Selective Attention in Consumer Finance: Evidence from a Randomized Intervention in the Credit Card Market*

Paolina C. Medina†

January 16, 2017

Abstract

In partnership with a Personal Finance Platform in Brazil, I implement a randomized intervention to measure the effect of reminders for timely payment of credit cards. While I find an 13% reduction in the cost of late payment fees paid, 31% of the users that avoid credit card late payments, incur instead checking account overdraft fees. This behavior leads to heterogeneous gains from the intervention with some users saving 15% in total fees paid, and others incurring increased fees of 5%. I analyze these results using theories of selective attention, and argue that when multiple tasks need to be performed, reminders that increase information about one task may crowd out attention from other less salient but equally important tasks. The results of this experiment suggest that when designing policy interventions, one size may not fit all, and targeting nudges to those who are more likely to benefit has the potential to significantly increase the gains from the intervention.

*I am grateful to Effi Benmelech, Meghan Busse, Leemore Dafny, Brian Melzer and Richard Thaler for their very valuable guidance and support. I also thank Hasat Cakkalkurt, Anthony DeFusco, Jeff Ely, Sean Hundtofte, Cynthia Kinnan, Daniel Martin, David Matsa, Takeshi Murooka, John Mondragon, Paola Sapienza, Andrei Shleifer and seminar participants at Kellogg-Finance, Kellogg-Strategy, Northwestern Econ and LMU Munich lunch seminars; 2016 Whitebox Conference; Banco de Mexico and 2016 IPA Financial Inclusion Gathering, for their comments. All remaining mistakes should be blamed on my own selective attention. IRB ID: STU00201364.

†Northwestern University. Email: p-medina@kellogg.northwestern.edu

1 Introduction

The role of limited attention has been studied in economics in a variety of settings ranging from monetary policy and labor search, to public finance and income inequality. The typical application assumes that attention is optimally allocated to acquire and process information from different stimuli, and that, from an individual decision making perspective, more information can do no harm. The free-disposal property of information has motivated the use of informational nudges which have proven effective in influencing behavior in a wide range of applications such as consumer finance (Beshears et al., 2015), health (Vervloet et al., 2012), and environmental policy (Allcott and Kessler, 2015). But, could nudging with information have unintended consequences? Recent findings from psychology and neuroscience show that attention affects consumers perception of what is important, and that consumers' allocation of attention can be distorted by the salience of alternative options. I argue that when consumers have to attend to multiple tasks, informational nudges focusing on one task may crowd out attention from other equally important, but less salient tasks.

This paper explores the effects of single-task reminders in the context of consumer finance. Specifically, I focus on the effect of reminders for timely credit card payments across two high stakes margins: credit card late payment fees, which represent 41% of the total cost of credit card fees (Agarwal et al., 2015); and overdraft fees, which make up a more than $23 billion business (Melzer and Morgan, 2015). To do so, I design a field experiment in which credit cardholders are exposed to "attention shocks" about future credit card payments and due dates. These attention shocks take place at different stages of the credit card billing cycle and take the form of smart-phone push notifications. The experimental pool is made up of users of a personal finance management platform in Brazil that collects transactional level data from both checking accounts and credit cards, for each user over time.

When evaluating the effects of the intervention on credit card performance, the treatment lowers the probability of being charged a late payment fee by 2.6 percentage points from a basis of 29.1%, and reduces the cost of fees paid by

12.9% from a basis of $R32.94. However, I also find that the cost of overdraft fees paid increases by 10.5% from a basis of $R25.96. This effect is driven by users that have paid overdraft in the past, and whose probability of paying overdraft during the treatment period increases by 2.9 percentage points, from a basis of 76%. Furthermore, for this group of users, substituting late payments for overdraft leads on average to a 5% increase in the total amount of fees paid, from a basis of $R129.48. On the other hand, for the 71% of users that did not incur overdraft at baseline, the intervention reduced the total cost of fees paid by 15% from a basis of $R30.48. These gains and loses average to a 2.5% reduction in the total cost of contingent fees for the treatment group.

This study constitutes, to my knowledge, the first evidence showing the existence of trade-offs when influencing consumer behavior with nudges. As opposed to Chetty et al. (2014) who find that nudging to save for retirement with automatic contributions does not crowd out savings from other financial instruments, I find that nudging individuals to avoid late payments in their credit card leads to increased overdraft fees in their checking account. I argue that the no-crowd-out result found by Chetty et al. (2014) using Danish data cannot be generalized to settings where the resource affected by the nudge is close to a relevant trade-off margin. One such margin is negative income: for individuals on a tight budget, the opportunity cost of essential consumption is high, and allocating resources for additional savings (or debt payback) would have to come at the expense of savings in other instruments (or increased debt in other products). Liquidity constraints in Brazil are indeed more likely to bind compared to Denmark: according to the World Bank, among the richest 60% of Brazilians, one third would not be able to come up with funds for an emergency, while the corresponding figure is 5% for Denmark. [1]

To investigate the nature and optimality of the trade-off between late payments and overdraft fees, Section 6 argues that under certain assumptions, this trade-off can be considered an overreaction to reminders caused by a distortion in consumers' allocation of attention towards the most salient tasks. This argu-

[1] Source: Global Findex 2014, The World Bank.

ment is based on findings from psychology and neuroscience, which differ from the standard treatment of attention in economics in two main ways. First, attention is shown to affect consumer choice by altering the weights assigned to different product attributes in consumer's decision utility: attributes that receive more attention are considered more important. Second, consumers' allocation of attention towards different product attributes is allowed to be distorted by variables such as priming, salience or availability (Bordalo, Gennaioli and Shleifer (2013), Hare, Malmaud and Rangel (2011), Fehr and Rangel (2011),Thaler and Sunstein (2003)). Specifically, I interpret the empirical findings of this study using the intuition from Bordalo, Gennaioli and Shleifer (2015) theory of selective attention under which reminders have both an informational and psychological effect. In their framework some consumers are "Forgetful But Otherwise Rational" and others are "Forgetful And Salient Thinkers". While the former group can only benefit from receiving a reminder, the latter may instead overreact, distorting their allocation of attention towards the most salient option in their choice set.

The results of this experiment suggest that when designing policy interventions, one size may not fit all, and targeting nudges to those who are more likely to benefit from them, given their observable characteristics, can lead to larger welfare gains. One such targeting strategy consists of establishing differentiated defaults under which users whose observable characteristics predict overreaction to reminders are enrolled in an opt-in program, while the remaining fraction of users who are likely to benefit from the intervention are enrolled in an opt-out reminder based program.

I complement the analysis by presenting repayment and spending patterns that reveal details about the channels through which attention shocks affect consumer behavior. Even though some reminders were sent early in the billing cycle to encourage users to set some money aside for their next payment, the treatment effect on late payments is not driven by changes in spending patterns or changes in resources available at the time the payments are due. Instead, it is driven by increased attention towards credit card payments conditional on available liquidity. Furthermore, the effect almost fully disappears in the billing cycle immediately after the intervention when no additional messages were sent, and seems to be driven by "short-term delinquents" who, in the absence of the treatment, would have paid their bill within 30 days of their due date.

Finally, while various papers have tested the effect of attention shocks and reminders across settings ranging from overdraft use (Stango and Zinman, 2014), microcredit loan payments (Cadena and Schoar, 2011), savings (Karlan et al., Forthcoming), and Islamic credit cards (Bursztyn et al., 2015), this paper is, to my knowledge, the first test of the effect of reminders in credit card repayment decisions when card holders are in good standing, and with a control group that receives no message. Importantly, and as opposed to previous work using reminders, this paper looks at a richer set of outcomes that span not only direct effects on credit card performance, but also indirect effects on checking account outcomes.

The rest of the paper is organized as follows: In Section 2, I provide background on credit card late payment fees, and review the related literature. Section 3 describes the data and the Partner company. I present the experimental design in Section 4 and the empirical results in Section 5. Sections 6 and 7 discuss potential psychological mechanisms and policy implications. Section 8 concludes.

2 Background and literature review

2.1 Institutional background

Credit cards offer a form of revolving credit, in which consumers have the flexibility of borrowing and paying back continuously as long as the total amount borrowed is less than the individually assigned credit limit, and as long as at least the minimum payment is paid back no later than an individually assigned due date.

The billing cycle for the typical credit card lasts approximately 30 days. After this period, the minimum payment required to stay current is calculated taking into account the amount of money borrowed during the cycle. Such minimum payment needs to be covered in the so called grace period, which ends on the credit card due date. Failure to cover the minimum payment by the due date leads to a late payment fee.

Figure 1: Typical credit card billing cycle

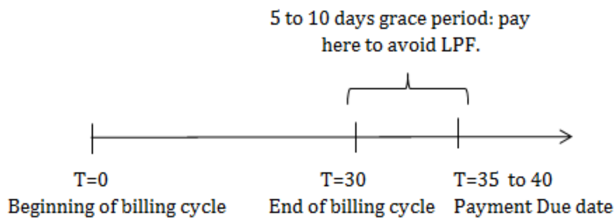

As for the specific case of Brazil, the grace period on a credit card lasts 5 to 10 days.[2] Credit cards late payment fees are regulated to be no more than 2% of the full balance in the account, plus an additional moratory interest of 1% per month. Importantly these fees and interest are charged on top of the regular monthly revolving interest which is usually higher than 10% per month.[3] For example, an individual with a balance of $R2,000 that misses a payment, and pays 15 days after his due date will be charged a fee of $R40, plus moratory interest of $R10 for the 15 days overdue. On top of that, the regular interest charge would be $R200. Not surprisingly, given these very high interest rates, the fraction of individuals that pay in full is close to 93% conditional on making a payment. Among those that do not pay in full, 78% are paying late[4]. In this example, had the individual paid in full and on time, he would have saved $R250.

Despite these very high penalties as of the end of 2015, the rate of delinquent credit cards in the Brazilian market was 39%.[5] What can explain this high delinquency rate? In the context of high liquidity constraints, we would expect that some people would simply not be able to make their payment, or would strategically choose not to pay. However, it is important to notice that, due to the revolving nature of credit cards, for individuals with a sufficiently large probability of eventually paying back their credit card balances, liquidity constraints are not enough to justify a late payment: when an individual makes a payment on his credit card, he frees up a fraction of credit limit equal to the payment made and, therefore,

[2] In the US, grace periods for the major banks are approximately 27 days.

[3] Brazilian Central Bank and Ferman (Forthcoming).

[4] https://www.spcbrasil.org.br/uploads/st_imprensa/analise_inadimplencia_e_cartao_de_credito.pdf

[5] http://agenciabrasil.ebc.com.br/economia/noticia/2015-11/inadimplencia-com-rotativo-do-cartao-de-credito-atingiu-389-em-setembro

that same amount of money can still be used for whatever liquidity needs the user may have, as long as credit cards are accepted as a payment form. Therefore, it is not only high liquidity constraints, but more precisely high cash needs that may, in fact, justify the need to miss a credit card payment.

The nature of the data used in this project allows me to observe not only information on credit card balances and due dates for the selected pool of experimental users, but also the corresponding checking account balances and transactions. I find that 36% of users charged a late payment fee had enough balance in their checking account, at the time of their credit card due date, to cover their credit card balance in full. This observation alone motivates the widespread idea that some of these late payments may be the result of consumers' mistakes, potentially in the form of limited attention to late payment fees or to credit card due dates.

On the one hand, some recent observations, both empirical and from the structure of credit card contracts, support the idea that consumers do not pay attention to late payment fees. Gabaix and Laibson (2006), and Heidhues, Koszegi and Murooka (Forthcoming) present models in which not advertising certain product attributes (e.g., credit card late payment fees) when a fraction of myopic consumers do not pay attention to these attributes when shrouded, is in fact an equilibrium outcome. Furthermore, recent empirical evidence shows that decreasing the level of late payment fees, does not lead to an increase in late payments (Agarwal et al., 2015).

Instead, attention shocks coming from experiencing credit card fees in the past, can affect the incidence of credit card fees in the future. Agarwal et al. (2013) present a model in which consumers build a stock of "knowledge" (that depreciates over time) on how to avoid late payment fees. Using a sample of US cardholders, they find that incurring a late payment fee, reduces the probability of incurring a late payment fee during the subsequent month by 44 percent.

In this paper, I argue that a fraction of credit card late payments is caused by consumers' limited attention. I do that by exogenously imposing attention shocks in the form of reminders and studying the reaction of consumers in a variety of outcomes.

2.2 Selective attention and reminders in consumer financial markets

The benchmark framework for the analysis of limited attention in economics is based on the rational inattention model in which attention is a scarce resource that is allocated towards different signals as a result of a frictionless optimization process. This framework has been applied in a variety of settings ranging from monetary policy (Sims, 2003) and labor markets (Maćkowiak and Wiederholt, 2015) to income inequality (Banerjee and Mullainathan, 2008).

An alternative view based on results from psychology and neuroscience has two main components. First, attention is thought to play a role in the computation of value at the time of choice, affecting the weights assigned to different attributes in the decision utility. Second, the allocation of attention is allowed to be distorted by supposedly irrelevant factors through priming, salience or availability (Bordalo, Gennaioli and Shleifer (2013), Hare, Malmaud and Rangel (2011), Fehr and Rangel (2011), Thaler (2015), Thaler and Sunstein (2003)). These models also emphasize the selective nature of attention by which mental effort is allocated to the processing of some stimuli in preference to others (Kahneman, 1973).

