$ or in total. When the outcome is grocery-line purchases (y_i^G), α_G is the same-category MPC and α_R is the cross-category MPC; the interpretations reverse for retail-line purchases. The experiment separately identifies the same- and cross-category MPCs for both categories by inducing variation in neither limit, one limit, or both limits. The exclusion restriction requires that treatment assignment affects spending only through credit limits, which is satisfied by design. In prespecified heterogeneity analysis, we fully interact each regression with group indicators (e.g., indicators for tercile of baseline credit utilization) to estimate coefficients separately for each group.

We construct purchase outcomes by first summing purchase data from the experiment's launch on October 3, 2025 through October 2, 2026 and then dividing by the number of months in the post-period. All administrative outcomes are winsorized at the 0.1st and 99.9th percentiles. We report heteroskedasticity-robust standard errors, clustered by user in specifications that pool the two follow-up survey waves.

5.2 Treatment Effects by Arm

We additionally estimate treatment effects separately for each treatment arm using the following specification:

$$y_i^\ell = \sum_{j=1}^{6} \tau_j T_i^j + \theta_{s(i)} + \varepsilon_i, \tag{8}$$

where τ_j is the effect of assignment to arm j relative to the control group and $\theta_{s(i)}$ are again randomization strata fixed effects. This specification includes a separate indicator for each arm, so it does not assume that effects scale linearly with the size of the limit increase or that increases on the two lines are additive. Comparing the τ_j coefficients across treatment arms enables us to test both assumptions directly.

6 Results

For both credit lines, category-specific credit-limit increases cause large same-category consumption responses but near-zero spillovers to other-category spending. Figure 3 presents these central findings, estimated using specifications (6) and (7). For each $1 increase in the grocery-line credit

limit, participants consume an additional 9.8 cents per month on their grocery line and consume 1.1 cents *less* per month on their retail line. For each $1 increase in the retail-line credit limit, participants consume an additional 8.1 cents per month on their retail line and a not statistically significant 0.2 cents per month on their grocery line.[15]

Figure 4 reports treatment effects separately by arm, estimated using specification (8). The large same-category increases and near-zero spillovers across categories hold in all four treatment arms in which we increased only one credit line. Increasing both credit limits raises spending on each line by roughly the same amount as only increasing that line's limit. Furthermore, the responses scale nearly linearly with the size of the credit-limit increase. The 100% arms generate almost twice the additional spending of the 50% arms ($14.1 versus $8.0 on the grocery line; $28.5 versus $16.8 on the retail line).[16] This near-linearity supports summarizing the responses with a single MPC per line, as we do throughout the paper.

The remainder of this section presents five additional pieces of evidence that support a mental-accounting interpretation of our main results and rule out alternative explanations: (1) the large same-category MPCs and near-zero cross-category spillovers persist in every post-treatment month (Section 6.1); (2) the same-category responses represent increases in total consumption rather than payment-method substitution (Section 6.2); (3) binding constraints cannot explain the results, which hold for both constrained and unconstrained borrowers (Section 6.3); (4) machine-learning heterogeneity analysis finds significant heterogeneity in same-category effects but no subgroup with a meaningful positive spillover (Section 6.4); and (5) the pattern holds within every merchant sub-category, supporting the mapping of credit lines to mental accounts (Section 6.5). Section 7 synthesizes the implications for consumers.

6.1 Persistence of Consumption Targets

In the model, consumers reset their category-specific consumption targets when a salient cue—in the experiment, this is the category-ℓ credit-limit increase—prompts them to solve the local optimization problem. Because full rebudgeting is costly, the new category-ℓ target, $c^*_{\ell t}$, persists until a future cue triggers re-optimization. Under this model, we expect the same-category MPC to respond immediately and persist over time. We also expect a zero cross-category MPC on impact and over time if consumers do not rebudget.

Our central results are reproduced in every single one of the 12 post-treatment months. Fig-

[15] Results are similar using average balances rather than monthly purchases as the outcome variable (Appendix Figure A6).

[16] Appendix Figure A7 formalizes the comparison, showing similar marginal effects of the first and second 50% increments.

ure 5 presents month-specific same- and cross-category MPCs. Both same- and cross-category MPCs are zero for the three months leading up to the experiment, as expected. For both credit lines, the same-category MPCs in the first month are close to the 12-month averages in Figure 3 (0.079 vs. 0.098 for the grocery line, 0.064 vs. 0.081 for the retail line).[17] The same-category effects are thus both immediate and durable, remaining large and statistically significant in every post-treatment month.

The cross-category spillovers are likewise near zero in the first month of the experiment and in every one of the subsequent 11 months. Both cross-category MPCs decline over time and eventually become slightly negative; we discuss implications of this result for how people finance their increased consumption in Section 7.

The persistence of the same-category MPCs also helps rule out the alternative explanation that our results are caused by some goods being indivisible (Zorzi and Beraja, 2026). If credit-limit increases primarily facilitate indivisible durable purchases, we could observe positive same-category MPCs and zero spillovers because users were statically constrained in the sense that their available credit limit plus savings was less than the cost of the indivisible good. We would then expect the same-category MPCs to spike upon treatment as people purchase previously unattainable indivisible goods, but to quickly erode over time. Instead, they do not erode, consistent with a persistent new consumption target. We discuss indivisibility further in Section 6.5.

6.2 Effects on Total Consumption: Evidence from Survey Data

The model predicts that a category-specific credit-limit increase raises total consumption in the corresponding category, rather than just spending on that category's credit line. An alternative explanation for our results in administrative data is payment-method substitution. Under this alternative, consumers would buy the same total amount but move purchases they would otherwise have paid for in cash onto the expanded credit line. Because the administrative data capture only spending on Cashea, we supplement them with survey measures of category-specific spending on and off Cashea. In our baseline survey, users report making 47.1% of their grocery-category purchases and 59.6% of their retail-category purchases through Cashea.

Figure 6 presents the same- and cross-category MPCs separately for spending on Cashea, off Cashea, and combined, using survey data for all outcomes. First, we validate the quality of the survey data by estimating MPCs for on-Cashea spending, which we directly observe in administrative data. We find very similar same-category MPCs of 0.065 for the grocery line (compared to 0.084 in administrative data over the same period) and 0.050 for the retail line (compared to 0.062

[17] Appendix Figure A8 shows treatment effects by arm over time.

in administrative data).

Second, the survey reproduces the near-zero cross-category MPCs we find in administrative data over the same period (Appendix Table A7). Turning to off-Cashea spending, which we observe only in the survey, we find no evidence of payment-method substitution. For both lines, the same- and cross-category MPCs are near zero and not statistically significant. Thus, the MPCs we estimate for total spending on each category are also remarkably consistent with the on-Cashea MPCs we estimate from administrative data, at 0.071 and 0.068 for the same-category grocery and retail MPCs, while we cannot reject zero spillover for the cross-category total spending MPCs. Collectively, these results suggest that the on-Cashea spending increases we observe in the administrative data represent increases in total consumption rather than payment-method substitution.

6.3 The Role of Constraints

In the absence of mental accounting, our model generates positive same-category MPCs and zero cross-category MPCs *only* if a borrower is statically constrained. That is, in the counterfactual without a credit-limit increase, the borrower would have immediately maxed out their category-ℓ credit line *and* have no cash in savings or spending on other categories which they could reallocate to category ℓ. In contrast, a consumer who is dynamically but not statically constrained—one who would not max out their category-ℓ credit limit unless hit by a negative income shock in the future—will respond to a category-ℓ credit-line increase by increasing their current consumption in *both* categories, unless they engage in mental accounting.

Thus, an alternative explanation for the large same-category and zero cross-category MPCs could be that users are statically constrained. We test this by estimating MPCs separately by terciles of baseline credit-line utilization, which was one of our two main prespecified heterogeneity tests. If static constraints explain the main result, positive same-category MPCs with near-zero spillovers should appear only among users who are already using nearly all of their category-ℓ credit limit.

Figure 7 shows that even for users who are in the bottom two terciles—with baseline utilization ratios of less than 10.0% and 24.9%, respectively—same-category MPCs are economically meaningful and we can rule out positive cross-category MPCs above 0.014 across all terciles and both lines. For example, for the middle tercile with baseline utilization ratios between 10.0% and 24.9%, the same-category MPCs are 0.087 and 0.069 for the grocery and retail lines, and we cannot reject that the cross-category MPCs are zero. As we would expect if the most-constrained tercile is either the most dynamically constrained or faces binding static constraints, same-category MPCs are largest for this group. We find the same core pattern in every tercile and on both lines, despite the

fact that the two lines differ in baseline limits, repayment terms, and merchant types. Appendix Figure A9 shows that results by terciles of baseline utilization of each line separately (rather than total utilization) are similar. Appendix Figure A10 shows that we find similar results when we test for heterogeneity using a baseline survey measure of *self-reported* credit constraints, as we prespecified: even the users who self-report being unconstrained have positive and economically meaningful same-category MPCs but zero cross-category MPCs.[18]

Descriptive evidence confirms that few users are statically constrained. Median baseline utilization is only 16.8% on the grocery line and 12.0% on the retail line, 74.0% of users report typically spending cash on both categories, and 79.6% report having liquid savings. Appendix Figure A13 shows that, in the control group, less than 20% of users were constrained on the grocery line (panel (a)) and less than 10% on the retail line (panel (c)) in the month before the experiment or in the first month of the experiment—when we already see an immediate effect of the credit-line increases. Nevertheless, many users do use most of their credit limit *at some point* over the course of the experiment (panels (b) and (d)). This is consistent with dynamic constraints, which create a precautionary motive that explains why consumers respond to a credit-line increase at all but not why they increase consumption on only one category.

6.4 Heterogeneity by Machine-Learning Predicted Treatment Effects

To test whether some groups of consumers behave according to standard models without mental accounting, we use machine learning to estimate heterogeneity in treatment effects by observables (Chernozhukov, Demirer, Duflo, and Fernández-Val, 2025). Specifically, we use all available baseline characteristics in administrative data and estimate four separate models to predict heterogeneity in the treatment effects of an increase in each line on same-category and cross-category spending (see Appendix C). If some groups treat money fungibly, we should be able to identify a group with a large, positive spillover from an increase in the credit limit on one line to consumption on the other category.

We sort users into quintiles of their predicted treatment effect separately for each of four effects, corresponding to the same-category and cross-category treatment effects of an increase in each credit line. We find substantial heterogeneity in same-category consumption responses, with positive effects in all five quintiles on both lines.

In contrast to standard models in which money is fungible, we do not find a meaningful pos-

[18] We also find positive and meaningful same-category MPCs and zero or near-zero cross-category MPCs for all groups in the other two heterogeneity tests we prespecified: heterogeneity by credit limit on each line (Appendix Figure A11) and by the proportion of consumption in each category that the user reported spending through Cashea in the baseline survey (Appendix Figure A12).

itive spillover in any quintile of predicted spillover on either line, despite finding significant heterogeneity in same-category responses (Figure 8). When grocery-line limits increase, the spillover is negative or not statistically distinguishable from zero in every quintile. When retail-line limits increase, the same holds in every quintile except the one with the highest predicted spillover, and even there the spillover is small. Its point estimate implies a 3.5% increase in grocery consumption, compared to an 18.5% increase in retail consumption for the same consumers. A standard model in which money is fungible would predict similar proportional increases across categories, so even this group appears to engage in mental accounting. We conclude that consumers in every subgroup we can identify from observable baseline characteristics form mental accounts.

6.5 Sub-Category MPCs and Mapping Credit Lines to Mental Accounts

An important assumption throughout our analysis is that the grocery and retail credit lines correspond to consumers' mental accounts as specified in the model. Each line bundles several merchant sub-categories—grocery stores, pharmacies, and restaurants on the grocery line and home goods, clothing, services, technology, tools, health, and sports on the retail line—yet we model each bundle as a single account with a single consumption target. If consumers formed different mental accounts (e.g., combining health and pharmacy spending) and we misallocated merchant sub-categories, we would expect positive cross-category MPCs, biasing us *against* finding zero spillovers.

We test our assumptions about the composition of consumers' mental accounts by estimating MPCs separately for each merchant sub-category, normalized by its baseline share of spending on its corresponding credit line (Figure 9). Remarkably, these MPCs coincide with Cashea's assignment of merchant sub-categories to credit lines. Specifically, we estimate positive and statistically significant same-category MPCs for every single one of the 10 merchant sub-categories. Further, the normalized cross-category MPCs are near zero for every merchant sub-category, ranging from -0.014 to 0.013, and we can rule out any larger than 0.028. We view these results as compelling evidence that the two credit lines map cleanly to consumers' mental accounts. This is consistent with evidence from Read, Loewenstein, and Rabin (1999) that consumers accept the brackets imposed upon them.

The sub-category MPCs help further rule out indivisibility as a sufficient alternative explanation. Same-category MPCs are positive and statistically significant even for the most divisible sub-categories on each line, such as groceries and clothing, and the responses to the two lines are symmetric.

7 Discussion

Our credit-limit changes caused large changes in the level and allocation of consumption. Doubling the credit limit on the retail line increased retail consumption by $29 per month, or 5.9% of median monthly income, with essentially no change in grocery consumption. Doubling the grocery-line credit limit increased grocery consumption by $14 per month (slightly larger per dollar of credit-line increase than the same-category retail-line response), or 2.9% of median monthly income, again with essentially no change in retail consumption. This dramatic shift in the allocation of consumption is especially striking given that the vast majority of consumers were not statically constrained. The ubiquity of this effect across consumers is similarly striking: we find no subgroup of users with the meaningful positive cross-category spillover that a standard model without mental accounting would predict.

Given these facts, what are the implications for consumers? The money for the increased consumption must come from somewhere. Unlike credit cards, these are short-maturity loans, so consumers cannot just delay repaying their debt into the future. The results are also persistent and occur not only for on-credit-line spending but also for total spending, so they cannot be explained by consumers pulling forward consumption or shifting payment methods. Our second follow-up survey asked several questions about where this money comes from.[19]

Appendix Table A8 shows that we do not observe a statistically significant change across our six treatment arms in any of the variables we measure—including other spending on categories not covered by the grocery and retail lines, the flow or stock of savings, other debt, income, or transfers/loans to others. While we are unable to conclusively trace any universal source of money used to fund the increased spending, we can rule out large negative effects of the increased spending due to credit-limit increases: participants do not substantially decrease their savings or increase their other debt. The negative cross-category MPCs that emerge over time could indicate that people fund some share of their persistently higher same-category consumption by decreasing savings in the short run, and once their budget constraint binds in the future, reducing spending on the other category.

A related concern is whether expanded access to credit increases default. Panel (a) of Appendix Figure A14 shows effects on the probability that a user has an installment past due as of 12

[19]These questions capture other spending on categories not covered by the grocery and retail lines (e.g., rent or mortgage payments, utilities, transportation, education, health, and entertainment); the flow of savings over the last 30 days and a question on whether this was higher, about the same, or lower than their flow of savings in September 2025; the current stock of savings and a question on whether this was higher, about the same, or lower than in September 2025; debt with other lenders including personal loans, credit cards, store financing, installment payments, or loans from informal moneylenders or family and friends; income (as more access to credit could in theory increase income); and transfers or loans to others. Only 26.4% of participants had any other debt from any of these sources.

months post-intervention by treatment arm. There are no statistically significant effects for the retail treatment arms, and only small increases for the grocery-line and both-line treatment groups. The largest increase is just a 1.7 percentage-point increase (13.0%) in the probability that a user has an installment 90+ days past due for the treatment arm in which the credit limit was doubled for both lines. The credit-limit increases were profitable for Cashea in all treatment arms (Panel (b) of Appendix Figure A14). The results suggest that consumers are not taking on more debt than they can repay and even the large credit-limit increases induced by the experiment did not reflect an over-extension of credit by Cashea.

Finally, we asked participants directly in the second follow-up survey whether they wished they had spent less on Cashea over the prior six months. Just 13.1% of treated participants reported wishing they had spent less. While we document large changes in both the level and allocation of consumption, the vast majority of consumers seem content with their borrowing decisions and we see only modest increases in delinquency considering the large magnitude of the credit-limit changes.

8 Conclusion

We test whether mental accounting for credit affects borrowing decisions. In our model of mental accounting for credit, consumers respond to an increase in a category-specific credit line by increasing consumption *only* on that category, whereas without mental accounting, consumers who are not statically constrained increase consumption on both categories (or neither). We conduct an experiment with the largest consumer credit provider in Venezuela to test this theory: all of this lender's users have access to two category-specific credit lines, one for grocery stores, pharmacies, and restaurants, and the other for all other retail merchants. We randomized large 50% or 100% increases to one credit line, the other line, or both.

We find strong evidence in support of mental accounting. Our central finding is that, for each $1 of category-specific credit-limit increase, consumers increase their spending on that category by an additional 8–10 cents per month with near-zero cross-category spillover. The grocery- and retail-line results are mirror images. We present five additional pieces of supporting evidence: (1) the results persist in every single month in the post-period; (2) the increase in credit-line spending is not offset by reduced cash spending on the category; (3) the results hold for both constrained and unconstrained consumers; (4) machine-learning techniques fail to identify any subgroup of consumers who treat the relaxation of the budget constraint fungibly; and (5) the central result replicates in every single merchant sub-category of spending within the credit lines.

Mental accounting has a prominent position in the behavioral economics canon but causal empirical evidence for it is comparatively limited. This is especially true in credit markets, despite many credit products being category-specific (e.g., mortgages, auto loans, trade credit, and "buy now, pay later" credit available only at partnering merchants) or potentially priming consumers to form mental accounts (e.g., credit cards with much higher rewards for specific spending categories such as travel or dining).

Understanding whether people form mental accounts around credit products matters for policy. Our findings suggest that consumers do form mental accounts for credit products, and that they accept the mental accounts imposed on them. This gives policymakers power both to shape consumers' mental accounts and to alter the composition of spending through subsidies and regulation. This is increasingly relevant as financial technology dramatically expands the provision of category-specific credit (Acree, Barnes, Bruce, and Hannon, 2026). Furthermore, mental accounting may dampen the pass-through of monetary policy, since shifts in different parts of the yield curve change the supply of different credit products, which may then raise consumption only in the categories those products finance.

References

Acree, Nina R., Kayleigh Barnes, Alexander Bruce, and Simona Hannon (2026). "'Buy Now, Pay Later' Beyond 'Pay in 4', A Comprehensive Product Overview." *FEDS Notes*.

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel (2018). "Do Banks Pass through Credit Expansions to Consumers Who Want to Borrow?" *Quarterly Journal of Economics* 133(1), 129–190.

Aiyagari, S. Rao (1994). "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics* 109(3), 659–684.

Alok, Shashwat, Pulak Ghosh, Ankit Kalda, and Chirasree Mitra (2026). "The Buy Now Pay Later Divide: Merchant Heterogeneity and Market Structure." Working Paper.

Argyle, Bronson S., Taylor D. Nadauld, and Christopher J. Palmer (2020). "Monthly Payment Targeting and the Demand for Maturity." *Review of Financial Studies* 33(11), 5416–5462.

Aydin, Deniz (2022). "Consumption Response to Credit Expansions: Evidence from Experimental Assignment of 45,307 Credit Lines." *American Economic Review* 112(1), 1–40.

Aydin, Deniz and Olivia Kim (2024). "Precautionary Debt Capacity." Working Paper.

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

Berg, Tobias, Valentin Burg, Ana Gombović, and Manju Puri (2020). "On the Rise of FinTechs: Credit Scoring Using Digital Footprints." *Review of Financial Studies* 33(7), 2845–2897.

Berg, Tobias, Valentin Burg, Jan Keil, and Manju Puri (2025). "The Economics of 'Buy Now, Pay Later': A Merchant's Perspective." *Journal of Financial Economics* 171, 104093.

Berg, Tobias, Andreas Fuster, and Manju Puri (2022). "FinTech Lending." *Annual Review of Financial Economics* 14(1), 187–207.

Bian, Wenlong, Lin William Cong, and Yang Ji (2023). "The Rise of E-Wallets and Buy-Now-Pay-Later: Payment Competition, Credit Expansion, and Consumer Behavior." NBER Working Paper 31202.

Boehm, Johannes, Etienne Fize, and Xavier Jaravel (2025). "Five Facts about MPCs: Evidence from a Randomized Experiment." *American Economic Review* 115(1), 1–42.

Bordalo, Pedro, Nicola Gennaioli, Giacomo Lanzani, and Andrei Shleifer (2026). "A Cognitive Theory of Reasoning and Choice." *Quarterly Journal of Economics* 141(3), 1921–1963.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2020). "Memory, Attention, and Choice." *Quarterly Journal of Economics* 135(3), 1399–1442.

Bosshardt, Joshua, Marco Di Maggio, Ali Kakhbod, and Amir Kermani (2024). "The Credit Supply Channel of Monetary Policy Tightening and its Distributional Impacts." *Journal of Financial Economics* 160, 103914.

Chava, Sudheer, Rohan Ganduri, Nikhil Paradkar, and Linghang Zeng (2023). "Shocked by Bank Funding Shocks: Evidence from Consumer Credit Cards." *Review of Financial Studies* 36(10). Ed. by Itay Goldstein, 3906–3952.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Iván Fernández-Val (2025). "Fisher–Schultz Lecture: Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments, with an Application to Immunization in India." *Econometrica* 93(4), 1121–1164.

Choi, James J., David Laibson, and Brigitte C. Madrian (2009). "Mental Accounting in Portfolio Choice: Evidence from a Flypaper Effect." *American Economic Review* 99(5), 2085–2095.

Chorniy, Anna, Amy Finkelstein, and Matthew J. Notowidigdo (2025). "Paternalistic Social Assistance: Evidence and Implications from Cash vs. In-Kind Transfers." NBER Working Paper 34506.

