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

Sean Higgins† Raymond Kluender‡

September 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 (RCT) with Cashea, Venezuela's largest consumer credit provider. Cashea provides two distinct credit lines to all of its borrowers: one for grocery stores and pharmacies, and one for all other retail merchants. We randomized 0%, 50%, or 100% credit-limit increases on one or both credit lines for 210,000 active users. We estimate large, symmetric same-category monthly marginal propensities to consume (MPCs) out of credit-limit increases 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) there are no subgroups who treat 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 mentally account for credit, with economically large effects on borrowing and the allocation of spending across categories, as well as 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, Columbia New Empirical Finance Conference, 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 mediates the effects of expansionary monetary and fiscal policy (Agarwal, Chomsisengphet, Mahoney, and Stroebel, 2018; Auclert, 2019).

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 decisions (Thaler, 1985, 1990, 1999), but there is limited work on whether mental accounting shapes borrowing decisions.

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 likely heterogeneous across consumers. Second, even in an analogous setting for credit with category-specific credit lines, testing fungibility requires exogenous variation in category-specific credit limits. Without mental accounting, credit-limit increases can cause an increase in consumption due to binding constraints (Gross and Souleles, 2002; Zeldes, 1989) or precautionary motives (Aydin and Kim, 2024), but spending should increase in all categories; in contrast, mental accounting predicts a consumption increase only for the spending categories associated with the credit line. Third, it 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 210,000 consumers in partnership with Cashea, the largest consumer lender in Venezuela. Cashea has more than eight 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, electronics, 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 210,000 active users. We estimate category-specific marginal propensities to consume out of these credit-limit increases using 11 months of administrative data supplemented with baseline and follow-up surveys which mea-

sure 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 on the effects of these category-specific credit-limit increases with and without mental accounting. Thus, the model illustrates how the experiment is a test of mental accounting. Our model features a two-category consumption-savings model with separate credit lines for each category, mirroring our institutional setting. Future income is subject to exogenous shocks, and as a result people may be dynamically constrained, which can lead them to increase consumption in response to a credit-limit increase even if they are not statically constrained (i.e., already using their entire credit limit in at least one category). 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 change in response to a change in the category-specific credit limits. We treat a category-1 credit-limit increase as a salient cue about category-1 consumption (*à la* Bordalo, Gennaioli, and Shleifer, 2020). The change in credit limits causes the consumer to re-optimize, but because re-optimizing the full consumption portfolio is difficult, consumers solve a simpler single-good re-optimization problem to set their new consumption targets, holding fixed their target in the other category—in the spirit of similarity-based attention (Bordalo, Gennaioli, Lanzani, and Shleifer, 2025).

The model predicts that with mental accounting, consumers have a high psychological cost of deviating from their targets and thus either do not change consumption (if they are not dynamically constrained) or increase their consumption in the category whose credit limit increased, while holding fixed their consumption in the other category. Without mental accounting, as long as the user is not statically constrained, there is no cost of deviating from the category-specific targets and thus the user will re-optimize their consumption across both categories. Thus, in the absence of mental accounting, a user will either not change consumption (if dynamic constraints have no chance of binding) or will increase consumption on both categories in response to a category-specific credit line increase.

We find overwhelming evidence of mental accounting for credit: 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 9.6 cents, a $1 increase in the retail-line credit limit increases monthly retail-line purchases by 7.9 cents, and spillover effects to the other categories are, in each case, an order of magnitude smaller (-1.0 cents for the grocery-line-to-retail spillover and 0.4 cents for the retail-line-to-grocery spillover). Comparing results across treatment

arms, same-category MPCs are approximately linear: increasing a credit limit by 100% leads to approximately twice as large of a spending response as increasing it by 50%. Increasing both credit limits leads to similar spending responses as the same-category responses to increasing only one credit line. Given that our experiment increased credit limits by a large amount (50% or 100%) and that roughly half of users' grocery and other retail spending is done through Cashea, the same-line MPCs represent meaningful changes to consumption: for example, increasing the retail-line limit by 100% caused an increase in monthly consumption equivalent to 5.8% of 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 and durable over time—consistent with consumers permanently updating only 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 nearly identical to the ten-month averages (0.079 vs. 0.096 for the grocery line and 0.064 vs. 0.079 for the retail line), which shows that the same-line effects are both immediate and persistent. The cross-line spillovers are persistently an order of magnitude smaller, consistent with mental accounting. The point estimate of the grocery-line-to-retail spillover is 0.009 and statistically insignificant in the first month of the experiment, while the retail-line-to-grocery spillover is 0.013. Both cross-line MPCs are persistently near-zero and decline slightly over time. We cannot reject that the grocery-line-to-retail spillover equals zero for the first eight months, after which it is slightly negative and statistically significant. The retail-line-to-grocery spillover is slightly positive and significant for the first four months, after which we cannot reject that it equals zero.

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. To distinguish these explanations, we use our follow-up surveys, which measure spending in each category both on and off Cashea. The survey accurately replicates the outcomes estimated using administrative data: we estimate same-category, on-Cashea MPCs of 0.065 for the grocery line and 0.048 for the retail line (versus 0.096 and 0.079 in the administrative data, respectively). Turning to the measures of spending off-Cashea, we find no evidence of payment-method substitution: for both the grocery and retail lines, we estimate statistically insignificant, near-zero effects of the credit-limit increases on both same- and cross-category off-Cashea spending. We therefore interpret our results throughout as increases in overall consumption.

Third, the results cannot be explained by binding credit constraints. Without mental accounting, a same-category response but no cross-category spillover can only occur if borrowers are

statically constrained (i.e., they exhaust the credit line and have no cash to reallocate toward that category). Our participants are not: median baseline utilization is 16.8% on the grocery line and 12.0% on the retail line, and 74.0% of users typically spend cash on both categories. Given the possibility that static constraints could explain a zero cross-category spillover, we prespecified heterogeneity by tercile of credit utilization (separately by line) at baseline. While the same-category MPC is higher for the most-constrained tercile, as we would expect, the cross-category spillover is near-zero for all three terciles across both lines. The strong same-line response and near-zero spillover for users in the first and second credit-limit utilization terciles, who are very far from their constraints, is strong evidence of mental accounting. We additionally test heterogeneity by self-reported credit constraints among baseline survey respondents, and again see large same-category and zero cross-category MPCs even among users who report they are unconstrained.

Fourth, we find no subgroups with large cross-category spillovers in response to an increase in one line's credit limit, which contradicts standard models without mental accounting. We use machine learning and all baseline characteristics available in administrative data to test whether there is any detectable heterogeneity in the zero or near-zero cross-category spillovers (Chernozhukov, Demirer, Duflo, and Fernández-Val, 2025), and fail to reject the null hypothesis of no heterogeniety. In contrast, we do find substantial heterogeneity in the large same-category consumption responses. 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 merchant-category-specific MPCs separately for each of these ten sub-categories for both same-line and cross-line credit-limit increases. We estimate positive and statistically significant same-category MPCs for *every single sub-category* out of an own-line credit-limit increase. The cross-line MPCs are near zero for all ten of the merchant categories, with no cross-line MPC exceeding 0.003. 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 $28 per month (5.8% of 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 past-due loans are null or small, with the largest being a 1.1 percentage-point (5.5%) increase in 90+ day delinquency in the arm that doubled both limits, and the credit-limit increases were profitable for Cashea. 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" loans, with particular concern over the provision of credit for basic necessities like groceries and healthcare.[1] Our results suggest that expanding category-specific credit induces re-budgeting 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 predictions of standard models that 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; Auclert, 2019), but shifts in the yield curve in response to monetary policy are sometimes non-parallel, and changes to rates on different parts of the yield curve affect rates of different credit products. Short-term rates pass through to credit-card interest rates (Ausubel, 1991; Bräuning and Stavins, 2025), while longer-term rates pass through to auto-loan and mortgage rates (Beraja, Fuster, Hurst, and Vavra, 2019; Di Maggio, Kermani, Keys, Piskorski, Ramcharan, Seru, and Yao, 2017; Scharfstein and Sunderam, 2016; Tran, 2025). Without mental accounting, a decrease in auto-loan rates should increase consumption of both cars and groceries, whereas with mental accounting, cheaper auto loans may only increase car consumption. Furthermore, because utility is concave in the consumption of each good, the optimal increase in overall consumption in response to a decrease in auto-loan rates is smaller when

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

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 what goods cash is spent on. We randomize category-specific credit limits and measure spending on and off the platform, which allows us to estimate how expansions to category-specific credit limits 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 no more than 6%.[2] Over the next 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 2% of its value a decade earlier.[3]

Cashea filled this void in 2022 by launching dollar-indexed, interest-free installment loans through merchant partners at the point of sale. Cashea has more than eight million users and 20,000 merchant partners spanning grocery stores, pharmacies, retail outlets, and service providers. The Cashea app has been downloaded on more than half of smartphones in the country, processes

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

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 granular and 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 (described in more detail below), 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 up through 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, and also 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.[4] The consumer makes the down payment at purchase and 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 access even after the borrower becomes current, by lowering the user's level and associated limit. Conversely, as users accumulate purchases and repay on time, 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.[5] 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 prevents dollar-indexed loan amounts and local-currency inflation from distorting intertemporal consumption decisions or choices about

[4]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.
[5]Embassy of Switzerland in Venezuela (2025), "Economic Report 2024/2025 Venezuela."

how much to consume on credit.[6] 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 clothing, home goods, electronics, services, and other retail. We refer to these two credit lines as the "grocery line" and "retail line" to streamline exposition. Appendix Table A1 reports baseline spending (number of transactions and dollars) by store subcategory 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 disbursed across other retail categories.

The grocery and retail lines have different terms to mitigate risk and because grocery stores and other retail have different margins. In particular, the maturity of grocery-line loans is shorter than that of the retail-line loans in part because the merchants are financing the credit, and grocery stores have thinner margins than other retail stores, and thus less ability to absorb a delay in cash flows. 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: a down payment followed by three installments at 14-day intervals, with up to 12 installments available at select merchants for higher-level users (Guttman-Kenney and Zhang, 2026). 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 (Appendix Figure A1), relative to a median monthly income 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 categories of goods, 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

[6] 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."

mental accounting following Farhi and Gabaix (2019), in which consumers face a psychological cost of deviating from category-specific consumption targets. Finally, we microfound the 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 of other retail goods. 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 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: $m_t \equiv y_t + x_t - (1 + r)\,(a_{G,t-1} + a_{R,t-1})$. 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. 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. It also states that the dollar amount spent on credit on a good cannot exceed the total dollar amount spent on that good (on credit and with cash), which is merely an accounting restriction. 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 $\gamma_\ell > 0$ 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.

The formulation in equation (5) and in Farhi and Gabaix (2019) is agnostic about where the category-specific consumption targets come from. The source of these targets is critical in our setting, as the experiment changes category-specific credit limits and the model's predictions will depend on how the targets respond to those changes in credit limits. Thus, we microfound 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, 2025).

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), these new consumption targets will determine their new consumption paths, since deviating from the new target is costly; in Appendix B, we formalize the thresholds $\bar{\gamma}_G$ and $\bar{\gamma}_R$ that define "high-γ" consumers. In contrast, for low-γ consumers, deviating is not costly and the consumer will re-optimize their consumption across both categories in spite of the targets only changing for category ℓ. Thus, our microfoundation of the consumption targets 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.

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 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 Aydin (2022), Aydin and Kim (2024), and D'Acunto, Rauter, Scheuch, and Weber (2020). Survey evidence also supports the relevance of dynamic liquidity 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, even in the presence of dynamic liquidity constraints that bind in the bad state in period 2, without mental accounting a borrower would always equate their marginal utilities of consumption across categories in period 1, and thus would never respond to a credit-limit increase by increasing period-1 consumption only on one category—unless they are *statically* constrained (i.e., a credit limit binds in period 1). Being statically constrained requires two conditions to hold

in period 1 in the absence of a credit-limit increase: first, they 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 is the only case in which the agent would not have equal marginal utilities of consumption across categories in the absence of the credit-line increase, which is a necessary condition for them to respond to a credit-limit increase by changing consumption only on one category without mental accounting.

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. Borrowers are not 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. Nearly all Cashea users (74.0%) also typically spend cash on both categories. Thus, testing whether they respond to a category-specific credit-line increase by increasing consumption only on that category or on both categories is a test of mental accounting.

4 Experimental Design and Data

4.1 Experimental Design

The experiment sample includes 210,000 active users. We defined active users as those who had 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.[7] Table A14 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 A2 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 A2 shows that even 11 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, and (3) the average utilization ratio for the retail line, resulting in 27 strata.[8] Appendix Figures A1 and A3 show the distributions of baseline limits and baseline utilization by line.

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

[7] 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.

[8] Misfits within a stratum were randomly assigned, which generated some small differences in the number of users allocated to each treatment arm, as in McKenzie and Bruhn (2011).

control group received an otherwise comparable email promoting the usage of both Cashea lines and showing users examples of goods they could purchase on each line, in order to mitigate the concern that the messages could have salience effects, which could explain our treatment effects if the control group had not received similar messages. Each message contained a graphics interchange format (GIF) animation that cycled through the images shown, as well as additional images shown in Appendix Figure A4. Users also received a push notification notifying them of the limit increase, with similar language for the control group.[9]

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 that their post-experiment credit limit is within 2.5% of their actual observed limit, and most of the other half are close (Appendix Figure A5). 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 A6).

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 11 months from the launch of the experiment on October 3, 2025, through September 3, 2026. Table A14 reports baseline characteristics of the sample and demonstrates balance using both joint F-tests across treatment arms and omnibus F-tests across covariates.

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.

[9] 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."

We fielded a baseline survey the week before the experiment launched. The baseline survey elicited spending in each category on and off Cashea, why they made 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.[10] 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%.[11]

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 for their stock of savings, the flow of savings over the last 30 days, and whether their stock and flow of savings had increased or decreased since September 2025.[12]

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 treatment arms within survey respondents. Appendix Table A3 compares covariates for survey respondents and non-respondents and Appendix Table A4 tests for balance across treatment arms within the follow-up survey respondents. Appendix Table A5 reports demographics for the baseline survey respondents.

[10] 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 190,000 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.

[11] 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.

[12] 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.

5 Empirical Strategy

5.1 Marginal Propensities to Consume

Our empirical strategy is designed to test the model in Section 3 by exploiting the randomized credit-limit increases to recover the MPC out of a same-category limit increase and the MPC out of a cross-category limit increase. Under any model without mental accounting in which money is fungible and agents are not statically constrained, agents should increase spending in both categories or neither. With mental accounting, we predict agents will increase spending *only* on the same-line category for which they received a credit-limit increase (or on neither category).

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. While the magnitude of natural changes in credit limits is small relative to the experimentally induced credit-limit changes, this is important for the purposes of estimating MPCs because the 50% and 100% increases apply to heterogeneous baseline limits. Larger baseline credit limits also see larger dollar-denominated ordinary growth in credit limits which creates an identification problem.

We address this challenge with an instrumental variables (IV) strategy that isolates the experimentally induced variation in credit limits by instrumenting for post-treatment credit limits with the treatment assignment indicators. The first-stage equations are:

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

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^k is an indicator for assignment to treatment arm $k \in \{1, \ldots, 6\}$ (control omitted), and $\delta_{s(i)}$ 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:

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

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, α_G is the same-category MPC and α_R is the cross-category MPC;

the interpretations reverse for retail-line purchases. The structure of 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 and is satisfied by design. In prespecified heterogeneity analysis, we fully interact each regression with strata indicators (e.g., indicators for tercile of baseline credit utilization) to estimate strata-specific coefficients.

We construct purchase outcomes by first summing purchase data from the experiment's launch on October 3, 2025 through September 3, 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 at the user level.

5.2 Treatment Effects by Arm

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

$$y_i^{\ell} = \sum_{k=1}^{6} \alpha_k T_{ik} + \delta_{s(i)} + \varepsilon_i, \tag{8}$$

where α_k is the effect of assignment to arm k relative to the control group. By including a separate indicator for each arm, this specification imposes no assumptions about the linearity of effects in the size of the limit increase or the interaction of increases across lines, and comparisons of the α_k coefficients provide nonparametric tests of the dose-response and additivity effects of the credit-limit changes.

6 Results

Our model predicts that without mental accounting, individuals will either not respond to a category-ℓ credit-limit increase or will respond by increasing consumption on *both* categories. In contrast, with mental accounting, individuals will either not respond to the limit increase or will respond by increasing consumption *only* on category ℓ. As our main result, we estimate same- and cross-category monthly marginal propensities to consume out of a credit-limit increase using specifications (6) and (7).

Figure 3 presents the central findings of the paper. For each $1 increase in the grocery-line credit limit, participants consume an additional 9.6 cents per month on their grocery line, with a near-zero spillover to retail-line spending of -1.0 cents. For each $1 increase in the retail-line credit limit, participants consume an additional 7.9 cents per month on their retail line again with

a near-zero spillover to grocery line spending of 0.4 cents per month.[13] These large same-category spending responses, with virtually no cross-category spillovers, are consistent with our model of mental accounting and cannot be explained by alternative models.

Figure 4 reports treatment effects separately by arm. The large same-category increases and near-zero spillovers across categories hold across all four treatment arms in which we increased only one credit line, and consumption increases in both categories when we increase both credit limits. 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 ($13.8 versus $7.9 on the grocery line; $27.9 versus $16.4 on the retail line).[14] 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 pieces of evidence in support of a mental-accounting interpretation of these main results: (1) the large same-category MPCs and near-zero cross-category spillovers persist in every post-treatment month (Section 6.1); (2) the results on credit-line spending represent increases in total consumption rather than payment-method substitution (Section 6.2); (3) the main results holds separately for constrained and unconstrained borrowers (Section 6.3); (4) machine-learning heterogeneity analysis finds significant same-category MPCs and near-zero spillovers in every quintile of predicted treatment effects (Section 6.4); and (5) the pattern holds within every sub-category of consumption, 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 re-set 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 would expect the same-category MPC to respond immediately and persist over time. Under mental accounting, we expect zero cross-category MPC on impact and over time if consumers do not rebudget.

We estimate month-specific same- and cross-category MPCs and present them in Figure 5. Both same- and cross-category MPCs are zero for the three months leading up to the experiment, as expected. For both credit lines, we see large on-impact same-category MPCs that are remarkably similar to the average monthly MPCs presented in Figure 3 (0.079 vs. 0.096 for the grocery line,

[13] Results are similar using average balances rather than monthly purchases as the outcome variable (Appendix Figure A7).

[14] Appendix Figure A8 formalizes the comparison, showing similar marginal effects of the first and second 50% increments.

0.064 vs. 0.079 for the retail line).[15] There is predictable seasonality that shows up in the same-category MPCs for both credit lines around the December holidays.

Our main result of near-zero cross-category spillovers for either line increase is present immediately in the first month of the experiment and generally persists over the subsequent 11 months. Both cross-category MPCs decline slightly over time, such that by cross-category spillover MPCs eventually become slightly negative; we discuss implications of this result for how people finance their increased consumption in Section **??**.

In addition to providing evidence supporting our model of mental accounting, the persistence of the large same-category and near-zero cross-category MPCs also helps rule out the alternative explanation that our results are caused by some goods being indivisible (Zorzi and Beraja, forthcoming). If credit-limit increases primarily facilitate indivisible durable purchases, we could observe 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 MPCs to spike upon treatment as people purchase previously unattainable indivisible goods, but to quickly erode over time. Instead, they are stable over time, 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 an increase in consumption on the category with a credit-limit increase, not just an increase in the on-credit-line consumption on that category. An alternative explanation for our results in administrative data is payment-method substitution: for example, users would have consumed the same total amount of groceries even in the absence of a grocery-line limit increase, but after the increase they shift cash consumption of groceries onto their credit line. Because we observe only on-credit-line consumption in the administrative data, we supplement the administrative data 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 through Cashea and 59.6% of their retail-category purchases.