The policy relevance of limited attention either in its rational or quasi-rational form has been discussed by Fishman and Hagerty (2003), Hirshleifer and Teoh (2003), and Dranove and Jin (2010) among others. A growing literature looks at the role of attention shocks on consumer decision making in the telecommunications (Grubb, 2015), fitness (Calzolari and Nardotto, 2011), medical (Vervloet et al., 2012) and electricity (Allcott and Rogers, 2014) markets.

In the context of consumer financial markets, Stango and Zinman (2014) is the first paper to look at behavioral responses to attention shocks using micro-data. Specifically, they focus on the overdraft market. In their context, consumers' attention is affected when taking a marketing survey that in a quasi-random fashion includes questions related to overdraft fees. They find that consumers receiving overdraft related questions are about 12.3% less likely to incur overdraft during the survey month. This effect lasts and depreciates over the following 2 years.

A number of papers have looked at the effect of attention shocks in the form of cell-phone text messages in randomized control trials. Karlan et al. (Forthcoming) remind individuals about future lumpy expenses (such as health emergencies or a durable good to be purchased), and find significant increases in saving rates on commitment savings products in Bolivia, Peru and the Philipines. Alan et al. (2015) find that when advertising a discount in the cost of overdraft via text message, consumer demand for overdraft decreases, as predicted by shrouded equilibrium models. As for reminders to repay existing debt, Karlan, Morten and Zinman (2015) investigate the effect of messages with different content reminding individuals to make their micro-loan payment on time, and find the largest effect for messages containing "a personal touch" in which the name of the loan officer is included in the text. These results are consistent with Cadena and Schoar (2011) who find that the size of the effect of reminders for repayment of micro-loans in Uganda is in the same order of magnitude as the size of the effect of a 25% reduction in the cost of capital. However, so far only Bursztyn et al. (2015), and this paper have looked at reminders for repayment in the context of revolving credit.

In Bursztyn et al. (2015), the key manipulation tries to disentangle the roles of morality and religion among holders of an Islamic credit card. Their experimental pool is made up of users who were already late in their payments and who had already been asked with a text message to pay back to their bank. They find that one additional message emphasizing the morality of repayment increases the fraction of people meeting their minimum payment by nearly 20%. This message is more effective than substantial cash rebates and is also more effective than one additional neutral (non-moral) reminder. In contrast, in this paper the key manipulation across treatments is in the timing and number of reminders sent over the billing cycle. The experimental pool is made up of users that are both current or late in their payments, and the control group receives no message. Importantly, this paper looks at a richer set of outcomes that span not only direct effects on credit card performance, but also indirect effects on checking account outcomes.

This paper contributes to a better understanding of the role of nudges to redirect attention towards financial tasks in two main ways. On the one hand, given the nature of the data, I am able to look at a wider set of outcomes, including individual spending, checking account balances, and spillovers to other types of fees.[6]

On the other hand, in terms of the broader literature on the allocation of attention towards different stimuli and the effect of reminders (inclusive of non-financial settings), this paper provides some light in to the psychological mechanisms through which reminders affect behavior. Under certain assumptions, the results are not consistent with the fully-rational inattention model, and instead can be interpreted through theories of selective attention under which the choice of which stimulus to pay attention to is affected by its salience. Specifically, I find that the taxonomy proposed by Bordalo, Gennaioli and Shleifer (2015) matches and explains the heterogeneous overreaction to reminders, in which the salience of an option may affect the behavior of some users in a way that Forgetful But Otherwise Rational users would not be affected.[7]

In that sense, while most of the results presented in this paper take an agnostic approach on the nature and sources of inattention, Section 6 argues that when jointly considered, and under certain reasonable assumptions, the results provide supporting evidence to theories of selective attention where individuals' allocation of attention is disproportionally drawn to tasks that are more salient or stand out in their choice set.

3 Partner company and data

The Partner company (the "Partner" from now on) is a personal finance management (PFM) platform operating in Brazil since August 2014. It currently has a user base of over 2 million users. The Partner operates on two platforms, a website and a smart phone app, and through these platforms, offers various services.

[6]While other papers such as Baker (2015), Kuchler (2015), Gelman et al. (2014) and Pagel and Vardardottir (2016) have used data from similar providers, this is the first among those to introduce experimental variation to the analysis of data from a Personal Finance Management company.

[7]Importantly, this paper does not use the definition of salience introduced in Bordalo, Gennaioli and Shleifer (2013) and Bordalo, Gennaioli and Shleifer (2015), but instead given the context of the application in which a prospective memory task is receiving an attention shock I consider a stimulus to be more salient if it is more vivid, visible or available than other stimuli. This is similar to the definition used by Chetty, Looney and Kroft (2009) that in the context of taxation "define salience in terms of the visibility of the tax-inclusive price."

The main services consist of aggregation of information across accounts and banks: when a user signs up, he registers his accounts on the platform and provides the Partner with his credentials to log in to his online banking services. The partner has access to as many credit cards, checking accounts or savings accounts as the users register, and provides aggregate statistics of individuals' financial position (e.g., net balance (debit balances – credit balances), total value of deposits, and total cost of expenses related to groceries, leisure, etc.).

In addition, it allows for classification and tracking of expenses: for each registered account, the partner collects the date of the transaction, a string description of the transaction, and the monetary value of the transaction. It later classifies these transactions into one of 35 predefined categories, such as utilities, transportation, grocery, bars and restaurants, tuition, etc. Users can subsequently look up their total expenses in each of these categories.

Finally, the Partner also offers a budgeting tool that allows individuals to input into a friendly template, how much they plan to spend in each of the 35 predefined categories, or any other user-defined category. The Partner then classifies expenses in real time into the corresponding categories, and reports consumers' "net position" in each of the budget categories. Notably, only around 12% of the user base utilizes the budgeting tool.

The main channels through which the Partner interacts with its users are mobile surveys and messages sent through mobile push notifications. The surveys are displayed on the main screen of the phone once the user logs in to the app and are used for various market research purposes. The push notifications are routinely sent to invite consumers to log in to the app to track their expenses, or to make a budget at the beginning of the month.

In terms of data availability, the raw data for the analysis consists of transactional level data, login information, and snapshots of account balances. Specifically, for each transaction originating from a registered account, I observe its value, date, string description, account _id and user_id. Notably, the string description of each transaction allows identification of bank fees. Similarly, for each user I observe the

Figure 2: Partner website

Figure 3: Survey and Push notification. Mobile-phone screen examples.

a) Survey

b) Push notification

date and time of account creation,[8] the login channel (mobile or website), and date and time of login.[9] Finally, for each account I observe the type of account (credit card, or checking account), bank_id, user_id, and a snapshot of balances. A drawback of the dataset is that, for privacy reasons, it is not possible to get demographic information at the user level.

The data is collected from online banking platforms through the following protocol: when signing up, users provide the Partner with their credentials to access their own online banking services. With such information, and with appropriately encrypted protocols, the usual operational procedure consists of retrieving the transactional information of the user accounts whenever the user logs in to the PFM app or website. When such logins take place the Partner is able to retrieve three months of information. Importantly, in addition to the usual operational procedure, the Partner has access to a special "brute force" tool, that allows them to establish the connection with the banks and retrieve the same three months of information even when the user does not log in to the app. This "brute force" tool is crucial to ensure that observations in treatments and control groups are comparable.

4 Experimental design

4.1 Experimental pool and sample allocation

To run the experiment, the Partner company agreed to provide access to the information of 26,069 users, to send them push notifications containing information on credit card due dates and late payment fees over two billing cycles (September 25, 2015 to December 13, 2015), and to send each participant a survey ten days after receiving the last push notification. They also agreed to retrieve information with the "brute force" method at the end of October and November 2015, as well as in early January and February 2016.

[8]The date and time of account creation refers to the account created with the Partner. I do not observe the date in which bank accounts were created.

[9]More than 80% of interactions occur through the mobile app.

Figure 4: Experimental pool and timing of the intervention

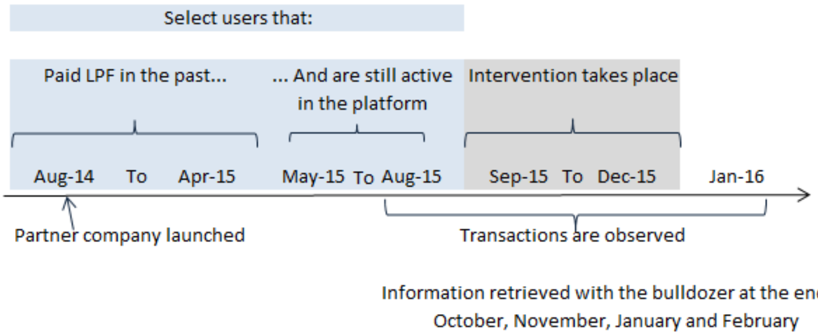

The experimental pool corresponds to the universe of users that were charged a late payment fee between August 2014 and April 2015, and whose last login was on or after May 2015. Access is granted for available transactions taking place between August 1, 2015 and January 31, 2016 (See Figure 4).

The main design consists of randomly assigning 26,069 users into four groups: one Control Group, two main Treatment Groups, and one additional Robustness or alternative-control Treatment Group. This last group is used to investigate explanations alternative to the main limited attention channel proposed in this paper.

Sample size (SS) is allocated to maximize the minimum detectable effect across the following comparisons: Each of the two main Treatments with each other and with the Robustness Treatment Group, and the pool of the two main Treatments with the Control Group. Ideally we would like to have:

$$SS_{t1} = SS_{t2} = SS_{robustness} = \frac{SS_{control}}{2}.$$

Due to some logistical limitations, SS_{t2} is 33% smaller than SS_{t1} and $SS_{robustness}$, deviating the sample allocation from the optimal, and leading to the following final allocation of users: Treatment Group 1 received 5,987 users, Treatment Group 2 received 4,038 users, the Robustness Treatment Group received 5,954, and finally the Control Group received 10,090 users (See Figure 5).

The sample is stratified according to the following variables: a dummy variable indicating whether the user had enough balance in his checking accounts at the time when he was charged an LPF, a dummy variable indicating whether or not

Figure 5: Sample distribution

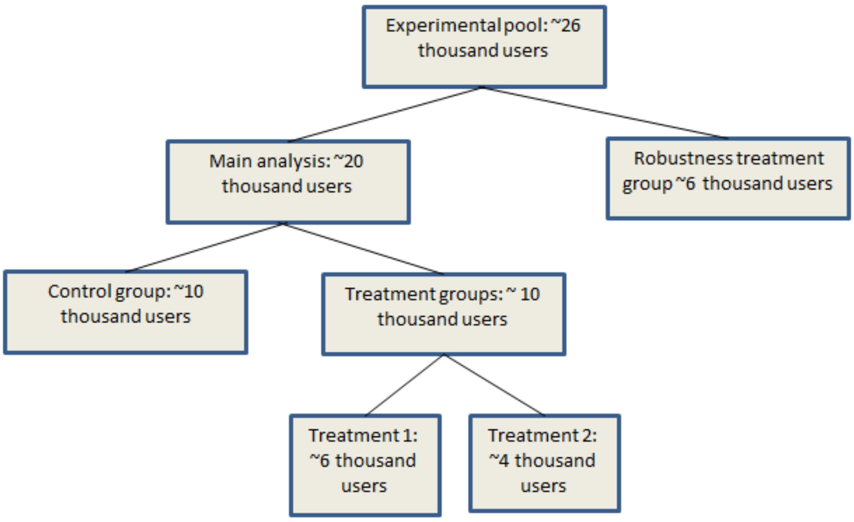

the user receives periodic paychecks, a discrete variable indicating the due date (DD from now on) on the main credit card of every user, and another dummy variable indicating whether or not the gap between DD and paycheck date is between 0 and 15.

4.2 Main Treatment Groups

For this intervention, five messages were defined to be sent at different moments of the credit card billing cycle. Three of these messages were designed to be sent early in the billing cycle, to encourage planning. In an attempt to strengthen individuals' mental accounts for credit card payments, the messages encourage users to set some money aside for their next credit card payment. These messages are referred to as "Planning" messages. Two more messages were designed to be sent close to the credit card due date, to invite consumers to make their credit card payment on time. This last set of messages is referred to as "DD Alert" messages. The specific content of the message can be found in Table A1. The different moments of the billing cycle in which the messages were sent can be found in Figure 6.

Treatment Group 1 received both Planning and DD Alert messages and is referred from now on as the Full-Treatment Group. Treatment Group 2 received only DD Alert messages and is referred to as the DD Alert Group from now on. Finally,

Figure 6: Timing of messages over the billing cycle. (Messages were sent in red periods.)

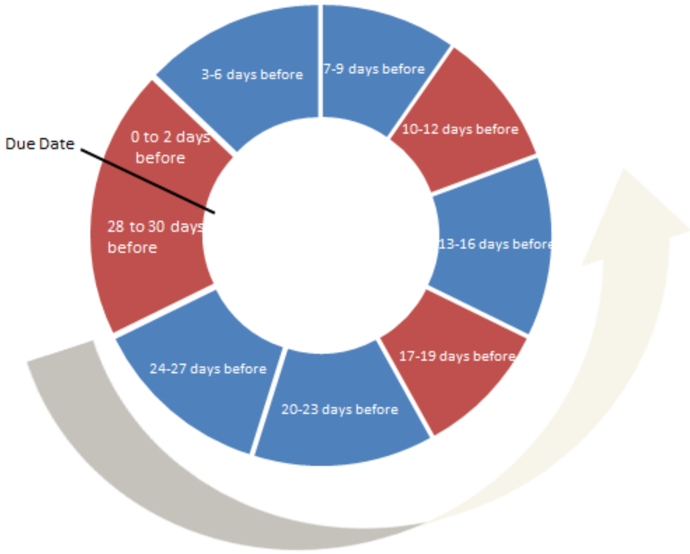

the Control Group received no messages. All users in the experimental pool were removed from regular message lists sent by the Partner, and therefore received no additional messages.