D'Acunto, Francesco, Thomas Rauter, Christoph K. Scheuch, and Michael Weber (2020). "Perceived Precautionary Savings Motives: Evidence from FinTech." NBER Working Paper 26817.

deHaan, Ed, Jungbae Kim, Ben Lourie, and Chenqi Zhu (2024). "Buy Now Pay (Pain?) Later." *Management Science* 70(8), 5586–5598.

Di Maggio, Marco, Justin Katz, and Emily Williams (2022). "Buy Now, Pay Later Credit: User Characteristics and Effects on Spending Patterns." NBER Working Paper 30508.

Farhi, Emmanuel and Xavier Gabaix (2019). "Optimal Taxation with Behavioral Agents." NBER Working Paper 21524.

Garber, Gabriel, Atif Mian, Jacopo Ponticelli, and Amir Sufi (2024). "Consumption Smoothing or Consumption Binging? The Effects of Government-Led Consumer Credit Expansion in Brazil." *Journal of Financial Economics* 156, 103834.

Gelman, Michael and Nikolai Roussanov (2024). "Managing Mental Accounts: Payment Cards and Consumption Expenditures." *Review of Financial Studies* 37(8), 2586–2624.

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

Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang (2020). "The Marginal Propensity to Consume over the Business Cycle." *American Economic Journal: Macroeconomics* 12(2), 351–384.

Guttman-Kenney, Benedict, Chris Firth, and John Gathergood (2023). "Buy Now, Pay Later (BNPL) ...on Your Credit Card." *Journal of Behavioral and Experimental Finance* 37, 100788.

Hastings, Justine S. and Jesse M. Shapiro (2013). "Fungibility and Consumer Choice: Evidence from Commodity Price Shocks." *Quarterly Journal of Economics* 128(4), 1449–1498.

Hastings, Justine S. and Jesse M. Shapiro (2018). "How Are SNAP Benefits Spent? Evidence from a Retail Panel." *American Economic Review* 108(12), 3493–3540.

Katz, Justin (2023). "Saving and Consumption Responses to Student Loan Forbearance." Working Paper.

Katz, Justin, Dominic Russel, and Claire Shi (2026). "The Supply Side of Consumer Debt Repayment." Working Paper.

Kőszegi, Botond and Filip Matějka (2020). "Choice Simplification: A Theory of Mental Budgeting and Naive Diversification." *Quarterly Journal of Economics* 135(2), 1153–1207.

Laudenbach, Christine, Elin Molin, Kasper F. Roszbach, and Talina Sondershaus (2025). "Buy Now Pay (Less) Later: Leveraging Private BNPL Data in Consumer Banking." Working Paper.

Lian, Chen (2021). "A Theory of Narrow Thinking." *Review of Economic Studies* 88(5), 2344–2374.

Mian, Atif, Amir Sufi, and Emil Verner (2017). "Household Debt and Business Cycles Worldwide." *Quarterly Journal of Economics* 132(4), 1755–1817.

Milkman, Katherine L. and John Beshears (2009). "Mental Accounting and Small Windfalls: Evidence from an Online Grocer." *Journal of Economic Behavior & Organization* 71(2), 384–394.

Müller, Karsten and Emil Verner (2024). "Credit Allocation and Macroeconomic Fluctuations." *Review of Economic Studies* 91(6), 3645–3676.

Ponce, Alejandro, Enrique Seira, and Guillermo Zamarripa (2017). "Borrowing on the Wrong Credit Card? Evidence from Mexico." *American Economic Review* 107(4), 1335–1361.

Rabin, Matthew and Georg Weizsäcker (2009). "Narrow Bracketing and Dominated Choices." *American Economic Review* 99(4), 1508–1543.

Read, Daniel, George Loewenstein, and Matthew Rabin (1999). "Choice Bracketing." *Journal of Risk and Uncertainty* 19(1), 171–197.

Sussman, Abigail B. and Rourke L. O'Brien (2016). "Knowing When to Spend: Unintended Financial Consequences of Earmarking to Encourage Savings." *Journal of Marketing Research* 53(5), 790–803.

Thaler, Richard (1985). "Mental Accounting and Consumer Choice." *Marketing Science* 4(3), 199–214.

Thaler, Richard (1990). "Anomalies: Saving, Fungibility, and Mental Accounts." *Journal of Economic Perspectives* 4(1), 193–205.

Thaler, Richard (1999). "Mental Accounting Matters." *Journal of Behavioral Decision Making* 12(3), 183–206.

Tversky, Amos and Daniel Kahneman (1981). "The Framing of Decisions and the Psychology of Choice." *Science* 211(4481), 453–458.

Yin, Xiao (2026). "Learning in the Limit: Income Inference from Credit Extensions." *Journal of Finance* 81(5), 2729–2772.

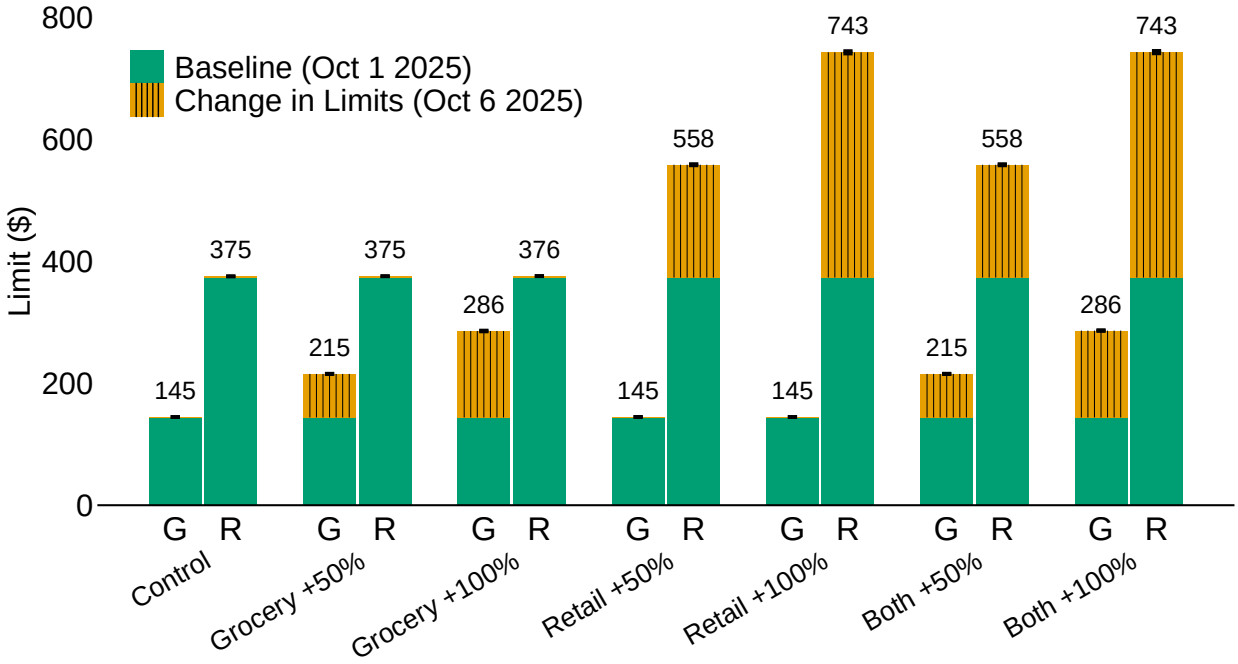
Zeldes, Stephen P. (1989). "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97(2), 305–346.

Zhang, C. Yiwei, Abigail B. Sussman, Nathan Wang-Ly, and Jennifer K. Lyu (2022). "How Consumers Budget." *Journal of Economic Behavior & Organization* 204, 69–88.

Zorzi, Nathan and Martin Beraja (2026). "Durables and the Marginal Propensity to Spend." *American Economic Review* 116(11), 4077–4115.

Figures

Figure 1. Credit-Limit Intervention



Notes: Bars show mean credit limits in USD by treatment arm and credit line; numbers above bars report the total. Whiskers show 95% confidence intervals (heteroskedasticity-robust) for each arm's mean change in limits. "G" denotes grocery line and "R" denotes retail line. Green bars reflect the baseline mean limit on each line as of October 1, 2025, pooled across all arms. Orange bars reflect the change in limits as of October 6, 2025, three days after the experiment launched. Arm-level baseline means and mean changes are reported in Appendix Table A3. See Appendix Figure A1 for credit-limit changes over the entire experimental period. $N = 209{,}996$.

Figure 2. Experiment Outreach

A. Control





B. Grocery Line +50%





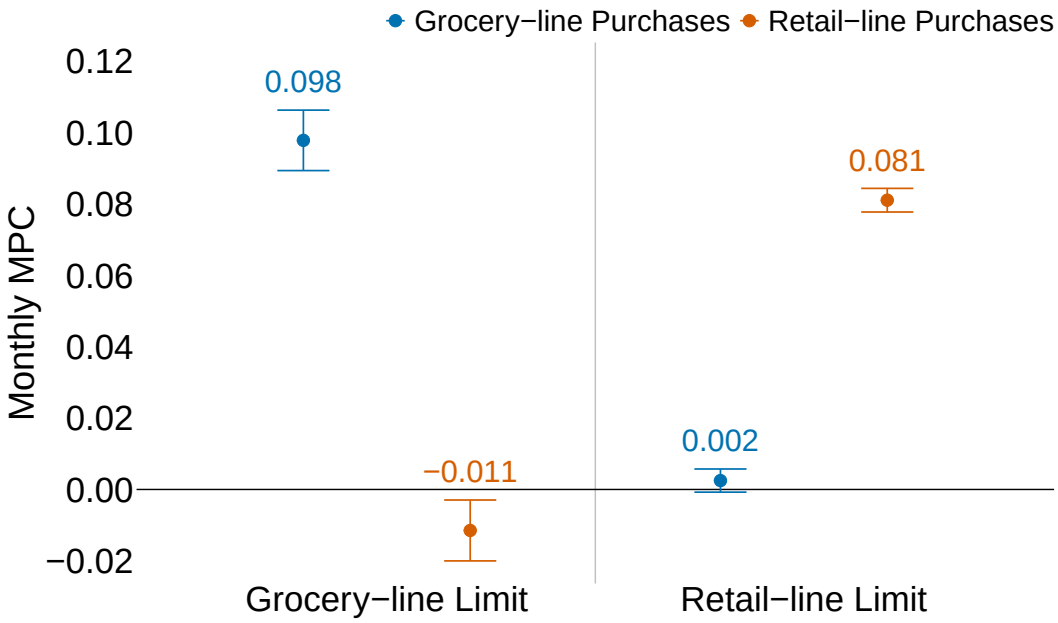
C. Both Lines +100%





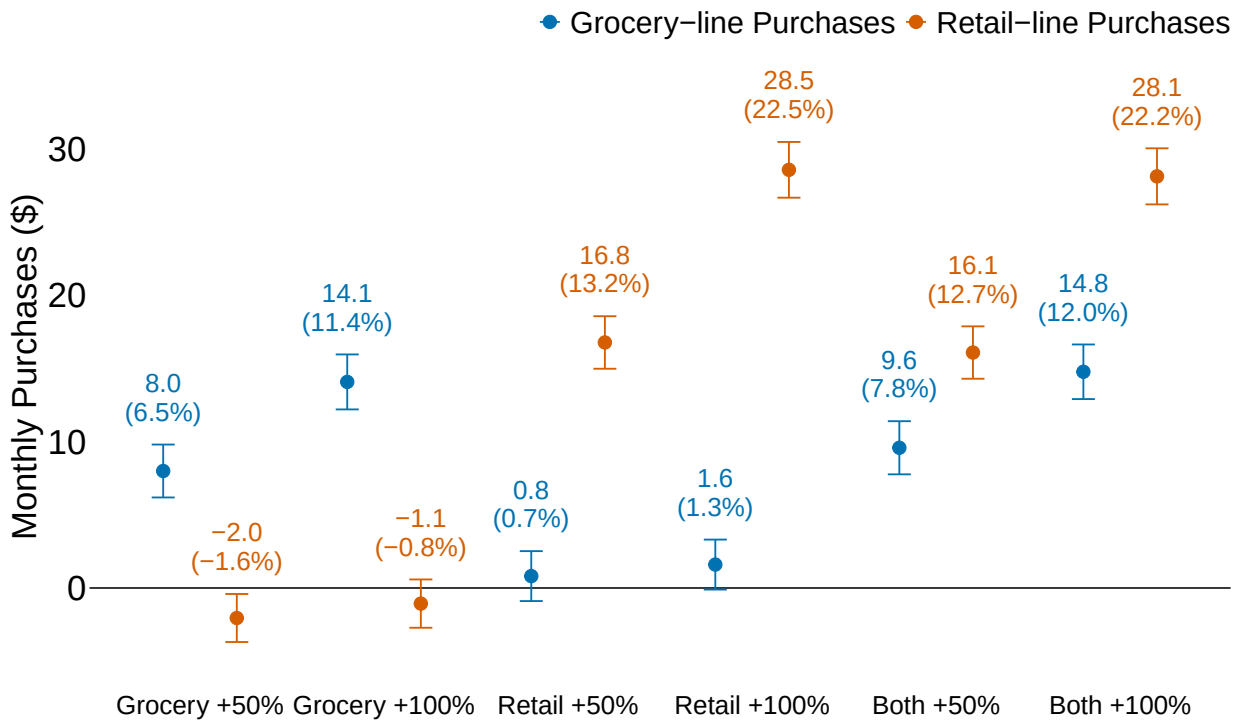
Notes: The figure displays screenshots of email outreach sent to participants, showing three of seven arms as examples. Panel A shows the control-group message, which promotes Cashea without mentioning a credit-limit increase. Panel B shows the treatment message for a 50% grocery-line increase. Panel C shows the treatment message for a 100% increase on both lines. All treatment and control arms received analogous outreach. The message contained a GIF animation that cycled through the images shown, as well as additional images shown in Appendix Figure A3.

Figure 3. Marginal Propensities to Consume out of Credit-Limit Increases



Notes: This figure plots instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line, estimated via equations (6) and (7). The dependent variable is average monthly purchases on each line from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Post-treatment credit limits are measured on October 6, 2025, and are also winsorized at the 0.1st and 99.9th percentiles. Whiskers represent 95% confidence intervals. $N = 209{,}996$. Full estimates are reported in Appendix Table A9.

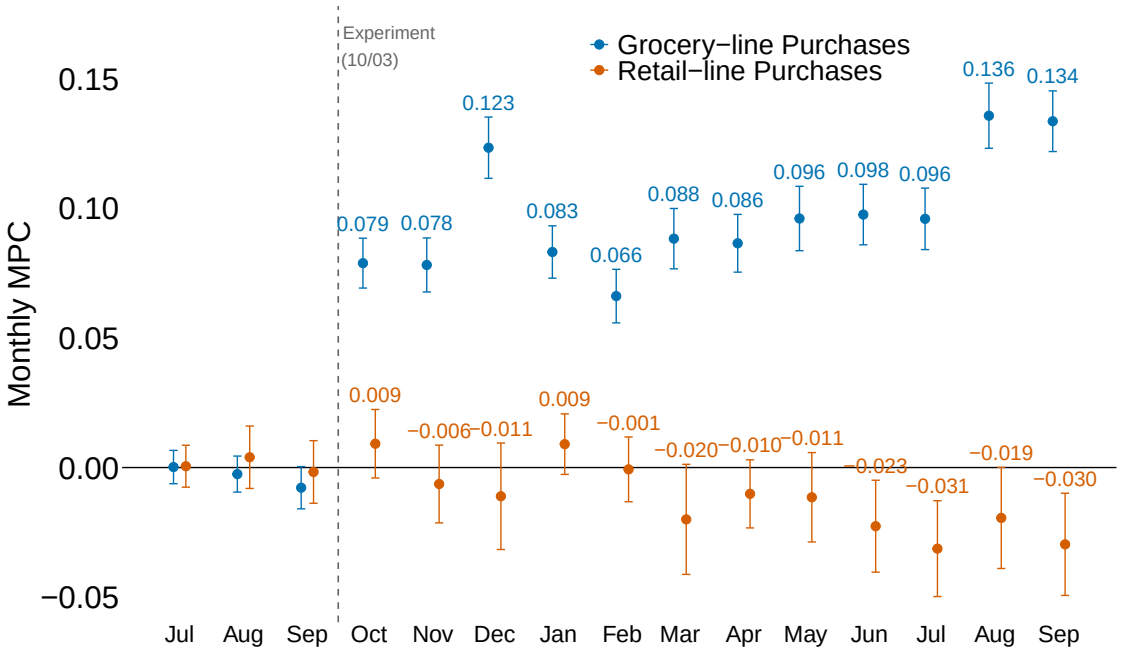
Figure 4. Treatment Effect Estimates of Monthly Purchase Effects



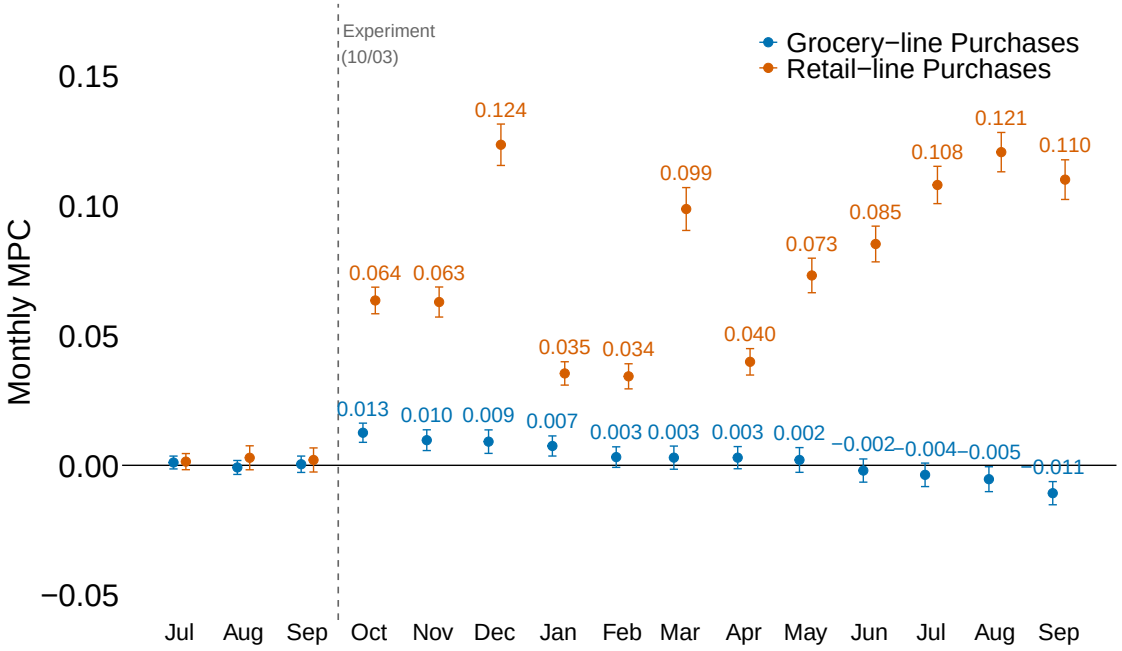
Notes: This figure plots treatment effect estimates of credit-limit increases on each line, estimated via equation (8). The dependent variable is average monthly purchases on each line from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Whiskers represent 95% confidence intervals. Labels report each point estimate in dollars and, in parentheses, as a percentage of the control-group mean of $123 for the grocery line and $127 for the retail line. $N = 209{,}996$. Full estimates are reported in Appendix Table A10.

