Do consumers borrow on their cheapest credit card? Evidence from Mexico*†

Alejandro Ponce‡
World Bank

Enrique Seira
Secretaria de Economia, Mexico

Guillermo Zamarripa
SHCP, Mexico

First version: November 2007

This version: April 2009

Abstract

Most cardholders have more than one credit card, yet, it is not evident how these individuals manage their accounts. In this paper we construct a novel data set that includes information on all the credit cards held by more than 10,000 consumers in Mexico in 2004 and 2005 and empirically study the intra-temporal allocation of debt, payments and purchases among the credit cards consumers already hold. We find that the difference in the interest rates between homogeneous cards is not an important determinant of allocations. On average, cardholders forego potential savings for a sum that amounts to 16% of their financing cost. We show that non-price determinants of allocations have more explanatory power than interest rates. We find that consumers tend to put a larger fraction of their monthly payments and purchases on the card they spent more on during the preceding billing period, regardless of their interest rate ranking. The most compelling explanation relates to mental accounting and financial unsophistication. The low price sensitivity can explain why high interest rates prevail in this market, regardless of any search or switching cost.

KEYWORDS: Credit cards; household finance; debt heuristics; consumer behavior; Mexico.

JEL Classifications: D12 , D14, G20

---

*We are grateful to Jon Levin, for numerous discussions, revisions and comments. We thank Susan Athey, Tim Bresnahan, Simeon Djankov, Liran Einav, Josh Frank, Caroline Hoxby, Soohyung Lee, Brigitte Madrian, Ulrike Malmendier, Pedro Miranda, Sriniketh Nagavarapu, Roger Noll, Adriano Rampini, Felix Reichling, Joanne Yoong, and participants at Stanford University IO, the World Bank Research Group, and the Federal Reserve Bank of Boston for comments and suggestions. Ponce and Seira acknowledge the support of the Dissertation Grant of the Federal Reserve Bank of Boston's Research Center on Behavioral Economics and Decision-Making. Ponce appreciates the support of the Kohlhagen Fellowship Fund at Stanford University and Mexico's National Council for Science and Technology (CONACYT). All errors remain our own.

†The views expressed in this paper are those of the authors and do not necessarily reflect the positions or policies of the World Bank, the Secretaria de Hacienda y Credito Publico or the Secretaria de Economia of Mexico.

‡**Corresponding author:** P4-423, World Bank, 1818 H St. NW, Washington DC, 20433. E-mail: aponcerodriguez@worldbank.org

1 Introduction

Credit cards are the most popular source of consumer credit. In fact, most individuals hold not one, but many of them. Yet there are no studies examining how people with multiple credit cards manage these accounts. In this paper, we fill this gap by investigating how consumers with several credit cards allocate their debt, payments and purchases among them. Our primary interest is to assess the extent to which interest rates influence individuals' decisions regarding which credit card to borrow on. The theory behind this decision is simple. Almost any model of consumer choice predicts that in the presence of perfect information and homogeneous goods, individuals should borrow primarily on the credit card with the lowest interest rate. More specifically, if individuals are aware of the difference in prices, credit cards are homogeneous in all relevant respects except for the interest rates, and consumers are not constrained by the contractual features of their cards (e.g. credit limits and minimum payments), we should observe individuals allocating their debt, if any, to the card with the lowest interest rate. Alternatively, if consumers are unsophisticated, have biased preferences, are poorly informed, or have to perform costly calculations, they could make their financial decision on the basis of factors other than prices. This view is related to the idea of financial unsophistication [Bernheim and Garrett (2003), Hastings and Tejeda-Ashton (2008), Lusardi and Tufano (2008)], unawareness, and mental accounting [Thaler (1985), Thaler (1999), Prelec and Loewenstein (1998), Ranyard, Hinkley, Williamson and McHugh (2006), and Choi, Laibson and Madrian (forthcoming)]. While the reasons for such behavior could be numerous, the point is that borrowers could make their allocation decisions using different criteria than monetary costs. This is not only costly, but could have market equilibrium implications as high and disperse prices could prevail in the market even without any switching or search cost.

We investigate which factors, including interest rate differentials, have a significant influence on the fraction of debt, payments, and purchases that consumers allocate to a certain card. We concentrate on the intra-temporal allocation across cards and take the total monthly levels of debt, purchases, and payments as given. This is to simplify our analysis at little cost, as dynamics seems to play a minor role on the way consumers manage their debt among their different cards.[1] More specifically, there are three cases in which dynamics may influence which card consumers select to borrow on: Non-divisible purchases, low realizations of the income stream, and anticipated changes in the credit limit. For example, consumers may pay off their cheap credit card to free-up credit for a large purchase, or they may borrow from expensive credit cards today because they believe they will be able to pay the debt back in the future, although ex-post, they may not be able to do

[1] Conceptually, we can separate the consumer's decision problem into two stages: In the first stage, for each period, consumers select the total amount of purchases and payments to be made with their credit cards on that month, and as a result, the total amount of debt to carry on. In the second stage, once individuals have selected the totals, they decide on the allocation of those purchases and repayments between the cards they already hold.

so [Ausubel (1991)]. Alternatively, if banks consider borrowing in a non-linear way when deciding to boost a credit line, consumers could anticipate such rule and borrow in all their credit cards, regardless of the interest rates, in order to relax their liquidity constraints in the future. Despite these concerns, we believe these considerations are quantitatively small.

