Does Saving Cause Borrowing? Implications for the Co-Holding Puzzle

Paolina C. Medina* and Michaela Pagel†

May 31, 2023

Abstract

We analyze an experiment involving 3.1 million bank customers who were encouraged to save through SMS messages. We first theoretically show that by examining their spending, saving, and borrowing responses we can distinguish between the leading explanations for co-holding liquid savings and credit card debt. Using a machine learning algorithm, we then predict individual-level treatment effects and find that the most responsive individuals reduce spending and increase their savings by 5.1% (225 USD PPP per month), while their credit card debt remains unchanged. We argue that these joint findings suggest people co-hold because they mentally separate savings and debt accounts.

Keywords: saving nudges, credit card borrowing, co-holding puzzle, heterogeneous treatment effects, causal forest

JEL codes: G5, D14

*Mays Business School of Texas A&M University. E-mail: pmedina@tamu.edu

†Columbia GSB, NBER, and CEPR. E-mail: mpagel@columbia.edu

We thank seminar and conference participants at University of Houston, AEA Meetings, Darthmouth 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. We also thank 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 thank Jose Antonio Murillo for support and access to the data, Isidoro Garcia for helping with preliminary analysis, and Sofía Terrazas for outstanding research assistance. Portions of this research were conducted with computing resources provided by Texas A&M High Performance Research Computing Cluster. The experiment was pre-registered under RCT number AEARCTR-0003047. IRB ID:IRB2018-0597. At the time of the experiment, Medina was a full-time employee of Banorte.

1 Introduction

A significant number of US households hold both liquid savings and credit card debt (Gross and Souleles, 2002), which appears to be an 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 mental accounting or agency problems within households, where individuals do not use their savings to repay credit card debt because they want to control their own impulsive behavior or the financial decisions of their spouses. If they were to pay off their debt, they would just accumulate debt again in the future and effectively spend their savings (Bertaut et al., 2009; Vihriälä, 2019).

Despite the prevalence of co-holding liquid savings and credit card debt, there are numerous policies in place to promote savings, many of which use nudges to encourage people to save (Benartzi et al., 2017). However, these policies do not fully understand where the money for savings comes from, as they often focus on immediate savings outcomes (Thaler, 1994; Beshears and Kosowsky, 2020).

This paper aims to investigate the impact of saving nudges on spending, saving, and high-interest debt, specifically credit card debt. By doing so, we gain insights into the mechanisms behind the widespread co-holding of savings and credit card debt, which is important for researchers and policymakers alike.

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, 2,679,545 customers were treated with (bi-)weekly SMS messages encouraging them to save for 7 weeks, while a distinct group of 374,893 customers received no messages.

Co-holding liquid savings and credit card debt is a common practice not only in the US but also in Mexico. In our sample, the average credit card interest rate is 35.2%, while checking accounts do not pay any interest. Despite the significant difference in rates, we observe that 13.5% of individuals who pay credit card interest keep balances higher than 50% of their income in their checking accounts (measured as the minimum balance over 6 months). If these individuals were to use their savings to repay their credit card debt, they could reduce their interest charges by an amount equivalent to 5% of their monthly income.

We first theoretically examine the effects of saving nudges on individual behavior in the presence of different underlying economic mechanisms for co-holding. We find that, if the nudges make individuals more patient, then models based on a liquidity premium for cash 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, then these models would predict an increase in cash holdings, which are financed by increases in debt to smooth consumption.

On the other hand, models of mental accounting make different predictions. If the nudges made individuals more patient and increase the cash holdings that are separated and hidden from the impatient self or spouse, then the models would predict a decrease in spending, an increase in savings, and no change in credit card debt. Therefore, individuals may keep their credit card debt unchanged despite having more savings.

We then 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 some idiosyncratic factors or noise, which we would spuriously attribute to the treatment. These idiosyncratic factors might then also influence other outcomes, such as borrowing, and invalidate our inference.[1]

In our main analysis, we focus 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 7.8% decrease in spending and a 5.1% increase in savings, as measured by transactions 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 0.71% per month. This effect is small and precisely estimated with a standard error of 1.76%. To put this in perspective, we compare the upper and lower confidence intervals of our estimate to the decrease in spending or increase in saving. We find that for every 1 MXN in decreased spending or increased saving, we can rule out increases or decreases in borrowing cost of more than 0.2 to 0.3 cents MXN.[2] For comparison, the average

[1] We illustrate the potential pitfalls of overfitting that the causal forest approach overcomes, by comparing our results with the estimates for individuals who belong to the strata blocks with the largest observed treatment effects.

[2] MXN stands for Mexican Pesos. As of the end of 2019, 1 MXN corresponded to 0.107 USD

credit card interest rate is 35.2% 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 cents in one month and compound to 41 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 6 months prior to 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 0.5 to 0.7 cents. 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 conclude that the observed changes in spending, savings, and credit card debt are more consistent with models of mental accounting than transaction convenience. In turn, we provide four additional pieces of evidence to support this conclusion. First, as mentioned, we are not able to predict any treatment effect heterogeneity on borrowing. Second, we show that individuals who co-hold are among the most responsive to the saving nudge. Third, one of the savings message 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. Finally, the reduction in spending appears discretionary, as ATM withdrawals are reduced more than credit and debit card spending. We thus argue that, while transaction convenience and credit limit variability are very likely motives for some subjects to co-hold at baseline, mental accounting is likely the primary driver for a large fraction of individuals who co-hold and respond strongly to the saving nudge.

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 conduct a randomized controlled trial on a large scale 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 latest machine learning techniques for causal inference.

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 transaction demand for liquidity may contribute to it. Maki (2002) studied whether households run up credit card debt strategically 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 liquid-

PPP, based on OECD coversion rates available at https://www.oecd-ilibrary.org/economics/data/aggregate-national-accounts/ppps-and-exchange-rates_data-00004-en.

ity premium. In Telyukova (2013), this premium is generated in a transaction convenience model and in Fulford (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 mental accounting drive the co-holding. Consistent with our findings, Gathergood and Weber (2014) provide evidence in favor of limited financial literacy among those who co-hold and, more recently, Gathergood and Olafsson (2020) show that while co-holding is less prevalent in Iceland than previously thought, it is more consistent with mental accounting than with rational explanations.

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, 2020; Akbaş 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). To the best of our knowledge, only two research papers have examined the effects of these nudges on other positions of household balance sheets, such as borrowing: Beshears et al. (2019) and Chetty et al. (2014).

In both of these studies, credit card borrowing is measured through snapshots (through biannual credit card balances from a credit bureau or through 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 as well as 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 in interest is charged. Furthermore, we focus on a different type of nudge than Beshears et al. (2019) and Chetty et al. (2014). They study the consequences of automatic enrollment, while we study the consequences of informational nudges. This softer intervention has become widely popular and effective (Halpern, 2015), and studying unintended effects of informational nudges is important for both public and private stakeholders.

We are also adding 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, 2020), 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 allow 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. 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.

2 Conceptual Framework

In this section, we will formally discuss two categories of 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. These two categories of models cover the three most common explanations for the co-holding puzzle found in the literature: transactions-convenience (Telyukova, 2013), credit limit variability (Fulford, 2015; Gorbachev and Luengo-Prado, 2019; Druedahl and Jørgensen, 2018), and mental accounting (Bertaut et al., 2009; Vihriälä, 2019). Appendix C 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. Furthermore, by adhering to a straightforward approach, we aim to demonstrate the broad applicability of our findings.

2.1 Liquidity premia: transactions-convenience and credit limit variability

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 due to variability in liquidity needs and a precautionary motive. This cash can be used for future consumption and individuals may smooth consumption by borrowing, effectively paying an insurance premium in the form of credit card interest.

To illustrate the intuitions of this class of models, we use a simple model with two periods, one consumption good, and log utility. The agent consumes in periods 1 and 2, denoted as $c_{1,2}$. In period 1, the agent may borrow to consume, denoted as b_1, because they must hold a certain

[3] In the context of microfinance, causal forests and other machine learning methods for causal inference have been used by Beaman et al. (2021) 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.

amount of cash, denoted as x, for transactions purposes or due to limited access to credit in the future.

The maximization problem is

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

subject to $x_1 - c_1 > x$. The parameter δ reflects the agent's patience. We assume the agent is not credit-constrained. The optimal solution for c_1^* is

$$\text{if } x_1 - c_1^* \geqq x \Rightarrow b_1 = 0 \text{ then } c_1^* = \frac{1}{\delta+1}x_1$$

$$\text{and if } x_1 - c_1^* < x \Rightarrow b_1 > 0 \text{ then } c_1^* = \frac{1}{\delta+1}x_1 - \frac{r}{(\delta+1)(1+r)}x \text{ and } b_1 = c_1^* - x_1 + x.$$

In this model, an agent may co-hold if they are patient and anticipate needing cash in the future. 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 liquidity premia model co-hold and become more patient, then they decrease their consumption and repay their debt by the same amount, i.e.,* $\frac{\partial b_1}{\partial \delta} = \frac{\partial c_1^*}{\partial \delta} < 0$.

Proof. See comparative statics with respect to δ when $x_1 - c_1^* < x$. All details and derivations can be found in Appendix C. □

To understand this proposition, we look at the comparative statics with respect to the patience to save, δ for both agents that optimally co-holds and agents that do not:

If $x_1 - c_1^* \geqq x$ (no co-holding) then $\frac{\partial c_1^*}{\partial \delta} = -\frac{1}{(\delta+1)^2}x_1 < 0$ and if $x_1 - c_1^* < x$ (co-holding) then $\frac{\partial c_1^*}{\partial \delta} = \frac{-r}{(\delta+1)(1+r)} < 0$.

If $x_1 - c_1^* \geqq x$ (no co-holding) then $\frac{\partial b_1}{\partial \delta} = 0$ and if $x_1 - c_1^* < x$ (co-holding) then $\frac{\partial b_1}{\partial \delta} = \frac{\partial(c_1^* - x_1 + x)}{\partial \delta} = \frac{\partial c_1^*}{\partial \delta} < 0$.

The intuition behind this result is straightforward: if agents become more patient, they will consume less and save more. However, the desire to insure against future cash needs does not change, and therefore x remains unchanged. If agents co-hold and have credit card debt, they will want to reduce existing debt instead of increase their cash holdings since credit card debt is more expensive. As a result, they will consume less and borrow equally less, i.e., $\frac{\partial b_1}{\partial \delta} = \frac{\partial c_1^*}{\partial \delta} < 0$.[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 card. This leads us to the following proposition:

Proposition 2. *If agents in the liquidity-premia model co-hold and their cash needs, x, increase, then they will increase their debt by almost the same amount (a bit less due to the increased costs of interest they pay), i.e.,* $\frac{\partial b_1}{\partial x} = 1 - \frac{r}{(\delta+1)(1+r)}$.

Proof. See comparative statics with respect to x when $x_1 - c_1^* < x$. All details and derivations can be found in Appendix C. □

To understand the proposition, we look at the comparative statics for consumption and debt with respect to the cash needs, x:

If $x_1 - c_1^* \geqq x$ (no co-holding) then $\frac{\partial c_1^*}{\partial x} = 0$ and if $x_1 - c_1^* < x$ (co-holding) then $\frac{\partial c_1^*}{\partial x} = \frac{-r}{(\delta+1)(1+r)} < 0$.

If $x_1 - c_1^* \geqq x$ (no co-holding) then $\frac{\partial b_1}{\partial x} = 0$ and if $x_1 - c_1^* < x$ (co-holding) then $\frac{\partial b_1}{\partial x} = 1 + \frac{\partial c_1^*}{\partial x} > 0$.

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, but 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, while an increase in the required amount of cash holdings leads to an increase in borrowing.

Credit-Limit-Chasing

A related class of models assume 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

[4] A microfounded modeling approach for a transaction 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.

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 liquidity-premia models 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 and cash holdings unchanged.

In Appendix C, 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.

Mental Accounting

To illustrate the intuitions of the mental-accounting or dual-spouse models (Bertaut et al., 2009; Vihriälä, 2019), we now outline a simple model with two periods, one consumption good, and log utility. The agent can consume in periods 1 and 2, denoted by $c_{1,2}$. In period 1, the agent can borrow to consume, denoted by b_1. In addition, in some period 0, the agent's previous patient self and/or their patient spouse locks away a certain amount, x, in an inaccessible savings account that is only for period 2 consumption. A fraction a of this amount is forgotten or hidden from the impatient self/spouse. In turn, the maximization problem of the impatient self/spouse in period 1 is

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

and the optimal solution for c_1^* is

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

$$\text{and if } x_1 - x - c_1^* < 0 \Rightarrow b_1 > 0 \text{ then } c_1^* = \frac{1}{\beta+1}x_1 - \frac{r+a}{(\beta+1)(1+r)}x \text{ and } b_1 = c_1^* - x_1 + x.$$

The parameter β reflects the agent's patience. We assume the agent is not credit-constrained.

Now, the amount withheld, x, is decided by a more patient self/spouse ($\delta > \beta$ is governing their patience) in period 0 with the following problem

$$\text{if } x_1 - x - c_1^* \geqq 0 \Rightarrow b_1 = 0$$

$$\text{then } max\{log(\frac{1}{\beta+1}(x_1 - ax)) + \delta log(x_1 - \frac{1}{\beta+1}(x_1 - ax))\}$$

$$\Rightarrow x^* = \frac{(\delta-\beta)}{(\delta+1)a} x_1 > 0 \text{ as } a, \beta \in [0,1] \text{ and } \delta > \beta$$

this equals $x = \frac{\delta}{\delta+1} x_1$ if $a = 1$ and $\beta = 0$ (all cash can be hidden and the other self/spouse is perfectly impatient). In this case, the final consumption of the impatient party in both periods is exactly the same as what a rational agent in a unitary household would allocate to first and second period consumption.

$$\text{Instead if } x_1 - x - c_1^* < 0 \Rightarrow b_1 > 0$$

$$\text{then } max\{log(c_1^*) + \delta log(x_1 - c_1^* - rb_1)\}$$

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

this equals $x = x_1 \frac{1}{1+\delta}(\delta - r)$ if $a = 1$ and $\beta = 0$ (all cash can be hidden and the other self/spouse is perfectly impatient). In this case, the final consumption of the impatient party in both periods is similar to what a rational agent in a univariate household would decide. It is not exactly the same because the patient party takes the interest costs into account that the impatient self/spouse incurs. The patient party knows that the impatient party can borrow to increase their consumption in period 1, and the impatient party will do so. This problem becomes more relevant when the impatient party knows about a fraction of the hidden cash. Therefore, if the interest costs increase, the patient self or spouse will hide less cash, i.e., if $r \uparrow$ then $x \downarrow$.

Now, suppose a nudge increases the patient self's patience, we can summarize the theoretical take-away in the following proposition.

Proposition 3. *If agents in the mental-accounting model co-hold and the patient self becomes more patient, then they increase their hidden assets, i.e., if $a \in (0, 1]$ then $\frac{\partial x}{\partial \delta} > 0$.*

Proof. See comparative statics with respect to δ when $x_1 - x - c_1^* < 0$. All details and derivations can be found in Appendix C. □

To understand this proposition, we now look at the comparative statics with respect to the discount factor of the patient party δ:

If $x_1 - x - c_1^* \geqq 0$ (no co-holding) then $\frac{\partial x}{\partial \delta} = \frac{a(\delta+1) - a(\delta-\beta)}{((\delta+1)a)^2} x_1 > 0$ and if $x_1 - x - c_1^* < 0$ (co-holding) then $\frac{\partial x}{\partial \delta} = \frac{1}{(1+\delta)}(-x + x_1 \underbrace{\frac{1+r}{r+a}}_{>1}) \in (\underbrace{\frac{1}{1+\delta}(-x + x_1)}_{>0 \text{ if } x < x_1}, \frac{1}{1+\delta}(-x + x_1 \frac{1+r}{r})]$ if $a \in (0, 1]$.

The intuition is the following: if the patient self/spouse becomes more patient, then they withdraw more money especially if they can hide it.

A change in the patient self's patience, δ, changes the hidden cash x. Therefore, we now look at comparative statistics with respect to the hidden cash x. Our main theoretical take-away can be summarized in the following proposition:

Proposition 4. *If agents in the mental-accounting model co-hold and the patient self increases the hidden assets, then the impatient agent consumes less, especially when more of the assets are hidden, i.e.,* $\frac{\frac{\partial c_1^*}{\partial x}<0}{\partial a} < 0$*. If the agent is very impatient and all assets are hidden,* $\beta = 0$ *and* $a = 1$*, they decrease their consumption by the same amount as the hidden assets and their borrowing is unchanged, i.e.,* $\frac{\partial c_1^*}{\partial x} = -1$ *and* $\frac{\partial b_1}{\partial x} = 0$*.*

Proof. See comparative statics with respect to x when $x_1 - x - c_1^* < 0$. All details and derivations can be found in Appendix C. $\square$

To understand this proposition, we look at the comparative statics with respect to x:

If $x_1 - x - c_1^* \geqq 0$ (no co-holding) then $\frac{\partial c_1^*}{\partial x} = -\frac{1}{\beta+1}a < 0$ and if $x_1 - x - c_1^* < 0$ (co-holding) then $\frac{\partial c_1^*}{\partial x} = -\frac{r+a}{(\beta+1)(1+r)} < 0$ and $\frac{\frac{\partial c_1^*}{\partial x}}{\partial a} < 0$ in both cases.