Figure 5. MPC Estimates by Month

(a) Grocery-line Limit

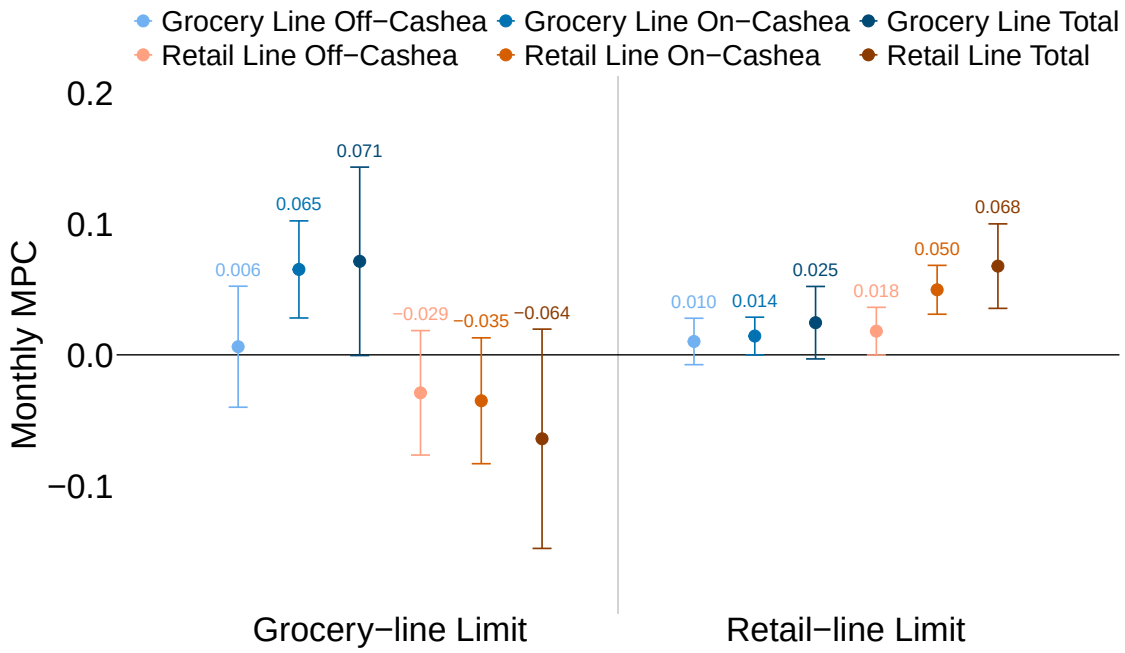


(b) Retail-line Limit



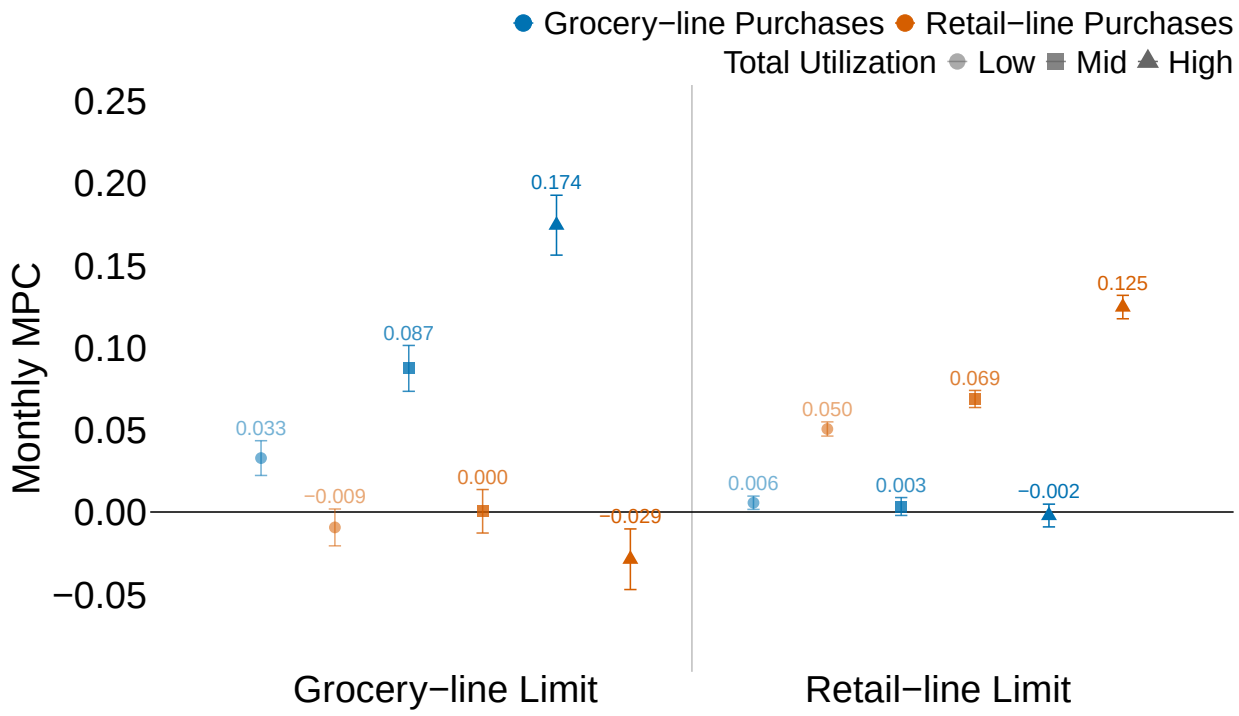
Notes: This figure plots monthly instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line, estimated via equations (6) and (7), separately for each calendar month, with each user's credit limits on October 6, 2025 as the post-treatment limits. The dependent variables are monthly purchases on each line from July 2025 through September 2026, winsorized at the 0.1st and 99.9th percentiles within each month. Whiskers represent 95% confidence intervals. $N = 209{,}996$. Full estimates are reported in Appendix Table A11.

Figure 6. Marginal Propensities to Consume: Cashea and Non-Cashea Spending



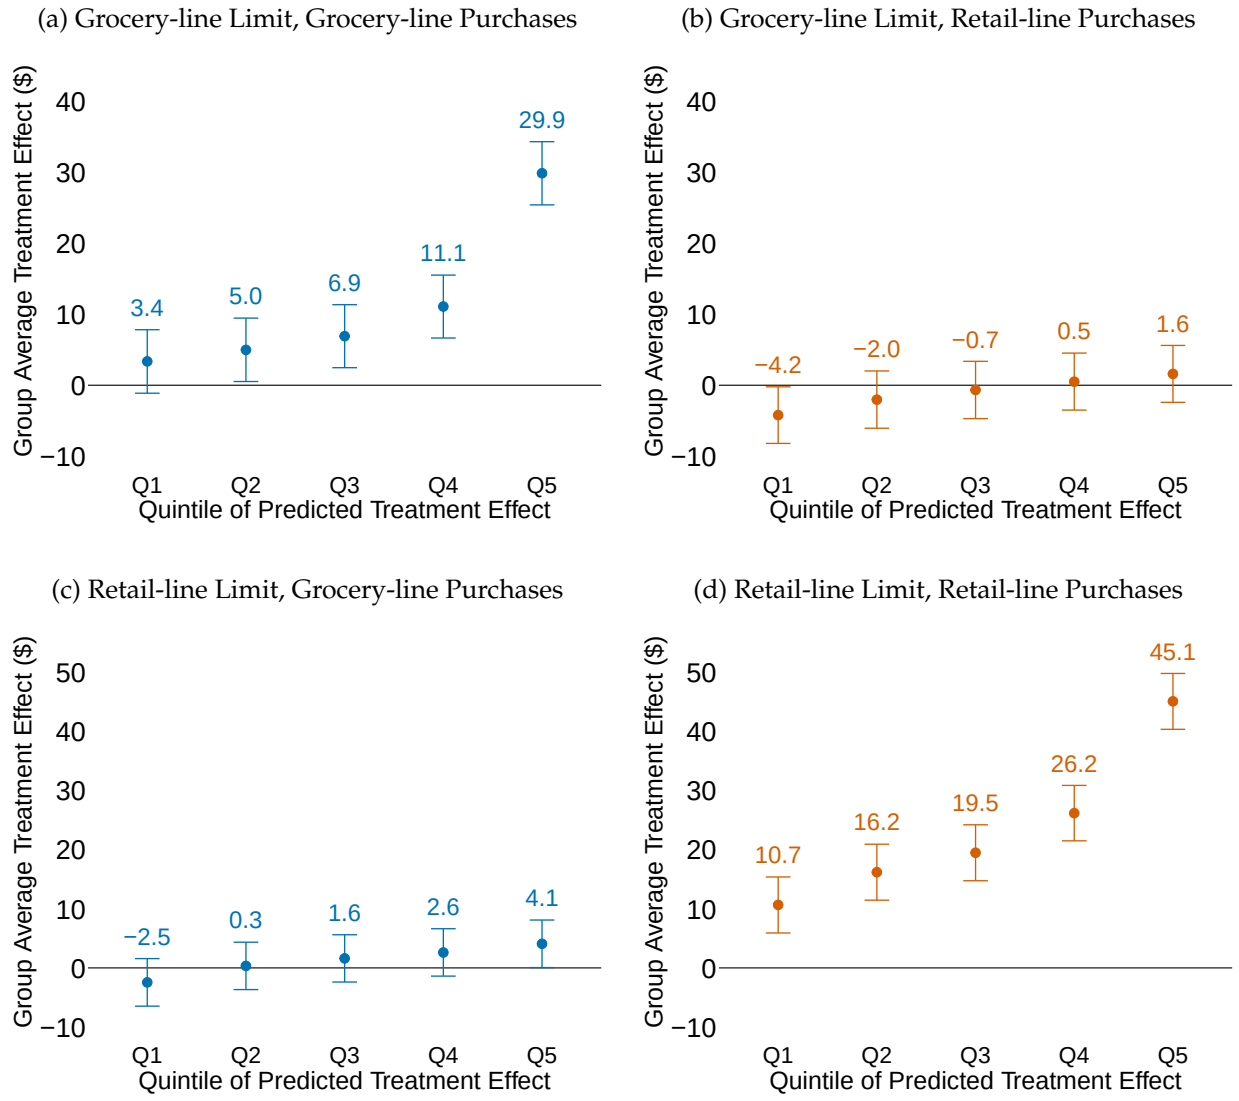
Notes: The figure plots instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line, estimated via equations (6) and (7). The dependent variables are self-reported spending outcomes from a pooled sample of first- and second-wave follow-up respondents, normalized to their 30-day equivalent and winsorized at the 95th percentile. Whiskers represent 95% confidence intervals. $N = 23{,}541$ respondent-waves; standard errors are clustered by user. Full estimates are reported in Appendix Table A7.

Figure 7. Marginal Propensity to Consume by Total Baseline Utilization



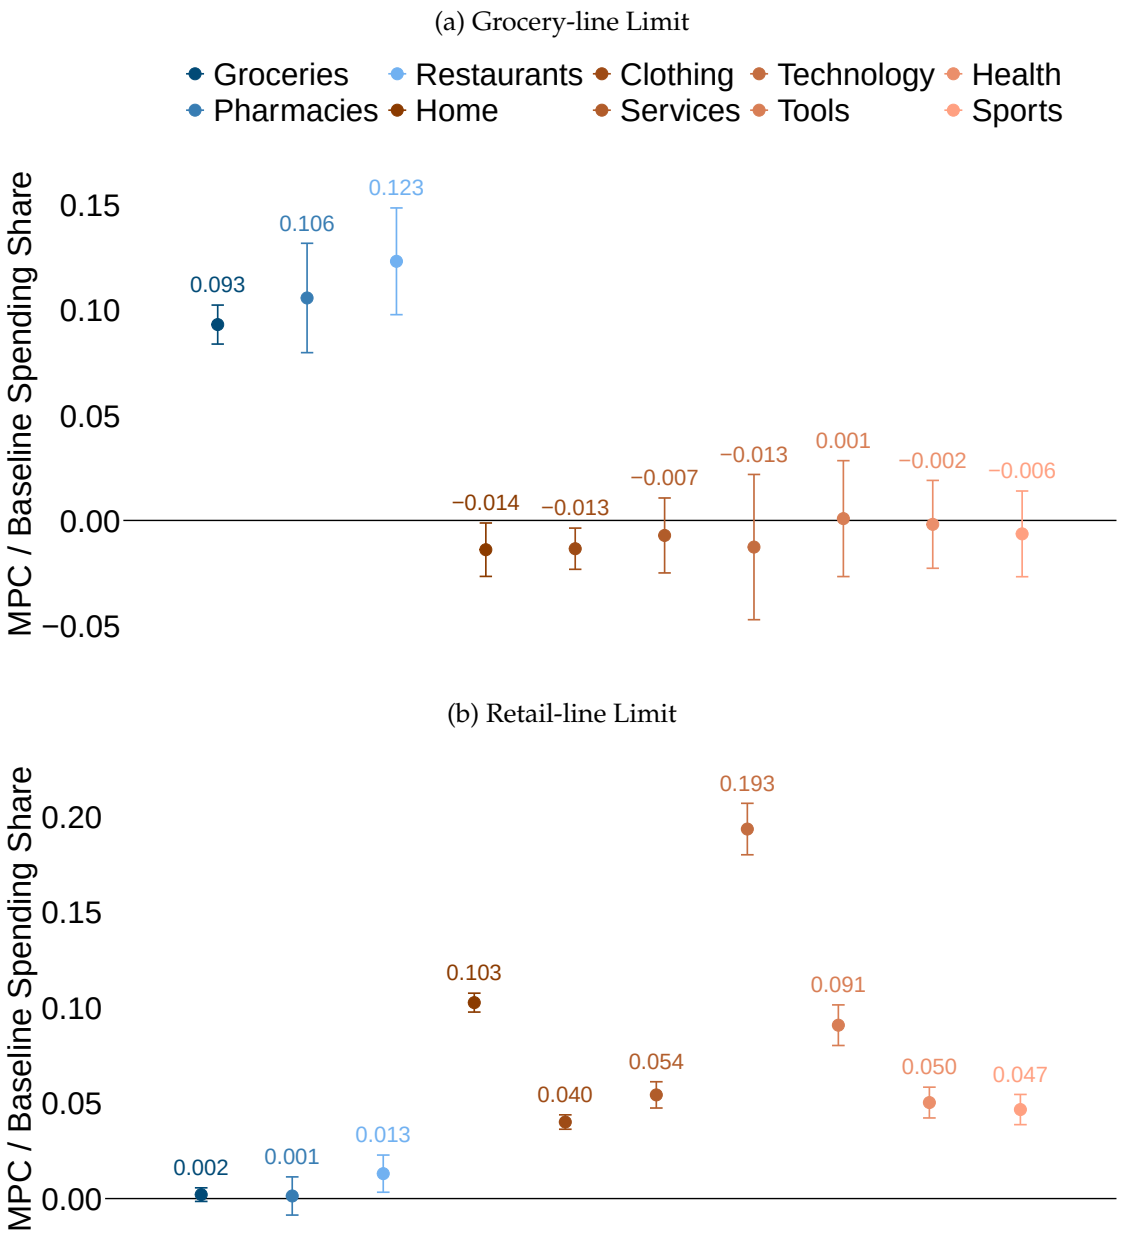
Notes: The figure plots instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line by tercile of baseline total utilization, estimated via versions of equations (6) and (7) in which every term except the strata fixed effects is interacted with tercile indicators. The dependent variable is average monthly purchases on each line from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Whiskers represent 95% confidence intervals. We conduct joint Wald tests of equality across the three tercile coefficients. For cross-line effects, $p = 0.163$ for the retail-line limit $\rightarrow$ grocery-line purchases pair and $p = 0.042$ for the grocery-line limit $\rightarrow$ retail-line purchases pair. $N = 209{,}996$.

Figure 8. Treatment Effects by Quintile of Predicted Individual Treatment Effect



Notes: The figure plots Group Average Treatment Effects (GATES) from the machine-learning heterogeneity procedure of Chernozhukov, Demirer, Duflo, and Fernández-Val (2025). Panels (a) and (b) restrict the sample to the control group and the two grocery-line arms (+50% and +100%); panels (c) and (d) to the control group and the two retail-line arms. The dependent variable is average monthly purchases on each line from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Point estimates are medians over 100 data splits; whiskers represent medians across sample splits of the 95% confidence bounds. $N = 87{,}578$ in panels (a) and (b) and $N = 87{,}631$ in panels (c) and (d). Corresponding average and heterogeneous treatment effects are reported in Appendix Table A12.

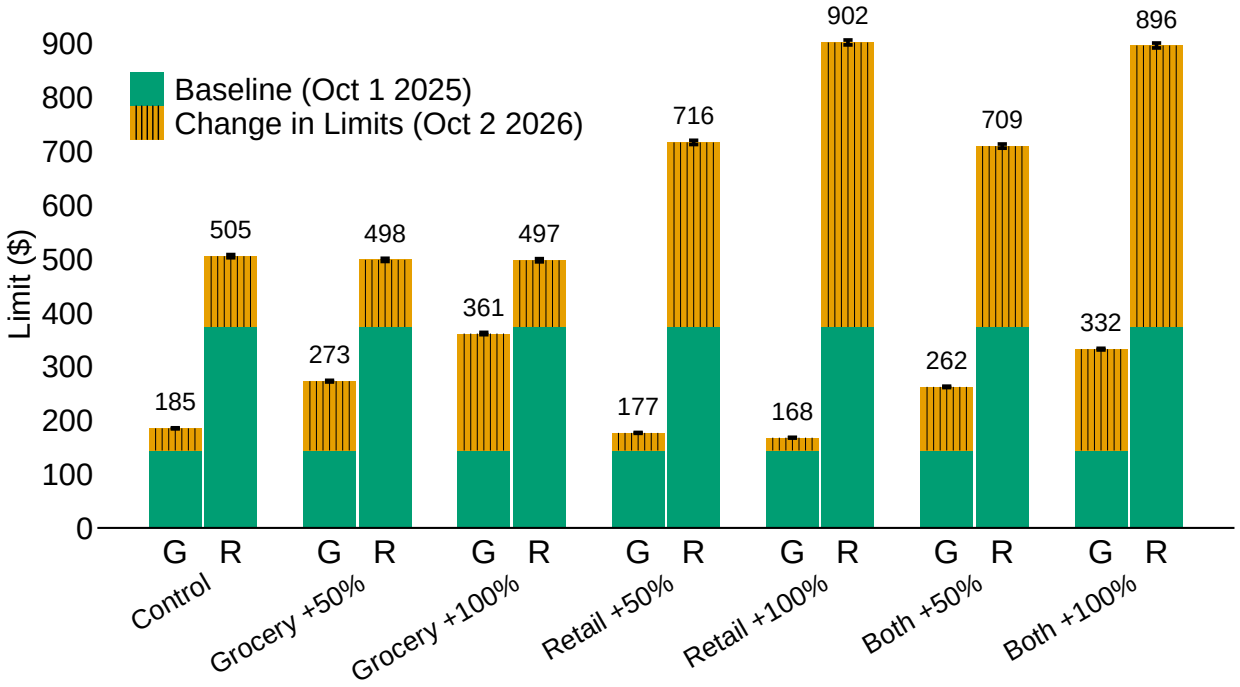
Figure 9. Marginal Propensities to Consume by Merchant Sub-Category



Notes: The figure plots instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line by sub-category, estimated via equations (6) and (7). The dependent variables are average monthly purchases in each sub-category from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Estimates have been normalized by each sub-category's share of baseline spending on its credit line (mean sub-category purchases over the 30 days ending October 1, 2025 divided by mean purchases on that line). Whiskers represent 95% confidence intervals. $N = 209{,}996$. Full estimates are reported in Appendix Table A13.