4.3 Robustness Treatment Group

In addition to the main experimental design, I set aside six thousand randomly selected users as part of a Robustness Treatment Group. Users in this group received different placebos in each of the treated billing cycles. In the first treated billing cycle, they received only the first planning message (24-27 days before the DD). In the second treated billing cycle, they received five messages from the pool of messages regularly sent by the Partner. These messages were chosen for having the same login rates as the main treatment messages, but contain no information about late payment fees or credit card due dates (See Table A2). The order of the messages was selected at random and one of these four messages was randomly selected to be sent twice and to ultimately keep the number of messages equal to the number of messages sent to the Full-Treatment Group.

The robustness group is used to investigate the alternative hypothesis that attention shocks are affecting delinquency only through an informational channel, by informing consumers of the cost of late payment fees. It is also used to investigate the possibility that the effect is coming from the fact that consumers are more informed in general about their finances as they log in more often to the Partner app as a result of the treatment. The details of how the robustness group is used to investigate these two explanations is described in section 6, which deals with the psychology and economics behind the treatment effects.

4.4 Randomization test and covariate balance

For the 26,069 users originally considered, transactional information was retrieved with the "brute force" method described in section 3, at the end of October and November 2015, as well as in early January and February 2016, with each retrieval providing information for the three prior months. Out of the 26,069 users provided by the partner company, only 13,538 users had transactional information available during the billing cycles considered in the analysis.[10] The large difference between the number of users originally considered in the study and the final number of users for which information is available seems to be caused by users canceling their accounts with the Partner company, changing their online banking passwords, connectivity issues between the Partner company and individual banks, and, importantly, changes in the security protocols of the second largest bank in the sample that no longer allows the Partner company to retrieve transactional information.

Fortunately, despite high attrition levels, unbiased inference only requires attrition to be uncorrelated with the treatment. To show that is the case, in Table 1, I present a simple comparison of means between treatments and control groups for the main variables of interest. Specifically I compare the mean value of these variables at baseline, but considering only observations for which information is available both at baseline and during the treatment period. These comparisons suggest that randomization worked and unbiased inference can be drawn from the experiment. To further familiarize the reader with the data, Table 2 presents additional descriptive statistics for the main variables used in the analysis.

[10] I consider that a user has available information when the total value of his monthly transactions is different from zero, and it has at least one transaction in the first third of the period and at least one transaction in the last third of that period.

5 Results

I organize the results of the analysis in three categories. First I describe the effect of the treatment on the incidence of credit card late payments. Specifically, I explore the channels through which the treatment leads to reductions in late payments. I find that the treatment effect is contingent on available liquidity at the time the payment is due: the treatment effect on late payments is only present for users that had enough balance in their checking account at the time their credit card payment was due, and the treatment has no effect on spending patterns nor on balances at the end of the billing cycle. I also find that the effect disappears almost fully in the billing cycle immediately after the intervention, and it seems to be driven by "short-term delinquents" that in the absence of the treatment would have paid their bill within 30 days of their due date.

Then, I explore the treatment effect of the intervention on other types of fees. Specifically, I focus on overdraft fees and find that in the intervention does not lead to increased overdraft use in the extensive margin, but it does lead to a 10.5% increase in the cost of overdraft fees paid (intensive margin). These aggregate results mask strong heterogeneities across users with different baseline overdraft patterns, with a large and significant increase for individuals that have incurred overdraft in the past; as well as slight but not significant differential treatment effects across users with different login patterns at baseline.

Finally, I consider the total amount of contingent fees paid by consumers in credit card and checking accounts and find heterogeneous benefit gains from the intervention. Specifically I explore the login-activity margin, and the overdraft-activity margin, and find that while users with a clean history benefit significantly from the intervention, those that have paid overdraft in the past end up with a negative effect from the intervention. On average, the results are close to canceling out, leading to a not statistically significant decrease on the total amount of fees paid by consumers.

These results highlight the importance of considering a rich set of outcomes when implementing treatment evaluations of policies designed to empower consumers.

5.1 Late payment fees

Table 3 shows the results of regressing an indicator for late payment fees, on each of the treatment indicators, with and without individual fixed effects. In columns 1-3, the specification consists of a simple linear regression without controls and without individual fixed effects. As can be seen, the average treatment effect of the pooled treatments is strongly significant. Furthermore, the Full-Treatment Group has a significantly larger treatment effect, compared to the DD Alert Group.

Taking advantage of the availability of baseline data, for columns 4-6, I introduce the preferred specification for the rest of the analysis that includes individual level fixed effects, to absorb any remaining variation not captured by the randomized assignment of individuals to treatment and Control Groups. The resulting specification is a difference in difference one, in which the coefficient of interest is given by the interaction of the variable *During* which separates treatment periods from baseline periods, and the corresponding Treatment Group indicator, denoted below by T:

$$Outcome_{it} = \alpha + \alpha_i + \beta_1 * During_t + \beta_2 * T_i + \gamma * During_t * T_i + \epsilon_{it} \tag{1}$$

Under the preferred specification, reminders reduce late payments by a magnitude between 3.42 and 1.39 percentage points (pp hereafter) from a basis of 29.1% depending on the treatment. These reductions correspond to 11.75% and 4.77% from control levels. Furthermore, the Full-Treatment leads to an effect that is statistically significantly larger than the effect generated by the DD Alert Treatment. Back of the envelope calculations show that the average marginal benefit of sending up to five reminders over one month is positive and decreasing, with the first two messages having an average marginal effect of 0.695pp, and the subsequent three messages having an average marginal effect of 0.676pp, corresponding to a 2.73% reduction in the average effect of one additional message.

To investigate the channels through which attention shocks over the billing cycle affect consumer behavior, I explore two hypotheses:

Hypothesis 1: The messages induce individuals to change their spending patterns to reach the time for payment with larger balances in their checking account, making them better able to fulfill their repayment responsibilities.

Hypothesis 2: The messages increase individuals' attention towards credit card payments, and operate only conditional on having enough balance in their checking account at the time when the payment is due.

I find precisely estimated null treatment effects on balances and spending patterns. Instead, the evidence suggests that the treatment effect operates conditional on available liquidity. Furthermore, a more intensive treatment in which more messages are sent leads to larger treatment effects on login rates into the app, which I use to proxy for increased attention towards credit card payments.

Table 4 presents the results of estimating the same basic regression, using the log of total expenses in selected categories as the outcome variable. The basis of the analysis are the 35 automatic categories in which the Partner company classifies transactions, grouped in to four economically meaningful classes. This classification does not pretend to be exhaustive, and is not always mutually exclusive, but is intended to represent different degrees of discretion in spending.

Under "Essential expenses" I group home expenses (rent and/or mortgage payments), residential bills, health, education, TV/internet/phone, transportation, work expenses, groceries and services. I consider house keeping, gifts/donations, leisure, personal care, bars/restaurants and shopping as "Discretionary expenses." And finally, under "Short Run Consumables" I include coffee shops, bars and alcohol, restaurants and fast food. Even when credit card payments account for $R1,318 on average (conditional on making a payment), and overall monthly spending accounts to R$8,859, I do not find a significant reduction in consumption across the main spending categories. This result is robust to different grouping of categories. The appendix provides a test for treatment effects of spending across all 29 non-financial categories automatically created by the Partner company.

Furthermore, even when three of the Full-Treatment messages were sent early in the billing cycle to encourage users to set some money aside for their next credit card payment, there is no effect in credit card and checking account balances, measured at the time when credit card payments are due.

Finally, to look at the interaction between the checking account and credit card margins, I construct a variable of relative liquidity (EB) defined as an indicator variable that takes the value of one when an individual has enough balance in his checking account to cover the total balance in his credit card, and zero otherwise. The results show that instead of changing available liquidity, the magnitude of the treatment effect depends on existing levels of liquidity: To explore the role of liquidity, I interact individuals' relative liquidity with each treatment indicator, to test whether there are differences in the magnitude of the effect for individuals that have enough money in their checking account, and those who do not.[11]

The resulting specification is the following:

$$\begin{aligned} Outcome_{it} = \alpha + \alpha_i + \beta_1 * During_t + \beta_2 * T_i + \beta_3 * EB_{it} + \gamma_1 * During_t * T_i \\ + \gamma_2 * During_t * EB_{it} + \gamma_3 * T_i * EB_{it} + \phi During_t * T_i * EB_{it} + \epsilon_{it}, \end{aligned} \tag{2}$$

where the triple interaction ϕ is the coefficient of interest.

As can be seen in Table 6, the treatment effect is significantly larger for individuals that carry enough liquidity in their checking accounts, with the pool of treatments leading to a reduction of 3.63pp in the fraction of individuals paying late payment fees. This number should be benchmarked against the baseline fraction of individuals that have enough liquidity but still pay late their credit card, which is 21.86%. Furthermore, conditional on existing balances, the difference between treatment effectiveness accounts for 2.2pp. While this difference is not statistically different from zero, I argue that the effect is robust across a variety of specifications.

The results are qualitatively the same when exploring the very intuitive notion that the fraction of people charged a late payment fee while having enough bal-

[11]Notice that the concerns for potential endogeneity of balances can be ruled out on two grounds: first, the section above just showed that the three messages sent before the grace period had no effect on liquidity. Second, the DD Alert Group receives messages only at the end of the billing cycle, i.e. during the grace period in which payments are due, and therefore by construction, the balances at the end of the billing cycle are unaffected by the treatment.

ance in their checking account, should decrease depending on the intensity of the attention shocks. To conduct such a test, I construct as an outcome variable the interaction between having been charged a late fee, and having enough checking balance to cover the full balance due on the corresponding credit card. This evidence is presented in Table 7. It can be seen that the fraction of users with enough balance in their checking account to cover their credit card balance in full, reduces by 2.17pp when considering the pooled treatments. The persistent differences between the Full-Treatment and the DD Alert treatment even conditional on liquidity, suggest that individuals that receive more messages are more likely to pay, partly because they allocate more attention to their credit card payments.

To find evidence supporting the hypothesis that the effect is driven by a reallocation of attention towards credit card related tasks, I proxy for attention towards credit card payments with the number of logins in to the PFM. Table 8 shows that individuals in the Full-Treatment Group are 60% more likely to log into the app during the treatment period, compared to the DD Alert Group.

Given that login in more often in to the app leads to increasing attention towards overall finances in general and not exclusively towards credit card payments, I explore the alternative explanation that the results are driven by increased attention towards overall finances and not by increased attention towards credit card payments. When analyzing the effect of the placebo messages on the Robustness Treatment Group, I find in Table A8 that individuals receiving a simple credit card message in the first period, and five additional status quo messages in the second period, logged in as often as the Full-Treatment Group (column 1), but however have a significantly larger fraction of late payers (column 2). In fact, the fraction of late payers for this placebo group is not significantly different from that of the Control Group (column 3).

The Partner also gave access to data corresponding to the billing cycle immediately after the intervention, when no more messages were sent. Using this data, I find no statistically significant difference between treatment and Control Groups, in the fraction of people incurring late payment fees (See Table A7).

To complete the analysis of the treatment effect on late payment fees, I look for heterogeneous treatment effects depending on repayment behavior at baseline. I find suggestive evidence that the effect is larger for individuals that are current in their payments, compared to those that are carrying arrears at the time when they receive the message. I first look at the differential effect of the treatment on individuals that experienced a late payment fee at baseline.[12]

In Table 10 column 1, I look at heterogeneous effects using three indicator variables, the first one (current 1) taking the value of 1 when the individual is current in his payments at the time when the intervention started; the second variable (current 2) takes the value of 1 when the individual paid a late payment fee in the period previous to the intervention and did not make a payment by the time when the messages were sent (these users, by construction have been late for less than 30 days). The omitted category indicates when an individual had been late in his payments by more than 30 days, at the time when the messages were sent.

The resulting specification is the following:

$$\begin{aligned} Outcome_{it} =& \alpha + \alpha_i + \sum_{j=1}^{3} \phi_j During_t * T_i * Current_{jit} \\ & + \beta_1 * During_t + \beta_2 * T_i + \sum_{j=1}^{3} \beta_{3j} * Current_{jit} \\ & + \gamma_1 * During_t * T_i + \sum_{j=1}^{3} \gamma_{2j} * During_t * Current_{jit} \\ & + \sum_{j=1}^{3} \gamma_{3j} * T_i * Current_{it} + \epsilon_{it} \end{aligned} \tag{3}$$

Where $Current_1$ indicates that an individual is current in his payments, $Current_2$ indicates that the individual has been late for less than 30 days in his payment, and finally $Current_3$ (omitted category) indicates that an individual has been late for more than 30 days. The triple interactions are the coefficients of interest. I find

[12]Note that while all users in the experimental pool have paid a late payment fee in the past, only 37.7% of them paid a late payment fee in the two periods previous to the intervention, which are here defined as "the baseline period". The probability of paying a late fee conditional on paying one at baseline is 44.4%.

that the effect is the largest for individuals that are current, however there are no statistically significant differences.

