NBER WORKING PAPER SERIES

DOES SAVING CAUSE BORROWING? IMPLICATIONS FOR THE CO-HOLDING PUZZLE

Paolina C. Medina
Michaela Pagel

Working Paper 28956
http://www.nber.org/papers/w28956

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2021, Revised September 2025

We are indebted to the Editor Stefan Nagel, an anonymous AE, and two anonymous referees. We thank seminar and conference participants at the University of Houston, AEA Meetings, Dartmouth HF Seminar, IPA Financial Inclusion Gathering, SITE, CFPB Research Conference, FIRS, MFA, CEPR HF seminar, Australian National University (ANU), Shanghai Advanced Institute of Finance, SJTU (SAIF), and Lone Star Conference, as well as Taha Choukhmane (discussant), James Choi (discussant), John Beshears, Doug Bernheim, David Laibson, Antonio Gargano, John Gathergood, Zack Liu, Anton Lines, Catherine Evans, Nhan Le, Erkki Vihriälä (discussant) and Han Xia (discussant) for providing valuable comments. We also thank Jose Antonio Murillo for invaluable support and access to the data, Isidoro Garcia for helping with preliminary analysis, and Sofia Terrazas for outstanding research assistance. Portions of this research were conducted with computing resources provided by the Research Computing Data Core at the University of Houston and Texas A&M High Performance Research Computing Cluster. The experiment was pre-registered under RCT number AEARCTR-0003047. IRB ID:IRB2018-0597. Authors do not have direct conflicts of interest. At the time of the experiment, Medina was a full-time employee of Banorte; however, the bank does not reserve the contractual right to review or withhold any of the research results. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Does Saving Cause Borrowing? Implications for the Co-Holding Puzzle
Paolina C. Medina and Michaela Pagel
NBER Working Paper No. 28956
June 2021, Revised September 2025
JEL No. D14, G5

ABSTRACT

Using an experiment in which 3.1 million bank customers were encouraged to save, we explore the mechanisms behind co-holding liquid savings and credit card debt. Theoretically, we first show that the joint responses of spending, saving, and borrowing to the nudge differ for different economic models of co-holding. Using machine learning techniques, we then find that the most responsive individuals reduce spending and increase their savings by 4.9% (206 USD PPP per month) while their credit card debt remains unchanged. For them, the marginal responses to the nudge are consistent with our model of co-holding for the purpose of self- or partner control.

Paolina C. Medina
University of Houston
pcmedina@uh.edu

Michaela Pagel
Washington University in St Louis
Olin Business School
and NBER
mpagel@wustl.edu

A significant number of US households hold both liquid savings and credit card debt (Gross and Souleles, 2002), which appears to be unnecessary and costly financial behavior. The leading explanations for this "co-holding puzzle" can be classified into two categories. The first category is based on assigning a liquidity premium to cash (Zinman, 2007), where cash is preferred due to the limited acceptance of credit cards (Telyukova, 2013) or restricted access to credit in times of need (Gorbachev and Luengo-Prado, 2019; Fulford, 2015; Druedahl and Jørgensen, 2018). The second category is based on self- or partner control, where individuals do not use their savings to repay credit card debt because they want to control their spouses' impulsive spending. If they were to pay off the debt, they would just accumulate debt again in the future and effectively spend their savings (Bertaut et al., 2009; Vihriälä, 2019).

Numerous policies are in place to promote savings, many of which use nudges (Benartzi et al., 2017). In this paper, we draw upon a unique, large-scale experiment designed to promote liquid savings and use it to examine the causes of co-holding. To do so, we not only focus on the immediate savings outcomes (Thaler, 1994; Beshears and Kosowsky, 2020) but also investigate the impact of saving nudges on high-interest debt, specifically credit card debt.

We analyze the joint responses of spending, savings, and credit card debt to saving nudges using data from a large-scale field experiment paired with comprehensive and accurate panel data of individual bank accounts and credit cards originating from a large bank in Mexico called Banorte. In our experiment, 2,696,936 customers were treated with (bi-)weekly SMS messages encouraging them to save for seven weeks, while a distinct group of 357,567 customers received no messages.

We first theoretically examine the effects of saving nudges on individual behavior in the presence of different underlying economic mechanisms for co-holding. We show that if the nudges make individuals more patient, then models based on transaction convenience or credit-limit chasing would predict a decrease in spending, repayment of credit card debt (as it is the most expensive item on the household balance sheet), and constant cash balances. Instead, if the nudges increased individual liquidity needs, these models would predict an increase in cash holdings financed by increases in debt to smooth consumption.

On the other hand, models of self- or partner-control make different predictions. If the nudges made individuals more patient and increased the cash holdings that are separated and hidden from the impatient party, then the models would predict a decrease in spending and an increase in savings but less of a change in credit card debt.

To empirically analyze the results of the experiment, we first note that the average treatment effect might conceal large individual differences in responses to the nudge. Some people might not respond at all, while others might respond strongly. The large size of our experiment allows us to study treatment effect heterogeneity across the margins that are relevant to distinguish the different models of co-holding: spending, saving, and credit card borrowing. We do so by training causal forests for these three outcomes following Athey et al. (2019). Using the omnibus test for the presence of heterogeneity proposed by Chernozhukov et al. (2018), we find evidence of treatment effect heterogeneity in spending and saving but not in credit card borrowing.

The causal forests allow us to predict treatment effects at the individual level and sort individuals into those who are highly responsive to the treatment and those who are not while avoiding overfitting. Overfitting would occur if our individual-level predictions for treatment effects were influenced by idiosyncratic factors, which we would spuriously attribute to the treatment. These idiosyncratic factors might also affect other outcomes, such as borrowing, and invalidate our inference.

Our primary analysis focuses on individuals in the top quartile of the distribution of predicted treatment effects on spending who have a credit card. This group decreased their spending and increased their checking account balances significantly, with a 7.2% decrease in spending and a 4.9% increase in savings, as measured by transactions (ATM withdrawals, card payments, and electronic transfers) and checking account balances over the treatment period. We then examine the effect on credit card borrowing. On average, this group decreased their credit card interest charges by 1.01% per month. This effect is small and precisely estimated with a standard error of 2.2%. To put this in perspective, we compare our estimate's upper and lower confidence intervals to the decrease in spending or increase in saving. For every 1 MXN in decreased spending or increased saving, we can rule out increases or decreases in borrowing cost of more than 0.3 or 0.5 cents MXN, respectively.[1] For comparison, the average credit card interest rate for this group is 38.5% per annum. If the credit card balance were to be 1 MXN larger or smaller for a billing cycle and the credit card bill is not repaid in full, the difference in interest incurred would be 3.2 cents in one month and compound to 46 cents over one year (given the auto-correlation of monthly credit card balances not fully repaid of 0.8).

We then turn our attention to the subset of individuals who rolled over credit card debt and paid interest on it in the six months before the intervention. Here, we also see increases in savings of similar magnitude, and once more, for every 1 MXN in savings, we can rule out increases or decreases in borrowing costs of more than 0.6 or 0.9 cents, respectively. Additionally, we document that these individuals did not use their new savings to pay off their existing credit card debt after the intervention. This implies that the nudge to save exacerbated the co-holding of low-interest savings and high-interest debt.

We argue that the observed marginal responses in spending, savings, and credit card debt to the nudge are more consistent with our model of self- or partner control than transaction convenience. In addition, we document three patterns that suggest self- or partner-control motivations to co-hold for these bank clients. First, one of the savings messages alluded to "locking away your savings." This message had a statistically significantly larger effect than all the other messages, while there is no difference between messages alluding to short-term versus long-term saving goals. Second, reductions in spending come largely from reductions in ATM withdrawals, which are more susceptible to sharing pressures from family members than funds held in bank accounts and used for payments via credit or debit cards. Third, if the borrowing response to the treatment were driven by financial transaction needs (instead of self-or-spousal control motives), our

[1] MXN stands for Mexican Pesos. As of the end of 2019, 1 MXN corresponded to 0.107 USD PPP, based on OECD conversion rates available at https://www.oecd-ilibrary.org/economics/data/aggregate-national-accounts/ppps-and-exchange-rates_data-00004-en.

rich set of financial covariates, measured at baseline, would likely capture heterogeneous effects on borrowing. However, as mentioned, we do not find evidence of treatment effect heterogeneity on that outcome.

Co-holding liquid savings and credit card debt is a common practice in the US and Mexico. In our sample, the average credit card interest rate is 35.8%, while checking accounts do not pay any interest. Despite the significant difference in rates, we observe that 26% of individuals who pay credit card interest maintain sufficient balances in their checking accounts to pay off their credit card balances in full. Moreover, 16% of individuals who pay credit card interest keep checking account balances that amount to at least 50% of their monthly after-tax income.[2] If they were to use these balances to repay their credit card debt, they could reduce their interest charges by an amount equivalent to 4% of their monthly after-tax income.

Our study makes three key contributions to the existing literature. First, we provide new evidence on the economic mechanisms behind the co-holding puzzle. Second, we analyze a large-scale randomized controlled trial and examine a variety of outcome variables, which demonstrate that nudges aimed at increasing savings do not lead to increased borrowing. Third, we carefully apply and discuss the latest machine-learning techniques for causal inference in a way that allows us to test the hypotheses of our theoretical frameworks.

We are interested in the interaction of saving and borrowing because many households co-hold credit card debt and perfectly liquid assets. Gross and Souleles (2002) first documented the phenomenon and noted that the transactions demand for liquidity may contribute to it. Maki (2002) studied whether households strategically run up credit card debt in preparation for a bankruptcy filing. However, most puzzle households are unlikely to file for bankruptcy (Telyukova, 2013). Zinman (2007) argues that credit cards and demand deposits are different assets and carry a liquidity premium. In Telyukova (2013), this premium is generated in a transaction convenience model, and in Fulford

[2]For this calculation we consider the minimum checking account balance observed on any given day over the six months preceding the experiment.

(2015), Gorbachev and Luengo-Prado (2019), and Druedahl and Jørgensen (2018), it arises due to variability in credit limits. In contrast, Bertaut et al. (2009) and Vihriälä (2019) argue that intra-household agency conflicts or self- or partner-control drive the co-holding.

Related to our study, Gathergood and Weber (2014) provide evidence of limited financial literacy among those who co-hold. More recently, Gathergood and Olafsson (2024) use high-frequency transaction-level data from Iceland to show that most co-holding is modest and relatively short-lived, but some individuals co-hold persistently. They then show that co-holding is related to the allocation of different categories of expenditure to credit and debit accounts. Batista et al. (2023) use bank account data and a large-scale field experiment to show that informing customers about their co-holding behavior and associated costs does not significantly alter their debt repayment behavior. They also propose a psychological explanation for co-holding, rooted in the principles of mental accounting.

Second, this paper contributes to a large literature on saving nudges, which documents positive treatment effects on savings of varying magnitude. Previous studies have found that interventions such as automatic enrollment in 401(k) savings plans, SMS messages, and Fintech apps can increase savings (Choi et al., 2004; Karlan et al., 2016; Gargano and Rossi, 2024; Akbas et al., 2016; Rodríguez and Saavedra, 2015). However, when people save more in response to nudges, the additional saving may be offset by changes in other positions of the household balance sheets or by future dissaving (Choukhmane, 2019). A handful of research papers examine the effects of pension savings on other positions of household balance sheets, such as borrowing. In Beshears et al. (2019) and Chetty et al. (2014), credit card borrowing is measured through snapshots (through biannual credit card balances from a credit bureau or annual measures of unsecured debt from tax administrates). These snapshots of balances do not provide accurate information on how much high-interest unsecured debt is rolled over. Credit card balances reflect spending in a given month and debt held. By comparison, we directly observe the average daily credit card balances in a billing cycle, how much of that bill is repaid, and how much interest is charged. Two more recent studies use bank account data from the UK to study the effects of increased pension savings. Choukhmane and Palmer (2023) look at increased minimum retirement contributions and find that consumers cut their spending, lower their deposit account balances, and increase credit card debt levels. Similarly, Beshears et al. (2024) find that increases in unsecured and mortgage debts partially offset pension automatic enrollment. Our study differs from the papers above by focusing on informational nudges rather than automatic enrollment interventions. This softer approach has become widely popular (Halpern, 2015), and exploring its unintended effects is important for both public and private stakeholders. We also add to the body of research that examines unintended consequences of nudges in various areas, such as financial accounts (Beshears et al., 2015; Goldin et al., 2017; Medina, 2021), health outcomes (Wisdom et al., 2010), and energy conservation (Costa and Kahn, 2013; Allcott and Kessler, 2019).

Third, our paper is one of the first to apply causal forests in the household finance literature, along with Burke et al. (2020). Causal forests have been successfully used in other fields, such as education (Carlana et al., 2022), labor (Davis and Heller, 2020), and development economics (Ashraf et al., 2020).[3] In our setting, a substantially larger sample size allows us to use these methods in two novel ways. First, we are powered enough to study treatment effects on sub-populations of interest identified by the causal forest, which allows us to test the predictions of competing theories. Second, we are able to compare causal forests and other methods for treatment effect heterogeneity based on experimental strata to illustrate the risk of overfitting bias.

In Section I we lay out two simple models to theoretically examine the effects of saving nudges on individual behavior in the presence of different underlying economic mechanisms for co-holding. In Section II we describe our empirical setting and data. Section III presents and discusses our empirical results and Section IV concludes.

[3] In the context of microfinance, causal forests and other machine learning methods for causal inference have been used by Beaman et al. (2023) to study selection into agricultural credit, Afzal et al. (2019) to study commitment features of microfinance loans, and by Breza et al. (2020) to study the impact of bank accounts and mobile money among previously unbanked factory workers.

I. Conceptual Framework

In this section, we will formally discuss two models that explain the co-holding puzzle. Within each framework, we will then look at the impact of the saving nudge on spending, saving, and borrowing decisions. The two models are meant to broadly cover the three most common explanations for the co-holding puzzle found in the literature: transactions-convenience (Telyukova, 2013), credit limit variability or chasing (Fulford, 2015; Gorbachev and Luengo-Prado, 2019; Druedahl and Jørgensen, 2018), and self- or partner-control (Bertaut et al., 2009; Vihriälä, 2019). Appendix IA.IV contains all the detailed derivations.

Our objective is to keep the theoretical exposition as transparent and simple as possible to most effectively convey the underlying intuitions driving our empirical tests of the models. Also, by following a straightforward approach, we aim to show the broad applicability of our findings.

A. Liquidity Premia Models: Transactions-Convenience and Credit-Limit-Chasing

The literature offers several reasons why having cash may be more valuable than an open line of credit (Zinman, 2007). For instance, certain transactions, like rent or mortgage payments, cannot be made with credit cards (Telyukova, 2013). Additionally, credit limits may be volatile, and an open line of credit may not be available when it is needed the most (Fulford, 2015). As a result, individuals may hold a certain amount of cash for insurance purposes. This desire to hold cash may come hand-in-hand with borrowing and individuals effectively paying an insurance premium in the form of credit card interest.

We use a simple two-period model to illustrate the intuitions behind transactions-convenience frameworks. The agent allocates their initial endowment, denoted by x_1, to consumption in period 1, denoted by c_1, and period 2, denoted by c_2. Their utility function is a log specification, and we thus restrict $c_1 > 0$ and $c_2 > 0$. The agent has to carry over a certain amount of cash, $x < x_1$, for transaction purposes. When the agent wants to consume more in period 1 than what their endowment and required cash holdings allow for, they can borrow, denoted as b_1, which costs them interest, denoted by r (while holding cash pays no interest).

In period 1, the agent's maximization problem is given by:

$$\max_{c_1}\{log(c_1) + \delta log(x_1 - c_1 - rb_1)\}$$

$$\text{where } b_1 := f_{b_1}(c_1) = \begin{cases} c_1 - x_1 + x & \text{if } c_1 - x_1 + x > 0 \\ 0 & \text{otherwise.} \end{cases}$$

When the agent is required to hold a certain amount of cash, $x > 0$, they may end up borrowing and therefore co-holding. The parameter $\delta \in (0, 1)$ denotes the agent's discount factor and thus reflects their patience. We assume the agent is not credit-constrained. The optimal solution for c_1, denoted by c_1^*, is discontinuous and determines their borrowing choice, which we denote by $b_1^* := f_{b_1}(c_1^*)$. Either the agent ends up borrowing or, if they do not borrow, they may or may not hold savings in addition to x. The conditions and consumption functions for each case are derived in Appendix IA.IVA. If the following condition holds

$$x_1 - \frac{1}{\delta + 1}x_1 + \frac{r}{(\delta + 1)(1 + r)}x < x \tag{1}$$

the agent borrows and thus co-holds. Their optimal consumption and borrowing are then given by

$$c_1^* = \frac{1}{\delta + 1}x_1 - \frac{r}{(\delta + 1)(1 + r)}x \text{ and } b_1^* = c_1^* - x_1 + x.$$

Now, suppose an agent co-holds and receives a nudge, such as a saving message, that increases their patience. The first theoretical takeaway can be summarized in the following proposition:

PROPOSITION 1: *If agents in the transactions-convenience model co-hold (Condition 1 is satisfied) and become more patient, then they decrease their consumption and repay their debt by the same amount, that is,* $\frac{\partial b_1^*}{\partial \delta} = \frac{\partial c_1^*}{\partial \delta} < 0$.

Proof. See comparative statics with respect to δ for the co-holding case. All details and derivations can be found in Appendix IA.IVA. □

To understand this proposition, we look at the comparative statics with respect to their patience, δ, for when the agent co-holds

$$\frac{\partial c_1^*}{\partial \delta} = -\frac{1}{(\delta+1)^2}(x_1 - \underbrace{\frac{r}{1+r}}_{<1} x) < 0 \text{ and } \frac{\partial b_1^*}{\partial \delta} = \frac{\partial (c_1^* - x_1 + x)}{\partial \delta} = \frac{\partial c_1^*}{\partial \delta} < 0.$$

The intuition behind this result is straightforward: if agents become more patient, they will consume less and save more. If agents co-hold and have credit card debt, they will want to reduce existing debt instead of increasing their cash holdings since credit card debt is more expensive. As a result, they will consume less and borrow equally less, that is, $\frac{\partial b_1^*}{\partial \delta} = \frac{\partial c_1^*}{\partial \delta} < 0$.[4]

Next, we consider what would happen if, instead, the nudge leads to an exogenous increase in cash holdings x. We can think of the nudge as making individuals more cautious, increasing their concern about the risk of losing access to credit or being unable to pay certain expenses with their credit cards. This leads us to the following proposition:

PROPOSITION 2: *If agents in the transactions-convenience model co-hold (Condition 1 is satisfied) and their cash needs increase, then they will increase their debt by the same amount minus a factor reflecting the increased costs of interest they pay, that is,* $\frac{\partial b_1^*}{\partial x} = 1 - \frac{r}{(\delta+1)(1+r)}$.

Proof. See comparative statics with respect to x for the co-holding case in Appendix IA.IVA. □

[4] An alternative modeling approach for a transactions-convenience model would involve the agent allocating consumption across goods that can be paid with cash and goods that can be paid with cash or credit. In this case, x, the cash holdings, would be endogenously determined as a function of how much the individual wants to consume of the good that can only be paid for with cash. If the individual becomes more patient in this scenario, their consumption of both goods in period 1 would decrease, and their consumption of both goods would increase in period 2.

To understand the proposition, we look at the comparative statics for consumption and debt with respect to the cash needs, x: $\frac{\partial c_1^*}{\partial x} = \frac{-r}{(\delta+1)(1+r)} < 0$ and $\frac{\partial b_1^*}{\partial x} = 1 + \frac{\partial c_1^*}{\partial x} > 0$

In response to increased cash needs, when the interest rate is zero, $r = 0$, the agent's consumption remains unchanged, but borrowing increases by the same amount as the required cash holdings. If the interest rate is positive, there is a small decrease in consumption due to the additional interest on the debt the agent chooses to hold. However, borrowing still increases by almost the same amount as the desired cash holdings.

To summarize, we have established that in liquidity-premia models where agents co-hold, an increase in patience leads to a decrease in borrowing. In contrast, an increase in the required amount of cash holdings leads to an increase in borrowing.

Credit-Limit-Chasing Models

A related class of models assumes that individuals co-hold cash and debt as a precautionary measure to ensure access to credit in the future (Druedahl and Jørgensen, 2018; Fulford, 2015; Gorbachev and Luengo-Prado, 2019). In other words, individuals hold debt as a way to keep their credit lines open. The comparative statics predictions of these models are the same as those in the transactions-convenience model just discussed. The models predict that if individuals become more patient, they will reduce their consumption and increase their debt repayment, while the amount of cash they hold will remain unchanged. On the other hand, if individuals' cash needs increase, they will increase their borrowing while keeping consumption remains unchanged.

In Appendix IA.IVB, we use a simple three-period model with one consumption good and log utility to demonstrate these comparative statics. Here, we follow the modeling approach of Druedahl and Jørgensen (2018) and assume that credit lines need to be used to stay open. An alternative interpretation of the model would be that individuals need to maintain existing outstanding loans because of the possible inability to secure new loans if the existing ones were repaid.

In an economy where certain goods, like mortgage payments, mandate payment through checking accounts or in which credit lines need to be used to remain open, other existing economic models find parallels. The transactions demand for cash model in (Tobin, 1958) aligns with the necessity for physical cash in daily transactions, while precautionary demand resonates with maintaining checking account buffers for essential payments. The Baumol-Tobin model optimizes the balance of transaction convenience and interest costs (Tobin, 1956), and inventory models draw a comparison to managing financial inventories (Baumol, 1952). Buffer stock models that mandate cash holdings for precautionary reasons generate a liquidity premium for cash as a transactions-convenience or credit-limit-chasing model would when there is uncertainty about the need for cash payments or credit limits. At the same time, technological constraints on payment methods and advancements in payment technologies (Ashraf et al., 2003; Bachas et al., 2021; Higgins, 2020) are likely to influence the demand for cash for transaction purposes or the desire to keep credit lines open.

B. Self- or Partner-Control Models

To illustrate the intuitions of self- or partner-control models (Bertaut et al., 2009; Vihriälä, 2019), we now outline a simple two-period model with two agent selves or spouses: an impatient party and a patient party. The impatient party allocates their endowment to consumption in period 1, denoted by c_1, and period 2, denoted by c_2. Their utility function is a log specification, and we thus restrict $c_1 > 0$ and $c_2 > 0$. When the agent wants to consume more in period 1 than their available resources, they can borrow, denoted as b_1, which costs them interest, denoted by r (while holding cash pays no interest). In period 0, the impatient agent's previous patient self or their patient spouse takes a certain amount of cash, denoted by $x \geq 0$, of the overall initial endowment, x_1, and deposits it in an inaccessible savings account that is only for period-2 consumption. A fraction $a \in (0, 1]$ of this amount is invisible to the impatient party or forgotten by them.

Therefore, the impatient party's perceived endowment in period 1 is $x_1 - ax$, and they perceive the inaccessible savings in period 1 as $(1 - a)x$. The perceived savings are determined by $1 - a$ because that is the fraction of inaccessible resources, x, that is visible to the impatient party. Now, if the impatient party wants to consume more than $x_1 - x$ in period 1, they must borrow because the entire amount x is inaccessible to them; that is, they borrow if $c_1 > x_1 - x$. Thus, the maximization problem of the impatient party in period 1 is

$$\underset{c_1}{max}\{log(c_1) + \beta log(x_1 - ax - c_1 - rb_1)\}$$

$$\text{where } b_1 := f_{b_1}(c_1) = \begin{cases} c_1 - x_1 + x & \text{if } c_1 - x_1 + x > 0 \\ 0 & \text{otherwise.} \end{cases}$$

$\beta \in (0, \delta)$ denotes the impatient party's discount factor, with $\delta \in (0, 1)$ denoting the patient party's discount factor.

The impatient agent's optimal solution for c_1, denoted by c_1^*, is discontinuous and determines their borrowing choice, which we denote by $b_1^* := f_{b_1}(c_1^*)$. Either they end up borrowing, or they end up not borrowing with or without additional savings. If they borrow, they also co-hold because they would only borrow when they know about a fraction of cash that is kept for future consumption. The conditions for each case are shown in Appendix IA.IVC. If the impatient agent ends up co-holding, their optimal consumption is

$$c_1^* = \frac{1}{\beta + 1} x_1 - \frac{r + a}{(\beta + 1)(1 + r)} x \text{ and } b_1^* = f_{b_1}(c_1^*) = f_{c_1^*}(x) - x_1 + x.$$

The impatient agent's optimal consumption and borrowing depend on the amount withheld, x, which is decided by a more patient party in some period 0. We thus define $c_1^* := f_{c_1^*}(x)$ and $b_1^* := f_{b_1^*}(x)$. In turn, the patient agent maximizes the following problem

$$\underset{x}{max}\{log(f_{c_1^*}(x)) + \delta log(x_1 - f_{c_1^*}(x) - r f_{b_1^*}(x))\}$$

to ensure that the patient agent does not want to hide more cash than what is available, $x < x_1$, we further assume that r is small and $\delta - \beta < (1 + \delta)a$, that is, the disagreement in discount factors is not too large relative to the patient agent's ability to hide cash when interest costs are small.