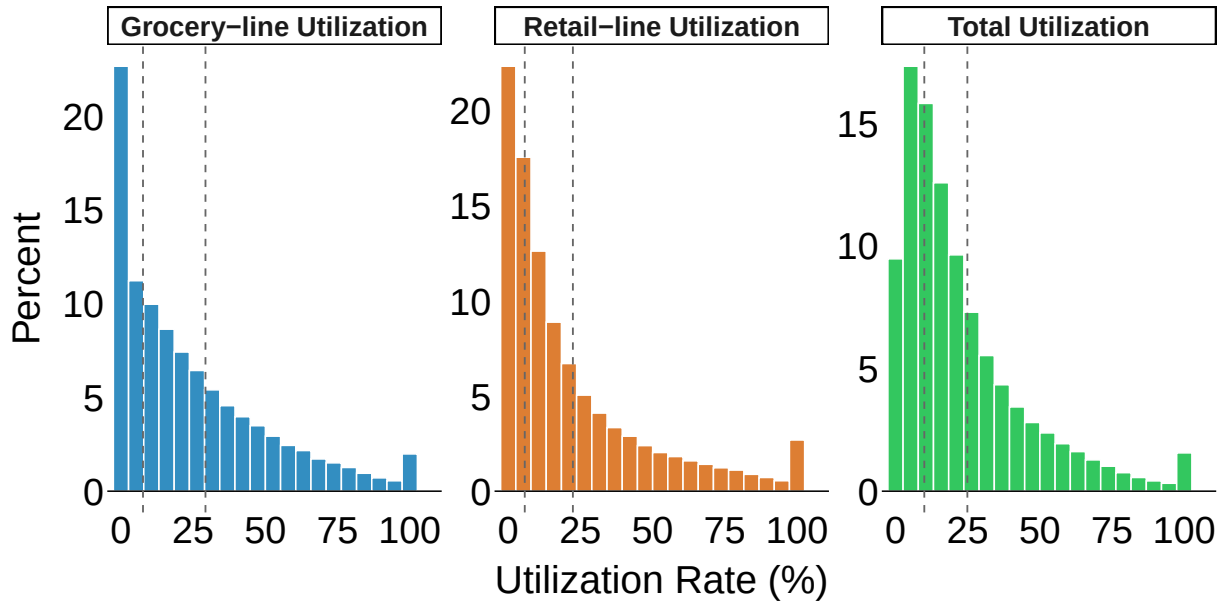
A Appendix Exhibits

Figure A1. Credit-Limit Intervention over Experimental Period



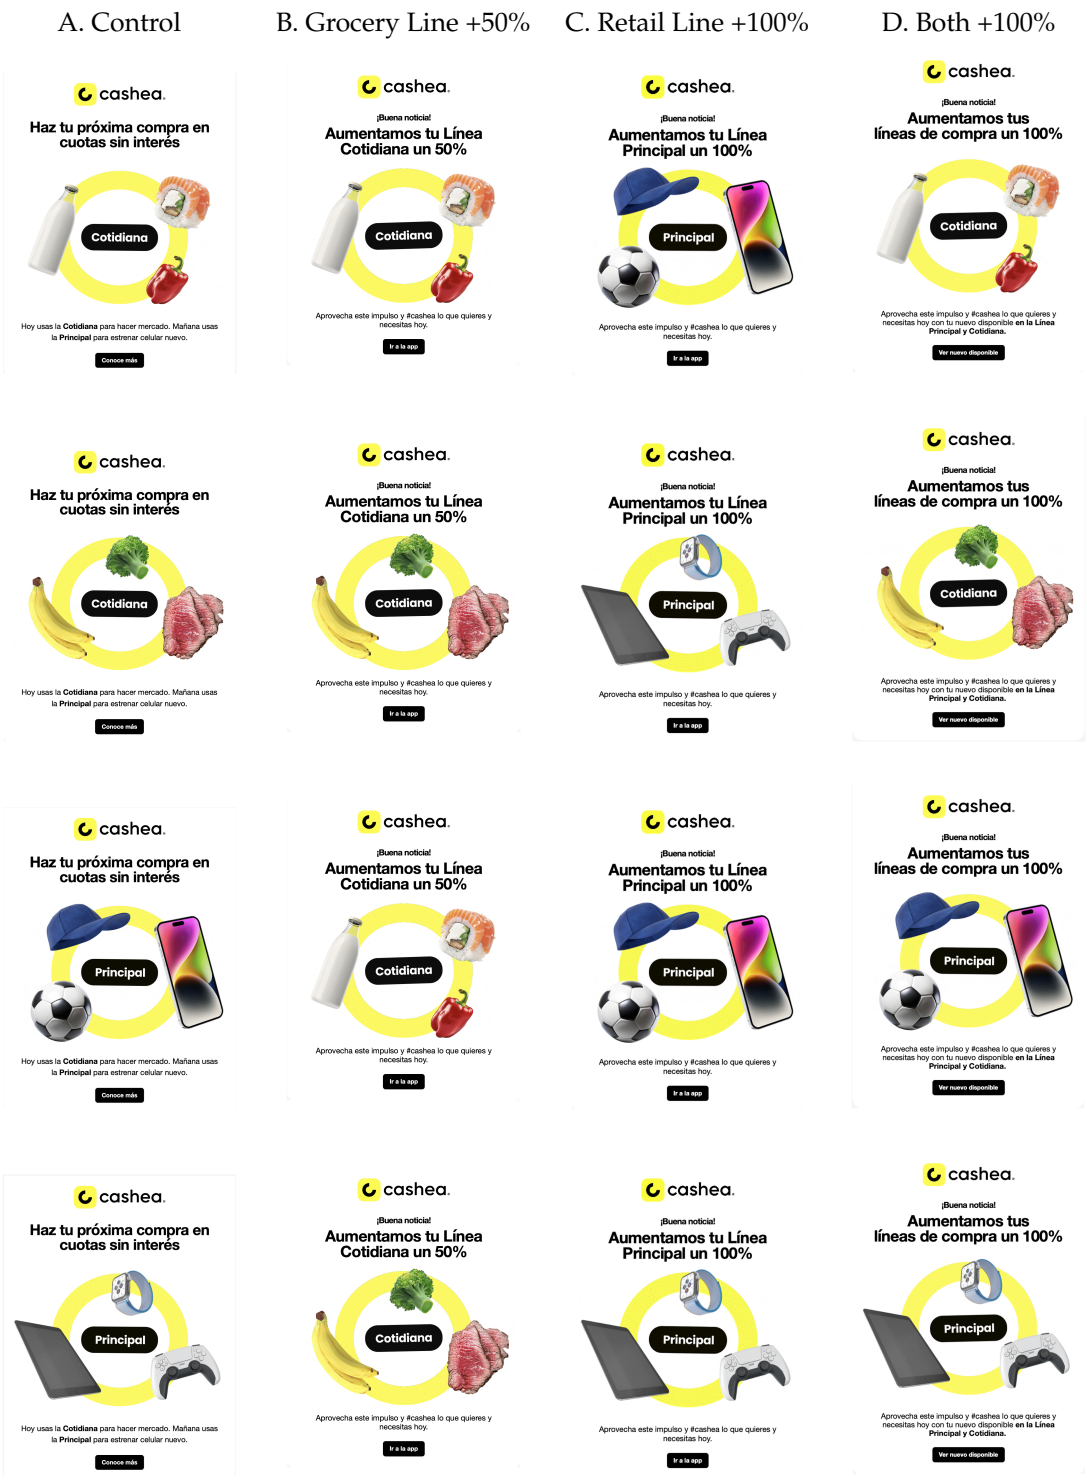
Notes: Bars show mean credit limits in USD by treatment arm and credit line; numbers above bars report the total. Whiskers show 95% confidence intervals (heteroskedasticity-robust) for each arm's mean change in limits. "G" denotes grocery line and "R" denotes retail line. Green bars reflect the baseline mean limit on each line as of October 1, 2025, pooled across all arms. Orange bars reflect the change in limits as of October 2, 2026, one year after the experiment launched. $N = 209{,}988$.

Figure A2. Distribution of Baseline Utilization



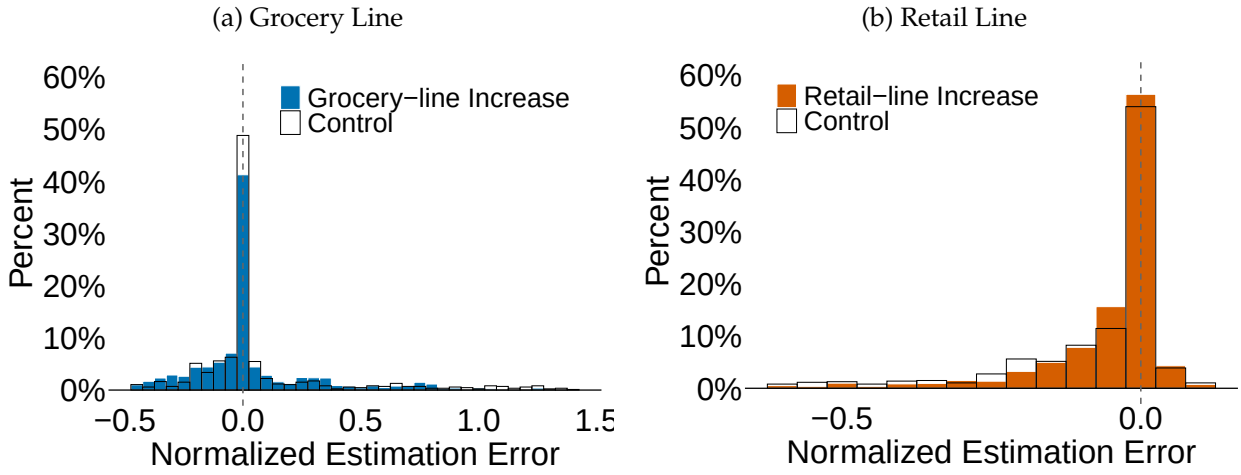
Notes: The figure plots the distribution of baseline credit utilization rates across users in the experimental sample, separately for the grocery line, the retail line, and the combined total. Dashed vertical lines mark tercile cutoffs. Utilization is defined as the 30-day average utilization rate for each line as of October 1, 2025. $N = 209{,}996$.

Figure A3. Full Outreach Messaging



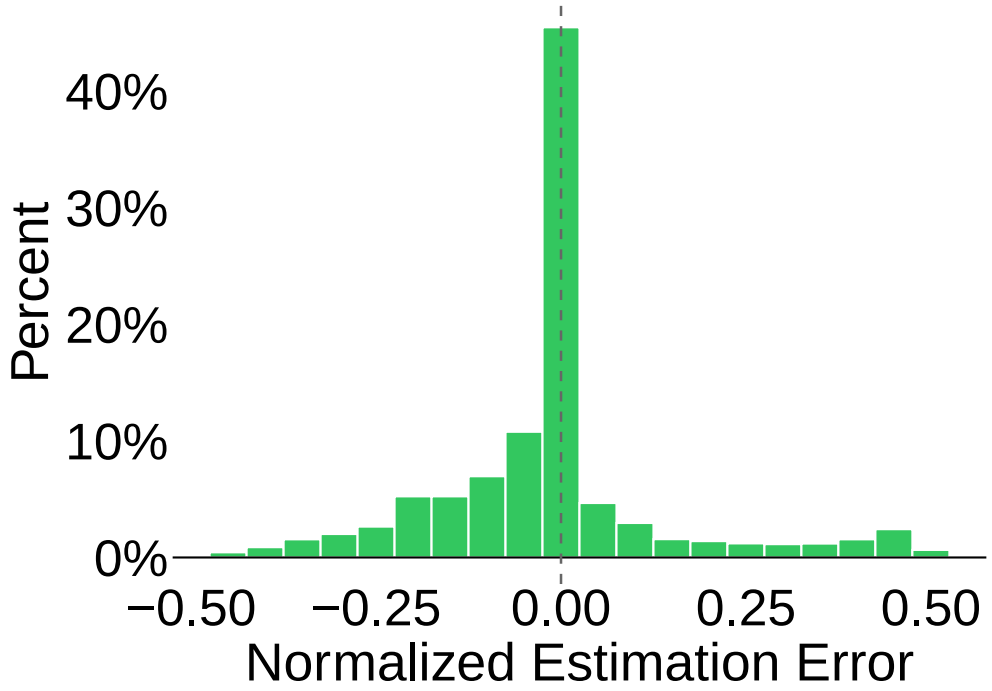
Notes: The figure displays screenshots of email outreach sent to participants, showing four of seven arms as examples. Panel A shows the control-group messages, which promote Cashea without mentioning a credit-limit increase. Panel B shows the treatment messages for a 50% grocery-line increase. Panel C shows treatment messages for a 100% retail-line increase. Panel D shows the treatment messages for a 100% increase on both lines. All treatment and control arms received analogous outreach. The message contained a GIF animation that cycled through the images shown.

Figure A4. Knowledge of Post-Experiment Credit Limits



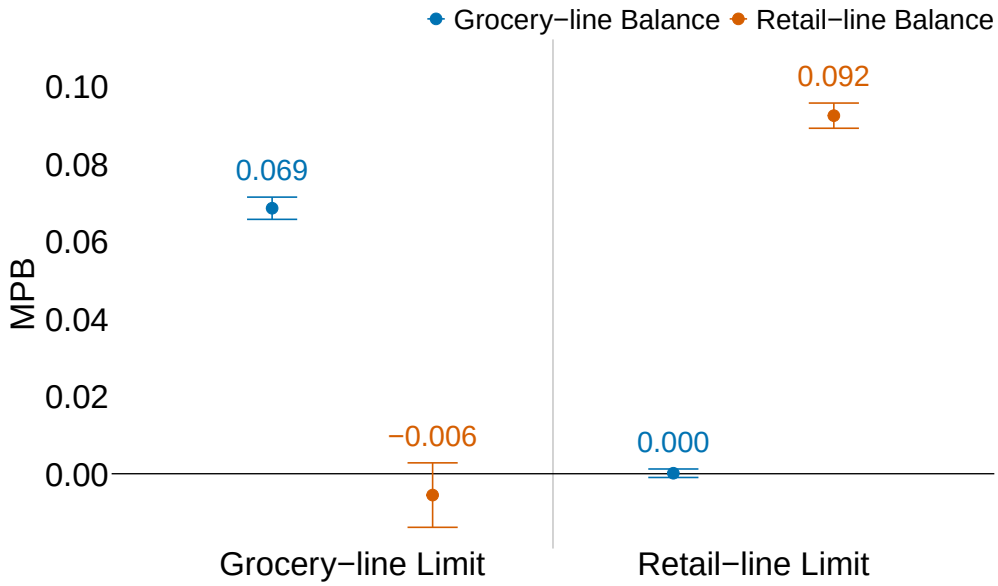
Notes: The figure plots the distribution of normalized credit-limit estimation error in the first follow-up survey, defined as $(\hat{L}_i^{\ell} - L_i^{\ell})/L_i^{\ell}$, where $\hat{L}_i^{\ell}$ is the respondent's self-reported limit on line ℓ and L_i^{ℓ} is the corresponding administrative limit on the date the respondent started the survey. Each panel is trimmed at the 5th and 95th percentiles of the pooled treatment and control sample; bin width is 0.05. $N = 4{,}229$ for the grocery-line panel and $N = 4{,}422$ for the retail-line panel.

Figure A5. Baseline Total Limit Estimation Error



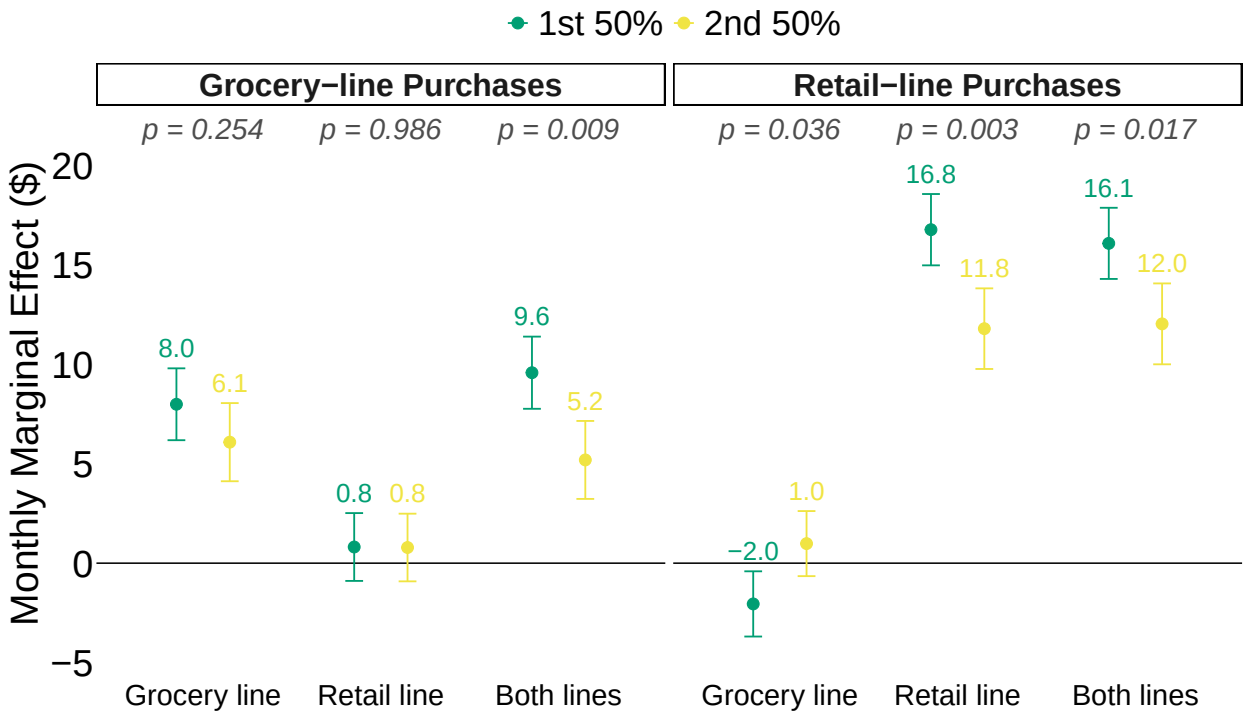
Notes: The figure plots the distribution of normalized credit-limit estimation error in the baseline survey, defined as $(\hat{L}_i - L_i)/L_i$, where $\hat{L}_i$ is the respondent's self-reported total credit limit summing across both lines, and L_i is the corresponding administrative limit recorded on October 1, 2025. The distribution is trimmed at the 5th and 95th percentiles; bin width is 0.05. $N = 7{,}709$.

Figure A6. Marginal Propensities to Borrow out of Credit-Limit Increases



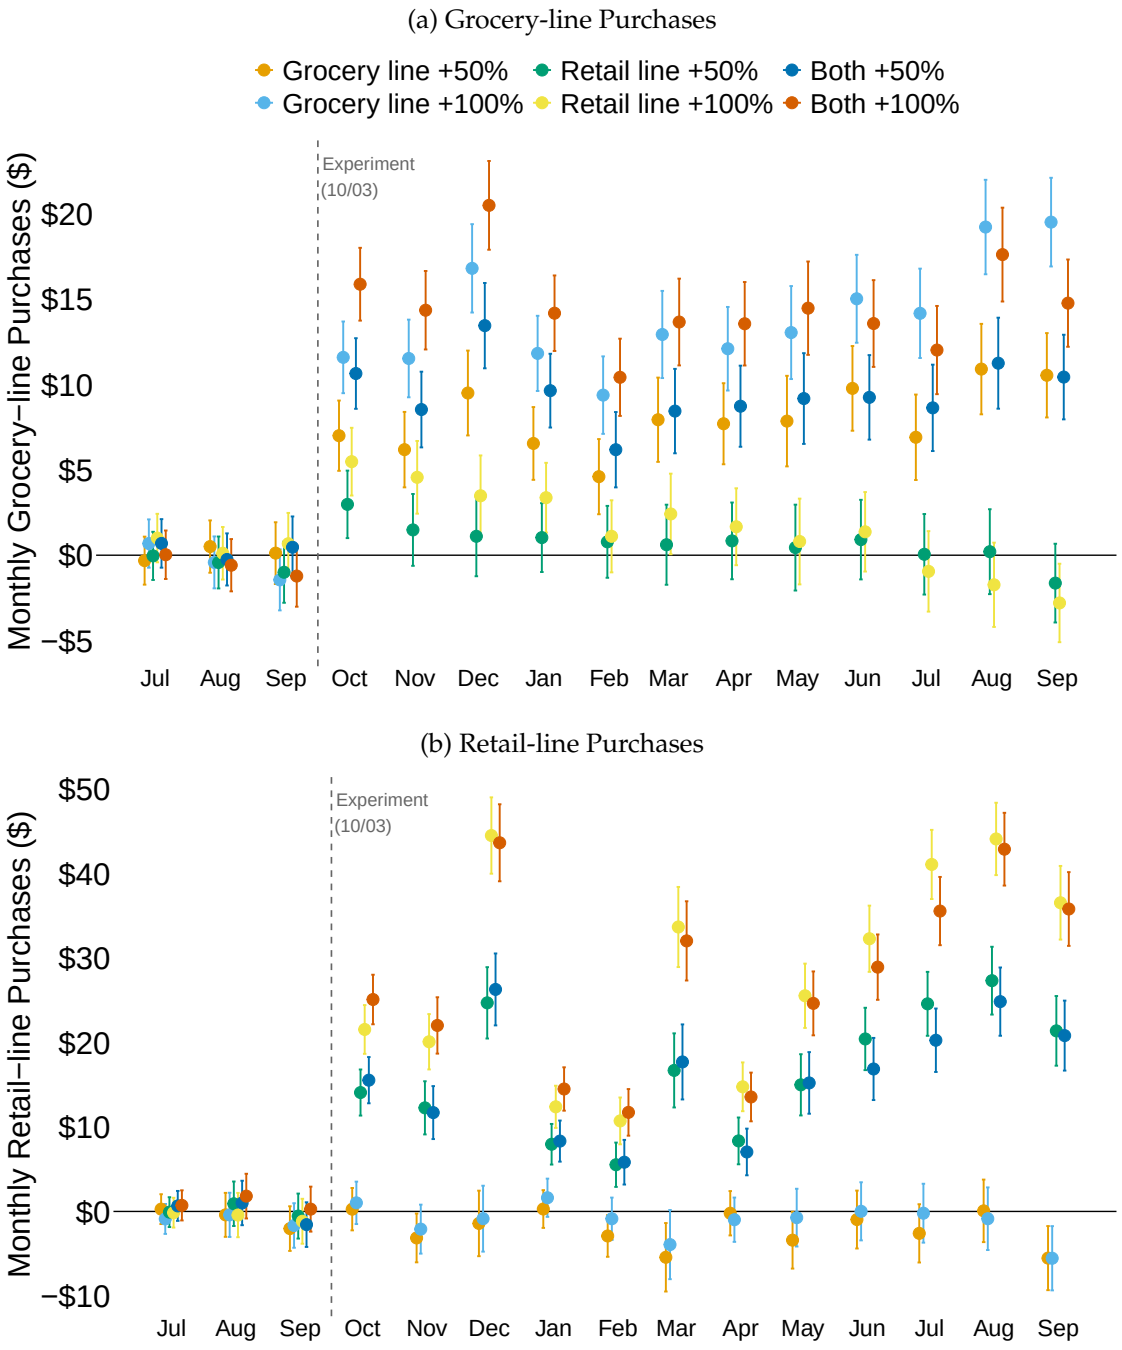
Notes: The figure plots the instrumental-variables estimates of the marginal propensity to borrow out of credit-limit increases on each line, estimated via equations (6) and (7). The dependent variable is average daily balances on each line from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Post-treatment credit limits are measured on October 6, 2025, and are also winsorized at the 0.1st and 99.9th percentiles. Whiskers represent 95% confidence intervals. $N = 209{,}996$. Full estimates are reported in Appendix Table A14.

Figure A7. Marginal Treatment Effects



Notes: The figure plots treatment effect estimates of the marginal effect of the first 50% limit increase (τ_{50}) and the second 50% increase ($\tau_{100} - \tau_{50}$) on each line's purchases, by which line's limit was increased, estimated via equation (8). The dependent variable is average monthly purchases on each line from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. The p-value annotated above each pair of estimates tests $H_0 : \tau_{50} = \tau_{100} - \tau_{50}$. Whiskers represent 95% confidence intervals. $N = 209{,}996$.