If $x_1 - x - c_1^* \geqq 0$ (no co-holding) then $\frac{\partial b_1}{\partial x} = 0$ and if $x_1 - x - c_1^* < 0$ (co-holding) then $\frac{\partial b_1}{\partial x} = \frac{\partial c_1^*}{\partial x} + 1$.

To understand the intuition, we will describe two mechanisms. First, when additional cash can be hidden, i.e., the separation-or-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, i.e., $\frac{\partial c_1^*}{\partial x} < 0$. Now, if accounts can be fully separated and cash can be hidden, $a \to 1$, the sensitivity of consumption to x becomes more negative, $\frac{\partial c_1^*}{\partial x} \in (-\frac{r+1}{(\beta+1)(1+r)} < -\frac{r}{(\beta+1)(1+r)}]$ for $a \in (0, 1]$. The sensitivity of consumption to x is negative because the impatient agent consider 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., $a = 0$), they would take the interest costs that borrowing incurs into account.

If $a = 1$ and $\beta = 0$ (full separation of accounts, i.e., all cash can be hidden, and the agent is perfectly impatient), then the sensitivity of consumption to hidden cash x equals -1, $\frac{\partial c_1^*}{\partial x} = -1$. In turn, when the agent's consumption goes down a lot in response to an increase in x then the effect on borrowing is zero, if $\beta = 0$ and $a = 1$, then $\frac{\partial b_1}{\partial x} = 0$, i.e., borrowing does not respond at all.

In summary, the class of mental-accounting models predict that, when the patient agent becomes more patient and hides more cash, then the impatient agent does not respond with borrowing more or less.

3 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 achieve this objective, we utilize data from a large-scale field experiment conducted in Mexico.

3.1 The Mexican Credit Card Market: Basic Facts

The credit card market in Mexico is dominated by the five largest banks, with Banorte being one of the largest providers. As of June 2017, there were 17.9 million active credit card accounts with a positive balance, on a population of 124 million. The five largest banks control 85% of the market, with the top two products accounting for over 25% and the top six products covering just over 50%. Credit card debt represents 22% of the consumer credit portfolio, including mortgage debt, as of the end of 2015.[5] The average number of credit cards per cardholder are 1.27 cards, according to a nationally representative survey conducted in 2018. Among individuals with at least one credit card, 79% have only one card, 15% have two, and the rest have more than two cards.[6] Interest rates on credit cards in Mexico are high compared to those in the US, with the average rate in 2017 being 26.4% above the federal short-term interest rate of 7.17%.

3.2 Experiment: Sample Population and Experimental Treatments

We analyze data from a large-scale field experiment conducted by the bank, which involved 3,054,438 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 3,054,438 customers, a random sample of 374,893 clients was

[5]Banco de Mexico, multiple reports, including: https://www.banxico.org.mx/publicaciones-y-prensa/rib-tarjetas-de-credito/rib-tarjetas-credito--tasas-i.html and https://www.banxico.org.mx/publicaciones-y-prensa/reportes-sobre-las-condiciones-de-competencia-en-l/%7B9A9ADEB4-7D4E-8307-B645-DB78A8A91ADE%7D.pdf

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

[7]Payroll accounts are a type of deposit account in Mexico that are specifically designed for employees to receive their paychecks. They are commonly offered by banks 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, the data on salaries is likely to be very accurate.

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, which had previously proven to be effective in experiments run by the bank. 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 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.

[8]The first set of messages was sent on Friday, September 13th, the same day they were scheduled. However, 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.

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

3.3 Descriptive Statistics

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 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 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 1 presents descriptive statistics for individuals in both the treatment and control groups. The average age of these individuals is 44 years, with an average monthly after-tax income of approximately 13,508 MXN (1,445 USD), and an average of 7 years of banking history with Banorte. Their average checking account balance is approximately 18,123 MXN, and around 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. On average, these individuals have more average income and higher average checking account balances than the sample of all clients. Their average credit card balance is 17,998 MXN (1,926 USD), with a median of 10,458 MXN (1,119 USD). The average monthly interest paid by individuals with credit cards is 266 MXN (28 USD). Note that, this average includes individuals who do not pay any interest. Furthermore, individuals with credit cards have substantial borrowing capacity, with an average limit of 83,801 MXN (8,978 USD) and a median of 45,000 MXN. All continuous variables are winsorized at the 2nd and 98th percentiles. Credit card related variables (namely interest charges, credit card balances and credit limits) are winsorized based on only individuals who have a credit card.

3.3.1 Prevalence of Co-Holding

Table 2 provides information on credit card interest payments, checking account balances, credit card balances, and interest payments for individuals with credit cards, by deciles of checking account balances over income. The sample is limited to individuals who have a credit card. The table

[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).

shows that even among individuals in the higher deciles of checking account balances, a significant fraction, ranging from 17% to 24%, pay credit card interest. We observe that the top 30% of individuals with the highest checking account balances could pay off their entire credit card debt and save approximately 700 MXN per month (75 USD), as Banorte's average credit card interest rate is 35.2% and the return on checking accounts is 0%. These potential savings correspond to around 5% of individuals' monthly incomes.

Focusing on individuals who roll over credit card debt, we define the co-holding puzzle group as those who have more than 50% of their income held in their checking accounts and are paying credit card interest. Specifically, we require that the minimum checking account balance observed over the last 6 months be equal to or greater than 50% of their income. Approximately 13.5% of individuals who are paying credit card interest are in the puzzle group. In Table B2, we compare individuals in the puzzle group to the rest of those who are paying credit card interest. The puzzle group is slightly older, but they have similar monthly income and slightly longer tenure with the bank.[12] The two groups differ mainly in their checking account and credit card balances.

It is also worth noting that individuals appear to carry their debt persistently, with an 80% correlation between rolling over debt in any given month and doing so in the next month.

3.3.2 Randomization Checks

Table 3 shows that the randomization process was successful in balancing observable characteristics across treatment and control groups.[13] There are no significant differences in the characteristics of individuals assigned to the treatment and control groups.

3.4 Methodology to Analyze Experiment

In our analysis of the experiment, we first analyze the impact of the saving nudges on spending, saving, and borrowing outcomes using regression models to compare means between the treatment and control groups. We conduct this analysis for the entire population, as well as for subgroups of individuals who had a credit card or paid interest at baseline. In addition, we investigate heterogeneity in the response to the saving nudges on the three main outcomes of interest using causal forests.

[12] For more details on the differences, refer to Appendix B.1.

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

3.4.1 Outcome Variables

Our main variables of interest are spending, saving, and borrowing. 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 ones as well as those initiated by user rules) to capture all potential responses to the treatment.

To measure saving, 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.[14]

To have a monthly measure of spending and borrowing, our main results are based on treatment effects over the first month, i.e., the first four weeks, of the intervention. In Appendix B.5, we show that the week-by-week treatment effects over the 7-weeks of the treatment are not statistically significantly different from each other. Thus considering the four weeks of the intervention is akin to a monthly normalization.

We measure the cost of credit card borrowing by looking at credit card interest charges. 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, if they do not repay the statement balance in full by the due date, they are charged interest on the entire average daily balances observed during the corresponding billing cycle.[15]

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

[14] 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.

[15] 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.

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, i.e., the cycles that end 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 individual-level due dates, we identify individuals with a billing cycle fully covered by the intervention and for them we look at saving and spending over the weeks that overlap with that billing cycle (Section 4.5). In addition, Appendix 4.5 shows the distribution of due dates over the month. Due dates are pretty evenly distributed over the month, which is reassuring.

We are also interested in measuring the amount of debt that individuals rolled 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.[16]

We also calculate the actual interest-bearing credit card balances for each individual by dividing the credit card interest by the monthly interest rate. 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 two billing cycles that intersected with the treatment period (September and October) and the subsequent two billing cycles.

For all continuous variables that are non-negative, we take the natural logarithm of one plus the variable in our main analysis.[17] To alleviate concerns about log specifications, in Appendix B5 we replicate the analysis expressing the outcome variables in levels (MXN pesos).

[16]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.

[17]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 3.2.

3.4.2 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 (1):

$$Y_i = \alpha_s + \beta * treatment_i + \epsilon_i \tag{1}$$

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.

3.4.3 Heterogeneous Effects

We then analyze heterogeneous treatment effects of the intervention on spending, saving, and borrowing outcomes for two reasons. First, prior research has suggested multiple explanations for the co-holding or credit card debt puzzle, and we showed that these different explanations have distinct implications for how people respond to the saving nudge. It's possible that certain explanations are more relevant to certain individuals, which would result in heterogeneity in their responses to the treatment. This heterogeneity could thus provide valuable information about the diverse reasons why people co-hold. Second, from a policy perspective, it is important to consider the unintended consequences of encouraging savings, such as an increase in borrowing among individuals who save the most.

However, identifying individuals with the largest response to a treatment is a task prone to overfitting. The standard practice consists of interacting a variable of interest with the treatment indicator. The resulting coefficient is indicative of larger or smaller effects for the corresponding group. However, if there are multiple variables of interest, the number of interactions (and groups) increases exponentially. When sorting based on those coefficients, the largest treatment effects may correspond to groups that experience idiosyncratic shocks that lead to spuriously large treatment effects (overfitting). To address this issue, we use the causal forest algorithm which is designed to predict individual treatment effects for a wide range of sub-populations, without risking invalid inference due to overfitting. It accomplishes this through three distinctive features: sample splitting, orthogonalization, and optimization of an objective function that captures treatment effect heterogeneity. In Appendix A, we provide a detailed description of the method.

Theoretically, individuals might respond in their spending, saving, and/or borrowing, we thus estimate separate forests for these three relevant outcome variables (described in Subsection 3.4.1).

4 Results

4.1 Aggregate Effects on Spending, Saving, and Borrowing

Table 4 presents the aggregate treatment effects on spending, saving, and credit card borrowing. The 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 0.9% reduction from a base of 16,732 MXN. Column (2) shows a significant 0.6% increase in savings from a base of 17,394 MXN.

We now move on to individuals with credit cards and credit card debt (Panels B and C). 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 0.4% and any decreases greater than 1.2% from a starting point of 214 MXN. This means that for every 1 MXN reduction in spending, the corresponding change in credit card borrowing costs is not more than 0.006 MXN (0.004*214/(16732*0.009)) and not less than 0.017 MXN. In the aggregate, we thus find that the majority of the reduction in spending is reflected in increased savings.

The reason for 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 have been responsive to the intervention. The inclusion of such individuals was necessary to successfully train a model to predict treatment effect heterogeneity. 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.

4.2 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 on only variables with a variable importance larger than 1%, following the approach of Athey and Wager (2019) for feature selection.[18] To test for treatment effect heterogeneity, we follow Burke et al. (2020) and perform the calibration test of Chernozhukov et al. (2018). Here, we fit a linear model to residualized outcomes using the average of individual forest predictions and the difference between each prediction and the average prediction as the only two regressors. The main coefficient of interest is the estimate for the differential forest prediction, which is considered evidence of treatment effect heterogeneity if it is positive and significant.

We present the results in Table 5. 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 show any 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 with variable importance greater than 1% from the causal forest model used in Column (2) for saving. In Column (6), the model was trained using only the variables with variable importance greater than 1% from the causal forest model used in Column (3) for borrowing. The coefficients for the differential forest prediction for borrowing are not significantly different from zero in any of these specifications. Therefore, based on the variables considered in the analysis, we cannot find evidence of treatment effect heterogeneity in borrowing outcomes.

Based on the evidence of treatment effect heterogeneity in spending and saving, we now aim to determine the magnitude of the differences in treatment effects across the population. We visualize the distribution of predicted treatment effects in Figure 1.

We then split individuals into quartiles of predicted treatment effects and for each quartile we calculate actual treatment effects, using a cross-fitted ranking over five folds.[19] For exposition, the top quartile of predicted treatment effects on spending contains the most negative predictions whereas the top quartile of predicted treatment effects on savings contains the most positive predictions. Figure 2 displays 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

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

[19] More specifically, 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 variable observed in each fold are not used when assigning observations in that same fold to a specific quartile.

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.

Table 6 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). The table shows that there is a significant overlap between the two, with around 14% to 16% of individuals in both the top quartiles of predicted treatment effect for spending and saving. If the overlap was perfect, the proportion would be 25%.

4.3 Treatment Effects on Spending, Saving, and Borrowing for the Top Quartile of Predicted Treatment Effects

In Subsection 4.2, we demonstrated that there is strong heterogeneity in the spending and saving responses to the treatment. We also observed that, for individuals outside of the top quartile of the predicted treatment effect distribution, the nudge did not lead to significant changes in behavior. Therefore, for this group of individuals, the question of how reductions in spending and increases in saving relate to borrowing decisions is not well defined. Consequently, we now shift our focus to investigate the magnitude of the treatment effects among the subset of individuals in the top quartile of predicted treatment effects who have a credit card. This group exhibits a significant change in spending, providing an opportunity to test the hypotheses formulated in the models developed in Section C.

If individuals became more patient or cautious about their future cash needs in response to receiving saving messages, both the liquidity-premia and mental-accounting 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 mental-accounting models, borrowing remains unchanged.

Table 7 displays the treatment effects on spending, saving, and borrowing for individuals in the top quartile of the distribution of predicted treatment effects on spending. Panel A includes all individuals who own a credit card, while Panel B focuses on individuals who pay credit card interest.

We will first discuss the results in Panel A. Column (1) shows the spending outcomes for credit cardholders in the top quartile of the predicted treatment effect distribution. The estimated reduction in spending is 7.82% on a baseline of 33,485 MXN, which corresponds to a decrease of 2,619 MXN. Column (2) indicates that almost all of the decrease in spending is attributable to an increase in checking account balances by 5.08% on a basis of 41,463 MXN, equivalent to an increase of 2,106 MXN.

In Column (3), we analyze interest payments for the two billing cycles affected by the intervention, and we observe a decrease of 0.71% from a basis of 207 MXN, with a standard error of 1.76%. 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 8.63 (5.68) MXN can be ruled out. Dividing this value by the decrease in spending, we can state with 95% statistical confidence that at most 0.0033 (0.0022) 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*0.3756/12 = 0.0313.[20] Thus, we can conclude with statistical confidence that less than 10% of the reduction in spending is reflected in reduced credit card interest payments.

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 somewhat less tightly estimated, we can see that the treatment does not seem to have a significant positive or negative impact on this measure of rolled-over debt.

Moving on to Panel B, the same regressions are run for the subset of individuals that 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 0.71 (0.53) MXN.

[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 (207.37/6625.23)*12 = 0.3756 (see Tables 7 and B3). As mentioned earlier, when considering the entire experimental pool the interest rate is 35.4%.

Table B3 in Appendix B.3 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, and the findings are very similar to the ones on paid credit card interest, indicating that less than 10% (21%) of the reduction in spending is reflected in interest-bearing balances (for individuals that 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 to 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. 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 10% (and 6% 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. Table B4 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. Again, we find very similar results.

Tables 8 and B4 shows our main Tables 7 and B3 but for individuals in the top quartile of predicted treatment effects on saving instead of spending. The results are very similar. This table is effectively showing our program evaluation results. As in Table 7, 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 0.29 cents and a decrease of more than 0.48 cents for each extra 1 MXN in saving. The same is true for the other outcome variables such as balances and repayments.

4.4 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, as shown in Table 7, 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. As we discussed, the literature proposes two main explanations for this behavior.

As we formally showed, these explanations make different predictions about the joint responses

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

of spending, saving, and borrowing to the saving messages. When the messages are interpreted as either a shock to preferences or a shock to cash needs or hidden assets, our empirical findings align with the class of mental-accounting models. People tend to view their spending, saving, and borrowing as distinct accounts, so the saving messages affect spending and saving, but borrowing remains unaffected.

We now present four additional pieces of evidence that support a preference-based explanation for the credit card debt puzzle.

Firstly, as outlined in Subsubsection 3.4.3 and Table 5, our analysis reveals that the variables examined in our study are unable to predict any treatment effect heterogeneity in borrowing behavior. 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 self or spouse.

Secondly, Figure 3 illustrates the distribution of individuals with credit cards who pay interest and co-hold debt, across quartiles of predicted treatment effects. We observe that the majority of co-holding individuals are also in the highest quartile of predicted spending and saving treatment effects. This reinforces the idea that co-holding debt is a psychological mechanism for exercising self-control that makes individuals more receptive to saving nudges as well.

Thirdly, we can look at how deposits as well as different types of spending, ATM withdrawals, credit/debit card purchases, or transfers, are affected. The results are shown in Table 9 and indicate that individuals increase their saving primarily by cutting down their discretionary spending, i.e., ATM withdrawals which likely reflects smaller everyday items that individuals are tempted to buy here and there.

Finally, 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 mental accounting, wherein individuals constrain themselves to save more. We now further discuss the varying treatment effects of different messages.

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 149,561 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 the spending, saving, and borrowing of these individuals. This approach allows us to compare the effects of each message specifically for those individuals who are most likely to respond to the nudge.

Table 10 presents the borrowing, saving, and spending effects of each individual message. It shows that the reduction in spending and increase in savings effect is large and significant for all individual messages, except for Message 1, which has a non-significant effect (the messages are displayed in Subsection 3.2). 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 savings and goal-reaching. The pairwise comparison between the short-term and long-term messages shows no statistically significant difference in treatment effects, suggesting 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 than both the short-term and the long-term message groups. This message emphasizes the safeness saving on a bank account and reaching general goals, which is consistent with the mental accounting hypothesis and the idea of constraining oneself to save more. Finally, it is noteworthy that in the specification that restricts to individuals that received Message 4 as well as in all others, we find a tightly estimated zero effect on borrowing.