In a dual-self model, we can think of the impatient self as the future self (whose period-1 consumption is governed by β). In contrast, the patient self is a period-0 self trying to maximize overall welfare governed by the discounting parameter δ. This is in line with the framework for welfare analysis, where overall welfare is governed by some period-0 self who discounts utility in a non-hyperbolic manner, outlined in Laibson (1997). If we interpret the model in a dual-partner framework, we can think of c_1 and c_2 as joint consumption in a household in which the impatient and patient parties have different discount factors.

As optimal consumption by the impatient party, the optimal amount of hidden cash is discontinuous. The specific conditions for each case are derived in Appendix IA.IVC. If the following condition holds

$$x_1 - x_1 \frac{1}{1+\delta}(\delta \frac{1+r}{r+a} - \frac{\beta + r\beta}{a - r\beta}) - \frac{1}{\beta+1} x_1 + \frac{r+a}{(\beta+1)(1+r)} x_1 \frac{1}{1+\delta}(\delta \frac{1+r}{r+a} - \frac{\beta + r\beta}{a - r\beta}) < 0 \tag{2}$$

the impatient agent will co-hold and the optimal solution for x, denoted by x^*, is

$$x^* = x_1 \frac{1}{1+\delta}(\delta \frac{1+r}{r+a} - \frac{\beta + r\beta}{a - r\beta}).$$

When a is high, that is, cash can be hidden effectively, the patient party will hide more cash if the impatient party is more impatient, that is, β is low, and the resulting consumption in periods 1 and 2 is similar to what a rational agent in a univariate household would decide. When the impatient party knows about a fraction of the hidden cash and is impatient, that is, a and β are low, then they are more likely to borrow, while the patient party trades off the consumption-smoothing benefits of hiding cash with the interest costs resulting from the impatient party's debt.

Now, suppose that the difference between the patient and impatient parties' discount factors, that is, δ and β, increases, for example, because a nudge increases the patient self's patience, we can summarize the theoretical takeaway in the following proposition.

PROPOSITION 3: *If agents in the self- or partner-control model co-hold (Condition 2 is* *satisfied) and if the patient self becomes more patient or the impatient self becomes more impatient, then the patient self increases the hidden assets, that is,* $\frac{\partial x^*}{\partial \delta} > 0$ *and* $\frac{\partial x^*}{\partial \beta} < 0$.

Proof. Proof. See comparative statics of x^* with respect to δ and β when the agent co-holds. All details and derivations can be found in Appendix IA.IVC. □

To understand this proposition, we look at the comparative statics with respect to the discount factor of the patient party δ for the co-holding case

$$\frac{\partial x^*}{\partial \delta} = \frac{1}{(1+\delta)}(-x^* + x_1 \underbrace{\frac{1+r}{r+a}}_{>1}) \in (\underbrace{\frac{1}{1+\delta}(-x^* + x_1)}_{>0 \text{ if } x^* < x_1}, \frac{1}{1+\delta}(-x^* + x_1 \frac{1+r}{r})] \text{ for } a \in (0, 1].$$

The intuition is the following: if the patient party becomes more patient, then they withhold more money, especially if they can hide it.

A change in the patient self's patience, δ, changes the optimal hidden cash x^*. Therefore, we now look at comparative statistics of the impatient agent's spending with respect to the hidden cash. Our main theoretical takeaway can be summarized in the following proposition.

PROPOSITION 4: *If agents in the self- or partner-control model co-hold (Condition 2 is satisfied) and the patient self increases the hidden assets, x, then the impatient agent consumes less, especially when more of the assets can be hidden, that is,* $\frac{\partial c_1^*}{\partial x} < 0$ *and* $\frac{\partial c_1^*}{\partial x \partial a} < 0$.

Proof. See comparative statics with respect to x for the co-holding case in Appendix IA.IVC. □

To understand this proposition, we look at the comparative statics with respect to x for the co-holding case

$$\frac{\partial c_1^*}{\partial x} = -\frac{r+a}{(\beta+1)(1+r)} < 0, \; \frac{\partial c_1^*}{\partial x \partial a} < 0, \text{ and } \frac{\partial b_1^*}{\partial x} = \frac{\partial c_1^*}{\partial x} + 1.$$

When cash can be hidden, that is, the hiding-of-accounts friction a is not zero, the sensitivity of consumption to changes in x is negative, meaning that as x increases, consumption in period 1 decreases, that is, $\frac{\partial c_1^*}{\partial x} < 0$. Now, if cash can be hidden well, $a \to 1$, the sensitivity of consumption to x becomes more negative. The sensitivity of consumption to x is negative because the impatient agent considers the amount x as taken away, and they are not aware of it. The sensitivity of consumption is not equal to -1 because the agent distributes the overall loss in resources to consumption in periods 1 and 2. Additionally, even if they are fully aware of x (i.e., $a \to 0$), they would take the interest costs that borrowing incurs into account, so long as $\beta > 0$.

Technically, the impatient agent is unconstrained in their borrowing. What nevertheless constrains them is (1) the overall resources they think are available for period 1 and 2 consumption and (2) their log utility function, which would make no consumption infinitely painful. To illustrate the inner workings of the model for different levels of a (the invisibility of the hidden cash) and β (the impatient agent's discount factor), Figure 1 shows the amount of cash that is hidden by the patient party, x^*, and the amount of debt that the impatient party decides to borrow, b_1^*. The discount factors δ and β are calibrated using the estimates in Laibson et al. (2024), and the interest rate for borrowing is set to $r = 0.1059 * 2/12$, the first-stage estimate in Laibson et al. (2024), scaled to a 2-months period as in the experiment. x_1 is normalized to 1, so the values of x^* and b_1^* can be interpreted as fractions of the agents' initial endowment.

In this figure, we can see that if the impatient agent is becoming more patient (β is increasing), the patient party hides less cash, and the impatient party borrows less as their preferences are more aligned. We can also see that if the patient party cannot hide as much cash (a is low), then the impatient party borrows more (b_1^* is higher). The reason is that the impatient party knows about the hidden cash, which makes it less effective. For that reason, if a decreases, the patient agent has to hide more cash. That said, the patient agent also takes the interest costs into account, which makes them hide less cash when doing so is less effective, and the benefits of consumption smoothing are not large enough. For our calibration, however, the first effect dominates.

In Figure 2, we display the resulting consumption in periods 1 and 2. We can see that when β is low and a is increasing, the patient agent is better able to reduce consumption in period 1 and increase consumption in period 2, by hiding cash (x). Instead, when β and a are both low, the impatient agent forces a lot of consumption in period 1 and less consumption in period 2. The discontinuities in the consumption functions result from the discrete change in the interest rate when the agent starts borrowing and from the fact that hidden cash costs interest, on the one hand, but, on the other hand, reduces period-1 consumption. In the asymptotes, when the impatient agent is patient enough (β is high), they allocate the same amounts of consumption to both periods. However, when β is low, period-2 consumption is reduced relative to period-1 consumption, especially when the patient agent cannot hide cash effectively (a is low).

Table I compares the borrowing response to changes in the amount of hidden savings, x^*, in the self- or partner-control model to the borrowing response to same-magnitude changes in liquidity needs, x, in the transactions-convenience model. The amount of hidden savings, x^*, depend on the patient agent's patience δ. We first generate 100 random normal values for δ using the point estimate from Laibson et al. (2024) as the mean and the corresponding standard error as the standard deviation. For each value of δ, we then obtain the resulting hidden cash x^* in the self- or partner-control model. The value of x^* is then assumed to be the cash needed for transaction purposes in the transactions-convenience model. In turn, we obtain the resulting borrowing b_1^* in both models. Table I presents the average values of hidden/required cash and debt for both models. For the hiding-of-accounts parameter a we use four intermediate values ranging from 0.3 to 0.6. As before, x_1 is normalized to 1 and the interest rate for borrowing is set to $r = 0.1059 * 2/12$, the first-stage estimate in Laibson et al. (2024), scaled to a 2-months period as in the experiment. The impatient party's impatience is set to a value in the lower confidence range value of the estimate in Laibson et al. (2024).

Table I shows two key differences between the two models. First, the level of debt and thus co-holding is larger in the self- or partner-control model, ranging from 4% to 83% more. Second, there is less co-movement between borrowing and savings in the self- or partner-control model as measured by the slope estimate from regressing borrowing, b_1^*, on savings, x^*. In the transactions-convenience model, the co-movement is always close to

1, whereas in the self/partner-control model, it ranges from 0.55 to 0.77. To generate our results, we chose values for the hidden-cash friction factor that are relatively low. We thus do not need much invisibility or forgetting, which is reassuring. We have to acknowledge, though, that we chose a relatively low value for the impatient party's discount factor, β. That said, when we have lower values for a, we can choose higher values for β, for example, when we pick the point estimate for β from Laibson et al. (2024), we get the same results for values of a less than 0.5. Additionally, we have a very stylized model, so it is hard to draw quantitative conclusions about what would happen in a fully-fledged life-cycle model with liquid and illiquid assets such as the one in Laibson et al. (2024).

II. Experimental Design and Descriptive Statistics

In this analysis, we aim to investigate the impact of a preference shock that increases individuals' patience on their consumption, saving, and borrowing behavior to understand the economic mechanisms behind co-holding liquid savings and credit card debt. To do so, we utilize data from a large-scale field experiment conducted by the Mexican bank Banorte.

A. *The Mexican Credit Card Market: Basic Facts*

The bank Banorte is one of the 5 largest providers in the Mexican credit card market. As of December 2021, the top two largest banks control 56.5% of the market, and the top five largest banks control 87.0% of the market, with credit card debt representing 37.5% of the consumer credit portfolio (excluding mortgage debt).[5] The average number of credit cards per cardholder is 1.27 cards, according to a nationally representative survey conducted in 2018.[6] Interest rates on credit cards in Mexico are high compared to those in the US, with an average rate in December 2021 of 32.9% while the federal short-term interest rate was 5.71%.

[5]Refer to Banco de Mexico https://www.banxico.org.mx/publicaciones-y-prensa/rib-tarjetas-de-credito/%7BF4772BE9-C2BB-D24A-D104-A36851A01AF7%7D.pdf.

[6]Refer to INEGI, Encuesta Nacional de Inclusion Financiera, 2018.

B. Experiment: Sample and Experimental Treatments

We analyze data from a large-scale field experiment, which involved 3,054,503 customers. The customers were randomly selected from the universe of the bank's customers that satisfied three requirements. First, individuals needed to have a payroll account with Banorte.[7] Second, participants had to maintain an average daily balance of at least 50 MXN over the 2 months before the intervention. Third, individuals had a valid cell phone number to receive SMS messages.

From the experimental pool of 3,054,503 customers, a random sample of 357,567 clients was selected to be in the control group and received no messages. The remaining clients were assigned to the treatment group and randomly received one of seven messages designed to encourage savings. The treated customers were further divided into two groups, with half receiving the messages on a weekly basis and the other half receiving them bi-weekly (i.e., one message every other week). The intervention lasted 7 weeks, from September 13 to November 1, 2019.[8]

The experiment was stratified along the following dimensions: income quartiles, age quartiles, median of tenure with the bank, quartiles of baseline savings, dummy for clients for whom Banorte is the main bank, dummy for clients considered predominantly digital (30% or less of checking account charges made through cash withdrawals), median of the number of ATM transactions, dummy for clients with a credit card, and terciles of debit

[7] Payroll accounts are a type of deposit account in Mexico that are specifically designed for employees to receive their paychecks. Banks commonly offer them in partnership with companies to disburse salary payments. Unlike regular deposit accounts, payroll accounts often waive minimum balance requirements and offer access to credit products with special terms. Holders of a payroll account can also access all other products offered by the bank through standard application procedures without any restrictions. Since the bank has access to payroll information, our data on salaries is likely to be very accurate.

[8] The first set of messages was sent on Friday, September 13th, the same day they were scheduled. For all subsequent weeks, the messages were scheduled at the beginning of the week with the Customer Relationship Management (CRM) group to be sent out on Fridays whenever possible, but subject to other scheduling requests. The final set of messages was scheduled with the CRM group on October 27th and sent out on November 1st, 2019.

card transactions. The baseline refers to the 6 months previous to the intervention.

Table IA.I shows that the randomization process was successful in balancing observable characteristics across treatment and control groups.

The treatment messages were as follows:

Message 1: "Congratulations. Your average balance over the last 12 months has been great! Continue to increase your balance and strengthen your savings."

Message 2: "Increase the balance in your Banorte Account and get ready today for year-end expenses!"

Message 3: "Join customers your age who already save 10% or more of their income. Commit and increase the balance in your Banorte Account by $XXX this month."[9]

Message 4: "In Banorte, you have the safest money box! Increase your account balance by $XXX this payday and reach your goals."

Message 5: "Increase your balance this month by $XXX and reach your dreams. Commit to it. You can do it by saving only 10% of your income."

Message 6: "The holidays are coming. Commit to saving $XXX in your Banorte Account and avoid money shortfalls at year-end!"

Message 7: "Be prepared for an emergency! Commit to leaving 10% more in your account. Don't withdraw all your money on payday."

The messages used in the experiment can be classified into three categories. The first category includes messages that refer to short-term savings goals, specifically Messages 2, 6, and 7. The second category includes messages about savings in general, which are Messages 1, 3, and 5. Finally, there is one message, Message 4, which addresses self-control issues and suggests locking away money as a solution.

C. Descriptive Statistics and Prevalence of Co-Holding

Banorte routinely collects information on balances and transactions for all of their customers' accounts. They also perform bi-monthly credit checks on the credit bureau for customers who have a valid credit check authorization, which includes those with at

[9] XXX was a personalized amount representing 10% of the balance in the last 3 months.

least one credit product, such as a Banorte credit card.[10]

We thus have access to 161 pre-treatment variables for each individual, including financial behavior such as checking and credit card balances, as well as interest payments over the past six months. We also have demographic variables and geographic dummies. During the treatment period, we have access to all relevant variables, such as average daily balances for checking and credit card accounts, interest charges, deposits, outgoing transfers, credit card payments, balances as reported to credit bureaus, ATM withdrawals, and total card spending (transactions of debit and credit cards).

Table II presents descriptive statistics for individuals in both the treatment and control groups. The average age of these individuals is 45 years, with an average monthly after-tax income of 13,499 MXN (1,441 USD) and an average of 7 years of banking history with Banorte. Their average checking account balance is 20,006 MXN, and 12% of them hold at least one Banorte credit card.[11]

We also present these descriptive statistics separately for individuals who have a Banorte credit card. These individuals have higher average monthly after-tax income and higher average checking account balances than the sample of all clients. Their average credit card balance is 18,282 MXN (1,951 USD), with a median of 10,458 MXN (1,116 USD). The average monthly interest paid by individuals with credit cards is 270 MXN (29 USD). Note that this average includes individuals who pay zero interest. Furthermore, individuals with credit cards have substantial borrowing capacity, with an average limit of 84,244 MXN (8,991 USD) and a median of 45,000 MXN. All continuous variables are winsorized at the 1st and 99th percentiles. Interest charges, credit card balances, and credit limits are winsorized based only on individuals who have a credit card.

Table III provides information on checking account balances, credit card balances, and interest payments for individuals who have a credit card and who incurred credit card

[10]The periodic credit checks are used to offer personalized credit options and do not affect credit scores (analogous to "soft credit pulls" in the US).

[11]We also observe savings account balances, but savings accounts are rarely used, with less than 1% of users in our sample having a savings account, and the average balance on them being 57 MXN (conditional on being positive).

interest during the month previous to the intervention. Individuals are split by deciles of checking account balances over income, where checking account balances correspond to the minimum checking account balance observed for each individual over the 6 months previous to the intervention, and income refers to monthly after-tax income. On average, interest charges represent 4% of monthly after-tax income. Nevertheless, across all groups, 26% of individuals have sufficient savings to pay off their credit card balances in full, and 16% maintain a checking account balance of at least 50% of their monthly after-tax income.

D. Empirical Strategy

D.1. Outcome Variables

Theoretically, individuals might respond to the nudge in their spending, saving, and borrowing decisions. These are the main variables of interest for the analysis, and we now describe precisely how these variables are constructed.

To measure spending, we add up ATM withdrawals, credit/debit card transactions, and outgoing transfers. We include all outgoing transfers (irregular as well as recurring ones) to capture all potential responses to the treatment.

To measure savings, we use checking account balances. To avoid looking at balances at one arbitrary date, we calculate average daily balances on user's accounts. Average daily balances capture how much individuals are holding in their accounts at different points in time, weighted by the duration of the deposits.[12]

To measure the cost of credit card borrowing, we use credit card interest charges.

[12] Due to its volume, we did not receive daily balances from the partner bank. Instead, several aggregate measures were calculated using SQL queries on the dataset stored on the bank's servers, including average daily balances over each week of the intervention and the minimum balance observed over the 6 months previous to the intervention on each user's account. In addition, we also received monthly observations for the average daily balances kept by each user on each of the 6 months previous to the intervention. The bank routinely collects average daily balances at the monthly level as part of its normal operations.

Credit card interest is charged over the average daily balances on a credit card account observed over a given billing cycle. At the end of a billing cycle, individuals receive a credit card bill showing a due date (20 days after the last day of the billing cycle), all transactions incurred during the billing period, and the balance on the last day of the billing cycle (this balance is known as the statement balance). If individuals pay the statement balance in full by the due date, no interest is charged. However, suppose they do not repay the statement balance in full by the due date. In that case, they are charged interest on the entire average daily balances observed during the corresponding billing cycle.[13]

In contrast to spending and saving balances, which we can measure with arbitrary granularity (daily, weekly, monthly, etc.), credit card interest is only defined on a monthly basis since individuals are charged once per month on the basis of the average daily balances observed over a given billing cycle and depending on whether they are repaid in full or not.

Because billing cycles start at different days of the month for different individuals, our 7-week treatment period partially affected two billing cycles, with starting dates in September and October 2019. For our analysis, we take the average of the interest charges for the two billing cycles affected by the intervention. We also look at interest charges corresponding to average daily balances over the two billing cycles subsequent to the treatment, that is, the cycles that start in November and December, to account for potential carryover effects of the treatment on credit card interest charges.

We note that our measure of borrowing does not overlap exactly with the treatment window for all individuals, which, in principle, could bias our estimate for the treatment effect on borrowing. To alleviate this concern, we provide a robustness check for individuals whose billing cycle corresponds closely to the treatment window. Here, using

[13]As in the U.S., even if the individual is short of the full balance by a very small amount, interest charges are not calculated on the basis of this shortage. Instead, interest charges are calculated on the basis of the average of the balances observed by the end of each day for each day of the billing cycle. Average daily balances are a measure of how much individuals are borrowing and for how many days they do so.

individual-level billing cycle dates, we identify individuals with a cycle fully covered by the intervention, and for them, we look at saving, spending and borrowing over the weeks that overlap with the treatment(Section E). In addition, Figure IA.2 shows the distribution of billing cycle end dates. Dates are pretty evenly distributed over the month, which is reassuring.

To have a corresponding monthly measure of spending, our main results are based on treatment effects over the first month, that is, the first four weeks of the intervention. In Appendix IA.IIID, we show that the week-by-week treatment effects over the seven weeks of the treatment are not statistically significantly different from each other. Thus, the first four weeks of the intervention are akin to a monthly normalization without loss of generality.

We are also interested in measuring the amount of debt that individuals roll over. We define rolled-over debt as the credit card balance on the last day of the billing cycle (i.e., the statement balance) minus the payments received between the end of the billing cycle and the corresponding due date. This measure tells us how much of the consumption of the previous billing cycle will still need to be borrowed over the following billing cycle. We focus on the last billing cycle that intersects with the treatment period, which is the October billing cycle. This measure provides us with a snapshot of individuals' rolled-over debts towards the end of the treatment period.[14]

We also calculate the actual interest-bearing credit card balances for each individual by multiplying average daily balances by a binary variable flagging when an individual was charged credit card interest on the corresponding billing cycle.[15] To address the concern that individuals may substitute to non-Banorte credit cards, we also use end-of-the-month balances for all other credit cards that individuals have, using data from the credit bureau. Finally, we examine payments received on a credit card account during the

[14]Note that this measure of rolled-over debt is not the basis for calculating credit card interest because credit card interest is calculated on the basis of average daily balances.

[15]This would be equivalent to dividing monthly interest charges by the corresponding monthly interest rate.

two billing cycles that intersected with the treatment period (September and October) and the subsequent two billing cycles.

D.2. Aggregate Effects of the Intervention

To examine how the intervention affects spending, saving, and borrowing for the entire group of study participants, we estimate Specification (3) on all continuous non-negative variables:

$$Y_i = exp\{\alpha_s + \beta * treatment_i + \epsilon_i\} \tag{3}$$

Here, α_s refers to fixed effects for the randomization blocks, and β represents the treatment effect of the intervention, which is calculated as the difference in outcomes between the treatment group and the control group. We follow Chen and Roth (2024) and Cohn et al. (2022) in using an exponential (Poisson) specification to consistently estimate proportional treatment effects. Proportional treatment effects are defined as $\frac{E[Y|T=1]-E[Y|T=0]}{E[Y|T=0]}$. Therefore, using the coefficients in Equation 3, we calculate proportional treatment effects as $exp\{\beta\} - 1$. For binary variables, we use a linear probability model.[16]

D.3. Heterogeneous Effects

We also study treatment effect heterogeneity in spending, saving, and borrowing. We do so for several reasons. First, as with most communication-based nudges, not all individuals might have seen the treatment message on their cell phones, or they may not have paid attention to it or deemed it relevant. Mechanically, we would expect null treatment effects across all outcomes for them. Identifying individuals affected by the nudge in their spending, saving, and borrowing is important for distinguishing between competing models and evaluating policies for the following reasons.

[16] In principle, credit card balances can be negative if borrowers pay more than the outstanding balance at the end of each month, but this occurs in less than 1% of observations in our data. We replaced these negative values with zero when we winsorized, as described in Subsection B.

In terms of distinguishing between competing models, one of the premises in Propositions 1-4 of our theoretical framework is that individuals experience a shock to their preferences or liquidity needs. Regardless of the model, it is unlikely that the nudge was an effective shock for individuals who do not change their behavior at all as a response to the treatment. Further, changes in spending could be financed with debt or with savings, which would point to the relevance of self- or partner-control models versus other types of models. Thus, we focus on individuals for whom we expect the largest treatment effects on spending and for whom the nudge was likely to be an effective shock to their preferences or liquidity needs. From an econometric perspective, to avoid selecting on an outcome variable (e.g., spending), we use causal forest models to predict individual-level treatment effects using baseline observable characteristics (Athey et al., 2019). Sorting groups based on their observed treatment effects is prone to overfitting since the largest observed treatment effect may correspond to groups that experience idiosyncratic shocks, which lead to spuriously large treatment effects. Causal forests are designed to identify individuals with the largest response to the treatment without risking invalid inference due to overfitting (see Appendix IA.IIA for a detailed description). We then focus on individuals with the largest predicted treatment effect on spending, according to the causal forest.

In terms of policy evaluation, the potential of financing new savings with debt is more relevant for those individuals who actually increase their savings as a result of the treatment. Thus, we also study the borrowing behavior of individuals with the largest predicted treatment effects on savings.

In addition, whether the hypothesized null effect on borrowing applies uniformly to all individuals or if it represents the average of positive, null, and negative effects would also point to the relevance of different models and would lead to different policy conclusions.

III. Results

A. Aggregate Effects on Spending, Saving, and Borrowing

Table IV presents the aggregate treatment effects on spending, saving, and credit card borrowing. Here, all message treatments are combined into a single dummy variable that equals 1 if an individual received any of the treatment messages. Panel A presents results for the entire sample. In Column (1), the average effect of the treatment on spending is shown, indicating a significant 0.8% reduction from a base of 17,870 MXN. Column (2) shows a significant 0.6% increase in savings from a base of 19,913 MXN.

We now move on to individuals with credit cards (Panel B). Columns (1) and (2) show treatment effects on spending and saving of -1.9% and 1.1% from bases of 31,998 MXN and 35,657 MXN, respectively. Column (3) of the table shows the effect of the treatment on the interest paid on credit card debt during the two billing cycles covered by the intervention. The results are very precisely estimated and show a change that is statistically indistinguishable from zero. More specifically, we can rule out any increases greater than 1.4% and any decreases lower than 1.8% from a starting point of 216 MXN. This means that for every 1 MXN reduction in spending, the corresponding change in credit card borrowing costs is not more than 0.0049 MXN (0.014*216/(31998*0.019)) and not less than -0.0064 MXN. Panel C shows a similar pattern for individuals carrying credit card debt at baseline.