Figure 6 estimates 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 restricted to the time period we asked about in the surveys) and 0.048 for the retail line (compared to 0.062 in administrative data).

[15] Appendix Figure A9 shows treatment effects by arm over time.

Second, we cannot reject zero grocery-limit-to-retail spillover for on-Cashea spending in survey data but we do find a small spillover from retail-line increases to grocery spending (statistically significant at the 10% level), again consistent with the corresponding findings in administrative data (see Appendix Table A6).

Turning to off-Cashea spending, which can only be observed in the survey, we cannot reject zero same-category or cross-category MPCs in either category. 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 0.073 and 0.066 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—consistent with mental accounting—the on-Cashea spending increases we observe in the administrative data represent increases in consumption, not just 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 also 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 baseline terciles of credit-line utilization, which was one of our two main prespecified heterogeneity tests. The logic of the test is that if static constraints explain the main result, we would expect to observe the positive same-category MPC for category ℓ and zero or very low spillover to the other category only among those who are already constrained, i.e., 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 cross-category MPCs are zero or very small. For example, for the middle tercile with baseline utilization ratios between 10.0% and 24.9%, the same-category MPCs are 0.086 and 0.066 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. Appendix Figure A10 shows that results by terciles of baseline utilization of each line separately (rather than total utilization) are similar. Appendix Figure A11 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 self-reported unconstrained have positive and economically meaningful same-category MPCs but zero cross-category MPCs.[16]

Further evidence that most users are not statically constrained is shown in Appendix Figure A14: 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)), consistent with dynamic constraints creating the precautionary motive that explains why consumers do increase consumption at all in response to a credit line increase, but that does not explain 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 (and thus respond to a category-ℓ credit-limit increase by increasing consumption on *both* categories if they increase consumption at all), we use machine-learning to estimate heterogeneity in treatment effects by observables (Chernozhukov, Demirer, Duflo, and Fernández-Val, 2025). Specifically, we use all available 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.

In contrast to standard models without mental accounting, we find zero or near-zero spillovers for all five quintiles grouped by the predicted treatment effect of a retail-line limit increase on grocery spending and of a grocery-line limit increase on retail spending (Figure 8) and formally cannot reject no heterogeneity in the near-zero spillovers (Appendix Table A7). For same-category consumption responses, we do find substantial heterogeneity, but also find that even for the lowest

[16] 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 presepcified: heterogeneity by credit limit on each line (Appendix Figure A12) and by the proportion of consumption in each cateogry that the user reported spending through Cashea in the baseline survey (Appendix Figure A13).

quintile of predicted same-category treatment effects, there is a positive, statistically significant, and economically meaningful same-category consumption response. We conclude that our main results of a large same-category consumption response and a zero or near-zero spillover hold across any subgroups that can be identified using observable characteristics.

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 coincide with the consumers' mental accounts as specified in the model. Each line bundles several spending categories—groceries, pharmacies, and restaurants on the grocery line; clothing, home goods, electronics, services, tools, health, sporting goods, and other retail on the other—and we represent 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 sub-categories of consumption, we would expect to see higher cross-category MPCs and we would be biased *against* finding zero cross-category MPCs.

We test our assumptions about the composition of consumers' mental accounts by estimating sub-category-specific MPCs, which we present in Figure 9. Remarkably, we estimate MPCs that are perfectly coincident with the assignment of sub-categories to mental accounts and the approximate share of consumption each sub-category comprises within the mental account. Specifically, we estimate positive and statistically significant same-category MPCs for every single one of the ten sub-categories of spending we observe. Further, we can rule out cross-category MPCs onto any sub-category of spending of more than 0.006 and the MPCs are generally precise zero estimates (ranging between -0.005 and 0.003). 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. In addition to the fact that we document completely symmetric MPCs for the grocery- and retail-line increases, we also see large same-category MPCs for the most divisible sub-categories on both lines like groceries and clothing.

7 Discussion

Our credit limit changes caused large changes in the level and allocation of consumption: for example, doubling the credit limit on the retail line led to a $28 per month increase in retail consumption with essentially no change to grocery consumption, which amounts to 5.8% of median

monthly income. In contrast, doubling the grocery-line credit limit led to a $14 per month increase in grocery consumption with essentially no change to retail consumption, which amounts to 2.9% of monthly income. 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: there are no subgroups of users for whom we can identify the significant and positive cross-category spillover that would be predicted by a standard theory model that excludes mental accounting.

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 it cannot be explained by consumers pulling forward their consumption. We included several questions to assess where the money is coming from in our second follow-up survey.[17]

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 may be funding some share of their persistently higher same-line 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 A15 shows effects on the probability a loan payment is past due by treatment arm as of 11 months post-intervention. There are no statistically significant effects on the probability a loan is past due for the retail treatment arms, and only small increases in the probability for the grocery-line and both-line treatment groups. The largest increase is just a 1.1 percentage-point increase (5.5%) in the likelihood a loan is 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 (see Panel B of

[17] 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.

Appendix Figure A15 for the effects on user-level profits by treatment arm). 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 loans past due 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 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 the relaxation of the budget constraint fungibly; and, (5) the central result replicates in every single sub-category of spending within the credit lines.

Mental accounting has a prominent position in the behavioral economics canon, but comparatively limited causal empirical evidence. 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). Yet we know little about whether consumers form mental accounts around credit products and, if so, how this affects their consumption and borrowing decisions.

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 to both 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, it affects how central banks should think about monetary policy, as shifts in interest rates on different parts of the yield curve pass through to different credit products, and with mental accounting these changes may pass through only to the subset of consumption categories financed with those credit products.

References

Acree, Nina R., Kayleigh Barnes, Alexander Bruce, and Simona M. 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?*." *The Quarterly Journal of Economics* 133(1), 129–190.

Aiyagari, S. Rao (1994). "Uninsured Idiosyncratic Risk and Aggregate Saving." *The 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."

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

Auclert, Adrien (2019). "Monetary Policy and the Redistribution Channel." *American Economic Review* 109(6), 2333–2367.

Ausubel, Lawrence M. (1991). "The Failure of Competition in the Credit Card Market." *American Economic Review* 81(1), 50–81.

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." SSRN Scholarly Paper. Rochester, NY.

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

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra (2019). "Regional Heterogeneity and the Refinancing Channel of Monetary Policy." *The Quarterly Journal of Economics* 134(1), 109–183.

Berg, Tobias, Valentin Burg, Ana Gombović, and Manju Puri (2020). "On the Rise of FinTechs: Credit Scoring Using Digital Footprints." *The 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." Working Paper 31202. National Bureau of Economic Research.

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 (2025). "A Cognitive Theory of Reasoning and Choice." Working Paper 33466. National Bureau of Economic Research.

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

Bräuning, Falk and Joanna Stavins (2025). "The Credit Card Spending Channel of Monetary Policy: Micro Evidence from Account-level Data." Working Paper 25-10. Federal Reserve Bank of Boston.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Iván Fernández-Val (2025). "Fisher–Schultz Lecture: Generic Machine Learning Inference on Heterogeneous Treatment Effects in Random-

ized 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." Working Paper 34506. National Bureau of Economic Research.

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

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."

Di Maggio, Marco, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao (2017). "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging." *American Economic Review* 107(11), 3550–3588.

Farhi, Emmanuel and Xavier Gabaix (2019). "Optimal Taxation with Behavioral Agents." Working Paper 21524. National Bureau of Economic Research.

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." *The Review of Financial Studies* 37(8). Ed. by Lauren Cohen, 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." *The 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.

Guttman-Kenney, Benedict and Walter Zhang (2026). "Guttman-Kenney and Zhang (2026)." Tech. rep.

Hastings, Justine S. and Jesse M. Shapiro (2013). "Fungibility and Consumer Choice: Evidence from Commodity Price Shocks." *The 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." SSRN Scholarly Paper. Rochester, NY.

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

Kőszegi, Botond and Filip Matějka (2020). "Choice Simplification: A Theory of Mental Budgeting and Naive Diversification*." *The 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."

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

McKenzie, David and Miriam Bruhn (2011). "Tools of the Trade: Doing Stratified Randomization with Uneven Numbers in some Strata." World Bank Blogs: Development Impact. Blog post, published November 6, 2011.

Mian, Atif, Amir Sufi, and Emil Verner (2017). "Household Debt and Business Cycles Worldwide." *The 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." *The 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.

Scharfstein, David and Adi Sunderam (2016). "Market Power in Mortgage Lending and the Transmission of Monetary Policy." Working Paper. Harvard Business School.

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.

Tran, Tu-Uyen (2025). "What Drives Consumer Interest Rates?" https://www.minneapolisfed.org/article/2025/v
drives-consumer-interest-rates.

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." *The Journal of Finance*.

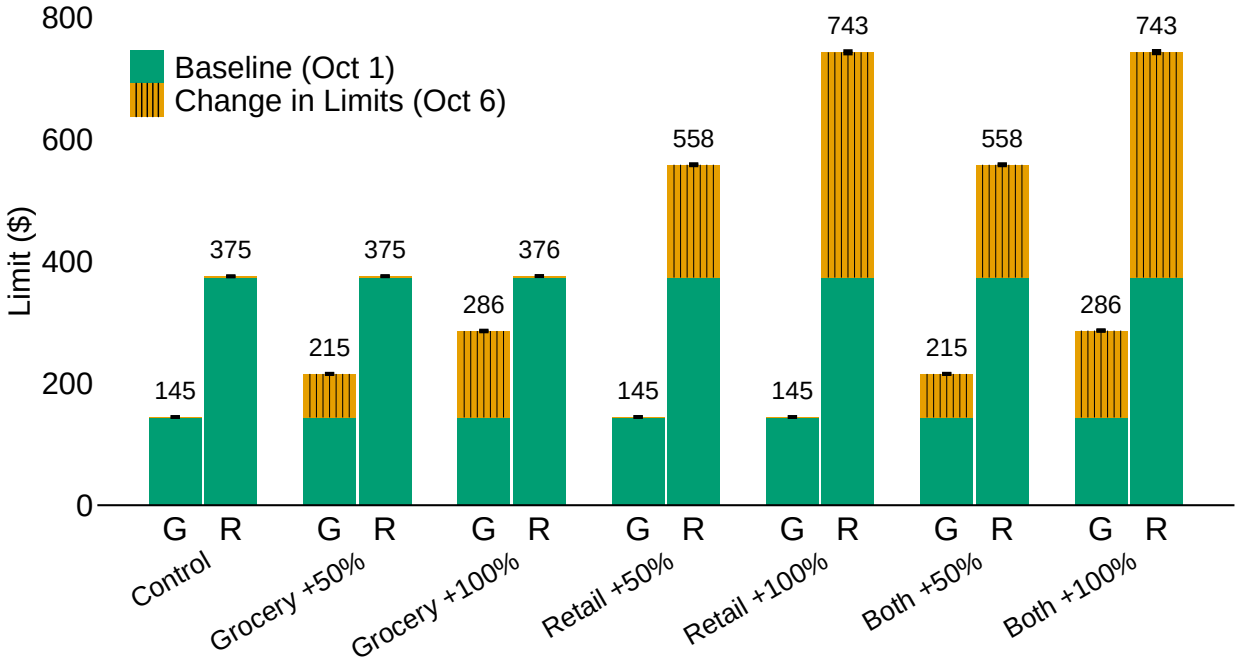
Zeldes, Stephen P. (1989). "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97(2), 305–346.

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 (forthcoming). "Durables and the Marginal Propensity to Spend." *American Economic Review*.

Figures

Figure 1. Credit-Limit Intervention



Notes: Bars show mean credit limits in USD by treatment arm and credit line. "G" denotes grocery line and "R" denotes retail line. Green bars reflect baseline means as of October 1, 2025. Orange bars reflect the change in limits as of October 6, 2025, three days after the experiment launched. Small deviations from exact 0%, 50%, or 100% increases reflect natural changes in credit limits according to Cashea's levels system. Full estimates reported in Appendix Table A2. See Figure A2 for credit-limit changes over the entire experimental period. N = 209,996.

Figure 2. Experiment Outreach

A. Control





B. Grocery Line +50%





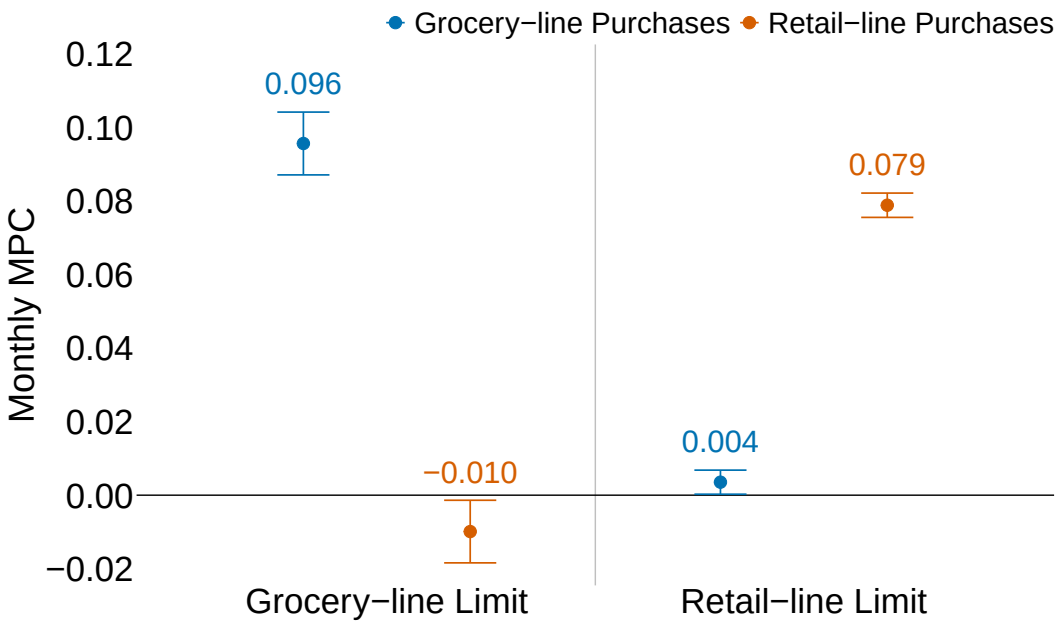
C. Both Lines +100%





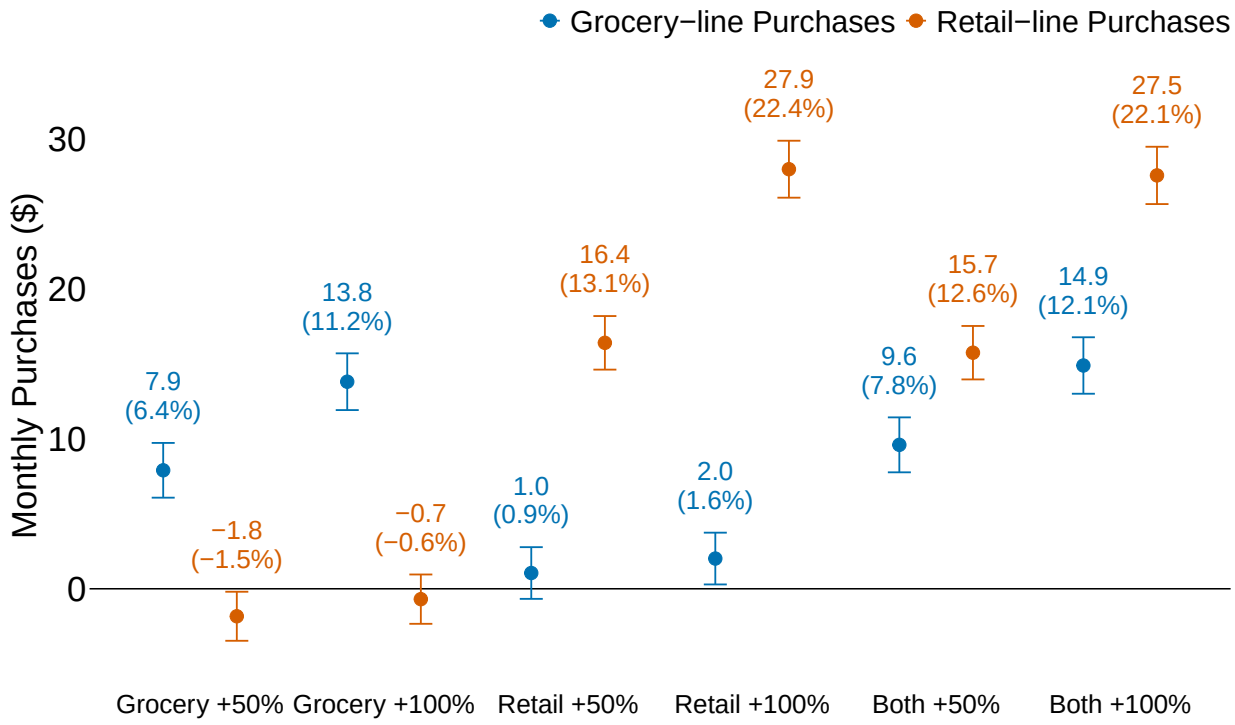
Notes: Screenshots of email outreach sent to participants, showing three of seven treatment arms as examples. Panel A shows the control-group message, which promotes Cashea without mentioning a credit-limit increase. Panel B shows the treatment message for a 50% grocery-line increase. Panel C shows the treatment message for a 100% increase on both lines. All treatment 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 A4.

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



Notes: Figure plots instrumental-variables estimates of the marginal propensity to consume out of credit-limit increases on each line, estimated via the specification described in Section 5. The dependent variable is average monthly purchases on each line from October 3, 2025 to September 3, 2026, winsorized at the 0.1st and 99.9th percentiles. Post-experiment credit limits as of October 6, 2025 are instrumented with the six treatment assignment indicators; regressions control for randomization strata fixed effects. Bars represent 95% confidence intervals. N = 209,996 for the grocery-line model and N = 209,996 for the retail-line model. Full estimates are reported in Appendix Table A9.