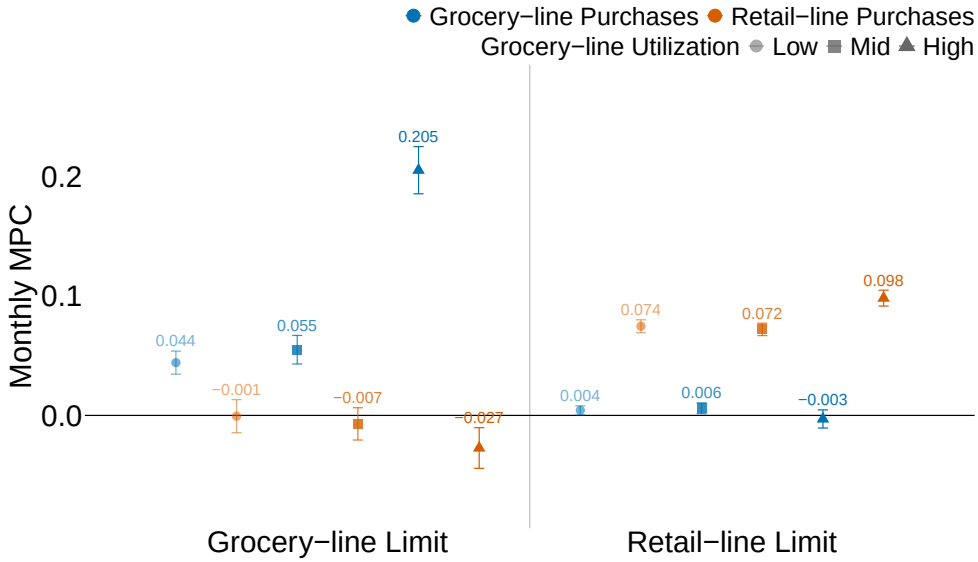
Figure A8. Dynamic Effects by Treatment Arm on Purchases



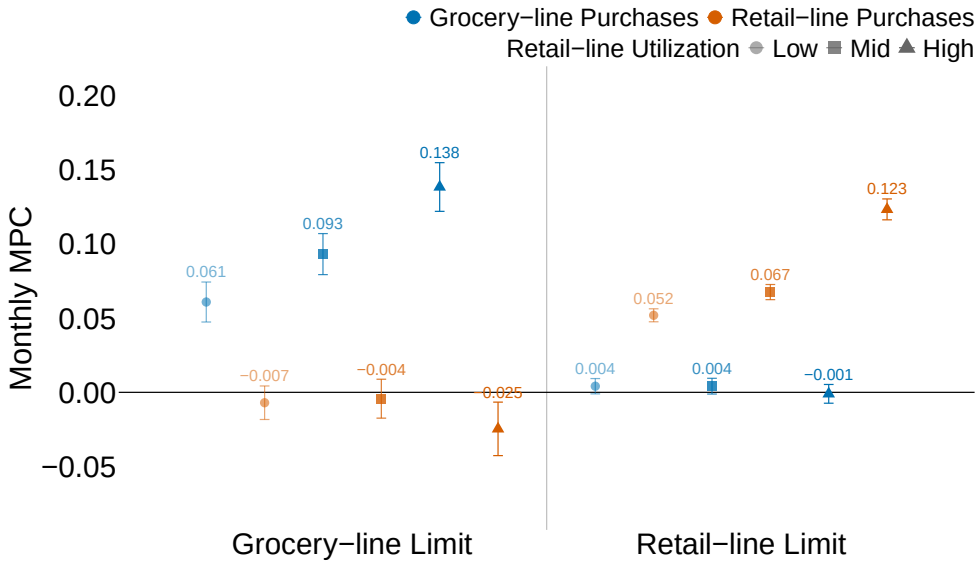
Notes: The figure plots monthly treatment effect estimates of credit-limit increases on each line, estimated via equation (8) separately for each calendar month. The dependent variables are monthly purchases on each line from July 2025 through September 2026, winsorized at the 0.1st and 99.9th percentiles within each month. Whiskers represent 95% confidence intervals. $N = 209{,}996$.

Figure A9. Marginal Propensity to Consume by Baseline Grocery- and Retail-line Utilization

(a) By Grocery-line Utilization Tercile

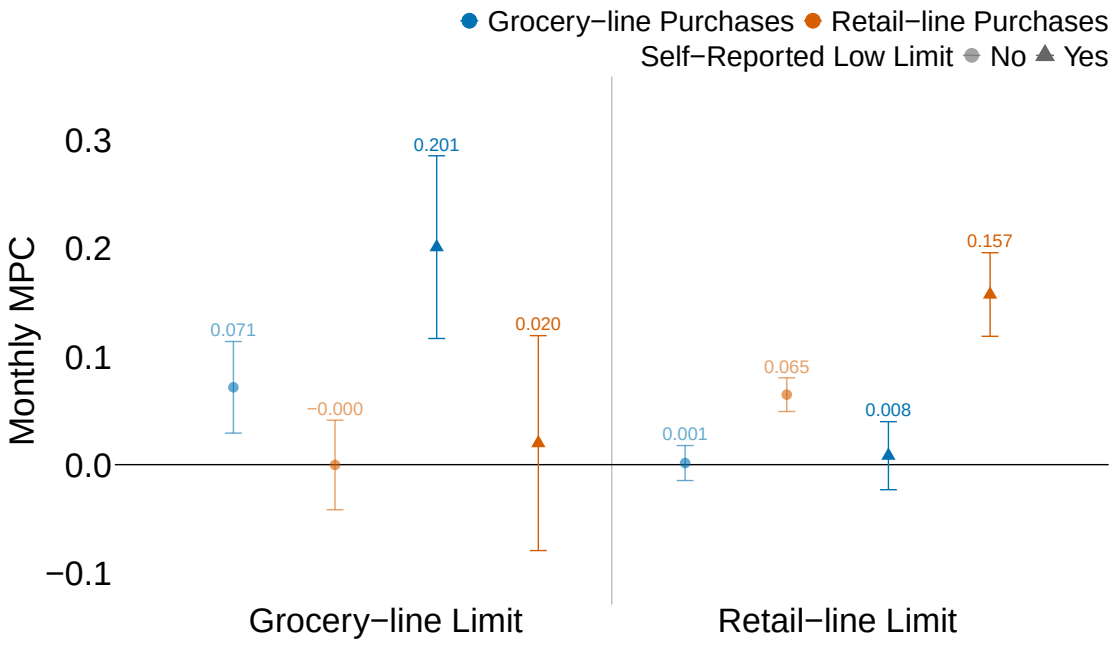


(b) By Retail-line Utilization Tercile



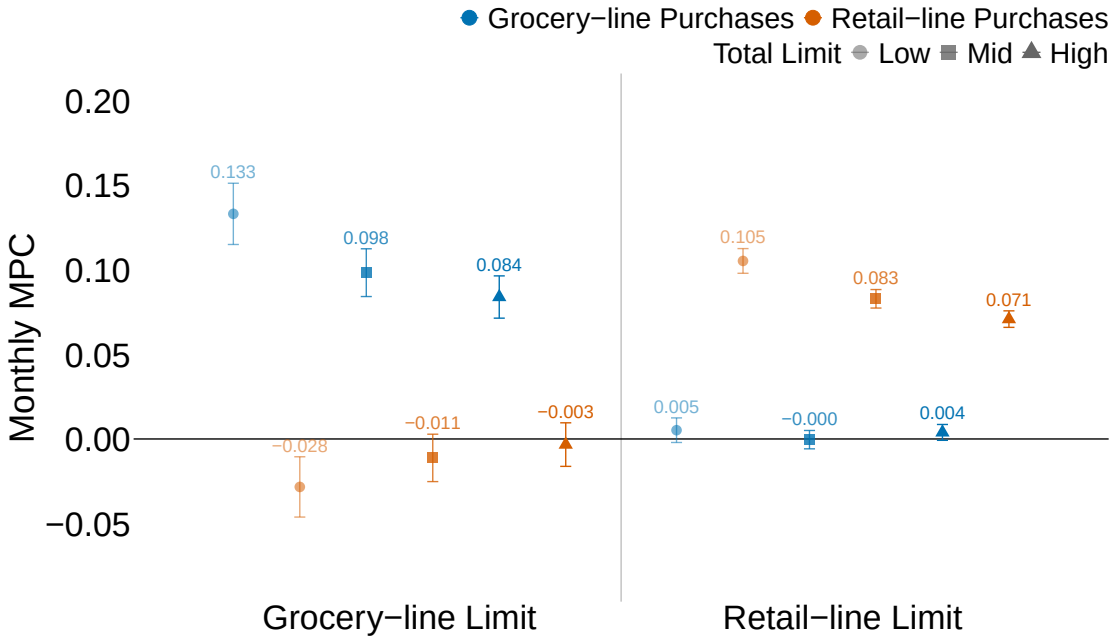
Notes: The figure plots instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line by tercile of baseline utilization on each line, estimated via versions of equations (6) and (7) in which every term except strata fixed effects is interacted with tercile indicators. The dependent variable is average monthly purchases on each line from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Whiskers represent 95% confidence intervals. We conduct joint Wald tests of equality across the three tercile coefficients. For cross-line effects in panel (a): $p = 0.139$ for retail-line limit $\rightarrow$ grocery-line purchases and $p = 0.053$ for grocery-line limit $\rightarrow$ retail-line purchases. In panel (b): $p = 0.385$ for retail-line limit $\rightarrow$ grocery-line purchases and $p = 0.174$ for grocery-line limit $\rightarrow$ retail-line purchases. $N = 209{,}996$.

Figure A10. Marginal Propensity to Consume by Self-Reported Credit Constraint



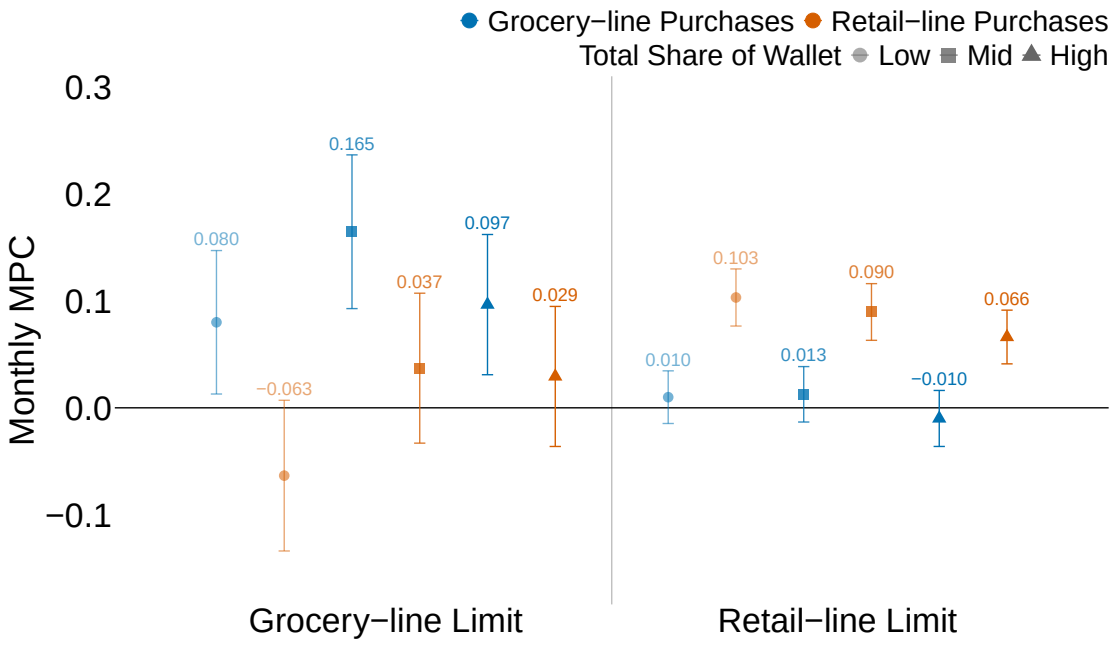
Notes: The figure plots instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line by self-reported credit constraint status in the baseline survey, estimated via versions of equations (6) and (7) in which every term except strata fixed effects is interacted with binary indicators. The dependent variable is average monthly purchases on each line from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Whiskers represent 95% confidence intervals. $N = 8{,}728$.

Figure A11. Marginal Propensity to Consume by Baseline Credit Limit



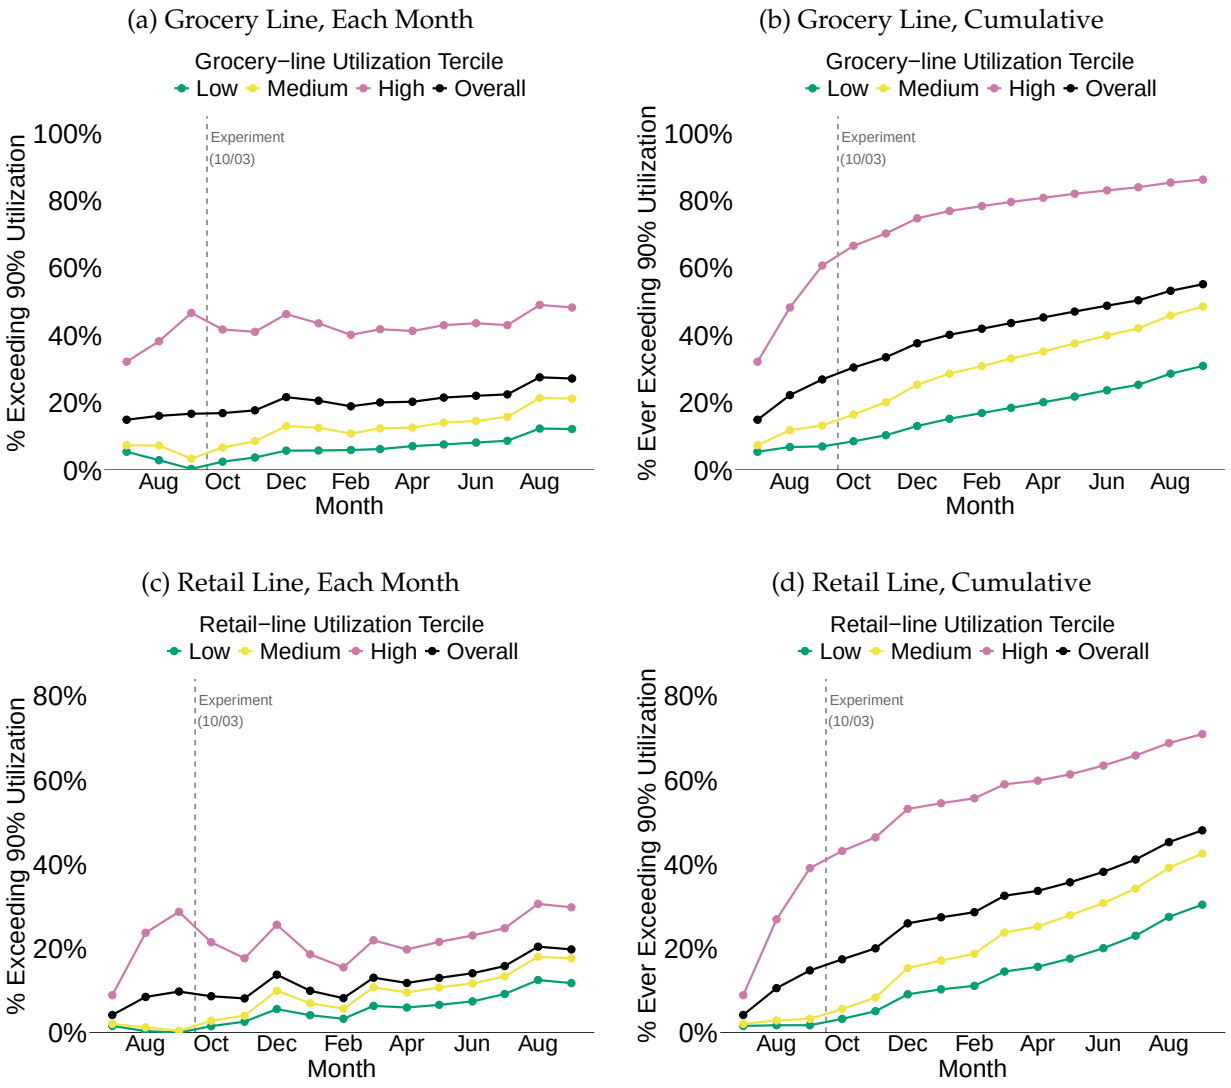
Notes: The figure plots instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line by tercile of total credit limit (grocery line plus retail line) on October 1, 2025, estimated via versions of equations (6) and (7) in which every term except strata fixed effects is interacted with tercile indicators. The dependent variable is average monthly purchases on each line from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Whiskers represent 95% confidence intervals. $N = 209{,}996$.

Figure A12. Marginal Propensity to Consume by Baseline Cashea Share of Wallet



Notes: The figure plots instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line by baseline Cashea share of wallet, estimated via versions of equations (6) and (7) in which every term except strata fixed effects is interacted with tercile indicators. Share of wallet is defined as total Cashea spending as a fraction of total spending on and off Cashea in the grocery and retail categories, and computed from self-reported spending over the previous 30 days in the baseline survey. The dependent variable is average monthly purchases on each line from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Whiskers represent 95% confidence intervals. $N = 8{,}711$.

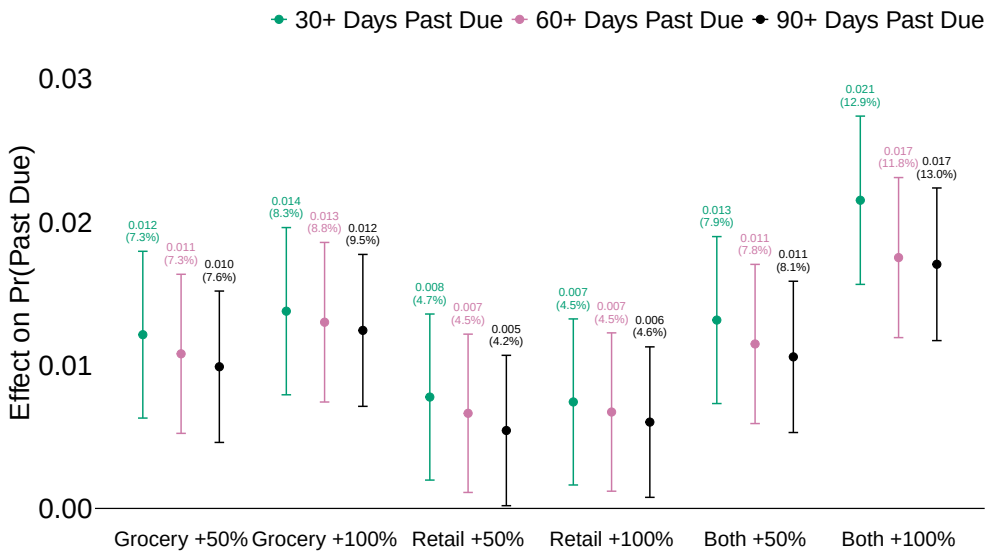
Figure A13. Monthly Constraint-Hitting Rates by Baseline Utilization



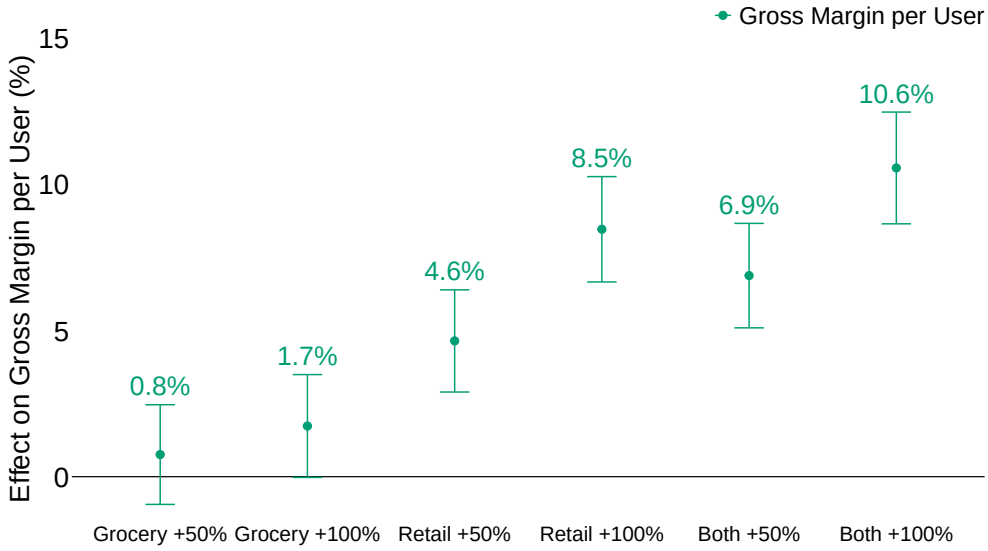
Notes: The figure plots the share of control-group users whose utilization on each line reached 90% on at least one day within each calendar month, among users observed in that month (panels (a) and (c)) and the share whose utilization had ever reached 90% since July 2025, among all control-group users observed at any point from July 2025 through September 2026 (panels (b) and (d)), by tercile of average utilization on that line over the 30 days ending October 1, 2025. $N = 30{,}001$.

Figure A14. Treatment Effects by Arm on Installments Past Due and Monthly Gross Margin

(a) Installments Past Due



(b) Gross Margin per User



Notes: The figure plots treatment effect estimates of credit-limit increases on each line, estimated via equation (8). In panel (a), the dependent variables are indicators for holding at least one installment 30 or more, 60 or more, or 90 or more days past due as of October 2, 2026. In panel (b), the dependent variable is average monthly gross margin per user from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Estimates and confidence intervals in panel (b) are reported as a percent of the control-group mean. Whiskers represent 95% confidence intervals. $N = 209{,}996$.

Table A1. Baseline Spending by Credit Line

Category	% Total $	% Total N
Panel A: Retail line		
Clothing	31.0%	46.4%
Home	35.3%	25.3%
Services	8.9%	8.7%
Tools	6.4%	5.9%
Sports	4.1%	4.6%
Health	4.4%	4.3%
Technology	8.7%	3.6%
Other	1.2%	1.2%
Total	100.0%	100.0%
Panel B: Grocery line		
Groceries	82.2%	75.0%
Pharmacies	12.5%	13.6%
Restaurants	4.2%	8.6%
Other	1.1%	2.8%
Total	100.0%	100.0%

Notes: The table reports the distribution of purchases across store categories on each line in the 90 days before the experiment (July 5 to October 2, 2025). *% Total $* is each category's share of purchase value on its line and *% Total N* is its share of the number of orders. $N = 209{,}996$.