If the 1 MXN reduction in spending had been used entirely to pay off or increase credit card debt, the change in credit card interest payments would have been 0.030 MXN. Therefore, we conclude that less than 21% of the reduction in spending is reflected in reduced credit card debt, and no more than 16% is reflected in new debt. In the aggregate, we thus find that the majority of the reduction in spending is reflected in increased savings.

The relatively low average treatment effect on spending and saving is not surprising, as the experimental pool included individuals with minimal constraints, some of whom may not be responsive to the intervention. The inclusion of such individuals is necessary to train a model to predict treatment effect heterogeneity successfully. The sufficiently diverse experimental pool allows us to overcome the implicit selection of experimenting only with individuals for which the treatment is expected to work (Athey et al., 2021), which often leads to unsuccessful applications of the causal forests that fail to detect treatment effect heterogeneity. In order for the algorithm to learn who responds to the nudge and who does not, there must be enough individuals who do not respond to the treatment.

B. Heterogeneity Analysis

We now use causal forests to analyze how different individuals respond to the treatment in terms of spending, saving, and borrowing. For each outcome, we first train a pilot forest with 2,000 trees using all 161 pre-treatment variables. We then train a second forest using only variables with variable importance larger than 1%, following the approach of Athey and Wager (2019) for feature selection.[17]

To test for treatment effect heterogeneity, we follow Burke et al. (2020) and perform the calibration test of Chernozhukov et al. (2018). We conduct the tests on all 3.1 million observations, as well as separately on observations from individuals who have credit cards and those who paid credit card interest at baseline. Across all sub-groups, we find evidence consistent with the presence of treatment effect heterogeneity in the spending and saving responses to the saving nudge. In contrast, various models trained on borrowing outcomes fail to capture any significant heterogeneity in treatment effects, which we interpret as evidence suggesting that there is indeed no heterogeneity in the borrowing response. A detailed description of the test and corresponding results and interpretation can be found in Appendix IA.IIB.

For outcomes for which the algorithm detects treatment effect heterogeneity (i.e., spending and saving), we show the distribution of predicted treatment effects in Figure 3 (Panels A and B). We then split individuals into quartiles of predicted treatment effects.

[17] Variable importance indicates how often a variable was used to select splits across the multiple trees of the causal forest.

For each quartile, we calculate actual treatment effects, using a cross-fitted ranking over five folds.[18] For exposition, the top quartile of predicted treatment effects on spending contains the most negative predictions. In contrast, the top quartile of predicted treatment effects on savings contains the most positive predictions. Panels C and D of Figure 3 display the treatment effects on spending and saving for each quartile of predicted treatment effects. The figure indicates that the actual treatment effects are larger for individuals with larger predicted treatment effects, thus confirming the validity of predicted treatment effects as a sorting score for actual treatment effects.

Notably, although some observations have a positive (negative) predicted treatment effect on spending (saving), none of the quartile splits show a positive (negative) actual treatment effect on spending (saving). Essentially, the forests identify two distinct groups of individuals: a large first group with a treatment effect of zero (quartiles 1 to 3 of predicted treatment effects) and a smaller second group with a strong and statistically significant treatment effect (the top quartile of predicted treatment effects). The predictions for the first group exhibit a high degree of noise, as the predicted treatment effects span a wide range of negative and positive values that all result in an actual treatment effect of zero. In contrast, individuals in the top quartile of the predicted treatment effect distribution have a statistically significant and economically meaningful treatment effect.

C. *Treatment Effects on Spending, Saving, and Borrowing for the Top Quartile of Predicted Treatment Effects*

We now shift our focus to investigate the magnitude of the treatment effects among the subset of individuals in the top quartiles of predicted treatment effects on spending and saving. Individuals in these groups include a non-negligible segment of the population of credit card holders, as well as of those who paid credit card interest at baseline or those

[18] We use a cross-fitted ranking of predicted treatment effects where we split the sample into five folds and train a causal forest for each fold to predict treatment effects on the remaining folds (Chernozhukov et al., 2018; Abadie et al., 2018). We then separately rank the cross-fitted predictions for each fold and split them into quartiles. This approach ensures that the values of the outcome variables observed in each fold are not used when assigning observations in that same fold to a specific quartile.

who simultaneously hold credit card debt and high checking account balances. Panel A of Figure 4 shows that 41% of individuals who have a credit card, 49% of individuals who paid credit card interest at baseline, and 47% of individuals with savings higher than 50% of their income and who are carrying credit card interest belong to the top quartile of predicted treatment effects on spending. Panel B shows a similar pattern for the top quartile of predicted treatment effects on saving.

We first look at credit card holders in the top quartile of predicted treatment effects on spending. Individuals in this group exhibit a significant change in spending, providing an opportunity to test the hypotheses formulated in the models developed in Section I.

If individuals become more patient or cautious about their future cash needs in response to receiving saving messages, both the liquidity-premia and self- or partner-control models predict a decrease in spending. However, these two models differ in their predictions regarding borrowing decisions. In liquidity-premia models, borrowing either decreases or increases, depending on whether the messages represent a shock to patience or cash needs. In contrast, in self- or partner-control models, borrowing responds much less.

Table V displays the treatment effects on spending, saving, and borrowing for individuals in the top quartile of the distribution of predicted treatment effects on spending who have a credit card. Panel A includes all individuals who own a credit card, while Panel B focuses on individuals who pay credit card interest at baseline.[19]

We will first discuss the results in Panel A. Column (1) shows the spending outcomes for credit card holders in the top quartile of the predicted treatment effect distribution. The estimated reduction in spending is 7.22% on a baseline of 34,969 MXN, which corresponds to a decrease of 2,524 MXN. Column (2) indicates that almost all of the decrease in spending is attributable to an increase in checking account balances by 4.93% on a basis of 39,174 MXN, equivalent to an increase of 1,932 MXN.

[19]Note that the treatment effects on spending reported in Panels A and B of Table V differ from the treatment effect reported for the top quartile in Panel C of Figure 3, as the latter considers all individuals in the top quartile of predicted treatment effects, including those with and without credit cards.

In Column (3), we analyze interest payments for the two billing cycles affected by the intervention, and we observe a decrease of 1.01% from a basis of 210.16 MXN, with a standard error of 2.21%. By taking the lower confidence bound of the estimate and multiplying it by the baseline, we can conclude that a decrease (increase) in borrowing costs of more than 11.23 (6.98) MXN can be ruled out. Dividing this value by the decrease in spending, we can state with 95% statistical confidence that at most 0.0044 (0.0028) MXN for each 1 MXN decrease in spending was reflected in reduced (increased) interest costs.

If, on the other hand, each additional 1 MXN reduction in spending results in a decreased credit card balance that incurs interest, then the interest payments would decrease by 1*0.385/12 = 0.0321.[20] Thus, we can conclude with statistical confidence that less than 14% of the reduction in spending is reflected in reduced credit card interest payments. Furthermore, the results of the calibration test in Appendix IA.IIB suggest the absence of treatment effect heterogeneity in borrowing, implying that the null effect on borrowing applies uniformly to all potential subgroups of individuals in the top quartile of predicted treatment effects.

In Column (4), we observe that the same holds for interest payments for the two billing cycles following the two covered by the treatment. Furthermore, Columns (5) and (6) show that the indicator variable for whether an individual pays interest for either of the two billing cycles covered by the treatment or the two cycles after the treatment is estimated to be very close to zero with high precision. Finally, Column (7) presents the effects on the statement balance at the end of the billing cycle fully covered by the treatment, minus the repayments made towards that bill. This provides a snapshot of the rolled-over debt an individual has towards the end of the treatment. While the estimates are less tightly estimated, the treatment does not have a significant positive or negative

[20] To determine the effective interest rate of this group, we divide the interest payments by the interest-bearing balances. For individuals in the top quartile of the treatment effect distribution, the interest rate is (210.16/6557.54)*12 = 0.385 (see Tables V and IA.VI). As mentioned earlier, when considering the entire experimental pool, the interest rate is 35.8%.

impact on this measure of rolled-over debt.

Moving on to Panel B, the same regressions are run for the subset of individuals who have a credit card and paid interest at baseline. The findings are very similar to those in Panel A. We can rule out that interest costs decreased (increased) by more than 0.0086 (0.0055) MXN.

Table IA.VI presents the same specifications as before but with alternative outcome variables. In Columns (1) and (2), we use credit card balances that accrue interest (are not fully repaid) during and after treatment. The findings are very similar to the ones on paid credit card interest, indicating that less than 15% (27%) of the reduction in spending is reflected in interest-bearing balances (for individuals who paid credit card interest at baseline). Columns (3) and (4) use credit card balances independent of whether they are fully repaid, which makes it difficult to learn much from the results since the effect is confounded by individuals who have credit card balances but repay them in full. Nevertheless, we find similar results. Columns (5) and (6) use credit card balances from the credit card bureau, which includes non-Banorte credit cards. While this measure does not necessarily reflect interest payments either, it is reassuring that the effects are centered around zero, suggesting that individuals do not substitute other cards. This is consistent with a national survey that reports that 79% of individuals who have at least one credit card have only one.[21] Additionally, a robustness check that restricts the sample to individuals without other credit lines confirms this finding (as we discuss in Subsection E). Finally, in Columns (7) and (8), we use credit card repayments as the outcome variable, and we find that the reduced spending is not reflected in increased repayments by more than 15% (and 17% for individuals with credit card debt at baseline). This is true for the two months covering the intervention as well as the two months after the intervention.

Tables VI and IA.VII replicate the analysis in our main Tables V and IA.VI but focus on individuals in the top quartile of predicted treatment effects on saving rather than

[21]INEGI, Encuesta Nacional de Inclusion Financiera, 2018.

on spending.[22] The results are very similar. As in Table V, we show that the effect on borrowing is a very tightly estimated zero; we can rule out an increase in credit card borrowing costs of more than 0.30 cents and a decrease of more than 0.54 cents for each extra 1 MXN in saving. The same is true for the other outcome variables, such as balances and repayments.

Table IA.III presents the overlap between the treatment effect predictions of the causal forests for spending and saving (again, calculating quartiles with a cross-fitted ranking over five folds). If the overlap were perfect, the proportion of individuals at the intersection of both top quartiles would be 25%. The table shows that there is a significant overlap between the two, with around 15% to 17% of individuals in both the top quartiles of predicted treatment effects for spending and saving. This suggests that the same people reduce their spending and increase their savings (instead of repaying existing debt) as a result of the treatment.

D. Lessons for the Co-Holding Puzzle

Our empirical findings indicate that a significant number of individuals who pay credit card interest respond to saving nudges with notable reductions in spending and increases in savings. However, we also find that these additional savings are not utilized to pay off credit card debt during the current or subsequent billing cycles after receiving the nudge, exacerbating the co-holding of low-interest savings and high-interest debt.

In principle, the worsening of co-holding induced by the treatment could be driven by savers who increase their borrowing or borrowers who increase their savings. Given that we found no treatment effect on borrowing, the way in which the nudge exacerbates co-holding is through changes in savings among those who were already paying interest. In turn, the increase in savings among those carrying credit card interest could be coming from individuals who already had substantial savings and further increased them

[22] Note that the treatment effects on savings reported in Panels A and B of Table VI differ from the treatment effect reported for the top quartile in Panel D of Figure 3, as the latter considers all individuals in the top quartile of predicted treatment effects, including those with and without credit cards.

(intensive margin of co-holding) or from individuals who were not saving much to begin with but who, as a result of the treatment, may have started to save (extensive margin of co-holding). Table IA.IV shows that individuals most responsive to the treatment tend to carry higher checking account balances than those who respond the least. In addition, Table IA.V compares treatment effects for individuals in the top quartile of baseline checking account balances who were paying credit card interest at baseline to the rest of individuals paying credit card interest. Consistent with Table IA.IV, individuals in the top quartile of baseline checking account balances have a significantly higher treatment effect; however, this significance is weak, only reaching the 10% level. That said, given that these two independent tests both point towards a concentration of treatment effects on individuals with high checking account balances at baseline, we consider this suggestive evidence the treatment effect is concentrated on individuals who were already saving and an exacerbation of co-holding on the intensive margin.

The theoretical literature proposes two main explanations for co-holding: transaction-convenience and self- or partner-control. As shown, when the messages are interpreted as a shock to patience or required cash holdings, the self- or partner-control model predicts less co-movement between borrowing and saving than the transactions-convenience model. Because our empirical findings indicate no co-movement, they are more aligned with the self- or partner-control model. We note that, however, the self- or partner-control model also falls short of fitting the empirical finding quantitatively.

We now present three additional empirical patterns that we interpret in support of a self- or partner-control explanation for co-holding driving the marginal responses in spending, saving, and borrowing that we observe in response to the treatment.

First, as we will discuss in the next subsubsection, the largest treatment effect is associated with Message 4, which emphasizes the safety of savings and achieving goals. This message and its effect align with the behavioral hypothesis of self- or partner-control, wherein individuals constrain themselves to save more.

Second, in Table VII, we look at the types of spending that were reduced as a result of the treatment: ATM withdrawals, credit/debit card purchases, or transfers. In the development literature, cash-on-hand has been found to be subject to stronger intra-household sharing pressures, vis-a-vis cash in a bank account or digital wallet (Riley, 2024; Jakiela and Ozier, 2016). We find that about half of the reduction in spending comes from decreased ATM withdrawals (52% for all individuals with a credit card, and 55% for those who paid interest at baseline), which is consistent with individuals saving in a way that protects their cash from other household members. Nonetheless, we recognize that the other half of the reduction comes from decreases in card spending (and non-significant effects on transfers), implying that the reduction in spending may also be driven by other factors.[23]

Third, as outlined in Subsubsection D.3 and Table **??**, our analysis reveals that the variables examined in our study cannot predict any treatment effect heterogeneity in borrowing. This finding aligns with a model that suggests borrowing is not primarily driven by household needs but rather by the desire to constrain the spending capacity of the impatient party.

Effects by Treatment Message

Our next goal is to investigate whether the effects on saving and borrowing vary across the different treatment messages. We begin by narrowing our focus to the 150,177 individuals who are in the top quartile of the distribution of predicted treatment effects on spending and who have a credit card. We then calculate the treatment effect of each specific message on spending, saving, and borrowing.

Table VIII presents the effects on spending, saving, and borrowing of each message. It shows that the reduction in spending and increase in saving effect is large and significant

[23] In principle, ATM withdrawals could also be considered a proxy for discretionary spending, and reducing such expenditures could be a reasonable strategy for anyone trying to preserve liquidity. We note, however, that ATM withdrawals are a very coarse proxy for discretionary spending. On the one hand, discretionary spending can also take place with payment cards. In Mexico, the two most popular types of establishments for payments with cards are grocery stores and retail stores (Expansion, 2023). On the other hand, cash is also frequently used for rent payments (Instituto Nacional de Estadística y Geografía (INEGI), 2018).

for all individual messages, except for Messages 5 and 6, which have a non-significant effect on savings (the messages are displayed in Subsection B). We then group messages into short-term and long-term saving messages, keeping Message 4 separate, which has a distinct focus on the safeness of saving and goal-reaching. The pairwise comparison between the short-term and long-term messages shows no statistically significant difference in treatment effects. This suggests that our results apply to settings that aim to increase savings for the shorter and longer run.

Interestingly, Message 4 has a statistically significantly larger treatment effect on spending than both the short-term and the long-term message groups and a statistically significantly larger effect on saving than the short-term message groups. Message 4 message emphasizes the safeness of saving on a bank account and reaching goals. Although the differences are statistically significant only at the 10% level, this finding is consistent with the self- or partner-control hypothesis and the idea of constraining oneself to save more. Finally, it is noteworthy that we find null effects on borrowing for all individual messages.

E. Discussion and Robustness

Credit Constraints

It is possible that individuals do not increase their credit card borrowing because they are credit-constrained. Table IA.VIII and IA.IX repeat the main analysis for individuals with a credit card utilization below the median for individuals in the top quartile of the distribution of predicted treatment effects on spending and saving, respectively. Even among individuals who have enough credit limit available, we do not find changes in borrowing as a result of the treatment. We note that our descriptive statistics also confirm that individuals have ample space until they hit their credit limits.

Customers with Banorte as their Main Bank

Individuals who have accounts outside of Banorte may change their borrowing on those accounts after the treatment. To rule out this possibility, we replicate the analysis for individuals for whom Banorte is likely their main bank. Specifically, we look at the subsample of individuals for whom the following three conditions are satisfied: they receive their payroll on a Banorte payroll account, they have a credit card with Banorte, and they have no credit (of any type) outside of Banorte, according to the credit bureau records. Tables IA.X and IA.XI show the saving and borrowing results for this group of individuals. We find very similar results.

Overlap of Treatment Period and Billing Cycles

We note that since each individual has a different credit card due date and the due dates are distributed over the calendar month, their credit card billing cycles do not coincide perfectly with our 7-week treatment period. To alleviate the concern that this mismatch could drive the null effect on credit card borrowing, we run our main specifications considering only individuals for whom the treatment covers one entire billing cycle. For them, we look at treatment effects on spending and saving over the weeks that are perfectly aligned with their billing cycles (the one that was fully covered by the intervention). The results can be found in Table IA.XII. We can see that the main results continue to hold, and we do not find large changes in borrowing. We also show in Figure IA.2 that individuals' billing cycle end dates are pretty evenly distributed over the calendar month.

Dynamics of Treatment Effects on Spending and Saving

We also explore the possibility that individuals respond to the treatment only at the beginning of the intervention and that by the end of the observation period, the effect has disappeared. In Figures IA.3 and IA.4, we compare weekly treatment effects to the treatment effect on the first week of the intervention. We do so for all individuals who have a credit card, individuals in the top quartile of the distribution of predicted treatment effects, and individuals in the top quartile who also paid credit card interest at baseline. We can see that there are no statistically significant differences in the effectiveness of the treatment across treatment weeks. This finding confirms the validity of our main outcome measure, spending and saving over the first month of the treatment. As discussed, we use spending and saving over the first month of the treatment in order to have a monthly measure that we can compare to the monthly credit card interest charges.

We note that our design does not allow us to infer whether the treatment's effectiveness would have been retained had it lasted longer or if the increases in saving (and reductions in spending) would have persisted after the intervention. Consequently, we cannot determine whether the preference shock induced by the nudge was permanent or temporary.

Characterizing Individuals with the Largest Predicted Responses to the Treatment

To characterize individuals with the largest predicted responses to saving nudges, Table IA.IV compares the baseline characteristics of individuals in the top and bottom quartiles of the distribution of predicted treatment effects. Compared to individuals in the bottom quartile of the distribution of the predicted treatment effects, individuals with the highest predicted response are about one year younger and have slightly higher income, longer tenure with the bank, larger checking account balances, larger credit card balances, and larger credit card limits. All these can be thought of as proxies for trust and usage of the bank. Therefore, the algorithm appears to load on a very basic mechanism: when people use the bank more, they are more likely to pay attention to communications issued by the bank.

Interestingly, individuals with high predicted treatment effects have larger checking account balances. On the one hand, individuals in this group may be in better financial positions, which allows them to respond to the nudge. However, we note that while checking balances are substantially higher in the top quartile of the distribution of predicted treatment effects, the difference in average income is much smaller. Other characteristics we do not observe, such as higher financial literacy, frugality, or patience (among others), may jointly lead to higher baseline balances and a higher propensity to respond to saving nudges. If this is the case, then nudges would be more effective at the intensive margin, successfully leading to more savings among those already saving, but less so at the extensive margin, turning non-savers into savers. In Figure IA.1, we also see that checking account balances carries large variable importance. This is broadly consistent with preference-based explanations for higher saving rates among the rich (Dynan et al., 2004; Carroll, 1998) and the literature on poverty and self-control (Mullainathan and Shafir, 2013; Bernheim et al., 2015; Schilbach et al., 2016; Carvalho et al., 2016).

F. *Lessons for Heterogeneity Analysis in Academic and Policy Settings*

Our approach to identifying individuals with the largest response to the treatment is designed to avoid the overfitting bias that could arise when comparing treatment effects of multiple subpopulations while retaining high predictive power. In the following, we discuss the lessons we learned when comparing causal forests to traditional approaches.

Comparisons of Limited Power: Heterogeneity by Experimental Strata

The traditional way to analyze heterogeneous treatment effects is to split the sample based on strata from the experimental design. However, this approach relies on very coarse partitions of the covariate space, and it may be under-powered to identify the group with the largest treatment effect. In Table IX, we show the treatment effects on spending (Panel A) and saving (Panel B) across experimental strata. We find limited heterogeneity across the sub-populations that were pre-selected for heterogeneity analysis before the experiment was run. The largest treatment effects seem to be in the top quartile of the distribution of baseline checking account balances, although none of the differences are statistically significant.

To search for the group with the largest treatment effect, it might be tempting to further split the sample by overlaying strata dimensions and ultimately calculate and compare the treatment effects for each strata block.[24] However, as we now show, this

[24] We note that this is not the standard way in which people calculate heterogeneous treatment effects (and we are not aware of any study that has done so), but we use this as a limiting case of what would happen when trying to find heterogeneous treatment effects with a rich set of explanatory variables without considering the risk of overfitting.

approach can lead to substantial bias if not adjusted for overfitting.

The Pitfalls of Overfitting: Sorting by Observed Treatment Effects at the Strata-Block Level

To demonstrate the bias arising from overfitting, we divided the sample into 6,104 distinct and mutually exclusive blocks based on the interaction of all experimental strata. We then compute treatment effects on spending for each block and assign each observation the treatment effect of its corresponding group. Next, we split the sample into quartiles based on the assigned treatment effects and identify the top quartile, which consists of observations from strata blocks with the highest observed treatment effects. We then estimate the treatment effects on checking account balances and credit card interest in this top quartile.

Table X displays the results of this analysis. Columns (1) to (3) report treatment effects for individuals in strata blocks with the largest observed treatment effects. The findings reveal significant reductions in spending, of 38%. Additionally, these individuals experienced substantial increases in savings and decreases in borrowing, as evidenced in Columns (2) and (3).

In comparison, the results obtained from the causal forest are shown in Columns (4) to (6). Column (4) shows that the increases in spending are smaller, in the order of 7%. Columns (5) and (6) show the corresponding treatment effects on saving and borrowing. These estimates are not subject to overfitting bias. The coefficients on borrowing in Column (6) are significantly closer to zero than those in Column (3).

Policy and Business Applications

This analysis has important implications for policy-makers and businesses alike. Many policies aim to target individuals for whom the treatment's impact is expected to be the largest in order to be cost-effective. For instance, in the case of a policy distributing cash transfers to stimulate the economy within a budget constraint, it would be optimal to first reach individuals with the highest marginal propensity to spend. Our analysis suggests that if policymakers try to identify these individuals using observational data from previous fiscal stimulus payments and compare the marginal propensities to consume of different subpopulations through a manual search, they may encounter two problems.

First, suppose the group they identify is based on a coarse partition of the covariate space, such as low-income individuals. In that case, they may treat individuals who do not need the money because they have high liquid wealth despite their low income. Alternatively, suppose they split and sort the sample by many dimensions, such as income, wealth, education, age, etc. In that case, they may overfit the data, resulting in identifying a subpopulation that has a seemingly high spending rate that would not replicate out of sample.

From a business perspective, estimating individual treatment effects enables firms to maximize profits by treating only clients for whom the expected revenue from the treatment exceeds the cost. For instance, in the experiment analyzed here, the expected revenue is the monetary value of incremental deposits, and the cost of the treatment is the cost of sending SMS messages. By searching over pre-specified strata, Panel B of Table IX shows that we have little guidance as to what is the most responsive group since none of the interactions are statistically significant. If anything, the most promising partition appears to be quartiles of baseline checking account balances (see Column (1) of Table IX, Panel B), since for these partitions we see increasing effects. For individuals in the top quartile of baseline checking account balances, we find a treatment effect of 1.2% (-0.002 + 0.014). Targeting this group would yield a total increase in balances of approximately 192.8 million MXN by sending messages to 750,000 clients (1.2% increase from a base of 21,423 MXN for 750,000 clients).

However, treating only the top 5% of individuals in the distribution of predicted treatment effects would reduce costs by 95%, as messages would be sent only to the 150,000 clients most likely to respond. For the top 5% of individuals in the distribution of predicted treatment effects, we find a 1,164.47 MXN increase in balances, which aggregate to 174.64 million MXN (150,000*1,164.47 MXN). This results in a more efficient cost-revenue ratio. With a causal forest, firms can sort clients based on their predicted revenue and treat them until marginal revenue equals marginal cost.

In addition, there may be an interaction between business and policy implications. Using the two examples above, a bank may target individuals most responsive to the savings message to maximize profits. At the same time, a government targets individuals for fiscal transfers. The implications of this interaction will likely depend on the overlap of the two target populations and on whether the private and public interventions are complements or substitutes. If the two target populations were to overlap and the interventions were substitutes, the treatment effects on spending and saving could counteract each other. An estimate of the overall impact of the interventions would require training the models with data from a period in which both programs are in place or structural assumptions about their interactions.