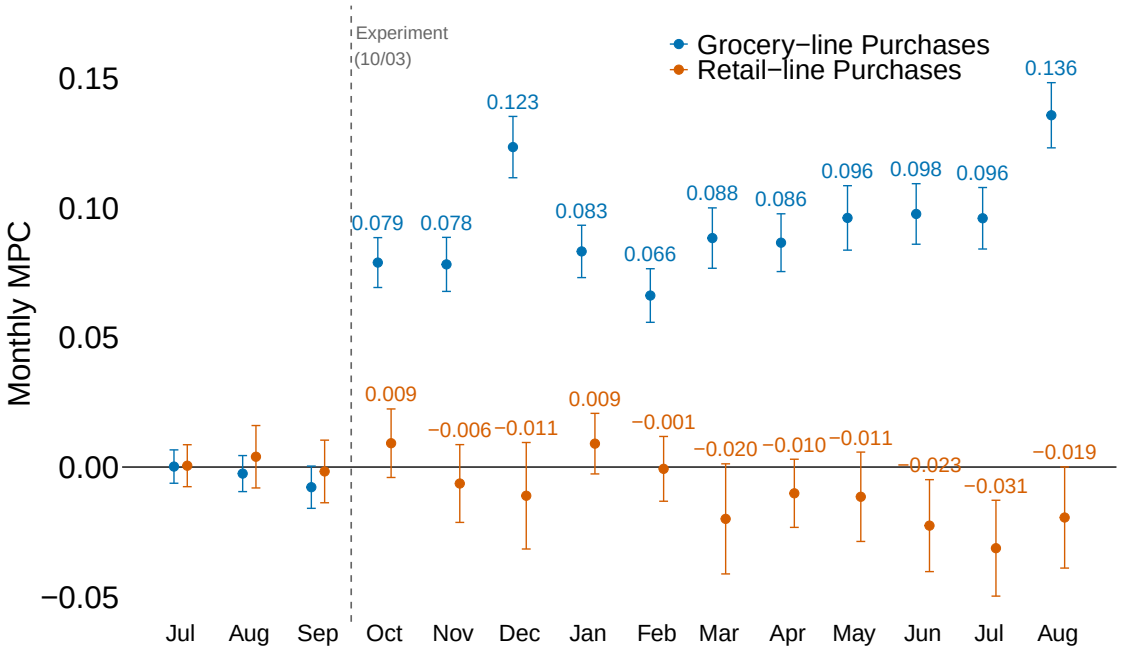
Figure 4. Treatment Effect Estimates of Cumulative Purchase Effects



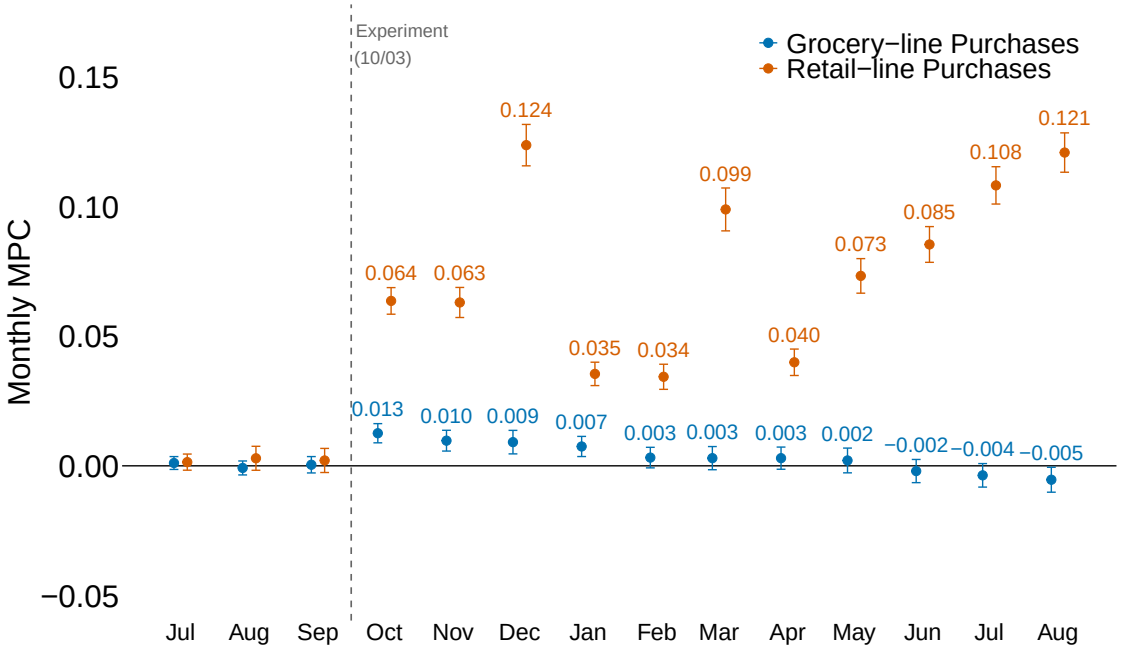
Notes: Figure plots treatment effect estimates of average monthly purchases in USD over the post-treatment period from October 3, 2025 to September 3, 2026, by treatment arm. Blue denotes grocery-line purchases; orange denotes retail-line purchases. Point estimates are labeled with dollar amounts and percent changes relative to the control-group mean of $123 for the grocery line and $125 for the retail line. Bars represent 95% confidence intervals. All estimates are from OLS regressions of the form described in Section 5, controlling for randomization strata fixed effects. Standard errors are clustered at the user level. N = 209,996 for the grocery-line model and N = 209,996 for the retail-line model. Full estimates are reported in Appendix Table A11

Figure 5. Dynamic MPC Estimates on Purchases

(a) Grocery-line Limit

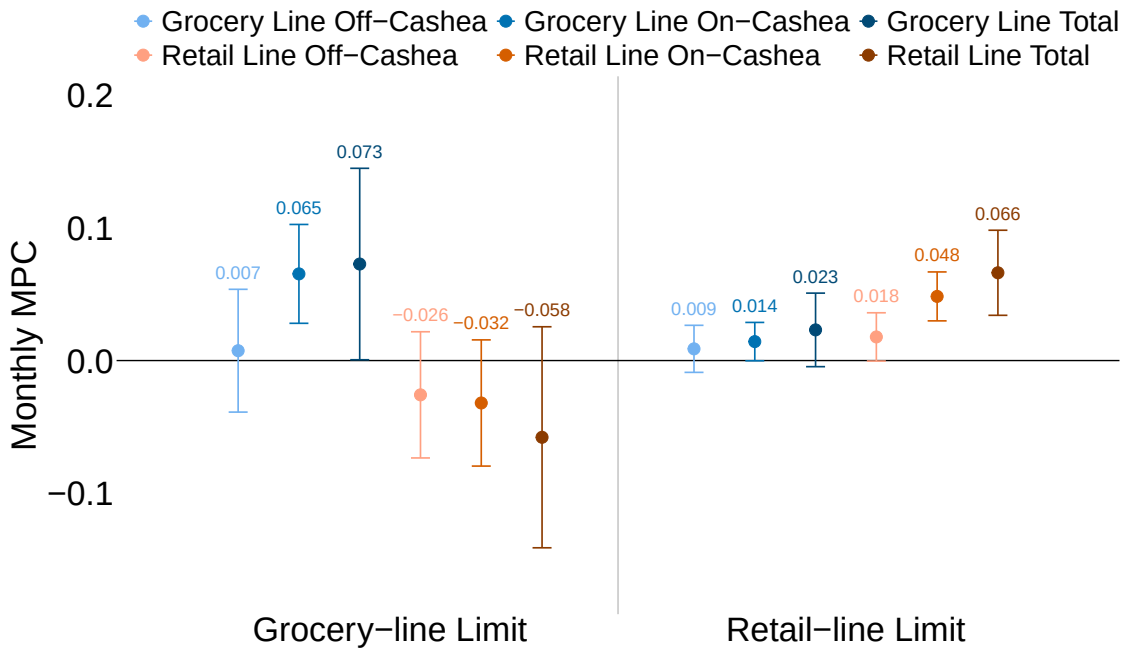


(b) Retail-line Limit



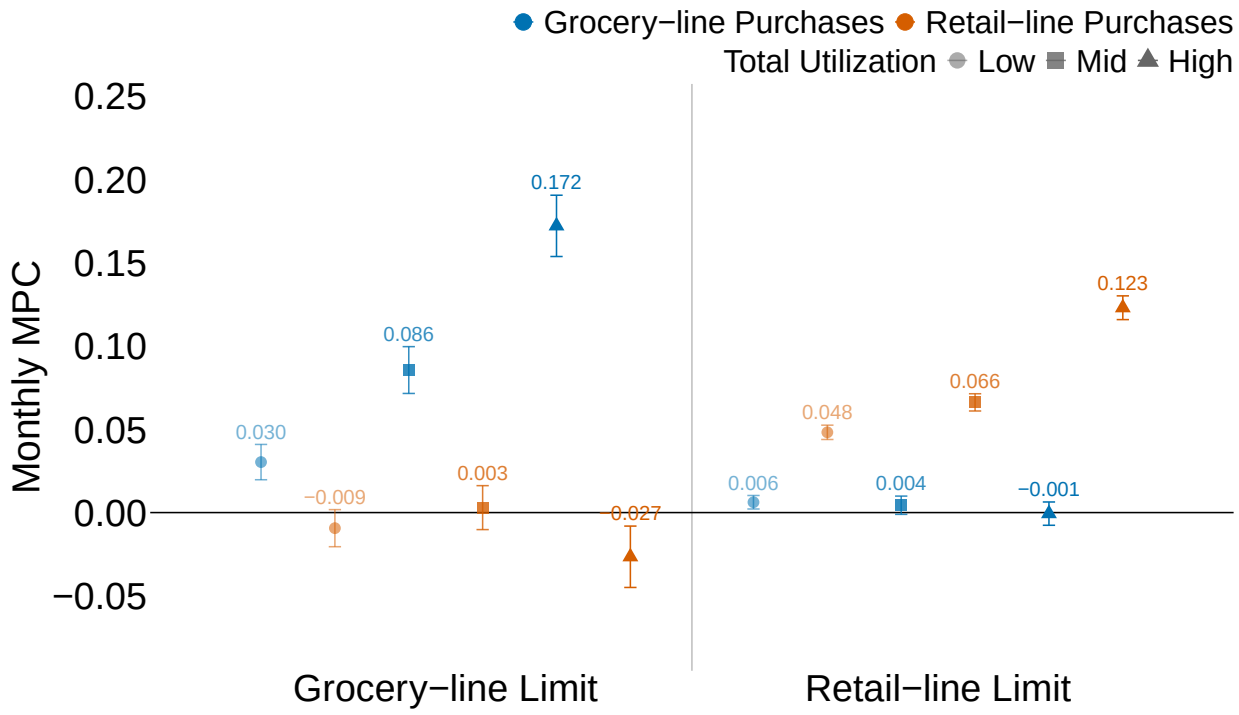
Notes: Figure plots monthly OLS treatment effect estimates from Equations 6 and 7, estimated separately for each month from September 2025 through August 2026, each using a user's limit on October 6th as the instrument. The top panel shows grocery-line purchases and the bottom panel shows retail-line purchases. Each point represents the estimated MPC for each month's purchases; vertical bars indicate 95% confidence intervals. Regressions include randomization strata fixed effects. Users with no purchasing activity in a given month are assigned a value of zero. All continuous variables are winsorized at the 0.1st and 99.9th percentiles within each month. Standard errors are clustered at the user level. The vertical dashed line marks the start of the experiment. N = 209,996. Full estimates reported in Appendix Table A12.

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



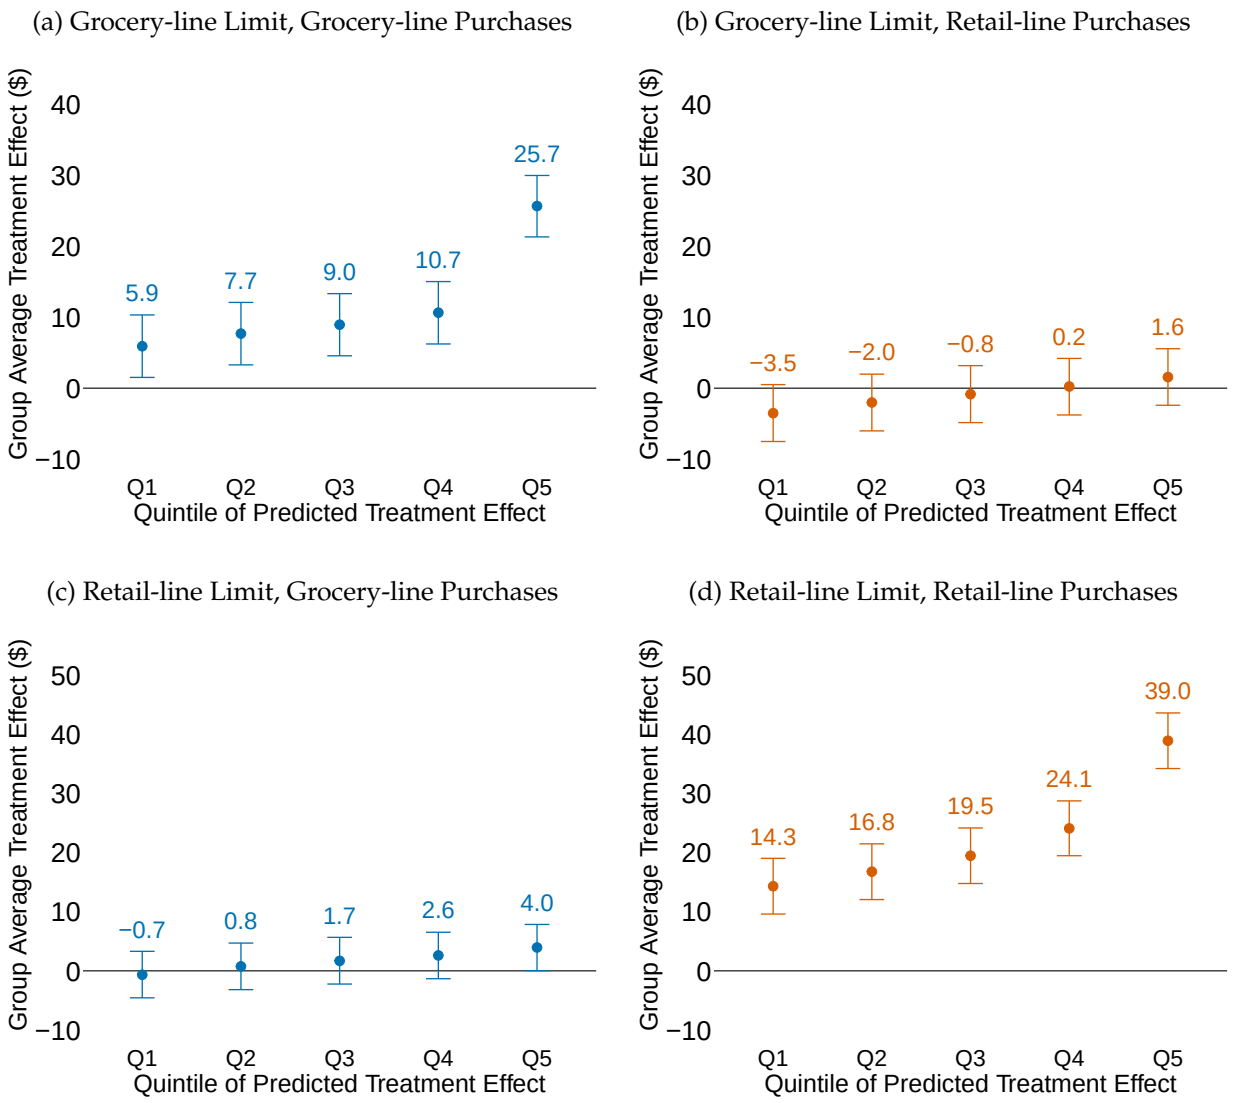
Notes: Figure plots 2SLS estimates of Equations 6–7 using self-reported spending outcomes from a pooled sample of first- and second-wave follow-up survey respondents. The endogenous regressors are post-experiment LC and LP credit limits, instrumented by the six treatment-arm assignment dummies. Each point represents the estimated MPC—dollars of spending per dollar of credit limit—for a given outcome–endogenous regressor pair. The x-axis groups estimates by endogenous regressor (grocery-line limit, retail-line limit); within each group, colors distinguish six outcomes: grocery-line on-Cashea spending, grocery-line off-Cashea spending, grocery-line total spending, retail-line on-Cashea spending, retail-line off-Cashea spending, and retail-line total spending. First-wave survey spending is normalized to 30-day equivalents based on elapsed days since randomization; second-wave spending is used as reported (30-day). All continuous spending outcomes are winsorized at the 95th percentile (a more conservative level than outcomes from administrative data, since the data come from survey responses subject to misreporting). Vertical bars indicate 95% confidence intervals constructed from standard errors clustered at the user level. All specifications include randomization strata fixed effects. N = 23,436. Full estimates reported in Appendix Table A6.

Figure 7. Marginal Propensity to Consume by Total Baseline Utilization



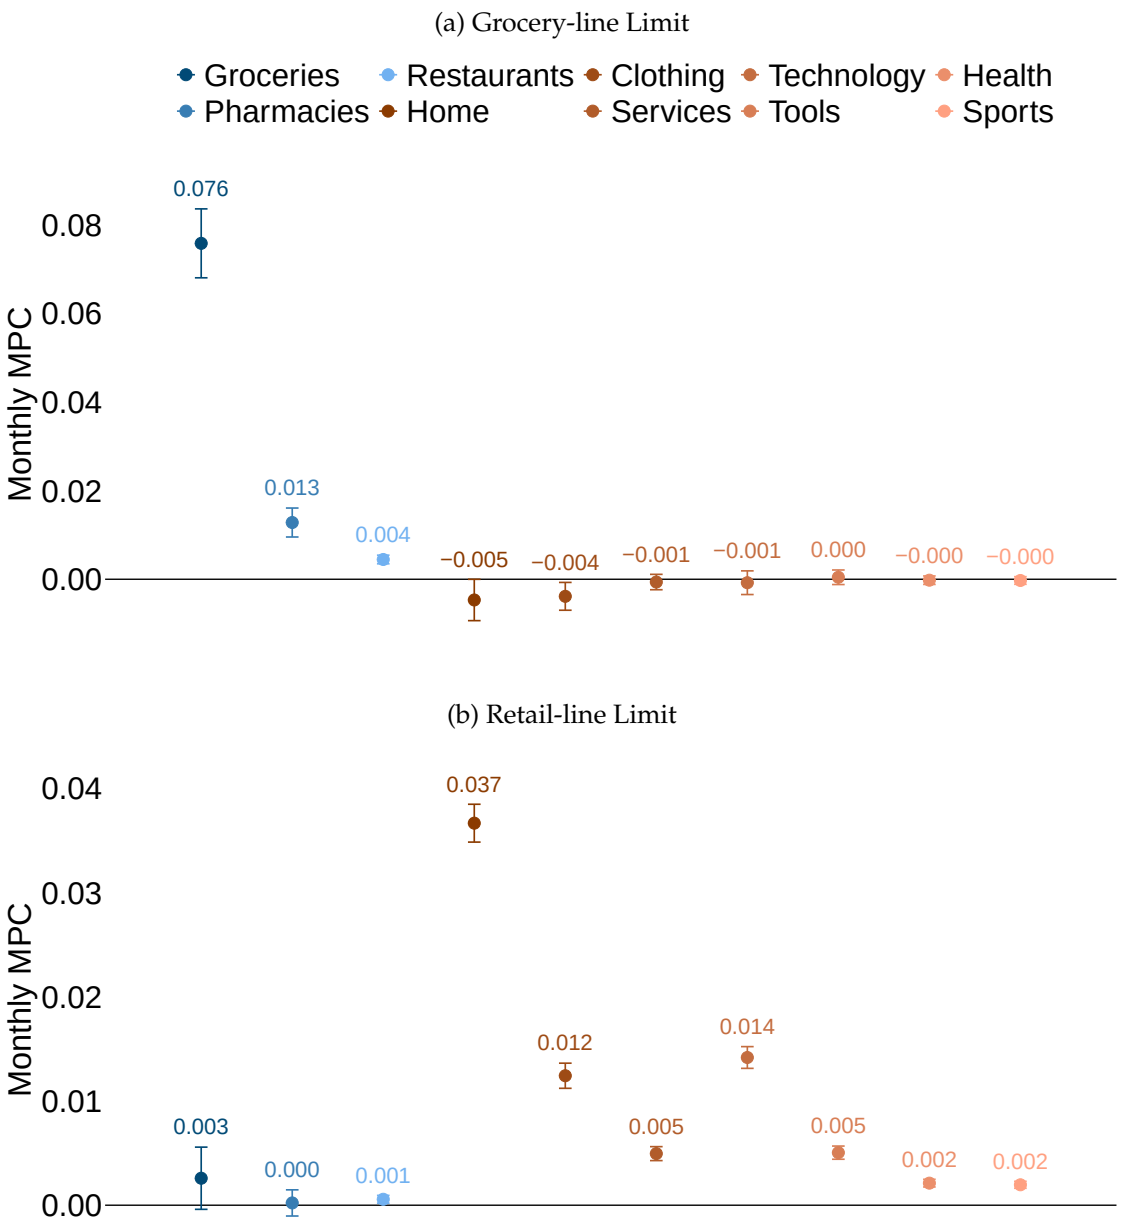
Notes: Figure plots heterogeneous 2SLS estimates of Equations 6–7, allowing the marginal propensity to consume to vary by tercile of baseline average total credit utilization (across both Línea Cotidiana and Línea Principal) over the 30 days prior to October 1, 2025. Panel (a) uses average monthly Línea Cotidiana purchases from October 3, 2025 to September 3, 2026 as the dependent variable; panel (b) uses average monthly Línea Principal purchases over the same time horizon. The endogenous regressors are post-treatment LC and LP credit limits (and their interactions with tercile indicators), instrumented by treatment-arm assignments (and their interactions with tercile indicators). Each point represents the estimated MPC—dollars of spending per dollar of credit limit—for a given tercile-endogenous regressor pair, with the "Low" tercile as the base group. Vertical bars indicate 95% confidence intervals constructed from standard errors clustered at the user level. All specifications include randomization strata fixed effects. Outcome and limit variables are winsorized at the 0.1st and 99.9th percentiles. Joint Wald tests of equality across the three tercile coefficients reject at $p < 0.001$ for both same-line effects. For cross-line effects, $p = 0.253$ for the LP Limit $\rightarrow$ LC Purch pair and $p = 0.037$ for the LC Limit $\rightarrow$ LP Purch pair. N = 209,996.