4.5 Discussion and Robustness

Alternative Outcome Variables and Specifications

Table B5 and B6 repeat the main analysis expressing the dependent variables in MXN, instead of the natural log, to address the concerns in Cohn et al. (2022). The results are consistent with our previous findings likely due to the low number of zeroes in our sample and the orthogonality of the treatment.

Credit Constraints

It is possible that individuals do not increase their credit card borrowing because they are credit constrained. Table B7 and B8 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 would 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 to be 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 B9 and B10 show the saving and borrowing results for this group of individuals. We find very similar results.

Prediction Error and Persistence of Credit Card Debt

Predicted treatment effects estimate actual treatment effects with error, and it is possible that some individuals with a large predicted treatment effect on saving may not respond to the nudge. To rule out the possibility that the null effects on borrowing outcomes shown before are driven by individuals for whom the predictions of the causal forest are not accurate, we investigate the relation between “prediction errors" and the treatment effect of the intervention on borrowing outcomes. In contrast to standard prediction exercises, we do not observe individual level prediction errors since actual treatment effects are never observed at the individual level. We define “prediction errors" at the group level as the difference between the simple average of individual-level predicted treatment effects of observations in a given group and the (average) treatment effect of observations in the same group, calculated ensuring covariate balance.

We implement the same five-fold cross-fitted procedure described in Subsection 4.2 to assign observations in the top quartile of predicted treatment effects to ten decile groups.[22] For each group, we focus on individuals who had a credit card and paid interest at baseline and calculate the corresponding prediction error on saving and the average treatment effect of the intervention on credit card interest. Figure B2 shows a scatter plot of these two variables. We can see that, as expected, prediction errors are uncorrelated with treatment effects on borrowing outcomes. The prediction errors are thus the result of noise, which is uncorrelated with the treatment.

We also note that, while individuals who paid interest on their credit cards during the baseline period have a 80% probability of incurring interest during the treatment period, it is possible that the treatment effect on spending, saving, and borrowing is driven by those for whom interest payments were not persistent. To investigate this possibility, we examine the correlation between the fraction of individuals actually paying credit card interest during the treatment period and the magnitude of the treatment effect on saving, across different groups of observations. As before, we split individuals in the top quartile of predicted treatment effects into decile groups. For each group, we focus on individuals who had a credit card and paid interest at baseline, and then calculate the treatment effect on checking account balances. Figure B3 shows a scatter plot of these two variables. We can see that there is no clear relationship between them, suggesting that indeed, individuals increased their savings as a result of the nudge independent of whether or not they carried credit card debt.

[22] Specifically, we rank observations into 40 groups, and focus the analysis on the top 10 groups, which we interpret as deciles within the top quartile of the distribution.

Characterizing Individuals with the Largest Predicted Responses to the Treatment

To characterize individuals with the largest predicted responses to saving nudges, Table A1 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 older 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 are proxies for trust and usage of the bank. Therefore, the algorithm appears to just 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, it is possible that individuals in this group are in better financial positions which allow 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. It is then possible that other characteristics that we do not observe, such as higher financial literacy, frugality, or patience (among others) may jointly lead to higher baseline balances and and to a higher propensity to respond to saving nudges. If this is the case, then nudges would be more effective in the intensive margin, successfully leading to more savings among those who were already saving, but less so in the extensive margin to turn non-savers into savers. In Figure A1, we can also see that checking account balances carry 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).

Dynamics of Treatment Effects on Spending, Saving, and Borrowing

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 B5 and B6, we can see the weekly point estimates of the treatment effect on spending and saving, for all individuals that 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. Overall, we can see that the coefficients are similar across all treatment weeks. It does not seem to be the case that the treatment effects vary over the course of the intervention. This finding confirms the validity of our main outcome measure, spending and saving over the first month of the treatment. We use spending and saving of the first four weeks of the treatment, in order to have a monthly measure that we can compare to the monthly credit card interest charges.

Overlap of Treatment Period and Billing Cycles

We note that since each individual has a different due date 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 the null effect on credit card borrowing could be driven by this mismatch, we run our main specifications considering only individuals for whom one entire billing cycle is covered by the treatment. For them, we look at treatment effects on spending and saving over the weeks that are perfectly aligned with their individual billing cycle (the one that was fully covered by the intervention). The results can be found in Table B11. We also show in Figure B4 that individuals' billing cycle end dates are pretty evenly distributed over the calendar month.

4.6 Lessons for Heterogeneity Analysis in Academic and Policy Settings

We use causal forests (Athey et al., 2019) to study how treatment effects vary across different sub-populations. This approach allows us to obtain accurate estimates of the treatment effects of the intervention for various sub-groups and to identify the sub-group that benefits the most from the intervention without worrying about overfitting. In contrast, traditional methods of identifying heterogeneous treatment effects rely on heterogeneity across pre-specified strata groups. We now discuss the lessons we learn when comparing the two approaches.

Comparisons of Limited Power: Heterogeneity by Experimental Strata

One way to analyze heterogeneous treatment effects is to split the sample based on strata from the experimental design. However, this approach has limitations because it 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 11, we show the treatment effects on saving across the experimental strata and find limited heterogeneity across the sub-populations that were pre-selected for heterogeneity analysis before the experiment was run. We only find a significant treatment effect for individuals with pre-treatment checking account balances in the top quartile, who had the largest treatment effect of a 1.3% increase in savings (this coefficient corresponds to the intercept, -0.004, plus the coefficient on the group of people with the largest checking account balances, 0.017).

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 the treatment effects for each strata block.[23] However, this approach can lead to substantial bias if not adjusted for overfitting, as

[23] 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

we now show.

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 compute treatment effects 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 treatment effects on checking account balances and credit card interest for this top quartile.

Table 12 displays the results of this analysis. Columns (2) to (4) report treatment effects for individuals in strata blocks with the largest observed treatment effects. The findings reveal significant reductions in spending, of 44%. Additionally, these individuals experienced substantial increases in savings and decreases in borrowing, as evidenced in Columns (3) and (4).

In comparison, the results obtained from the causal forest are shown in Columns (5) to (8). Column (5) displays the number of observations included in this part of the analysis. Column (6) shows that the increases in spending is smaller, in the order of 8%. Columns (7) and (8) show the corresponding treatment effects on saving and borrowing. These estimates are not subject to overfitting bias. The coefficients in Column (8) are significantly closer to zero than those in Column (4).

Policy and Business Applications

This analysis has important implications for policy-making. 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 policy-makers 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.

Firstly, if the group they identify is based on a coarse partition of the covariate space, such as low-income individuals, they may end up treating individuals who do not need the money because they have high liquid wealth despite their low income. Alternatively, if they split the sample over many dimensions such as income, wealth, education, age, etc., then they may overfit the data,

heterogeneous treatment effects with a rich set of explanatory variables without considering the risk of overfitting.

resulting in identifying a subpopulation that has a seemingly high spending rate that would not replicate out of sample.

From a business perspective, being able to estimate individual treatment effects enables firms to maximize profit 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, the most promising partition is quartiles of checking account balances, which would yield an average increase in balances of 203.3 million MXN (1.3% from a base of 20,893 MXN for 750K clients) by sending messages to 750,000 clients (see Column (1) of Table 11).

However, treating only the top 5% of individuals in the distribution of predicted treatment effects would lead to a reduction in costs of 80% by sending messages only to 150,000 clients who are more likely to respond. For the top 5% of individuals in the distribution of predicted treatment effects we find a 1,056.82 MXN increase in balances, which aggregate to 158.52 million MXN (150,000*1,056.82 MXN) which leads to a more efficient cost-revenue ratio. With a causal forest, firms can sort clients based on their predicted revenue and treat them until marginal revenues equal marginal costs.

5 Conclusion

Our study analyzes 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, an important and widely debated topic in the field of household finance. Our results provide evidence that encouraging individuals who are paying credit card interest to save can lead to increased saving, regardless of existing levels of debt.

We then compare our findings with the predictions of existing models of co-holding. These models can be categorized into two groups: those that explain co-holding through a liquidity premium for cash, and those that attribute it to mental accounting or intra-household agency constraints.

When our saving messages are interpreted as a shock to patience, we show that the first class of models predict a decrease in spending that is fully reflected in a decrease in borrowing, while the second class of models predict that spending decreases, saving increases, and borrowing is unchanged. On the other hand, when our saving messages are interpreted as a shock to cash needs or hidden cash, the first class of models predict that spending is unchanged but saving is increased, which is fully reflected in an increase in borrowing. The second class of models again predict that spending decreases, saving increases, and borrowing is unchanged.

We then demonstrate that the results of our experiment are consistent with theories based on mental accounting, which posit that patient selves or spouses set savings aside in a non-fungible mental account, resulting in a tightly estimated null effect on borrowing.

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 in a cost effective way.

Overall, our study contributes to a deeper understanding of the complex interactions between spending, saving, and borrowing, and provides valuable insights for policymakers and financial institutions seeking to promote financial well-being among their clients.

References

Abadie, A., M. M. Chingos, and M. R. West (2018). Endogenous stratification in randomized experiments. *Review of Economics and Statistics 100*(4), 567–580.

Afzal, U., G. d'Adda, M. Fafchamps, S. R. Quinn, and F. Said (2019). Implicit and explicit commitment in credit and saving contracts: A field experiment. Technical report, National Bureau of Economic Research.

Akbaş, M., D. Ariely, D. A. Robalino, and M. Weber (2016). How to help the poor to save a bit: Evidence from a field experiment in kenya.

Allcott, H. and J. B. Kessler (2019). The welfare effects of nudges: A case study of energy use social comparisons. *American Economic Journal: Applied Economics 11*(1), 236–76.

Ashraf, N., N. Bau, C. Low, and K. McGinn (2020). Negotiating a better future: How interpersonal skills facilitate intergenerational investment. *The Quarterly Journal of Economics 135*(2), 1095–1151.

Athey, S., V. Hadad, N. Keleher, O. Medina, R. Nissan, R. Rosemberg, M. Schaelling, J. Spiess, and J. Wright (2021). Computational applications to behavioral science. Technical report.

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

Athey, S., J. Tibshirani, S. Wager, et al. (2019). Generalized random forests. *The Annals of Statistics 47*(2), 1148–1178.

Athey, S. and S. Wager (2019). Estimating treatment effects with causal forests: An application. *arXiv preprint arXiv:1902.07409*.

Beaman, L., D. Karlan, B. Thuysbaert, and C. Udry (2021). Self-selection into credit markets: Evidence from agriculture in mali. Technical report, National Bureau of Economic Research.

Benartzi, S., J. Beshears, K. L. Milkman, C. R. Sunstein, R. H. Thaler, M. Shankar, W. Tucker-Ray, W. J. Congdon, and S. Galing (2017). Should governments invest more in nudging? *Psychological science 28*(8), 1041–1055.

Bernheim, B. D., D. Ray, and Ş. Yeltekin (2015). Poverty and self-control. *Econometrica 83*(5), 1877–1911.

Bertaut, C. C., M. Haliassos, and M. Reiter (2009). Credit card debt puzzles and debt revolvers for self control. *Review of Finance 13*(4), 657–692.

Beshears, J., J. J. Choi, D. Laibson, B. C. Madrian, and K. L. Milkman (2015). The effect of providing peer information on retirement savings decisions. *The Journal of finance 70*(3), 1161–1201.

Beshears, J., J. J. Choi, D. Laibson, B. C. Madrian, and W. L. Skimmyhorn (2019). Borrowing to save? the impact of automatic enrollment on debt. Technical report, National Bureau of Economic Research.

Beshears, J. and H. Kosowsky (2020). Nudging: Progress to date and future directions. *Organizational Behavior and Human Decision Processes 161*, 3–19.

Breza, E., M. Kanz, and L. F. Klapper (2020). Learning to navigate a new financial technology: Evidence from payroll accounts. Technical report, National Bureau of Economic Research.

Burke, J., J. Jamison, D. Karlan, K. Mihaly, and J. 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, M., E. La Ferrara, and P. Pinotti (2022). Goals and gaps: Educational careers of immigrant children. *Econometrica 90*(1), 1–29.

Carroll, C. D. (1998). Why do the rich save so much?

Carvalho, L. S., S. Meier, and S. W. Wang (2016). Poverty and economic decision-making: Evidence from changes in financial resources at payday. *American Economic Review 106*(2), 260–284.

Chernozhukov, V., M. Demirer, E. Duflo, and I. Fernandez-Val (2018). Generic machine learning inference on heterogenous treatment effects in randomized experiments. Technical report, National Bureau of Economic Research.

Chetty, R., J. N. Friedman, S. Leth-Petersen, T. H. Nielsen, and T. Olsen (2014). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from denmark. *The Quarterly Journal of Economics 129*(3), 1141–1219.

Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick (2004). For better or for worse: Default effects and 401 (k) savings behavior. In *Perspectives on the Economics of Aging*, pp. 81–126. University of Chicago Press.

Choukhmane, T. (2019). Default options and retirement saving dynamics. In *112th annual conference on taxation. NTA*.

Cohn, J. B., Z. Liu, and M. I. Wardlaw (2022). Count (and count-like) data in finance. *Journal of Financial Economics 146*(2), 529–551.

Costa, D. L. and M. E. Kahn (2013). Energy conservation "nudges" and environmentalist ideology: Evidence from a randomized residential electricity field experiment. *Journal of the European Economic Association 11*(3), 680–702.

Davis, J. M. and S. B. Heller (2020). Rethinking the benefits of youth employment programs: The heterogeneous effects of summer jobs. *Review of Economics and Statistics 102*(4), 664–677.

Druedahl, J. and C. N. Jørgensen (2018). Precautionary borrowing and the credit card debt puzzle. *Quantitative Economics 9*(2), 785–823.

Dynan, K. E., J. Skinner, and S. P. Zeldes (2004). Do the rich save more? *Journal of political economy 112*(2), 397–444.

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

Gargano, A. and A. G. Rossi (2020). There's an app for that: Goal-setting and saving in the fintech era. *Available at SSRN*.

Gathergood, J. and A. Olafsson (2020). The co-holding puzzle: New evidence from transaction-level data.

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

Goldin, J., T. Homonoff, and W. Tucker-Ray (2017). Retirement contribution rate nudges and plan participation: Evidence from a field experiment. *American Economic Review 107*(5), 456–61.

Gorbachev, O. and M. J. Luengo-Prado (2019). The credit card debt puzzle: The role of preferences, credit access risk, and financial literacy. *Review of Economics and Statistics 101*(2), 294–309.

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

Halpern, D. (2015). *Inside the nudge unit: How small changes can make a big difference*. Random House.

Karlan, D., M. McConnell, S. Mullainathan, and J. Zinman (2016). Getting to the top of mind: How reminders increase saving. *Management Science 62*(12), 3393–3411.

Maki, D. M. (2002). The growth of consumer credit and the household debt service burden. In *The impact of public policy on consumer credit*, pp. 43–68. Springer.

Medina, P. (2020). Side-Effects of Nudges: Evidence from a Randomized Intervention in the Credit Card Market. *Review of Financial Studies*.

Mullainathan, S. and E. Shafir (2013). Decision making and policy in contexts of poverty. *Behavioral foundations of public policy*, 281–300.

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

Schilbach, F., H. Schofield, and S. Mullainathan (2016). The psychological lives of the poor. *American Economic Review 106*(5), 435–440.

Telyukova, I. A. (2013). Household need for liquidity and the credit card debt puzzle. *Review of Economic Studies 80*(3), 1148–1177.

Thaler, R. H. (1994). Psychology and savings policies. *The American Economic Review 84*(2), 186–192.

Vihriälä, E. (2019). Untangling the credit card debt puzzle. Technical report, Working Paper.

Wisdom, J., J. S. Downs, and G. Loewenstein (2010). Promoting healthy choices: Information versus convenience. *American Economic Journal: Applied Economics 2*(2), 164–78.

Zinman, J. (2007). Household borrowing high and lending low under no-arbitrage. *Dartmouth University*.

Figures and Tables

Table 1: Descriptive Statistics

All Individuals (N= 3,054,503)					
	Mean	Std dev	P25	P50	P75
Age (Years)	44.31	15.98	31.00	43.00	56.00
Monthly Income	13,508.46	13,101.24	6,116.67	9,866.88	15,005.78
Tenure (Months)	80.52	72.68	22.00	59.33	125.37
Monthly Spending	16,122.10	40,352.17	3,100.00	9,034.20	13,278.36
Checking Account Balance	18,122.86	50,830.78	729.00	2,295.69	10,402.39
Fraction with Credit Card	0.12	0.32	0.00	0.00	0.00
Monthly Credit Card Interest	31.53	128.88	0.00	0.00	0.00
Credit Card Balance	2,132.81	6,018.13	0.00	0.00	0.00
Ending Card Balance - Payments	585.86	704.27	0.00	0.00	0.00
Credit Card Limit	9,930.49	20,050.48	0.00	0.00	0.00
Individuals with a Credit Card (N= 362,223)					
	Mean	Std dev	P25	P50	P75
Age (Years)	41.82	12.47	33.00	42.00	53.00
Monthly Income	19,632.27	17,983.48	9,071.32	13,912.75	22,718.28
Tenure (Months)	102.71	72.29	43.27	86.43	148.53
Monthly Spending	28,532.08	65,871025	6,181.81	18,063.10	21,145.28
Checking Account Balance	32,212.66	69,364.31	1,581.29	5,157.02	23,069.07
Monthly Credit Card Interest	266.07	389.71	0.00	0.00	170.01
Credit Card Balance	17,998.39	29,741.04	104.21	10,457.89	27,137.36
Ending Card Balance - Payments	5,073.91	6,736.91	0.00	0.00	2,980.34
Credit Card Limit	83,801.60	108,109.54	15,000.00	45,000.00	100,000.00

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 3,054,438 customers is taken as a random sample from the universe of the bank's customers satisfying three requirements: individuals have 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 are salary transfers flagged as such in the payroll accounts. Tenure with the bank corresponds to the number of months since the accounts were opened. 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 reflect the interest charges over a given billing cycle on Banorte credit cards. Ending Card Balance - Payments correspond to the statement balance for a given billing cycle minus the payments received on the card over the following billing cycle (i.e., the payments made towards the bill of that previous billing cycle). Spending, balances, interest charges, payments, and credit limits are all in Mexican Pesos (MXN). 1 MXN = 0.107 USD PPP.