We also note that a profit-maximizing targeting policy may fail to prioritize consumers with the highest benefit from the treatment. In our context, Table IA.IV shows that the group most responsive to saving nudges had higher checking account balances at baseline. Targeting them is optimal from the perspective of a bank balancing operational costs and revenue from deposits. However, they may already have sufficient buffer to face financial shocks and may not derive the highest benefits from the additional savings. Other machine learning models with more complex objective functions could be used to find socially optimal targeting policies (Rolf et al., 2020).

IV. Conclusion

Our study uses a large-scale field experiment involving 3.1 million individuals to examine the economic mechanisms behind co-holding low-interest savings and credit card debt. Our results provide evidence that encouraging individuals who are paying credit card interest to save can lead to increased savings, regardless of existing debt levels.

We then compare our findings with the predictions of existing models of co-holding. These models can be categorized into two groups: those explaining co-holding through a liquidity premium for cash and attributing it to self- or partner control.

When our saving messages are interpreted as a shock to patience, we show that the transactions-convenience model predicts a decrease in spending, reflected in a decrease in borrowing. On the other hand, when our saving messages are interpreted as a shock to cash needs or hidden cash, the transactions-convenience model predicts that spending is unchanged but saving is increased, reflected in an increase in borrowing. In contrast, the self- or partner-control model predicts less co-movement between saving and borrowing in both situations. When the patient party becomes more patient and increases hidden cash, then the impatient party decreases their spending rather than increases their borrowing.

We then demonstrate that the results of our experiment are more consistent with theories based on self- or partner-control, as individuals decrease their spending and increase their savings without changing their borrowing very much.

Finally, we highlight the benefits of state-of-the-art machine-learning methods for causal inference and discuss their implications for targeting that can allow both public and private institutions to influence consumer behavior cost-effectively.

REFERENCES

Abadie, Alberto, Matthew M Chingos, and Martin R West, 2018, Endogenous stratification in randomized experiments, *Review of Economics and Statistics* 100, 567–580.

Afzal, Uzma, Giovanna d'Adda, Marcel Fafchamps, Simon R Quinn, and Farah Said, 2019, Implicit and explicit commitment in credit and saving contracts: A field experiment, Technical report, National Bureau of Economic Research.

Akbas, Merve, Dan Ariely, David A Robalino, Michael Weber, et al., 2016, How to help poor informal workers to save a bit: Evidence from a field experiment in kenya, *IZA Discussion Papers* 10024.

Allcott, Hunt, and Judd B Kessler, 2019, The welfare effects of nudges: A case study of energy use social comparisons, *American Economic Journal: Applied Economics* 11, 236–76.

Ashraf, Nava, Natalie Bau, Corinne Low, and Kathleen McGinn, 2020, Negotiating a better future: How interpersonal skills facilitate intergenerational investment, *The Quarterly Journal of Economics* 135, 1095–1151.

Ashraf, Nava, Nathalie Gons, Dean S Karlan, and Wesley Yin, 2003, A review of commitment savings products in development countries, *ADB Economics Working Paper Series* 45.

Athey, Susan, Vitor Hadad, Niall Keleher, Octavio Medina, Rebecca Nissan, Rachel Rosemberg, Matt Schaelling, Jann Spiess, and Josh Wright, 2021, Computational applications to behavioral science, Technical report.

Athey, Susan, and Guido W Imbens, 2016, Recursive partitioning for heterogeneous causal effects, *Proceedings of the National Academy of Sciences* 113, 7353–7360.

Athey, Susan, Julie Tibshirani, Stefan Wager, et al., 2019, Generalized random forests, *The Annals of Statistics* 47, 1148–1178.

Athey, Susan, and Stefan Wager, 2019, Estimating treatment effects with causal forests: An application, *arXiv preprint* arXiv:1902.07409.

Bachas, Pierre, Paul Gertler, Sean Higgins, and Enrique Seira, 2021, How debit cards enable the poor to save more, *The Journal of Finance* 76, 1913–1957.

Batista, Rafael M, Ella Mao, and Abigail B Sussman, 2023, A psychological account of co-holding: Why some choose to simultaneously save and borrow, *Available at SSRN* 4558490.

Baumol, William J., 1952, The transactions demand for cash: An inventory theoretic approach, *The Quarterly Journal of Economics* 66, 545–556.

Beaman, Lori, Dean Karlan, Bram Thuysbaert, and Christopher Udry, 2023, Selection into credit markets: Evidence from agriculture in mali, *Econometrica* 91, 1595–1627.

Benartzi, Shlomo, John Beshears, Katherine L Milkman, Cass R Sunstein, Richard H Thaler, Maya Shankar, Will Tucker-Ray, William J Congdon, and Steven Galing, 2017, Should governments invest more in nudging?, *Psychological Science* 28, 1041–1055.

Bernheim, B Douglas, Debraj Ray, and Şevin Yeltekin, 2015, Poverty and self-control, *Econometrica* 83, 1877–1911.

Bertaut, Carol C, Michael Haliassos, and Michael Reiter, 2009, Credit card debt puzzles and debt revolvers for self control, *Review of Finance* 13, 657–692.

Beshears, John, Matthew Blakstad, James J Choi, Christopher Firth, John Gathergood, David Laibson, Richard Notley, Jesal D Sheth, Will Sandbrook, and Neil Stewart, 2024, Does pension automatic enrollment increase debt? evidence from a large-scale natural experiment, Technical report, National Bureau of Economic Research.

Beshears, John, James J Choi, David Laibson, Brigitte C Madrian, and Katherine L Milkman, 2015, The effect of providing peer information on retirement savings decisions, *The Journal of Finance* 70, 1161–1201.

Beshears, John, James J Choi, David Laibson, Brigitte C Madrian, and William L Skimmyhorn, 2019, Borrowing to save? the impact of automatic enrollment on debt, Technical report, National Bureau of Economic Research.

Beshears, John, and Harry Kosowsky, 2020, Nudging: Progress to date and future directions, *Organizational Behavior and Human Decision Processes* 161, 3–19.

Breza, Emily, Martin Kanz, and Leora F Klapper, 2020, Learning to navigate a new financial technology: Evidence from payroll accounts, Technical report, National Bureau of Economic Research.

Burke, Jeremy, Julian Jamison, Dean Karlan, Kata Mihaly, and Jonathan Zinman, 2020, Credit building or credit crumbling? a credit builder loan's effects on consumer behavior, credit scores and their predictive power, Technical report, National Bureau of Economic Research.

Carlana, Michela, Eliana La Ferrara, and Paolo Pinotti, 2022, Goals and gaps: Educational careers of immigrant children, *Econometrica* 90, 1–29.

Carroll, Christopher D, 1998, Why do the rich save so much?, Technical report.

Carvalho, Leandro S, Stephan Meier, and Stephanie W Wang, 2016, Poverty and economic decision-making: Evidence from changes in financial resources at payday, *American Economic Review* 106, 260–284.

Chen, Jiafeng, and Jonathan Roth, 2024, Logs with zeros? some problems and solutions, *The Quarterly Journal of Economics* 139, 891–936.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Ivan Fernandez-Val, 2018, Generic machine learning inference on heterogenous treatment effects in randomized experiments, Technical report, National Bureau of Economic Research.

Chetty, Raj, John N Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen, 2014, Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from denmark, *The Quarterly Journal of Economics* 129, 1141–1219.

Choi, James J, David Laibson, Brigitte C Madrian, and Andrew Metrick, 2004, For better or for worse: Default effects and 401 (k) savings behavior, in *Perspectives on the Economics of Aging*, 81–126 (University of Chicago Press).

Choukhmane, Taha, 2019, Default options and retirement saving dynamics, in *112th Annual Conference on Taxation. NTA.*

Choukhmane, Taha, and Christopher Palmer, 2023, How do consumers finance increased retirement savings, Technical report, Working Paper.

Cohn, Jonathan B, Zack Liu, and Malcolm I Wardlaw, 2022, Count (and count-like) data in finance, *Journal of Financial Economics* 146, 529–551.

Costa, Dora L, and Matthew E Kahn, 2013, Energy conservation "nudges" and environmentalist ideology: Evidence from a randomized residential electricity field experiment, *Journal of the European Economic Association* 11, 680–702.

Davis, Jonathan MV, and Sara B Heller, 2020, Rethinking the benefits of youth employment programs: The heterogeneous effects of summer jobs, *Review of Economics and Statistics* 102, 664–677.

Druedahl, Jeppe, and Casper Nordal Jørgensen, 2018, Precautionary borrowing and the credit card debt puzzle, *Quantitative Economics* 9, 785–823.

Dynan, Karen E, Jonathan Skinner, and Stephen P Zeldes, 2004, Do the rich save more?, *Journal of Political Economy* 112, 397–444.

Expansion, 2023, El uso de tarjetas de crédito y débito anota récord en mayo.

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

Gargano, Antonio, and Alberto G Rossi, 2024, Goal setting and saving in the fintech era, *The Journal of Finance* 79, 1931–1976.

Gathergood, John, and Arna Olafsson, 2024, The coholding puzzle: New evidence from transaction-level data, *The Review of Financial Studies* hhae016.

Gathergood, John, and Jörg Weber, 2014, Self-control, financial literacy & the co-holding puzzle, *Journal of Economic Behavior & Organization* 107, 455–469.

Goldin, Jacob, Tatiana Homonoff, and Will Tucker-Ray, 2017, Retirement contribution rate nudges and plan participation: Evidence from a field experiment, *American Economic Review* 107, 456–61.

Gorbachev, Olga, and María José Luengo-Prado, 2019, The credit card debt puzzle: The role of preferences, credit access risk, and financial literacy, *Review of Economics and Statistics* 101, 294–309.

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, 149–185.

Halpern, David, 2015, Inside the nudge unit: How small changes can make a big difference, *Random House* ISBN-10, â0753556537.

Higgins, Sean, 2020, Financial technology adoption: Network externalities of cashless payments in mexico, *American Economic Review* forthcoming.

Instituto Nacional de Estadística y Geografía (INEGI), 2018, Encuesta nacional de inclusión financiera 2018, Online.

Jakiela, Pamela, and Owen Ozier, 2016, Does africa need a rotten kin theorem? experimental evidence from village economies, *The Review of Economic Studies* 83, 231–268.

Karlan, Dean, Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman, 2016, Getting to the top of mind: How reminders increase saving, *Management Science* 62, 3393–3411.

Laibson, D, 1997, Golden Eggs and Hyperbolic Discounting, *Quarterly Journal of Economics* 112, 443–477.

Laibson, David, Sean Chanwook Lee, Peter Maxted, Andrea Repetto, and Jeremy Tobacman, 2024, Estimating discount functions with consumption choices over the lifecycle, *The Review of Financial Studies* hhae035.

Maki, Dean M, 2002, The growth of consumer credit and the household debt service burden, in *The Impact of Public Policy on Consumer Credit*, 43–68 (Springer).

Medina, Paolina C, 2021, Side effects of nudging: Evidence from a randomized intervention in the credit card market, *The Review of Financial Studies* 34, 2580–2607.

Mullainathan, Sendhil, and Eldar Shafir, 2013, Decision making and policy in contexts of poverty, *Behavioral Foundations of Public Policy* 281–300.

Riley, Emma, 2024, Resisting social pressure in the household using mobile money: Experimental evidence on microenterprise investment in uganda, *American Economic Review* 114, 1415–1447.

Robinson, Peter M, 1988, Root-n-consistent semiparametric regression, *Econometrica* 931–954.

Rodríguez, Catherine, and Juan Saavedra, 2015, Nudging youth to develop savings habits: Experimental evidence using sms messages, *CESR-Schaeffer Working Paper* 2015-018.

Rolf, Esther, Max Simchowitz, Sarah Dean, Lydia T Liu, Daniel Bjorkegren, Moritz Hardt, and Joshua Blumenstock, 2020, Balancing competing objectives with noisy data: Score-based classifiers for welfare-aware machine learning, in *International Conference on Machine Learning*, 8158–8168, PMLR.

Schilbach, Frank, Heather Schofield, and Sendhil Mullainathan, 2016, The psychological lives of the poor, *American Economic Review* 106, 435–440.

Telyukova, Irina A, 2013, Household need for liquidity and the credit card debt puzzle, *Review of Economic Studies* 80, 1148–1177.

Thaler, Richard H, 1994, Psychology and savings policies, *The American Economic Review* 84, 186–192.

Tibshirani, Julie, Susan Athey, and S Wager, 2020, grf: Generalized random forests. r package version 120.

Tobin, James, 1956, The interest-elasticity of transactions demand for cash, *The Review of Economics and Statistics* 38, 241–247.

Tobin, James, 1958, Liquidity preference as behavior towards risk, *The Review of Economic Studies* 25, 65–86.

Vihriälä, Erkki, 2019, Untangling the credit card debt puzzle, Technical report, Working Paper.

Wisdom, Jessica, Julie S Downs, and George Loewenstein, 2010, Promoting healthy choices: Information versus convenience, *American Economic Journal: Applied Economics* 2, 164–78.

Zinman, Jonathan, 2007, Household borrowing high and lending low under no-arbitrage, *Dartmouth University* Citeseer.

Figures and Tables

Figure 1. Self- or Partner-Control Model: Patient Party's Hidden Cash and the Impatient Party's Debt as a Function of the Impatient Party's Impatience and the Fraction of Cash that can be Hidden

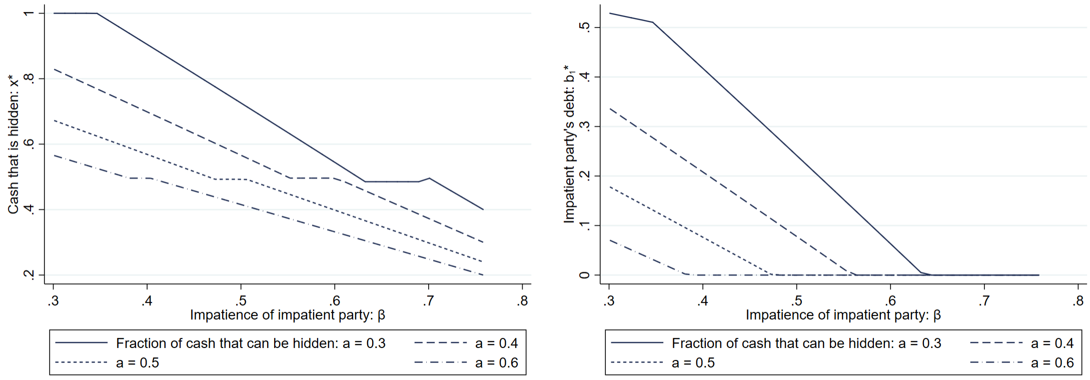

This graph shows the amount of cash that the patient party hides, x^*, as well as the debt that is borrowed by the impatient party in period 1, b_1^*, as a function of the impatient party's impatience, β, for different values of the fraction of cash that can be hidden, a. The range of the impatient party's discount factor equals the estimate ($\beta = 0.9891^{2/12} * 0.5305$) +/- two times the standard error estimate (0.1143) from the structural estimation in Laibson et al. (2024). The patient party's discount factor is assumed to be $\delta = 0.9891^{2/12}$, corresponding to the exponential estimate in Laibson et al. (2024) (of the hyperbolic model) scaled to a 2-months period as in the experiment, the initial endowment is normalized to 1, $x_1 = 1$, and the interest rate for borrowing is set to $r = 0.1152 * 2/12$, the first-stage estimate in Laibson et al. (2024), scaled to a 2-months period as in the experiment.

Figure 2. Self- or Partner-Control Model: Consumption in Periods 1 and 2 as a Function of the Impatient Party's Impatience and the Fraction of Cash that can be Hidden

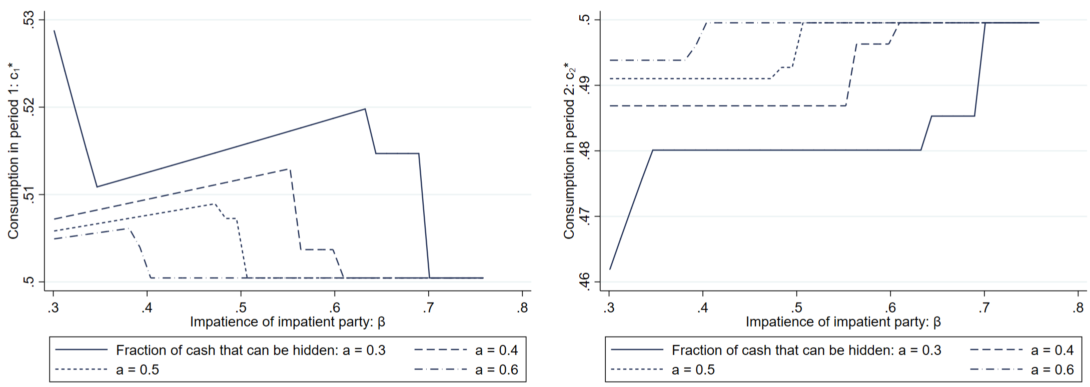

This graph shows consumption in periods 1 and 2, c_1^* and c_2^*, as a function of the impatient party's impatience, β, for different values of the fraction of cash that can be hidden, a. The calibration is the same as in Figure 1.

Table I. Comparison of Transactions-Convenience and Self- or Partner-Control Models: Simultaneous Cash and Debt Holdings for different Values of Patience and Cash Visibility

	Fraction of cash that can be hidden: a			
	a = 0.3	a = 0.4	a = 0.5	a = 0.6
Hidden cash, x^*, or cash needed for transactions, x	0.979	0.754	0.613	0.516
Debt in self- or partner-control model, b^*_1	0.491	0.262	0.120	0.0217
Co-movement b^*_1 and x^* in self- or partner-control model	0.770	0.698	0.625	0.553
Debt in transactions-convenience model	0.471	0.248	0.108	0.0119
Co-movement b^*_1 and x in transactions-convenience model	0.991	0.991	0.991	0.991
Additional debt in self- or partner-control model (in percent)	4.089	5.882	10.88	82.95
Decrease in co-movement in self- or partner-control model (in percentage points)	22.12	29.35	36.59	43.82

This table shows simulation results of the self-or partner-control model and the transaction-convenience model for 100 random realizations of the patient party's patience parameter delta, δ, and four intermediate values of the fraction of cash that can be hidden, a. δ is drawn from a random normal distribution with mean 0.9891 and a standard deviation of 0.0053, corresponding to the estimate and standard error for this parameter in Laibson et al. (2024). Each realization is then scaled by $\delta^{2/12}$ to match the 2-month period of the intervention we study. Rows 1 to 3 pertain to the self- or partner-control model. The first row shows the resulting hidden cash, x^*, averaged over the 100 different realizations of δ. Row 2 shows the resulting amount of debt held by the impatient party averaged over the 100 different model simulations. Row 3 depicts the level of co-movement between borrowing and saving, which equals the estimate obtained by regressing borrowing, b_1^*, on cash holdings, x^*. Rows 4 and 5 pertain to the transaction-convenience model. The cash needed for transaction purposes is assumed to equal the hidden cash, x^*, in the self- or partner-control model. Row 4 shows the resulting amount of debt held by the transactions-convenience agent again averaged over the 100 different model simulations, and Row 5 depicts the level of co-movement between borrowing and saving. Rows 6 and 7 compare the levels of debt and co-movement in the two models. The impatient party's impatience is set to a lower confidence range value of the estimate in Laibson et al. (2024), $\beta = 0.9891 * 0.5305 - 1.5 * 0.1143$, the transaction-convenience agent's patience is set to $0.9891^{2/12}$, the initial endowment is normalized to 1, $x_1 = 1$, and the interest rate for borrowing is set to $r = 0.1152 * 2/12$.

Table II. Descriptive Statistics

All Individuals (N= 3,054,503)	Mean	Std. Dev.	P25	P50	P75
Age (Years)	45	16	31	43	56
Monthly Income	13,499	13,712	6,117	9,867	15,006
Tenure (Months)	82	73	22	59	125
Monthly Spending	17,940	41,755	3,107	9,034	13,278
Checking Account Balance	20,006	53,714	742	2,477	11,360
Has Credit Card	0.12	0.32	0.00	0.00	0.00
Monthly Credit Card Interest	32	180	0	0	0
Credit Card Balance	2,194	6,507	0	0	0
Statement Balance - Payments	603	777	0	0	0
Credit Card Limit	10,109	20,188	0	0	0
Individuals with a Credit Card (N= 362,223)	**Mean**	**Std. Dev.**	**P25**	**P50**	**P75**
Age (Years)	43	13	33	42	53
Monthly Income	19,745	18,654	9,071	13,913	22,718
Tenure (Months)	104	73	43	86	149
Monthly Spending	31,012	71,746	6,182	18,063	21,145
Checking Account Balance	33,963	70,437	1,581	5,157	23,069
Monthly Credit Card Interest	270	412	0	0	170
Credit Card Balance	18,282	31,480	104	10,458	27,137
Statement Balance - Payments	5,027	7,138	0	0	2,980
Credit Card Limit	84,244	110,546	15,000	45,000	100,000

This table presents summary statistics for all individuals in the experiment, and for the subset of individuals who have a credit card. The experimental pool of 3,054,503 customers is taken as a random sample from the universe of the bank's customers satisfying three requirements: individuals had a payroll account, kept an average daily balance of at least 50 MXN over the 2 months previous to the intervention, and had a valid cell phone number to receive SMS messages. The statistics are calculated with monthly information at the individual level covering the 6 months previous to the intervention. Monthly Income refers to salary transfers flagged as such in the payroll accounts. Tenure with the bank corresponds to the number of months since the accounts were opened. Monthly Spending is defined as the sum of ATM withdrawals, credit and debit card spending, and outgoing transfers. Checking Account Balances are the average daily balances in the checking (payroll) accounts. Credit Card Interest reflects interest charges over a given billing cycle on Banorte credit cards. Statement Balance - Payments refers to the statement balance for the given billing cycle minus the payments received on the credit card account over the following billing cycle i.e., the payments made towards the corresponding bill. Spending, balances, interest charges, payments, and credit limits are in Mexican Pesos (MXN). 1 MXN = 0.107 USD PPP.

Table III. Checking and Credit Card Account Balances for Individuals Who Have a Credit Card and Carry Credit Card Interest - By Deciles of Average Daily Balance on Checking Accounts, Over Income

Decile	Minimum Checking Balance Divided by Income (Minimum)	Minimum Checking Balance (Average)	Monthly Income (Average)	Credit Card Balance (Average)	Monthly Credit Card Interest (Average)	Credit Card Interest over Income (Average)	Fraction with Savings Larger than Credit Card Balance	Fraction with Savings Larger than 50% of Income
1	0.0000	0	17,752	20,031	601	0.04	0.00	0.00
2	0.0002	11	20,065	18,459	549	0.03	0.01	0.00
3	0.0006	37	18,171	18,942	554	0.04	0.01	0.00
4	0.0027	78	19,466	19,556	590	0.03	0.02	0.00
5	0.0048	243	20,639	23,136	702	0.06	0.01	0.00
6	0.0165	934	19,629	20,913	643	0.04	0.19	0.00
7	0.0575	2,750	22,203	21,651	633	0.03	0.34	0.00
8	0.1479	6,451	21,684	23,007	693	0.05	0.46	0.00
9	0.3508	18,723	22,454	20,442	611	0.04	0.63	0.57
10	0.9493	117,510	23,827	22,585	660	0.04	0.89	1.00
All		14,673	20,589	20,872	627	0.04	0.26	0.16

This table presents statistics about credit card borrowing and checking account balances for 111,999 individuals who have a credit card and paid credit card interest in the month previous to the intervention. For checking account balances, we consider the minimum balance observed for each individual over the 6 months previous to the intervention, which we refer to as the minimum checking balance. All other variables correspond to the month before the intervention. Individuals are split into deciles of minimum checking balance divided by income. For observations in each decile group, we present the minimum (across individuals) of the minimum checking balance divided by income, as well as the average (across individuals) of said minimum checking balance, monthly income, monthly credit card balances, and monthly interest charges. We also show the fraction of individuals with checking account balances larger than or equal to their credit card balances, the ratio of monthly interest charges to monthly income, and the fraction of individuals with checking account balances larger than 50 percent of their monthly income. Income refers to after tax monthly income. Balances, income, and interest charges are in Mexican Pesos (MXN). 1 MXN = 0.107 USD PPP.