Column 2 follows the same classification of users, but instead of three indicator variables, I use a single discrete variable that takes values 0 to 2 with 0 indicating that a user is current in his payments. The resulting econometric specification is analogous to the one described in equation 2 where the treatment indicator is interacted with the variable indicating the treatment period, and another variable indicating the heterogeneity of interest, which in this case corresponds to the repayment status which is labeled below as $CurrentSummary$:

$$\begin{aligned} Outcome_{it} = &\alpha + \alpha_i + \beta_1 * During_t + \beta_2 * T_i + \beta_3 * CurrentSummary_{it} \\ &+ \gamma_1 * During_t * T_i + \gamma_2 * During_t * CurrentSummary_{it} \\ &+ \gamma_3 * T_i * CurrentSummary_{it} \\ &+ \phi During_t * T_i * CurrentSummary_{it} + \epsilon_{it} \end{aligned} \tag{4}$$

As before, the coefficient of interest is the triple interaction ϕ.

Table 10 column 2 shows that going from being current to late by less than 30 days reduces the treatment effect by 2.3 percentage points, and it is statistically significant. Finally, when I compare in column 3 individuals that are current, vs all non-current individuals (more and less than 30 days late), I find that current individuals have a treatment effect that is larger in magnitude by 2.69 percentage points compared to individuals that are not current in their payments that have a reduction of 0.0597pp. As a benchmark notice that the average effect of the pooled treatments on credit card late fees is a reduction of 2.6pp (See Table 3). These results seem consistent with standard collection practices in the credit industry: typically when an individual is late in his payments, banks initiate collection processes by reaching out directly to consumers asking for repayment and eventually reporting to the corresponding credit bureaus. For users carrying arrears that are already undergoing a collection process, we would expect that the marginal effect of receiving a reminder by a company other than the bank, would be lower.

Finally, to complete the analysis of the dynamics of consumer credit card repayments, I look at the treatment effect of the intervention on long delinquencies, defined as those that last for more than 30 days; and short delinquencies, defined as those that are paid within 30 days of the corresponding payment deadline. In Table 11, I find that the treatment has a significant effect on short delinquencies, with a reduction of 1.85pp (notice that the reduction in overall delinquencies is 2.6pp, and this is therefore a significant fraction of the effect), and a non-significant reduction in the fraction of users incurring "long delinquencies". This supports the idea that the treatment effect is operating through an attention channel, as it is likely that the driver of people paying late for more than 30 days is financial hardship and not necessarily limited attention: the treatment is affecting the behavior of individuals that in the absence of the treatment would have in fact paid, but would have done it late, after being charged a late payment fee.

5.2 Other types of fees

The richness of the data allows me to extend the analysis to outcomes related to contingent fees in checking accounts. Specifically, in this section I look at the treatment effect on overdraft fees. During the two billing cycles previous to the intervention, the average fraction of users incurring overdraft each month was 22.5%, and the average cost was $R89 per month (conditional on paying). To estimate the effect of the treatment on overdraft, I use the specification described in equation 1, with two outcomes. The first outcome is defined as an indicator variable taking the value of one when a user incurs overdraft in a given period (extensive margin). The second outcome is the cost of overdraft fees paid (intensive margin). I look at heterogeneities across two dimensions: overdraft use and login activity, both identified at baseline. The resulting econometric specification is analogous to the one used in equation 2 where the treatment indicator is interacted with the variable indicating the treatment period, and another variable indicating the heterogeneity of interest.

In columns 1 to 3 of Table 12 I find that the intervention does not lead on average to a detectable increase in the fraction of people using overdraft in the sample. However, such aggregate results hide important heterogeneities, since among users that have incurred overdraft at baseline there is a of 2.89pp increase in the fraction of people paying overdraft fees, from a basis of 76%.[13] Similarly, increased login activity at baseline leads to a (statistically non-significant) reduction in the magnitude of this effect: going from the first to third quartile of the login distribution (1 to 5 logins) changes the treatment effect from an increase of 0.952 pp to a smaller increase of 0.8588pp in the fraction of people paying overdraft fees. In columns 4 to 6 I find that on average, the total cost of overdraft increased \$R2.73 from a basis of \$R25.98. This effect is driven by users who paid overdraft at baseline, who represent 29% of the experimental pool. For this subset of users, there is an increase of \$R6.62 from a basis of \$R98, whereas for the rest of users that have not paid overdraft fees in the past, the treatment has no effect on overdraft use.

It is important to notice the strong persistence of overdraft use over time: after paying overdraft in the pre-treatment period, the probability of using overdraft goes up to 76%. It seems reasonable that the strong persistence of overdraft use together with an attention shock towards an alternative task can be the slight “nudge” that in this case pushes consumers towards the unintended outcome of overdraft use.

Appendix 3 provides some more details on the mechanics of the treatment effect on overdraft fees.

5.3 Total cost of contingent fees

I now look at the overall effect in the amount of total contingent fees paid by users in the Treatment Group compared to the Control Group. Specifically I look at the sum of the total cost of credit card late fees and the total cost of overdraft fees paid by consumers. I focus on the average treatment effect for the overall experimental pool, but also look across two relevant heterogeneities, individuals that have incurred overdraft in the past, as well as individuals with different baseline login frequency.

[13] The set of users that paid overdraft at baseline represents 29% of the user base. A 2.89 percent increase in the probability of incurring overdraft among the member of this subgroup, represents a 0.8 percentage point increase in the fraction of the overall treatment group paying overdraft, which in turn represents 31% of the 2.6% of users that avoid credit card late payment fees as a result of the intervention. Furthermore, under the standard assumption of decreasing marginal utility of money, this is the most vulnerable group, as they are paying the most in fees, and running out of money in their checking account routinely.

In Table 13 I find that there is a 2.5% decrease in the total cost of fees paid. However, this effect is not statistically significant.[14] However, relevant heterogeneities are again hidden in the averages: Among those that do not paid overdraft fees at baseline, there are savings that account for \$R4.62 from a basis of \$R30.31, while for those that did pay overdraft at baseline, there are losses of \$R6.59 from a basis of \$R129.98.[15] Similarly, login at baseline has an increasing effect on the benefits of the intervention, but these effects are not statistically significant: going from the first to third quartile of the login distribution (1 to 5 logins) changes the treatment effect from a decrease of 62 cents in the total amount of fees paid, to a larger decrease of 80 cents.

The appendix presents the analogous regression in logs, as well as details of the effects on the cost of overdraft and credit cards contingent fees separately.[16]

6 The psychology and economics behind the treatment effects

6.1 Salience distortions in consumer choice

In this section, I argue that reminders are affecting consumer behavior, not only through an informational channel that brings to memory a task potentially forgotten, but also by affecting how much attention consumers allocate to such a task. The distinction between information and attention is relevant because for

[14]The change in benefits gains from 10.8% in the credit card market alone, to 2.5% when considering both credit and checking outcomes, comes mechanically from two sources: first, the numerator decreases as part of the credit card savings is offset by increased overdraft. Second, the denominator increases, as now we are considering both credit card fees and overdraft fees. An alternative comparison would keep the denominator fixed in both calculations. If we were to do that, the benefits from the intervention would decrease from 10.8% to 4.6%.

[15]The effect on overdraft fees for the group of users that paid overdraft before the intervention is larger than the effect on the group that did not pay overdraft at baseline by \$R11.21, as can be seen in the triple interaction regression coefficient of Table 13.

[16]Overdraft in Brazil has a very non-linear pricing structure in which withdrawals above existing balances an up to a pre-specified limit on an individual checking account are offered at competitive interest rate in the same range or lower than credit card monthly interest rates. However, intensive use of overdraft, above the pre-established limit, leads to hefty increases in prices, that are unambiguously higher than interest rate credit cards, for the major companies. The first level of overdraft is known in Brazil as "cheque especial", the second level is know as "adiantamento a depositante". See appendix Appendix 3 for a description of overdraft prices for the largest banks in Brazil.

a Forgetful But Otherwise Rational agent, reminders can do no harm. However, for a Forgetful And Salient Thinker agent, reminders can lead to distortions and suboptimal choices.

Consider a setting in which individuals choose from a choice set whose elements are made of all the alternatives that a consumer can remember. For a given choice set, consumers assign different attention weights to the elements in the set, and choose the element that maximizes their decision utility. Under this setting, the effect of reminders is potentially two-fold: On the one hand, a reminder can bring to memory an element that was previously being excluded from the choice set. But on the other hand a reminder can also (potentially simultaneously) affect the weights that each element receives, by making one element more salient than the others.

For a given choice set, a salience shock may disproportionally affect the weight that a certain element receives, leading to an overreaction to reminders. In my setting, two elements that should be present in a consumer choice set are overdraft fees (which are widespread in Brazil) and late payment fees. Sometimes overdraft fees are more expensive than late payment fees. Consider an individual for whom overdraft fees are in the choice set, but due to lack of memory late payment fees are not present. If this user receives a treatment message about late payment fees, it could be that the message brings late payment fees into the choice set without distorting the rational weights that overdraft and LPFs should have, or alternatively, it could be that the message not only brings LPFs in to the choice set, but also makes LPFs more salient.

To test for an attention distortion in the weights that consumers allocate to different elements of a choice set, I assume that if a consumer pays an overdraft at baseline, then overdraft fees are in his evoked choice set, or in other words he remembers overdraft fees.

Table 14 shows that the pattern of not paying a late payment fee and instead incurring an overdraft fee, occurs 11% more often in the Treatment Group, compared to the Control Group. This change in behavior leads to a 5% increase in the total cost of fees incurred by these users. I try to argue that this behavior is an overreaction to information caused by salience.

An alternative explanation is that consumers are rationally choosing to incur overdraft instead of paying late on their credit cards, because credit scores might be differentially affected by late payments and overdraft. Information from the major credit scoring agency in Brazil show that when a fee is missed, it takes on average 60 days for it to affect a credit score, and this occurs typically after several attempts to collect debts. In general, banks initiate the reporting process only after a credit card has been delinquent for 30 days or more. This practice is also common in other countries, including the US.

Therefore, the relevant question to assess the credit score consequences of the trade-off between credit card late fees and overdraft use, depends on whether or not the compliers of the treatment (in the statistical sense) that are substituting away from a credit card LPF and instead incurring overdraft, are affecting their credit scores. For there to be a credit score effect, we would need to have a decrease in the probability of spending more than 30 days without a payment. However, as can be seen on Table 11, it seems that the treatment is operating on the margin of short-term delinquents, that may incur a payment but nevertheless pay within 30 days of their payment deadline. Putting together the institutional background and the results on long term delinquency, it seems that the treatment would leave credit scores unaffected. In particular, it seems that the substitution between credit card late fees and more expensive overdraft is not leading to additional benefits in terms of credit scores.

The heterogeneity in benefits from the intervention leads to a natural taxonomy of consumers taken from Bordalo, Gennaioli and Shleifer (2015), where some consumers are Forgetful But Otherwise Rational (FBOR) and will not overreact to attention shocks; but some other consumers are Forgetful And Salient Thinkers (FAST) and the weights they assign to elements in their choice set will be distorted by salience, potentially leading to overreaction. I think the results of this section provide supportive evidence of the existence of FAST thinkers.[17]

[17] In an unreleased 2016 working paper with the same title, Bordalo, Gennaioli and Shleifer develop an alternative model of memory, founded on two psychological properties: similarity and interference (Kahana, 2012). Under this formulation, the repetition of a reminder increases the probability of remembering credit card tasks, but crowds out the probability of remembering checking account tasks (interference property), effectively distorting the weights assigned to the two tasks at hand in consumers'

6.2 Alternative explanations

In this subsection, I explore two explanations under which the effect is not driven by changes in the allocation of attention and instead is driven by information or a signaling/contextual inference channel.

First, I discuss as a potential alternative explanation, a purely informational channel under which, when receiving the message, individuals are informed about the contractual terms of their credit card: if they do not pay on time, they will be charged a late payment fee that averages $R40. Consider a fully rational individual that somehow does not know (or underestimates) the price of a late payment. Suppose that as a result of a cost-benefit analysis, this individual chooses not to pay his credit card this month. For this individual, receiving and seeing a message containing information about the real LPF could lead to a revised decision, and decrease the fraction of late payers through a purely informational channel.

To investigate the hypothesis that all the effects are coming from such a purely informational channel, I look at the first period of the intervention, and compare the effect of the Full Treatment, to the effect of the Robustness Treatment, considering only individuals that logged in (and therefore saw) the first message sent. If the Full-Treatment has a larger effect than the robustness treatment, I consider it as evidence that information is not the only channel through which attention shocks are having an effect. The results are in Table 15. I find that the Full-Treatment Group has a significantly larger effect, which means that receiving the same message more than once has an incremental effect, which is not consistent with the "information-only channel."

I now investigate the possibility that the treatment effect is taking place through a signaling or contextual inference channel, under which users infer that the PFM is giving them Financial Advice. Under this alternative hypothesis, users infer that a specialized party that has rich information about their finances is telling them to

decision utility. In the specific setting of this paper, the model of memory in the unrealeased working paper, leads to the same prediction as that of the 2015 formulation for FAST thinkers, when two conditions are met: recall is a binary function, and salience is defined as availability or visibility of a product. For this analysis, I will use the simpler memory formulation of 2015, and interpret the distortion in consumer choice as being driven by attention instead of memory. The policy implications, and the rest of the analysis remains unchanged.

pay their credit card on time, because that is the task that requires their attention the most. First, it is worth noticing that the Partner does not provide financial advice to their users, nor does it advertises financial advise services. The main service provided by the PFM is to allow tracking and categorization of transactions. Furthermore the message is framed as a reminder that a payment is coming up, and not as a suggestion, advice or demand.