Table 2: Checking and Credit Card Account Balances for Individuals Who Have a Credit Card - By Deciles of Average Daily Balance on Checking Accounts, Over Income

	All Individuals with a Credit Card				*Individuals with a Credit Card Who Paid Interest*				
Decile	N	Checking Account Balance over Income (Average)	Fraction of Clients with Non-Zero Credit Card Balance	Fraction of Clients Paying Credit Card Interest	Checking Account Balances (Average)	Monthly Income (Average)	Credit Card Balances (Average)	Monthly Credit Card Interest (Average)	Credit Card Interest over Income (Average)
1	36,223	0.00	0.72	0.50	0.01	16,019.88	28,804.16	571.35	0.05
2	36,223	0.00	0.58	0.36	9.05	20,713.47	23,654.68	500.35	0.03
3	36,223	0.00	0.56	0.35	45.02	19,226.49	24,039.50	506.01	0.03
4	36,222	0.01	0.59	0.34	160.47	18,871.20	25,794.53	535.75	0.04
5	36,222	0.02	0.60	0.33	523.51	21,579.45	29,258.95	603.34	0.04
6	36,222	0.05	0.61	0.31	1,420.75	22,544.68	31,026.73	619.37	0.04
7	36,222	0.12	0.64	0.29	3,525.20	23,440.66	34,996.86	683.40	0.04
8	36,222	0.39	0.62	0.24	10,852.61	23,067.15	38,223.50	717.47	0.05
9	36,222	1.45	0.59	0.20	35,875.11	23,129.84	36,077.00	669.31	0.05
10	36,222	8.25	0.55	0.17	128,245.90	18,009.11	33,025.35	623.27	0.05

This table presents statistics about credit card borrowing and checking account balances for individuals who have a credit card and pay interest. Individuals are split into deciles of checking account balances over income. For checking account balances, we consider the minimum balance observed for each user in the 6 months previous to the intervention. For observations in each decile group, we first present the average of checking account balances over income as well as the fraction of individuals with a non-zero credit card balance and the fraction of individuals paying credit card interest in the month previous to the intervention. We then focus on individuals who are paying credit card interest in the month previous to the intervention. For them, we present average checking account balances (for each user we consider the minimum over the 6 months previous to the intervention and report the average across users), as well as average credit card balances, average monthly interest charges, and the average (across users) of the ratio of monthly interest charges to monthly income in the month previous to the intervention. Balances and interest charges are in Mexican Pesos (MXN). 1 MXN = 0.107 USD PPP.

Table 3: Covariate Balance

	Control	Treatment	p-value of Difference
Age (Years)	44.28	44.31	0.2157
Monthly Income	13,495.60	13,510.17	0.6892
Tenure (Months)	84.16	80.04	0.5219
Monthly Spending	16,232.41	16,107.47	0.5602
Ln Monthly Spending +1	8.18	8.17	0.3290
Checking Account Balance	18,221.77	18,096.49	0.2951
Ln Checking Account Balance +1	8.03	8.02	0.3210
Monthly Credit Card Interest	32.04	31.46	0.2489
Ln Monthly Credit Card Interest +1	0.26	0.26	0.4283
Credit Card Balance	3,914.83	3,935.19	0.4124
Ln Credit Card Balance +1	1.33	1.34	0.5973
Ending Card Balance - Next Payment	579.17	586.75	0.3151
Ln Ending Card Balance - Next Payment +1	6.34	6.34	0.7027
Credit Card Limit	17,973.16	17,924.83	0.6176
N	357,567	2,696,936	

This table presents a covariate balance test in which we estimate Equation (1) regressing the dependent variable specified in each row on strata fixed effects and a treatment indicator. We present the average value of each dependent variable for treatment and control groups and the p-value of regressing the corresponding outcome on the 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 0.3788. The statistics are calculated with monthly information at the individual level covering the 6 months previous to the intervention. Monthly Income are salary transfers flagged as such in the payroll accounts. Tenure with the bank corresponds to the number of months since the accounts were opened. 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 reflect the interest charges over a given billing cycle on Banorte credit cards. Ending Card Balance - Payments correspond to the statement balance for a given billing cycle minus the payments received on the card over the following billing cycle (i.e., the payments made towards the bill of that previous billing cycle). Spending, balances, interest charges, payments and credit limits are all in Mexican Pesos (MXN). 1 MXN = 0.107 USD PPP.

Table 4: Aggregate Effects of the Intervention

	(1)	(2)	(3)	(4)	(5)
	Ln Spending +1	Ln Checking Account Balance +1	Ln Credit Card Interest +1 During Treat.	Paid Interest During Treat. {0,1}	Ln Ending Statement Balance - Payments After Treat. +1
Panel A: All Individuals					
Any Treatment	-0.009*	0.006*			
	(0.005)	(0.003)			
Observations	3,054,503	3,054,503			
Mean of Dep.Var. in Control Group	16,732.41	17,393.63			
Panel B: Individuals with a Credit Card					
Any Treatment	-0.021***	0.012**	-0.004	-0.001	-0.003
	(0.006)	(0.006)	(0.004)	(0.004)	(0.005)
Observations	362,223	362,223	362,223	362,223	362,223
Mean of Dep.Var. in Control Group	29,960.75	34,586.21	213.84	0.41	4,981.45
Panel C: Individuals with a Credit Card Who Paid Interest at Baseline					
Any Treatment	-0.019**	0.017**	-0.004	-0.001	0.002
	(0.007)	(0.007)	(0.005)	(0.005)	(0.005)
Observations	152,016	152,016	152,016	152,016	152,016
Mean of Dep.Var. in Control Group	31,818.77	31,940.83	479.14	0.81	10,219.67

This table presents the results of estimating Equation (1) for spending, saving, and borrowing. 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 over the first month of the intervention. Credit Card Interest corresponds to the average of the monthly interest charges for interest-bearing balances 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 Card Balance - Payments After Treat. correspond to the statement balance for the second billing cycle affected by the treatment minus the payments received on the card over the following billing cycle (i.e., the payments made towards the bill of that previous billing cycle). 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 during at least one of the 6 months previous to 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. Observations are at the user level. Robust standard errors in parentheses.

Table 5: Calibration Test for Evaluation of The Causal Forests

Dep.Var	(1) Ln Monthly Spending +1	(2) Ln Checking Account Balance +1	(3) Ln Monthly Credit Card Interest +1	(4) Ln Monthly Credit Card Interest +1	(5) Ln Monthly Credit Card Interest +1	(6) Ln Monthly Credit Card Interest +1
Panel A: All Individuals (N = 3,054,503)						
Mean Forest Prediction	1.881***	1.529***				
	(0.424)	(0.473)				
Differential Forest Prediction	0.592***	0.347***				
	(0.198)	(0.128)				
Panel B: Individuals with a Credit Card (N = 362,223)						
Mean Forest Prediction	1.971***	1.831**	1.770*	2.148**	3.106*	2.542**
	(0.591)	(0.664)	(0.911)	(0.612)	(0.701)	(0.656)
Differential Forest Prediction	0.718**	0.446**	-0.22	0.076	-0.050	-0.329
	(0.312)	(0.254)	(0.299)	(0.285)	(0.418)	(0.499)
Panel C: Individuals with a Credit Card Who Paid Interest at Baseline (N= 152,016)						
Mean Forest Prediction	1.881***	1.981**	2.835*	1.937***	2.917**	3.516***
	(0.720)	(0.896)	(0.994)	(0.890)	(0.846)	(0.796)
Differential Forest Prediction	0.738**	0.812***	-0.477	0.309	-0.139	0.140
	(0.391)	(0.321)	(0.566)	(0.652)	(0.544)	(0.466)

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. In Column (1) the dependent variable is Spending, defined as the sum of ATM withdrawals, credit and debit card spending, and outgoing transfers over the first month of the intervention. The dependent variable in Column (2) corresponds to Checking Account Balances, measured as average daily balances during the first month of the intervention. Columns (3) to (5) use Credit Card Interest as the dependent variable. Credit Card Interest corresponds to the average of the monthly interest charges for interest-bearing balances during the two billing cycles affected by the intervention. Column (3) trains the model with all 161 available covariates. Column (4) trains the model with variables with importance larger than 1% according to the model in Column (3). Column (5) trains the model with variables used to train the model for checking account balances. Column (6) trains the model with variables used to train the model for spending. Panel A considers all individuals. Panel B considers all individuals who have a credit card. Panel C considers individuals with a credit card who paid credit card interest during at least one of the 6 months previous to the intervention. All dependent variables are residualized.

Figure 1: Distribution of Predicted Treatment Effects

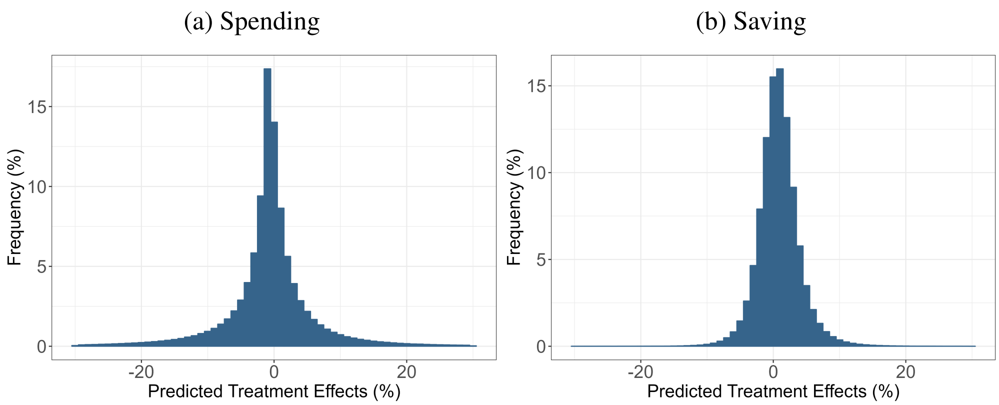

This graph shows the distribution of predicted treatment effects, at the individual level, stemming from the causal forests with 2,000 trees. The outcome variables are Ln Monthly Spending +1 and Ln Checking Account Balances +1, respectively. We use the explanatory variables described in Subsection 4.2 and Figure A1. Spending is defined as the sum of ATM withdrawals, credit and debit card spending, and outgoing transfers over the first month of the intervention. Our measure of checking account balances correspond to average daily balances during the first month of the intervention.

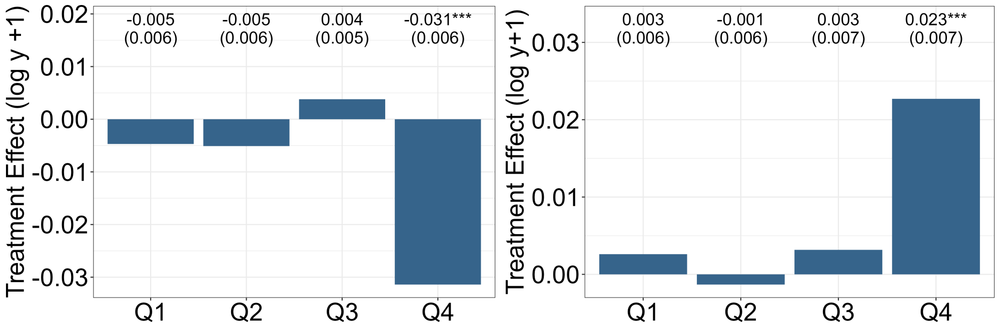

(a) Spending (b) Saving

Figure 2: Treatment effects on spending and checking account balances, as a function of predicted treatment effects. Individuals are split in to quartiles of predicted treatment effects on spending and savings, based on the score generated by the causal forest. Treatment effects are estimated using Ln Monthly Spending +1 and Ln Checking Account Balances +1 as the dependent variables. Spending is defined as the sum of ATM withdrawals, credit and debit card spending, and outgoing transfers over the first month of the intervention. Our measure of checking account balances corresponds to average daily balances during the first month of the intervention. 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.

Table 6: 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	0.1106	0.0619	0.0596	0.0179
2	0.0616	0.1004	0.0450	0.0430
3	0.0418	0.0743	0.0855	0.0484
4	0.0360	0.0133	0.0599	0.1408

(b) Individuals with a Credit Card

	1	2	3	4
1	0.1107	0.0406	0.0610	0.0377
2	0.0599	0.1173	0.0473	0.0253
3	0.0615	0.0649	0.0939	0.0301
4	0.0179	0.0272	0.0480	0.1570

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

	1	2	3	4
1	0.1017	0.0576	0.0571	0.0332
2	0.0592	0.0901	0.0526	0.0478
3	0.0571	0.0669	0.1070	0.0195
4	0.0319	0.0355	0.0339	0.1492

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.

Table 7: Treatment Effects on Spending, Saving, and Borrowing (Top Quartile of Predicted Treatment Effects on Spending)

Dep.Var	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Ln Spending +1	Ln Checking Account Balance +1	Ln Credit Card Interest +1 During Treat. (Banorte)	Ln Credit Card Interest +1 After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ln Ending Statement Balance - Payments After Treat. +1
Panel A: Individuals with a Credit Card							
TE	-0.0782***	0.0508***	-0.0071	-0.0077	0.0087	-0.0061	0.0033
	(0.0120)	(0.0145)	(0.0176)	(0.0178)	(0.0094)	(0.0095)	(0.0198)
Mean of Dep. Var in Control Group (MXN)	33,485.48	41,463.01	207.37	210.91	0.47	0.46	5,088.41
Change in Spending or Saving (MXN)	-2,618.56	2,106.32					
Upper Confidence Interval (MXN)			5.68	5.73	0.01	0.01	214.26
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0022	0.0022	0.0000	0.0000	0.0818
Lower Confidence Interval (MXN)			-8.63	-8.98	-0.00	-0.01	-180.68
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0033	-0.0034	-0.0000	-0.0000	-0.0690
N= 149561							
Panel B: Individuals with a Credit Card Who Paid Interest at Baseline							
TE	-0.0739***	0.0537**	-0.0056	-0.0053	0.0091	-0.0084	-0.0019
	(0.0182)	(0.0255)	(0.0205)	(0.0201)	(0.0102)	(0.0104)	(0.0202)
Mean of Dep. Var in Control Group (MXN)	35,190.08	36,471.10	400.85	415.03	0.87	0.89	11,186.20
Change in Spending or Saving (MXN)	-2,600.55	1,958.50					
Upper Confidence Interval (MXN)			13.86	14.15	0.03	0.01	421.63
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0053	0.0054	0.0000	0.0000	0.1621
Lower Confidence Interval (MXN)			-18.35	-18.55	-0.01	-0.03	-464.14
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0071	-0.0071	-0.0000	-0.0000	-0.1785
N= 72365							

This table shows treatment effects on spending, saving, and borrowing for individuals in the top quartile of predicted treatment effects on spending. 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. Credit Card Interest corresponds to the average of the monthly interest charges for interest-bearing balances during the two billing cycles affected by the intervention. Paid Interest is a binary variable flagging when an individual was charged credit card interest on a given billing cycle. These variables are measured on a monthly basis and averaged over the two months of the intervention or the two months following the intervention. Statement Balance - Payments After Treat. correspond to the balance on a credit card on the last day of the last billing cycle affected by the intervention minus payments received on the credit card account during the billing cycle immediately following the intervention (i.e., the payments made towards the bill of that previous billing cycle). Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest during at least one of the 6 months previous to the intervention. Average treatment effects are calculated with the AIPW method. 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.[1] 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. *p<0.1; **p<0.05; ***p<0.01.