Table IV. Aggregate Effects of the Intervention

	(1)	(2)	(3)	(4)	(5)
Dep. Var.	Monthly Spending	Checking Account Balance	Monthly Credit Card Interest During Treat.	Paid Interest During Treat. {0,1}	Ending Statement Balance - Payments
Panel A: All Individuals					
Any Treatment	-0.008*	0.006**			
	(0.005)	(0.003)			
Observations	3,054,503	3,054,503			
Mean of Dep. Var. in Control Group	17,870.06	19,913.02			
Panel B: Individuals with a Credit Card					
Any Treatment	-0.019***	0.011**	-0.002	-0.001	-0.002
	(0.006)	(0.005)	(0.008)	(0.004)	(0.006)
Observations	362,223	362,223	362,223	362,223	362,223
Mean of Dep. Var. in Control Group	31,997.80	35,657.27	215.91	0.41	5,028.69
Panel C: Individuals with a Credit Card Who Paid Interest at Baseline					
Any Treatment	-0.018***	0.014**	-0.003	-0.001	-0.003
	(0.007)	(0.005)	(0.008)	(0.005)	(0.006)
Observations	152,016	152,016	152,016	152,016	152,016
Mean of Dep. Var. in Control Group	33,982.15	32,929.96	483.77	0.81	10,317.99

This table presents the results of estimating Equation (3) for spending, saving, and borrowing. Variable definitions are provided in Table 2. Monthly Spending and Checking Account Balances are measured over the first month of the intervention. Credit Card Interest is measured as the average of the monthly interest charges during the two billing cycles affected by the intervention. Paid Interest is a binary variable flagging whether an individual was charged credit card interest on any of the two billing cycles affected by the intervention. Ending Statement Balance - Payments is measured for the last billing cycle affected by the intervention. Panel A considers all individuals. Panel B considers all individuals who have a credit card. Panel C considers only individuals who have a credit card and incurred interest at baseline. Any Treatment is a binary variable that takes the value of one if an individual is assigned to receive any of the treatment messages. Treatment effects are calculated using Poisson regressions with strata fixed effects. We report percentage treatment effects calculated as exponentiated coefficients -1. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

Figure 3. Distribution of Predicted Treatment Effects and Actual Treatment Effects

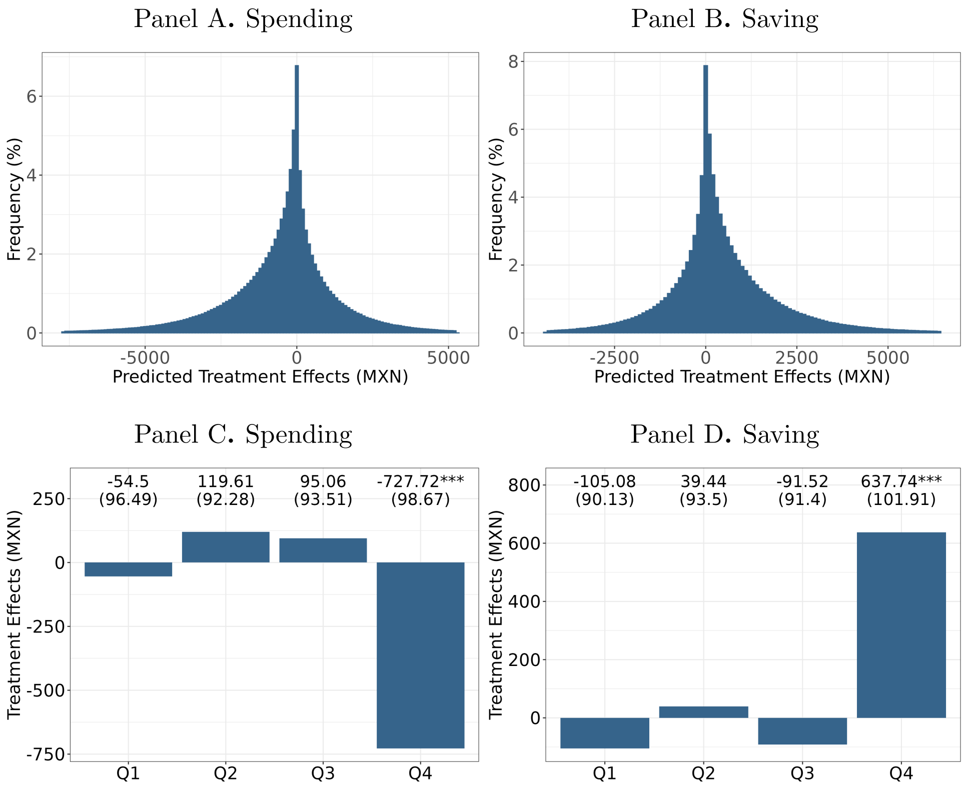

This figure considers all individuals in the experimental pool, including those with and without a credit card. Panels A and B show the distribution of predicted treatment effects at the individual level obtained from causal forests trained with 2,000 trees. The plots are truncated at the 1 and 99 percentiles. The outcome variables are Monthly Spending and Checking Account Balances, respectively. Explanatory variables are described in Subsection B and Figure IA.1. Spending is defined as the sum of ATM withdrawals, credit and debit card spending, and outgoing transfers over the first month of the intervention. Checking account balances corresponds to average daily balances during the first month of the intervention. Panels C and D show treatment effects on spending and checking account balances, as a function of predicted treatment effects. All 3,054,503 individuals are split into quartiles of predicted treatment effects on spending and savings based on cross-fitted rankings over five folds. For spending, the top quartile corresponds to the most negative effect. For saving, the top quartile corresponds to the most positive effect.

Figure 4. Distribution of Individuals Who Have a Credit Card, Who Pay Credit Card Interest or Who Co-hold Credit Card Debt and Liquid Savings, Across Quartiles of Predicted Treatment Effects

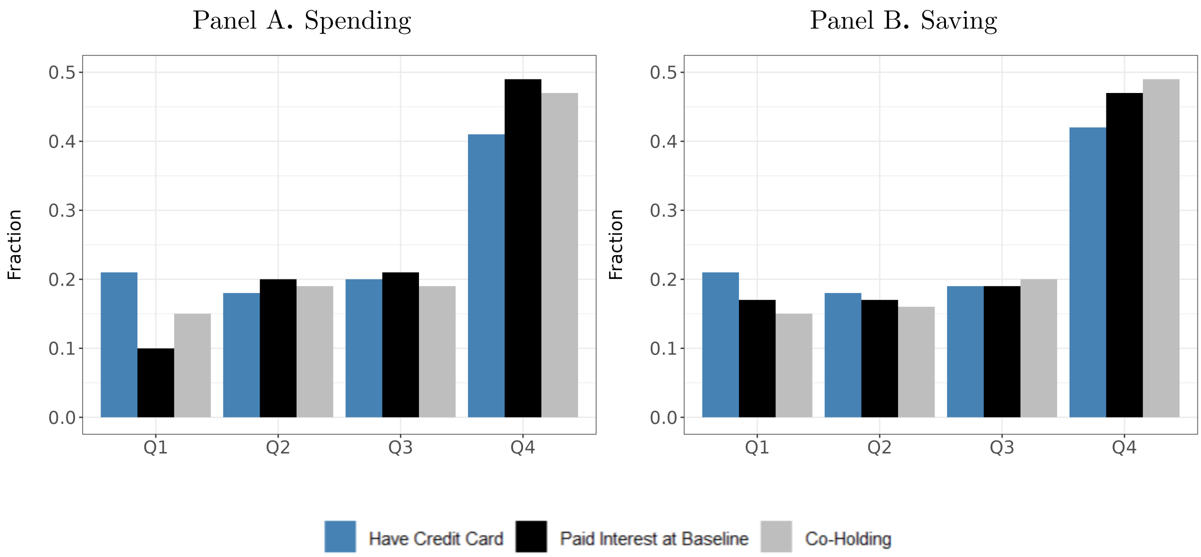

This graph considers three distinct groups of individuals: individuals who have a credit card, individuals who paid credit card interest at baseline, and individuals who paid credit card interest at baseline while keeping at least 50% of their monthly after-tax income on their checking accounts over the 6 months previous to the intervention. Panels A and B plot their distributions across quartiles of predicted treatment effects on spending and saving, respectively. For spending, the top quartile corresponds to the most negative effect. For saving, the top quartile corresponds to the most positive effect. Bars of the same color add up to 1.

Table V. Treatment Effects on Spending, Saving, and Borrowing (Top Quartile of Predicted Treatment Effects on Spending)

Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Monthly Spending	Checking Account Balance	Monthly Credit Card Interest During Treat. (Banorte)	Monthly Credit Card Interest After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ending Statement Balance - Payments
Panel A: Individuals with a Credit Card							
% TE	-0.0722***	0.0493***	-0.0101	-0.0110	-0.0037	-0.0041	-0.0016
	(0.0133)	(0.0128)	(0.0221)	(0.0223)	(0.0094)	(0.0094)	(0.0308)
Mean of Dep. Var. in Control Group (MXN)	34,968.92	39,174.23	210.16	208.55	0.47	0.46	5,148.73
Change in Spending or Saving (MXN)	-2,524.28	1,932.02					
Upper Confidence Interval (MXN)			6.98	6.83	0.01	0.01	302.58
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0028	0.0027	0.0000	0.0000	0.1199
Lower Confidence Interval (MXN)			-11.23	-11.43	-0.01	-0.01	-319.06
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0044	-0.0045	-0.0000	-0.0000	-0.1264
N= 150177							
Panel B: Individuals with a Credit Card Who Paid Interest at Baseline							
% TE	-0.0688***	0.0555***	-0.0099	-0.0102	-0.0054	-0.0051	-0.0038
	(0.0170)	(0.0173)	(0.0229)	(0.0232)	(0.0102)	(0.0102)	(0.0180)
Mean of Dep. Var. in Control Group (MXN)	35,550.10	34,913.86	384.09	385.84	0.85	0.86	10,613.58
Change in Spending or Saving (MXN)	-2,444.22	1,936.38					
Upper Confidence Interval (MXN)			13.42	13.58	0.01	0.01	334.37
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0055	0.0056	0.0000	0.0000	0.1368
Lower Confidence Interval (MXN)			-21.03	-21.46	-0.02	-0.02	-414.67
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0086	-0.0088	-0.0000	-0.0000	-0.1697
N= 73946							

This table shows treatment effects for individuals in the top quartile of predicted treatment effects on spending who have a credit card. Variable definitions are provided in Table 2. Monthly Spending and Checking Account Balances are measured over the first month of the intervention. Credit Card Interest is measured as the average of the monthly interest charges during the two billing cycles affected by the intervention or immediately following it. Paid Interest is a binary variable flagging whether an individual was charged credit card interest on any of the two billing cycles affected by the intervention or immediately following it. Ending Statement Balance - Payments is measured for the last billing cycle affected by the intervention. Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest at baseline. Treatment effects are calculated using Poisson regressions with strata fixed effects. We report percentage treatment effects calculated as exponentiated coefficients -1. The change in spending (saving), expressed in MXN, is calculated by multiplying the % TE and the Mean of Spending (Checking Account Balances) in the Control Group. Upper (lower) confidence intervals, expressed in MXN, are calculated as (point estimate +/- 1.96*Standard Error)*Mean of Dep. Var. in Control Group. The upper (lower) confidence interval for the variable Paid Interest is calculated as (point estimate +/- 1.96*Standard Error). Abs. Value stands for absolute value. The number of observations is the same across all columns in the same panel. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

Table VI. Treatment Effects on Spending, Saving, and Borrowing (Top Quartile of Predicted Treatment Effects on Saving)

Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Monthly Spending	Checking Account Balance	Monthly Credit Card Interest During Treat. (Banorte)	Monthly Credit Card Interest After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ending Statement Balance - Payments
Panel A: All Individuals with a Credit Card							
% TE	-0.0504***	0.0601***	-0.0115	-0.0097	-0.0038	-0.0040	-0.0015
	(0.0167)	(0.0148)	(0.0206)	(0.0211)	(0.0065)	(0.0065)	(0.0302)
Mean of Dep. Var. in Control Group (MXN)	39,242.16	34,697.03	216.53	219.12	0.46	0.44	4,817.55
Change in Spending or Saving (MXN)	-1,978.95	2,084.12					
Upper Confidence Interval (MXN)			6.25	6.94	0.00	0.00	277.93
Upper Confidence Interval Divided by Abs. Value of Change in Saving			0.0030	0.0033	0.0000	0.0000	0.1334
Lower Confidence Interval (MXN)			-11.23	-11.19	-0.01	-0.01	-292.39
Lower Confidence Interval Divided by Abs. Value of Change in Saving			-0.0054	-0.0054	-0.0000	-0.0000	-0.1403
N= 151834							
Panel B: Individuals with a Credit Card Who Paid Interest at Baseline							
% TE	-0.0639***	0.0622***	-0.0101	-0.0099	0.0011	-0.0036	-0.0138
	(0.0214)	(0.0220)	(0.0225)	(0.0223)	(0.0095)	(0.0094)	(0.0263)
Mean of Dep. Var. in Control Group (MXN)	34,453.77	35,069.81	407.71	406.25	0.73	0.74	10,211.85
Change in Spending or Saving (MXN)	-2,201.33	2,179.94					
Upper Confidence Interval (MXN)			13.87	13.78	0.01	0.01	386.77
Upper Confidence Interval Divided by Abs. Value of Change in Saving			0.0064	0.0063	0.0000	0.0000	0.1774
Lower Confidence Interval (MXN)			-22.13	-21.81	-0.01	-0.02	-667.66
Lower Confidence Interval Divided by Abs. Value of Change in Saving			-0.0102	-0.0100	-0.0000	-0.0000	-0.3063
N= 72015							

This table shows treatment effects for individuals in the top quartile of predicted treatment effects on saving who have a credit card. Variable definitions are provided in Table 2. Monthly Spending and Checking Account Balances are measured over the first month of the intervention. Credit Card Interest is measured as the average of the monthly interest charges during the two billing cycles affected by the intervention or immediately following it. Paid Interest is a binary variable flagging whether an individual was charged credit card interest on any of the two billing cycles affected by the intervention or immediately following it. Ending Statement Balance - Payments is measured for the last billing cycle affected by the intervention. Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest at baseline. Treatment effects are calculated using Poisson regressions with strata fixed effects. We report percentage treatment effects calculated as exponentiated coefficients -1. The change in spending (saving), expressed in MXN, is calculated by multiplying the % TE and the Mean of Spending (Checking Account Balances) in the Control Group. Upper (lower) confidence intervals, expressed in MXN, are calculated as (point estimate +/- 1.96*Standard Error)*Mean of Dep. Var. in Control Group. The upper (lower) confidence interval for the variable Paid Interest is calculated as (point estimate +/- 1.96*Standard Error). Abs. Value stands for absolute value. The number of observations is the same across all columns in the same panel. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

Table VII. Treatment Effects on Deposits, ATM Whithdrawals, Card Spending, and Transfers (Top Quartile of Predicted Treatment Effects on Spending)

	(1)	(2)	(3)	(4)
Dep. Var.	Monthly Deposits	Monthly ATM Withdrawals	Monthly Spending with Cards	Monthly Outgoing Transfers
Panel A: Individuals with a Credit Card				
% TE	0.014	-0.087***	-0.063**	-0.012
	(0.013)	(0.032)	(0.027)	(0.021)
Mean of Dep. Var. in Control Group (MXN)	31,446.39	14,613.65	18,340.23	2,014.60
N= 150177				
Panel B: Individuals with a Credit Card who Paid Interest at Baseline				
% TE	0.008	-0.096**	-0.054*	0.040
	(0.021)	(0.045)	(0.031)	(0.037)
Mean of Dep. Var. in Control Group (MXN)	29,200.91	13,843.29	20,691.33	1,015.48
N= 73946				

This table considers individuals in the top quartile of predicted treatment effects on spending who have a credit card. Spending with Cards is defined as the sum of debit and credit card purchases in stores or online. Outgoing Transfers refers to all outgoing electronic transfers originated from a given account. Incoming transfers are classified as Deposits. All variables are measured over the first month of the intervention. Treatment effects are calculated using Poisson regressions with strata fixed effects. We report percentage treatment effects calculated as exponentiated coefficients -1. Mean of the Dependent Variable is reported in MXN. 1 MXN = 0.107 USD PPP. The number of observations is the same across all columns in the same panel. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

Table VIII. Treatment Effects on Spending, Saving, and Borrowing by Message (Top Quartile of Predicted Treatment Effects on Spending)

Dep.Var.	(1)	(2)	(3)
	Monthly Spending	Checking Account Balance	Monthly Credit Card Interest
Short-term messages			
Msg 2	-0.070***	0.051**	-0.008
Year-end Expenses	(0.023)	(0.024)	(0.027)
Msg 6	-0.041*	0.012	-0.018
Avoid Shortfalls	(0.023)	(0.023)	(0.028)
Msg 7	-0.075***	0.049**	-0.015
Emergency	(0.023)	(0.024)	(0.028)
All Short-term	-0.062***	0.037*	-0.014
Msgs. Pooled	(0.019)	(0.020)	(0.025)
Long-term messages			
Msg 1	-0.050**	0.045*	-0.019
Congratulations	(0.023)	(0.024)	(0.029)
Msg 3	-0.077***	0.055**	0.004
Others your Age	(0.023)	(0.024)	(0.028)
Msg 5	-0.054**	0.034	-0.008
Reach Dreams	(0.023)	(0.023)	(0.028)
All Long-term	-0.060***	0.045**	-0.008
Msgs. Pooled	(0.019)	(0.019)	(0.025)
Mental Accounting			
Msg 4	-0.090***	0.069***	-0.006
Money Box	(0.023)	(0.024)	(0.028)
Differences Across Types of Messages			
Short-term	-0.002	-0.007	-0.006
- Long-term	(0.014)	(0.014)	(0.020)
Short-term	0.028*	-0.032*	-0.008
- Mental Accounting	(0.016)	(0.017)	(0.021)
Long-term	0.030*	-0.024	-0.002
- Mental Accounting	(0.016)	(0.016)	(0.021)

This table shows treatment effects for individuals in the top quartile of predicted treatment effects on spending who have a credit card. Variable definitions are provided in Table 2. Monthly Spending and Checking Account Balances are measured over the first month of the intervention. Credit Card Interest is measured as the average of the monthly interest charges during the two billing cycles affected by the intervention. Treatment effects are calculated using Poisson regressions with strata fixed effects. We report percentage treatment effects calculated as exponentiated coefficients -1. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

Table IX. Heterogeneous Treatment Effects, by Experimental Strata

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
				Panel A: Dep. Var. Monthly Spending					
Any Treatment	-0.005	-0.005	-0.007	-0.008	-0.007	-0.009	-0.008	-0.008	-0.007
	(0.009)	(0.009)	(0.010)	(0.007)	(0.007)	(0.007)	(0.006)	(0.006)	(0.006)
Incremental: Group1	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted
Incremental: Group2	0.005	-0.007	0.003	0.001	-0.002	0.003	0.001	-0.002	-0.003
	(0.015)	(0.014)	(0.015)	(0.011)	(0.010)	(0.011)	(0.013)	(0.013)	(0.012)
Incremental: Group3	-0.002	0.001	-0.006						
	(0.015)	(0.015)	(0.014)						
Incremental: Group4	-0.017	-0.006	-0.003						
	(0.015)	(0.015)	(0.014)						
				Panel B: Dep. Var. Checking Account Balance					
Any Treatment	-0.002	-0.001	0.005	0.004	0.006	0.008	0.005	0.006	0.005
	(0.007)	(0.007)	(0.008)	(0.006)	(0.006)	(0.006)	(0.005)	(0.006)	(0.005)
Incremental: Group1	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted
Incremental: Group2	0.009	0.007	0.003	0.004	-0.001	-0.002	0.001	-0.001	0.001
	(0.011)	(0.011)	(0.011)	(0.010)	(0.010)	(0.009)	(0.012)	(0.011)	(0.012)
Incremental: Group3	0.006	0.009	-0.002						
	(0.011)	(0.012)	(0.010)						
Incremental: Group4	0.014	0.009	0.005						
	(0.011)	(0.011)	(0.011)						
Group Definition	Quartiles of Checking Acct. Balance	Quartiles of Income	Quartiles of Age	Median of Tenure with Bank	Median of ATM Withdrawals	Median of Debit Card Transactions	Is Digital?	Main Bank?	Has Credit Card?

This table presents heterogeneous treatment effects by experimental strata for all individuals in the experiment. In Panel A the dependent variable is Monthly Spending. In Panel B the dependent variable is Checking Account Balance. Both variables are defined in Table 2 and are measured over the first month of the intervention. Treatment effects are estimated in each column with the following specification: $y_i = exp(\alpha_s + \beta_1 * Treatment_i + \beta_2 * Group_i + \beta_3 * Treatment_i * Group_i + error_i)$ where α_s represents strata fixed effects and $Group_i$ is a vector of dummy variables taking the value of 1 when individual i belongs to the corresponding group-category. For the proportional treatment effect of any treatment on the omitted category we report $exp(\beta_1) - 1$. For incremental effects, we report exponentiated coefficients of the form $exp(\beta_1 + \beta_3) - exp(\beta_1)$, interpreted as the difference (in percentage points) of the proportional treatment effects of each group relative to the omitted category. In each column, *Group* is defined over a different variable: quartiles of checking account balances, quartiles of income, quartiles of age, median of tenure with Banorte, median of ATM withdrawals, median of debit card transactions, a binary variable indicating if a given individual is digital (i.e., no more than 20% of checking account outflows are ATM withdrawals), and a binary variable indicating if a given user has a credit card, respectively. In all cases, we consider all 3,054,503 observations at the user level. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

Table X. Comparison of Treatment Effects for Users in Strata Blocks with the Highest Observed Treatment Effect on Spending and for Users with the Highest Predicted Treatment Effect on Spending

	Top Quartile of Individuals' Observed Treatment Effects			Top Quartile of Individuals' Predicted Treatment Effects		
	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.	Monthly Spending	Checking Account Balance	Monthly Credit Card Interest	Monthly Spending	Checking Account Balance	Monthly Credit Card Interest
Panel A: Individuals with a Credit Card						
% TE	-0.3811***	0.3787***	-0.3462***	-0.0722***	0.0493***	-0.0101
	(0.0127)	(0.0226)	(0.0245)	(0.0133)	(0.0128)	(0.0221)
Mean of Dep. Var. in Control Group (MXN)	36,707.18	38,859.97	242.92	34,968.92	39,174.23	210.16
N	149,647	149,647	149,647	150,177	150,177	150,177
Panel B: Individuals with a Credit Card Who Paid Interest at Baseline						
% TE	-0.2721***	0.1892***	-0.1910***	-0.0688***	0.0555***	-0.0099
	(0.0195)	(0.0289)	(0.0242)	(0.0170)	(0.0173)	(0.0229)
Mean of Dep. Var. in Control Group (MXN)	39,400.25	30,647.70	414.98	35,550.10	34,913.86	384.09
N	68,751	68,751	68,751	73,946	73,946	73,946

This table shows treatment effects for individuals in the top quartile of either observed or predicted treatment effects on spending and who have a credit card. For Columns (1) to (3), we split the sample into 6,104 mutually exclusive groups (blocks) defined by the interaction of all experimental strata. For each group we calculate treatment effects on spending with a linear model. We assign to each observation in the group the treatment effect of its group. We then split the sample into quartiles based on the treatment effect assigned to each observation. For columns (4) to (6) we calculate quartiles based on predicted treatment effects on spending obtained with the causal forest for spending. In both cases, the top quartile corresponds to the most negative effect, and for individuals in the top quartile who have a credit card, we calculate treatment effects using Poisson regressions with strata fixed effects. We report percentage treatment effects calculated as exponentiated coefficients -1. Variable definitions are provided in Table 2. Monthly Spending and Checking Account Balances are measured over the first month of the intervention. Credit Card Interest is measured as the average of the monthly interest charges during the two billing cycles affected by the intervention. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

Internet Appendix for "Does Saving Cause Borrowing? Implications for the Co-Holding Puzzle"[25]

Paolina C. Medina and Michaela Pagel

This Internet Appendix provides additional tables and figures supporting the study.

IA.I. Randomization Check

Table IA.I shows that the randomization process was successful in balancing observable characteristics across treatment and control groups.

Table IA.I. Covariate Balance

	Control	Treatment	p-value of Difference
Age (Years)	45	45	0.1604
Monthly Income	13,506	13,497	0.7030
Tenure (Months)	80	82	0.3950
Monthly Spending	17,771	17,963	0.5186
Checking Account Balance	20,079	19,996	0.3629
Monthly Credit Card Interest	32	33	0.4065
Credit Card Balance	3,985	4,009	0.7742
Statement Balance - Payments	599	604	0.8153
Credit Card Limit	18,012	18,103	0.4571
N	357,567	2,696,936	

This table presents a covariate balance test, reporting the average value of each dependent variable for treatment and control groups, along with the p-value of regressing the corresponding outcome on a treatment indicator with strata fixed effects and robust standard errors. The p-value of an F-test from regressing the treatment indicator on all of the covariates with strata fixed effects is 0.4608. The statistics are calculated with monthly information at the individual level covering the 6 months previous to the intervention. Variable definitions are provided in Table 2.

[25]Citation format: Medina, Paolina C., and Michaela Pagel, Internet Appendix for "Does Saving Cause Borrowing? Implications for the Co-Holding Puzzle," *Journal of Finance* [DOI String]. Please note: Wiley-Blackwell is not responsible for the content or functionality of any additional information provided by the authors. Any queries (other than missing material) should be directed to the authors of the article.