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



Notes: Figure plots Group Average Treatment Effects (GATES) from the machine-learning heterogeneity procedure of Chernozhukov, Demirer, Duflo, and Fernández-Val, 2025, using a random forest to construct the proxy predictor of the conditional average treatment effect. Users are sorted into quintiles of the predicted effect and the treatment effect is estimated within each quintile. 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 two combined arms are excluded from all panels. Panels (a) and (d) are the same-line response and panels (b) and (c) the cross-line spillover. The outcome is average monthly purchases from the start of the experiment on October 3, 2025 through September 3, 2026. Points are medians over 100 data splits; vertical bars indicate 95% confidence intervals. Regressions include randomization strata fixed effects and a stratum-specific propensity score. Quintiles are formed separately within each panel from that panel's own proxy, so a given quintile index does not identify the same users across panels. N = 87,578 in panels (a) and (b) and 87,631 in panels (c) and (d). Corresponding average and heterogeneous treatment effects are reported in Table A7.

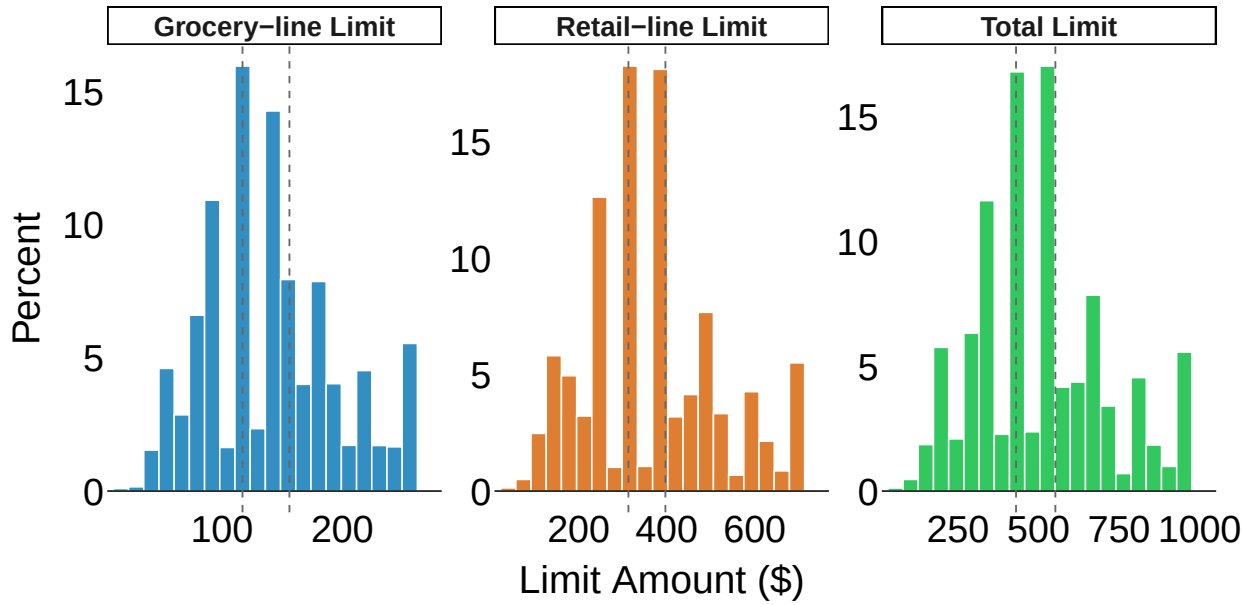
Figure 9. Marginal Propensities to Consume by Product Category



Notes: Figure plots 2SLS estimates of Equations 6–7 with merchant-category purchase outcomes. The dependent variable is cumulative total purchases in each merchant category. The endogenous regressors are post-treatment grocery-line and retail-line credit limits as of October 6, 2025, instrumented by the six treatment-arm assignment dummies. Each point represents the estimated MPC—cents of total spending per dollar of credit limit—for a given category–endogenous regressor pair. The panels group estimates by endogenous regressor (grocery-line limit, retail-line limit); within each group, colors distinguish merchant categories. grocery-line categories (groceries, pharmacies, restaurants) are shown first in the legend, followed by retail-line categories (home, clothing, services, technology, tools, health, sports). Vertical bars indicate 95% confidence intervals constructed from standard errors clustered at the user level. All specifications include randomization strata fixed effects. Outcome and limit variables are winsorized at the 0.1st and 99.9th percentiles. Purchases are measured as average monthly spending from October 3, 2025 to September 3, 2026. N = 209,996. Full estimates reported in Appendix Table A13.

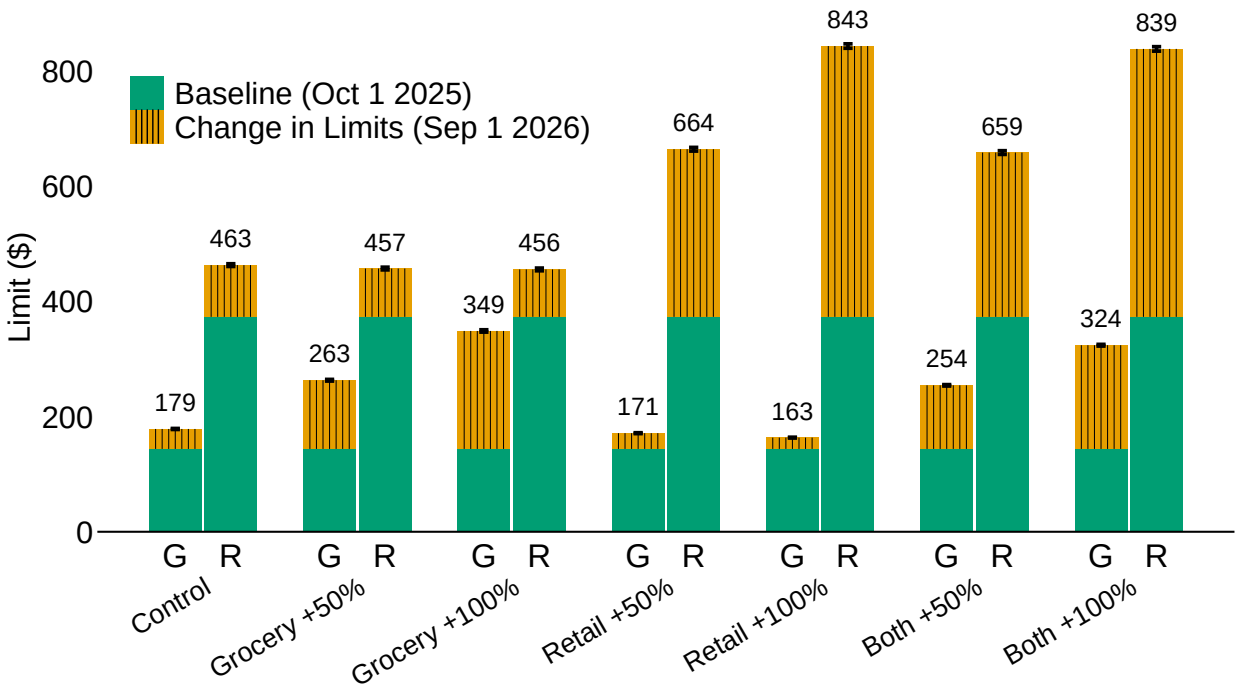
A Appendix Exhibits

Figure A1. Distribution of Baseline Limits



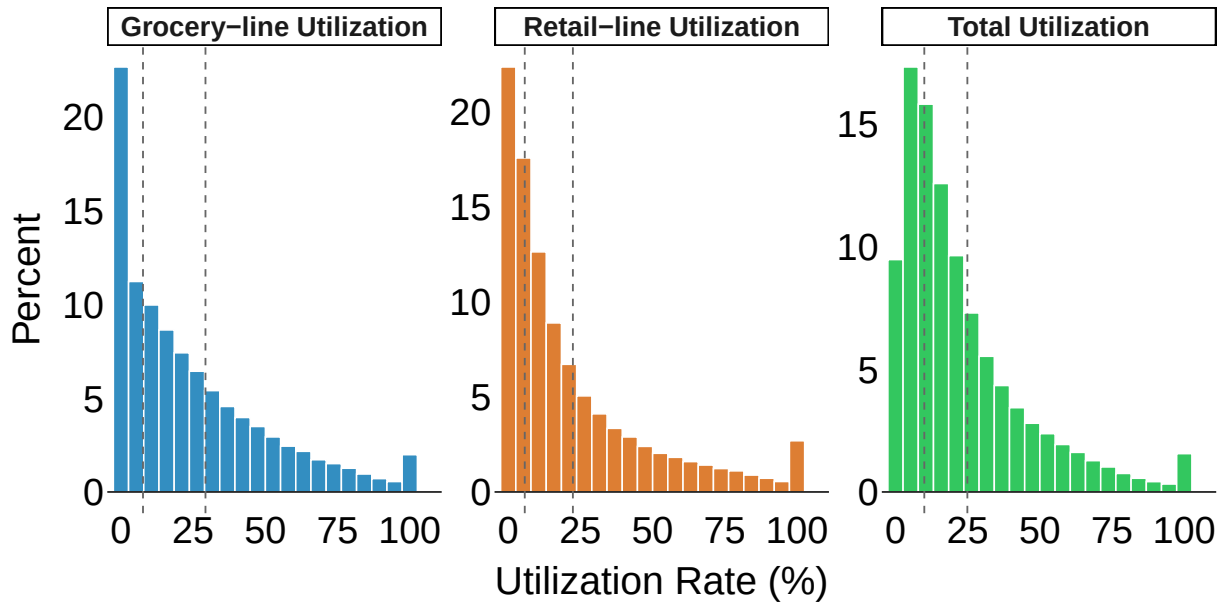
Notes: Figure plots the distribution of baseline credit-limit amounts across Cashea users, separately for the grocery line, the retail line, and the combined total (Total = Grocery Line + Retail Line). Each panel shows a histogram with 20 bins, where the y-axis reports the share of observations in each bin as a percentage of the total sample. Limits are measured in USD as of October 1, 2025, and are winsorized at the 95th percentile within each limit type. N = 209,996 for the grocery-line panel, N = 209,996 for the retail-line panel, and N = 209,996 for the Total panel.

Figure A2. Credit-Limit Intervention over Experimental Period



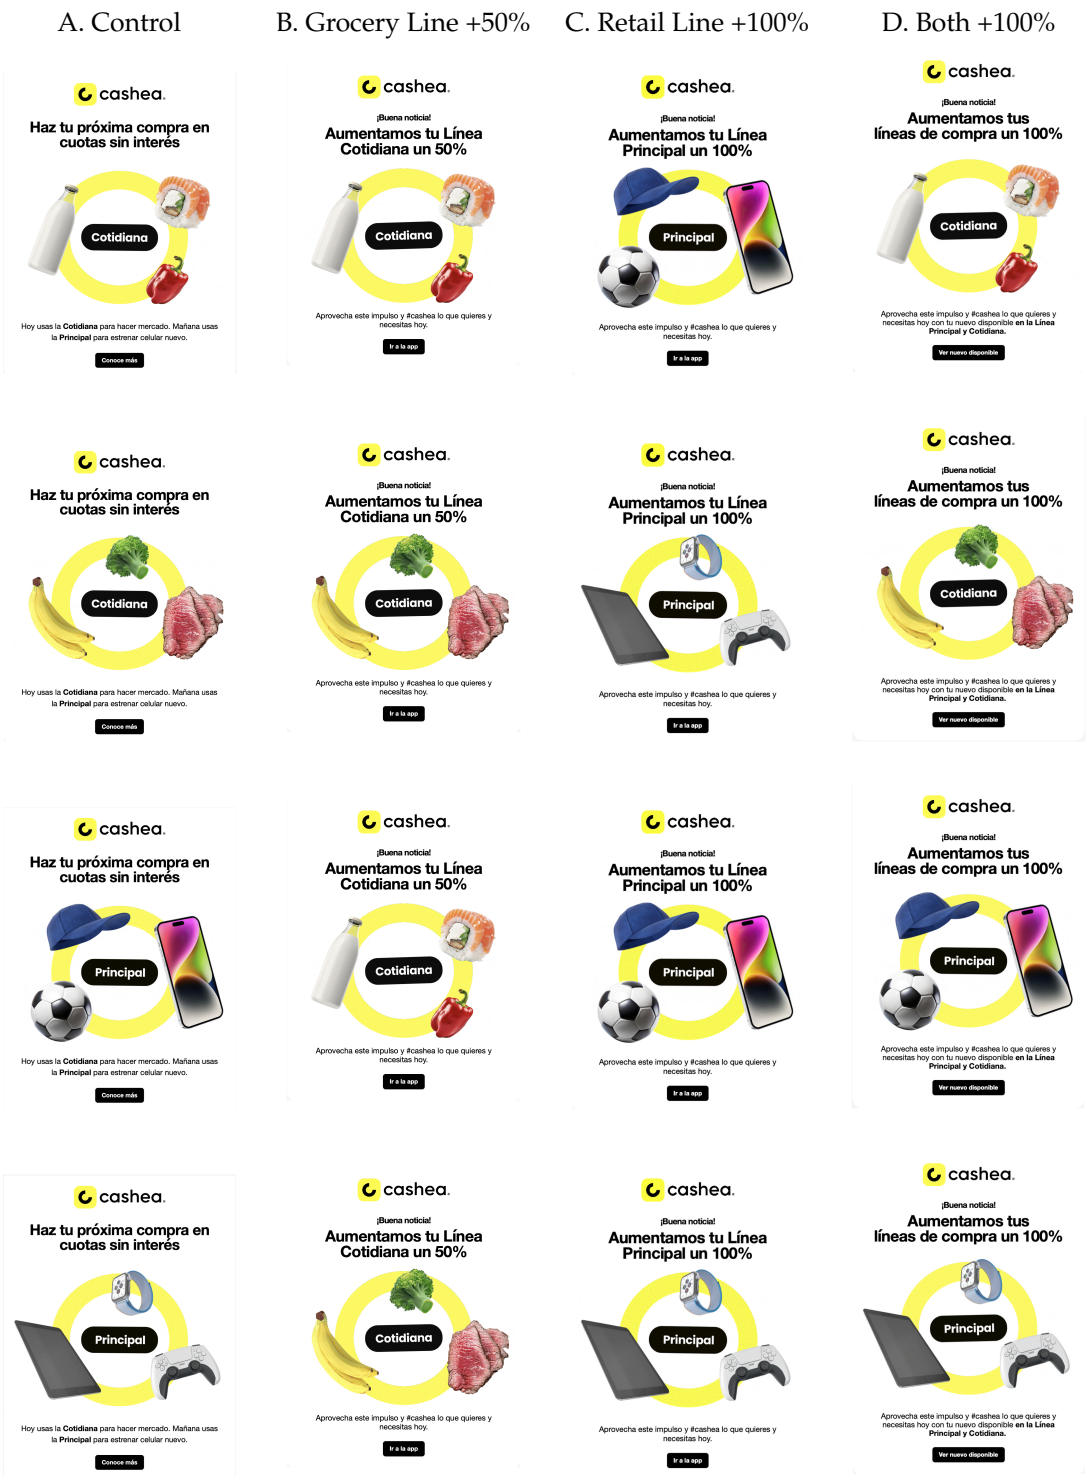
Notes: Bars show mean credit limits in USD by treatment arm. "G" denotes grocery line and "R" denotes retail line. Green bars reflect baseline means as of October 1, 2025. The sum of orange and green bars reflect limits as of September 1, 2026, ten months after the experiment launched. Small deviations from exact 0%, 50%, or 100% increases reflect natural changes in credit limits according to Cashea's levels system. N = 209,988.

Figure A3. Distribution of Baseline Utilization



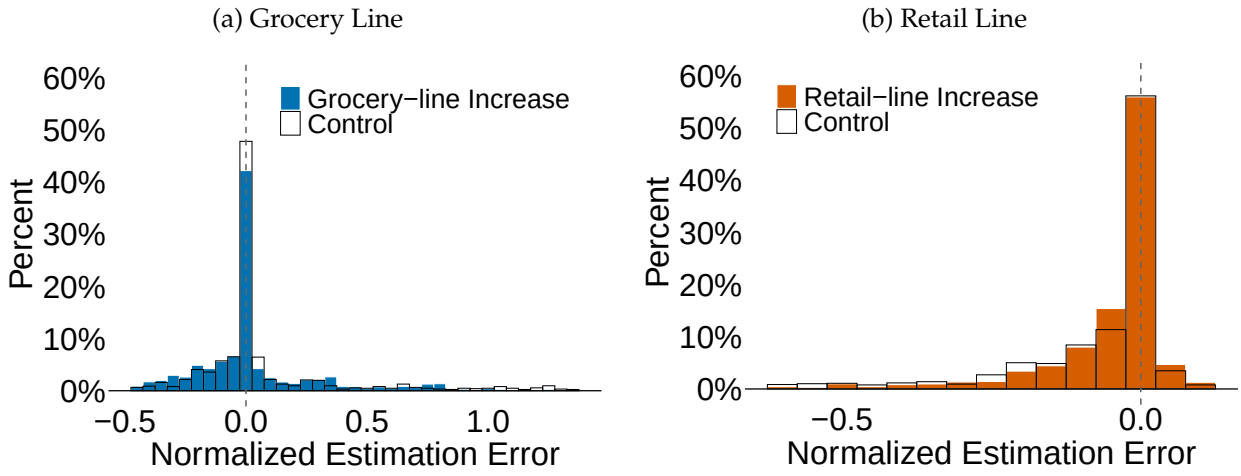
Notes: Figure plots the distribution of baseline credit utilization rates across Cashea users, separately for the grocery line, the retail line, and the combined total. Each panel shows a histogram with 20 bins, where the y-axis reports the share of observations in each bin as a percentage of the total sample. Utilization is defined as the 30-day average utilization rate for each line as of October 1, 2025, expressed as a percentage, and is bounded between 0 and 100. N = 209,996 for the grocery-line panel, N = 209,996 for the retail-line panel, and N = 209,996 for the Total panel.

