Mental Accounting for Credit: Experimental Evidence from Venezuela

Sean Higgins[1], Raymond Kluender[2]

Date: October 2, 2025

Introduction

This analysis plan pre-specifies our planned design and analysis of a randomized controlled trial measuring the impact of credit limit increases through a buy now, pay later (BNPL) app on consumer behavior.

Departments

1. Kellogg School of Management
2. Harvard Business School

Motivation

Recent evidence indicates that consumers increase borrowing when their credit limits rise, even when they are far from fully utilizing their existing credit lines. This behavior remains a puzzle in household finance: why do unconstrained consumers respond to limit increases? Existing experiments find evidence for various mechanisms, including precautionary motives (Aydin, 2022; Aydin and Kim, 2024) or interpreting the limit increase as a signal about future income (Yin, 2025). Another potential mechanism is mental accounting: the idea that consumers separate financial decisions based on how money is categorized (Thaler, 1995; Hastings and Shapiro, 2018), but this hypothesis has not been tested experimentally.

We partner with Cashea, a Venezuelan buy now, pay later (BNPL) provider with two distinct credit lines: one for food and drugstore purchases, and one for clothing, electronics, furniture, and other retail goods. We will conduct a randomized controlled trial (RCT) to study the effect of increasing one or both credit lines. By comparing borrowing responses across credit lines, we will provide an experimental test of mental accounting for credit. Using administrative contract-level data, we will estimate the marginal propensity to borrow (MPB) overall and on each line in response to the randomized credit line increases.

Partner

We are partnering with Cashea, a buy now, pay later (BNPL) provider in Venezuela. Cashea offers interest-free installment loans. With over six million users, Cashea is the largest source of consumer credit in Venezuela, facilitating an estimated 8% of total consumer spending.

Cashea users have two credit lines: a Línea Cotidiana that can be spent on food and drugstore purchases, and a Línea Principal that can be spent on clothing, furniture, electronics, and other retail products.

Research Question

Our primary research questions are: (1) What are the marginal propensities to borrow out of credit line increases on separate credit lines? (2) How do these effects vary based on initial credit limits and utilization, the size of the limit increase, and the purchase category associated with the limit increase?

Our experiment randomly assigns users to receive limit increases, allowing us to assess the causal effects on consumer behavior. We randomly vary the size of the limit increase and which credit line(s) it applies to. Randomizing across the two distinct credit lines provides an experimental test for mental accounting: for example, consider a consumer who is unconstrained on both credit lines (i.e., has substantial unused borrowing capacity on both lines) and randomly receives a limit increase on Line A. As long as the consumer is not already at their satiation point for the goods they can purchase with Line B and also spends some cash on the goods they can purchase with Line A, if mental accounting is not at play the borrower should either not increase borrowing on either line, or increase borrowing on both lines. If, on the other hand, they increase borrowing on only Line A, this behavior would be consistent with mental accounting.

Data

Cashea maintains an administrative database that records users' purchase and payment history. This database covers each user's sign-up date and credit limits; the dollar amount, date, category, and associated credit line of each purchase; and the dollar amount, due date, payment date, and late fees for each installment.

For a subset of our sample, we plan to run a baseline survey to elicit the amount of money in each category they spend using their Cashea line of credit and the amount they spend using other sources (e.g., cash). We'll also elicit some basic demographic information and knowledge of their credit limits.

Experimental Design

Sample selection: Our experiment will include 210,000 "active" Cashea users as of August 1, 2025. We define active users as users who (1) have made a purchase on both credit lines in the 40 days preceding the sample definition date and (2) made their first purchase using Cashea at least 62 days before the sample definition date.

Treatment: We will randomly assign users to receive a credit limit increase on the Línea Cotidiana, the Línea Principal, both lines, or neither line (control). Among treated users, we will randomize whether the limit increases by 50% or 100%. All limit increases will be permanent, though credit limits will otherwise evolve (from the new levels) as normal based on user behavior. Treated users will be notified of changes to their credit limit(s) via push notifications and emails. To isolate the effect of limit increases from the effect of the messaging, users in the control group will receive push notifications and emails reminding them of their credit limits at the same time as the treatment groups.

Upon notifying users of treatment, Cashea will automatically adjust the credit lines in the Cashea app. Treated users will receive a credit limit increase of 50% or 100%, and this increase will apply to the Línea

Principal only, the Línea Cotidiana only, or both lines. The sample will be evenly split across the treatment arms and control group.