IA.II. Heterogeneity Analysis

IA.IIA. Methodology

Causal forests are based on causal trees, and their relation is analogous to the relation between widely known random forests and regression trees. Regression trees predict an individual outcome Y_i using the mean Y of observations that share similar covariates, X. To define what counts as similar, regression trees partition the covariate space into disjoint groups of observations called 'leaves.' Within each leaf, all observations share values (or belong to the same value interval) of certain Xs. A tree starts with a training sample that is treated first as a single group and then recursively partitioned. For each value $X_j = x$, the algorithm forms candidate splits, placing all observations with $X_j \leq x$ in a left leaf and all observations with $X_j > x$ in a right leaf. The split is implemented if it minimizes a certain loss criterion, such as mean squared error $(\sum_{i=1}^{n}(\hat{y}_i - y_i)^2)$. This criterion is evaluated in the sample; that is, the same observations used to define where to split are also used to calculate the mean value of the outcome in each leaf. The algorithm then repeats the process for each of the two new leaves and so on until it reaches a stopping rule. Using the final set of leaves, the tree provides out-of-sample predictions by figuring out in which terminal leaf a certain observation falls based on its covariate values and assigning a predicted value equal to the average value of all observations in that leaf in the training sample.

Random forests are an ensemble of n trees in which n random subsamples of the data are taken, and each subsample is used to train a tree. Predictions for each observation in a test sample (which could be the full original dataset) are defined as the average across n predictions obtained by pushing that one observation down each of the n trees.

In contrast to regular random forests that predict individual outcomes Y_i, causal forests want to predict conditional treatment effects ($E[Y_1 - Y_0|X = x]$ in a potential outcomes framework) to measure how causal effects vary for different subpopulations. Standard loss criteria such as goodness-of-fit measures are not available because we do not observe the treatment effect $Y_1 - Y_0$ for any one individual. Athey and Imbens (2016) show that maximizing the expected mean squared error of predicted treatment effects instead of the infeasible mean squared error itself is basically equivalent to maximizing the variance of treatment effects across leaves. Thus, this defines a new criterion for sample splitting that is specifically designed to identify treatment effect heterogeneity. They also show that to reduce overfitting bias, the training sample should be further split into a splitting and an estimation sample so that the observations used to choose where to create new leaves are different from the ones used to calculate treatment effects within each leaf. Causal forests are different from off-the-shelf machine learning methods in three ways.

First, in addition to dividing the data into training and test samples, causal forests divide the training data further into two sub-samples: a splitting sample and an estimation sample. The splitting sample is used to grow trees (2,000 in our case), and the estimation sample is used to estimate the treatment effects. This honesty is crucial for causal forests to attain consistent estimation of treatment effects, and similar strategies are implemented in other recently developed methods for causal inference with machine learning (Chernozhukov et al., 2018).

Second, causal forests use a splitting rule that tackles treatment effect heterogeneity directly. This is, each tree splits into two children nodes where heterogeneity in treatment effects is maximized. Thus, causal forests are tailored to find sub-populations with different treatment effects.

Finally, causal forests calculate treatment effects, ensuring that the treatment indicator is orthogonal to all covariates for all observations. The algorithm computes estimates of propensity scores and outcomes for treatment and control groups by training separate regression forests. Then, on residual treatments and outcomes. To calculate the treatment effect on a sub-population of interest, the algorithm plugs the individual predictions of the causal forest into an Augmented Inverse Probability Weighting Estimate (AIPW) that combines models of outcome regressions with models of treatment propensity to estimate causal effects.[26]

We use the generalized random forest (grf) package in R, to estimate our causal forests. Hyperparameters are optimally selected to maximize predictive power. This package allows for the estimation of causal forests but also allows for the estimation of other forest-based methods for causal inference. To do so efficiently, this package involves an approximate gradient-based loss criterion (instead of the exact loss criterion described above to build intuition) and aggregates the results of the n trees with one single weighted estimation of the treatment effect instead of averaging n estimations of treatment effects. The mechanics of the algorithm are described as follows:

1. The first step is to compute estimates of propensity scores for the treatment and marginal outcomes conditional on covariates by training separate regression forests and performing predictions (fitted values) for each observation. These predictions are used to calculate residuals, which will be referred to as orthogonalized outcomes and orthogonalized treatment status.

2. For each tree, a random subsample with 50% of the database is drawn (training sample).

3. The training sample is further split into a splitting subsample and an estimation subsample (50-50 by default).

4. A single initial root node is created for the splitting sample, and child nodes are split recursively to form a tree. Each node is split using the following algorithm:

   (a) A random subset of variables are selected as candidates to split on.

   (b) For each of these variables, we look at all of their possible values and consider splitting into two child nodes based on a measure of goodness of split, determined to maximize the heterogeneity in treatment effect estimates across nodes.

   (c) All observations with values for the split variable that are less than or equal to the split value are placed in a new left child and all examples with values greater than the split value are placed in a right child node.

5. The estimation sample is used to populate the leaf nodes of the tree. Each observation is 'pushed down' the tree and assigned to the leaf in which it falls.

6. Steps 2 to 5 are repeated 2,000 times; that is, we estimate 2,000 trees.

[26] This estimator is locally efficient and is known as a "doubly robust estimator" since it is consistent whenever the model of treatment propensity *or* the model of expected outcomes are correctly specified.

7. Treatment effects are predicted for each observation on a test dataset (potentially the full dataset) as follows:

   (a) Each test observation is pushed down each tree to determine what leaf it falls in. Given this information, a list with neighboring observations in each tree leaf is created (the neighbors come from the estimation sample of each tree). Each neighbor observation is weighted by how many times it fell in the same leaf as the test observation.

   (b) Treatment effects are calculated using orthogonalized outcomes and treatment status of the neighbor observations.

8. In addition to personalized treatment effects, the package allows for estimation of treatment effects across all observations in a dataset, or arbitrary subsamples of it. This is done with an AIPW estimator that ensures balance across all covariates in the group, using the treatment propensities estimated in step 1.

IA.IIB. Calibration Tests

To test for treatment effect heterogeneity, we follow Burke et al. (2020) and perform the calibration test of Chernozhukov et al. (2018). Using the forest predictions, this test computes the best linear fit of a target estimand on two regressors: the average of the individual forest predictions (which we refer to as the Mean Forest Prediction) and the difference between each individual prediction and the Mean Forest Prediction (which we refer to as the Differential Forest Prediction). We proceed to define the components of this linear regression.

The target estimand is the residualized outcome of interest. Residualized outcomes refer to observed outcomes (e.g., spending, checking account balances, or credit card interest) minus predicted outcomes, where individual-level predictions of outcomes come from separate models (e.g., a random forest) that use all available covariates. Importantly, the model for predicted outcomes does not use the treatment indicator as a predictor. If there is a positive treatment effect, observations in the treatment group will have a higher residualized outcome. Robinson (1988) shows that we can estimate treatment effects by regressing residualized outcomes on a treatment dummy multiplied by residualized treatment propensities. In the calibration test, both regressors (Mean Forest Prediction and Differential Forest Prediction) are multiplied by residualized treatment propensities. For each observation, residualized treatment propensities are given by the difference between the treatment indicator and the predicted probability of being in the treatment group, where the predicted probability comes from a separate model (e.g., a probability forest) based on baseline observable characteristics (other than the treatment dummy). Multiplying by residualized propensities is basically a tool to control for residual covariate imbalance in doubly robust estimators. The orthogonalization of individual-level outcomes and treatment propensities is one of the building blocks of the causal forest algorithm of Athey et al. (2019).

The coefficient for Mean Forest Prediction tests whether the model accurately predicts the outcome. A substantial deviation from one would imply that the model is not well calibrated. The coefficient for Differential Forest Prediction acts as an omnibus test for the presence of heterogeneity. If the predicted treatment effects were purely spurious, this coefficient would be indistinguishable from zero. A significantly positive coefficient implies that the predicted treatment effects are positively correlated with actual treatment effects, and we can use them as a score that is positively correlated with treatment effects. The standard practice is to use a one-sided hypothesis test, where the null hypothesis is a slope coefficient (Differential Forest Prediction) less than or equal to zero (Tibshirani et al., 2020)

We present the results in Table IA.II. Columns (1) and (2) use spending and saving, respectively, as the outcome variables. We conduct the tests on all 3.1 million observations, as well as separately on observations from individuals who have credit cards and those who paid credit card interest at baseline. Across all sub-groups, we find evidence consistent with the presence of treatment effect heterogeneity in the spending and saving responses to the saving nudge.

In contrast, the causal forest model trained on borrowing outcomes does not capture significant heterogeneity in treatment effects. This is shown in Column (3), where the model was trained using all 161 explanatory variables. In Column (4), the model was trained using only the covariates with variable importance greater than or equal to 1% from the model in Column (3). In Column (5), the model was trained using only the variables used in the model for checking account balances in Column (2). In Column (6), the model was trained using only the variables used in the model for spending in Column (1). The coefficients for the Differential Forest Prediction are not significantly positive in any of these specifications. Therefore, based on the variables considered in the analysis, we do not find evidence of treatment effect heterogeneity in borrowing.[27]

We note that the failure to reject the hypothesis that the coefficient for Differential Forest Prediction is less than or equal to zero does not guarantee the absence of treatment effect heterogeneity. It simply implies that the specific model we used failed to capture it. Nevertheless, we interpret these results as suggestive evidence for the absence of treatment effect heterogeneity. While the same model specification successfully predicted treatment effect heterogeneity for other outcomes, various attempts (shown in Columns 3-5 above) using borrowing as the target variable failed to do so despite ample statistical power, a rich set of covariates, and a focus on relevant sub-populations who have credit cards and who paid credit card interest at baseline.

[27]Note that the p-values for a one-sided test for $X \leq 0$ are lower than the p-values for the two-sided test of $X \neq 0$. Therefore, based on the results in the table, we find no evidence for heterogeneity.

Table IA.II. Calibration Tests for Evaluation of Causal Forests

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.	Monthly Spending	Checking Account Balance	Monthly Credit Card Interest	Monthly Credit Card Interest	Monthly Credit Card Interest	Monthly Credit Card Interest
Panel A: All Individuals (N = 3,054,503)						
Mean Forest Prediction	1.416***	0.992**				
	(0.449)	(0.504)				
Differential Forest Prediction	1.083***	0.922***				
	(0.202)	(0.167)				
Panel B: Individuals with a Credit Card (N= 362,223)						
Mean Forest Prediction	1.029*	1.106*	2.180***	-0.515	1.225*	1.845***
	(0.589)	(0.648)	(0.626)	(0.761)	(0.682)	(0.629)
Differential Forest Prediction	1.245***	0.951***	0.243	-0.136	-0.463	0.014
	(0.356)	(0.207)	(0.260)	(0.359)	(0.329)	(0.237)
Panel C: Individuals with a Credit Card Who Paid Interest at Baseline (N= 152,016)						
Mean Forest Prediction	1.328*	1.459*	1.190	1.077	1.939**	0.840
	(0.739)	(0.868)	(0.885)	(0.756)	(0.794)	(0.851)
Differential Forest Prediction	0.917**	1.207***	-0.067	-0.432	0.241	0.179
	(0.416)	(0.390)	(0.505)	(0.549)	(0.674)	(0.613)

This table presents the results of calibration tests as in Athey and Wager (2019), based on Chernozhukov et al. (2018). The tests fit a linear model for the residualized dependent variables using the mean forest prediction and the deviation between individual-level predictions and the mean prediction as the sole two regressors (both multiplied by residualized treatment status). Residualized variables correspond to the difference between fitted values predicted from covariates and observed values. A coefficient of 1 for Mean Forest Prediction suggests that the forest correctly predicts treatment effects, on average. A positive coefficient for Differential Forest Prediction acts as an omnibus test for the presence of heterogeneity: If the coefficient is significantly greater than 0, then we can reject the null of no heterogeneity. The p-values are one-sided for the null of a coefficient smaller than or equal to 0. Variable definitions are provided in Table 2. Monthly Spending and Checking Account Balances are measured over the first month of the intervention. Credit Card Interest is measured as the average of the monthly interest charges during the two billing cycles affected by the intervention. Columns (1) and (2) train the model with the variables in Figure IA.1 panels (a) and (b), respectively. Column (3) uses all 161 available covariates. Column (4) uses variables with importance larger than 1 percent according to the model in Column (3). Column (5) uses the same variables as in Column (2). Column (6) uses the same variables as in Column (1). Panel A considers all individuals. Panel B considers all individuals who have a credit card. Panel C considers only individuals who have a credit card and incurred interest at baseline. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

IA.IIC. Overlap Between the Treatment Effect Predictions of the Causal Forests for Spending and Saving

Table IA.III presents the overlap between the treatment effect predictions of the causal forests for spending and saving (again, calculating quartiles with a cross-fitted ranking over five folds).

Table IA.III. Distribution of Observations According to their Predicted Treatment Effects on Spending and Saving

Columns: Quartiles of Predicted Treatment Effects on Spending
Rows: Quartiles of Predicted Treatment Effects on Saving

(a) All Individuals

	1	2	3	4
1	0.0980	0.0757	0.0604	0.0158
2	0.0646	0.0896	0.0587	0.0371
3	0.0675	0.0493	0.0903	0.0430
4	0.0191	0.0356	0.0412	0.1549

(b) Individuals with a Credit Card

	1	2	3	4
1	0.1016	0.0764	0.0564	0.0156
2	0.0690	0.0947	0.0596	0.0273
3	0.0583	0.0585	0.0939	0.0398
4	0.0213	0.0202	0.0409	0.1678

(c) Individuals with a Credit Card Who Paid Interest at Baseline

	1	2	3	4
1	0.0968	0.0816	0.0584	0.0137
2	0.0859	0.0862	0.0548	0.0235
3	0.0523	0.0655	0.0925	0.0392
4	0.0151	0.0176	0.0436	0.1738

This table shows the distribution of observations according to the predicted treatment effects on spending and saving. Panel (a) considers all 3,054,503 individuals. Panel (b) considers 362,223 individuals who have a credit card. Panel (c) considers 152,016 individuals who have a credit card and who paid credit card interest in at least one of the 6 months previous to the intervention. In each panel the sum of the fractions presented in each cell of the corresponding table equals one. The sorting into quartiles is based on cross-fitted rankings over five folds. For spending the top quartile corresponds to the most negative effect. For saving the top quartile corresponds to the most positive effect.

IA.IID. Comparison of Individuals in the Top and Bottom Quartiles of the Distribution of Predicted Treatment Effects

To understand the differences between individuals who respond to nudges and those who don't, we can compare the descriptive statistics of individuals in the top and bottom quartiles of the distribution of predicted treatment effects. Note that, by design, we would not expect these to be balanced. Instead, significant differences imply that those who respond to the treatment are significantly different than those who do not. The comparison can be found in Table IA.IV.

Table IA.IV. Differences Between Top and Bottom Quartiles of the Distribution of Predicted Treatment Effects

	Spending		Saving	
	Bottom Quartile	Top Quartile	Bottom Quartile	Top Quartile
Monthly Spending	18,854	19,332	20,909	21,881
Checking Account Balance	15,323	24,667	15,829	22,835
Monthly Income	14,441	16,736	14,049	17,338
Age (Years)	45	45	46	45
Tenure With Banorte (Months)	69	89	74	89
Credit Card Balance	986	4,184	864	3,161
Credit Card Limit	8,193	18,834	9,176	16,166

This table presents means of each variable for individuals in the top and bottom quartiles of the distribution of predicted treatment effects on spending or saving, respectively. The statistics are calculated with monthly information at the individual level covering the 6 months previous to the intervention. Variable definitions are provided in Table 2.

IA.IIE. Variable Importance

Figure IA.1 illustrates the variable importance for the causal forests with either spending or savings as the outcome variables. Variable importance indicates how often a given pre-treatment variable was used to select splits across the multiple trees of the causal forests.

Figure IA.1. Variable Importance: Causal Forests for Spending and Saving

(a) Spending (b) Savings

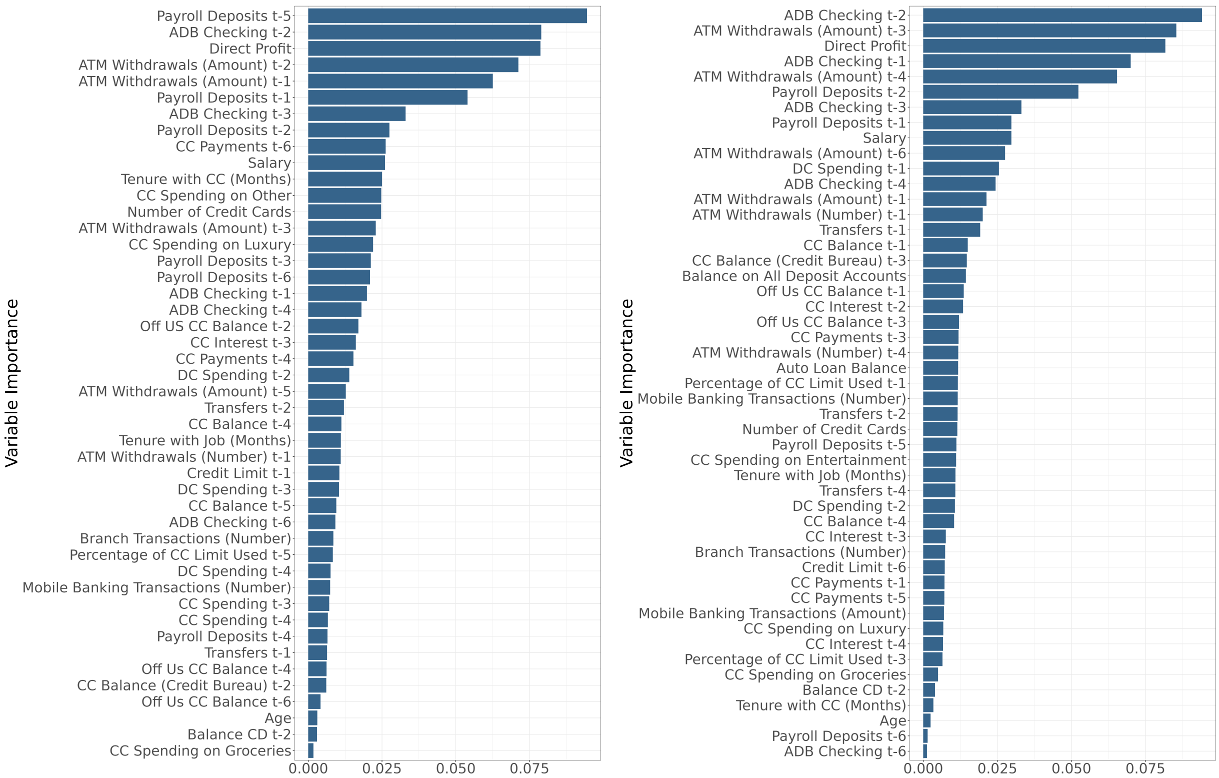

This graph shows the variable importance of the variables used in the estimation of the final causal forests for both spending and saving. Spending is defined as the sum of ATM withdrawals, credit and debit card spending, and outgoing transfers over the first month of the intervention. Checking Account Balances correspond to average daily balances during the first month of the intervention. Following Athey and Wager (2019), we first estimate pilot causal forests using all available pre-treatment variables (161 variables) and then re-estimate final models using only those with variable importance larger than 1%. The resulting variables are listed in the vertical axis of the graphs. Variable importance indicates how often a variable was used to select splits in the multiple trees of the causal forest. ADB refers to average daily balances. Off Us, CC Balance refers to credit card balances reported to the credit bureau on credit cards outside of Banorte. DC refers to debit card, CD refers to certificate of deposit, t refers to the month of the treatment; in turn, t-1 to 6 refers to the 1 to 6 months before the treatment period. All variables are monthly.

IA.IIF. Heterogeneity by Checking Account Balances for Individuals Who Paid Credit Card Interest at Baseline

Table IA.V considers individuals who had a credit card and paid interest at baseline and compares the treatment effects of the top quartile of checking account balances to the treatment effect in lower quartiles. Consistent with Table IA.IV, there is suggestive evidence that the treatment effect on savings and spending is more prominent for individuals who were already saving and carried higher checking account balances at baseline.

Table IA.V. Treatment Effects by Quartiles of Checking Account Balances, for Individuals with a Credit Card Who Paid Interest at Baseline

	(1)	(2)	(3)	(4)	(5)
Dep. Var.	Monthly Spending	Checking Account Balance	Monthly Credit Card Interest During Treat.	Paid Interest During Treat. {0,1}	Ending Statement Balance - Payments
Any Treatment	-0.005	0.004	0.001	-0.001	-0.002
	(0.010)	(0.009)	(0.011)	(0.006)	(0.012)
Incremental Effect: Top Quartile	-0.044*	0.041*	-0.015	-0.001	-0.003
	(0.026)	(0.024)	(0.031)	(0.017)	(0.029)
Observations	152,016	152,016	152,016	152,016	152,016
Mean of Dep. Var. in Control Group	33,982.15	32,929.96	483.77	0.81	10,317.99

This table presents treatment effects for individuals who have a credit card and incurred interest at baseline. Treatment effects are estimated with the following equation: $y_i = exp(\alpha_s + \beta_1 * Any\ Treatment_i + \beta_2 * Top\ Quartile_i + \beta_3 * Any\ Treatment_i * Top\ Quartile_i + error_i)$ where $alpha_s$ represents strata fixed effects and $Top\ Quartile$ is a dummy taking the value of 1 when individual i is in the top quartile of the distribution of Checking Account Balances. For the proportional treatment effect of any treatment on the omitted category we report $exp(\beta_1) - 1$. For incremental effects we report $exp(\beta_1 + \beta_3) - exp(\beta_1)$, interpreted as the difference (in percentage points) of the proportional treatment effects for the Top Quartile, relative to the omitted category. Variable definitions are provided in Table 2. Monthly Spending and Checking Account Balances are measured over the first month of the intervention. Credit Card Interest is measured as the average of the monthly interest charges during the two billing cycles affected by the intervention or immediately following it. Paid Interest is a binary variable flagging whether an individual was charged credit card interest on any of the two billing cycles affected by the intervention or immediately following it. Ending Statement Balance - Payments is measured for the last billing cycle affected by the intervention. Any Treatment is a binary variable that takes the value of one if a given individual is assigned to receive any of the treatment messages. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

IA.III. Alternative Outcome Variables, Subsamples and Specifications

IA.IIIA. Alternative Outcome Variables

Table IA.VI shows results for a set of alternative outcome variables related to credit card borrowing. Table IA.VII shows results for the set of alternative outcome variables for individuals in the top quartile of predicted treatment effects on saving as opposed to spending.

IA.IIIB. Individuals with Low Credit Card Utilization

Tables IA.VIII and IA.IX show the main results for individuals who have a credit card and who have a credit line utilization lower than the median.

Tables IA.X and IA.XI show our main specification for individuals in the top quartile of predicted treatment effects on saving, for whom Banorte is their main bank (i.e., they receive their payroll at Banorte, have a credit card with Banorte, and, according to credit bureau records, do not have credits outside of Banorte).