Figure A4. Full Outreach Messaging



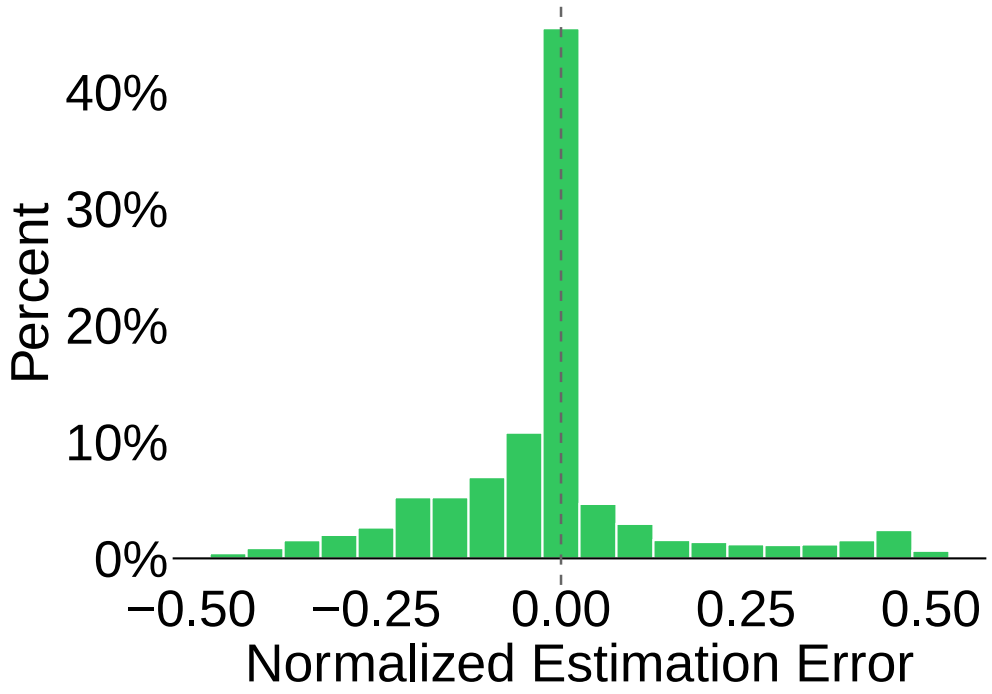
Notes: Screenshots of email outreach sent to participants, showing three of seven treatment arms as examples. Panel A shows the control-group messages, which promotes Cashea without mentioning a credit-limit increase. Panel B shows the treatment messages for a 50% grocery-line increase. Panel C shows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

Figure A5. Knowledge of Post-Experiment Credit Limits



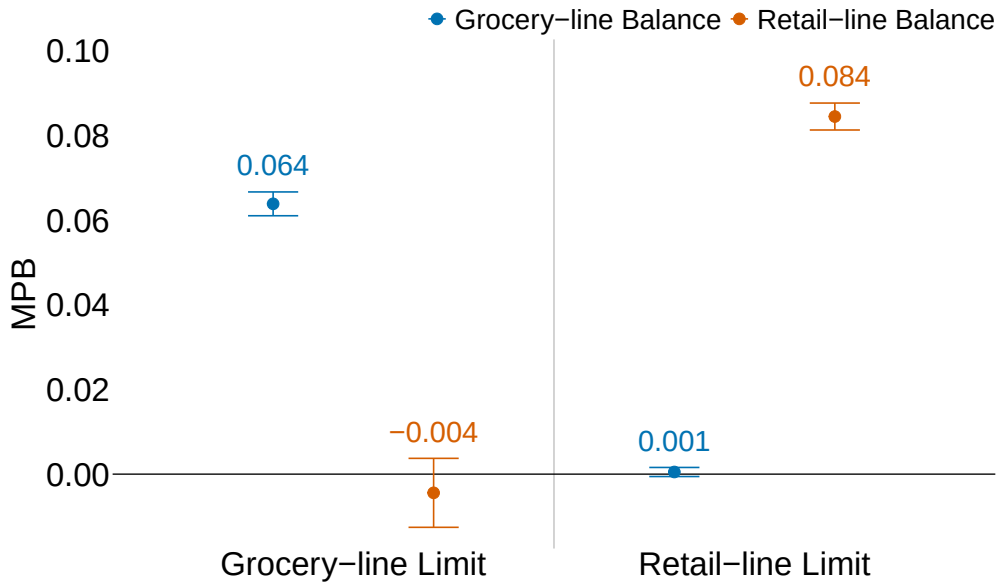
Notes: Figure plots the distribution of normalized credit-limit estimation error in the first follow-up survey, defined as $(\hat{L}_{ig} - L_{ig})/L_{ig}$, where $\hat{L}_{ig}$ is the respondent's self-reported limit on line g and L_{ig} is the corresponding administrative limit. Distributions are trimmed at the 5th and 95th percentiles; bin width is 0.05. The mass at zero indicates exact knowledge of the post-experiment credit limit. N = 4,229 for the grocery-line panel and N = 4,419 for the retail-line panel.

Figure A6. Baseline Total Limit Estimation Error



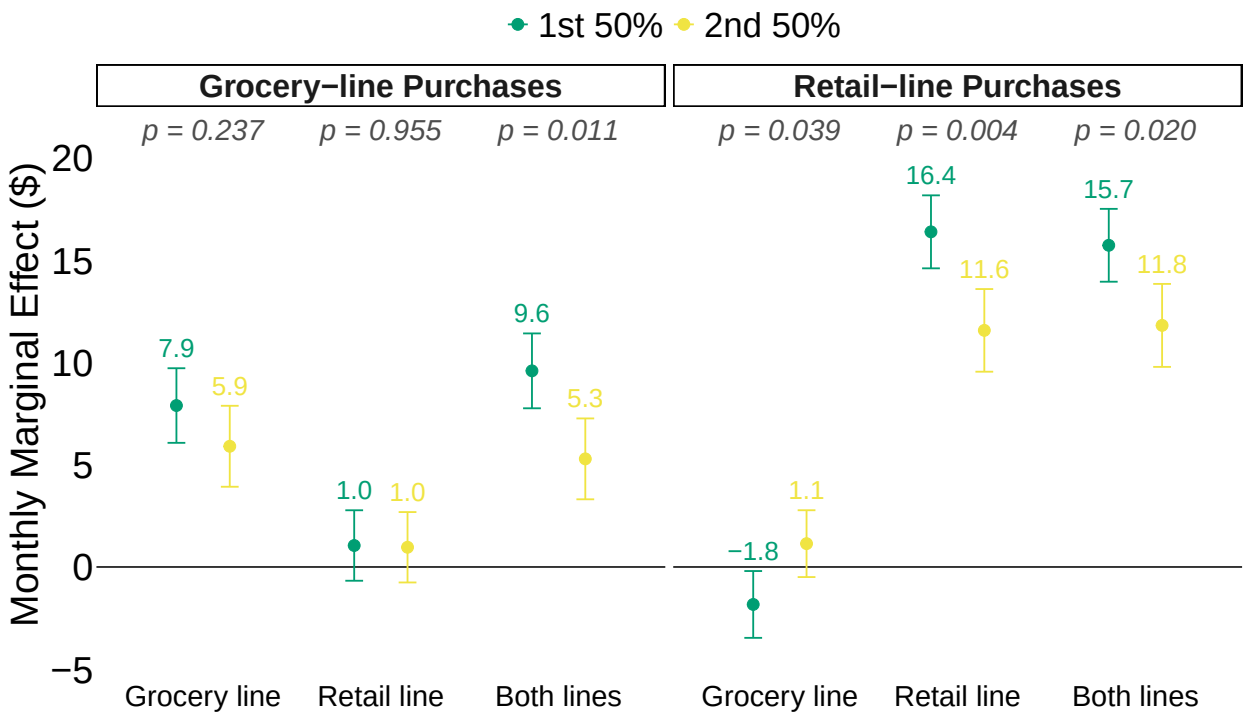
Notes: Figure plots the distribution of normalized credit-limit estimation error at baseline, defined as $(\hat{L}_i - L_i)/L_i$, where $\hat{L}_i$ is the respondent's self-reported total credit limit summing across both lines, and L_i is the corresponding administrative limit recorded on October 1, 2025. Total limit is the sum of the grocery-line and retail-line limits. The sample is restricted to baseline survey respondents who can be matched to administrative records and who report a positive administrative total limit. The distribution is trimmed at the 5th and 95th percentiles of the normalized error. Each bar represents a bin of width 0.05 (5 percentage points). The y-axis reports the share of observations within each bin as a percentage of the trimmed sample. The dashed red vertical line marks zero, indicating exact recall. Positive (negative) values indicate overestimation (underestimation) of one's credit limit. The baseline survey was fielded the week prior to the experiment's launch; administrative limits are drawn from administrative data as of October 1, 2025. N = 7,709.

Figure A7. Marginal Propensities to Consume out of Credit-Limit Increases: Average Balances



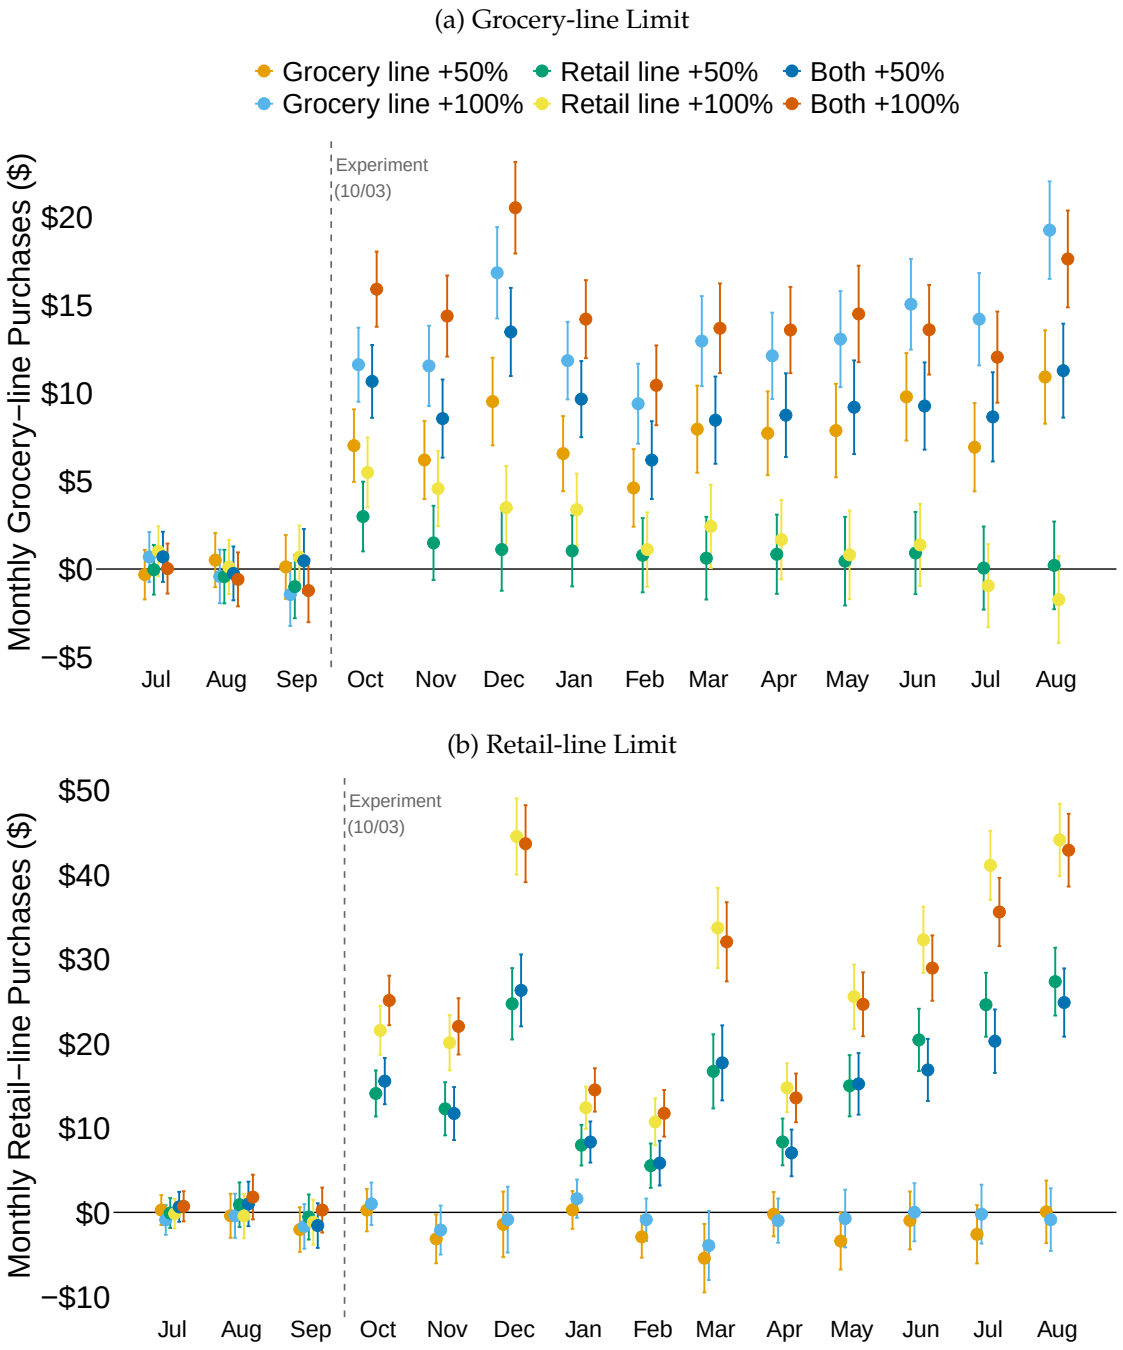
Notes: Figure plots instrumental-variables estimates analogous to Figure 3, with average daily balances on each line over the post-treatment period from October 3, 2025 to September 3, 2026 as the dependent variable. Post-experiment credit limits are instrumented with the six treatment assignment indicators; regressions control for randomization strata fixed effects. Bars represent 95% confidence intervals. N = 209,996 for the grocery-line model and N = 209,996 for the retail-line model. Full estimates are reported in Appendix Table A10.

Figure A8. Marginal Treatment Effects



Notes: Figure plots treatment effect estimates of the marginal effect of credit-limit increases on purchases, measured in dollars of spending per treatment arm, separately for Línea Cotidiana (LC) and Línea Principal (LP) purchases. Each panel corresponds to a different outcome: LC purchases (left) and LP purchases (right). Within each panel, the x-axis distinguishes by which limit was increased (LC line, LP line, or both lines), and the two bars show the *1st 50% marginal effect* (the effect of a 50% limit increase relative to control) and the *2nd 50% marginal effect* (the additional effect of increasing the limit by a further 50%, computed as the difference between the 100% and 50% treatment coefficients). Includes randomization strata fixed effects. Standard errors are clustered at the user level. The p-value annotated above each pair of bars tests $H_0 : \beta_{50} = \beta_{100} - \beta_{50}$ (i.e., equal marginal effects across the two 50-percentage-point increments), based on a two-sided t-test of the linear combination $2\beta_{50} - \beta_{100} = 0$. Standard errors for the 2nd 50% marginal effect are computed via the delta method. Post-experiment limit levels are winsorized at the 0.1st and 99.9th percentiles. Purchases are measured as average monthly spending from October 3, 2025 to September 3, 2026.

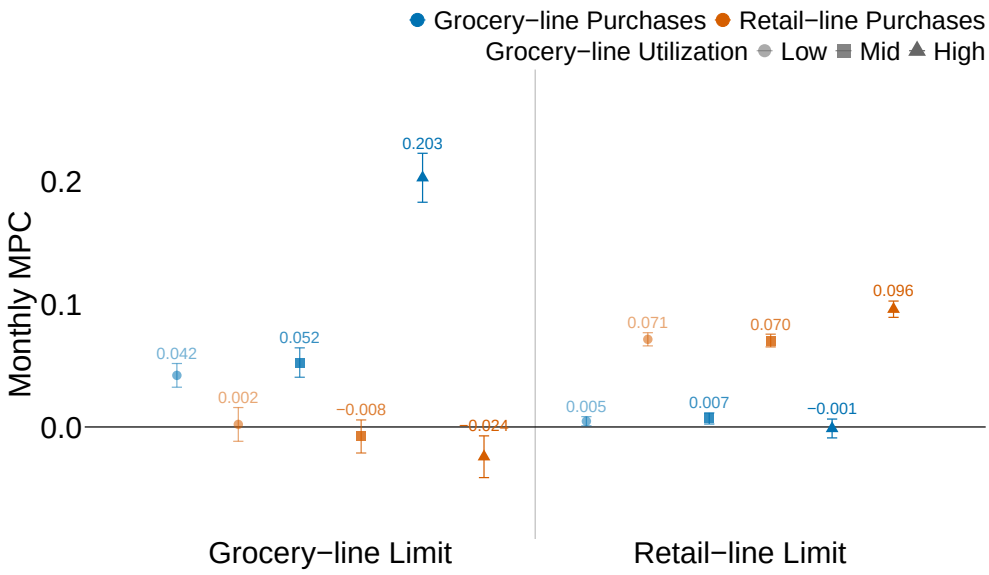
Figure A9. Dynamic Effects by Treatment Arm on Purchases



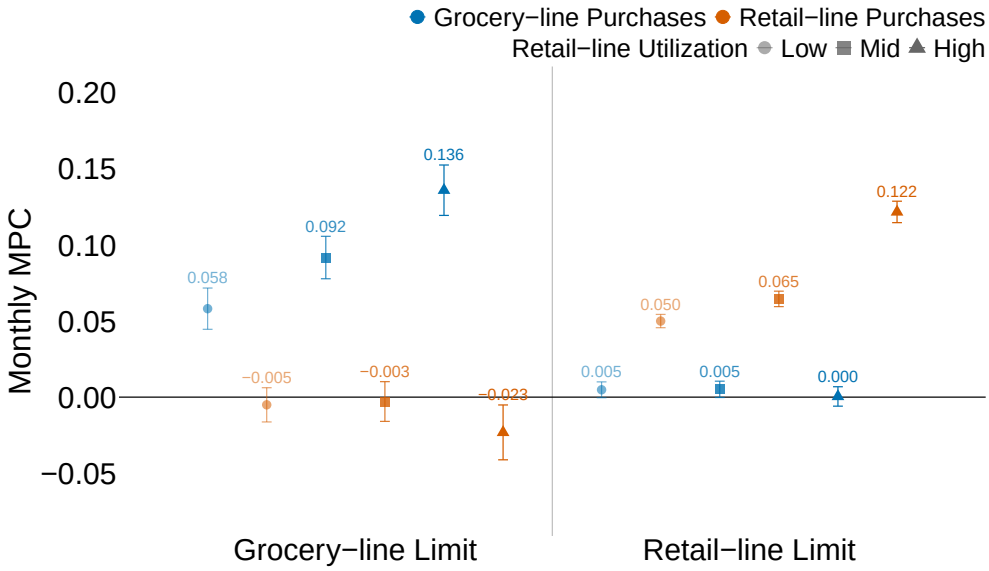
Notes: Figure plots monthly OLS treatment effect estimates from Equation 8, estimated separately for each month from July 2025 through August 2026. The left panel shows *Línea Cotidiana* (LC) purchases and the right panel shows *Línea Principal* (LP) purchases, both measured in dollars. Each point represents the estimated treatment effect for a given arm relative to the control group; vertical bars indicate 95% confidence intervals. Regressions include randomization strata fixed effects and the six treatment-arm indicators (LC +50%, LC +100%, LP +50%, LP +100%, Both +50%, Both +100%). Users with no purchasing activity in a given month are assigned a value of zero. All outcomes are winsorized at the 0.1st and 99.9th percentiles within each month. Standard errors are clustered at the user level. The vertical dashed line marks the start of the experiment. N = 209,996, based on the sample present in a representative post-experiment month.