Table A2. Summary Statistics and Balance Table

	Control mean	Grocery line +50%	Grocery line +100%	Retail line +50%	Retail line +100%	Both lines +50%	Both lines +100%	Joint F-statistic
Account age (days)	393.41	-1.10	0.37	0.90	0.79	-0.48	-0.86	0.62
		(1.42)	(1.42)	(1.42)	(1.42)	(1.42)	(1.42)	[0.717]
Total credit limit ($)	517.72	-0.81	-0.84	-2.07	-0.17	-0.12	-0.40	0.82
		(1.11)	(1.11)	(1.11)	(1.11)	(1.11)	(1.11)	[0.553]
Grocery-line limit ($)	144.01	-0.14	-0.19	-0.59	-0.03	0.05	-0.11	0.97
		(0.30)	(0.30)	(0.30)	(0.30)	(0.30)	(0.30)	[0.440]
Retail-line limit ($)	373.71	-0.66	-0.64	-1.48	-0.14	-0.17	-0.29	0.76
		(0.83)	(0.83)	(0.83)	(0.83)	(0.83)	(0.83)	[0.605]
Total purchases ($)	227.05	-2.37	-3.30	-1.90	-1.04	-1.14	-1.70	0.70
		(1.78)	(1.78)	(1.78)	(1.78)	(1.78)	(1.78)	[0.651]
Grocery-line purchases ($)	115.06	0.17	-1.50	-1.13	0.29	0.33	-1.34	1.55
		(0.94)	(0.94)	(0.94)	(0.94)	(0.94)	(0.94)	[0.157]
Retail-line purchases ($)	111.99	-2.54	-1.81	-0.77	-1.34	-1.47	-0.36	0.79
		(1.39)	(1.39)	(1.39)	(1.39)	(1.39)	(1.39)	[0.580]
Total balance ($)	115.90	-0.90	-0.75	-0.30	-0.81	-0.04	0.94	1.16
		(0.85)	(0.85)	(0.85)	(0.85)	(0.85)	(0.85)	[0.325]
Grocery-line balance ($)	33.59	0.15	-0.12	-0.27	0.11	0.15	-0.22	1.22
		(0.23)	(0.23)	(0.23)	(0.23)	(0.23)	(0.23)	[0.291]
Retail-line balance ($)	82.31	-1.05	-0.63	-0.02	-0.91	-0.19	1.16	1.71
		(0.80)	(0.80)	(0.80)	(0.80)	(0.80)	(0.80)	[0.114]
Omnibus F-statistic		0.82	0.65	0.78	0.35	0.39	0.76	
		[0.568]	[0.714]	[0.601]	[0.932]	[0.909]	[0.618]	
N	30,001	29,999	29,997	30,001	29,997	30,000	30,001	

Notes: The table reports the control-group mean and the coefficients from a regression of each baseline variable on treatment-arm indicators (control omitted) and strata fixed effects, with standard errors in parentheses and p-values in brackets. The joint *F*-statistic tests that the six arm coefficients, from a regression of the variable on arm indicators and strata fixed effects, are jointly zero. The omnibus *F*-statistic tests that the coefficients from a regression of an indicator for assignment to that arm (relative to control) on account age and the six line-specific characteristics, with strata fixed effects, are jointly zero, separately for each arm; the three totals are excluded because they are sums of the line-specific values. Credit limits and account age are measured as of October 1, 2025; purchases are summed over the 30 days ending October 1, 2025 and balances are average daily outstanding balances over the 30 days ending October 1, 2025.

Table A3. Experimental Setup

		Baseline Mean, Oct 1 ($)		Intended Increase		Realized Δ, Oct 1–Oct 6 ($)	
Treatment Arm	N	Retail line	Grocery line	Retail line	Grocery line	Retail line	Grocery line
Control	30,001	373.71	144.01	—	—	2.04	0.79
Grocery line +50%	29,999	373.06	143.87	—	+50%	2.26	71.44
Grocery line +100%	29,997	373.07	143.82	—	+100%	2.47	141.83
Retail line +50%	30,001	372.22	143.42	+50%	—	185.06	0.87
Retail line +100%	29,997	373.57	143.99	+100%	—	369.38	0.80
Both +50%	30,000	373.54	144.06	+50%	+50%	185.04	71.46
Both +100%	30,001	373.42	143.90	+100%	+100%	369.70	142.50

Notes: The table summarizes the credit-limit intervention by treatment arm. *Baseline Mean* reports average credit limits on each line as of October 1, 2025, *Intended Increase* reports the assigned percentage increase on each line, and *Realized* Δ reports the average change in each limit between October 1 and October 6, 2025.

Table A4. Predictors of Follow-Up Survey Response

	Non-survey mean	Survey	Joint F-statistic
Account age (days)	389.91	14.75	146.96
		(1.22)	[0.000]
Total credit limit ($)	507.58	30.37	874.78
		(1.03)	[0.000]
Grocery-line limit ($)	141.35	8.25	912.14
		(0.27)	[0.000]
Retail-line limit ($)	366.24	22.12	836.71
		(0.76)	[0.000]
Total purchases ($)	221.07	19.25	160.70
		(1.52)	[0.000]
Grocery-line purchases ($)	111.82	14.89	347.52
		(0.80)	[0.000]
Retail-line purchases ($)	109.25	4.35	13.06
		(1.20)	[0.000]
Total balance ($)	114.41	4.06	29.03
		(0.75)	[0.000]
Grocery-line balance ($)	33.37	0.03	0.03
		(0.19)	[0.866]
Retail-line balance ($)	81.04	4.03	31.35
		(0.72)	[0.000]
Omnibus F-statistic		205.39	
		[0.000]	
N	185,292	24,704	

Notes: The table compares baseline characteristics of users who responded to either follow-up survey against non-respondents. The joint *F*-statistic tests that the coefficient on the respondent indicator, from a regression of the variable on a respondent indicator and strata fixed effects, is zero, using heteroskedasticity-robust standard errors. The omnibus *F*-statistic tests that the coefficients from a regression of the survey-response indicator on account age and the six line-specific characteristics, with strata fixed effects, are jointly zero; the three totals are excluded because they are sums of the line-specific values. Credit limits and account age are measured as of October 1, 2025; purchases are summed over the 30 days ending October 1, 2025 and balances are average daily outstanding balances over the 30 days ending October 1, 2025.

Table A5. Follow-Up Survey Subsample Balance Table

	Control mean	Grocery line +50%	Grocery line +100%	Retail line +50%	Retail line +100%	Both lines +50%	Both lines +100%	Joint F-statistic
Account age (days)	421.56	2.45	7.19	2.46	-2.80	-1.10	-5.89	1.76
		(4.61)	(4.59)	(4.58)	(4.47)	(4.64)	(4.52)	[0.102]
Total credit limit ($)	592.46	1.76	4.01	1.57	0.93	6.29	1.54	0.62
		(3.93)	(3.84)	(3.87)	(3.77)	(3.86)	(4.19)	[0.717]
Grocery-line limit ($)	163.73	0.64	1.18	0.43	0.25	1.86	0.39	0.78
		(1.05)	(1.01)	(1.02)	(1.00)	(1.03)	(1.10)	[0.589]
Retail-line limit ($)	428.72	1.12	2.83	1.14	0.68	4.43	1.15	0.55
		(2.92)	(2.87)	(2.90)	(2.81)	(2.87)	(3.13)	[0.769]
Total purchases ($)	261.05	-7.79	0.29	2.06	5.77	-2.92	2.25	1.28
		(5.47)	(5.64)	(5.57)	(5.67)	(5.53)	(5.71)	[0.263]
Grocery-line purchases ($)	137.29	0.30	-0.41	0.06	2.37	-1.46	1.92	0.42
		(2.98)	(2.99)	(2.99)	(2.92)	(2.94)	(2.95)	[0.867]
Retail-line purchases ($)	123.77	-8.09	0.71	2.00	3.40	-1.46	0.34	1.57
		(4.23)	(4.38)	(4.43)	(4.45)	(4.32)	(4.45)	[0.152]
Total balance ($)	126.92	0.84	1.74	0.98	-0.48	-0.05	-0.24	0.14
		(2.90)	(3.12)	(2.96)	(2.69)	(2.83)	(2.78)	[0.990]
Grocery-line balance ($)	35.32	-0.12	0.34	-0.49	-0.07	-0.29	0.03	0.28
		(0.70)	(0.72)	(0.70)	(0.68)	(0.70)	(0.69)	[0.948]
Retail-line balance ($)	91.60	0.96	1.39	1.48	-0.41	0.24	-0.26	0.16
		(2.77)	(2.99)	(2.83)	(2.58)	(2.69)	(2.66)	[0.987]
Omnibus F-statistic		0.85	0.62	0.27	0.39	0.74	0.37	
		[0.550]	[0.738]	[0.965]	[0.909]	[0.636]	[0.918]	
N (person-wave)	3,754	3,966	3,955	4,198	4,310	4,067	4,261	

Notes: The table tests for balance across treatment arms among respondents to either follow-up survey, pooling the two waves (one observation per respondent per wave). The joint *F*-statistic tests that the six arm coefficients are jointly zero; this test and the omnibus test use standard errors clustered by user. The omnibus *F*-statistic tests that the coefficients from a regression of an indicator for assignment to that arm (relative to control) on account age and the six line-specific characteristics, with strata fixed effects, are jointly zero, separately for each arm; the three totals are excluded because they are sums of the line-specific values. Credit limits and account age are measured as of October 1, 2025; purchases are summed over the 30 days ending October 1, 2025 and balances are average daily outstanding balances over the 30 days ending October 1, 2025.

Table A6. Demographic Characteristics

Variable	N	Mean	SD	P1	P25	P50	P75	P99
Monthly Income ($)	8,728	616.77	518.98	5.27	250.00	480.00	800.00	2,000.00
Share Female	22,678	0.56	0.50	0.00	0.00	1.00	1.00	1.00
Age (Survey)	22,794	40.50	9.98	21.00	33.00	40.00	47.00	65.00
Age (Admin)	208,863	40.76	10.87	21.00	33.00	40.00	48.00	68.00

Notes: The table reports demographic characteristics of survey respondents and the full administrative sample (age). Monthly income is from the baseline survey only, and is winsorized at the 95th percentile. Age and share female use each respondent's most recent response across the baseline and two follow-up surveys; share female excludes respondents not reporting male or female.

Table A7. Marginal Propensities to Consume: On- and Off-Cashea Spending

	DV: Average Monthly Purchases			
	Survey		Admin	
	Grocery line	Retail line	Grocery line	Retail line
Panel A: Off Cashea				
Grocery-line Limit (Post)	0.006	-0.029		
	(0.024)	(0.024)		
Retail-line Limit (Post)	0.010	0.018		
	(0.009)	(0.009)		
Observations	23,541	23,541		
Panel B: On Cashea				
Grocery-line Limit (Post)	0.065	-0.035	0.084	-0.007
	(0.019)	(0.025)	(0.004)	(0.005)
Retail-line Limit (Post)	0.014	0.050	0.007	0.062
	(0.007)	(0.010)	(0.002)	(0.002)
Observations	23,541	23,541	209,996	209,996
Panel C: Total				
Grocery-line Limit (Post)	0.071	-0.064		
	(0.037)	(0.043)		
Retail-line Limit (Post)	0.025	0.068		
	(0.014)	(0.016)		
Observations	23,541	23,541		

Notes: The table reports instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line, estimated via equations (6) and (7). In the Survey columns, the dependent variables are self-reported spending outcomes from a pooled sample of first- and second-wave follow-up respondents, normalized to their 30-day equivalent and winsorized at the 95th percentile. In the administrative columns, the dependent variables are on-Cashea purchases on each line from administrative data for the full experimental sample, measured over windows matching the survey recall periods (October 1, 2025 to the median first-wave response date, December 4, 2025, normalized to 30 days; and the 30 days ending on the median second-wave response date, May 14, 2026), averaged across the two windows and winsorized at the 0.1st and 99.9th percentiles.

Table A8. Other Borrowing / Savings Behavior

	Control mean	Grocery line +50%	Grocery line +100%	Retail line +50%	Retail line +100%	Both lines +50%	Both lines +100%	Joint F-statistic [RW p]	N
Monthly income ($)	780.22	33.98	31.23	28.10	38.41	30.03	81.28	1.82	11,782
		(25.13)	(24.75)	(24.45)	(24.58)	(25.13)	(25.21)	[0.497]	
Other spending ($)	468.66	-6.58	20.28	1.87	16.47	36.12	21.25	1.53	9,021
		(16.98)	(17.58)	(17.13)	(17.28)	(17.54)	(17.41)	[0.655]	
Savings stock ($)	1,256.33	55.16	74.95	-96.30	82.37	-30.49	187.49	1.54	7,968
		(106.92)	(106.70)	(102.33)	(106.35)	(103.72)	(110.10)	[0.655]	
Less savings, stock (binary)	0.36	0.04	0.04	0.00	0.01	0.02	0.00	1.19	7,978
		(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	[0.757]	
Savings flow, last 30 days ($)	168.69	4.18	0.44	4.92	0.57	2.36	4.93	0.07	7,653
		(11.86)	(11.49)	(11.61)	(11.34)	(11.60)	(11.68)	[0.998]	
Less savings, flow (binary)	0.38	0.05	0.03	0.01	0.01	0.02	0.00	1.26	7,446
		(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	[0.757]	
Other debt ($)	131.69	15.64	14.57	19.93	15.80	11.48	30.05	1.24	7,574
		(11.14)	(11.12)	(11.19)	(11.08)	(11.03)	(11.54)	[0.757]	
Transfers/loans to others ($)	59.05	2.62	3.14	3.50	10.12	3.62	12.49	1.19	5,211
		(5.50)	(5.52)	(5.45)	(5.64)	(5.52)	(5.79)	[0.757]	

Notes: The table reports treatment effect estimates of credit-limit increases on self-reported outcomes from the second follow-up survey, estimated via equation (8). *Other spending* is spending on categories not covered by either line. *Less savings* indicators equal one if the respondent reports that their stock or flow of savings has decreased since September 2025. Continuous outcomes are winsorized at the 95th percentile. Standard errors, clustered at the user level, are reported in parentheses. The joint F-statistic tests that all six arm coefficients are zero; brackets report Romano–Wolf stepdown p-values adjusting for the eight outcomes (9,999 wild cluster bootstrap replications imposing the null). Sample sizes differ across rows because of item nonresponse; in particular, income was asked first and is observed for respondents who did not complete the survey, and transfers/loans to others was added to the questionnaire after the email invitation but before the push-notification reminder; neither varies with treatment assignment.

Table A9. MPC out of Credit-Limit Increases

	DV: Average Monthly Purchases through October 2, 2026		
	Retail-line Purchases	Grocery-line Purchases	Total Purchases
Retail-line Limit (Post)	0.081	0.002	0.083
	(0.002)	(0.002)	(0.003)
Grocery-line Limit (Post)	-0.011	0.098	0.087
	(0.004)	(0.004)	(0.007)
Observations	209,996	209,996	209,996

Notes: The table reports instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line, estimated via equations (6) and (7). The dependent variable is average monthly purchases on each line and in total from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Post-treatment credit limits are measured on October 6, 2025, and are also winsorized at the 0.1st and 99.9th percentiles. Standard errors, clustered at the user level, are reported in parentheses.

Table A10. Treatment Effects of Credit-Limit Increases

	DV: Average Monthly Purchases through October 2, 2026		
	Retail-line Purchases	Grocery-line Purchases	Total Purchases
Grocery line +50%	-2.05	7.98	5.95
	(0.84)	(0.92)	(1.49)
Grocery line +100%	-1.06	14.06	13.12
	(0.84)	(0.96)	(1.53)
Retail line +50%	16.75	0.81	17.63
	(0.91)	(0.87)	(1.51)
Retail line +100%	28.53	1.60	30.19
	(0.97)	(0.87)	(1.55)
Both +50%	16.07	9.57	25.74
	(0.91)	(0.92)	(1.56)
Both +100%	28.09	14.75	42.92
	(0.98)	(0.95)	(1.63)
Strat Bin FE	Yes	Yes	Yes
SE Clustered	Yes	Yes	Yes
Control Mean	126.74	122.90	249.71
Observations	209,996	209,996	209,996

Notes: The table reports treatment effect estimates of credit-limit increases on each line, estimated via equation (8). The dependent variable is average monthly purchases on each line and in total from October 3, 2025 to October 2, 2026, each winsorized separately at the 0.1st and 99.9th percentiles. Standard errors, clustered at the user level, are reported in parentheses.

Table A11. MPC Estimates by Month

	Panel A: Grocery-line (LC) Limit Increase		*Panel B: Retail-line (LP) Limit Increase*	
Month	Grocery-line Purchases	Retail-line Purchases	Grocery-line Purchases	Retail-line Purchases
Oct 2025	0.079	0.009	0.013	0.064
	(0.005)	(0.007)	(0.002)	(0.003)
Nov 2025	0.078	-0.006	0.010	0.063
	(0.005)	(0.008)	(0.002)	(0.003)
Dec 2025	0.123	-0.011	0.009	0.124
	(0.006)	(0.010)	(0.002)	(0.004)
Jan 2026	0.083	0.009	0.007	0.035
	(0.005)	(0.006)	(0.002)	(0.002)
Feb 2026	0.066	-0.001	0.003	0.034
	(0.005)	(0.006)	(0.002)	(0.002)
Mar 2026	0.088	-0.020	0.003	0.099
	(0.006)	(0.011)	(0.002)	(0.004)
Apr 2026	0.086	-0.010	0.003	0.040
	(0.006)	(0.007)	(0.002)	(0.003)
May 2026	0.096	-0.011	0.002	0.073
	(0.006)	(0.009)	(0.002)	(0.003)
Jun 2026	0.098	-0.023	-0.002	0.085
	(0.006)	(0.009)	(0.002)	(0.004)
Jul 2026	0.096	-0.031	-0.004	0.108
	(0.006)	(0.009)	(0.002)	(0.004)
Aug 2026	0.136	-0.019	-0.005	0.121
	(0.006)	(0.010)	(0.002)	(0.004)
Sep 2026	0.134	-0.030	-0.011	0.110
	(0.006)	(0.010)	(0.002)	(0.004)
Observations	209,996	209,996	209,996	209,996

Notes: The table reports monthly instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line, estimated via equations (6) and (7). The dependent variables are monthly purchases on each line from October 2025 through September 2026, winsorized at the 0.1st and 99.9th percentiles within each month. Standard errors, clustered at the user level, are reported in parentheses.

Table A12. Average and Heterogeneous Treatment Effects

	DV: Average Monthly Purchases through October 2, 2026			
	Grocery-line Limit Increase		Retail-line Limit Increase	
	Grocery-line Purchases	Retail-line Purchases	Grocery-line Purchases	Retail-line Purchases
ATE	11.26	-1.06	1.24	23.53
	[0.000]	[0.250]	[0.180]	[0.000]
	(9.26, 13.26)	(-2.86, 0.74)	(-0.57, 3.03)	(21.41, 25.64)
HTE	0.736	0.215	0.327	0.766
	[0.000]	[0.067]	[0.006]	[0.000]
	(0.576, 0.882)	(-0.011, 0.437)	(0.095, 0.544)	(0.635, 0.899)
Observations	87,578	87,578	87,631	87,631
Control Mean	122.90	126.74	122.90	126.74

Notes: The table reports Best Linear Predictor estimates from the machine-learning heterogeneity procedure of Chernozhukov, Demirer, Duflo, and Fernández-Val (2025). Columns 1 and 2 restrict the sample to the control group and the two grocery-line arms (+50% and +100%); columns 3 and 4 to the control group and the two retail-line arms. The dependent variable is average monthly purchases on each line from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Estimates are medians over 100 data splits; brackets report p-values and parentheses report medians across splits of the 95% confidence bounds. Quintile-level estimates are plotted in Figure 8.

Table A13. Marginal Propensities to Consume by Merchant Sub-Category

	DV: Monthly Purchases by Merchant Sub-Category / Baseline Share	
	Grocery-line Limit (Post)	Retail-line Limit (Post)
Panel A: Grocery-line Categories		
Groceries	0.093	0.002
	(0.005)	(0.002)
Pharmacies	0.106	0.001
	(0.013)	(0.005)
Restaurants	0.123	0.013
	(0.013)	(0.005)
Observations	209,996	209,996
Panel B: Retail-line Categories		
Home	-0.014	0.103
	(0.006)	(0.003)
Clothing	-0.013	0.040
	(0.005)	(0.002)
Services	-0.007	0.054
	(0.009)	(0.004)
Technology	-0.013	0.193
	(0.018)	(0.007)
Tools	0.001	0.091
	(0.014)	(0.005)
Health	-0.002	0.050
	(0.011)	(0.004)
Sports	-0.006	0.047
	(0.010)	(0.004)
Observations	209,996	209,996

Notes: The table reports instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line by sub-category, estimated via equations (6) and (7). The dependent variables are average monthly purchases in each sub-category from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Estimates are normalized by each sub-category's share of baseline spending on its line (mean sub-category purchases over the 30 days ending October 1, 2025 divided by mean total purchases on the line). Standard errors, clustered at the user level, are reported in parentheses.

Table A14. MPB out of Credit-Limit Increases

	DV: Average Balance through October 2, 2026		
	Retail-line Balance	Grocery-line Balance	Total Balance
Retail-line Limit (Post)	0.092	0.000	0.093
	(0.002)	(0.001)	(0.002)
Grocery-line Limit (Post)	-0.006	0.069	0.063
	(0.004)	(0.001)	(0.005)
Observations	209,996	209,996	209,996

Notes: The table reports instrumental-variables estimates of the marginal propensity to borrow out of credit-limit increases on each line, estimated via equations (6) and (7). The dependent variable is average daily balances on each line and in total from October 3, 2025 to October 2, 2026, winsorized at the 0.1st and 99.9th percentiles. Post-treatment credit limits are measured on October 6, 2025, and are also winsorized at the 0.1st and 99.9th percentiles. Standard errors, clustered at the user level, are reported in parentheses.