The implications of such an inference, would be that consumers act as if credit card payments were more important than say, overdraft fees. From a reduce form perspective, this is equivalent to saying that individuals pay more attention to credit cards, over checking accounts, and in that sense it does not reject the explanation that reminders make individuals act as if credit cards were more important than checking accounts. In that sense, we may say that the advice sent by the informed party is in fact making credit card related tasks more salient, than checking account tasks.

Furthermore, if one wants to consider the "advice" channel, as different from the "salience" channel, it is important to notice that under a fully rational model, receiving the same piece of advice once, or more than once should have no differential effect. I however find in my data that the treatment that receives more messages leads to consistently larger effects than the treatment that has only two messages, as can be seen through section 5.1. Furthermore, the results in Table 15 support the same hypothesis.

7 Policy implications

7.1 Targeted nudging

The results of this study show that there are heterogeneities in benefits gained from the intervention. Furthermore these heterogeneities are predictable. I propose a policy intervention with the potential to more than double the benefits found in this study, by sending reminders only to those that are more likely to benefit from them. Under a "targeted nudging" approach, individuals that given their observable characteristics are likely to benefit from the intervention would be enrolled into an opt-out reminder based program to reduce credit card late payments. The rest of the users would be enrolled into an opt-in program in which they would be free to self select to receive messages, but would not receive, by default, any message. The procedure to change from one program to another should have as low transaction costs as possible.

In this specific context, individuals that typically carry sufficiently large balances in their checking accounts, or individuals that did not incur overdraft at baseline would be enrolled in an opt-out reminder based program. All other users would be enrolled in an opt-in version of the program. If inertia is large enough and users do not switch from one program to another, back of the envelope calculations suggest that the benefits from the intervention would increase to account for a reduction of more than 5.6% in the total amount of contingent fees paid by consumers,[18] which corresponds to a 224% increase in net benefits, from the current level of 2.5%.[19]

More generally, the "targeted nudging" approach suggests that individuals with different expected responses from the treatment should get different versions of the treatment. One version of this policy could have different users receiving differentiated default treatments, another version could have some users not being treated at all, and yet another version could have some users sent different messages that, for example, shock their attention towards overdraft fees instead of credit card late payment fees. This targeted approach is aligned with the discussion of Beshears et al. (2009) that, in the context of retirement savings, proposes different default portfolios for individuals with different demographic characteristics. It is also closely related to the policy recommendation of Allcott and Kessler (2015) that in the context of energy saving policies, consider the potential dis-utility from behavioral changes induced by providing one-page letters comparing a household's energy use to that of its neighbors, and propose sending these informational nudges

[18]This 5.6% is given by the 15% reduction in fees experienced by 71% of users carrying 37.5% of the total value of fees paid, that would be enrolled in the program.

[19]An important caveat, is that these calculations do not take in to account that the existence of "annoyance" cost from receiving too many messages documented byDamgaard and Gravert (2016). To try to learn about consumers subjective evaluation of the intervention, Appendix 5 provides the results of a satisfaction survey sent treatment participants 10 days after they received the last treatment message.

only to users for which energy savings out-weight the dis-utility cost (measured by individuals' willingness to pay) of receiving the nudge.

The results of this paper show that while nudges have a huge potential to increase benefit at low costs, it is not always clear what the best way to nudge is, and the choice of a specific policy should be based on a careful evaluation, with a sufficiently rich set of outcomes.

8 Conclusion

The welfare evaluation of a nudge-based policy intervention will not be complete if outcomes that could be indirectly affected are not taken in to account: could more savings lead to increased debt? (Ashraf, Karlan and Yin, 2006); could buying more fertilizer, lead to lower pesticide use? (Duflo, Kremer and Robinson, 2011); or could grammar improvement crowd-out studying time from math or arts? (Mayer et al., 2015).

In this paper, I present new evidence of the effect of informational nudges on consumer behavior. As opposed to previous studies, I find that nudges have not only direct effects on outcomes of interest, but also indirect effects on unintended outcomes. Specifically, I look at the effect of reminders for credit card timely payment using credit card and checking account transactional data. The Brazilian personal finance market provides an appropriate setting to learn about the effect of informational nudges on a rich set of outcomes since a large fraction of users incur overdraft every month, and are therefore close to a relevant trade-off margin. I find that the effects of the intervention are different for consumers with different observable characteristics and while in general, credit card late payments decrease, checking account overdraft use increases for users that were close to the overdraft margin.

Informed by the results of this paper, I propose a second generation of nudge-based policies, that takes in to account that consumers response to nudges is heterogeneous, and makes use of the growing availability of rich data: targeting nudges to users whose observable characteristics predict strictly positive welfare effects.

Future research could use machine learning techniques to further characterize subpopulations with differential benefit gains from nudging (Athey and Imbens, 2016), and target different versions of the policy to users with different expected responses.

References

Agarwal, Sumit, John C Driscoll, Xavier Gabaix, and David Laibson. 2013. "Learning in the Credit Card Market."

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. 2015. "Regulating Consumer Financial Products: Evidence from Credit Cards." *The Quarterly Journal of Economics*, 130(1): 111–164.

Alan, Sule, Mehmet Cemalcılar, Dean Karlan, and Jonathan Zinman. 2015. "Unshrouding Effects on Demand for a Costly Add-on: Evidence from Bank Overdrafts in Turkey." National Bureau of Economic Research, Working Paper 20956.

Allcott, Hunt, and Dmitry Taubinsky. 2015. "Evaluating behaviorally motivated policy: experimental evidence from the lightbulb market." *The American Economic Review*, 105(8): 2501–2538.

Allcott, Hunt, and Judd B. Kessler. 2015. "The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons." National Bureau of Economic Research, Working Paper 21671.

Allcott, Hunt, and Todd Rogers. 2014. "The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation." *The American Economic Review*, 104(10): 3003–3037.

Ashraf, Nava, Dean Karlan, and Wesley Yin. 2006. "Tying Odysseus to the Mast: Evidence From a Commitment Savings Product in the Philippines." *The Quarterly Journal of Economics*, 121(2): 635–672.

Athey, Susan, and Guido Imbens. 2016. "Recursive partitioning for heterogeneous causal effects." *Proceedings of the National Academy of Sciences*, 113(27): 7353–7360.

Baker, Scott. 2015. "Debt and the consumption response to household income shocks."

Banerjee, Abhijit V, and Sendhil Mullainathan. 2008. "Limited attention and income distribution." *The American Economic Review*, 98(2): 489–493.

Beshears, John, James J Choi, David Laibson, and Brigitte C Madrian. 2009. "The importance of default options for retirement saving outcomes: Evidence from the United States." In *Social security policy in a changing environment.* 167–195. University of Chicago Press.

Beshears, John, James J Choi, David Laibson, Brigitte C Madrian, and Katherine L Milkman. 2015. "The effect of providing peer information on retirement savings decisions." *The Journal of finance*, 70(3): 1161–1201.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2013. "Salience and Consumer Choice." *Journal of Political Economy*, 121(5): 803–843.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2015. "Memory, attention and choice."

Bursztyn, Leonardo, Stefano Fiorin, Daniel Gottlieb, and Martin Kanz. 2015. "Moral Incentives: Experimental Evidence from Repayments of an Islamic Credit Card." National Bureau of Economic Research, Working Paper 21611.

Cadena, Ximena, and Antoinette Schoar. 2011. "Remembering to Pay? Reminders vs. Financial Incentives for Loan Payments." National Bureau of Economic Research, Working Paper 17020.

Calzolari, Giacomo, and Mattia Nardotto. 2011. "Nudging with information: A randomized field experiment."

Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and taxation: Theory and evidence." *The American Economic Review*, 99(4): 1145–1177.

Chetty, Raj, John N Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen. 2014. "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark." *The Quarterly Journal of Economics*, 129(3): 1141–1219.

Damgaard, Mette Trier, and Christina Gravert. 2016. "The hidden costs of nudging: Experimental evidence from reminders in fundraising."

Dranove, David, and Ginger Zhe Jin. 2010. "Quality disclosure and certification: Theory and practice." *Journal of Economic Literature*, 48(4): 935–963.

Duflo, Esther, Michael Kremer, and Jonathan Robinson. 2011. "Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya." *The American Economic Review*, 101(6): 2350–2390.

Fehr, Ernst, and Antonio Rangel. 2011. "Neuroeconomic foundations of economic choice: recent advances." *The Journal of Economic Perspectives*, 25(4): 3–30.

Ferman, Bruno. Forthcoming. "Reading the Fine Print: Information Disclosure in the Brazilian Credit Card Market." *Management Science*.

Fishman, Michael J, and Kathleen M Hagerty. 2003. "Mandatory versus voluntary disclosure in markets with informed and uninformed customers." *Journal of Law, Economics, and Organization*, 19(1): 45–63.

Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *The Quarterly Journal of Economics*, 121(2): 505–540.

Gelman, Michael, Shachar Kariv, Matthew D Shapiro, Dan Silverman, and Steven Tadelis. 2014. "Harnessing naturally occurring data to measure the response of spending to income." *Science*, 345(6193): 212–215.

Grubb, Michael D. 2015. "Consumer Inattention and Bill-Shock Regulation." *The Review of Economic Studies*, 82(1): 219–257.

Hare, Todd A, Jonathan Malmaud, and Antonio Rangel. 2011. "Focusing attention on the health aspects of foods changes value signals in vmPFC and improves dietary choice." *The Journal of Neuroscience*, 31(30): 11077–11087.

Heidhues, Paul, and Botond Kőszegi. 2010. "Exploiting naivete about self-control in the credit market." *The American Economic Review*, 2279–2303.

Heidhues, Paul, Botond Koszegi, and Takeshi Murooka. Forthcoming. "Inferior products and profitable deception." *Review of Economic Studies*.

Hirshleifer, David, and Siew Hong Teoh. 2003. "Limited attention, information disclosure, and financial reporting." *Journal of Accounting and Economics*, 36(1): 337–386.

Kahana, Michael Jacob. 2012. *Foundations of human memory.* Oxford University Press.

Kahneman, D. 1973. *Attention and effort. Prentice-Hall series in experimental psychology*, Prentice-Hall.

Kahneman, Daniel, and Richard H Thaler. 2006. "Anomalies: Utility maximization and experienced utility." *The Journal of Economic Perspectives*, 20(1): 221–234.

Karlan, Dean, Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman. Forthcoming. "Getting to the top of mind: How reminders increase saving." *Management Science*.

Karlan, Dean, Melanie Morten, and Jonathan Zinman. 2015. "A personal touch: Text messaging for loan repayment." *Behavioral Science and Policy Association*, 1(2): 25–31.

Kuchler, Theresa. 2015. "Sticking to Your Plan: Empirical Evidence on the Role of Present Bias for Credit Card Paydown."

Maćkowiak, Bartosz, and Mirko Wiederholt. 2015. "Business Cycle Dynamics under Rational Inattention." *The Review of Economic Studies*, 82(4): 1502–1532.

Mayer, Susan E., Ariel Kalil, Philip Oreopoulos, and Sebastian Gallegos. 2015. "Using Behavioral Insights to Increase Parental Engagement: The Parents and Children Together (PACT) Intervention." National Bureau of Economic Research, Working Paper 21602.

Melzer, Brian T, and Donald P Morgan. 2015. "Competition in a consumer loan market: Payday loans and overdraft credit." *Journal of Financial Intermediation*, 24(1): 25–44.

Pagel, Michaela, and Arna Vardardottir. 2016. "The Liquid Hand-to-Mouth."

Simon, Herbert A. 1955. "A behavioral model of rational choice." *The quarterly journal of economics*, 99–118.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of Monetary Economics*, 50(3): 665–690.

Stango, Victor, and Jonathan Zinman. 2014. "Limited and Varying Consumer Attention: Evidence from Shocks to the Salience of Bank Overdraft Fees." *Review of Financial Studies*, 27(4): 990–1030.

Thaler, R.H. 2015. *Misbehaving: The Making of Behavioral Economics. Business - economics*, W.W. Norton.

Thaler, Richard H, and Cass R Sunstein. 2003. "Libertarian paternalism." *The American Economic Review*, 93(2): 175–179.

Vervloet, Marcia, Annemiek J Linn, Julia CM van Weert, Dinny H De Bakker, Marcel L Bouvy, and Liset Van Dijk. 2012. "The effectiveness of interventions using electronic reminders to improve adherence to chronic medication: a systematic review of the literature." *Journal of the American Medical Informatics Association*, 19(5): 696–704.

Tables

Table 1: Covariate Balance at Baseline: Randomization test

	Control Group	Treatment group 1	Treatment group 2	Robustness treatment group	Bonferroni p-values
Number of users originally allocated	10,090	5,987	4,038	5,954	
Number of users with available information	5,181	3,124	2,130	3,103	
Fraction of missing users (attrition)	0.4865	0.4782	0.4725	0.4788	0.4425
Fraction of users with LPF	0.2301	0.2321	0.2279	0.2261	0.8935
Fraction of users with short delinquencies	0.1595	0.1661	0.1593	0.1573	0.5804
Fraction of users charged a LPF, who had enough balance in their checking account to cover their CC balance	0.3693	0.3675	0.3689	0.3776	0.8595
Fraction of users with OD	0.2258	0.2214	0.2211	0.2312	0.6787
Number of logins per month	5.05	4.95	5.39	5.18	0.4264
Number of credit cards	1.46	1.45	1.45	1.46	0.9244
Income ($R)	9,082.42	8,802.95	8,393.89	8,796.09	0.3622
Essential expenses ($R)	1,848.30	1,779.14	1,887.02	1,793.22	0.9966
Discretionary expenses ($R)	1,557.58	1,504.03	1,363.97	1,335.46	0.6260
Short run consumables ($R)	150.02	149.02	124.56	145.86	0.2244

This Table presents a test for covariate balance across treatments and control group. For all users with information available at the end of the intervention, I calculate the mean value at baseline, of a set of variables. The unit of observation is at the user-billing cycle level. Baseline is defined as the period compressing two billing cycles previous to the intervention. LPF stands for late payment fee. Short delinquencies indicate the presence of users that were charged a late payment fee but paid back on their credit card within 30 days of the missed deadline. OD stands for overdraft fee. Essential expenses include home expenses (rent and/or mortgage payments), residential bills, health, education, TV/internet/phone, transportation, work expenses, groceries and services. Discretionary expenses include house keeping, gifts/donations, leisure, personal care, bars/restaurants and shopping. Short Run Consumables include coffee shops, bars, restaurants and fast food.