Figure A10. Marginal Propensity to Consume by Baseline LC/LP Utilization

(a) By Línea Cotidiana Utilization Tercile

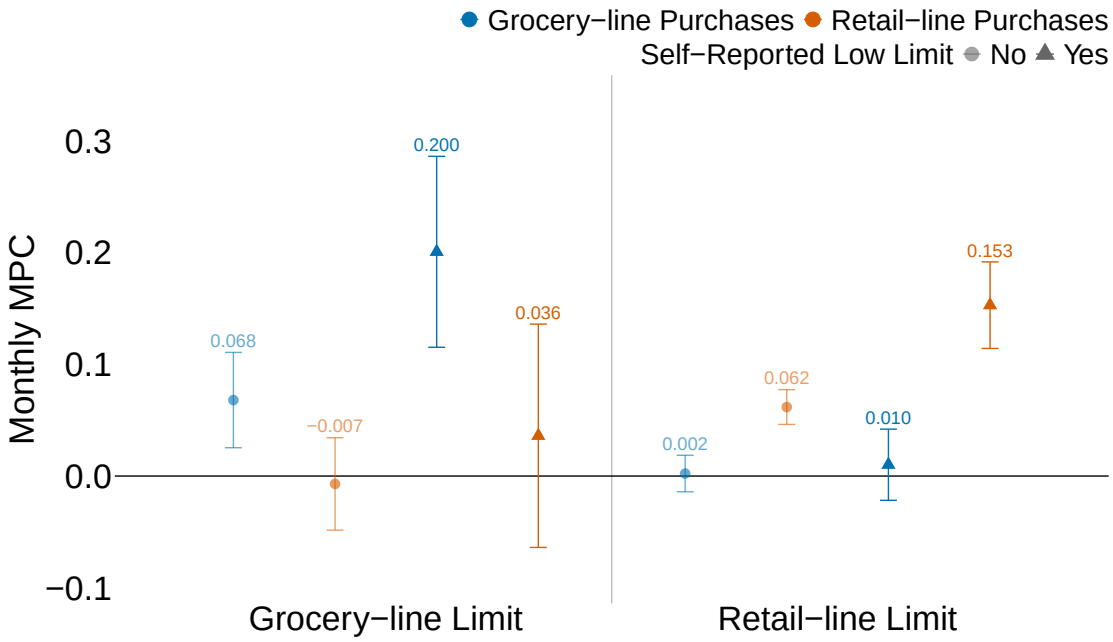


(b) By Línea Principal Utilization Tercile



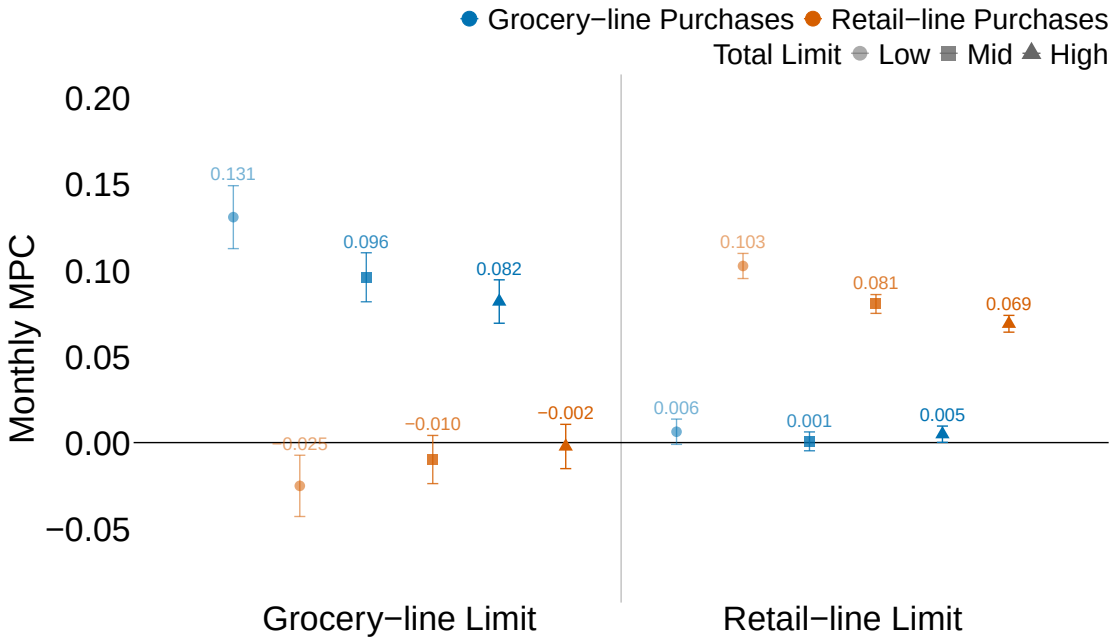
Notes: Figure plots heterogeneous 2SLS estimates of Equations 6–7, allowing the marginal propensity to consume to vary by tercile of baseline credit utilization over the 30 days prior to October 1, 2025. Panel (a) splits by tercile of baseline Línea Cotidiana utilization; panel (b) splits by tercile of baseline Línea Principal utilization. Within each panel, points correspond to LC and LP average monthly purchases from October 3, 2025 to September 3, 2026 as separate dependent variables. The endogenous regressors are post-treatment LC and LP credit limits (and their interactions with tercile indicators), instrumented by treatment-arm assignments (and their interactions with tercile indicators). Each point represents the estimated MPC—dollars of spending per dollar of credit limit—for a given tercile-endogenous regressor pair, with the "Low" tercile as the base group. Vertical bars indicate 95% confidence intervals constructed from standard errors clustered at the user level. All specifications include randomization strata fixed effects. Outcome and limit variables are winsorized at the 0.1st and 99.9th percentiles. Joint Wald tests of equality across the three tercile coefficients reject at $p < 0.001$ for all same-line effects. For cross-line effects in panel (a): $p = 0.199$ for LP Limit $\rightarrow$ LC Purch and $p = 0.061$ for LC Limit $\rightarrow$ LP Purch. In panel (b): $p = 0.486$ for LP Limit $\rightarrow$ LC Purch and $p = 0.168$ for LC Limit $\rightarrow$ LP Purch. N = 209,996 in panel (a) and N = 209,996 in panel (b).

Figure A11. Marginal Propensity to Consume by Self-Reported Credit Constraint



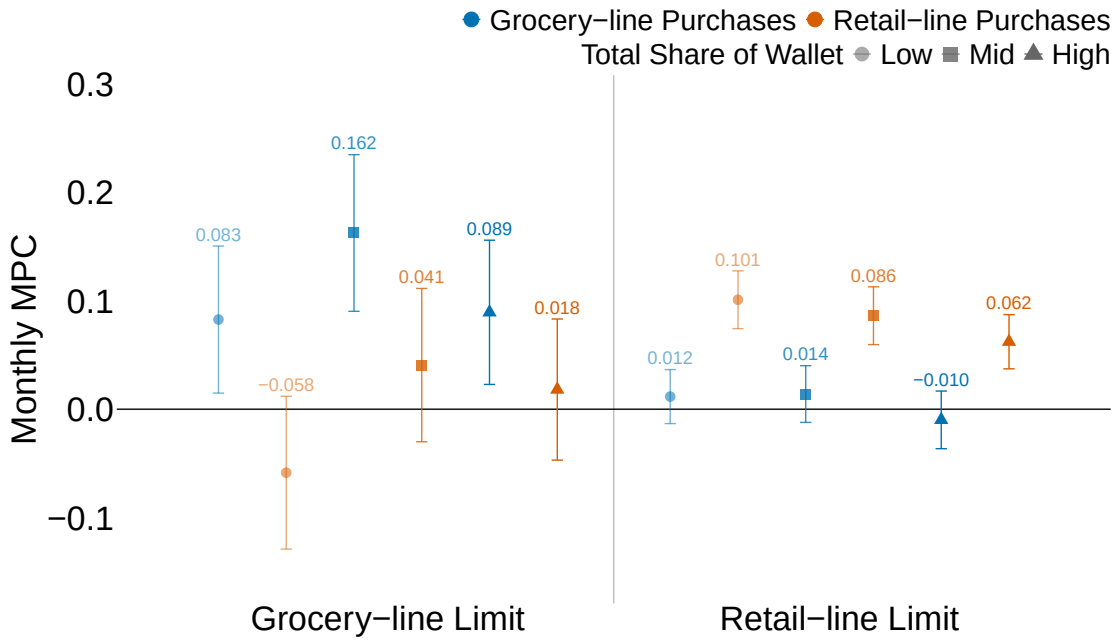
Notes: Figure plots heterogeneous 2SLS estimates of Equations 6–7, allowing the marginal propensity to consume to vary by whether the respondent self-reported their credit limit as a binding constraint in the baseline survey. Panel (a) uses average monthly Línea Cotidiana purchases from October 3, 2025 to September 3, 2026 as the dependent variable and conditions on whether the respondent cited a low LC limit as a reason for spending off Cashea; panel (b) uses average monthly Línea Principal purchases over the same time horizon and conditions on the analogous LP constraint indicator. The endogenous regressors are post-treatment LC and LP credit limits (and their interactions with the binary constraint indicator), instrumented by treatment-arm assignments (and their interactions with the constraint indicator). Each point represents the estimated MPC—dollars of spending per dollar of credit limit—for the constrained and unconstrained groups, with the unconstrained group ("No") as the base. Vertical bars indicate 95% confidence intervals constructed from standard errors clustered at the user level. All specifications include randomization strata fixed effects. Outcome and limit variables are winsorized at the 0.1st and 99.9th percentiles. N = 8,728 in panel (a) and N = 8,728 in panel (b).

Figure A12. Marginal Propensity to Consume by Baseline Credit Limit



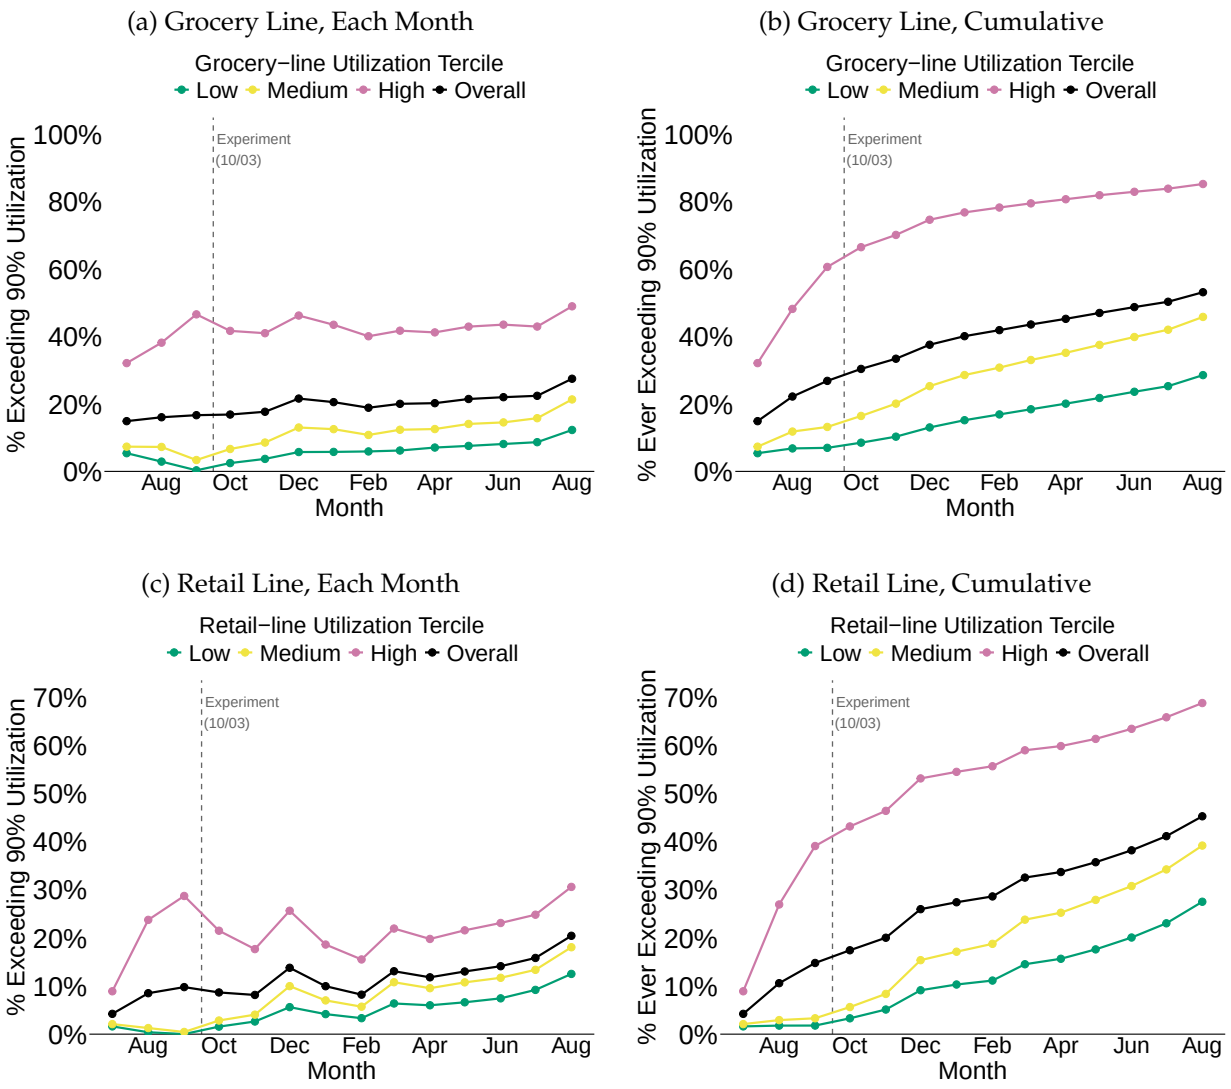
Notes: Figure plots heterogeneous 2SLS estimates of Equations 6–7, allowing the marginal propensity to consume to vary by tercile of baseline total credit limit (Línea Cotidiana plus Línea Principal) as of October 1, 2025. Panel (a) uses average monthly Línea Cotidiana purchases from October 3, 2025 to September 3, 2026 as the dependent variable; panel (b) uses average monthly Línea Principal purchases over the same time horizon. The endogenous regressors are post-treatment LC and LP credit limits (and their interactions with tercile indicators), instrumented by treatment-arm assignments (and their interactions with tercile indicators). Each point represents the estimated MPC—dollars of spending per dollar of credit limit—for a given tercile-endogenous regressor pair, with the "Low" tercile as the base group. Vertical bars indicate 95% confidence intervals constructed from standard errors clustered at the user level. All specifications include randomization strata fixed effects. Outcome and limit variables are winsorized at the 0.1st and 99.9th percentiles. N = 209,996.

Figure A13. Marginal Propensity to Consume by Baseline Cashea Share of Wallet



Notes: Figure plots heterogeneous 2SLS estimates of Equations 6–7, allowing the marginal propensity to consume to vary by tercile of baseline Cashea share of wallet. Share of wallet is defined as total Cashea spending (LC plus LP) as a fraction of total spending across Cashea and non-Cashea credit, expressed as a percentage, and is computed from self-reported spending in the baseline survey. Panel (a) uses average monthly Línea Cotidiana purchases from October 3, 2025 to September 3, 2026 as the dependent variable; panel (b) uses average monthly Línea Principal purchases over the same time horizon. The endogenous regressors are post-treatment LC and LP credit limits (and their interactions with tercile indicators), instrumented by treatment-arm assignments (and their interactions with tercile indicators). Each point represents the estimated MPC—dollars of spending per dollar of credit limit—for a given tercile-endogenous regressor pair, with the "Low" tercile as the base group. Vertical bars indicate 95% confidence intervals constructed from standard errors clustered at the user level. All specifications include randomization strata fixed effects. Outcome and limit variables are winsorized at the 0.1st and 99.9th percentiles. N = 8,711, smaller than the full experimental sample because share of wallet is only observed for baseline survey respondents.

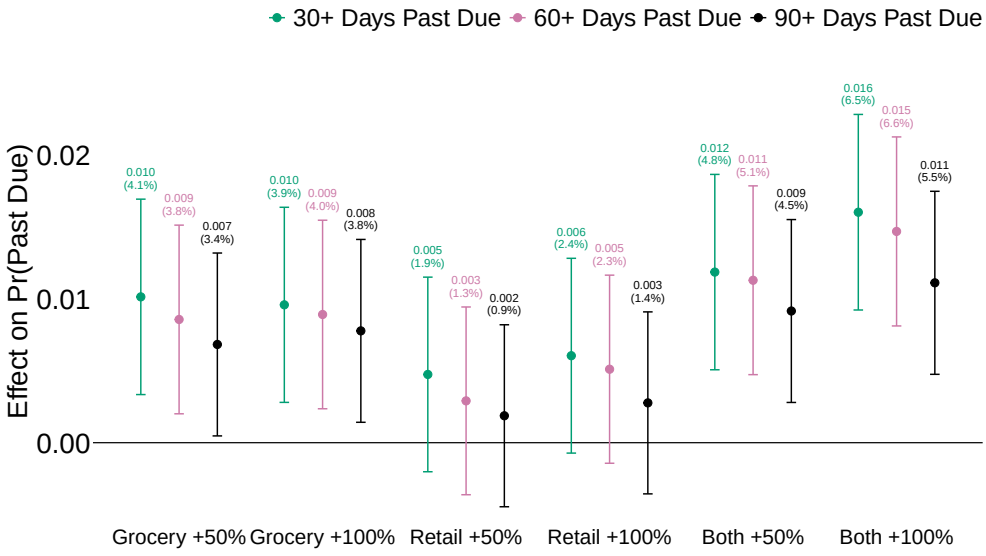
Figure A14. Monthly Post-Experiment Constraint-Hitting Rates by Baseline Utilization



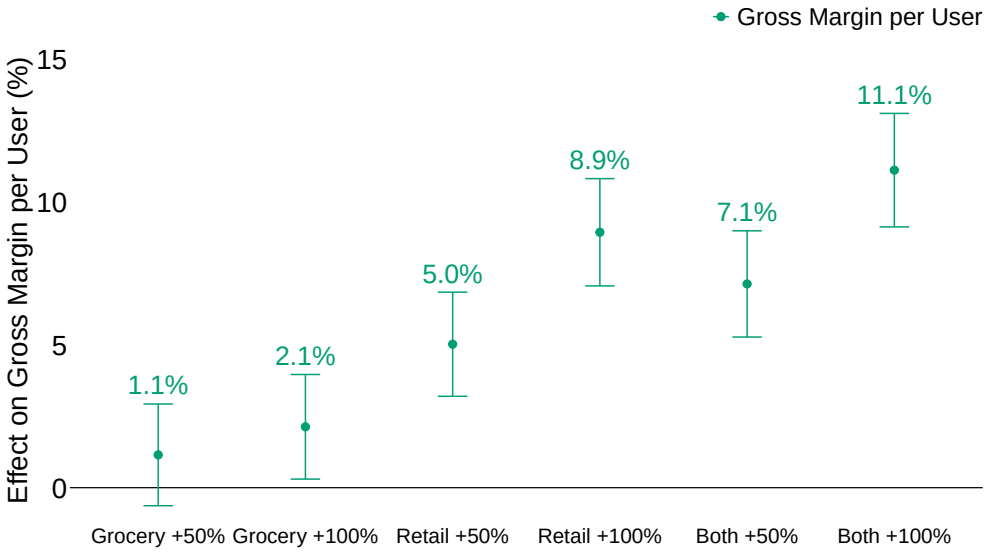
Notes: Figure plots the share reaching 90% utilization at some point *within* the calendar month in the left column and the share that has *ever* reached 90% since July 2025 in the right column. All panels use data spanning from July 2025 to August 2026. Within each row the two panels share a y-axis. Terciles are of baseline average 30-day utilization on the respective line as of October 1, 2025. Dashed line marks randomization (2025-10-03). The sample is restricted to the control group. N = 30,001.