Table IA.VI. Treatment Effects on Card Balances and Payments (Top Quartile of Predicted Treatment Effects on Spending)

Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Interest Bearing Balance During Treat.	Interest Bearing Balance After Treat.	Credit Card Balance During Treat. (Banorte)	Credit Card Balance After Treat. (Banorte)	Credit Card Balance During Treat. (Credit Bureau)	Credit Card Balance After Treat. (Credit Bureau)	Monthly Credit Card Payments During Treat.	Monthly Credit Card Payments After Treat.
Panel A: Individuals with a Credit Card								
% TE	-0.0102	-0.0117	-0.0105	-0.0116	0.0091	0.0109	-0.0159	-0.0176
	(0.0242)	(0.0258)	(0.0242)	(0.0243)	(0.0208)	(0.0209)	(0.0245)	(0.0230)
Mean of Dep. Var. in Control Group (MXN)	6,557.54	6,777.50	10,766.51	11,049.37	24,191.77	23,543.32	6,013.33	6,194.74
Upper Confidence Interval (MXN)	243.79	263.24	397.63	398.44	1,206.53	1,220.49	193.15	170.05
Upper Confidence Interval Divided by Abs. Value of Change in Spending	0.0966	0.1043	0.1575	0.1578	0.4780	0.4835	0.0765	0.0674
Lower Confidence Interval (MXN)	-377.16	-422.51	-623.73	-655.78	-765.97	-708.37	-384.37	-387.88
Lower Confidence Interval Divided by Abs. Value of Change in Spending	-0.1494	-0.1674	-0.2471	-0.2598	-0.3034	-0.2806	-0.1523	-0.1537
N= 150177								
Panel B: Individuals with a Credit Card Who Paid Interest at Baseline								
% TE	-0.0098	-0.0098	-0.0099	-0.0097	0.0062	-0.0056	-0.0148	-0.0288
	(0.0236)	(0.0226)	(0.0234)	(0.0233)	(0.0208)	(0.0209)	(0.0542)	(0.0536)
Mean of Dep. Var. in Control Group (MXN)	11,828.57	12,095.85	12,835.18	13,181.64	25,162.85	24,050.38	3,090.04	3,020.13
Upper Confidence Interval (MXN)	431.47	418.76	462.81	472.54	1,181.35	849.82	282.64	230.27
Upper Confidence Interval Divided by Abs. Value of Change in Spending	0.1765	0.1713	0.1893	0.1933	0.4833	0.3477	0.1156	0.0942
Lower Confidence Interval (MXN)	-664.39	-654.94	-715.80	-729.35	-870.33	-1,118.70	-373.83	-404.42
Lower Confidence Interval Divided by Abs. Value of Change in Spending	-0.2718	-0.2680	-0.2929	-0.2984	-0.3561	-0.4577	-0.1529	-0.1655
N= 73946								

This table shows treatment effects for individuals in the top quartile of predicted treatment effects on spending who have a credit card. Interest Bearing Balances correspond to average credit card daily balances over a billing cycle multiplied by a binary variable flagging if an individual was charged credit card interest on that billing cycle. Credit Card Balances at Banorte are average daily balances over a billing cycle. Credit Card Balances from the Credit Bureau are balances at the end of the month. Credit Card Payments correspond to the sum of all payments received against a credit card bill over a billing cycle. All variables are measured on a monthly basis and averaged over the two months during or after the treatment. Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest at baseline. Treatment effects are calculated using Poisson regressions with strata fixed effects. We report percentage treatment effects calculated as exponentiated coefficients -1. The change in spending (saving), expressed in MXN, is calculated by multiplying the % TE and the Mean of Spending (Checking Account Balances) in the Control Group. Upper (lower) confidence intervals, expressed in MXN, are calculated as (point estimate +/- 1.96*Standard Error)*Mean of Dep. Var. in Control Group. Abs. Value stands for absolute value. The number of observations is the same across all columns in the same panel. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

Table IA.VII. Treatment Effects on Card Balances and Payments (Top Quartile of Predicted Treatment Effects on Saving)

Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Interest Bearing Balance During Treat.	Interest Bearing Balance After Treat.	Credit Card Balance During Treat. (Banorte)	Credit Card Balance After Treat. (Banorte)	Credit Card Balance During Treat. (Credit Bureau)	Credit Card Balance After Treat. (Credit Bureau)	Monthly Credit Card Payments During Treat.	Monthly Credit Card Payments After Treat.
Panel A: Individuals with a Credit Card								
% TE	-0.0115	-0.0091	-0.0117	-0.0093	-0.0008	0.0012	-0.0200	-0.0174
	(0.0205)	(0.0213)	(0.0212)	(0.0213)	(0.0101)	(0.0103)	(0.0251)	(0.0247)
Mean of Dep. Var. in Control Group (MXN)	6,789.92	7,077.65	9,923.98	10,621.77	23,382.91	25,806.17	5,532.97	6,069.58
Upper Confidence Interval (MXN)	195.36	231.65	296.25	344.66	444.18	551.94	161.65	188.47
Upper Confidence Interval Divided by Abs. Value of Change in Saving	0.0937	0.1112	0.1421	0.1654	0.2131	0.2648	0.0776	0.0904
Lower Confidence Interval (MXN)	-351.25	-360.20	-528.47	-542.22	-481.59	-490.01	-382.97	-399.69
Lower Confidence Interval Divided by Abs. Value of Change in Saving	-0.1685	-0.1728	-0.2536	-0.2602	-0.2311	-0.2351	-0.1838	-0.1918
N= 151834								
Panel B: Individuals with a Credit Card Who Paid Interest at Baseline								
% TE	-0.0102	-0.0106	-0.0109	-0.0107	-0.0011	0.0017	-0.0138	-0.0141
	(0.0203)	(0.0205)	(0.0204)	(0.0204)	(0.0146)	(0.0147)	(0.0463)	(0.0461)
Mean of Dep. Var. in Control Group (MXN)	12,407.61	12,440.27	12,808.04	12,996.45	25,332.03	24,984.51	3,247.06	3,387.37
Upper Confidence Interval (MXN)	366.59	367.98	372.51	380.64	697.04	762.33	250.27	258.18
Upper Confidence Interval Divided by Abs. Value of Change in Saving	0.1682	0.1688	0.1709	0.1746	0.3198	0.3497	0.1148	0.1184
Lower Confidence Interval (MXN)	-618.56	-632.09	-650.63	-659.40	-752.77	-677.38	-339.58	-353.59
Lower Confidence Interval Divided by Abs. Value of Change in Saving	-0.2838	-0.2900	-0.2985	-0.3025	-0.3453	-0.3107	-0.1558	-0.1622
N= 72015								

This table shows treatment effects for individuals in the top quartile of predicted treatment effects on saving who have a credit card. Interest Bearing Balances correspond to average credit card daily balances over a billing cycle multiplied by a binary variable flagging if an individual was charged credit card interest on that billing cycle. Credit Card Balances at Banorte are average daily balances over a billing cycle. Credit Card Balances from the Credit Bureau are balances at the end of the month. Credit Card Payments refers to the sum of all payments received against a credit card bill over a billing cycle. All variables are measured on a monthly basis and averaged over the two months during or after the treatment. Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest at baseline. Treatment effects are calculated using Poisson regressions with strata fixed effects. We report percentage treatment effects calculated as exponentiated coefficients -1. The change in spending (saving), expressed in MXN, is calculated by multiplying the % TE and the Mean of Spending (Checking Account Balances) in the Control Group. Upper (lower) confidence intervals, expressed in MXN, are calculated as (point estimate +/- 1.96*Standard Error)*Mean of Dep. Var. in Control Group. Abs. Value stands for absolute value. The number of observations is the same across all columns in the same panel. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

Table IA.VIII. Treatment Effects on Spending, Saving, and Borrowing (Top Quartile of Predicted Treatment Effects on Spending — Individuals Below the Median Credit Line Utilization)

Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Monthly Spending	Checking Account Balance	Monthly Credit Card Interest During Treat. (Banorte)	Monthly Credit Card Interest After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ending Statement Balance - Payments
Panel A: Individuals with Credit Line Utilization Lower Than the Median							
% TE	-0.0711***	0.0619***	-0.0099	-0.0101	0.0029*	0.0024	0.0126
	(0.0200)	(0.0207)	(0.0317)	(0.0319)	(0.0017)	(0.0017)	(0.0478)
Mean of Dep. Var. in Control Group (MXN)	34,868.50	34,219.62	124.70	122.17	0.25	0.26	2,819.08
Change in Spending or Saving (MXN)	-2,478.57	2,117.85					
Upper Confidence Interval (MXN)			6.52	6.42	0.00	0.00	299.30
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0031	0.0030	0.0000	0.0000	0.1413
Lower Confidence Interval (MXN)			-8.99	-8.88	-0.00	-0.00	-228.52
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0042	-0.0042	-0.0000	-0.0000	-0.1079
N= 60289							

This table shows treatment effects for individuals in the top quartile of predicted treatment effects on spending who have a credit card and who are below the median of credit line utilization (0.2528). Credit line utilization is defined as the ratio of balances to credit line. Variable definitions are provided in Table 2 and the timing of measurement is described in Table 5. Treatment effects are calculated using Poisson regressions with strata fixed effects. We report percentage treatment effects calculated as exponentiated coefficients -1. The change in spending (saving), expressed in MXN, is calculated by multiplying the % TE and the Mean of Spending (Checking Account Balances) in the Control Group. Upper (lower) confidence intervals, expressed in MXN, are calculated as (point estimate +/- 1.96*Standard Error)*Mean of Dep. Var. in Control Group. The upper (lower) confidence interval for the variable Paid Interest is calculated as (point estimate +/- 1.96*Standard Error). Abs. Value stands for absolute value. The number of observations is the same across all columns in the same panel. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

Table IA.IX. Treatment Effects on Spending, Saving, and Borrowing (Top Quartile of Predicted Treatment Effects on Saving — Individuals Below the Median Credit Line Utilization)

Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Monthly Spending	Checking Account Balance	Monthly Credit Card Interest During Treat. (Banorte)	Monthly Credit Card Interest After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ending Statement Balance - Payments
Panel A: Individuals with Credit Line Utilization Lower Than the Median							
% TE	-0.0597***	0.0683***	-0.0100	-0.0102	-0.0011	-0.0012	-0.0017
	(0.0211)	(0.0221)	(0.0281)	(0.0281)	(0.0017)	(0.0017)	(0.0392)
Mean of Dep. Var. in Control Group (MXN)	34,204.65	33,955.43	152.53	150.36	0.27	0.28	3,226.36
Change in Spending or Saving (MXN)	-2,041.54	2,317.79					
Upper Confidence Interval (MXN)			6.87	6.75	0.00	0.00	242.18
Upper Confidence Interval Divided by Abs. Value of Change in Saving			0.0030	0.0029	0.0000	0.0000	0.1045
Lower Confidence Interval (MXN)			-9.91	-9.81	-0.00	-0.00	-253.47
Lower Confidence Interval Divided by Abs. Value of Change in Saving			-0.0043	-0.0042	-0.0000	-0.0000	-0.1094
N= 62492							

This table shows treatment effects for individuals in the top quartile of predicted treatment effects on saving who have a credit card and who are below the median of credit line utilization (0.2463). Credit line utilization is defined as the ratio of balances to credit line. Variable definitions are provided in Table 2 and the timing of measurement is described in Table 5. Treatment effects are calculated using Poisson regressions with strata fixed effects. We report percentage treatment effects calculated as exponentiated coefficients -1. The change in spending (saving), expressed in MXN, is calculated by multiplying the % TE and the Mean of Spending (Checking Account Balances) in the Control Group. Upper (lower) confidence intervals, expressed in MXN, are calculated as (point estimate +/- 1.96*Standard Error)*Mean of Dep. Var. in Control Group. The upper (lower) confidence interval for the variable Paid Interest is calculated as (point estimate +/- 1.96*Standard Error). Abs. Value stands for absolute value. The number of observations is the same across all columns in the same panel. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

Table IA.X. Treatment Effects on Spending, Saving, and Borrowing (Top Quartile of Predicted Treatment Effects on Spending — Main Bank)

Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Monthly Spending	Checking Account Balance	Monthly Credit Card Interest During Treat. (Banorte)	Monthly Credit Card Interest After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ending Statement Balance - Payments
Panel A: Individuals with a Credit Card							
% TE	-0.0864***	0.0573***	-0.0115	-0.0116	-0.0014	-0.0009	-0.0046
	(0.0173)	(0.0141)	(0.0251)	(0.0250)	(0.0114)	(0.0114)	(0.0264)
Mean of Dep. Var. in Control Group (MXN)	30,896.94	36,568.03	231.67	229.87	0.45	0.44	6,444.18
Change in Spending or Saving (MXN)	-2,668.58	2,094.51					
Upper Confidence Interval (MXN)			8.74	8.60	0.01	0.01	303.80
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0033	0.0032	0.0000	0.0000	0.1138
Lower Confidence Interval (MXN)			-14.06	-13.94	-0.01	-0.01	-363.09
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0053	-0.0052	-0.0000	-0.0000	-0.1361
N= 92708							
Panel B: Individuals with a Credit Card Who Paid Interest at Baseline							
% TE	-0.0814***	0.0639***	-0.0109	-0.0110	-0.0041	0.0015	-0.0029
	(0.0197)	(0.0213)	(0.0260)	(0.0258)	(0.0202)	(0.0205)	(0.0194)
Mean of Dep. Var. in Control Group (MXN)	31,255.67	34,619.81	378.97	375.65	0.87	0.86	11,847.35
Change in Spending or Saving (MXN)	-2,545.43	2,213.93					
Upper Confidence Interval (MXN)			15.16	14.88	0.03	0.04	415.58
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0060	0.0058	0.0000	0.0000	0.1633
Lower Confidence Interval (MXN)			-23.43	-23.13	-0.04	-0.03	-485.45
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0092	-0.0091	-0.0000	-0.0000	-0.1907
N= 45637							

This table shows treatment effects for individuals in the top quartile of predicted treatment effects on spending who have a credit card and for whom Banorte is their main bank (i.e., they receive their payroll at Banorte and, according to credit bureau records, they don't have credits outside of Banorte). Variable definitions are provided in Table 2 and the timing of measurement is described in Table 5. Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest at baseline. Treatment effects are calculated using Poisson regressions with strata fixed effects. We report percentage treatment effects calculated as exponentiated coefficients -1. The change in spending (saving), expressed in MXN, is calculated by multiplying the % TE and the Mean of Spending (Checking Account Balances) in the Control Group. Upper (lower) confidence intervals, expressed in MXN, are calculated as (point estimate +/- 1.96*Standard Error)*Mean of Dep. Var. in Control Group. The upper (lower) confidence interval for the variable Paid Interest is calculated as (point estimate +/- 1.96*Standard Error). Abs. Value stands for absolute value. The number of observations is the same across all columns in the same panel. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

Table IA.XI. Treatment Effects on Spending, Saving, and Borrowing (Top Quartile of Predicted Treatment Effects on Saving — Main Bank)

Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Monthly Spending	Checking Account Balance	Monthly Credit Card Interest During Treat. (Banorte)	Monthly Credit Card Interest After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ending Statement Balance - Payments
Panel A: Individuals with a Credit Card							
% TE	-0.0625***	0.0675***	-0.0098	-0.0097	0.0014	-0.0032	-0.0066
	(0.0164)	(0.0140)	(0.0238)	(0.0238)	(0.0120)	(0.0121)	(0.0311)
Mean of Dep. Var. in Control Group (MXN)	31,462.19	33,685.49	221.62	224.38	0.48	0.47	5,204.41
Change in Spending or Saving (MXN)	-1,964.83	2,274.30					
Upper Confidence Interval (MXN)			8.16	8.29	0.01	0.01	282.89
Upper Confidence Interval Divided by Abs. Value of Change in Saving			0.0036	0.0036	0.0000	0.0000	0.1244
Lower Confidence Interval (MXN)			-12.51	-12.65	-0.01	-0.01	-351.59
Lower Confidence Interval Divided by Abs. Value of Change in Saving			-0.0055	-0.0056	-0.0000	-0.0000	-0.1546
N= 90114							
Panel B: Individuals with a Credit Card Who Paid Interest at Baseline							
% TE	-0.0672***	0.0689***	-0.0108	-0.0109	-0.0037	-0.0035	-0.0043
	(0.0224)	(0.0193)	(0.0274)	(0.0275)	(0.0201)	(0.0201)	(0.0213)
Mean of Dep. Var. in Control Group (MXN)	30,726.52	32,947.75	361.23	359.52	0.74	0.75	10,239.92
Change in Spending or Saving (MXN)	-2,063.99	2,271.61					
Upper Confidence Interval (MXN)			15.48	15.45	0.03	0.03	383.11
Upper Confidence Interval Divided by Abs. Value of Change in Saving			0.0068	0.0068	0.0000	0.0000	0.1686
Lower Confidence Interval (MXN)			-23.30	-23.26	-0.03	-0.03	-470.88
Lower Confidence Interval Divided by Abs. Value of Change in Saving			-0.0103	-0.0102	-0.0000	-0.0000	-0.2073
N= 43907							

This table shows treatment effects for individuals in the top quartile of predicted treatment effects on saving who have a credit card and for whom Banorte is their main bank (i.e., they receive their payroll at Banorte and, according to credit bureau records, they don't have credits outside of Banorte). Variable definitions are provided in Table 2 and the timing of measurement is described in Table 5. Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest at baseline. Treatment effects are calculated using Poisson regressions with strata fixed effects. We report percentage treatment effects calculated as exponentiated coefficients -1. The change in spending (saving), expressed in MXN, is calculated by multiplying the % TE and the Mean of Spending (Checking Account Balances) in the Control Group. Upper (lower) confidence intervals, expressed in MXN, are calculated as (point estimate +/- 1.96*Standard Error)*Mean of Dep. Var. in Control Group. The upper (lower) confidence interval for the variable Paid Interest is calculated as (point estimate +/- 1.96*Standard Error). Abs. Value stands for absolute value. The number of observations is the same across all columns in the same panel. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

IA.IIIC. Individuals with Credit Card Billing Cycles Fully Covered by the Intervention

Figure IA.2 shows the distribution of the day of the month on which billing cycles end for all individuals with a credit card in our sample.

Figure IA.2. Distribution of Credit Card Billing Cycle End Dates

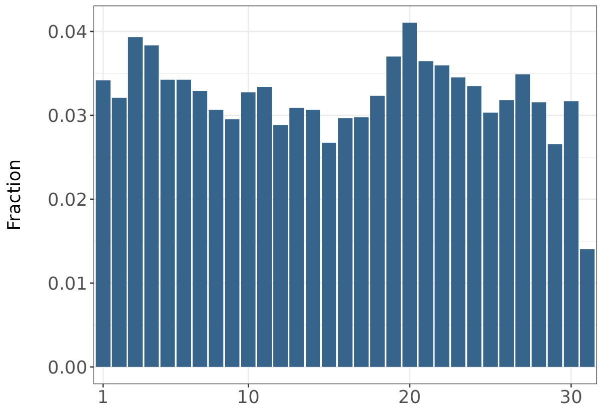

This figure shows the frequency distribution of credit card billing cycle end dates for August 2019, the month immediately preceding the intervention.

Table IA.XII shows the main result when we focus on individuals who have a billing cycle that lies entirely within the 7-week window of the intervention. For them, we measure monthly spending and savings during the billing cycle, which is fully covered by the intervention. The results are the same as those of our main specifications.

Table IA.XII. Treatment Effects on Spending, Saving, and Borrowing, for Individuals with a Billing Cycle Fully Covered by the Intervention (Top Quartile of Predicted Treatment Effects on Spending)

Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Monthly Spending	Checking Account Balance	Monthly Credit Card Interest During Treat. (Banorte)	Monthly Credit Card Interest After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ending Statement Balance - Payments
Panel A: Individuals with a Credit Card							
% TE	-0.0729***	0.0536***	-0.0086	-0.0089	-0.0038	-0.0018	-0.0030
	(0.0014)	(0.0141)	(0.0266)	(0.0265)	(0.0118)	(0.0118)	(0.0308)
Mean of Dep. Var. in Control Group (MXN)	35,826.12	40,296.45	208.14	207.12	0.46	0.46	5,179.32
Change in Spending or Saving (MXN)	-2,613.42	2,158.46					
Upper Confidence Interval (MXN)			9.06	8.91	0.01	0.01	297.13
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0035	0.0034	0.0000	0.0000	0.1137
Lower Confidence Interval (MXN)			-12.64	-12.61	-0.01	-0.01	-328.20
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0048	-0.0048	-0.0000	-0.0000	-0.1256
N= 96120							
Panel B: Individuals with a Credit Card Who Paid Interest at Baseline							
% TE	-0.0767***	0.0585**	-0.0092	-0.0092	-0.0031	-0.0025	-0.0031
	(0.0210)	(0.0235)	(0.0284)	(0.0284)	(0.0133)	(0.0133)	(0.0221)
Mean of Dep. Var. in Control Group (MXN)	34,629.69	35,780.30	379.33	377.11	0.87	0.86	10,675.84
Change in Spending or Saving (MXN)	-2,657.41	2,094.13					
Upper Confidence Interval (MXN)			17.61	17.51	0.02	0.02	428.17
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0066	0.0066	0.0000	0.0000	0.1611
Lower Confidence Interval (MXN)			-24.57	-24.42	-0.03	-0.02	-495.16
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0092	-0.0092	-0.0000	-0.0000	-0.1863
N= 47106							

This table shows treatment effects for individuals in the top quartile of predicted treatment effects on spending who have a credit card with a billing cycle fully covered by the intervention. Variable definitions are provided in Table 2. For each individual we identify the calendar weeks corresponding to said billing cycle and measure spending, checking account balances, and borrowing behavior over those weeks. Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest at baseline. Treatment effects are calculated using Poisson regressions with strata fixed effects. We report percentage treatment effects calculated as exponentiated coefficients -1. The change in spending (saving), expressed in MXN, is calculated by multiplying the % TE and the Mean of Spending (Checking Account Balances) in the Control Group. Upper (lower) confidence intervals, expressed in MXN, are calculated as (point estimate +/- 1.96*Standard Error)*Mean of Dep. Var. in Control Group. The upper (lower) confidence interval for the variable Paid Interest is calculated as (point estimate +/- 1.96*Standard Error). Abs. Value stands for absolute value. The number of observations is the same across all columns in the same panel. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level.

IA.IIID. Treatment Effects by Week

As discussed in the main text, credit card interest is not defined at the weekly level since it is calculated based on average daily balances over an entire billing cycle. To build a comparable monthly measure of spending and saving, we look at spending and saving over the first month of the treatment. This is without loss of generality if the effectiveness of the treatment does not change over time. To confirm that this is the case, we look at the weekly evolution of treatment effects.

The data for this analysis is at the week-by-user level. We estimate equation IA.III1:

$$y_{it} = exp\{\alpha_s + +\alpha_t + \beta_1 * Treatment_i + \sum_{t=2...7} \beta_{2,t} * Treatment_i * Week_t + error_{it}\} \quad \text{(IA.III1)}$$

Where α_s and α_t are strata and week fixed effects, respectively. Standard errors are clustered at the individual level. We define weekly spending as the sum of ATM withdrawals, credit and debit card spending, and outgoing transfers over the corresponding week. We calculate incremental effects on weekly spending relative to the first week as $exp\{\beta_1 + \beta_{2,t}\} - exp\{\beta_1\}$. For the savings analysis, we consider the average of daily checking account balances observed for each individual over the corresponding week. Since balances are a stock and not a flow, the treatment effect on weekly savings is captured by the difference between the cumulative effects on two consecutive weeks, which in an exponential regression corresponds to $exp\{\beta_1 + \beta_{2,t}\} - exp\{\beta_1 + \beta_{2,t-1}\}$ for t= 2,...,7. We then plot incremental effects relative to the percentage treatment effect of the first week. The results are presented in Figures IA.3 and IA.4.

We can see that there are no statistically significant differences in the effectiveness of the intervention over time. We note that we do not have data on saving and spending after the intervention. As a result, we do not know if the treatment would have retained its effectiveness had it lasted longer or if the increases in savings (or reductions in spending) would persist after the intervention. From a theoretical perspective, we cannot distinguish whether the preference shock induced by the nudge is permanent or temporary.

Figure IA.3. Weekly Treatment Effects on Spending, Relative to First Week

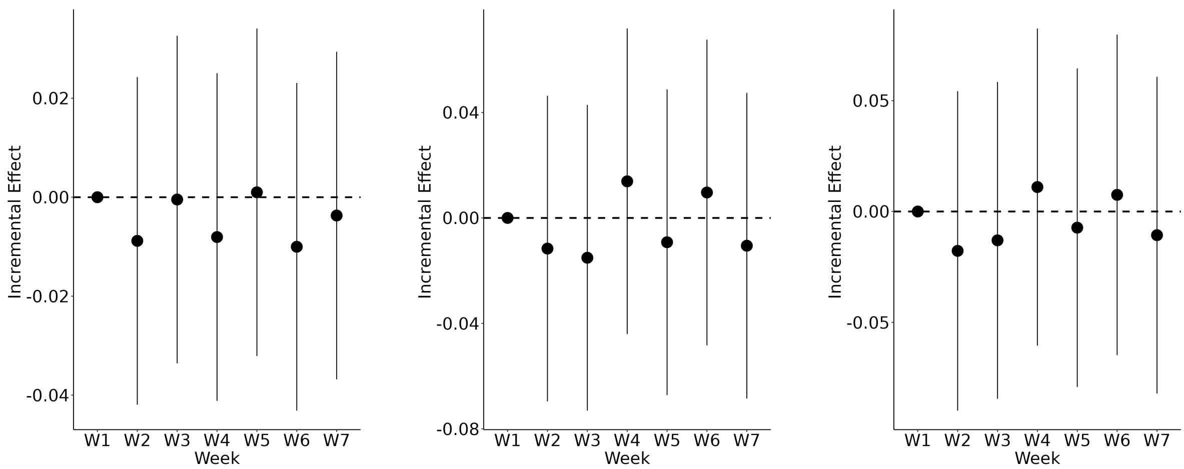

(a) All individuals with a Credit Card

(b) Individuals with Credit Cards in Top Quartile of the Predicted Treatment Effect Distribution

(c) Individuals with Credit Cards paying Interest in Top Quartile of the Predicted Treatment Effect Distribution

This figure shows the incremental effect of the treatment on spending on a weekly basis relative to the first week of the intervention. Spending is defined as the sum of ATM withdrawals, credit and debit card spending, and outgoing transfers. We estimate the following equation: $y_{it} = exp\{\alpha_s + \alpha_t + \beta_1 * Treatment_i + \sum_{t=2...7} \beta_{2,t} * Treatment_i * Week_t + error_{it}\}$. We plot incremental effects every week relative to the percentage treatment effect of week 1. The vertical lines correspond to 95% confidence intervals. Standard errors are clustered at the individual level. We consider all individuals who have a credit card, those in the top quartile of the distribution of treatment effects on spending, and those who also pay credit card interest at baseline.