Table 8: Treatment Effects on Spending, Saving, and Borrowing (Top Quartile of Predicted Treatment Effects on Saving)

Dep.Var	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Ln Spending +1	Ln Checking Account Balance +1	Ln Credit Card Interest +1 During Treat. (Banorte)	Ln Credit Card Interest +1 After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ln Ending Statement Balance - Payments After Treat. +1
Panel A: All Clients with a Credit Card							
TE	-0.0551***	0.0611***	-0.0082	-0.0080	-0.0045	-0.0041	0.0038
	(0.0182)	(0.0137)	(0.0175)	(0.0170)	(0.0067)	(0.0074)	(0.0181)
Mean of Dep. Var in Control Group (MXN)	37,265.33	31,737.78	218.54	220.34	0.44	0.45	4,739.24
Change in Spending or Saving (MXN)	-2,053.32	1,939.18					
Upper Confidence Interval (MXN)			5.70	5.58	0.00	0.00	186.14
Upper Confidence Interval Divided by Abs. Value of Change in Saving			0.0029	0.0029	0.0000	0.0000	0.0960
Lower Confidence Interval (MXN)			-9.29	-9.10	-0.01	-0.01	-150.12
Lower Confidence Interval Divided by Abs. Value of Change in Saving			-0.0048	-0.0047	-0.0000	-0.0000	-0.0774
N= 147647							
Panel B: Clients with a Credit Card Who Paid Interest at Baseline							
TE	-0.0639***	0.0559**	-0.0067	-0.0063	-0.0035	-0.0033	-0.0042
	(0.0201)	(0.0218)	(0.0200)	(0.0199)	(0.0097)	(0.0091)	(0.0209)
Mean of Dep. Var in Control Group (MXN)	31,034.19	27,809.32	403.93	405.33	0.74	0.76	10,414.98
Change in spending or savings (MXN)	-1,983.08	1,554.54					
Upper Confidence Interval (MXN)			13.12	13.26	0.01	0.01	382.90
Upper Confidence Interval Divided by Abs. Value of Change in Saving			0.0084	0.0085	0.0000	0.0000	0.2463
Lower Confidence Interval (MXN)			-18.54	-18.36	-0.02	-0.02	-470.38
Lower Confidence Interval Divided by Abs. Value of Change in Saving			-0.0119	-0.0118	-0.0000	-0.0000	-0.3026
N= 70912							

This table shows treatment effects on spending, saving, and borrowing for individuals in the top quartile of predicted treatment effects on 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. Credit Card Interest corresponds to the average of the monthly interest charges for interest-bearing balances during the two billing cycles affected by the intervention. Paid Interest is a binary variable flagging when an individual was charged credit card interest on a given billing cycle. These variables are measured on a monthly basis and averaged over the two months of the intervention or the two months following the intervention. Statement Balance - Payments After Treat. correspond to the balance on a credit card on the last day of the last billing cycle affected by the intervention minus payments received on the credit card account during the billing cycle immediately following the intervention (i.e., the payments made towards the bill of that previous billing cycle). Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest during at least one of the 6 months previous to the intervention. Average treatment effects are calculated with the AIPW method. 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.[1] 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. *p<0.1; **p<0.05; ***p<0.01.

Figure 3: Distribution of Individuals Who Have a Credit Card, Who Paid Interest at Baseline or Who are in the Puzzle Group, Across Quartiles of Predicted Treatment Effects

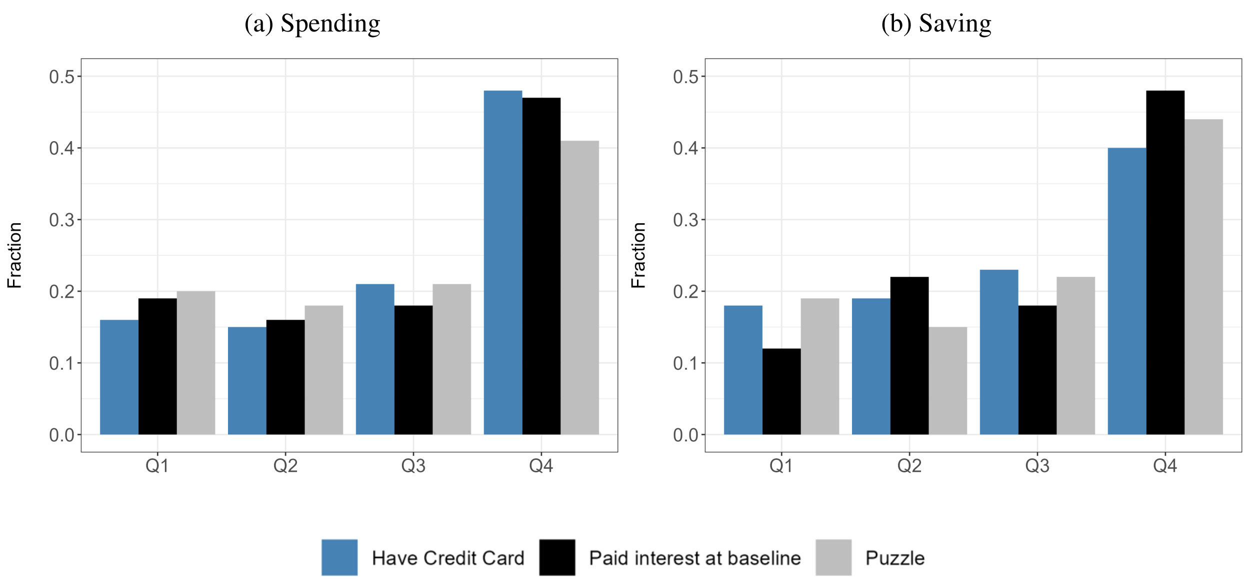

This graph shows the distribution of individuals who have a credit card, who paid credit card interest at baseline, or who belong to the puzzle group, across quartiles of predicted treatment effects on spending (left panel) or saving (right panel). Individuals belong to the puzzle group if they paid interest at baseline and the minimum balance in their accounts over the 6 months previous to the intervention represent more than 50% of their income. 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 9: Treatment Effects on Deposits, ATM Whithrawals, Card Spending, and Transfers (Top Quartile of Predicted Treatment Effects on Spending)

	(1)	(2)	(3)	(4)
Dep.Var	Ln Deposits +1	Ln ATM Withrawals +1	Ln Spending Debit or Credit Card +1	Ln Transfers +1
Panel A: All Clients with a Credit Card				
TE	-0.0102	-0.0878***	-0.0699***	0.0022
	(0.0103)	(0.0113)	(0.0117)	(0.0112)
Mean of Dep. Var in Control Group (MXN)	29,362.42	14,154.34	17,199.21	1,663.67
N= 149561				
Panel B: Clients with a Credit Card who Paid Interest at Baseline				
TE	-0.0095	-0.0991***	-0.0530***	0.0016
	(0.0124)	(0.0140)	(0.0128)	(0.0134)
Mean of Dep. Var in Control Group (MXN)	24,470.24	12,743.78	20,034.51	1,483.56
N= 72365				

This table considers all individuals with credit cards in the top quartile of the distribution of predicted treatment effects on spending. Deposits, withdrawals, spending with cards, and transfers are summed across the first month of the intervention. Spending with Credit or Debit Card is defined as the sum of debit or credit card store or online purchases. Transfers are defined as all outgoing electronic transfers originated from any of the accounts in the analysis. Incoming transfers are classified as deposits. The number of observations is the same across all columns in the same panel. The Mean of the Dependent Varible is reported in Mexican Pesos (MXN). 1 MXN = 0.107 USD PPP. Treatment effects are calculated with the AIPW method. Robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table 10: Treatment Effects on Spending, Saving, and Borrowing by Message (Top Quartile of Predicted Treatment Effects on Spending)

Dep.Var	(1)	(2)	(3)
	Ln Monthly Spending +1	Ln Checking Account Balance +1	Ln Credit Card Interest +1
Short-term messages			
Msg 2	-0.087***	0.052**	0.005
Year-end Expenses	(0.022)	(0.023)	(0.036)
Msg 6	-0.041*	0.022	-0.011
Avoid Shortfalls	(0.021)	(0.023)	(0.036)
Msg 7	-0.092***	0.055**	-0.011
Emergency	(0.021)	(0.023)	(0.036)
All Short-term	-0.074***	0.043**	-0.006
Msgs. Pooled	(0.017)	(0.018)	(0.028)
Long-term messages			
Msg 1	-0.035	0.021	-0.013
Congratulations	(0.022)	(0.023)	(0.036)
Msg 3	-0.122***	0.078***	0.006
Others your Age	(0.021)	(0.023)	(0.036)
Msg 5	-0.069***	0.037	-0.015
Reach Dreams	(0.022)	(0.023)	(0.036)
All Long-term	-0.075***	0.045**	-0.007
Msgs. Pooled	(0.017)	(0.018)	(0.028)
Mental Accounting			
Msg 4	-0.124***	0.081***	-0.009
Money Box	(0.021)	(0.023)	(0.036)
Differences Across Types of Messages			
Short-term	0.002	-0.003	0.002
- Long-term	(0.014)	(0.015)	(0.025)
Short-term	0.050***	-0.038*	0.003
- Mental Accounting	(0.019)	(0.020)	(0.032)
Long-term	0.049**	-0.036*	0.002
- Mental Accounting	(0.019)	(0.021)	(0.031)

This table shows treatment effects on spending, saving, and borrowing for individuals in the top quartile of predicted treatment effects on spending. 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 are average daily balances during the first month of the intervention. Credit Card Interest corresponds to the average of the interest charges for the interest-bearing balances during the two billing cycles covered by the intervention. Robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table 11: Heterogeneous Treatment Effects on Spending, by Experimental Strata

	Dep. Var: Ln Spending +1								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Any Treatment	-0.004	-0.007	-0.008	-0.008	-0.009	-0.009	-0.008	-0.009	-0.008
	(0.009)	(0.009)	(0.009)	(0.006)	(0.005)	(0.006)	(0.005)	(0.006)	(0.005)
Any Treatment*Group$_1$	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted
Any Treatment*Group$_2$	-0.004	0.002	0.004	-0.004	0.002	-0.001	-0.003	0.003	-0.003
	(0.009)	(0.008)	(0.009)	(0.006)	(0.004)	(0.006)	(0.008)	(0.007)	(0.007)
Any Treatment*Group$_3$	0.002	-0.006	-0.004						
	(0.009)	(0.009)	(0.009)						
Any Treatment*Group$_4$	-0.017*	-0.004	-0.005						
	(0.009)	(0.009)	(0.009)						
Group Definition	Quartiles of Checking Acct. Balance	Quartiles of Income	Quartiles of Age	Median of Tenure with Banorte	Median of ATM Withrawals	Median of Debit Card Transactions	Is Digital?	Main Bank?	Has Credit Card?
Observations	3054503	3054503	3054503	3054503	3054503	3054503	3054503	3054503	3054503

This table presents heterogeneous treatment effects on spending, by experimental strata. Spending is defined as the sum of ATM withdrawals, credit and debit card spending, and outgoing transfers over the first month of the intervention. Treatment effects are estimated in each column with the following OLS regression: $y_i = \alpha_s + Treatment_i + Group_{ij} + Treatment * Group_{ij}$ where $alpha_s$ represents strata fixed effects and $Group_{ij}$ is a dummy variable that takes the value of 1 when individual i belongs to Group j. In each column the $Group_{ij}$ is defined over a different variable which in turn defines the experimental strata. In Columns (1)-(9) $Group_{ij}$ is defined by quartiles of checking account balances, quartiles of income, quartiles of age, median of tenure with Banorte, Median ATM Withrawals, 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 withrawals), or a binary variable indicating if a given user has a credit card, respectively. In all cases we consider all 3.1 million observations at the user level. Robust standard errors in parenthesis. *p<0.1; **p<0.05; ***p<0.01.

Table 12: 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, based on the Causal Forest for Spending

	Top Quartile of Individuals' Observed Treatment Effects				Top Quartile of Individuals' Predicted Treatment Effects (Causal Forest)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep.Var.	N	Ln Spending +1	Ln Checking Account Balances +1	Ln Credit Card Interest +1	N	Ln Spending +1	Ln Checking Account Balances +1	Ln Credit Card Interest +1
Panel A: All Clients with a Credit Card								
	151,752				149,561			
TE		-0.4391***	0.3319***	-0.1544***		-0.0782***	0.0508***	-0.0071
		(0.0118)	(0.02021)	(0.0231)		(0.0120)	(0.0145)	(0.0176)
Mean of Dep. Var (MXN)		31,820.78	35,210.03	231.82		33,485.48	41,463.01	207.37
Panel B: Clients with a Credit Card Who Paid Interest at Baseline								
	68,107				72,365			
TE		-0.3176***	0.2083***	-0.1148***		-0.0739***	0.0537**	-0.0056
		(0.0173)	(0.0292)	(0.0250)		(0.0182)	(0.0255)	(0.0205)
Mean of Dep. Var (MXN)		26,678.11	34,178.63	380.31		35,190.08	36,471.10	400.85

This table shows treatment effects on spending and borrowing for clients in groups with the highest observed treatment effects or for clients with the highest treatment effects predicted at the individual level according to the causal forest for spending. For Columns (1) to (4), 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 savings. We then 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. The top quartile corresponds observations in strata blocks with the highest observed treatment effect on spending. For observations in that group, we calculate treatment effects on spending, saving and borrowing regressing the corresponding outcome variable on a treatment indicator and strata-block fixed effects. 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. Credit Card Interest corresponds to the average of the monthly interest charges for interest-bearing balances during the two billing cycles covered by the intervention. In Columns (5) to (8) we calculate treatment effects for the same set of dependent variables but for the top quartile of individuals with the highest individual treatment effects as predicted by the causal forest. Robust standard errors in parenthesis. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

Internet Appendix

Appendix A Causal Forests and The Generalized Random Forest Algorithm

A.1 Detailed Description

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 sub-populations. 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 not the same 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 in 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 group by training separate regression forests. Then the algorithm performs sample splits to identify heterogeneous treatment effects 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.[24]

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 estimation of causal forests, but also allows for 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.

[24] 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.

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.
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.

A.2 Variable Importance

Figure A1 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 A1: Variable Importance: Causal Forests for Spending and Saving

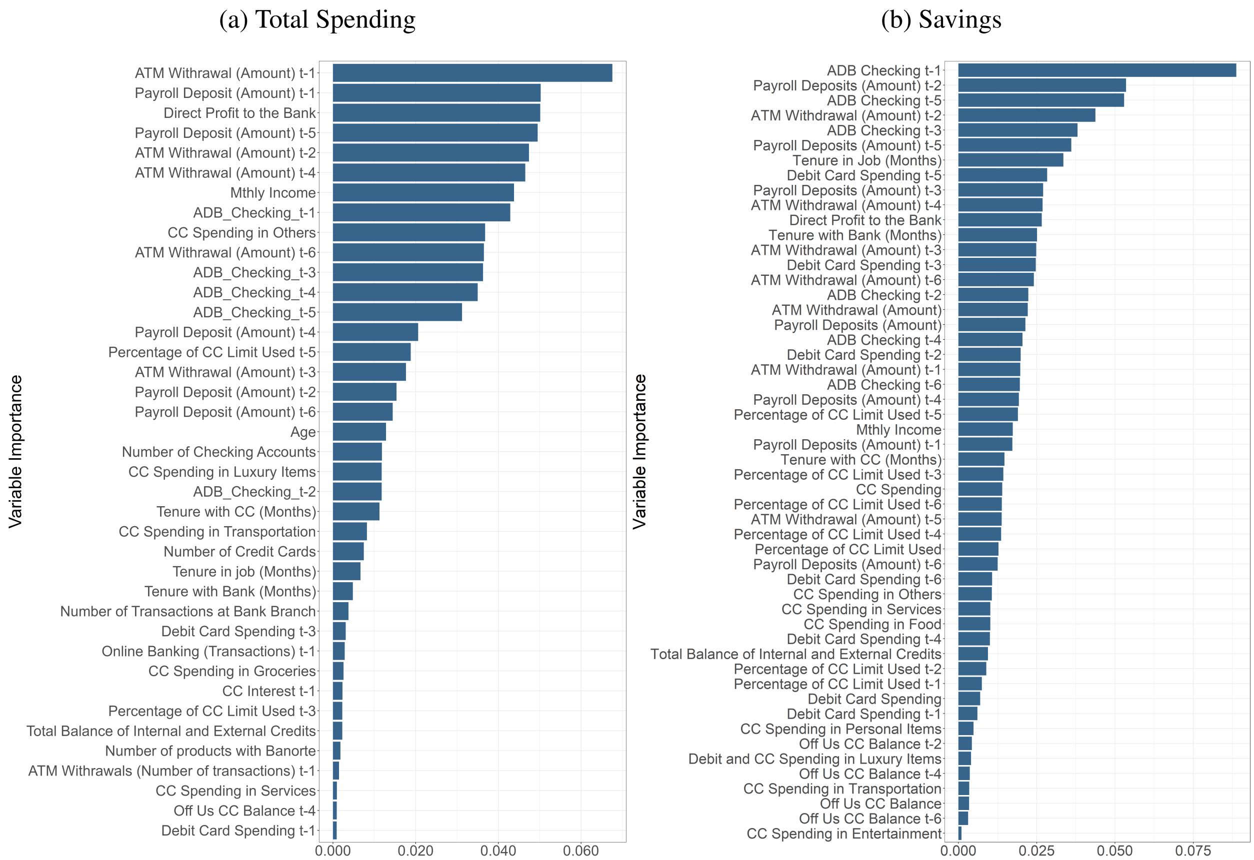

This graph shows the variable importance of the variables used in the estimation of the final causal forests for both spending and saving. Spending is defined as the sum of ATM withdrawals, credit and debit card spending, and outgoing transfers over the first month of the intervention. Checking Account 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. By construction, the variable importance of all variables used in a causal forests add up to one. 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. t refers to the month of the treatment, in turn, t-1 or 2 or 3 refer to the 1 or 2 or 3 months before the treatment period. All variables are monthly.

A.3 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 A1.

Table A1: 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	19,752.64	20,536.85	20,088.93	21,117.29
Checking Account Balance	15,464.43	23,636.47	16,017.05	21,338.30
Monthly Income	15,469.88	16,280.77	14,118.44	15,109.87
Age (Years)	44.83	46.96	44.18	46.35
Married	0.49	0.50	0.48	0.52
Women	0.49	0.51	0.49	0.51
Tenure With Banorte (Months)	72.05	90.52	74.60	88.69
Credit Card Balance	836.75	3,688.96	919.41	2,837.60
Credit Card Limit	6,734.76	16,476.65	8,302.55	15,873.31

This table presents means of each variable for individuals that fall into the top and bottom quartiles of the distribution of predicted treatment effects on spending or saving, respectively. 1 MXN=0.107 USD PPP.