Table 2: Descriptive Statistics: Baseline, all groups.

	Mean	St. Dev.	Median	p25	p75
Income ($R)	8,843.97	18,178.45	3,473.34	816.13	9,059.83
Checking account balance/Credit card balance	1.52	5.06	0.24	-0.33	2.09
Number of logins per month	5.11	10.21	2	0	5
Number of credit cards	1.46	0.75	1	1	2
Credit card monthly payments ($R)	934.76	1,672.35	227.61	0	1,138.66
Credit card balance ($R)	1,605.67	2,055.02	856.03	285.75	2,092.68
Fraction of users paying LPF	0.2298	0.4207	0	0	0
Fraction of users with short delinquencies	0.1646	0.3708	0	0	0
Fraction of users charged a LPF, who had enough balance in their checking account to cover their CC balance	0.3436	0.4749	0	0	1
Fraction of users paying overdraft fees	0.2185	0.4132	0	0	0

This Table contains summary statistics of some of the main variables used in the analysis. The unit of observation is at the user-billing cycle level. The calculations correspond to values at baseline, where baseline is defined as the period compressing two billing cycles before the intervention. Users from all treatment and control groups are pooled. LPF stands for late payment fee. Short delinquencies indicate the presence of users that were charged a late payment fee but paid back on their credit card within 30 days of the missed deadline. CC stands for credit card.

Table 3: Probability of paying a CC Late Fee

	Dependent variable: Paid CC Late Fee 0,1					
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
Pooled treatments	-0.0256***					
	(0.00691)					
Full-Treatment		-0.0322***				
		(0.00792)				
DD Alert			-0.0160*			
			(0.00909)			
During*Pooled treatments				-0.0260***		
				(0.00748)		
During*Full-Treatment					-0.0342***	
					(0.00846)	
During*DD Alert						-0.0139
						(0.00996)
During				0.0611***	0.0611***	0.0611***
				(0.00545)	(0.00545)	(0.00545)
Constant	0.291***	0.291***	0.291***			
	(0.00499)	(0.00499)	(0.00499)			
Observations	20,870	16,610	14,622	41,740	33,220	29,244
Number of clusters				10,435	8,305	7,311
Periods included	During 1&2	During 1&2	During 1&2	Baseline+ During 1&2	Baseline+ During 1&2	Baseline+ During 1&2
Individual fixed effects	No	No	No	Yes	Yes	Yes

This Table presents the results from OLS regressions estimating equation 1 with different treatments. The unit of observation is at the user-billing cycle level. *Full − treatment* takes the value of one if a given user belongs to the Full-treatment group, that received five messages. *DD Alert* takes the value of one if a given users belong to the DD Alert group, that received two messages. *Pooled treatments* takes the value of one if a given users belongs to either one of the two treatments. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. Standard errors clustered at user level in parenthesis.The difference in coefficients for column 2 and 3 is 0.0162 with a p-value of 0.099. The difference in coefficients between column 5 and 6 is 0.0203, with a p-value of 0.055. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 4: Treatment effect on spending, for selected categories

	Log of expenses by category +1			
	(1)	(2)	(3)	(4)
	Essential expenses	Discretionary expenses	Short run consumables	ATM withdrawals
During*Full-Treatment	-0.00453	0.00540	-0.0161	-0.00535
	(0.0384)	(0.0398)	(0.0382)	(0.05758)
During	0.785***	0.515***	0.718***	1.1653***
	(0.0239)	(0.0246)	(0.0237)	(0.03510)
Mean of dependent variable during intervention period in control group ($R)	1944.13	1409.19	165.80	1311.94
Observations	31,692	30,900	32,864	32,864
Number of clusters	8,298	8,294	8,284	8,284

This Table presents the results from OLS regressions estimating equation 1 with different outcomes. $Full - treatment$ takes the value of one if a given user belongs to the Full-treatment group, that received five messages. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. Observations are windsorized at the 1st and 99th percentile.
*** p<0.01, ** p<0.05, * p<0.1

Table 5: Treatment effect on balances at the end of the billing cycle

	(1) Log CC Balance	(2) Log Ch. Acct. Balance	(3) Rel. Liquidity
During*Full-Treatment	0.000865	-0.00216	-0.00331
	(0.0132)	(0.0124)	(0.00907)
During	0.0260***	-0.104***	0.0489***
	(0.00807)	(0.00764)	(0.00554)
Mean of dependent variable during intervention period in control group ($R)	1,572.34	2,163.23	1.39
Observations	33,220	32,887	33,220
Number of clusters	8,305	8,305	8,305

This Table presents the results from OLS regressions estimating equation 1 with different outcomes. $Full-treatment$ takes the value of one if a given user belongs to the Full-treatment group, that received five messages. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. Observations are windsorized at the 1st and 99th percentile. Negative balances are mapped to the positive domain by adding a constant equal to its minimum observed value. CC stands for credit card. *** p<0.01, ** p<0.05, * p<0.1

Table 6: Liquidity and treatment effects

	Dep. Var: Paid CC Late Fee {0,1}		
	(1)	(2)	(3)
During*Pooled treatments	-0.0125		
	(0.0105)		
During*EB*Pooled treatments	-0.0363**		
	(0.0163)		
During*Full-Treatment		-0.0173	
		(0.0120)	
During*EB*Full-Treatment		-0.0451**	
		(0.0186)	
During*DD Alert			-0.00560
			(0.0138)
During*EB*DD Alert			-0.0231
			(0.0215)
During	0.0760***	0.0760***	0.0760***
	(0.00755)	(0.00755)	(0.00755)
During*EB	-0.0318***	-0.0318***	-0.0318***
	(0.0119)	(0.0119)	(0.0119)
EB*DD Alert			0.00831
			(0.0190)
EB*Full-Treatment		0.00790	
		(0.0166)	
EB*Pooled treatments	0.00820		
	(0.0145)		
Enough balance (EB)	-0.0503***	-0.0503***	-0.0503***
	(0.0105)	(0.0105)	(0.0105)
Mean of dependent variable during intervention period for control group with EB=1	0.219	0.219	0.219
Observations	41,740	33,220	29,244
Number of clusters	10,435	8,305	7,311

This Table presents the results from OLS regressions estimating equation 2 with different treatments. $Full-treatment$ takes the value of one if a given user belongs to the Full-treatment group, that received five messages. $DD\ Alert$ takes the value of one if a given user belongs to the DD Alert treatment group, that received two messages. $Pooled\ Treatments$ takes the value of one when a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. $During$ takes the value of one for billing cycles during the treatment period. EB takes the value of one when a given individual has enough balance in her checking account to cover the full outstanding balance in her credit card. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. *** p<0.01, ** p<0.05, * p<0.1

Table 7: Fraction of users with enough checking balance when LPF was charged

	Dep. Var: Enough Balance* Late payment {0,1}		
	(1)	(2)	(3)
During*Pooled treatments	-0.0217***		
	(0.00471)		
During*Full-Treatment		-0.0281***	
		(0.00530)	
During*DD Alert			-0.0124**
			(0.00625)
During	0.0307***	0.0307***	0.0307***
	(0.00348)	(0.00348)	(0.00348)
Mean of dependent variable during intervention period in control group	0.0781	0.0781	0.0781
Observations	41,740	33,220	29,244
Number of clusters	10,435	8,305	7,311
Omitted Category	Control Group	Control Group	Control Group

This Table presents the results from OLS regressions estimating equation 1 for various treatments. $Full - treatment$ takes the value of one if a given user belongs to the Full-treatment group, that received five messages. *DD Alert* takes the value of one if a given user belongs to the DD Alert group, that received two messages. *Pooled Treatments* takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. *Enough Balance* takes the value of one when a given individual has enough balance in her checking account to cover the full outstanding balance in her credit card. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. *** p<0.01, ** p<0.05, * p<0.1

Table 8: Treatment effect on login activity

Dependent variable: Number of logins per period	(1)	(2)	(3)
During*Pooled treatments	0.930***		
	(0.124)		
During*Full-Treatment		1.098***	
		(0.141)	
During*DD Alert			0.684***
			(0.167)
During	-0.374***	-0.374***	-0.374***
	(0.0878)	(0.0878)	(0.0879)
Mean of dependent variable during intervention period in control group	4.67	4.67	4.67
Observations	41,740	33,220	29,244
Number of clusters	10,435	8,305	7,311

This Table presents the results from OLS regressions estimating equation 1 with different treatments. $Full-treatment$ takes the value of one if a given user belongs to the Full-treatment group, that received five messages. *DD Alert* takes the value of one if a given user belongs to the DD Alert group, that received two messages. $Pooled\ Treatments$ takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. *** p<0.01, ** p<0.05, * p<0.1

Table 9: Treatment effects and login into the app

	(1) Number of logins per month	(2) Paid CC Late Fee {0,1}	(3) Paid CC Late Fee {0,1}
During*Full-Treatment	-0.0259	0.0214*	
	(0.200)	(0.0116)	
During*Robustness Treatment			-0.00821
			(0.0106)
During	0.834***	0.0659***	0.0955***
	(0.140)	(0.00799)	(0.00656)
Omitted category	Robustness Treatment-group at baseline		Control Group at baseline
Observations	18,681	18,681	24,852
Number of clusters	6,227	6,227	8,284

This Table presents the results from OLS regressions estimating equation 2 with different outcomes. $Full - treatment$ takes the value of one if a given user belongs to the Full-treatment group, that received five messages. *Robustness* takes the value of one if a given user belongs to the Robustness treatment group, that received five placebo messages. The information corresponds to two billing cycles before the intervention, and the second billing cycle of the intervention. *During* takes the value of one for the second billing cycle of the treatment period. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. *** p<0.01, ** p<0.05, * p<0.1

Table 10: Heterogeneous treatment effects: Baseline repayment behavior

	Dependent variable: Paid CC Late Fee 0,1		
	(1)	(2)	(3)
During*Pooled treatments	0.00144	-0.0554***	-0.00597
	(0.0337)	(0.0181)	(0.0156)
During * Pooled treatments *Current 1	-0.0343		
	(0.0346)		
During * Pooled treatments* Current 2	-0.00813		
	(0.0380)		
During*Pooled treatments*Late (Linear)		0.0230*	
		(0.0136)	
During*Pooled treatments* Current (summary binary)			-0.0269
			(0.0174)
During*Current 1	0.380***		
	(0.0243)		
During* Current 2	0.0865***		
	(0.0267)		
During*Late (Linear)		-0.237***	
		(0.00969)	
During*Current (summary binary)			0.310***
			(0.0125)
During	-0.234***	0.375***	-0.165***
	(0.0236)	(0.0131)	(0.0111)
Observations	41,740	41,740	41,740
Number of clusters	10,435	10,435	10,435
Omitted category	More than 30 days delinquent	More than 30 days delinquent	All non-current users

This Table presents the results from OLS regressions estimating equations 3 and 4. *Pooled Treatments* takes the value of one when a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. For the specification in column 1, $Current_1$ indicates that an individual is current in his payments,$Current_2$ indicates that the individual has been late for less than 30 days in his payment, $Current_3$ (omitted category) indicates that an individual has been late for more than 30 days. For the specification in column two, *Late* (*linear*) is a discrete variable that takes values zero to two, with zero indicating that a user is current in his payments, one indicating that an individual has been late for less than 30 days, and two indicating that an individual has been late for more than 30 days. Finally, *Current* (*binary*) takes the value of one when an individual is current in his payments at the time when the messages were sent, and zero otherwise. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. *** p<0.01, ** p<0.05, * p<0.1

Table 11: Treatment effect on short and long delinquencies

	(1) Payment late for less than 30 days	(2) Payment late for more than 30 days
During*Pooled treatments	-0.0185***	-0.00744
	(0.00694)	(0.00492)
During	0.0394***	0.0217***
	(0.00500)	(0.00349)
Observations	41,740	41,740
Number of clusters	10,435	10,435

This Table presents the results from OLS regressions estimating equation 1 with different outcomes. *Pooled Treatments* takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions.
*** p<0.01, ** p<0.05, * p<0.1