Figure A15. Treatment Effects by Arm on Installments Past Due and Monthly Profit

(a) Installments Past Due



(b) Gross Margin per User



Notes: Figure plots treatment effects by arm on the probability of being past due as of September 3, 2026 in panel (a) and on average gross margin per user from October 3, 2025 through September 3, 2026 in panel (b). In panel (a), the dependent variable is an indicator for holding at least one installment 30 or more, 60 or more, or 90 or more days past due as of September 3, 2026; green denotes the 30-day threshold, pink the 60-day threshold, and black the 90-day threshold. Each point estimate is labeled with its coefficient and the percent change relative to the control-group mean in parentheses. In panel (b), the dependent variable is cumulative gross margin from randomization through that date. Gross margin is calculated using a proprietary internal formula at Cashea; it is winsorized at the 0.1st and 99.9th percentiles. Point estimates in panel (b) are shown as the percent change relative to the control-group mean. We cannot disclose the control means as this information is confidential. Bars represent 95% confidence intervals. All estimates are from OLS regressions of the form described in Section 5, with the control group as the omitted category, controlling for randomization strata fixed effects. Standard errors are clustered at the user level. N = 209,996 in panel (a) and N = 209,996 in panel (b).

Table A1. Baseline Spending by Credit Line

Category	N	$	% Total $	% Total N
Panel A: Retail line				
Clothing	411,518	22,228,356	30.9%	46.4%
Home	224,417	25,368,935	35.3%	25.3%
Services	77,455	6,389,242	8.9%	8.7%
Tools	52,001	4,563,430	6.4%	5.9%
Sports	40,389	2,947,698	4.1%	4.6%
Health	38,042	3,175,190	4.4%	4.3%
Technology	31,774	6,251,948	8.7%	3.6%
Other	11,000	897,291	1.2%	1.2%
Total	886,596	71,822,088	100.0%	100.0%
Panel B: Grocery line				
Groceries	970,746	59,799,814	82.2%	75.0%
Pharmacies	176,300	9,096,084	12.5%	13.6%
Restaurants	110,734	3,023,365	4.2%	8.6%
Other	36,094	804,861	1.1%	2.8%
Total	1,293,874	72,724,124	100.0%	100.0%

Notes: Table reports the distribution of orders by store category on each credit line for the experimental sample of 209,996 users during the 90 days prior to the experiment (July 5 to October 2, 2025).

Table A2. 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: Table summarizes the experimental design of the credit-limit intervention. Each row corresponds to one treatment arm. *N* reports the number of users assigned to that arm. *Retail line (Oct 1)* and *Grocery line (Oct 1)* report average baseline credit limits (in USD), respectively, measured on October 1, 2025, 2 days prior to treatment assignment. *Intended Retail line* and *Intended Grocery line* report the intended percentage increase in each credit limit for users in that arm; "—" indicates that the corresponding limit was not targeted for change. *Realized* Δ *Retail line* and *Realized* Δ *Grocery line* report the average realized dollar change in each credit limit between October 1 and October 6, 2025. Small deviations from exact 0%, 50%, or 100% increases reflect that some users also naturally received credit-limit increases or decreases due to moving up or down a level. All dollar amounts are in USD and are rounded to the nearest cent. Total N = 209,996.

Table A3. 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: Table shows coefficients for follow-up survey group (both surveys) relative to the non-survey group. All average variables in this table are calculated by taking the mean of the daily variable across the previous 30 days. Data as of October 1, 2025. P-value is derived from the joint F test with the following null hypotheses $H_0^k : \beta_k = 0$. On the other hand, Omnibus F-Test is derived from the regression $T_i^j = \alpha^j + \sum_{k=1}^{7} \beta_k^j X_{ki} + \gamma_s^j + \epsilon_i^j$ where the f-statistic results from $H_0 : \beta_1^j = \beta_2^j = ... = \beta_7^j = 0$. N = 209,996.

Table A4. 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: Table shows coefficients for each treatment group relative to the control group. All average variables in this table are calculated by taking the mean of the daily variable across the previous 30 days as of October 1, 2025 (pre-experiment launch). Sample restricted to follow-up survey respondents; a respondent to both follow-up waves is counted once per wave (N = 28,511 person-wave observations). P-value is derived from the joint F test with the following null hypotheses $H_0^k : \beta_{k1} = \beta_{k2} = ... = \beta_{k6} = 0$. On the other hand, Omnibus F-Test is derived from the regression $T_i^j = \alpha^j + \sum_{k=1}^{7} \beta_k^j X_{ki} + \gamma_s^j + \epsilon_i^j$ where the f-statistic results from $H_0 : \beta_1^j = \beta_2^j = ... = \beta_7^j = 0$ *** indicates $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

Table A5. Survey Demographics

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
Age	18,616	40.41	9.93	21.00	33.00	40.00	47.00	65.00
Share Female	18,532	0.56	0.50	0.00	0.00	1.00	1.00	1.00

Notes: Table reports summary statistics for demographic characteristics of Cashea users in the study sample. Monthly income is drawn from the baseline survey only and is winsorized at the 95th percentile; age and share female are computed by pooling the baseline and first follow-up surveys. Age is winsorized at 18 and 100. Share female is defined as the fraction of respondents identifying as female, restricting to binary male/female responses; respondents reporting another gender or preferring not to say are excluded from this row. Monthly income is reported in USD. N = 22,193.

Table A6. 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)	0.007	-0.026		
	(0.024)	(0.024)		
Retail-line Limit (Post)	0.009	0.018		
	(0.009)	(0.009)		
Observations	23,436	23,436		
Panel B: On Cashea				
Grocery-line Limit (Post)	0.065	-0.032	0.084	-0.007
	(0.019)	(0.024)	(0.004)	(0.005)
Retail-line Limit (Post)	0.014	0.048	0.007	0.062
	(0.007)	(0.009)	(0.002)	(0.002)
Observations	23,436	23,436	209,996	209,996
Panel C: Total				
Grocery-line Limit (Post)	0.073	-0.058		
	(0.037)	(0.042)		
Retail-line Limit (Post)	0.023	0.066		
	(0.014)	(0.016)		
Observations	23,436	23,436		

Notes: Table reports 2SLS estimates of Equations 6–7. The endogenous regressors are post-experiment LC and LP credit limits, instrumented by the six treatment-arm assignment dummies. Each row reports the estimated MPC for a given endogenous regressor (grocery-line limit, retail-line limit); columns distinguish the outcome line (grocery-line spending, retail-line spending). Survey columns use self-reported spending, normalized to a 30-day equivalent from a pooled sample of first- and second-wave follow-up survey respondents (observations are respondent-waves). Admin columns re-estimate the same specification on the full randomized sample (one observation per user), using administrative Cashea purchases over the survey recall windows: October 1, 2025 through the median first-wave survey date (December 4, 2025), normalized to 30 days as above, and the 30 days ending at the median second-wave survey date (May 14, 2026), averaged across the two windows. These definitions align perfectly with the survey recall windows. Administrative data only observe On-Cashea spending, so Admin columns appear in Panel B only. Panels A–C report estimates for Off-Cashea, On-Cashea, and total spending, respectively. Survey spending outcomes are winsorized at the 95th percentile, while administrative spending outcomes are winsorized at 0.1st and 99.9th percentiles. Standard errors, reported in parentheses, are clustered at the user level. All specifications include randomization strata fixed effects.

Table A7. Average and Heterogeneous Treatment Effects

	DV: Average Monthly Purchases through September 3, 2026			
	Grocery-line Limit Increase		Retail-line Limit Increase	
	Grocery-line Purchases	Retail-line Purchases	Grocery-line Purchases	Retail-line Purchases
ATE	11.63	-0.98	1.61	22.65
	[0.000]	[0.289]	[0.074]	[0.000]
	(9.67, 13.61)	(-2.78, 0.83)	(-0.16, 3.36)	(20.54, 24.77)
HTE	0.267	0.012	0.006	0.342
	[0.000]	[0.808]	[0.907]	[0.000]
	(0.190, 0.347)	(-0.090, 0.108)	(-0.096, 0.107)	(0.255, 0.429)
Observations	87,578	87,578	87,631	87,631
Control Mean	123.33	124.84	123.33	124.84

Notes: Table reports Best Linear Predictor estimates, following Chernozhukov, Demirer, Duflo, and Fernández-Val, 2025, using a random forest to construct the proxy predictor. The ATE row reports β_1, the average treatment effect; the HTE row reports β_2, the coefficient on the interaction between treatment and the demeaned proxy, so rejecting $\beta_2 = 0$ is evidence that the treatment effect varies with baseline characteristics. Columns 1 and 2 restrict the sample to the control group and the two grocery-line arms; columns 3 and 4 to the control group and the two retail-line arms. The two combined arms are excluded. Within each pair, the first column is the same-line response and the second the cross-line spillover. The outcome is average monthly purchases from October 3, 2025 through September 3, 2026. Estimates are medians over 100 data splits; brackets report p-values and parentheses 95% confidence intervals. Regressions include randomization strata fixed effects and a stratum-specific propensity score. Quintile-level estimates are plotted in Figure 8.

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 ($)	780.22	33.98	31.23	28.10	38.41	30.03	81.28	1.82	11,782
		(25.13)	(24.75)	(24.45)	(24.58)	(25.13)	(25.21)	[0.489]	
Other spending ($)	470.73	-8.11	21.23	1.26	16.26	35.69	22.00	1.59	8,814
		(17.21)	(17.84)	(17.39)	(17.51)	(17.76)	(17.66)	[0.578]	
Savings stock ($)	1,264.74	53.45	57.68	-120.12	67.71	-47.19	176.02	1.64	7,851
		(108.74)	(107.94)	(103.63)	(108.08)	(104.90)	(111.42)	[0.578]	
Less savings, stock (binary)	0.36	0.04	0.04	0.00	0.01	0.02	0.00	1.19	7,978
		(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	[0.761]	
Savings flow, last 30 days ($)	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)	[0.761]	
Other debt ($)	131.58	16.25	15.30	20.87	15.01	12.30	28.93	1.17	7,497
		(11.21)	(11.20)	(11.28)	(11.14)	(11.12)	(11.57)	[0.761]	
Transfers/loans to others ($)	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)	[0.761]	

Notes: Table reports OLS on the six treatment-arm dummies with randomization-stratum fixed effects; standard errors clustered on user in parentheses. Estimated using data from the second follow-up survey. Continuous outcomes are winsorized at the 95th percentile; the two binary indicators are linear probability models. The second-to-last column reports the cluster-robust joint Wald F-test that all six arm coefficients are zero, with, in brackets, the Romano–Wolf stepdown p-value adjusting for testing the eight outcomes jointly (9,999 wild cluster bootstrap replications imposing the null, with draws shared across outcomes). Sample sizes differ across columns for two reasons unrelated to treatment. Income was the first question asked and is therefore recorded for respondents who did not complete the remainder of the instrument, giving it the largest N. Transfers/loans to others was added after the survey was fielded by email but before the push-notification reminder, so it is observed only for respondents who completed the survey after the revision. Neither the ordering nor the timing of the addition varies with assignment, so both subsamples remain balanced across arms.

Table A9. MPC out of credit-line increase

	DV: Average Monthly Purchases through September 3, 2026		
	Retail-line Purchases	Grocery-line Purchases	Total Purchases
Retail-line Limit (Post)	0.079	0.004	0.082
	(0.002)	(0.002)	(0.003)
Grocery-line Limit (Post)	-0.010	0.096	0.086
	(0.004)	(0.004)	(0.007)
Observations	209,996	209,996	209,996

Notes: Table reports instrumental variables estimates of the marginal propensity to consume (MPC) out of credit-limit increases. The dependent variable in each column is average monthly purchases in USD on the indicated credit line (or total) over the post-treatment period from October 3, 2025 to September 3, 2026. Post-treatment retail-line and grocery-line limits on October 6, 2025 are instrumented with six treatment-arm dummies. All specifications include randomization strata fixed effects. Standard errors are reported in parentheses and are clustered at the individual level. See Section 5 for details on the IV specification.

Table A10. MPB out of credit-line increase

	Retail-line Balance	Grocery-line Balance	Total Balance
Retail-line Limit (Post)	0.084	0.001	0.085
	(0.002)	(0.001)	(0.002)
Grocery-line Limit (Post)	-0.004	0.064	0.060
	(0.004)	(0.001)	(0.005)
Observations	209,996	209,996	209,996

Notes: Table reports instrumental variables estimates of the marginal propensity to borrow (MPB) out of credit-limit increases. The dependent variable in each column is average daily balances in USD on the indicated credit line (or total) over the post-treatment period from October 3, 2025 to September 3, 2026. Post-treatment retail-line and grocery-line limits on October 6, 2025 are instrumented with six treatment-arm dummies. All specifications include randomization strata fixed effects. Standard errors are reported in parentheses and are clustered at the individual level. See Section 5 for details on the IV specification.

Table A11. Treatment effect of credit-line increase

	DV: Average Monthly Purchases through September 3, 2026		
	Retail-line Purchases	Grocery-line Purchases	Total Purchases
Grocery line +50%	-1.83	7.89	6.07
	(0.83)	(0.93)	(1.49)
Grocery line +100%	-0.69	13.79	13.20
	(0.84)	(0.96)	(1.53)
Retail line +50%	16.38	1.05	17.48
	(0.91)	(0.88)	(1.51)
Retail line +100%	27.94	2.01	30.03
	(0.97)	(0.88)	(1.55)
Both +50%	15.72	9.58	25.41
	(0.91)	(0.93)	(1.56)
Both +100%	27.53	14.87	42.48
	(0.97)	(0.96)	(1.63)
Strat Bin FE	Yes	Yes	Yes
SE Clustered	Yes	Yes	Yes
Control Mean	124.65	123.25	247.97
Observations	209,996	209,996	209,996

Notes: Table reports treatment effect estimates of credit-limit increase assignment on purchases. The dependent variable in each column is average monthly purchases in USD on the indicated credit line (or total) over the post-treatment period from October 3, 2025 to September 3, 2026. Regressors are six mutually exclusive treatment-arm dummies indicating random assignment to a retail-line and/or grocery-line credit-limit increase of 50% or 100%, relative to the control group. All specifications include randomization strata fixed effects. Standard errors are reported in parentheses and are clustered at the individual level. See Section 5 for details on this specification.

Table A12. Dynamic MPC Estimates on Purchases

	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)
Observations	209,996	209,996	209,996	209,996

Notes: Table reports monthly OLS treatment effect estimates from Equations 6 and 7, estimated separately for each month from October 2025 through August 2026, each using a user's limit on October 6th as the instrument. Panel A instruments with the grocery-line (LC) limit increase; Panel B instruments with the retail-line (LP) limit increase. Columns report the estimated MPC on grocery-line and retail-line purchases, respectively. Point estimates are reported in the top row of each panel, with standard errors clustered at the user level in parentheses below. Regressions include randomization strata fixed effects. Users with no purchasing activity in a given month are assigned a value of zero. All continuous variables are winsorized at the 0.1st and 99.9th percentiles within each month.

Table A13. Marginal Propensities to Consume by Merchant Category

	DV: Average Monthly Purchases by Merchant Category	
	Grocery-line Limit (Post)	Retail-line Limit (Post)
Panel A: Grocery-line Categories		
Groceries	0.076	0.003
	(0.004)	(0.002)
Pharmacies	0.013	0.000
	(0.002)	(0.001)
Restaurants	0.004	0.001
	(0.000)	(0.000)
Observations	209,996	209,996
Panel B: Retail-line Categories		
Home	-0.005	0.037
	(0.002)	(0.001)
Clothing	-0.004	0.012
	(0.002)	(0.001)
Services	-0.001	0.005
	(0.001)	(0.000)
Technology	-0.001	0.014
	(0.001)	(0.001)
Tools	0.000	0.005
	(0.001)	(0.000)
Health	-0.000	0.002
	(0.000)	(0.000)
Sports	-0.000	0.002
	(0.000)	(0.000)
Observations	209,996	209,996

Notes: Table reports 2SLS estimates of Equations 6–7 with merchant-category purchase outcomes. The dependent variable in each regression is cumulative total purchases in the merchant category listed in the row. The endogenous regressors are post-treatment grocery-line and retail-line credit limits as of October 6, 2025, instrumented by the six treatment-arm assignment dummies. Each cell reports the estimated MPC—cents of total spending per dollar of credit limit—for a given category–endogenous regressor pair, with standard errors clustered at the user level in parentheses. Panel A reports grocery-line categories (Groceries, Pharmacies, Restaurants); Panel B reports retail-line categories (Home, Clothing, Services, Technology, Tools, Health, Sports). All specifications include randomization strata fixed effects. Outcome and limit variables are winsorized at the 0.1st and 99.9th percentiles. Purchases are measured as average monthly spending from October 3, 2025 to September 3, 2026.

Table A14. 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)	[0.290]
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: Table shows coefficients for each treatment group relative to the control group. All average variables in this table are calculated by taking the mean of the daily variable across the previous 30 days. Data as of October 1, 2025. The joint F-test p-value is derived from regressing the covariate x_k against dummies for each treatment arm (with control as the omitted category) and testing the null hypotheses $\beta_k^1 = \beta_k^2 = ... = \beta_k^6 = 0$, where superscript numbers index treatment arms and subscript k indexes covariates. The omnibus F-statistic is derived from restricting the sample to one treatment arm and the control group, then regressing a treatment dummy for that treatment arm against all K covariates and testing the null hypothesis $\beta_1^j = \beta_2^j = ... = \beta_K^j = 0$, where superscript j indexes treatment arms and subscript numbers index covariates.

Table A15. Heterogeneous Treatment Effects: Classification Analysis (CLAN) for Selected Variables

	Same-line effect of a limit increase	
	Grocery line	Retail line
Personal characteristics		
Age	3.58	-0.26
	[0.000]	[0.103]
	(3.26, 3.89)	(-0.58, 0.05)
Days since initial application	9.69	16.76
	[0.000]	[0.000]
	(4.65, 14.79)	(11.49, 21.97)
Cashea Club Level 6	0.21	0.29
	[0.000]	[0.000]
	(0.19, 0.22)	(0.28, 0.30)
Credit position		
Pre-experiment total limit	64.42	119.76
	[0.000]	[0.000]
	(57.54, 71.19)	(112.90, 126.60)
High utilization frequency (5%)	0.48	1.68
	[0.000]	[0.000]
	(0.28, 0.68)	(1.45, 1.90)
Total utilization: top tercile	0.22	0.40
	[0.000]	[0.000]
	(0.21, 0.23)	(0.38, 0.41)
Spending behavior		
Average monthly purchases	169.80	213.21
	[0.000]	[0.000]
	(164.13, 175.40)	(207.58, 218.77)
Average monthly transaction count	1.34	1.52
	[0.000]	[0.000]
	(1.27, 1.42)	(1.45, 1.59)
Spending volatility	72.96	137.06
	[0.000]	[0.000]
	(68.71, 77.13)	(133.08, 140.90)
Repayment history		
Maximum days past due	-5.59	-13.64
	[0.000]	[0.000]
	(-6.23, -4.93)	(-14.51, -12.82)
On-time ratio	0.00	0.03
	[0.912]	[0.000]
	(-0.00, 0.01)	(0.02, 0.03)
History of default	-0.06	-0.16
	[0.000]	[0.000]
	(-0.07, -0.06)	(-0.17, -0.15)
Observations	87,578	87,631

Notes: Table compares mean baseline characteristics between the most and least affected quintiles of the predicted conditional average treatment effect, following Chernozhukov, Demirer, Duflo, and Fernández-Val, 2025, using a random forest proxy. Each entry is the difference in means between the top and bottom quintile, oriented so that a positive value indicates the more responsive group has more of the characteristic. Only the two same-line specifications are reported, because these are the specifications in which the test of $\beta_2 = 0$ in Table A7 rejects homogeneity. Column 1 restricts the sample to the control group and the two grocery-line arms with grocery-line purchases as the outcome; column 2 to the control group and the two retail-line arms with retail-line purchases as the outcome. Estimates are medians over 100 data splits; brackets report p-values and parentheses 95% confidence intervals.