Appendix B Additional Results

B.1 Characterization of Individuals in the Co-holding Puzzle Group

Table B2 compares individuals in the puzzle group to the rest of those who pay credit card interest.

Table B2: Individuals Paying Credit Card Interest With Checking Account Balances Over or Below 50% of their Income

	No-Puzzle (Less than 50%)	Puzzle (50% or more)	p-value of Difference
Age (Years)	40.36	49.42	0.000
Monthly Income	16,982.95	20,985.22	0.000
Tenure (Months)	115.78	142.62	0.000
Checking Account Balance	28,295.61	91,062.07	0.000
Credit Card Interest	459.08	660.87	0.000
Credit Card Balance	24,162.33	34,782.66	0.000
Credit Card Limit	81,305.90	171,661.03	0.000
P(Interest_t > 0| Interest_t-1 >0	0.84	0.85	0.000
N	136,907	15,109	

This table presents simple means of each variable for individuals that co-hold or not based on our co-holding puzzle definition. We consider all individuals who have a credit card and are paying credit card interest. We say that individuals co-hold if the minimum of the daily balances in her checking account over the last 6 months previous to the intervention is higher than 50% of her income and they are paying credit card interest. Monthly income and balances are in Mexican Pesos (MXN). 1 MXN=0.107 USD PPP. The last column presents robust standard errors of a t-test for differences in means.

B.2 Alternative Outcome Variables, Specifications and Subsamples

Table B3 shows results for a set of alternative outcome variables. Table B4 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.

Tables B5 and B6 present the impact of the intervention on spending, saving, and borrowing outcomes measured in MXN. Our results are very similar. These tables reassure that our results are not driven by the selection of the functional form of Log variable plus one (Cohn et al., 2022).

Tables B7 and B8 show the main results for individuals who have a credit card and who have a credit line utilization lower than the median.

Tables B9 and B10 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 recors, don't have credits outside of Banorte).

Table B3: 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)
	Ln Interest Bearing Balance +1 During Treat.	Ln Interest Bearing Balance +1 After Treat.	Ln Credit Card Balance +1 During Treat. (Banorte)	Ln Credit Card Balance +1 After Treat. (Banorte)	Ln Credit Card Balance +1 During Treat. (Credit Bureau)	Ln Credit Card Balance +1 After Treat. (Credit Bureau)	Ln Monthly Credit Card Payments +1 During Treat.	Ln Monthly Credit Card Payments +1 After Treat.
Panel A: Individuals with a Credit Card								
TE	-0.0067	-0.0074	-0.0078	-0.0081	-0.0048	0.0064	-0.0164	-0.0171
	(0.0176)	(0.0178)	(0.0198)	(0.0197)	(0.0108)	(0.0110)	(0.0168)	(0.0168)
Mean of Dep. Var in Control Group (MXN)	6,625.24	6,738.34	10,729.91	11,109.49	26,448.35	27,015.24	5,973.35	6,134.92
Upper Confidence Interval (MXN)	184.16	185.22	332.71	338.97	432.91	755.35	98.73	97.10
Upper Confidence Interval Divided by Abs. Value of Change in Spending	0.0703	0.0707	0.1271	0.1294	0.1653	0.2885	0.0377	0.0371
Lower Confidence Interval (MXN)	-272.93	-284.95	-500.10	-518.95	-686.81	-409.55	-294.65	-306.92
Lower Confidence Interval Divided by Abs. Value of Change in Spending	-0.1042	-0.1088	-0.1910	-0.1982	-0.2623	-0.1564	-0.1125	-0.1172
N= 149561								
Panel B: Individuals with a Credit Card Who Paid Interest at Baseline								
TE	-0.0049	-0.0045	-0.0047	-0.0049	-0.0069	0.0074	-0.0101	-0.0180
	(0.0202)	(0.0201)	(0.0203)	(0.0203)	(0.0164)	(0.0162)	(0.0201)	(0.0198)
Mean of Dep. Var in Control Group (MXN)	12,325.21	12,761.21	13,310.16	13,781.01	28,104.91	28,035.87	3,019.46	3,101.54
Upper Confidence Interval (MXN)	427.59	445.32	467.03	481.21	709.20	1,097.66	88.46	64.47
Upper Confidence Interval Divided by Abs. Value of Change in Spending	0.1644	0.1712	0.1796	0.1850	0.2727	0.4221	0.0340	0.0248
Lower Confidence Interval (MXN)	-548.37	-560.17	-592.14	-615.43	-1,097.61	-682.73	-149.45	-176.25
Lower Confidence Interval Divided by Abs. Value of Change in Spending	-0.2109	-0.2154	-0.2277	-0.2367	-0.4221	-0.2625	-0.0575	-0.0678
N= 72365								

This table shows treatment effects on credit card balances and payments for individuals in the top quartile of predicted treatment effects on spending. Interest Bearing Balances correspond to the average daily balances over a billing cycle multiplied by Paid Interest, a binary variable flagging when an individual was charged credit card interest on a given 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. Average treatment effects are calculated with the AIPW method. The change in spending, expressed in MXN, is calculated by multiplying the TE and the Mean of Spending 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.[1] 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. *p<0.1; **p<0.05; ***p<0.01.

Table B4: 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)
	Ln Interest Bearing Balance +1 During Treat.	Ln Interest Bearing Balance +1 After Treat.	Ln Credit Card Balance +1 During Treat. (Banorte)	Ln Credit Card Balance +1 After Treat. (Banorte)	Ln Credit Card Balance +1 During Treat. (Credit Bureau)	Ln Credit Card Balance +1 After Treat. (Credit Bureau)	Ln Monthly Credit Card Payments +1 During Treat.	Ln Monthly Credit Card Payments +1 After Treat.
Panel A: All Clients with a Credit Card								
TE	-0.0085	-0.0084	-0.0080	-0.0082	0.0024	0.0037	-0.0171	-0.0198
	(0.0170)	(0.0174)	(0.0171)	(0.0171)	(0.0107)	(0.0106)	(0.0189)	(0.0187)
Mean of Dep. Var in Control Group (MXN)	6,999.81	7,116.97	10,245.50	10,555.58	27,459.86	26,164.81	5,703.68	6,031.76
Upper Confidence Interval (MXN)	173.74	182.93	261.42	267.23	641.79	640.41	113.75	101.65
Upper Confidence Interval Divided by Abs. Value of Change in Saving	0.0896	0.0943	0.1348	0.1378	0.3310	0.3302	0.0587	0.0524
Lower Confidence Interval (MXN)	-292.73	-302.50	-425.35	-440.34	-509.98	-446.79	-308.82	-340.51
Lower Confidence Interval Divided by Abs. Value of Change in Saving	-0.1510	-0.1560	-0.2193	-0.2271	-0.2630	-0.2304	-0.1593	-0.1756
N= 147647								
Panel B: Clients with a Credit Card Who Paid Interest at Baseline								
TE	-0.0070	-0.0069	-0.0071	-0.0068	0.0020	-0.0028	-0.0176	-0.0202
	(0.0203)	(0.0199)	(0.0204)	(0.0202)	(0.0151)	(0.0152)	(0.0212)	(0.0211)
Mean of Dep. Var in Control Group (MXN)	12,325.83	12,412.07	12,799.39	12,966.99	28,068.87	27,181.31	3,404.09	3,379.69
Upper Confidence Interval (MXN)	404.14	398.48	420.90	425.21	886.86	733.68	81.53	71.47
Upper Confidence Interval Divided by Abs. Value of Change in Saving	0.2600	0.2563	0.2708	0.2735	0.5705	0.4720	0.0524	0.0460
Lower Confidence Interval (MXN)	-576.70	-569.76	-602.65	-601.56	-774.59	-885.89	-201.36	-208.07
Lower Confidence Interval Divided by Abs. Value of Change in Saving	-0.3710	-0.3665	-0.3877	-0.3870	-0.4983	-0.5699	-0.1295	-0.1338
N= 70912								

This table shows treatment effects on credit card balances and payments for individuals in the top quartile of predicted treatment effects on saving. Interest Bearing Balances correspond to the average daily balances over a billing cycle multiplied by Paid Interest, a binary variable flagging when an individual was charged credit card interest on a given 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. Average treatment effects are calculated with the AIPW method. The change in saving, expressed in MXN, is calculated by multiplying the TE and the Mean of 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.[1] 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. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

Table B5: Treatment Effects on Spending, Saving, and Borrowing in MXN (Top Quartile of Predicted Treatment Effects on Spending)

Dep.Var	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Ln Spending +1	Ln Checking Account Balance +1	Ln Credit Card Interest +1 During Treat. (Banorte)	Ln Credit Card Interest +1 After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ln Ending Statement Balance - Payments After Treat. +1
Panel A: Individuals with a Credit Card							
TE	-2,474.19***	1,857.12***	-2.55	-2.46	0.01	-0.01	-28.27
	(629.71)	(521.94)	(6.49)	(6.45)	(0.01)	(0.01)	(239.34)
Mean of Dep. Var in Control Group (MXN)	33,485.48	41,463.00	207.37	210.91	0.47	0.46	5,088.41
Upper Confidence Interval (MXN)			10.17	10.19	0.03	0.01	440.82
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0041	0.0041	0.0000	0.0000	0.1782
Lower Confidence Interval (MXN)			-15.27	-15.11	-0.01	-0.02	-497.37
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0062	-0.0061	-0.0000	-0.0000	-0.2010
N= 149561							
Panel B: Individuals with a Credit Card Who Paid Interest at Baseline							
TE	2,378.32***	1,701.02***	-4.86	-4.67	0.01	-0.01	-38.61
	(468.70)	(357.28)	(10.92)	(11.11)	(0.01)	(0.01)	(280.93)
Mean of Dep. Var in Control Group (MXN)	35,190.08	36,471.08	400.85	415.03	0.87	0.89	11,186.20
Upper Confidence Interval (MXN)			16.54	17.10	0.03	0.01	512.02
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0070	0.0072	0.0000	0.0000	0.2153
Lower Confidence Interval (MXN)			-26.25	-26.44	-0.01	-0.03	-589.23
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0110	-0.0111	-0.0000	-0.0000	-0.2478
N= 72365							

This table shows treatment effects on spending, saving, and borrowing for individuals in the top quartile of predicted treatment effects on spending with all outcome variables measured in MXN. 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. Credit Card Interest corresponds to the average of the monthly interest charges for interest-bearing balances during the two billing cycles affected by the intervention. Paid Interest is a binary variable flagging when an individual was charged credit card interest on a given billing cycle. These variables are measured on a monthly basis and averaged over the two months of the intervention or the two months following the intervention. Statement Balance - Payments After Treat. correspond to the balance on a credit card on the last day of the last billing cycle affected by the intervention minus payments received on the credit card account during the billing cycle immediately following the intervention (i.e., the payments made towards the bill of that previous billing cycle). Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest during at least one of the 6 months previous to the intervention. Average treatment effects are calculated with the AIPW method. 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.[1] 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. *p<0.1; **p<0.05; ***p<0.01.

Table B6: Treatment Effects on Spending, Saving, and Borrowing in MXN (Top Quartile of Predicted Treatment Effects on Saving)

Dep.Var	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Ln Spending +1	Ln Checking Account Balance +1	Ln Credit Card Interest +1 During Treat. (Banorte)	Ln Credit Card Interest +1 After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ln Ending Statement Balance - Payments After Treat. +1
Panel A: Individuals with a Credit Card							
TE	-2,189.60***	2,057.25***	-2.25	-2.12	-0.00	-0.00	15.02
	(414.99)	(298.17)	(4.23)	(4.27)	(0.01)	(0.01)	(179.46)
Mean of Dep. Var in Control Group (MXN)	37,265.33	31,737.78	218.54	220.34	0.44	0.45	4,739.24
Upper Confidence Interval (MXN)			6.04	6.25	0.01	0.01	366.76
Upper Confidence Interval Divided by Abs. Value of Change in Saving			0.0029	0.0030	0.0000	0.0000	0.1783
Lower Confidence Interval (MXN)			-10.53	-10.49	-0.02	-0.02	-336.73
Lower Confidence Interval Divided by Abs. Value of Change in Saving			-0.0051	-0.0051	-0.0000	-0.0000	-0.1637
N= 147647							
Panel B: Individuals with a Credit Card Who Paid Interest at Baseline							
TE	-2,174.25***	1,863.31***	-3.92	-3.73	-0.00	-0.00	-36.96
	(510.66)	(701.40)	(9.63)	(9.44)	(0.01)	(0.01)	(220.49)
Mean of Dep. Var in Control Group (MXN)	31,034.19	27,809.32	400.85	415.03	0.87	0.89	11,186.20
Upper Confidence Interval (MXN)			14.95	14.78	0.02	0.01	395.19
Upper Confidence Interval Divided by Abs. Value of Change in Saving			0.0080	0.0079	0.0000	0.0000	0.2121
Lower Confidence Interval (MXN)			-22.79	-22.24	-0.02	-0.02	-469.11
Lower Confidence Interval Divided by Abs. Value of Change in Saving			-0.0122	-0.0119	-0.0000	-0.0000	-0.2518
N= 70912							

This table shows treatment effects on spending, saving, and borrowing for individuals in the top quartile of predicted treatment effects on spending with all outcome variables measured in MXN. 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. Credit Card Interest corresponds to the average of the monthly interest charges for interest-bearing balances during the two billing cycles affected by the intervention. Paid Interest is a binary variable flagging when an individual was charged credit card interest on a given billing cycle. These variables are measured on a monthly basis and averaged over the two months of the intervention or the two months following the intervention. Statement Balance - Payments After Treat. correspond to the balance on a credit card on the last day of the last billing cycle affected by the intervention minus payments received on the credit card account during the billing cycle immediately following the intervention (i.e., the payments made towards the bill of that previous billing cycle). Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest during at least one of the 6 months previous to the intervention. Average treatment effects are calculated with the AIPW method. 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.[1] 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. *p<0.1; **p<0.05; ***p<0.01.

Table B7: 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)
	Ln Spending +1	Ln Checking Account Balance +1	Ln Credit Card Interest +1 During Treat. (Banorte)	Ln Credit Card Interest +1 After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ln Ending Statement Balance - Payments After Treat. +1
Panel A: Clients with Credit Line Utilization Lower Than the Median							
TE	-0.0794***	0.0571***	-0.0064	-0.0067	-0.0021**	-0.0022**	0.0019
	(0.0205)	(0.0205)	(0.0215)	(0.0215)	(0.0010)	(0.0011)	(0.0225)
Mean of Dep. Var in Control Group (MXN)	30,707.93	32,622.47	120.87	119.52	0.27	0.26	3,559.52
Change in Spending or Saving (MXN)	-2,437.65	1,864.37					
Upper Confidence Interval (MXN)			4.32	4.24	-0.00	-0.00	163.74
Upper Confidence Interval Divided by Abs. Value of Change in Saving			0.0023	0.0023	-0.0000	-0.0000	0.0878
Lower Confidence Interval (MXN)			-5.87	-5.84	-0.00	-0.00	-150.21
Lower Confidence Interval Divided by Abs. Value of Change in Saving			-0.0031	-0.0031	-0.0000	-0.0000	-0.0806
N= 62911							

This table shows treatment effects on spending, saving, and borrowing 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. Credit line utilization is defined as the ratio of balances to credit line. The median credit line utilization among this group of individuals is 24.17 percent. 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. Credit Card Interest corresponds to the average of the monthly interest charges for interest-bearing balances during the two billing cycles affected by the intervention. Paid Interest is a binary variable flagging when an individual was charged credit card interest on a given billing cycle. These variables are measured on a monthly basis and averaged over the two months of the intervention or the two months following the intervention. Statement Balance - Payments After Treat. correspond to the balance on a credit card on the last day of the last billing cycle affected by the intervention minus payments received on the credit card account during the billing cycle immediately following the intervention (i.e., the payments made towards the bill of that previous billing cycle). Average treatment effects are calculated with the AIPW method. 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.[1] 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. *p<0.1; **p<0.05; ***p<0.01.

Table B8: 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)
	Ln Spending +1	Ln Checking Account Balance +1	Ln Credit Card Interest +1 During Treat. (Banorte)	Ln Credit Card Interest +1 After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ln Ending Statement Balance - Payments After Treat. +1
Panel A: Clients with Credit Line Utilization Lower Than the Median							
TE	-0.0687***	0.0615***	-0.0071	-0.0074	-0.0030*	0.0012	-0.0045
	(0.0206)	(0.0204)	(0.0213)	(0.0211)	(0.0018)	(0.0018)	(0.0215)
Mean of Dep. Var in Control Group (MXN)	30,880.90	31,948.95	127.66	128.06	0.29	0.28	3,600.34
Change in Spending or Saving (MXN)	-2,120.65	1,966.12					
Upper Confidence Interval (MXN)			4.42	4.35	0.00	0.00	135.52
Upper Confidence Interval Divided by Abs. Value of Change in Saving			0.0022	0.0022	0.0000	0.0000	0.0689
Lower Confidence Interval (MXN)			-6.24	-6.24	-0.00	-0.00	-167.92
Lower Confidence Interval Divided by Abs. Value of Change in Saving			-0.0032	-0.0032	-0.0000	-0.0000	-0.0854
N= 61735							

This table shows treatment effects on spending, saving, and borrowing 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. Credit line utilization is defined as the ratio of balances to credit line. The median credit line utilization among this group of individuals is 25.81 percent. 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. Credit Card Interest corresponds to the average of the monthly interest charges for interest-bearing balances during the two billing cycles affected by the intervention. Paid Interest is a binary variable flagging when an individual was charged credit card interest on a given billing cycle. These variables are measured on a monthly basis and averaged over the two months of the intervention or the two months following the intervention. Statement Balance - Payments After Treat. correspond to the balance on a credit card on the last day of the last billing cycle affected by the intervention minus payments received on the credit card account during the billing cycle immediately following the intervention (i.e., the payments made towards the bill of that previous billing cycle). Average treatment effects are calculated with the AIPW method. 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.[1] 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. *p<0.1; **p<0.05; ***p<0.01.