B Model Details and Proofs

This appendix characterizes the equilibrium of the model described in Section 3 and includes formal statements of any technical assumptions, model predictions, and their proofs.

The consumer's problem. The consumer solves the following optimization problem:

$$\begin{aligned} \max_{\substack{\{c_{Gt}\},\,\{c_{Rt}\},\\ \{a_{Gt}\},\,\{a_{Rt}\}}} \quad & \mathbb{E}_{t=1}\left[\sum_{t=1}^{3}\beta^{t-1}U(c_{Gt},c_{Rt})\right] \\ \text{subject to} \quad & (c_{Gt}-a_{Gt})+(c_{Rt}-a_{Rt})\leq m_t \;\forall t, \\ & 0\leq a_{\ell t}\leq \min\left\{L_\ell,\; c_{\ell t}\right\} \text{ for } t\in\{1,2\} \text{ and } \ell\in\{G,R\}, \\ & a_{\ell 3}=0 \text{ for } \ell\in\{G,R\}, \end{aligned} \tag{A.1}$$

where $\beta > 0$ is the discount factor, $U(c_{Gt}, c_{Rt})$ is the instantaneous utility in period t, $m_t \equiv y_t + x_t - (1+r)\,(a_{G,t-1} + a_{R,t-1})$ denotes the net resources available in period t, and $x_{t+1} = m_t - [(c_{Gt} - a_{Gt}) + (c_{Rt} - a_{Rt})]$ is the law of motion for savings. Both credit limits are strictly positive ($L_G, L_R > 0$). Throughout, r is the interest rate on credit-line borrowing.

We model mental accounting following Farhi and Gabaix (2019):

$$U(c_{Gt}, c_{Rt}) = u(c_{Gt}, c_{Rt}) - \gamma_G\,|c_{Gt} - c^*_{Gt}| - \gamma_R\,|c_{Rt} - c^*_{Rt}|\,,$$

where $c^*_{\ell t}$ is the consumption target for category ℓ (described in more detail below) and γ_ℓ is the psychological cost of deviating from that target. We model the rational benchmark without mental accounting by setting $\gamma_G = \gamma_R = 0$.

Before a credit-limit increase, we denote the targets $c^*_{\ell t}$ as $\bar{c}_{\ell t}$. We treat an increase in the category-ℓ limit as a salient cue about category ℓ consumption (Bordalo, Gennaioli, and Shleifer, 2020), and assume that the consumer responds by solving for new category-ℓ targets $c^*_{\ell t}$ across all t and all states holding the category-k targets fixed at $\bar{c}_{kt}$. They do so because re-optimizing the full two-category problem is mentally costly and because the category-ℓ limit increase makes them think about category-ℓ consumption due to similarity-based attention (Bordalo, Gennaioli, Lanzani, and Shleifer, 2026).

Some consumers nevertheless solve the full re-optimization problem and deviate from these new targets if the cost of deviating is lower than the benefit of fully re-optimizing.

Timing. The timing is as follows. (i) Before the cue, the mental targets $\bar{c}_{\ell t}$ and $\bar{c}_{kt}$ of consumption equal the pre-cue plan, which the user had set by solving their optimization problem prior to receiving the credit-limit increase. (ii) At the beginning of period 1, the user experiences a credit-limit increase on line $\ell \in \{G, R\}$, i.e., L_ℓ increases. (iii) After the credit-limit increase, the consumer determines new consumption targets $c^*_{\ell t}$ for each time period t and each state by holding $\bar{c}_{kt}$ fixed for the other line $k \neq \ell$ and optimizing their discounted expected consumption utility as of period 1, subject to credit and budget constraints. (iv) The consumer then solves (A.1) at the new credit limit with these new consumption targets.

Assumption 1 (Preferences)**.** *Preferences over state-contingent consumptions are represented by time-additive expected utility, as specified in equation* (A.1)*. Utility $u : \mathbb{R}^2_{++} \to \mathbb{R}$ is C^2, strictly increasing and strictly concave, and represents a homothetic preference over deterministic within-period consumption bundles. Thus, the consumer is strictly risk averse over lotteries of within-period consumption bundles.*

Assumption 2 (Income states)**.** *Period-1 income y_1 and period-3 income $y_3 > y_1$ are certain, but period-2 income is risky:*

$$y_2 = \begin{cases} y_1 & \text{in the good state, with probability } p \\ y_1 - \delta & \text{in the bad state, with probability } 1 - p, \end{cases}$$

with $p \in (0, 1)$ and $\delta \in (0, y_1)$.

Assumption 3 (No static constraint in period 1)**.** *The period-1 optimum is interior, i.e., $a_{\ell 1} < L_\ell$ for $\ell \in \{G, R\}$.*

Assumption 4 (Regularity)**.** *All consumption levels are strictly positive; the good-state limits are strictly slack in period 2; in the bad state each limit is either strictly slack or strictly binding; and the binding regime is locally unchanged by the comparative static.*

Prediction 1. ***Without*** *mental accounting ($\gamma_G = \gamma_R = 0$), an increase in the category-ℓ credit limit will either not affect current consumption in* ***either*** *category or increase current consumption in* ***both*** *categories.*

More precisely, in response to an increase in L_ℓ, current consumption will not change in either category if the category-ℓ credit limit is slack in the bad state, while consumption will increase in both categories if the category-ℓ credit limit binds in the bad state.

Proof. Fix a category $\ell \in \{G, R\}$, and let $k \neq \ell$ denote the other category. Formally, we seek to

prove that under Assumptions 1–4,

$$\frac{\partial \hat{c}_{G1}}{\partial L_\ell}, \quad \frac{\partial \hat{c}_{R1}}{\partial L_\ell} \begin{cases} > 0, & \text{if line } \ell \text{ is strictly binding in the bad state,} \\ = 0, & \text{if line } \ell \text{ is slack in the bad state.} \end{cases}$$

Thus, absent a current static constraint (Assumption 3), a marginal increase in one category's credit limit can never raise current consumption in only that category. The proof collapses this problem into a period-1 choice between current consumption and a continuation value (Steps 1–2) and then signs the cross-derivative of the continuation value with respect to the treated limit (Step 3).

By homotheticity and concavity of u, there exist H and Φ such that $u = \Phi\left(H(c_G, c_R)\right)$, where H is increasing, concave, and degree-one homogeneous, and Φ is increasing and concave.

Define total consumption in period 1 as $C_1 \equiv c_{G1} + c_{R1}$ and define indirect utility as

$$v(C_1) \equiv \max_{c_{G1}, c_{R1}} u(c_{G1}, c_{R1}) \text{ s.t. } c_{G1} + c_{R1} = C_1.$$

Because Φ is strictly increasing, the allocation that maximizes u also maximizes H. Since u is strictly concave, the optimal allocation is unique; homogeneity of H then implies that it is proportional to C_1: $c_q(C_1) = \omega_q C_1$ for $q \in \{G, R\}$. Writing $\kappa \equiv H(\omega_G, \omega_R) > 0$, we obtain $v(C_1) = \Phi(\kappa C_1)$ and $v''(C_1) = \kappa^2 \Phi''(\kappa C_1) < 0$. Therefore, for each $q \in \{G, R\}$,

$$\frac{\partial c_q(C_1)}{\partial C_1} = \omega_q > 0.$$

In other words, an increase in total period-1 consumption implies increases in current consumption in both categories. Since

$$\frac{\partial c_q(C_1)}{\partial L_\ell} = \frac{\partial c_q(C_1)}{\partial C_1} \frac{\partial C_1}{\partial L_\ell},$$

to show that $\partial c_q(C_1)/\partial L_\ell > 0$ whenever L_ℓ binds in the bad state of period 2, it is sufficient to show that $\partial C_1/\partial L_\ell > 0$ whenever L_ℓ binds in the bad state of period 2.

Step 1: Continuation value.

Define b as: $b \equiv x_2 - (1+r)(a_{G1} + a_{R1})$. For given period-2 cash-on-hand $m_2 \equiv y_2 + x_2 - (1+r)(a_{G1} + a_{R1})$, define the period-2-and-3 value function

$$\begin{aligned} V(m_2; L_G, L_R) \equiv \max_{c_{G2}, c_{R2}, a_{G2}, a_{R2}, x_3 \geq 0} & \; u(c_{G2}, c_{R2}) + \beta v(x_3 + y_3 - (1+r)(a_{G2} + a_{R2})) \\ \text{s.t.} & \; (c_{G2} - a_{G2}) + (c_{R2} - a_{R2}) \leq m_2, \\ & \; 0 \leq a_{q2} \leq \min\{L_q, c_{q2}\} \text{ for } q \in \{G, R\}. \end{aligned} \tag{A.2}$$

Since period-2 income is $y_2 = y_1$ in the good state and $y_2 = y_1 - \delta$ in the bad state by Assumption 2, and since $m_2 = y_2 + b$, we can define the ex-ante continuation value as of period 1 as

$$W(b; L_G, L_R) = pV(y_1 + b; L_G, L_R) + (1-p)V(y_1 - \delta + b; L_G, L_R). \tag{A.3}$$

Since (A.2) is a concave program with m entering linearly, V and W are concave in b. Thus,

$$W_{bb} \leq 0$$

wherever the value function V is twice differentiable.[20]

Step 2: Choose C_1 and b in period 1 to maximize expected utility.

Let period-1 total borrowing be $A_1 \equiv a_{G1} + a_{R1}$. The period-1 budget constraint implies $x_2 = y_1 - C_1 + A_1$. Hence net liquidity carried into period 2 is

$$b = x_2 - (1+r)A_1 = y_1 - C_1 - rA_1. \tag{A.4}$$

Equation (A.4) defines b as a function of C_1, which we will now write as $b(C_1)$. For a fixed C_1, the consumer chooses the smallest feasible A_1, because current utility is unchanged while continuation utility is increasing in b. Under Assumption 3, $A_1(C_1) = \max\{0, C_1 - y_1\}$. Away from the kink where $C_1 = y_1$,

$$b(C_1) = \begin{cases} y_1 - C_1, & C_1 < y_1, \\ (1+r)(y_1 - C_1), & C_1 > y_1. \end{cases}$$

Thus, $b'(C_1) < 0$ and $b''(C_1) = 0$. Under Assumption 3, the period-1 problem can therefore be written as

$$\max_{C_1} \left\{ v(C_1) + \beta W(b(C_1); L_G, L_R) \right\}.$$

Its first-order condition is

$$v'(C_1) + \beta W_b(b(C_1); L_G, L_R) b'(C_1) = 0. \tag{A.5}$$

Let $\hat{C}_1$ denote the value of C_1 that solves equation (A.5). Then, defining D_1 as the derivative of the

[20] To see this, take m, m' and let z, z' denote choice vectors attaining $V(m)$ and $V(m')$. For $\lambda \in [0,1]$, the bundle $\lambda z + (1-\lambda)z'$ is feasible at cash-on-hand $\lambda m + (1-\lambda)m'$, since every constraint in (A.2) is linear in the choices and m (the budget constraint) or independent of m and linear in the choices ($a_{q2} \geq 0$, $a_{q2} \leq L_q$, $a_{q2} \leq c_{q2}$, $x_3 \geq 0$). The objective is concave in the choices, as u is concave and v is concave and composed with a linear function. Hence $V(\lambda m + (1-\lambda)m') \geq \lambda V(m) + (1-\lambda)V(m')$, so V is concave in m_2; W is then a convex combination of concave functions of b and is concave, so $W_{bb} \leq 0$ wherever it is twice differentiable.

left-hand side of equation (A.5) with respect to C_1,

$$D_1 = v''(\hat{C}_1) + \beta W_{bb}(b(\hat{C}_1); L_G, L_R)[b'(\hat{C}_1)]^2 + \beta W_b(b(C_1); L_G, L_R)b''(C_1) < 0,$$

where the last term equals zero since $b''(C_1) = 0$. Implicit differentiation gives

$$\frac{\partial \hat{C}_1}{\partial L_\ell} = -\frac{\beta b'(\hat{C}_1) W_{b,L_\ell}}{D_1}. \tag{A.6}$$

Since $b' < 0$ and $D_1 < 0$,

$$\operatorname{sgn}\left(\frac{\partial \hat{C}_1}{\partial L_\ell}\right) = -\operatorname{sgn}(W_{b,L_\ell}).$$

Therefore, to show that $\partial \hat{C}_1 / \partial L_\ell > 0$, it is sufficient that $W_{b,L_\ell} < 0$.

Step 3: Determine the sign of W_{b,L_ℓ}.

Since $m_2 = y_1 + b$ in the good state and $m_2 = y_1 - \delta + b$ in the bad state,

$$W_{b,L_\ell} = pV_{m_2,L_\ell}(y_1 + b; L_G, L_R) + (1-p)V_{m_2,L_\ell}(y_1 - \delta + b; L_G, L_R)\,.$$

Under the maintained regime, line ℓ is strictly slack in the good state in period 2, so the good-state optimum and value are locally unaffected by L_ℓ:

$$V_{m_2,L_\ell}(y_1 + b; L_G, L_R) = 0.$$

Consequently,

$$W_{b,L_\ell} = (1-p)V_{m_2,L_\ell}(y_1 - \delta + b; L_G, L_R)\,.$$

It therefore suffices to determine the sign of V_{m_2,L_ℓ} in the bad state under three exhaustive cases: (i) both lines bind in period 2, (ii) only line ℓ binds, (iii) line ℓ is slack, regardless of whether line k binds.

Case 1: Both lines bind. Suppose both lines bind, i.e., $a_{G2} = L_G < c_{G2}$ and $a_{R2} = L_R < c_{R2}$. Because $r > 0$, borrowing and positive saving cannot coexist, so when both lines bind $x_3 = 0$. The period-2 resource constraint when both lines bind is $c_{G2} + c_{R2} = m_2 + L_G + L_R$. Because $c_q > L_q$ for both categories (Assumption 4), the composition of consumption is not constrained. In the bad state, it follows that

$$V(m_2; L_G, L_R) = v(m_2 + L_G + L_R) + \beta v(y_3 - (1+r)(L_G + L_R))\,.$$

Thus for either category $q \in \{G, R\}$,

$$V_{m_2,L_q}(m_2; L_G, L_R) = v''(m_2 + L_G + L_R) < 0. \tag{A.7}$$

Case 2: Only line ℓ binds.

In this case,

$$a_\ell = L_\ell < c_\ell, \qquad a_k < L_k.$$

The Lagrangian of the continuation problem (A.2) is

$$\begin{aligned}\mathcal{L} = u(c_{G2}, c_{R2}) + \beta v(m_3) + \lambda \Big[m_2 - (c_{G2} - a_{G2}) - (c_{R2} - a_{R2}) - x_3 \Big] \\ + \sum_{q \in \{G,R\}} \Big[\eta_q \left(L_q - a_{q2}\right) + \zeta_q \, a_{q2} + \mu_q \left(c_{q2} - a_{q2}\right) \Big] + \nu x_3,\end{aligned} \tag{A.8}$$

where $m_3 = x_3 + y_3 - (1+r)(a_{G2} + a_{R2})$, λ is the multiplier on the period-2 budget constraint, and $\eta_q, \zeta_q, \mu_q, \nu \geq 0$ are the multipliers on $a_q \leq L_q$, $a_q \geq 0$, the accounting constraint $a_q \leq c_q$, and $x_3 \geq 0$, respectively. Since all consumption levels are strictly positive, the first-order conditions for c_{q2} and a_{q2} hold with equality:

$$\partial \mathcal{L} / \partial c_{q2} = u_q - \lambda + \mu_q = 0, \tag{A.9}$$

$$\partial \mathcal{L} / \partial a_{q2} = \lambda - \beta(1+r)\, v'(m_3) - \eta_q + \zeta_q - \mu_q = 0, \tag{A.10}$$

$$\partial \mathcal{L} / \partial x_3 = \beta v'(m_3) - \lambda + \nu = 0. \tag{A.11}$$

Adding (A.9) and (A.10), the multiplier μ_q on the accounting constraint cancels, leaving

$$u_q = \beta(1+r) v'(m_3) + \eta_q - \zeta_q, \qquad m_3 = x_3 + y_3 - (1+r)(a_{G2} + a_{R2}). \tag{A.12}$$

For the strictly binding line, $a_\ell = L_\ell > 0$, so $\eta_\ell > 0$ and $\zeta_\ell = 0$. For the other line, $\eta_k = 0$ and $\zeta_k \geq 0$. Hence

$$u_\ell > \beta(1+r) v'(m_3) \geq u_k,$$

and therefore

$$u_\ell > u_k. \tag{A.13}$$

Moreover, all consumption of category k must be credit-financed. If

$$c_k - a_k > 0,$$

one could shift a small amount of cash spending from category k to category ℓ, leaving debt and all intertemporal resources unchanged, and raise utility by (A.13). Hence

$$a_k = c_k < L_k. \tag{A.14}$$

The same no simultaneous saving and borrowing argument used above gives

$$x_3 = 0.$$

Equation (A.14) and the period-2 resource constraint imply

$$c_\ell = m_2 + L_\ell.$$

Since V is locally independent of L_k on this regime (A.14), we suppress L_k and write $V(m_2; L_\ell)$. The continuation problem therefore reduces to

$$V(m_2; L_\ell) = \max_{c_k > 0} \left\{ u(m_2 + L_\ell, c_k) + \beta v(y_3 - (1+r)(L_\ell + c_k)) \right\}. \tag{A.15}$$

We first establish that $u_{\ell k} > u_{\ell\ell}$, which will be needed to sign V_{m_2, L_ℓ} below. Since $u = \Phi \circ H$, we can differentiate u to obtain $u_{\ell k}$ and $u_{\ell\ell}$. By subtracting them,

$$u_{\ell k} - u_{\ell\ell} = \Phi''(H) H_\ell (H_k - H_\ell) + \Phi'(H)(H_{\ell k} - H_{\ell\ell}). \tag{A.16}$$

Because $\Phi' > 0$, condition (A.13) is equivalent to

$$H_\ell > H_k.$$

Since $\Phi'' < 0$, the first term in (A.16) is strictly positive.

Degree-one homogeneity of H implies

$$c_\ell H_{\ell\ell} + c_k H_{\ell k} = 0.$$

Therefore, concavity of H gives

$$\begin{aligned} H_{\ell k} - H_{\ell\ell} &= -\left(1 + \frac{c_\ell}{c_k}\right) H_{\ell\ell} \\ &\geq 0. \end{aligned}$$

It follows that

$$u_{\ell k} > u_{\ell \ell}. \tag{A.17}$$

Thus, the Hessian restriction we need is implied by risk aversion over the aggregate, convex homothetic intratemporal preferences, and the one-bind KKT condition (A.13).

Let $\hat{c}_k(m_2, L_\ell)$ solve (A.15), and define

$$\hat{m}_3 = y_3 - (1+r)(L_\ell + \hat{c}_k), \qquad \Gamma \equiv \beta(1+r)^2 v''(\hat{m}_3) < 0.$$

The first-order condition for c_k is

$$u_k(m_2 + L_\ell, \hat{c}_k) - \beta(1+r)v'(\hat{m}_3) = 0.$$

Implicit differentiation gives

$$\frac{\partial \hat{c}_k}{\partial L_\ell} = -\frac{u_{\ell k} + \Gamma}{u_{kk} + \Gamma}. \tag{A.18}$$

The envelope theorem gives

$$V_{m_2} = u_\ell.$$

Combining this with (A.18),

$$\begin{aligned} V_{m_2, L_\ell} &= u_{\ell\ell} + u_{\ell k}\frac{\partial \hat{c}_k}{\partial L_\ell} \\ &= \frac{u_{\ell\ell}u_{kk} - u_{\ell k}^2 + \Gamma(u_{\ell\ell} - u_{\ell k})}{u_{kk} + \Gamma}. \end{aligned} \tag{A.19}$$

The function $u = \Phi \circ H$ is concave, so

$$u_{\ell\ell}u_{kk} - u_{\ell k}^2 \geq 0.$$

Moreover,

$$u_{kk} + \Gamma < 0,$$

while (A.17) and $\Gamma < 0$ imply

$$\Gamma(u_{\ell\ell} - u_{\ell k}) > 0.$$

The numerator of (A.19) is therefore strictly positive and its denominator is strictly negative. Hence

$$V_{m_2, L_\ell} < 0. \tag{A.20}$$

Because $a_k = c_k < L_k$ holds strictly, a marginal change in L_k leaves the feasible set unchanged near the optimum. Concretely, the envelope theorem applied to (A.15) gives $V_{m_2} = u_\ell\big(m_2 + L_\ell,\, \hat{c}_k(m_2, L_\ell)\big)$, where $\hat{c}_k$ solves the first-order condition

$$u_k(m_2 + L_\ell, c_k) = \beta(1+r)\, v'\big(y_3 - (1+r)(L_\ell + c_k)\big),$$

in which L_k does not appear; hence $\partial \hat{c}_k / \partial L_k = 0$. Differentiating V_{m_2} in L_k,

$$V_{m_2,L_k} = u_{\ell\ell} \cdot \frac{\partial (m_2 + L_\ell)}{\partial L_k} + u_{\ell k} \frac{\partial \hat{c}_k}{\partial L_k} = 0.$$

Case 3: Line ℓ is slack. If $a_\ell < L_\ell$, a marginal increase in L_ℓ does not change the feasible set locally. Thus,

$$V_{m_2,L_\ell} = 0. \tag{A.21}$$

To show this, consider the Lagrangian and its first-order conditions from (A.2)–(A.11). Since line ℓ is strictly slack and the regime is locally unchanged (Assumption 4), $\eta_\ell = 0$ in a neighborhood of (m_2, L_ℓ), and therefore

$$V_{L_\ell} = \eta_\ell = 0$$

throughout that neighborhood, which implies (A.21). This argument does not use the status of line k, so it covers both the only-k-binds and neither-binds regimes.