Figure IA.4. Weekly Treatment Effects on Saving, Relative to First Week

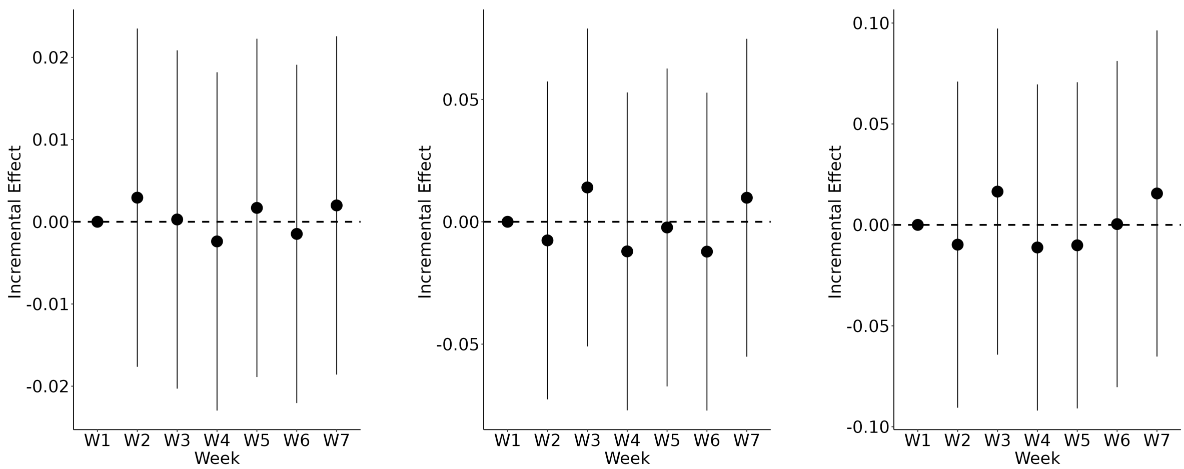

(a) All individuals with a Credit Card

(b) Individuals with Credit Cards in Top Quartile of the Predicted Treatment Effect Distribution

(c) Individuals with Credit Cards paying Interest in Top Quartile of the Predicted Treatment Effect Distribution

This figure shows the incremental effect of the treatment on weekly savings relative to the first week of the intervention. We estimate the following equation: $y_{it} = exp\{\alpha_s + \alpha_t + \beta_1 * Treatment_i + \sum_{t=2...7} \beta_{2,t} * Treatment_i * Week_t + error_{it}\}$. The outcome variable corresponds to the average of daily checking account balances observed over the corresponding week. We plot incremental effects on every week relative to the percentage treatment effect of week 1. The vertical lines correspond to 95% confidence intervals. Standard errors are clustered at the individual level. We consider all individuals who have a credit card, those in the top quartile of the distribution of treatment effects on spending, and those who also pay credit card interest at baseline.

IA.IV. Models: Formal Details and Derivations

IA.IVA. Transactions-convenience Model:

There are two periods, one consumption good and the agent's utility function is a log specification. The agent consumes in periods 1 and 2 (denoted by c_1 and c_2) where $c_1 > 0$ and $c_2 > 0$. In period 1, they may borrow, $0 \leq b_1 < b$, to consume because they must hold a certain amount of cash for transaction purposes x. Savings are denoted by s_1. The initial endowment is denoted by $x_1 > x$ and the interest for borrowing by r (cash pays no interest). The agent maximizes

$$max\{log(c_1) + \delta log(c_2)\}$$

subject to:

$$c_1 + s_1 \leq x_1 + b_1$$

$$c_2 + b_1(1 + r) \leq s_1$$

$$s_1 = x_1 - c_1 + b_1 \geq x$$

Combining the budget constraints for periods 1 and 2, we can rewrite the problem as:

$$\max_{c_1}\{log(c_1) + \delta log(x_1 - c_1 - rb_1)\}$$

$$\text{where } b_1 := f_{b_1}(c_1) = \begin{cases} c_1 - x_1 + x & \text{if } c_1 - x_1 + x > 0 \\ 0 & \text{otherwise.} \end{cases}$$

We assume the agent is not credit-constrained, $b = \infty$, and their discount factor is denoted by $\delta \in (0, 1)$.

The agent's optimal consumption in period 1 is discontinuous and determines whether or not they borrow. Either the agent is borrowing and thus co-holding in period 1 or not borrowing/co-holding. When they do not borrow, they may have savings greater than the required amount x or hold no additional savings in which case $c_1 = x_1 - x$.

The agent is not borrowing/co-holding if the following condition holds

$$x_1 - \frac{x_1}{1 + \delta} > x$$

and optimal consumption and borrowing, denoted by c_1^* and $b_1^* := f_{b_1}(c_1^*)$, are then

$$\frac{1}{c_1} = \delta \frac{1}{x_1 - c_1} \Rightarrow c_1^* = \frac{x_1}{1 + \delta}$$

and $b_1^* = 0$.

The agent is borrowing/co-holding if the following condition holds:

$$x_1 - \frac{1}{\delta+1}x_1 + \frac{r}{(\delta+1)(1+r)}x < x$$

and optimal consumption and borrowing are then given by

$$\frac{1}{c_1} = \delta \frac{1+r}{x_1 - c_1 - r(c_1 - x_1 + x)} \Rightarrow c_1^* = \frac{1}{\delta+1}x_1 - \frac{r}{(\delta+1)(1+r)}x$$

and $b_1^* = c_1^* - x_1 + x$.

Finally, the intermediate case when the agent just holds the required amount of cash x is given when

$$x_1 - \frac{x_1}{1+\delta} \leq x \leq x_1 - \frac{1}{\delta+1}x_1 + \frac{r}{(\delta+1)(1+r)}x$$

and optimal consumption equals $c_1^* = x_1 - x$ and $b_1^* = 0$.

Proof of Proposition 1. Comparative statics with respect to δ when the agent co-holds:

$\frac{\partial c_1^*}{\partial \delta} = \frac{-(1+r)x_1 + rx}{(\delta+1)^2(1+r)} < 0$ (since $x_1 > x$) the agent consumes less if they are more patient and want to save more.

$\frac{\partial b_1^*}{\partial \delta} = \frac{\partial(c_1^* - x_1 + x)}{\partial \delta} = \frac{\partial c_1^*}{\partial \delta} < 0$ the agent consumes less if they are more patient so they borrow less.

□

Proof of Proposition 2. Comparative statics with respect to x when the agent co-holds:

$\frac{\partial c_1^*}{\partial x} = \frac{-r}{(\delta+1)(1+r)} < 0$ this is because the agent is marginally poorer if they must borrow to maintain their larger cash needs, zero if $r = 0$.

$\frac{\partial b_1^*}{\partial x} = 1 + \frac{\partial c_1^*}{\partial x}$ the agent's cash needs are directly reflected in their borrowing (they consume slightly less because they are marginally poorer and thus, they borrow a bit less), the borrowing increase is equal to the increase in x if $r = 0$.

□

IA.IVB. Credit-Limit-Chasing Model:

Three periods, one consumption good, and log utility. The agent consumes only in periods 2 and 3, denoted by $c_2 > 0$ and $c_3 > 0$, in period 1 and 2 they may borrow to consume, denoted by $0 \leq b_1 < b$ and $0 \leq b_2 < b$, if they do not borrow in period 1, they cannot borrow in period 2 either, that is, $b_2 \leq b_1$. The agent expects to receive an income payment, denoted by $y_3 > 0$ in period 3, and holds cash $x_2 > 0$ at the beginning of period 2. The interest rate for borrowing is $0 \leq r < 1$, all interest is paid in period 3, and cash pays no interest. The agent is not credit-constrained $b = \infty$. b_2 and b_1 can be interpreted as one credit line that needs to be used to stay open or maintain the existing outstanding loan because of the possible inability to secure a new loan if the existing one were repaid. $x_2 > 0$ and $b_1 > 0$ denote cash-on-hand and debt-on-hand at the beginning of period 2, and the agent's maximization problem is:

$$\max_{c_2}\{log(c_2) + \delta log(y_3 + x_2 - c_2 - rb_1 - rb_2)\}$$

The agent's optimal consumption in period 1 is discontinuous. Either the agent is borrowing in periods 1 and 2 (and thus co-holding in period 1), or not borrowing. Because the only reason to borrow in period 1 is to ensure borrowing in period 2, $b_2 = b_1$.

The agent is not borrowing if the following condition holds:

$$\frac{1}{\delta+1}(y_3 + x_2) < x_2 \Rightarrow y_3 < x_2\delta$$

Then the optimal solutions for c_2, b_2, and b_1, denoted by c_2^*, b_2^*, and b_1^*, are determined by:

$$\frac{1}{c_2} = \delta\frac{1}{y_3 + x_2 - c_2} \Rightarrow c_2 = \frac{1}{\delta}(y_3 + x_2 - c_2) \Rightarrow c_2^* = \frac{1}{\delta+1}(y_3 + x_2) \text{ and } b_2^* = b_1^* = 0$$

The agent is borrowing if the following condition holds:

$$\frac{y_3 + x_2(1+2r)}{(1+2r)(\delta+1)} > x_2 \Rightarrow y_3 > x_2(1+2r)\delta$$

Then the optimal solutions for c_2, b_2, and b_1, denoted by c_2^*, b_2^*, and b_1^*, are determined by:

$$\frac{1}{c_2} = \frac{1+2r}{y_3 + x_2 - c_2 - 2r(c_2 - x_2)} \Rightarrow c_2\delta(1+2r) = y_3 + x_2(1+2r) - c_2(1+2r)$$

$$\Rightarrow c_2\delta(1+2r) = y_3 + x_2(1+2r) - c_2(1+2r) \Rightarrow c_2^* = \frac{y_3 + x_2(1+2r)}{(1+2r)(\delta+1)}$$

$$\text{and } b_2^* = b_1^* = c_2^* - x_2$$

Finally, the intermediate case when the agent is not borrowing or saving is given when:

$$x_2\delta \leq y_3 \leq x_2(1+2r)\delta$$

Then the optimal solutions for c_2, b_2, and b_1, denoted by c_2^*, b_2^*, and b_1^*, are:

$$c_2^* = x_2 \text{ and } b_2^* = b_1^* = 0$$

Comparative statics with respect to x_2 if the agent is borrowing:

$\frac{\partial c_2^*}{\partial x_2} = \frac{1}{(\delta+1)} > 0$ this is because the agent is richer and consumes more.

$\frac{\partial b_2^*}{\partial x_2} = \frac{1}{(\delta+1)} - 1 < 0$ this is because the agent has to borrow less in period 2 when they have more cash to consume.

Comparative statics with respect to δ if the agent is borrowing:

$\frac{\partial c_2^*}{\partial \delta} = -\frac{y_3+x_2(1+2r)}{(1+\delta)^2(1+2r)} < 0$ the agent consumes less if they are more patient and want to save more.

$\frac{\partial b_1^*}{\partial \delta} = \frac{\partial b_2^*}{\partial \delta} = \frac{\partial c_2^*}{\partial \delta} < 0$ and borrows less.

IA.IVC. Self-Control Model:

Two periods, one consumption good, and log utility. There is a patient and impatient party. The impatient party consumes positive amounts in periods 1 and 2, denoted by $c_1 > 0$ and $c_2 > 0$; in period 1, the impatient party may borrow to consume more, $0 \leq b_1 < b$. The initial endowment is denoted by x_1. In period 0, the impatient agent's previous patient self and/or their patient spouse takes a certain amount of cash, denoted by $x \geq 0$, of the overall initial endowment, x_1, and deposits it in an inaccessible savings account that is only for period-2 consumption. A fraction $a \in (0, 1]$ of this amount is invisible to the impatient party, or they forget about it. Therefore, the impatient agent's perceived endowment in period 1 is $x_1 - ax$ and they perceive the inaccessible savings in period 1 as $(1 - a)x$. The perceived savings are determined by $1 - a$ because that is the fraction of inaccessible resources, x that is visible to the impatient party. Therefore, if they want to consume more than perceived endowment minus perceived savings, $x_1 - ax - (1 - a)x = x_1 - x$, in period 1, they must borrow. The impatient party's period-1 maximization problem is

$$\max_{c_1}\{log(c_1) + \beta log(x_1 - ax - c_1 - rb_1)\}$$

$$\text{where } b_1 := f_{b_1}(c_1) = \begin{cases} c_1 - x_1 + x & \text{if } c_1 - x_1 + x > 0 \\ 0 & \text{otherwise.} \end{cases}$$

β denotes the impatient party's discount factor, and $\delta \in (0, 1)$ is the patient party's discount factor; we assume that $0 < \beta < \delta$. We further assume that the impatient party is not credit-constrained $b = \infty$, that r is small, and that $\delta - \beta < (1 + \delta)a$.

The impatient party's optimal consumption and borrowing, denoted by c_1^* and b_1^*, are discontinuous. Either the impatient agent is borrowing, or they choose not to borrow, in which case they may or may not hold positive savings. When the impatient agent borrows, the household co-holds (only if $x > 0$ and $a < 1$ the impatient agent will borrow as they would otherwise face negative consumption in period 2).

The agent is not borrowing (not co-holding) if the following condition holds:

$$x_1 - x - \frac{1}{\beta + 1}(x_1 - ax) > 0$$

Then the optimal solution for c_1 and b_1 is:

$$\frac{1}{c_1} = \beta \frac{1}{x_1 - ax - c_1} \Rightarrow c_1 = \frac{1}{\beta}(x_1 - ax - c_1)$$

$$\Rightarrow c_1^* = \frac{1}{\beta+1}(x_1 - ax) \text{ and } b_1^* = 0$$

The agent is borrowing and co-holding if the following condition holds:

$$x_1 - x - \frac{1}{\beta+1}x_1 + \frac{r+a}{(\beta+1)(1+r)}x < 0$$

Then the optimal solution for c_1 and b_1 is:

$$\frac{1}{c_1} = \beta \frac{1+r}{x_1 - ax - c_1 - r(c_1 - x_1 + x)} \Rightarrow c_1 = \frac{1}{\beta(1+r)}(x_1 - ax - c_1 - r(c_1 - x_1 + x))$$

$$\Rightarrow c_1(1+\frac{1}{\beta}) = \frac{1}{\beta(1+r)}(x_1 - ax - r(-x_1 + x)) \Rightarrow c_1^* = \frac{1}{(\beta+1)(1+r)}(x_1(1+r) - (r+a)x)$$

$$= \frac{1}{\beta+1}x_1 - \frac{r+a}{(\beta+1)(1+r)}x \text{ and } b_1^* = c_1^* - x_1 + x.$$

Finally, the intermediate case, when the impatient agent just consumes all of their available resources, is given when

$$x_1 - x - \frac{1}{\beta+1}x_1 + \frac{1}{\beta+1}ax \leq 0 \leq x_1 - x - \frac{1}{\beta+1}x_1 + \frac{r+a}{(\beta+1)(1+r)}x$$

Note that this is a possible parameter range as

$$\frac{1}{\beta+1}a < \frac{r+a}{(\beta+1)(1+r)} \Rightarrow a < \frac{r+a}{(1+r)} \Rightarrow a + ar < r + a$$

and optimal consumption equals $c_1^* = x_1 - x$ and $b_1^* = 0$.

Note that, for the transactions-convenience agent, optimal consumption was $c_1^* = \frac{1}{\delta+1}x_1 - \frac{r}{(\delta+1)(1+r)}x$ for the borrowing/co-holding case. Therefore, when $a = 0$ (no amount of the money can be hidden), we have the impatient agent fully taking into account the cash that is put away (as the transactions-convenience agent would if his discount factor were equal to β). Instead if $a = 1$, then the impatient agent will consume $\frac{1}{\beta+1}(x_1 - x)$ and split the entire lost amount x over their consumption in the two periods.

The impatient party's optimal consumption and borrowing depends on the amount withheld x, so we define $c_1^* := f_{c_1^*}(x)$ and $b_1^* := f_{b_1^*}(x)$. In turn, the optimal amount withheld x, denoted by x^*, is decided by a more patient party with the following problem:

$$\max_x \{log(f_{c_1^*}(x)) + \delta log(x_1 - f_{c_1^*}(x) - r f_{b_1^*}(x))\}$$

The patient party's optimal cash withheld, denoted by x^*, is discontinuous. They take into account whether the impatient agent is not borrowing or the impatient agent is borrowing, which depends on their choice of x.

First, the impatient agent is not borrowing if the following condition holds:

$$x_1 - x_1\frac{\delta-\beta}{(\delta+1)a} - \frac{1}{\beta+1}(x_1 - ax_1\frac{\delta-\beta}{(\delta+1)a}) > 0$$

Then the optimal solution for x is determined by:

$$\max_x\{log(\frac{1}{\beta+1}(x_1-ax)) + \delta log(x_1 - \frac{1}{\beta+1}(x_1-ax))\}$$

$$\Rightarrow \frac{-\frac{1}{\beta+1}a}{\frac{1}{\beta+1}(x_1-ax)} + \frac{\frac{\delta}{\beta+1}a}{(x_1-\frac{1}{\beta+1}(x_1-ax))} = 0 \Rightarrow x_1 - \frac{1}{\beta+1}(x_1-ax) = \frac{\delta}{\beta+1}(x_1-ax)$$

$$\Rightarrow x_1 - \frac{1}{\beta+1}x_1 - \frac{\delta}{\beta+1}x_1 = -\frac{1}{\beta+1}ax - \frac{\delta}{\beta+1}ax \Rightarrow -\frac{1+\delta}{\beta+1}ax = x_1(1-\frac{1+\delta}{\beta+1})$$

$$\Rightarrow x^* = x_1\frac{\delta-\beta}{(\delta+1)a}$$

Note that $0 < x^* < x_1$ as $a \in (0,1]$, $\beta \in (0,\delta)$, $\delta \in (0,1)$, and $\delta - \beta < (\delta+1)a$.

Second, the impatient agent is borrowing and co-holding, if the following condition holds:

$$x_1 - x_1\frac{1}{1+\delta}(\delta\frac{1+r}{r+a} - \frac{\beta+r\beta}{a-r\beta}) - \frac{1}{\beta+1}x_1 + \frac{r+a}{(\beta+1)(1+r)}x_1\frac{1}{1+\delta}(\delta\frac{1+r}{r+a} - \frac{\beta+r\beta}{a-r\beta}) < 0$$

Then the optimal solution for x is determined by:

$$\frac{1}{c_1^*}\frac{\partial c_1^*}{\partial x} + \frac{-\delta\frac{\partial c_1^*}{\partial x}(1+r) - \delta r}{(x_1 - c_1^* - r(c_1^* - x_1 + x))} = 0$$

$$\Rightarrow \frac{1}{\frac{1}{\beta+1}x_1 - \frac{r+a}{(\beta+1)(1+r)}x}\frac{r+a}{(\beta+1)(1+r)} = \delta\frac{\frac{r+a}{\beta+1} - r}{x_1(1+r) - (1+r)(\frac{1}{\beta+1}x_1 - \frac{r+a}{(\beta+1)(1+r)}x) - rx}$$

$$\Rightarrow \frac{1}{x_1\frac{1+r}{r+a} - x} = \delta\frac{\frac{a-r\beta}{\beta+1}}{x_1(1+r)\frac{\beta}{\beta+1} + (\frac{r+a}{\beta+1} - r)x} \Rightarrow \frac{1}{x_1\frac{1+r}{r+a} - x} = \delta\frac{1}{x_1(1+r)\frac{\beta}{a-r\beta} + x}$$

$$\Rightarrow (x_1\frac{1+r}{r+a} - x)\delta = x_1(1+r)\frac{\beta}{a-r\beta} + x$$

$$\Rightarrow x(1+\delta) = \delta x_1\frac{1+r}{r+a} - x_1\frac{(1+r)\beta}{a-r\beta} \Rightarrow x^* = x_1\frac{1}{1+\delta}(\delta\frac{1+r}{r+a} - \frac{\beta+r\beta}{a-r\beta})$$

And note that for $r \to 0$ we obtain $0 < x^* < x_1$ as $0 < \frac{1}{1+\delta}(\delta\frac{1}{a} - \frac{\beta}{a}) < 1$ as $a \in (0,1]$, $\beta \in (0,\delta)$, $\delta \in (0,1)$, and $\delta - \beta > (\delta+1)a$.

Finally, there is the intermediate case, when the impatient agent just consumes all of their available resources. This case is determined by the following conditions

$$x_1 - x_1\frac{\delta-\beta}{(\delta+1)a} - \frac{1}{\beta+1}x_1 + \frac{1}{\beta+1}ax_1\frac{\delta-\beta}{(\delta+1)a} \le 0$$

and $0 \le x_1 - x_1\frac{1}{1+\delta}(\delta\frac{1+r}{r+a} - \frac{\beta+r\beta}{a-r\beta}) - \frac{1}{\beta+1}x_1 + \frac{r+a}{(\beta+1)(1+r)}x_1\frac{1}{1+\delta}(\delta\frac{1+r}{r+a} - \frac{\beta+r\beta}{a-r\beta})$

(note that, these expressions can be simplified to see that the conditions can be satisfied by certain parameter values)

$$x_1\underbrace{\frac{\delta(a-1)+\beta}{(\delta+1)a}}_{\lesseqgtr 0} \le 0 \text{ and } x_1(\underbrace{\frac{\beta}{\beta+1}}_{<1} - \underbrace{\frac{1}{1+\delta}(\delta\frac{1+r}{r+a} - \frac{\beta+r\beta}{a-r\beta})}_{=\frac{x}{x_1}<1}\underbrace{\frac{(\beta+1)(1+r)-(r+a)}{(\beta+1)(1+r)}}_{\in[0,1]}) \ge 0.$$

Then the optimal solution for x can be found by solving:

$$\max_x\{log(x_1 - x) + \delta log(x_1 - x_1 + x)\}$$

$$\Rightarrow \frac{1}{x_1 - x} = \delta\frac{1}{x} \Rightarrow (x_1 - x)\delta = x \Rightarrow x^* = \frac{\delta}{\delta+1}x_1 \text{ and } 0 < x^* < x_1 \text{ for } \delta \in (0,1)$$

Proof of Proposition 3. Comparative statics with respect to δ (discount factor of patient party) when the agent co-holds:

$\frac{\partial x^*}{\partial \delta} = -x_1\frac{1}{(1+\delta)^2}(\delta\frac{1+r}{r+a} - \frac{\beta+r\beta}{a-r\beta}) + x_1\frac{1}{1+\delta}\frac{1+r}{r+a} = \frac{1}{1+\delta}(-x^* + x_1\underbrace{\frac{1+r}{r+a}}_{\ge 1})) > 0$ as $x_1 > x$ if the patient party is more patient then they withhold more money. As a result, the impatient party consumes less as $\frac{\partial c_1^*}{\partial x} \le 0$ and $\frac{\partial b_1^*}{\partial x} = \frac{\partial c_1^*}{\partial x} + 1$. When the impatient agent's consumption decreases a lot in response to an increase in x, the effect on borrowing is reduced.

Comparative statics with respect to β (discount factor of the patient party) when the agent co-holds:

$\frac{\partial x^*}{\partial \beta} = -x_1\frac{1}{1+\delta}\frac{(1+r)(a-r\beta)-(\beta+r\beta)(-r)}{(a-r\beta)^2} = -x_1\frac{(1+r)}{1+\delta}\frac{a-r\beta+\beta r}{(a-r\beta)^2} = -x_1\frac{(1+r)}{1+\delta}\frac{a}{(a-r\beta)^2} < 0$ if the impatient party becomes more impatient then the patient party withholds more money.

□

Proof of Proposition 4. Comparative statics of impatient party with respect to x when the agent co-holds:

$\frac{\partial c_1^*}{\partial x} = -\frac{r+a}{(\beta+1)(1+r)} < 0$ the moment cash is taken away then $\frac{\partial c_1^*}{\partial x} < 0$ (either because the hiding-of-accounts parameter a is positive or because borrowing costs interest that is taken into account), $\frac{\partial c_1^*}{\partial x} = -\frac{1+r}{(\beta+1)(1+r)} < -\frac{r}{(\beta+1)(1+r)}$, and $\frac{\partial c_1^*}{\partial x \partial a} < 0$ the sensitivity of consumption to hidden cash is negative and larger when a goes to 1 because the agent considers the cash x as taken away and they are not aware of it; it is not -1 because the agent distributes the overall loss in resources to consumption in periods 1 and 2.

$\frac{\partial b_1^*}{\partial x} = \frac{\partial c_1^*}{\partial x} + 1$ so when the agent's consumption goes down a lot in response to an increase in x, the effect on borrowing is reduced.

□