Table B9: Treatment Effects on Spending, Saving, and Borrowing (Top Quartile of Predicted Treatment Effects on Spending — Individuals for Whom Banorte is Their Main Bank)

Dep.Var	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Ln Spending +1	Ln Checking Account Balance +1	Ln Credit Card Interest +1 During Treat. (Banorte)	Ln Credit Card Interest +1 After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ln Ending Statement Balance - Payments After Treat. +1
Panel A: All Clients with a Credit Card							
TE	-0.0817***	0.0618***	-0.0081	-0.0077	-0.0081	-0.0073	-0.0028
	(0.0137)	(0.0151)	(0.0184)	(0.0184)	(0.0119)	(0.0120)	(0.0211)
Mean of Dep. Var in Control Group (MXN)	29,545.39	36,193.92	239.41	237.98	0.43	0.44	6,655.51
Change in Spending or Saving (MXN)	-2,413.74	2,235.08					
Upper Confidence Interval (MXN)			6.69	6.75	0.01	0.01	256.61
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0028	0.0028	0.0000	0.0000	0.1063
Lower Confidence Interval (MXN)			-10.57	-10.41	-0.01	-0.01	-293.88
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0044	-0.0043	-0.0000	-0.0000	-0.1218
N= 91625							
Panel B: Clients with a Credit Card Who Paid Interest at Baseline							
TE	-0.0792***	0.0651**	-0.0057	-0.0049	-0.0071	-0.0089	-0.0022
	(0.0191)	(0.0261)	(0.0227)	(0.0226)	(0.0214)	(0.0214)	(0.0231)
Mean of Dep. Var in Control Group (MXN)	31,503.58	33,124.14	402.42	401.13	0.72	0.72	12,790.04
Change in Spending or Saving (MXN)	-2,494.76	2,156.03					
Upper Confidence Interval (MXN)			15.61	15.79	0.02	0.02	550.65
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0063	0.0063	0.0000	0.0000	0.2207
Lower Confidence Interval (MXN)			-20.20	-19.74	-0.04	-0.04	-607.52
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0081	-0.0079	-0.0000	-0.0000	-0.2435
N= 43535							

This table shows treatment effects on spending, saving, and borrowing for individuals in the top quartile of predicted treatment effects on spending, for whom Banorte is their main bank (i.e., they receive their payroll at Banorte and, according to credit bureau recors, they don't have credits outside of Banorte). 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. Credit Card Interest corresponds to the average of the interest charges for the interest-bearing balances. Paid Interest is a binary variable flagging when an individual was charged credit card interest on a given billing cycle. These variables are measured on a monthly basis and averaged over the two months of the intervention or the two months following the intervention. Statement Balance - Payments After Treat. correspond to the balance on a credit card on the last day of the last billing cycle affected by the intervention minus payments received on the credit card account during the billing cycle immediately following the intervention (i.e., the payments made towards the bill of that previous billing cycle). Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest during at least one of the 6 months previous to the intervention. Average treatment effects are calculated with the AIPW method. 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.[1] 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. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

Table B10: Treatment Effects on Spending, Saving, and Borrowing (Top Quartile of Predicted Treatment Effects on Saving — Individuals for Whom Banorte is Their Main Bank)

Dep.Var	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Ln Spending +1	Ln Checking Account Balance +1	Ln Credit Card Interest +1 During Treat. (Banorte)	Ln Credit Card Interest +1 After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ln Ending Statement Balance - Payments After Treat. +1
Panel A: All Clients with a Credit Card							
TE	-0.0759*** (0.0194)	0.0643*** (0.0142)	-0.0064 (0.0187)	-0.0070 (0.0187)	-0.0085 (0.0128)	-0.0082 (0.0129)	-0.0021 (0.0218)
Mean of Dep. Var in Control Group (MXN)	31,278.02	34,615.44	228.79	227.35	0.46	0.46	5,076.50
Change in Spending or Saving (MXN)	-2,372.46	2,267.63					
Upper Confidence Interval (MXN)			6.93	6.74	0.01	0.01	206.49
Upper Confidence Interval Divided by Abs. Value of Change in Saving			0.0031	0.0030	0.0000	0.0000	0.0911
Lower Confidence Interval (MXN)			-9.84	-9.92	-0.02	-0.02	-227.33
Lower Confidence Interval Divided by Abs. Value of Change in Saving			-0.0043	-0.0044	-0.0000	-0.0000	-0.1002
N= 89899							
Panel B: Clients with a Credit Card Who Paid Interest at Baseline							
TE	-0.0720*** (0.0216)	0.0663*** (0.0195)	-0.0071 (0.0218)	-0.0077 (0.0219)	-0.0076 (0.0197)	-0.0082 (0.0198)	-0.0028 (0.0234)
Mean of Dep. Var in Control Group (MXN)	30,843.76	32,198.41	395.32	399.14	0.69	0.70	10,984.00
Change in Spending or Saving (MXN)	-2,219.46	2,133.70					
Upper Confidence Interval (MXN)			14.08	14.09	0.02	0.02	473.40
Upper Confidence Interval Divided by Abs. Value of Change in Saving			0.0066	0.0066	0.0000	0.0000	0.2219
Lower Confidence Interval (MXN)			-19.72	-20.23	-0.03	-0.03	-534.57
Lower Confidence Interval Divided by Abs. Value of Change in Saving			-0.0092	-0.0095	-0.0000	-0.0000	-0.2505
N= 41223							

This table shows treatment effects on spending, saving, and borrowing 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 and, according to credit bureau records, they don't have credits outside of Banorte). 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. Credit Card Interest corresponds to the average of the monthly interest charges for interest-bearing balances during the two billing cycles affected by the intervention. Paid Interest is a binary variable flagging when an individual was charged credit card interest on a given billing cycle. These variables are measured on a monthly basis and averaged over the two months of the intervention or the two months following the intervention. Statement Balance - Payments After Treat. correspond to the balance on a credit card on the last day of the last billing cycle affected by the intervention minus payments received on the credit card account during the billing cycle immediately following the intervention (i.e., the payments made towards the bill of that previous billing cycle). Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest during at least one of the 6 months previous to the intervention. Average treatment effects are calculated with the AIPW method. 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.[1] 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. *p<0.1; **p<0.05; ***p<0.01.

B.3 Prediction Errors and the Correlation between Carrying Credit Card Debt and Treatment Effects

Figure B2 shows that prediction errors are uncorrelated with treatment effects on borrowing outcomes. The prediction errors are thus the result of noise, which is uncorrelated with the treatment.

Figure B3 shows that there is no clear relationship between the fraction of individuals paying credit card interest and the treatment effect on spending and saving.

Figure B2: Correlation between Prediction Errors and Treatment Effects on Borrowing

(a) Prediction Error for Spending

(b) Prediction Error for Saving

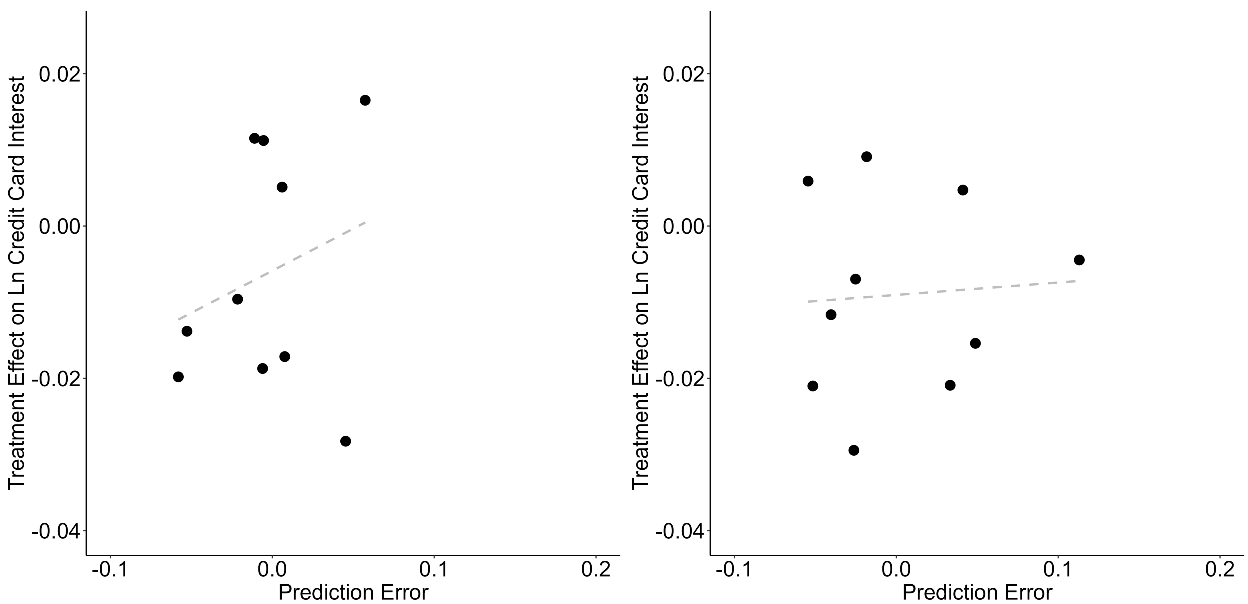

This graph shows the correlation between prediction errors and treatment effects on credit card interest. Prediction errors are defined as the difference between the simple average of individual-level predicted treatment effects on spending or saving, and the actual average treatment effect of observations in each group, as estimated with the AIPW method. The analysis considers observations in the top 25% of predicted treatment effects, which are further split into deciles. Observations are ranked with a cross-fitted procedure over five folds.

Figure B3: Correlation between the Fraction of Individuals Paying Credit Card Interest and the Treatment Effects on Spending and Saving

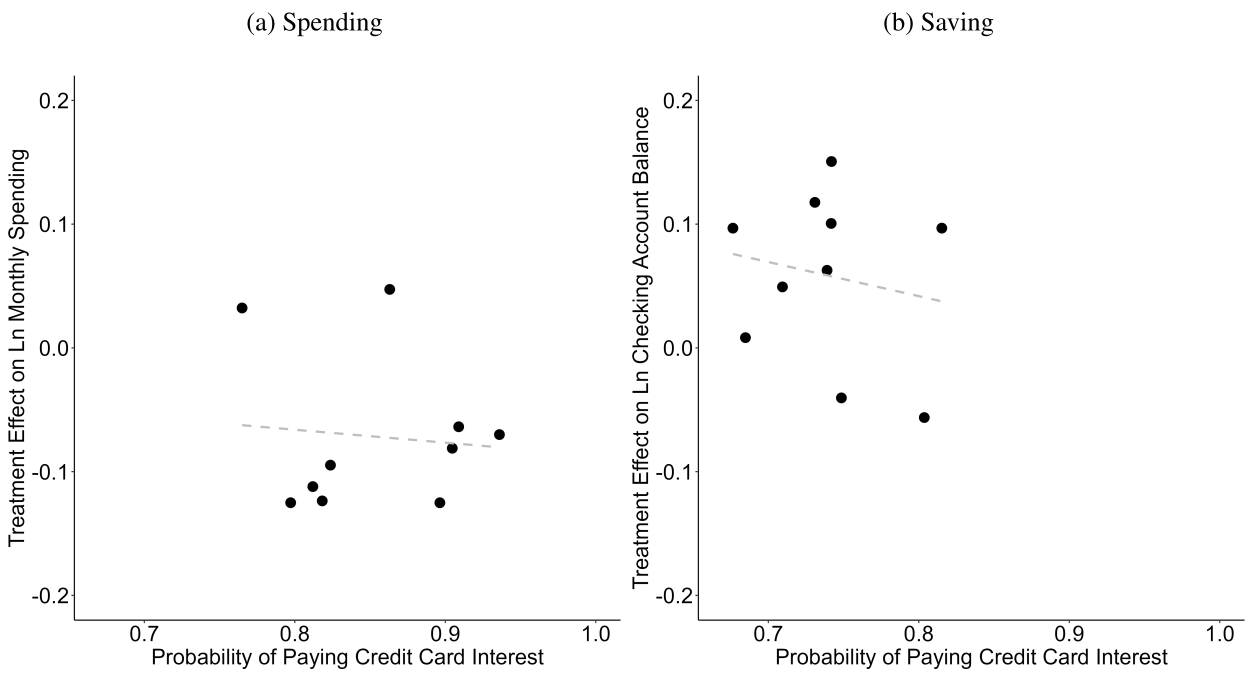

This graph shows the correlation between the fraction of individuals paying credit card interest during the treatment period and the treatment effect of the intervention on spending or saving. Treatment effects are calculated with the AIPW method. The analysis considers observations in the top 25% of predicted treatment effects on spending or saving, respectively, which are further split into deciles. Observations are ranked with a cross-fitted procedure over five folds.

B.4 Credit Card Due Dates and Billing Cycles

Figure B4 shows the distribution of the day of the month on which credit card payments are due, for all individuals who have a credit card in our sample.

Table B11 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 saving during the billing cycle fully covered by the intervention. The results are the same as in our main specifications.

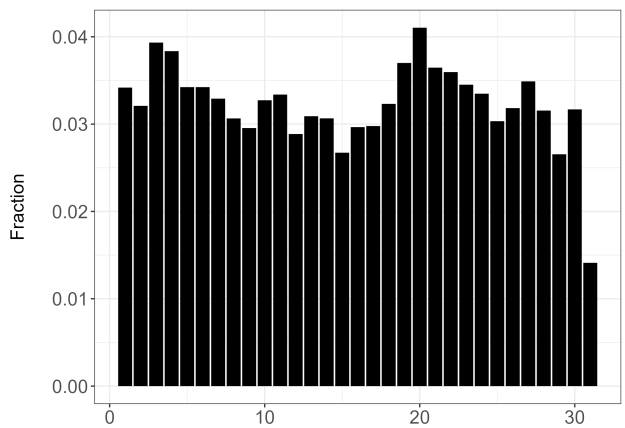

Figure B4: This figure shows the frequency distribution of the credit card due dates observed for the month of August 2019, immediately preceding the intervention.

Table B11: 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)
	Ln Spending +1	Ln Checking Account Balance +1	Ln Credit Card Interest +1 During Treat. (Banorte)	Ln Credit Card Interest +1 After Treat. (Banorte)	Paid Interest During Treat. {0,1}	Paid Interest After Treat. {0,1}	Ln Ending Statement Balance - Payments After Treat. +1
Panel A: Individuals with a Credit Card							
TE	-0.0817***	0.0514***	-0.0074	-0.0073	0.0085	0.0064	0.0047
	(0.0155)	(0.0189)	(0.0215)	(0.0212)	(0.0119)	(0.0120)	(0.0215)
Mean of Dep. Var in Control Group (MXN)	33,381.54	40,901.52	207.93	209.14	0.48	0.48	5,082.28
Change in Spending or Saving (MXN)	-2,726.85	2,100.77					
Upper Confidence Interval (MXN)			7.20	7.16	0.02	0.01	238.12
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0026	0.0026	0.0000	0.0000	0.0873
Lower Confidence Interval (MXN)			-10.29	-10.21	-0.01	-0.01	-190.74
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0038	-0.0037	-0.0000	-0.0000	-0.0699
N= 95719							
Panel B: Individuals with a Credit Card Who Paid Interest at Baseline							
TE	-0.0752***	0.0525*	-0.0058	-0.0055	0.0098	0.0096	-0.0057
	(0.0241)	(0.0297)	(0.0291)	(0.0294)	(0.0130)	(0.0131)	(0.0262)
Mean of Dep. Var in Control Group (MXN)	34,889.91	35,814.94	394.82	408.82	0.88	0.87	10,728.80
Change in Spending or Saving (MXN)	-2,622.92	1,881.95					
Upper Confidence Interval (MXN)			20.25	21.35	0.03	0.03	490.77
Upper Confidence Interval Divided by Abs. Value of Change in Spending			0.0077	0.0081	0.0000	0.0000	0.1871
Lower Confidence Interval (MXN)			-24.80	-25.83	-0.01	-0.01	-613.19
Lower Confidence Interval Divided by Abs. Value of Change in Spending			-0.0095	-0.0098	-0.0000	-0.0000	-0.2338
N= 45709							

This table shows treatment effects on spending, saving, and borrowing for individuals in the top quartile of predicted treatment effects on spending who have a credit card billing cycle fully covered by the intervention. For each individual we identify the calendar weeks corresponding to said billing cycle and we look at spending and saving during those weeks. Spending is defined as the sum of ATM withdrawals, credit and debit card spending, and outgoing transfers. Checking Account Balances correspond to average daily balances over the period. Credit Card Interest corresponds to the average of the monthly interest charges for interest-bearing balances during the billing cycle fully covered by the intervention. Paid Interest is a binary variable flagging when an individual was charged credit card interest during the billing cycle fully covered by the intervention. Statement Balance - Payments After Treat. correspond to the balance on a credit card on the last day of the last billing cycle affected by the intervention minus payments received on the credit card account during the billing cycle immediately following the intervention (i.e., the payments made towards the bill of that previous billing cycle). Panel A considers all individuals who have a credit card. Panel B considers only individuals who have a credit card and incurred interest during at least one of the 6 months previous to the intervention. Average treatment effects are calculated with the AIPW method. 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.[1] 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 ASD number of observations is the same across all columns in the same panel. Robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01.