- Treatment 1: 50% limit increase to the Línea Cotidiana only
- Treatment 2: 100% limit increase to the Línea Cotidiana only
- Treatment 3: 50% limit increase to the Línea Principal only
- Treatment 4: 100% limit increase to the Línea Principal only
- Treatment 5: 50% limit increase to both lines
- Treatment 6: 100% limit increase to both lines
- Control: No limit increases

Randomization: We will randomize at the user level and stratify based on three variables at baseline: (1) total credit limit across both lines, (2) average utilization ratio (balance divided by limit) over the last 30 days for the Línea Cotidiana, and (3) average utilization ratio over the last 30 days for the Línea Principal. For each variable, we will split users into terciles and form 27 strata from the interaction of these three binned variables. We will then randomize users within each stratum. Given the 27 strata and 7 treatment assignments (including control), we anticipate a number of "misfits." We will randomly assign these misfits within a stratum. Note that this is likely to result in some very small differences in our overall allocation of users to treatment arms (some may receive a few more than 30,000 while others may receive fewer).

Outcomes

Primary outcomes:
Purchases (total across both lines, and separately by line)
Balance (total across both lines, and separately by line)

Primary outcomes (details): Measured in USD. The outcome variables will be in either levels or changes depending on the specification (see below for specifics). We will also use number of purchases and a binary variable for whether any purchases are made and whether the balance is positive after our intervention.

Secondary outcomes: We will also test repayment behavior, measuring installments past due and in default (total across both lines, and separately by line).

Secondary outcomes (details): Measured both in USD and using a binary variable equal to 1 if the amount is positive.

Empirical Model

We will estimate two main specifications: one to estimate marginal propensities to borrow and one to estimate treatment effects across arms.

Our main specification will estimate the marginal propensity to borrow (MPB) out of a credit limit increase. To estimate the MPB, we will estimate

$$\Delta y_i^{\ell} = \alpha_p \Delta \mathrm{L}_i^p + \alpha_c \Delta \mathrm{L}_i^c + \alpha_{pc} \Delta \mathrm{L}_i^p \mathbf{1}[\Delta L_i^c > 0] + \alpha_{cp} \Delta \mathrm{L}_i^c \mathbf{1}\left[\Delta L_i^p > 0\right] + \delta_{s(i)} + \varepsilon_i$$

where Δy_i^ℓ denotes the change in purchases or balances in USD for line $\ell \in \{total, Cotidiana, Principal\}$ and ΔL_i^m denotes the change in the credit limit in USD on line $m \in \{Cotidiana, Principal\}$, and $\delta_{s(i)}$ are randomization strata fixed effects. $\mathbf{1}[\Delta L_i^c > 0]$ and $\mathbf{1}\left[\Delta L_i^p > 0\right]$ denote binary treatment indicators for increases in Linea Cotidiana and Principal, respectively. Meanwhile, α_{pc} and α_{cp} represent the additional MPB out of a Línea Principal credit limit increase when the Línea Cotidiana credit limit also increases and the additional MPB out of a Línea Cotidiana credit limit increase when the Línea Principal credit limit also increases, respectively. Thus, for individuals with both lines increased, the MPB out of Línea Principal will be $\alpha_p + \alpha_{pc}$, and the MPB out of a Línea Cotidiana will be $\alpha_c + \alpha_{cp}$. The coefficients α_p and α_c give the MPBs out of credit limit increases for those who had increases only on that line (not both).

Because credit limits can also change over time for reasons other than treatment, we will account for this by (i) using only the experimental change in credit limits as ΔL_i^m in an OLS regression using the specification above and/or (ii) using the observed change in credit limits for ΔL_i^m and using the experimental changes in credit limits (and their interactions with dummies for whether the other line also had an experimental increase) on each line as instruments in an instrumental variables (IV) regression. Which of these we use as our preferred specification will depend on how large the non-experimental changes in credit limits are during our experiment relative to the experimental changes; if the non-experimental changes are sufficiently small relative to the experimental changes, we will prefer option (i) and otherwise option (ii).