Table 12: Treatment effect on overdraft use

	Dep. var.: Paid OD Fee {0,1}			Dep. var.: Value of OD Fee ($R)		
	(1)	(2)	(3)	(4)	(5)	(6)
During*Pooled treatments	0.00849	0.000379	0.00952	2.729***	0.696	2.928***
	(0.00541)	(0.00527)	(0.00605)	(0.723)	(0.446)	(0.786)
During*Pooled treatments*OD at baseline		0.0289**		7.314***		
		(0.0145)		(2.288)		
During*Pooled treatments*Logins per month			-0.000233			
			(0.000542)			
During*Logins per month			0.00221***			-0.0414
			(0.000390)			(0.0600)
During*OD at baseline		0.0114		5.809***		
		(0.0104)		(1.520)		
During	0.0603***	0.0570***	0.0491***	5.432***	3.743***	4.579***
	(0.00387)	(0.00376)	(0.00435)	(0.482)	(0.293)	(0.530)
Mean of dependent variable during intervention period in control group		0.286			25.96	
Observations	41,740	41,740	41,740	41,740	41,740	41,740
Number of clusters	10,435	10,435	10,435	10,435	10,435	10,435

This Table presents the results from OLS regressions estimating equation 2 with different outcomes. *Pooled Treatments* takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. *OD at baseline* takes the value of one if a given user incurred overdraft at baseline. *Logins per month* indicated the number of logins per user and billing cycle Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions.
*** p<0.01, ** p<0.05, * p<0.1

Table 13: Treatment effect on the total cost of contingent fees paid

	Dep. var.: Total cost of contingent fees ($R)		
	(1)	(2)	(3)
During*Pooled treatments	-1.518	-4.625***	-1.299
	(1.417)	(1.286)	(1.551)
During*Pooled treatments*OD at baseline		11.21***	
		(3.943)	
During*Pooled treatments*Logins per month			-0.0458
			(0.113)
During*OD at baseline		10.14***	
		(2.778)	
During*Logins per month			0.216***
			(0.0815)
During	13.18***	10.24***	12.09***
	(1.015)	(0.946)	(1.114)
Mean of dependent variable during intervention period in control group ($R)	59.29	59.29	59.29
Observations	41,740	41,740	41,740
Number of clusters	10,435	10,435	10,435

This Table presents the results from OLS regressions estimating equation 1 for various outcomes. *Pooled Treatments* takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. *OD at baseline* takes the value of one if a given user incurred overdraft at baseline. *Logins per month* indicated the number of logins per user and billing cycle. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. *** p<0.01, ** p<0.05, * p<0.1

Table 14: Treatment effect on individuals that had overdraft at baseline

	Overdraft and No CC LPF	Total value of contingent fees
	(1)	(2)
During*Pooled treatments	0.0501***	6.582*
	(0.0158)	(3.728)
During	0.00863	20.38***
	(0.0111)	(2.612)
Observations	11,984	11,984
Number of clusters	2,996	2,996

This Table presents the results from OLS regressions estimating equation 1 with different outcomes. *Pooled Treatments* takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions.
*** p<0.01, ** p<0.05, * p<0.1

Table 15: Information treatment

	Dependent variable: Paid CC Late Fee {0,1}	
	(1)	(2)
During*Robustness	0.0396***	0.0591***
	(0.0119)	(0.0162)
During	-0.0122	-0.0333***
	(0.00812)	(0.0109)
Observations	18,681	10,083
Number of aux_statement_id	6,227	3,361
Comments	All users	Only users that logged-in after first message

This Table presents the results from OLS regressions estimating equation 1 with different outcomes. *Pooled Treatments* takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions.
*** p<0.01, ** p<0.05, * p<0.1

Appendices

Appendix 1 Treatment effect on spending

For robustness, this section presents the (null) average treatment effect of the treatment on spending categories, for each of the spending categories automatically constructed by the Partner company. The specification comes from Section 5.1:

$$Outcome_{it} = \alpha + \alpha_i + \beta_1 * During_t + \beta_2 * T_i + \gamma * During_t * T_i + \epsilon_{it} \quad (1)$$

Where outcome represents the log of spending +1, in each of the categories, and T represents an indicator variable that takes the value of 1 when a given user is a member of the Full-Treatment Group.

The results are presented in Table A3. Standard errors are clustered at the user level. The omitted category is the Control Group at baseline.

Appendix 2 Liquidity

Coming from Section 5.1, below I present an alternative specification for the role of relative liquidity, with a dummy variable indicating if an individual has enough balance in his checking account to cover the minimum payment on his credit card. I find that the treatment has a significant effect reducing the fraction of people being charged a late payment fee while having enough balance in their checking account, and furthermore, the Full-Treatment Group has a statistically significantly larger effect (p-value 0.040), when compared to the DD Alert Group.

I now continue exploring the role of liquidity, constructing a liquidity ratio under which I normalize balances in checking accounts at the end of the CC billing cycle, by the corresponding credit card balance. This variable has a mean of 1.52, with a median of 0.24, first quartile of -0.33 and 3rd quartile of 2.09.

To measure heterogeneous treatment effects for individuals with different levels of liquidity (according to the measure defined above) I run the following specification:

$$\begin{aligned}LateFee_{it} =&\alpha + \alpha_i + \beta_1 * During_t + \beta_2 * Treatment_i + \beta_3 * Liquidity \\ &+ \beta_4 * During_t * Treatment + \beta_5 * During_t * Liquidity \\ &+ \beta_6 * Treatment_i * Liquidity \\ &+ \mu * During_t * Treatment_i * Liquidity + \epsilon_{it}\end{aligned} \tag{2}$$

The main coefficients are presented in Table A6. The coefficients of interest correspond to the triple interactions of the treated periods, the treated group, and the liquidity ratio. It can be seen that for the Full-Treatment Group going from the first to third quartile of the liquidity ratio, leads to an increase of 0.59 percentage points, which is 27% larger than the treatment effect on individuals in the first quartile of the LR distribution.

Next, I explore a non-linear measure of liquidity, in which I construct indicator variables for whether an individual has enough money in his checking account, to cover the balance in full, to cover twice, three times, four or five or more times the balance in his credit card. For exposition purposes I present selected coefficients of interest. It can be seen that the treatment starts having an effect only when individuals have enough balance in their checking account to cover their full balance, and not as soon as they have enough liquidity to cover the minimum payment. Measures of the treatment effect when individuals can cover their credit card balance 2,3 and 4 times are very noisy, but consistently suggesting that reminders reduce late payments when enough liquidity is available.

Finally, I look at the interaction between the paycheck cycle and the credit card billing cycle, by looking at heterogeneous treatment effects for users that have a paycheck arriving up to 15 days before their corresponding credit card due date, and for users with paycheck arriving more than 15 days before the credit card payment is due. To reduce differences between comparison groups, I restrict the analysis to users that receive monthly payments, since users that receive more periodic payments can be different across other characteristics as well. I find that for individuals with a monthly paycheck for which the gap between due date and paycheck arrival is small, the treatment effect is more than twice as large compared to individuals with a large gap between paycheck arrival and credit card due date. However this difference is not statistically significant. See Table A10.

Appendix 3 Overdraft fees

In this section of the appendix, I cover two main aspects. First, I describe the non-linearities of overdraft prices, and secondly, I further look in to the mechanics of overdraft charges by looking at the interaction of the timing of overdraft charges, short term delinquencies and the timing of paycheck arrival.

To get a sense of contractual prices, Table A9 presents contractual information on credit card interest rates and other types of fees, taken from the websites of the largest consumer banks in Brazil.

Where:

Overdraft. AKA "Cheque especial", is a line of credit assigned to individuals that allows them to withdraw or spend more that their actual balances in their checking accounts up to a certain limit. This can be though as source of credit alternative to credit card debt.

Extended overdraft AKA "Adiantamiento a depositante", is form of "overdraft on overdraft" that allows consumers to spend above their overdraft limit, but at a very high cost. This fee seems to better match the intuition for an overdraft fees in other countries such as the US.

As can be seen, the strong non-linearities in overdraft prices make users subject to drastic changes in the prices they face when going above a certain spending limit.

Next, I present more detail on the mechanics of overdraft, specifically when it comes to the relation between avoiding payments that in the absence of the treatment would have been made less than 30 days late, and overdraft use. From the results of Section 5.2 we can see that the treatment is affecting individuals that in the absence of the treatment may have paid within 30 days anyways. To understand how they can be paying overdraft fees as well, I look at the timing of overdraft, and the timing of paycheck arrival. I find that the treatment effect on overdraft is larger for individuals with weekly or bi-weekly paychecks, and that for bi-weekly users overdraft is actually taking place in the bi-weekly cycle that contains the corresponding credit card due date and payment. See Table A11. These results emphasize the intricacies of liquidity management over the paycheck cycle: it's not just about having enough liquidity to fulfill financial obligations on a given point of time, it is about having liquidity at the right time.

Appendix 4 Total cost of contingent fees

Tables A12 to A14 provide decompositions of the total cost of contingent fees for individuals that have paid overdraft in the past and for individuals that have not paid. To do that, I split the sample according to a dummy variable indicating if an individual has paid overdraft in the past, and run two separate regressions. The coefficients presented in this appendix can also be obtained from the tables presented in Table13 by making the appropriate arithmetic.

As an additional robustness test, Table A15 presents the analogous specification as in table 13, with a log-transformed dependent variable. The coefficients now have a percentage change interpretation. While there are some differences in the magnitude, the results are qualitatively the same.

Appendix 5 Survey results

Ten days after receiving the last message, individuals in all treatment groups were sent a mobile survey. Only 1,028 individuals responded to the survey. Their answers are presented in Table A16.

As can be seen, the large majority of users (77.5%) want to continue receiving reminders. However, when asked about the number of messages they want to receive, 57.5% report a preference for getting only one or two messages, instead of the most effective treatment among the ones tested here consisting of five messages.This suggests that while there is potential for a profitable business proposition, there seems to be a mismatch between ex-post treatment effectiveness and ex-ante consumer satisfaction.

Appendix tables

Table A1: Treatment messages

Type of message	Time of message	Message content
DD Alert	3-6 days before DD	Your credit card due date is here. Pay today and avoid late payment fees (R$ 40 on average)! If you already paid, ignore this message.
DD Alert	6-9 days before DD	Your credit card due date is approaching. Pay today and avoid late payment fees (R$ 40 on average)! Ignore if you have already paid.
Planning	13-16 days before DD	Did you set money aside for your credit card payment this month? do not spend it somewhere else, and avoid late payment fees! (R$40 on average)
Planning	20-23 days before DD	Set some money aside for your next credit card payment, and do not spend it somewhere else. Avoid late payment fees (R$40 on average)
Planning	24-27 days before DD	Start planning for your next credit card payment today. Set some money aside and avoid late payment fees (R$40 on average)

This Table presents the content of the five messages defined as part of the intervention. Messages labeled as DD alert, were sent close the the corresponding credit card due date to encourage timely payment. Messages labeled as Planning were sent earlier in the billing cycle to encourage financial planning for the corresponding next credit card payment.

Table A2: Placebo messages

Message number	Message content
1	Make a budget and save money!
2	See where your money went!
3	How are your balances? Check it now!
4	Your finances want your attention. Log in now!

This Table presents the content of four messages routinely sent by the Partner to its users. This messages were selected for having the same login rates (as tested in a 500 user pilot), as the credit card messages defined specifically for this intervention. This messages were sent to the robustness treatment group during the second billing cycle of the intervention.

Table A3: Treatment effects on spending by category

	ATE Full-Treatment			
	Coefficient	Std. error	T-stat	P-value
Salary	0.06976	0.06227	1.12034	0.26260
Bonus	0.00110	0.01111	0.09939	0.92083
Return	-0.00960	0.04194	-0.22899	0.81888
Other income	-0.00650	0.06934	-0.09370	0.92535
House	0.01321	0.03861	0.34220	0.73221
Residential bills	0.00531	0.04815	0.11020	0.91226
Health	0.03233	0.04736	0.68261	0.49487
Education	-0.01034	0.03090	-0.33469	0.73787
Transport	-0.00198	0.04449	-0.04444	0.96456
Market	0.00078	0.04628	0.01692	0.98650
House keeper	0.00005	0.01177	0.00413	0.99671
TV / INTERNET / Phone	0.00848	0.04541	0.18662	0.85196
ATM Withdrawals	-0.00535	0.05758	-0.09284	0.92603
Bars/ Restaurants	0.00453	0.04063	0.11152	0.91121
Leisure	0.01196	0.03899	0.30686	0.75896
Shopping	0.02354	0.04311	0.54601	0.58507
Personal care	0.00904	0.05076	0.17804	0.85870
Services and insurance	0.06073**	0.02922	2.07844	0.03770
Travel	-0.00882	0.04874	-0.18087	0.85647
Gifts/ Donations	0.03451	0.03882	0.88893	0.37406
Family/ Children	0.00684	0.01979	0.34558	0.72967
Work expenses	0.00722	0.01389	0.51990	0.60315
Other expenses	0.00010	0.01535	0.00639	0.99490
Taxes	-0.00966	0.03940	-0.24515	0.80635
Loan type	0.00903	0.04036	0.22382	0.82290
Loan type	0.00545	0.02281	0.23890	0.81119
Loan type	0.00658	0.00935	0.70398	0.48147
Other loans	0.01533	0.03466	0.44242	0.65819
Card bill payment	-0.04817	0.07003	-0.68789	0.49154
Recovered investment	0.01841	0.04723	0.38980	0.69670
Investment	0.03175	0.05037	0.63038	0.52847
Transfer	0.02535	0.03935	0.64407	0.51955
No category	0.03056	0.05612	0.54461	0.58604
Other personalized category	-0.00357	0.02964	-0.12055	0.90405

This Table presents the results from OLS regressions estimating equation 1 for various spending categories when the treatment corresponds to the Full-treatment. Standard errors clustered at the user level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table A4: Fraction of users with enough checking balance when LPF was charged