B.5 Treatment Effects by Week

We now explore the treatment effect on saving and spending over the course of the treatment weeks. We calculate the treatment effects on spending and savings by week which are displayed in Figures B5 and B6. 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.

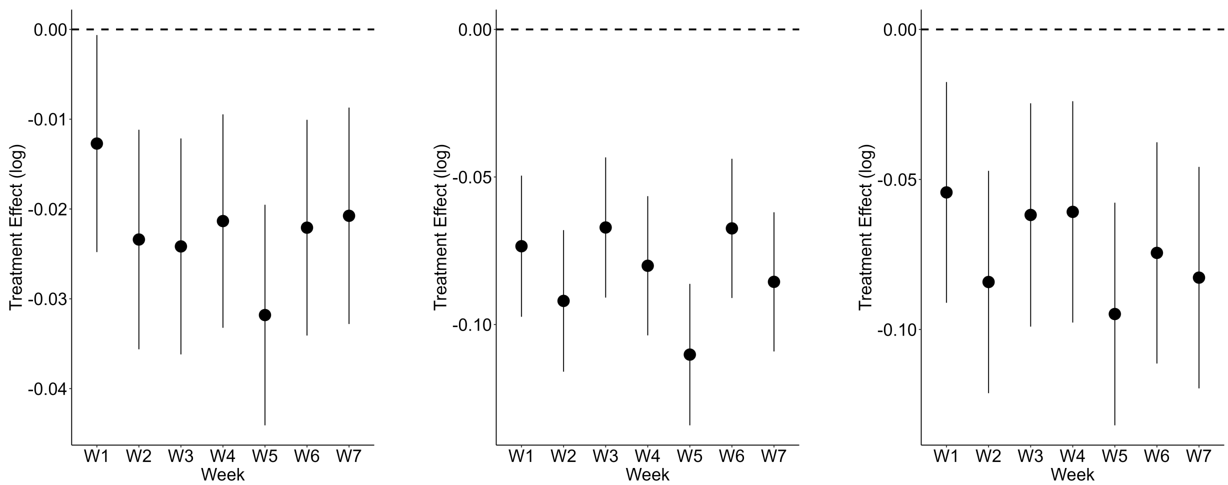

(a) All individuals with a Credit Card

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

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

Figure B5: This figure shows the treatment effect on spending during each of the seven weeks of the treatment. Average treatment effects are estimated using the natural log of spending, as the dependent variable. Spending is defined as the sum of ATM withdrawals, credit and debit card spending and outgoing transfers. For each week, we estimate the AIPW treatment effect for the subpopulation of interest and plot the resulting coefficients and 95% confidence intervals calculated using robust standard errors. We focus on individuals that have a credit card, those individuals in the top quartile of the distribution of treatment effects on spending as predicted by the causal forest, and those that also pay credit card interest at baseline.

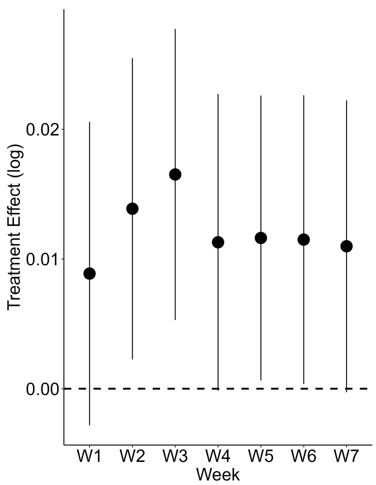

(a) All individuals with a Credit Card

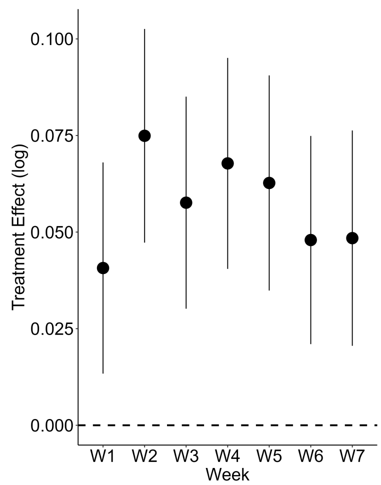

(b) Individuals with Credit Card in Top Quartile of the Predicted Treatment Effect Distribution

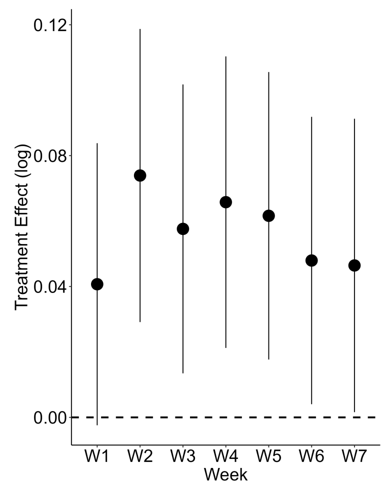

(c) Individuals with Credit Card paying Interest in Top Quartile of the Predicted Treatment Effect Distribution

Figure B6: This figure shows the treatment effect on savings during each of the seven weeks of the treatment. Average treatment effects are estimated using the natural log of average daily balances on a checking account during the corresponding week minus starting values, as the dependent variable. For each week, we estimate the AIPW treatment effect for the subpopulation of interest and plot the resulting coefficients and 95% confidence intervals calculated using robust standard errors. We focus on individuals that have a credit card, those individuals in the top quartile of the distribution of treatment effects on saving as predicted by the causal forest, and those that also pay credit card interest at baseline.

Appendix C Models: Formal Details and Derivations

Transaction-convenience model:

Two periods, one consumption good, log utility, individuals consume in periods 1 and 2, $c_{1,2}$, in period 1 they may borrow to consume b_1 because they must hold a certain amount of cash for transaction purposes x

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

subject to $x_1 - c_1 > x$ and $b_1 < b$. Suppose the agent is not credit-constrained $b = \infty$, if $x_1 - c_1^* \geqq x$, then $b_1 = 0$, if $x_1 - c_1^* < x$, then $b_1 = c_1^* - x_1 + x$, the optimal solution for c_1^* is:

$$\text{if } x_1 - c_1^* \geqq x \Rightarrow b_1 = 0 \text{ then } \frac{1}{c_1} = \delta \frac{1}{x_1 - c_1} \Rightarrow c_1 = \frac{1}{\delta}(x_1 - c_1) \Rightarrow c_1^* = \frac{1}{\delta + 1} x_1$$

$$\text{and if } x_1 - c_1^* < x \Rightarrow b_1 > 0 \text{ then } \frac{1}{c_1} = \delta \frac{1 + r}{x_1 - c_1 - r(c_1 - x_1 + x)} \Rightarrow c_1 = \frac{1}{\delta(1 + r)}(x_1 - c_1 - r(c_1 - x_1 + x))$$

$$\Rightarrow c_1(1 + \frac{1}{\delta(1 + r)}(1 + r)) = \frac{1}{\delta(1 + r)}(x_1 + rx_1 - rx) \Rightarrow c_1^* = \frac{1}{\delta + 1} x_1 - \frac{r}{(\delta + 1)(1 + r)} x \text{ and } b_1 = c_1^* - x_1 + x$$

Comparative statics with respect to x:

- if $x_1 - c_1^* \geqq x$ then $\frac{\partial c_1^*}{\partial x} = 0$ and if $x_1 - c_1^* < x$ then $\frac{\partial c_1^*}{\partial x} = \frac{-r}{(\delta+1)(1+r)} < 0$ this is because the agent is marginally poorer if they have to borrow to maintain their larger cash needs, zero if $r = 0$
- if $x_1 - c_1^* \geqq x$ then $\frac{\partial b_1}{\partial x} = 0$ and if $x_1 - c_1^* < x$ then $\frac{\partial b_1}{\partial x} = 1 + \frac{\partial c_1^*}{\partial x}$ the agent's cash needs are directly reflected in their borrowing but because they consume slightly less because they are marginally poorer, he borrows a bit less, borrowing increase is equal to increase in x if $r = 0$

Comparative statics with respect to δ:

- if $x_1 - c_1^* \geqq x$ then $\frac{\partial c_1^*}{\partial \delta} = -\frac{1}{(\delta+1)^2} x_1 < 0$ and if $x_1 - c_1^* < x$ then $\frac{\partial c_1^*}{\partial \delta} = \frac{-(1+r)x_1 + rx}{(\delta+1)^2(1+r)} < 0$ (see condition of model solvability, $x_1 > x$) the agent consumes less if they are more patient and wants to save more
- if $x_1 - c_1^* \geqq x$ then $\frac{\partial b_1}{\partial \delta} = 0$ and if $x_1 - c_1^* < x$ then $\frac{\partial b_1}{\partial \delta} = \frac{\partial (c_1^* - x_1 + x)}{\partial \delta} = \frac{\partial c_1^*}{\partial \delta} < 0$ the agent consumes less if they are more patient so they borrow less

Uncertainty-about-credit-limits model:

Three periods, one consumption good, log utility, individuals consume only in periods 2 and 3, $c_{2,3}$, in period 1 and 2 they may borrow to consume $b_{1,2}$, if they do not borrow in period 1, they cannot borrow in period 2 either, i.e., subject to $b_2 \leq b_1$ and $r < 1$. Backward induction, start with period 2, x_2 and b_1 are cash-on-hand and debt-on-hand:

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

note that, if $b_2 = b_1$ as the only reason to borrow in period 1 is to ensure borrowing in period 2, the optimal solution for c_2^* is:

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

(given the constraint that $b_2 \leq b_1$, for the case in which $b_2^* > 0$, then $b_1 = b_2^*$):

$$\text{if } b_2 = b_1 > 0 \text{ then } b_2 = c_2^* - x_2 + rb_1 = b_1 \Rightarrow b_2 = b_1 = \frac{1}{1-r}(c_2^* - x_2)$$

$$\text{if } b_2 = b_1 > 0 \text{ then } \frac{1}{c_2} = \delta\frac{1}{y_3 + x_2 - c_2 - rb_1 - rb_2} \Rightarrow c_2 = \frac{1}{\delta}(y_3 + x_2 - c_2 - 2\frac{r}{1-r}(c_2 - x_2))$$

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

Comparative statics with respect to x_2 if $b_1 = b_2 > 0$:

- $\frac{\partial c_2^*}{\partial x_2} = \frac{1+r}{\delta(1-r)+1+r} > 0$ this is because the agent is richer and consumes more
- $\frac{\partial b_2^*}{\partial x_2} = -\frac{1}{1-r} < 0$ 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 $b_1 = b_2 > 0$:

- $\frac{\partial c_2^*}{\partial \delta} = -\frac{(1-r)^2}{(\delta(1-r)+1+r)^2}y_3 - \frac{(1-r)(1+r)}{(\delta(1-r)+1+r)^2}x_2 < 0$ the agent consumes less if they are more patient and wants to save more
- $\frac{\partial b_1^*}{\partial \delta} = \frac{\partial b_2^*}{\partial \delta} = \frac{1}{1-r}\frac{\partial c_2^*}{\partial \delta} < 0$ and borrows less

Self-control model:

Two periods, one consumption good, log utility, individuals consume in periods 1 and 2, $c_{1,2}$, in period 1 they may borrow to consume b_1, but when they hold a certain amount of cash in a savings account x, they consider that money locked away for future consumption and it gets subtracted from x_1 that the agent then uses as their decision variable

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

subject to $b_1 < b$. Suppose the agent is not credit-constrained $b = \infty$, and in some period 0, they lock away x in an inaccessible savings account for period 2 consumption, a fraction a the agent forgot about. If $x_1 - x - c_1^* \geqq 0 \Rightarrow x_1 - c_1^* \geqq x$, then $b_1 = 0$, if $x_1 - x - c_1^* < 0$, then $b_1 = c_1^* - x_1 + x$, the optimal solution for c_1^* is:

$$\text{if } x_1 - x - c_1^* \geqq 0 \Rightarrow b_1 = 0 \text{ then } \frac{1}{c_1} = \beta \frac{1}{x_1 - ax - c_1} \Rightarrow c_1 = \frac{1}{\beta}(x_1 - ax - c_1)$$

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

$$\text{and if } x_1 - x - c_1^* < 0 \Rightarrow b_1 > 0 \text{ then } \frac{1}{c_1} = \beta \frac{1+r}{x_1 - ax - c_1 - r(c_1 - x_1 + x)}$$

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

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

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

Note that, for the transaction-convenience agent, optimal consumption was $c_1^* = \frac{1}{\delta+1}x_1 - \frac{r}{(\delta+1)(1+r)}x$. 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 and consuming the same as the transaction-convenience agent. 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. Now, the amount withheld x is decided by a more patient self/spouse with the following problem.

$$max\{log(c_1^*) + \delta log(x_1 - c_1^* - rb_1)\}$$

$$\text{if } x_1 - x - c_1^* \geqq 0 \Rightarrow b_1 = 0 \text{ then } \max_x\{log(\frac{1}{\beta+1}(x_1 - ax)) + \delta log(x_1 - \frac{1}{\beta+1}(x_1 - ax))\}$$

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

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

$$\Rightarrow x^* = x_1\frac{\delta - \beta}{(\delta+1)a} > 0 \text{ as } a, \beta \in [0,1] \text{ and } \delta > \beta$$

this equals $x = \frac{\delta}{\delta+1}x_1$ if $a = 1$ and $\beta = 0$ exactly same as what rational agent would allocate to first and second period consumption as then $c_2 = \frac{\delta}{\delta+1}x_1$ and $c_1 = (1 - \frac{\delta}{\delta+1})x_1 = \frac{1}{\delta+1}x_1$

$$\text{and if } x_1 - x - c_1^* < 0 \Rightarrow b_1 > 0 \text{ then } \max_x\{log(c_1^*) + \delta log(x_1 - c_1^* - rb_1)\}$$

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

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

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

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

this equals $x = \frac{\delta}{1+\delta}$ if $a = 1$ and $\beta = 0$ which is the same as what rational agent would allocate who faces cash needs; what if $a \to 0$ and $\beta = 0$: $x_1\frac{\delta}{1+\delta}\frac{1+r}{r}$, the agent takes the interest costs into account when they lock away some money, so if $r \uparrow$ then $x \downarrow$ as $\frac{\partial\frac{1+r}{r}}{\partial r} = -\frac{1}{r^2} < 0$.
Comparative statics of impatient self/spouse with respect to x:

- if $x_1 - x - c_1^* \geqq 0$ then $\frac{\partial c_1^*}{\partial x} = -\frac{1}{\beta+1}a$ and if $x_1 - x - c_1^* < 0$ then $\frac{\partial c_1^*}{\partial x} = -\frac{r+a}{(\beta+1)(1+r)} < 0$ there are two factors: one, the moment cash is taken away, and the separation-of-accounts friction a is not 0, then $\frac{\partial c_1^*}{\partial x} < 0$ and $\frac{\partial c_1^*}{\partial x} = -\frac{1+r}{(\beta+1)(1+r)} < -\frac{r}{(\beta+1)(1+r)}$ the separation-of-accounts sensitivity of consumption is negative and large because the agent consider the savings 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 (if $\beta = 0$ and $a = 1$ then it equals -1)

- if $x_1 - x - c_1^* \geqq 0$ then $\frac{\partial b_1}{\partial x} = 0$ and if $x_1 - x - c_1^* < 0$ then $\frac{\partial b_1}{\partial x} = \frac{\partial c_1^*}{\partial x} + 1$ so when the agent's consumption goes down a lot in response to an increase in x then the effect on borrowing is very little, if $\beta = 0$ (the agent always consumes everything immediately) and $a = 1$ then $\frac{\partial c_1^*}{\partial x} = -1$ and $\frac{\partial b_1}{\partial x} = 0$

Comparative statics with respect to δ (discount factor of patient self/spouse):

- if $x_1 - x - c_1^* \geqq 0$ then $\frac{\partial x}{\partial \delta} = \frac{a(\delta+1) - a(\delta-\beta)}{((\delta+1)a)^2}x_1 > 0$ and if $x_1 - x - c_1^* < 0$ then $\frac{\partial x}{\partial \delta} = -x_1\frac{1}{(1+\delta)^2}(\delta\frac{1+r}{r+a} - \frac{\beta+r\beta}{a-r\beta}) + x_1\frac{1}{1+\delta}\frac{1+r}{r+a} = \frac{1}{1+\delta}(-x + x_1\underbrace{\frac{1+r}{r+a}}_{>1})) \in [\frac{1}{1+\delta}(-x + x_1), \frac{1}{1+\delta}(-x + x_1\frac{1+r}{r})] > 0$ if the patient self/spouse is more patient then they withdraw more money, so the impatient self consumes less as $\frac{\partial c_1^*}{\partial x} \leqq 0$ and $\frac{\partial b_1}{\partial x} = \frac{\partial c_1^*}{\partial x} + 1$ so when the impatient agent's consumption goes down a lot in response to an increase in x then the effect on borrowing is very little