Our main specification allows for interaction effects between increases on the two lines: for example, if there is a decreasing MPB as total available credit increases, we would find $\alpha_{pc} < 0$ and $\alpha_{cp} < 0$. However, in the case where $\alpha_{pc} = \alpha_{cp} = 0$, the specification is over-specified which may hurt our statistical power to detect treatment effects through the α_p and α_c coefficients. Thus, we prespecify a decision rule that if we fail to reject the null hypothesis that $\alpha_{pc} = \alpha_{cp} = 0$ when estimating the above equation, we will use a specification without the interaction terms as the preferred specification for estimating the MPB:

$$\Delta y_i^\ell = \alpha_p \Delta \mathrm{L}_i^p + \alpha_c \Delta \mathrm{L}_i^c + \delta_{s(i)} + \varepsilon_i.$$

In this specification, the coefficients α_p and α_c give the MPBs out of credit limit increases on each respective credit line.

Our second main specification directly estimates treatment effects for each treatment arm; it is an OLS regression of the form

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

where y_i^ℓ is the outcome of interest for line $\ell \in \{total, Cotidiana, Principal\}$, T_{ik} are indicators for the treatment arms outlined above (control omitted), and $\delta_{s(i)}$ are randomization strata fixed effects. The coefficients of interest, α_k, measure the intent-to-treat (ITT) effect, which is the causal effect of treatment k on the outcome of interest relative to the control group. By including a separate indicator for each treatment arm, we do not impose parametric assumptions about the linearity or interaction of effects between the size of the limit increase and the associated credit line(s). We will also compare the α_k coefficients to each other to test whether the treatments have differential effects (e.g., whether a Línea Cotidiana increase has a different effect on spending on each line compared to a Línea Principal increase, or whether a 50% credit limit increase has a different effect than a 100% increase).

Heterogeneity

We will examine heterogeneity in the treatment effects across the following dimensions. To implement this, we will assign users to bins according to the specified sample splits across each variable, then interact these bins with each treatment indicator or change in credit limit (depending on the specification):

1. Baseline credit utilization ratio, total and on each line (in terciles)
2. Baseline credit limits, total and on each line (in terciles)
3. Share of spending through Cashea, total and on each line (baseline survey respondents only)
4. Self-reported credit constraints, comparing those who responded to the question “When you made purchases without using Cashea on [list of goods], why was that?” with “My Línea [Cotidiana/Principal] credit limit is too low” to those who responded any other response (baseline survey respondents only)

Sample Size and Power

Sample size: The overall experiment will include 210,000 users. These users will be evenly split across each treatment arm and the control arm, for a sample size of 30,000 users per arm.

Power calculations: We used baseline administrative data from Cashea to estimate sample means and standard deviations that we use to calculate minimum detectable effects (MDEs) for our primary outcomes for pairwise comparisons.

Table 1 shows the MDEs for our main outcomes. We can see that for total purchases our minimum detectable effect is $4.40 (2.0% of the baseline mean), for Línea Cotidiana purchases it is $2.60 (2.3%), and for Línea Principal purchases it is $3.10 (2.8%), all using a sample size of 30,000 per arm. Meanwhile, the MDE for total balance is $1.70 (2.1% of the baseline mean), for Línea Cotidiana it is $0.70 (2.3%), and for Línea Principal it is $1.40 (2.8%).

Table 1. Minimum detectable effects with 30,000 users per arm.

Variable	**Mean (USD)**	**Standard deviation (USD)**	**MDE (USD)**	**MDE as % of baseline mean**
Total purchases	224.3	193.3	4.4	2.0
Purchases on Línea Cotidiana	113.8	114.9	2.6	2.3
Purchases on Línea Principal	110.5	136.1	3.1	2.8
Total balance	79.6	74.6	1.7	2.1
Balance on Línea Cotidiana	29.3	29.8	0.7	2.3
Balance on Línea Principal	50.3	60.6	1.4	2.8

References

1. Aydin, Deniz. "Consumption response to credit expansions: Evidence from experimental assignment of 45,307 credit lines." *American Economic Review* 112, no. 1 (2022): 1-40.
2. Aydin, Deniz, and Olivia S. Kim. "Precautionary Debt Capacity." Working Paper (2024).
3. Hastings, Justine, and Jesse M. Shapiro. "How are SNAP benefits spent? Evidence from a retail panel." *American Economic Review* 108, no. 12 (2018): 3493-3540.
4. Thaler, Richard. "Mental accounting and consumer choice." *Marketing Science* 4, no. 3 (1985): 199-214.
5. Yin, Xiao. "Learning in the limit: Income inference from credit extension." Working Paper (2025).