In order to properly study allocation decisions, we require a panel of high-quality data with information on all the credit cards held by each consumer. We constructed such data set with the collaboration of the Credit Bureau and the three largest commercial banks in Mexico. First, we solicited the Credit Bureau to draw a random sample of 100,000 consumers who had at least one credit card by November 2004. Then, the Ministry of Finance of the Mexican Government asked to the three largest banks to provide monthly information for each one of the accounts in the sample. The result is a proprietary novel data set that enables us to link consumers with most of their monthly credit card information, provided that the credit cards were issued by one of the three cooperative banks. Since the industry is highly concentrated, by using information from these three banks, we were able to match the entire credit card activity of 69% of the card holders in our original sample.[2] The data are of very high quality, containing the Credit Bureau reports as well as the entire credit card records observed by our three banks, including monthly debt, purchases, cash advances, payments, interest rates, credit limits, and socio-demographic characteristics during part of 2004 and 2005.

Our empirical analysis exploits the variation provided by the difference in the interest rate across comparable credit cards already held by each consumer. This difference is, on average, 1.1% per month. Although interest rates are set by the product, there is substantial cross section and time series variation available. More specifically, there are three reasons why consumers could have credit cards with different interest rates in a given period. First, consumers may simply have chosen two cards with different interest rates. Second, the 'effective' interest rate of a given card could be lower because the bank granted the rebate as a complementary service when opening a checking or saving account or because the borrower performs well during the lifetime of the card. Third, the bank could send a temporary low interest rate offer. The cross section and time series variation in the difference in the interest rates allows us to identify whether consumers pay attention to these differentials when allocating their monthly payments and purchases. For simplicity, our main analysis focuses on individuals with two credit cards whose credit cards have a similar reward structure.

Our research looks mainly at the allocation of payments and purchases, rather than the allocation of debt. This is because, by focusing on flows rather than stocks, we are able to control for liquidity constraints and balance transfer costs.[3] In addition, by looking at immediate choices, we

[2] In recent years, the three largest banks have concentrated about 80% of the market in terms of number of accounts and total debt levels.

[3] Transferring balances in Mexico could be particularly expensive. Before 2004, balance transfers were not allowed.

are in better shape to unravel specific decision rules than we would be if studying a state variable such as debt. Finally, throughout our analysis, we control for factors such as credit limits and minimum monthly payments that could influence consumers' allocation decisions.

We find that the difference in the interest rate between consumers' credit cards is a very poor predictor of how individuals allocate their debt, purchases and payments between accounts. Among individuals borrowing on credit cards, half of them borrow on their more expensive card. These results hold when we study the allocation of purchases and payments, and control for the contractual features of credit cards, namely credit limits and minimum payments. These results are robust to alternative samples designed to exclude observations with small differences in interest rates, distinct billing cycles or small levels of debt. We conclude that individuals do not borrow on their cheaper card despite it could be feasible for them to do so. On average, these cardholders forego potential savings for a sum that amounts to 16% of their financing cost.

We explore several explanations to our findings including unawareness of interest rates, small stakes, and mental accounting/financial unsophistication. We study first whether consumers are aware of the interest rate charged on their credit cards by examining their response to low-interest rate offers. We find that individuals respond substantially to temporary reductions in the interest rate, suggesting that they are aware of prices and care about small potential savings. Interestingly, this result confirms that individuals make the most of low-interest rate offers, but leave money on the table by borrowing on their more expensive cards. Although this is inconsistent with the standard cost-minimizing framework, it is consistent with the theory of mental accounting. According to this theory, people tend to categorize their consumption into mental accounts and make decisions within a narrow spending context (represented here by credit cards). In section 5.3, we provide additional evidence suggesting that this is indeed the case. First, we find that individuals tend to purchase with the card which has less outstanding balances while repaying the card which has more debt. Consumers who engage in mental accounting might prefer to repay cards that have been spent up to its limits to 'mentally' reduce the stress that a single large debt put on them. Second, we show that individuals stick to their initial choices and repeat purchases with the same card. This behavior is related to the concept of inertia in the marketing literature and is referred to as "front of the wallet" behavior in the credit card jargon. If individuals categorize their purchases, they are likely to stick to a certain card and repeat purchases. In this sense, inertial tendencies are consistent with the idea of categorization and mental accounting.