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, $m_t \equiv y_t + x_t - (1+r)\left(a_{G,t-1} + a_{R,t-1}\right)$ 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. Throughout, r is the interest rate on credit-line borrowing.

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

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

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. The no-mental-accounting benchmark is nested in this model with $\gamma_G = \gamma_R = 0$.

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 new category-ℓ targets $c^*_{\ell t}$ across all t and all states by solving a local single-category reoptimization problem, 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, 2025). 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.

Assumption 1 (Preferences)**.** *Preferences over state-contingent consumption streams 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 in-*

creasing, 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:*

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

with $p \in (0,1)$.

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,

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

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: $c_q(C_1) = s_q C_1$ for $q \in \{G, R\}$. Writing $\kappa \equiv H(s_G, s_R) > 0$, we obtain $v(C_1) = \Phi(\kappa C_1)$ and $v''(C_1) = \kappa^2 \Phi''(\kappa C_1) < 0$. Therefore, for each $q \in \{G, R\}$,

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

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

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

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.[18]

[18] 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

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

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

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 C_1^* denote the value of C_1 that solves equation (A.5). Then, defining D_1 as the derivative of the left-hand side of equation (A.5) with respect to C_1,

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

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

$$\frac{\partial C_1^*}{\partial L_\ell} = -\frac{\beta b'(C_1^*) W_{b,L_\ell}}{D_1}. \tag{A.6}$$

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

$$\operatorname{sgn}\left(\frac{\partial C_1^*}{\partial L_\ell}\right) = -\operatorname{sgn}(W_{b,L_\ell}).$$

Therefore, to show that $\partial C_1^* / \partial L_\ell > 0$, it is sufficient that $W_{b,L_\ell} < 0$.

$\lambda z + (1-\lambda)z'$ 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.

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, the composition of consumption is not constrained. It follows that

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

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

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

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

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

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

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

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

$$x_3 = 0.$$

Equations (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 (shown below), we suppress L_k and write $V(m_2; L_\ell)$. The continuation problem therefore reduces to

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

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

Let $c_k^*(m_2, L_\ell)$ solve (A.15), and define

$$m_3^* = y_3 - (1+r)(L_\ell + c_k^*), \qquad \Gamma \equiv \beta(1+r)^2 v''(m_3^*) < 0.$$

The first-order condition for c_k is

$$u_k(m_2 + L_\ell, c_k^*) - \beta(1+r) v'(m_3^*) = 0.$$

Implicit differentiation gives

$$\frac{\partial c_k^*}{\partial L_\ell} = -\frac{u_{\ell k} + \Gamma}{u_{kk} + \Gamma}. \tag{A.18}$$

The envelope theorem gives

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

Combining this with (A.18),

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

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 $V_{m_2} = u_\ell\big(m_2 + L_\ell,\, c_k^*(m_2, L_\ell)\big)$, where 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 $\partial c_k^* / \partial L_k = 0$. Differentiating V_{m_2} in L_k,

$$V_{m_2, L_k} = u_{\ell\ell} \cdot \frac{\partial (m_2 + L_\ell)}{\partial L_k} + u_{\ell k} \frac{\partial c_k^*}{\partial L_k} = 0.$$

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,

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

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, 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 exhaust the bad-state regimes permitted by Assumption 4. Returning to the state decomposition derived at the start of this step

$$W_{b,L_\ell} = (1-p)\, V^B_{m_2,L_\ell},$$

Cases 1 and 2 give $V^B_{m_2,L_\ell} < 0$ whenever line ℓ is strictly binding in the bad state (equations (A.7) and (A.20)), while Case 3 gives $V^B_{m_2,L_\ell} = 0$ 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:

$$\frac{\partial C_1^*}{\partial L_\ell} \begin{cases} > 0, & \text{if line } \ell \text{ binds in the bad state,} \\ = 0, & \text{if line } \ell \text{ is slack in the bad state.} \end{cases}$$

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

$$c^*_{q1} = s_q C_1^*, \qquad s_q > 0.$$

Therefore,

$$\frac{\partial c^*_{q1}}{\partial L_\ell} = s_q \frac{\partial C_1^*}{\partial L_\ell}, \qquad q \in \{G, R\},$$

and we arrive at the conclusion:

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

□

We now turn to Prediction 2, where we now assume consumers may exhibit mental accounting ($\gamma_k > 0$). 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

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

Thus 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. For period-2 cash-on-hand m_2, define the frozen-k continuation value as the following:

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

This is the program (A.2) with the category-k consumption path frozen: in period 2 the consumer buys $\bar{c}_{k2}$, in period 3 she buys $\bar{c}_{k3}$, and category ℓ absorbs the terminal budget. Define the ex-ante continuation value

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

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 is

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

where $b(\cdot)$ 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 policies inside $\widetilde{V}$, and the category-k targets stay at $\bar{c}_{kt}$.

Two properties of the target reset problem are important for the analysis below. First, at the original credit limits, the reset problem returns the pre-cue plan itself: freezing the category-k path at $(\bar{c}_{k1}, \bar{c}_{k2}, \bar{c}_{k3})$ and re-optimizing the category-ℓ path recovers $c^*_{\ell 1} = \bar{c}_{\ell 1}$ and the pre-cue continuation plan, since the pre-cue plan solves the unrestricted problem and the solution is unique within the maintained regime. The comparative statics in Step 2 are therefore evaluated at the pre-cue plan. Second, and consequently, the binding pattern inside $\widetilde{V}$ at the original limits is the pre-cue plan's own regime, which is fixed by Assumption 4. In particular, a marginal increase in L_ℓ does not switch the regime.

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

$$\frac{\partial c^*_{\ell 1}}{\partial L_\ell} = -\frac{\beta b'\, \widetilde{W}_{b,L_\ell}}{\widetilde{D}_1}, \qquad \text{so} \qquad \operatorname{sgn}\left(\frac{\partial c^*_{\ell 1}}{\partial L_\ell}\right) = -\operatorname{sgn}(\widetilde{W}_{b,L_\ell}). \tag{A.26}$$

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 one-for-one in both states, so

$$\widetilde{W}_{b,L_\ell} = p\, \widetilde{V}^{\,G}_{m,L_\ell} + (1-p)\, \widetilde{V}^{\,B}_{m,L_\ell},$$

where the superscripts denote the good and bad states, and the slackness of the good-state limits gives $\widetilde{V}^{G}_{m,L_\ell} = 0$. Two bad-state cases remain.

Case 1: Line ℓ binds. Suppose $a_\ell = L_\ell < c_\ell$, with a strictly positive multiplier $\eta_\ell > 0$. Attach 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

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

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 frozen problem,

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

Three consequences. First, $x_3 = 0$: the x_3 condition gives $\nu = \lambda - \beta u_{\ell 3} = u_{\ell 2} - \beta u_{\ell 3} > \beta r\, u_{\ell 3} > 0$, so $x_3 \geq 0$ binds (borrowing at gross rate $1 + r > 1$ and saving cannot coexist). Second, a_k sits at its upper corner: an interior a_k ($\eta_k = \zeta_k = \mu_k = 0$) would force $\lambda = \beta(1+r)\, u_{\ell 3}$, contradicting $\lambda = u_\ell > \beta(1+r)\, u_{\ell 3}$, and $a_k = 0$ would force $\zeta_k = \beta(1+r)\, u_{\ell 3} - \lambda < 0$; hence

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

with multiplier $\eta_k + \mu_k = \lambda - \beta(1+r)\, u_{\ell 3} = \eta_\ell > 0$ (every unit of category-k spending that can be financed on line k is, freeing period-2 cash for category-ℓ consumption at a period-3 cost of $1 + r$). Third, since every active multiplier is strictly positive ($\eta_\ell, \nu, \eta_k + \mu_k > 0$), the active set is locally constant in (m, L_ℓ), and the binding constraints pin down the solution: the budget constraint gives

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

in which L_ℓ enters only through the binding constraint $a_{\ell 2} = L_{\ell 2}$, 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 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,

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

Note the contrast with (A.19): the term $u_{\ell k}\, \partial c_k^*/\partial L_\ell$ is absent because the category-k path does not adjust. Neither the Hessian restriction (A.17) nor homotheticity is used. If line k also binds ($a_{k2} = L_{k2} \le \bar{c}_{k2}$), the same expressions apply with $a_{k2} = L_{k2}$; no separate case is needed.

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,

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

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

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

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 (**??**) at the new credit limit with the new mental targets. We show the optimal plan is the target plan itself. Three observations.

First, the target plan is feasible: the reset problem (A.23)–(A.25) enforces every budget constraint, and the frozen category-k path was feasible before the cue and remains feasible at the weakly larger limit.

Second, the objective is concave: u is concave, each penalty $-\gamma_q |c_{qt} - c_{qt}^*|$ is concave, and all constraints are linear. Therefore any plan satisfying the first-order (subgradient) conditions is a global maximum.

Third, the target plan satisfies these conditions when the penalties are large. Consider first the period-1 choice. Suppose for the moment that in periods 2–3 the consumer keeps every c_{kt} at its target; this is verified at the future nodes below, and that verification does not use the present paragraph. Under this supposition the periods-2–3 value of (**??**) coincides with $\widetilde{W}$ for every buffer near the target plan's (the future-node inequalities below are strict under Assumption 4 and vary

continuously with the buffer), so period-1 consumption solves

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

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 $u_q - \gamma_q g_q + \beta \widetilde{W}_b b' = 0$ for both categories, that is, if and only if

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

The first inequality holds with its left-hand side equal to *zero*: that is exactly the first-order condition of the reset problem (A.25). The cued category sits at its own target, and its penalty is inert for any $\gamma_\ell \ge 0$. Substituting this first-order condition, $\beta \widetilde{W}_b b' = -u_\ell$, into the second inequality gives

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

Not deviating 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 require known marginal utilities prior to reoptimizing, so the consumer is able to decide whether it is worthwhile to fully rebudget without needing to first solve the full reoptimization problem.

The future nodes have the same structure, with one refinement. Inside the reset problem, category-ℓ consumption at each node (period 2 and period 3, in each state) is chosen freely, so its first-order condition equates u_ℓ to the shadow price λ of that node's resources. The subgradient condition for holding c_{kt} at its target at that node is:

$$\big|\, u_k(c_{\ell t}, \bar{c}_{kt}) - \lambda + \mu_k \,\big| \le \gamma_k,$$

where $\mu_k \ge 0$ is that node's multiplier on the accounting constraint $a_k \le c_k$: exactly as in (A.8), μ_k enters the c_k first-order condition, and it does not drop out here because c_{kt} carries a penalty while a_{kt} does not. At period 3 ($a_{q3} = 0$), at the good-state period-2 node (limits slack), and at the bad-state node whenever $a_{kt} < c_{kt}$ — in particular when line k binds, $a_{k2} = L_k < \bar{c}_{k2}$ — we have $\mu_k = 0$ and the condition takes the same form as in period 1:

$$\big|\, u_k(c_{\ell t}, \bar{c}_{kt}) - u_\ell(c_{\ell t}, \bar{c}_{kt}) \,\big| \le \gamma_k.$$

At the bad-state node in the corner established in Case 1 of Step 3, however, all category-k spending is credit-financed on a slack line ($a_{k2} = \bar{c}_{k2} < L_k$). There a marginal unit of c_{k2} relaxes the accounting cap $a_{k2} \leq c_{k2}$ one-for-one, so it is financed by borrowing rather than by period-2 cash: the a_{k2} first-order condition gives $\mu_k = \lambda - \beta(1+r)\, u_\ell(M_3 - \bar{c}_{k3}, \bar{c}_{k3}) > 0$, and the condition becomes

$$\left|\, u_k(c_{\ell 2}, \bar{c}_{k2}) - \beta(1+r)\, u_\ell(M_3 - \bar{c}_{k3}, \bar{c}_{k3}) \,\right| \leq \gamma_k.$$

Borrowing and saving carry no penalties and are chosen optimally in the reset problem, so their Kuhn–Tucker conditions hold exactly.

Define $\bar{\gamma}$ as the largest of these gaps across the (at most five) nodes of the target plan,

$$\bar{\gamma}_k \equiv \max\left(\left|\, u_k(c_{\ell t}, \bar{c}_{kt}) - \lambda_t + \mu_{kt} \,\right|\right)$$

across all time periods and states, which is a finite number, where $\lambda_t - \mu_{kt}$ equals $u_\ell(c_{\ell t}, \bar{c}_{kt})$ at every node except the bad-state credit-financed corner, where it equals $\beta(1+r)\, u_\ell(M_3 - \bar{c}_{k3}, \bar{c}_{k3})$; the period-1 bundle is $(c^*_{\ell 1}, \bar{c}_{k1})$. If $\gamma_k > \bar{\gamma}_k$, every subgradient condition above holds: the future nodes stay at their targets, the periods-2–3 continuation of (**??**) therefore coincides with $\widetilde{W}$ as supposed, and the period-1 conditions hold as shown. By concavity the target plan is a global maximum of (**??**). The consumer's actual consumption equals the targets in every period and state:

$$c_{\ell 1} = c^*_{\ell 1}, \qquad c_{k1} = \bar{c}_{k1}.$$

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

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

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

C Testing for Heterogeneous Treatment Effects

To test for heterogeneous treatment effects in a structured manner, we implement the machine-learning methodology proposed by Chernozhukov, Demirer, Duflo, and Fernández-Val (2025). We focus on whether there is predictable heterogeneity in the treatment effect of our monthly purchases outcome. We measure purchases in the same way as equations 7 and 8.

Following the notation in Chernozhukov, Demirer, Duflo, and Fernández-Val (2025), we consider the conditional average treatment effect (CATE):

$$s_0(Z) := E[Y(1)|Z] - E[Y(0)|Z],$$

where $Y(1)$ and $Y(0)$ are the potential outcomes with and without treatment on the given credit line. The best linear predictor of $s_0(Z)$ given $S(Z)$ is the solution to

$$\min_{b_1,b_2} E[s_0(Z) - b_1 - b_2 S(Z)]^2,$$

which, if it exists, is

$$BLP[s_0(Z)|S(Z)] = \beta_1 + \beta_2(S(Z) - \mathbb{E}[S(Z)]),$$

where $\beta_1 = \mathbb{E}[s_0(Z)]$ and $\beta_2 = Cov[s_0(Z), S(Z)]/\text{Var}[S(Z)]$. If $S(Z)$ is a perfect proxy for $s_0(Z)$, then $\text{Cov}[s_0(Z), S(Z)] = \text{Var}[S(Z)]$ and thus $\beta_2 = 1$. In general, $\beta_2 \neq 1$ due to noisy predictions $S(Z)$. If $S(Z)$ is purely noise and uncorrelated with $s_0(Z)$, then $\beta_2 = 0$. Furthermore, if there is no heterogeneity, i.e., $s_0(Z) = s$, then $\beta_2 = 0$ because $\text{Cov}[s_0(Z), S(Z)] = \text{Cov}[s, S(Z)] = 0$. Rejecting the hypothesis that $\beta_2 = 0$ therefore indicates that both (i) $S(Z)$ is a relevant predictor and, more importantly for our purposes, (ii) three is heterogeneity in $s_0(Z)$. Therefore, by estimating the β_2 coefficient of $\text{BLP}[s_0(Z)|S(Z)]$ and testing whether the coefficient is significantly different from zero, we can empirically test for heterogeneity in the treatment effect across individuals. [19]

In the ML step, we consider supported vector machine, elastic net, random forest, and neural net, and choose the model that maximizes $\Lambda \equiv |\beta_2|^2 \text{Var}(S(Z)) \propto \text{Corr}^2(s_0(Z), S(Z))$ for the BLP analysis.[20] Maximizing Λ is equivalent to maximizing the correlation between between the ML proxy $S(Z)$ and the CATE $s_0(Z)$, or equivalent to maximizing the R^2 in the regression of $s_0(Z)$ on $S(Z)$. The model that maximizes Λ for each of our analyses of purchase behavior is random

[19] The coefficient estimates will depend on the data split that takes place before training the machine learning models. The coefficient estimates, p-values, and lower and upper bounds of their confidence intervals that we report in the tables in this appendix are the medians of the estimates and confidence intervals from 100 such distinct splits, as recommended by Chernozhukov, Demirer, Duflo, and Fernández-Val (2025).

[20] Since each of the 100 data splits leads to a different Λ, we maximize the median Λ from the 100 data splits.

forest.

Table A7 shows that we are able to reject the null hypothesis of no heterogeneity in treatment effects for same-line responses and unable to do so for either cross-line effect.

Given that we detected heterogeneous treatment effects in same-line purchase responses, we follow Chernozhukov, Demirer, Duflo, and Fernández-Val (2025) to determine which characteristics predict the strongest response in monthly purchases. To do this, we split the data into training and estimation samples, and after training the ML model on the training sample, we use the model to estimate predicted individual treatment effects for each individual in the estimation sample. We then group individuals int he estimation sample into quintiles based on their predicted individual treatment effects and estimate "group average treatment effects" (GATES) for each quintile. Then, we estimate differences in average characteristics Z across groups with the highest and lowest predicted treatment effects, the most-affected and least-affected groups, in a classification analysis (CLAN).

For both the GATES and CLAN analyses, we choose the model that maximizes $\bar{\Lambda} = \sum_{k=1}^{K} \gamma_k^2 \mathbb{P}(S \in I_k)$, where $\gamma_k \equiv \mathbb{E}[s_0(Z)|G_k]$ is sorted group treatment effect for group k, and $G_k \equiv \{S \in I_k\}$ is the group defined by the interval $I_k \equiv [l_{k-1}, l_k)$; in our case these intervals correspond to quintile bins of S. Maximizing $\bar{\Lambda}$ is equivalent to maximizing R^2 in a regression of $s_0(Z)$ on $\bar{S}(Z)$, where $\bar{S} = \sum_{k=1}^{K} \gamma_k \mathbb{1}(S(Z) \in I_k)$. The model that maximizes $\bar{\Lambda}$ is random forest for all four outcomes.

Figure 8 presents the result of the GATES analysis, estimating treatment effects for each of the five quintiles, by predicted treatment effects. Consistent with our findings in Table A7, we find that the confidence intervals between Q1 and Q5 do not overlap for either same-line comparison, indicating that there are significant heterogeneous treatment effects. In Table A15, we present a CLAN analysis to understand which individual traits are associated with the largest impacts on monthly purchases for the two same-line outcomes where we observe statistically significant heterogeneous treatment effects.