	Dep. Var: Enough balance * Late payment fee {0,1}		
	(1)	(2)	(3)
During*Pooled treatments	-0.0205***		
	(0.00585)		
During*Full-Treatment		-0.0274***	
		(0.00662)	
During*DD Alert			-0.0105
			(0.00778)
During	0.0314***	0.0314***	0.0314***
	(0.00427)	(0.00427)	(0.00427)
Observations	41,740	33,220	29,244
Number of clusters	10,435	8,305	7,311

This Table presents the results from OLS regressions estimating equation 1 for various treatments. $Full - treatment$ takes the value of one if a given user belongs to the Full-treatment group, that received five messages. *DD Alert* takes the value of one if a given user belongs to the DD Alert group, that received two messages. *Pooled Treatments* takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. *Enough Balance* takes the value of one when a given user has enough balance in her checking account to cover the full outstanding minimum payment on her credit card. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions.
*** p<0.01, ** p<0.05, * p<0.1

Table A5: A linear measure of liquidity

	Dependent variable: Paid CC Late Fee {0,1}		
	(1)	(2)	(3)
During*Pooled treatments	-0.0219***		
	(0.00842)		
During*Pooled treatments*LR	-0.00247		
	(0.00166)		
During*Full-Treatment		-0.0256***	
		(0.00960)	
During*LR*Full-Treatment		-0.00406**	
		(0.00185)	
During*DD Alert			-0.0161
			(0.0111)
During*LR*DD Alert			-0.000132
			(0.00222)
Liquidity ratio (LR)	-0.00129	-0.00129	-0.00129
	(0.00101)	(0.00101)	(0.00101)
During*LR	-0.000992	-0.000992	-0.000992
	(0.00123)	(0.00123)	(0.00123)
Pooled treatments*LR	0.000376		
	(0.00140)		
LR*Full-Treatment		0.000365	
		(0.00156)	
LR*DD Alert			0.000432
			(0.00189)
During	0.0568***	0.0568***	0.0568***
	(0.00610)	(0.00610)	(0.00610)
Observations	39,696	31,625	27,793
Number of clusters	10,415	8,289	7,296

This Table presents the results from OLS regressions estimating equation 2 for various treatments. $Full-treatment$ takes the value of one if a given user belongs to the Full-treatment group, that received five messages. $DD\ Alert$ takes the value of one if a given user belongs to the DD Alert group, that received two messages. $Pooled\ Treatments$ takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. $During$ takes the value of one during the treatment period. LR is defined as the ratio of checking account balance to credit card balance, for each given user. Standard errors clustered at the user level. *** p<0.01, ** p<0.05, * p<0.1

Table A6: A non-linear measure of liquidity

	Dependent variable: Paid CC Late Fee {0,1}		
	(1)	(2)	(3)
During*Pooled treatments*EB0	-0.0103		
	(0.0227)		
During*Pooled treatments*EB1	-0.0557*		
	(0.0293)		
During*Pooled treatments*EB2	-0.0282		
	(0.0274)		
During*Pooled treatments*EB3	-0.0350		
	(0.0522)		
During*Pooled treatments*EB4	-0.0204		
	(0.0563)		
During*Pooled treatments*EB5	-0.0408		
	(0.0271)		
During*Full-Treatment*EB0		-0.00821	
		(0.0262)	
During*Full-Treatment*EB1		-0.0488	
		(0.0334)	
During*Full-Treatment*EB2		-0.0398	
		(0.0313)	
During*Full-Treatment*EB3		-0.0600	
		(0.0604)	
During*Full-Treatment*EB4		-0.0104	
		(0.0649)	
During*Full-Treatment*EB5		-0.0556*	
		(0.0303)	
During*DD Alert*EB0			-0.0138
			(0.0297)
During*DD Alert*EB1			-0.0664*
			(0.0388)
During*DD Alert*EB2			-0.0112
			(0.0363)
During*DD Alert*EB3			0.00239
			(0.0664)
During*DD Alert*EB4			-0.0350
			(0.0740)
During*DD Alert*EB5			-0.0184
			(0.0364)
Observations	41,740	33,220	29,244
Number of clusters	10,435	8,305	7,311

This Table presents the results from OLS regressions estimating equation 2 for various treatments. $Full-treatment$ takes the value of one if a given user belongs to the Full-treatment group, that received five messages. $DD\ Alert$ takes the value of one if a given user belongs to the DD Alert group, that received two messages. $Pooled\ Treatments$ takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. $During$ takes the value of one during the treatment period. EB_i takes the value of one when the ration of checking account balance to credit balance is between i and $i+1$. Standard errors clustered at the user level. *** p<0.01, ** p<0.05, * p<0.1

Table A7: Persistence: Treatment effects after the intervention

	Dependent variable: Paid CC Late Fee {0,1}		
	(1)	(2)	(3)
Post*Pooled treatments	-0.00864		
	(0.00929)		
Post*Full-Treatment		-0.0164	
		(0.0107)	
Post*DD Alert			0.00271
			(0.0124)
Post	0.0125*	0.0125*	0.0125*
	(0.00658)	(0.00658)	(0.00658)
Observations	31,305	24,915	21,933
Number of clusters	10,435	8,305	7,311

This Table presents the results from OLS regressions estimating equation 1 for various treatments. $Full-treatment$ takes the value of one if a given user belongs to the Full-treatment group, that received five messages. *DD Alert* takes the value of one if a given user belongs to the DD Alert group, that received two messages. $Pooled\ Treatments$ takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and one billing cycle after the intervention, when no further messages were sent. *Post* takes the value of one for the one billing cycle after the intervention. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. *** p<0.01, ** p<0.05, * p<0.1

Table A8: Treatment effects and login into the app

	(1) Number of logins per month	(2) Paid CC Late Fee {0,1}	(3) Paid CC Late Fee {0,1}
During*Full-Treatment	-0.0259	0.0214*	
	(0.200)	(0.0116)	
During*Robustness Treatment			-0.00821
			(0.0106)
During	0.834***	0.0659***	0.0955***
	(0.140)	(0.00799)	(0.00656)
Omitted category	Robustness Treatment-group at baseline		Control Group at baseline
Observations	18,681	18,681	24,852
Number of clusters	6,227	6,227	8,284

This Table presents the results from OLS regressions estimating equation 2 with different outcomes. $Full - treatment$ takes the value of one if a given user belongs to the Full-treatment group, that received five messages. *Robustness* takes the value of one if a given user belongs to the Robustness treatment group that received five placebo messages. The information corresponds to two billing cycles before the intervention, and the second billing cycle of the intervention. *During* takes the value of one for the second billing cycle of the treatment period. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. *** p<0.01, ** p<0.05, * p<0.1

Table A9: Contractual information on the levels of selected fees for the three largest banks

	Credit Card Monthly Interest Rate (%)	Overdraft Monthly Interest Rate (%) a.k.a. "Cheque especial"	Extended Overdraft Fee and Monthly Interest Rate (%) a.k.a. "Adiantamiento"
Banco do Brasil	1.98 - 15.60	4.50 -12.30	\$R51.95 + 15.69%
Itau	1.90-9.90	3.58 -12.61	\$R54.90 + 17.61%
Bradesco	1.90-15.99	9.39 -12.63	\$R51.75 + (9.39 to 12.63)%

This Table presents some examples of the most common contingent fees linked to individuals' checking accounts and credit cards. Source: Banks' websites and Brazilian Central Bank.

Table A10: Interaction of paycheck arrival and attention shocks

	(1) Dep. Var: Late payment fee {0,1}
During*Pooled treatments	-0.0122
	(0.0136)
During*Pooled treatments*Paycheck close to CC DD	-0.0256
	(0.0182)
During*Paycheck close to CC DD	0.0184
	(0.0134)
During	0.0596***
	(0.0101)
Mean of Dep.var. when Paycheck close to CC DD =1 and During=1	0.221***
	(0.00227)
Observations	27,924
Number of aux_statement_id	6,981

This Table presents the results from OLS regressions estimating equation 2. *Pooled Treatments* takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. Only information from users that receive periodic paychecks is included. *Paycheck close to DD* takes the value of one for any given users whose paycheck is received no more than 15 days before her credit card due date. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. *** p<0.01, ** p<0.05, * p<0.1

Table A11: Overdraft and periodicity of paychecks

	(1) Dep. Var: Overdraft fee {0,1}
During*Pooled treatments* Bi-weekly or weekly paycheck	0.0234**
	(0.0113)
During*Bi-weekly or weekly paycheck	0.0231***
	(0.00814)
During*Pooled treatments	-0.000157
	(0.00668)
During	0.0518***
	(0.00477)
Observations	41,740
Number of clusters	10,435

This Table presents the results from OLS regressions estimating equation 2. *Pooled Treatments* takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. Only information from users that receive periodic paychecks is included. $Bi-weekly\ or\ Weekly\ paycheck$ takes the value of one for any given users whose paycheck is received with bi-weekly or weekly periodicity. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. *** p<0.01, ** p<0.05, * p<0.1

Table A12: Total value of OD fees and credit card fees

	Dependent variable: Value of fees ($R)					
	(1)	(2)	(3)	(4)	(5)	(6)
	OD	CC	OD	CC	OD	CC
During*Pooled treatments	2.729***	-4.247***	0.696	-5.320***	2.928***	-4.228***
	(0.723)	(1.150)	(0.446)	(1.193)	(0.786)	(1.251)
During*Pooled treatments*OD at baseline			7.314***	3.893		
			(2.288)	(2.943)		
During*Pooled treatments*Logins per month					-0.0414	-0.00443
					(0.0600)	(0.0887)
During	5.432***	7.753***	3.743***	6.494***	4.579***	7.516***
	(0.482)	(0.845)	(0.293)	(0.890)	(0.530)	(0.921)
During*OD at baseline			5.809***	4.332**		
			(1.520)	(2.126)		
During*Logins per month					0.169***	0.0470
					(0.0413)	(0.0645)
Observations	41,740	41,740	41,740	41,740	41,740	41,740
Number of clusters	10,435	10,435	10,435	10,435	10,435	10,435

This Table presents the results from OLS regressions estimating equations 1 and 2 for various outcomes. *Pooled Treatments* takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. *OD at baseline* takes the value of one if a given user incurred overdraft at baseline. *Logins per month* indicated the number of logins per user and billing cycle. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table A13: Treatment effect on the total value of fees paid by users that paid overdraft at baseline

	Dep. var.: Cost of contingent fees ($R)		
	(1)	(2)	(3)
	OD	CC	Total
During*Pooled treatments	8.009***	-1.427	6.582*
	(2.245)	(2.691)	(3.728)
During	9.552***	10.83***	20.38***
	(1.492)	(1.931)	(2.612)
Observations	11,984	11,984	11,984
Number of clusters	2,996	2,996	2,996

This Table presents the results from OLS regressions estimating equation 1 with different outcomes. *Pooled Treatments* takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions.
*** p<0.01, ** p<0.05, * p<0.1

Table A14: Treatment effect on the total value of fees paid by users that did not incurred overdraft at baseline

	Dep. var.: Cost of contingent fees ($R)		
	(1)	(2)	(3)
	OD	CC	Total
During*Pooled treatments	0.696	-5.320***	-4.625***
	(0.446)	(1.193)	(1.286)
During	3.743***	6.494***	10.24***
	(0.293)	(0.890)	(0.946)
Observations	29,756	29,756	29,756
Number of clusters	7,439	7,439	7,439

This Table presents the results from OLS regressions estimating equation 1 with different outcomes. *Pooled Treatments* takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions.
*** p<0.01, ** p<0.05, * p<0.1

Table A15: Treatment effect on the log of total contingent fees

Dependent variable: Log of value of total fees {0,1}	(1)	(2)	(3)
During*Pooled treatments*OD at baseline		0.150**	
		(0.0735)	
During*Pooled treatments*Logins per month			-0.00348
			(0.00298)
During*Pooled treatments	-0.0333	-0.0761**	-0.0162
	(0.0317)	(0.0361)	(0.0350)
During*OD at baseline		-0.0409	
		(0.0522)	
During*Logins per month			0.00996***
			(0.00217)
During	0.409***	0.421***	0.359***
	(0.0226)	(0.0259)	(0.0251)
Observations	41,740	41,740	41,740
Number of clusters	10,435	10,435	10,435

This Table presents the results from OLS regressions estimating equation 1 for various outcomes. *Pooled Treatments* takes the value of one if a given user belongs to either of the two treatment groups. The information corresponds to two billing cycles before the intervention, and two billing cycles during the intervention. *During* takes the value of one for billing cycles during the treatment period. *OD at baseline* takes the value of one if a given user incurred overdraft at baseline. *Logins per month* indicated the number of logins per user and billing cycle. Standard errors clustered at the user level in parenthesis. Individual fixed effects included in all regressions. *** p<0.01, ** p<0.05, * p<0.1

Table A16: Survey results

	N	Yes (%)	No (%)		
You recently received reminders about your credit card DD Do you want to continue receiving them?	1028	77.5	22.5		
	N	1 or 2 (%)	3 or 4 (%)	5 or more (%)	None (%)
How many reminders would you like to receive?	1013	57.5	16.1	14.6	11.9
	N	A few days before the DD	A few days after I receive my salary (%)	In both dates (%)	In none of those dates (%)
When would you like to receive the messages?	1003	46.9	16.2	28.2	8.7

This Table presents the results of a survey sent to each user in the treatment group 10 days after they received the last message. The survey was sent via mobile phone, as described in section 3. The order of the questions and options was randomized across different users.