Completing Step 3.

Cases 1–3 cover the bad-state regimes permitted by Assumption 4. Returning to the state decomposition derived at the start of this step

$$W_{b,L_\ell} = (1-p) V_{m_2,L_\ell}(y_1 - \delta + b; L_G, L_R)\,,$$

Cases 1 and 2 give $W_{b,L_\ell} < 0$ whenever line ℓ is strictly binding in the bad state (equations (A.7) and (A.20)), while Case 3 gives $W_{b,L_\ell} = 0$ when line ℓ is slack (A.21). Hence

$$W_{b,L_\ell} \begin{cases} < 0, & \text{if line } \ell \text{ binds in the bad state,} \\ = 0, & \text{if line } \ell \text{ is slack in the bad state.} \end{cases} \tag{A.22}$$

Substituting (A.22) into (A.6) gives us:

$$\frac{\partial \hat{C}_1}{\partial L_\ell} \begin{cases} > 0, & \text{if line } \ell \text{ binds in the bad state,} \\ = 0, & \text{if line } \ell \text{ is slack in the bad state.} \end{cases}$$

Finally, homotheticity and the absence of a current static constraint imply

$$\hat{c}_{q1} = \omega_q \hat{C}_1, \qquad \omega_q > 0.$$

Therefore,

$$\frac{\partial \hat{c}_{q1}}{\partial L_\ell} = \omega_q \frac{\partial \hat{C}_1}{\partial L_\ell}, \qquad q \in \{G, R\},$$

and we arrive at the conclusion:

$$\frac{\partial \hat{c}_{q1}}{\partial L_\ell} \begin{cases} > 0, & \text{if line } \ell \text{ binds in the bad state,} \\ = 0, & \text{if line } \ell \text{ is slack in the bad state.} \end{cases}$$

□

We now turn to Prediction 2, where we assume consumers may exhibit mental accounting ($\gamma_G, \gamma_R > 0$). We continue to use Assumptions 1–4.

Prediction 2. ***With*** *sufficient mental accounting ($\gamma_G > \bar{\gamma}_G > 0$ and $\gamma_R > \bar{\gamma}_R > 0$), an increase in the category-ℓ credit limit will either not affect current consumption in* ***either*** *category or lead to an increase in current consumption* ***only*** *in category ℓ with no spillover to the other category.*

Proof. Fix a line $\ell \in \{G, R\}$, and let $k \neq \ell$ denote the other category. Formally, we seek to prove that under the assumptions above, actual period-1 consumption satisfies

$$\frac{\partial c_{\ell 1}}{\partial L_\ell} \begin{cases} > 0, & \text{if line } \ell \text{ is strictly binding in the bad state,} \\ = 0, & \text{if line } \ell \text{ is slack in the bad state,} \end{cases}$$
$$\frac{\partial c_{k1}}{\partial L_\ell} = 0 \text{ in both cases.}$$

A category-ℓ limit increase either moves consumption in neither category, or moves category-ℓ consumption alone.

Step 1: The target-reset problem.

The consumer re-optimizes only the category-ℓ consumption path, holding the category-k path fixed at the pre-cue plan. Define the period-2 continuation value with cash-on-hand m_2 as the following:

$$\begin{aligned}\widetilde{V}(m_2; L_\ell, L_k) \equiv \max_{c_\ell, a_\ell, a_k, x_3 \geq 0} \; & u(c_\ell, \bar{c}_{k2}) + \beta\, u(m_3 - \bar{c}_{k3},\ \bar{c}_{k3}) \\ \text{s.t.} \;\; & (c_\ell - a_\ell) + (\bar{c}_{k2} - a_k) + x_3 = m_2, \\ & m_3 = x_3 + y_3 - (1+r)(a_\ell + a_k), \\ & 0 \leq a_\ell \leq \min\{L_\ell, c_\ell\}, \\ & 0 \leq a_k \leq \min\{L_k, \bar{c}_{k2}\}.\end{aligned} \tag{A.23}$$

Then define the ex-ante period-2 continuation value as:

$$\widetilde{W}(b; L_\ell, L_k) = p\, \widetilde{V}(y_1 + b; L_\ell, L_k) + (1-p)\, \widetilde{V}(y_1 - \delta + b; L_\ell, L_k). \tag{A.24}$$

Following (A.4), b is the net liquidity carried into period 2. As with (A.2), program (A.23) is concave and m enters linearly, so $\widetilde{V}$ and $\widetilde{W}$ are concave in cash-on-hand: $\widetilde{W}_{bb} \leq 0$.

The new period-1 category-ℓ target with $L'_\ell > L_\ell$ is

$$c^*_{\ell 1} = \arg\max_{c_\ell} \left\{ u(c_\ell, \bar{c}_{k1}) + \beta\, \widetilde{W}\big(b(c_\ell + \bar{c}_{k1});\, L'_\ell, L_k\big) \right\}, \tag{A.25}$$

where $\bar{c}_{k1}$ is the pre-cue consumption plan of category k in period 1, and $b(\cdot)$ is constructed in Step 2 of the proof of Prediction 1 ($b' < 0$, $b'' = 0$ away from the kink). The future category-ℓ targets are the optimal consumptions inside $\widetilde{V}$, and the category-k targets stay at $\bar{c}_{kt}$.

Step 2: Response of the category-ℓ target.

The first-order condition of (A.25) is

$$u_\ell(c^*_{\ell 1}, \bar{c}_{k1}) + \beta\, \widetilde{W}_b\, b' = 0.$$

Let

$$\widetilde{D}_1 \equiv u_{\ell\ell}(c^*_{\ell 1}, \bar{c}_{k1}) + \beta\, \widetilde{W}_{bb}\, [b']^2 < 0.$$

Implicit differentiation gives $\frac{\partial c^*_{\ell 1}}{\partial L_\ell} = -\frac{\beta b'\, \widetilde{W}_{b,L_\ell}}{\widetilde{D}_1}$, so

$$\operatorname{sgn}\left(\frac{\partial c^*_{\ell 1}}{\partial L_\ell}\right) = -\operatorname{sgn}\big(\widetilde{W}_{b,L_\ell}\big). \tag{A.26}$$

It remains to sign $\widetilde{W}_{b,L_\ell}$.

Step 3: The sign of $\widetilde{W}_{b,L_\ell}$.

As in Step 3 of Prediction 1, the buffer enters period-2 cash-on-hand in both states, so

$$\widetilde{W}_{b,L_\ell} = p\,\widetilde{V}_{m_2,L_\ell}(y_1+b;L_\ell,L_k) + (1-p)\,\widetilde{V}_{m_2,L_\ell}(y_1-\delta+b;L_\ell,L_k),$$

and the slackness of the good-state limits gives $\widetilde{V}_{m_2,L_\ell}(y_1+b;L_\ell,L_k)=0$. Then consider the following two cases of the bad state.

Case 1: Line ℓ binds. Suppose $a_\ell = L_\ell < c_\ell$, with a strictly positive multiplier $\eta_\ell > 0$. Apply Lagrange multipliers to program (A.23) as in (A.8): λ on the budget constraint, $\eta_q \geq 0$ on $a_{q2} \leq L_q$, $\zeta_q \geq 0$ on $a_{q2} \geq 0$, $\mu_\ell \geq 0$ on $a_{\ell 2} \leq c_{\ell 2}$, $\mu_k \geq 0$ on $a_{k2} \leq \bar{c}_{k2}$, and $\nu \geq 0$ on $x_3 \geq 0$. Writing $u_{\ell 3} \equiv u_\ell(m_3 - \bar{c}_{k3}, \bar{c}_{k3})$, the first-order conditions are

$$\begin{aligned}
\partial\mathcal{L}/\partial c_{\ell 2} &= u_\ell(c_{\ell 2},\bar{c}_{k2}) - \lambda + \mu_\ell = 0,\\
\partial\mathcal{L}/\partial a_{\ell 2} &= \lambda - \beta(1+r)\,u_{\ell 3} - \eta_\ell + \zeta_\ell - \mu_\ell = 0,\\
\partial\mathcal{L}/\partial a_{k 2} &= \lambda - \beta(1+r)\,u_{\ell 3} - \eta_k + \zeta_k - \mu_k = 0,\\
\partial\mathcal{L}/\partial x_3 &= \beta u_{\ell 3} - \lambda + \nu = 0.
\end{aligned}$$

Since $a_{\ell 2} = L_\ell < c_{\ell 2}$ implies $\mu_\ell = \zeta_\ell = 0$, the $c_{\ell 2}$ condition gives $\lambda = u_\ell(c_{\ell 2},\bar{c}_{k2})$, and adding the $c_{\ell 2}$ and $a_{\ell 2}$ conditions gives the analogue of (A.12) for the mental-accounting problem,

$$u_\ell(c_{\ell 2},\bar{c}_{k2}) = \beta(1+r)\,u_{\ell 3} + \eta_\ell > \beta(1+r)\,u_{\ell 3}.$$

We then derive $\partial c_{\ell 2}/\partial m_2 = \partial c_{\ell 2}/\partial L_\ell = 1$. First, the first-order condition for x_3 gives $\nu = \lambda - \beta u_{\ell 3} = u_\ell(c_{\ell 2},\bar{c}_{k2}) - \beta u_{\ell 3} > \beta r\,u_{\ell 3} > 0$, so complementary slackness implies $x_3 = 0$. Also, a_k is binding: an interior a_k ($\eta_k = \zeta_k = \mu_k = 0$) gives $\lambda = \beta(1+r)\,u_{\ell 3}$, contradicting $\lambda = u_\ell > \beta(1+r)\,u_{\ell 3}$, and $a_k = 0$ would let $\zeta_k = \beta(1+r)\,u_{\ell 3} - \lambda < 0$. Hence we have

$$a_{k2} = \min\{L_k, \bar{c}_{k2}\},$$

with multiplier $\eta_k + \mu_k = \lambda - \beta(1+r)\,u_{\ell 3} = \eta_\ell > 0$.

Since the multipliers on these three constraints are strictly positive ($\eta_\ell, \nu, \eta_k + \mu_k > 0$), we have

$$c_{\ell 2} = m_2 + L_\ell + a_{k2} - \bar{c}_{k2},$$

in which L_ℓ enters only through the binding constraint $a_{\ell 2} = L_\ell$, so $\partial c_{\ell 2}/\partial m_2 = \partial c_{\ell 2}/\partial L_\ell = 1$. The

continuation value on this regime is

$$\widetilde{V}(m_2; L_\ell) = u\big(m_2 + L_\ell + a_{k2} - \bar{c}_{k2},\ \bar{c}_{k2}\big) + \beta\, u\big(y_3 - (1+r)(L_\ell + a_{k2}) - \bar{c}_{k3},\ \bar{c}_{k3}\big). \tag{A.27}$$

Differentiating (A.27) in m_2 gives $\widetilde{V}_{m_2} = u_\ell(c_{\ell 2}, \bar{c}_{k2})$. Since $\partial c_{\ell 2}/\partial L_\ell = 1$ while the period-3 term contains L_ℓ but not m_2,

$$\widetilde{V}_{m_2, L_\ell} = u_{\ell\ell}(c_{\ell 2}, \bar{c}_{k2}) < 0. \tag{A.28}$$

If line k also binds ($a_{k2} = L_k \le \bar{c}_{k2}$), the same expressions apply with $a_{k2} = L_k$.

Case 2: Line ℓ is slack. If $a_\ell < L_\ell$, a marginal increase in L_ℓ does not change the feasible set locally, so $\widetilde{V}_{m_2, L_\ell} = 0$.

Combining the two cases with the state decomposition gives

$$\widetilde{W}_{b, L_\ell} \begin{cases} = (1-p)\, u_{\ell\ell}(c_{\ell 2}, \bar{c}_{k2}) < 0, & \text{if line } \ell \text{ binds in the bad state,} \\ = 0, & \text{if line } \ell \text{ is slack in the bad state,} \end{cases} \tag{A.29}$$

and substituting (A.29) into (A.26)

$$\frac{\partial c^*_{\ell 1}}{\partial L_\ell} \begin{cases} > 0, & \text{if line } \ell \text{ binds in the bad state,} \\ = 0, & \text{if line } \ell \text{ is slack in the bad state.} \end{cases}$$

The category-k targets are unchanged by construction: $\partial c^*_{kt}/\partial L_\ell = 0$ for all t.

Step 4: With large penalties, consumption equals the targets.

After the reset, the consumer solves problem (A.1) at the new credit limit with the new mental targets. We show the optimal plan is the target plan itself.

In period 1, the consumer chooses (c_G, c_R) to solve

$$\max_{c_G, c_R}\ u(c_G, c_R) - \gamma_G |c_G - c^*_{G1}| - \gamma_R |c_R - c^*_{R1}| + \beta\, W\big(b(c_G + c_R)\big)$$

where $W(b)$ is the optimal continuation value of (A.1) given net liquidity b. At a kink of the penalty, the derivative of $|c_q - c^*_{q1}|$ may take any value $g_q \in [-1, 1]$. The candidate $(c^*_{\ell 1}, \bar{c}_{k1})$ is optimal if and only if there exist $g_\ell, g_k \in [-1, 1]$ with $u_q - \gamma_q g_q + \beta W_b b' = 0$ for both categories, that is, if and only if

$$\big|\, u_\ell(c^*_{\ell 1}, \bar{c}_{k1}) + \beta W_b b' \,\big| \le \gamma_\ell \qquad \text{and} \qquad \big|\, u_k(c^*_{\ell 1}, \bar{c}_{k1}) + \beta W_b b' \,\big| \le \gamma_k. \tag{A.30}$$

The first inequality holds with its left-hand side equal to *zero*: for γ_k above the threshold $\bar{\gamma}_k$ constructed below, consumption in periods 2–3 stays at the targets, so $W_b = \widetilde{W}_b$ and the condition is exactly the first-order condition of the reset problem (A.25). Substituting this first-order condition, $\beta W_b b' = \beta \widetilde{W}_b b' = -u_\ell$, into the second inequality gives

$$\left| u_k(c^*_{\ell 1}, \bar{c}_{k1}) - u_\ell(c^*_{\ell 1}, \bar{c}_{k1}) \right| \leq \gamma_k. \tag{A.31}$$

Not deviating from the new targets in period 1 requires only that the marginal-utility gap between the two categories, evaluated at the target plan, not exceed the penalty.

Note that evaluating conditions (A.30) and (A.31) only requires known marginal utilities prior to re-optimizing, so the consumer is able to decide whether it is worthwhile to fully rebudget without needing to first solve the full re-optimization problem.

It remains to find conditions under which the target plan $(c^*_{\ell,t}, \bar{c}_{k,t})$ in the reset problem (A.23)–(A.25) is optimal in every period t and state s; we now explicitly include a state subscript s for clarity. By construction, only the category-k consumptions require checking since category-ℓ consumption, borrowing, and saving are chosen optimally in the reset problem. Also, in the reset problem, the first-order condition of category-ℓ consumption sets u_ℓ equal to the shadow value $\lambda_t(s)$ of the resource constraints in period t and state s. In particular, $\lambda_1 = u_\ell(c^*_{\ell 1}, \bar{c}_{k1})$ by (A.25). Holding c_{kt} at its target in period t and state s is optimal in the penalized problem (A.1) if and only if it satisfies the following subgradient condition:

$$\left| u_k\big(c^*_{\ell t}(s), \bar{c}_{kt}(s)\big) - \lambda_t(s) + \mu_{kt}(s) \right| \leq \gamma_k, \tag{A.32}$$

where $\mu_{kt}(s) \geq 0$ is the multiplier on the constraint $a_k \leq c_k$ as in (A.8) (each credit line can be used only for same-category consumption). We have $\mu_{kt}(s) = 0$ in period 3 ($a_{q3} = 0$), in the good state in period 2 (slack limits), and in the bad state in period 2 ($t = 2, s = \text{bad}$) whenever $a_{k2}(\text{bad}) < \bar{c}_{k2}(\text{bad})$. The exception is the period-2 bad state corner in Case 1 of Step 3, where all category-k spending is credit-financed on a slack line ($a_{k2}(\text{bad}) = \bar{c}_{k2}(\text{bad}) < L_k$): a marginal unit of $c_{k2}(\text{bad})$ relaxes the cap $a_{k2}(\text{bad}) \leq c_{k2}(\text{bad})$ and is financed by borrowing, so $\mu_{k2}(\text{bad}) = \lambda_2(\text{bad}) - \beta(1+r)u_{\ell 3} > 0$ and the condition becomes $|u_{k2}(\text{bad}) - \beta(1+r)u_{\ell 3}| \leq \gamma_k$. Define

$$\bar{\gamma}_k \;\equiv\; \max\left\{ \left\{ \left| u_k\big(c^*_{\ell t}(s), \bar{c}_{kt}(s)\big) - \lambda_t(s) + \mu_{kt}(s) \right| \right\}_{t \in \{1,2,3\},\, s} \right\}$$

as the maximum of the left-hand side of equation (A.32) across time periods and states. If $\gamma_k > \bar{\gamma}_k$, every subgradient condition holds with strict inequality: all consumption plans stay at their

targets, so the periods-2–3 continuation of (A.1) coincides with $\widetilde{W}$ in the reset problem (A.24). Since (A.1) is a concave program, the target plan is a global maximum, so $c_{\ell 1} = c^*_{\ell 1}$ and $c_{k1} = \bar{c}_{k1}$.

Step 5: Conclusion.

Combining Steps 3 and 4: actual category-ℓ consumption moves one-for-one with its target, and actual category-k consumption stays at its unchanged target. Therefore

$$\frac{\partial c_{\ell 1}}{\partial L_\ell} \begin{cases} > 0, & \text{if line } \ell \text{ binds in the bad state,} \\ = 0, & \text{if line } \ell \text{ is slack in the bad state,} \end{cases}$$
$$\frac{\partial c_{k1}}{\partial L_\ell} = 0 \text{ in both cases.}$$

A category-ℓ limit increase either moves consumption in neither category, or raises category-ℓ consumption alone, with no spillover effect across consumption categories. □

C Testing for Heterogeneous Treatment Effects

In addition to the heterogeneity analyses along our pre-specified dimensions (i.e., baseline credit utilization ratio, baseline credit limit, share of spending through Cashea, and self-reported credit constraints), we are interested in testing whether our average treatment effects mask heterogeneity in same-line and cross-line MPCs over other observables. We are particularly interested in whether there are subgroups we can identify who treat credit limit increases as fungible and increase spending on both lines, as standard models without mental accounting would predict. We follow Chernozhukov, Demirer, Duflo, and Fernández-Val (2025) to estimate heterogeneous treatment effects without imposing any particular functional form assumptions or making assumptions about which variables to include.

We estimate heterogeneous treatment effects separately for each of the four treatment-outcome combinations. Treatment is defined as either a grocery line or a retail line limit increase only (we omit the "both" treatment groups from this analysis), and the outcomes are average monthly grocery-line or retail-line purchases over the 12 months following treatment. For each treatment-outcome pair, we estimate heterogeneous treatment effects using the control and the two treatment arms (+50% and +100%) which only increase the limit on the relevant line. We classify assignment to either treatment arm as a binary treatment variable.

The Chernozhukov, Demirer, Duflo, and Fernández-Val (2025) method randomly divides the sample into two equal-sized subsamples. We train a machine-learning model to predict individual treatment effects using all available baseline characteristics in the first subsample. The predicted

values are subsequently applied to the second subsample in two separate analyses. First, we regress purchases on treatment and on treatment interacted with a (demeaned) predicted effect. The coefficient on the interaction, which we label "HTE" in Table A12, is 1 if the predictions capture true heterogeneity perfectly, and 0 if there is no heterogeneity or the predictions are entirely uninformative. Therefore, rejecting $H_0: \text{HTE} = 0$ indicates that the treatment effect varies across users based on our set of observable characteristics. Second, we quantify the degree of heterogeneity by sorting users into quintiles based on their predicted treatment effect. We estimate the average treatment effect within each quintile, following the GATES approach (see Figure 8).

The results depend on each random subsample split, so we repeat the procedure for 100 different random splits. Following the recommendations of Chernozhukov, Demirer, Duflo, and Fernández-Val (2025), we report the median point estimates, p-values, and confidence intervals across these 100 splits.

We implement four different machine-learning models to identify the model whose predictions most closely correspond to the observed treatment effects according to the two selection criteria specified in Chernozhukov, Demirer, Duflo, and Fernández-Val (2025). The four we consider are elastic-net, random-forest, support-vector-machine, and neural-network models. Chernozhukov, Demirer, Duflo, and Fernández-Val (2025) do not require that the HTE test in Table A12 and the GATES analysis in Figure 8 use the same machine-learning model but, in practice, the elastic net model performs better on both criteria across all four treatment-outcome pairs.

We discuss the results in Section 6.4 of the paper. Table A12 shows that we can reject the null of no heterogeneity for both same-line effects ($p < 0.001$). For cross-line effects, we cannot reject the null hypothesis of no heterogeneity in the grocery-line-to-retail spillover, but we can reject it for the retail-line-to-grocery spillover ($p = 0.006$). Figure 8 shows treatment effects by quintile for each treatment-outcome pair.