Mental accounting, however, is not the only explanation. An alternative story is that individuals do not understand how credit cards work. For example, individuals may pay the credit card

In January 2004, the Central Bank of Mexico issued a new law mandating commercial banks to allow consumers to transfer balances from other credit cards [Ley para la Transparencia y el Ordenamiento de los Servicios Financieros (2004)]. Nonetheless, despite the high expectations raised by this law, balance transfers are not frequent, and interest rates have reportedly not changed.

with more debt, with the 'wrong' idea that it is more expensive because it accrues more interest. Unfortunately, because of data limitations, we cannot fully distinguish between these hypotheses. A more reasonable explanation is that financially unsophisticated consumers rely on mental accounts if these processes serve as substitutes for otherwise complex financial decisions. In this regard, lack of financial literacy and mental accounting are not mutually exclusive.

To our knowledge, this is the first paper to present field evidence on how consumers periodically allocate their debt among the financial instruments available to them. While previous literature has looked at consumers' departure from standard utility maximization, it is fair to say that most studies have focused on inter-temporal decisions and biased expectations about future preferences and future consumption [Miravete (2003), Agarwal, Chomsisengphet, Liu and Souleles (2005), Shui and Ausubel (2005), DellaVigna and Malmendier (2004), and DellaVigna and Malmendier (2006)]. Our analysis departs from these studies by focusing on 'intra-temporal' instead of 'inter-temporal' decisions. As a result, our analysis does not involve any uncertainty, allowing us to be agnostic about the form of the utility function and time varying shocks. The closest analysis to ours is Benartzi and Thaler (2001), who show that consumers use naive diversification strategies and divide their contribution evenly across the funds offered in their saving plans. Our paper complements their work in at least three ways. First, we study consumers' allocation of short term liabilities instead of long term assets. Second, in our setting, consumers make decisions every month and not just once, thus increasing the scope for learning [Agarwal, Driscoll, Gabaix and Laibson (2008)]. Third, it is likely that the complexity and uncertainty of our environment is lower, as the planning horizons are shorter.

Overall, our paper contributes to the literature and policy debate in three ways. First, we add to the household finance literature by documenting not only one of the most frequent financial choices individuals actually make, but also the discrepancies between observed and optimal behaviors when making these decisions [Campbell (2006)]. Second, we contribute to the so-called "credit card puzzle" debate. Several papers have argued that the fact that consumers borrow high and lend low at the same time represents a puzzle for the neoclassical models of consumer choice [Gross and Souleles (2002) and Bertaut and Haliassos (2006)].[4] Here, we address a similar problem but in a simpler setting and find that consumers in fact, leave money on the table. Finally, the paper adds to the behavioral literature on heuristics and market outcomes [Ellison (forthcoming)]. From a policy perspective, our analysis is relevant for at least two reasons. First, not paying attention to interest rates can be quite costly for consumers. Second, as already mentioned, if individuals allocate their debt based on rules other than interest rates, high and disperse prices could prevail in the market even without switching or search costs. This argument could explain why the effect

[4]Zinman (2007) disagrees with this literature by arguing that credit card debt is not a close substitute for cash. His critique does not apply here though, as we study arbitrage across homogeneous cards differing only by the interest rates

of the law that allowed consumers to transfer balances in Mexico was limited.

The remainder of the paper is organized as follows. In Section 2, we present a description of the Mexican credit card industry. In Section 3, we introduce the data used in our analysis. Section 4 study whether individuals minimize their interest cost. In Section 5, we explore possible explanations for the results. Section 6 analyzes some of the supply side implications. Section 7 concludes.

2 Credit cards in Mexico

2.1 The credit card industry in Mexico

As of December 2005, 12.2 million credit cards were active in Mexico. The value of transactions equaled 185 billion pesos and the outstanding balances on revolving accounts surpassed 148 billion pesos up from 47 billion in 2002. Among consumers who had at least one credit card in December 2005, 61% carried more than one. This is consistent with a small fraction of people having access to this market. For example, the number of credit cards per individual of working age has ranged between 0.19 and 0.28 over these years, and the number of households who have access to at least one credit card has barely reached 5% of the population.[5]

The credit card market in Mexico is highly concentrated. Three national banks have consistently controlled more than 80% of the market (measured by the number of cards or outstanding debt) from 2001 to 2007. The average credit card has had an associated interest rate of 28 percentage points above the federal discount rate [Banxico (2006)] and fees and service charges have been ranked as the highest in Latin America [Avalos and Hernandez (2006)].

During all this time, the credit card business has operated without ceilings on interest rates and fees. Most regulation has focused on the disclosure of information to consumers at the time of contracting. In recent years though, regulation to improve the competitive environment of the industry has been enacted. In January 2004, the Central Bank of Mexico issued a new law mandating commercial banks to allow consumers to transfer balances from other credit cards.[6] The introduction of balance transfers has allowed banks, particularly entrants, to compete with low interest rate products. Nonetheless, despite the expectations raised by this law, interest rates on most cards are still high, balance transfers are reportedly rare and consumers still use the cards with higher interest rates.

[5]The average Peso-Dollar exchange rate for December 2005 was 10.62. This rate fluctuated between 11.42 pesos per dollar in January 2005 and 10.54 in January 2006. Source: Banco de Mexico (http://www.banxico.org.mx/PortalesEspecializados/tiposCambio/indicadores.html)

[6]Ley para la Transparencia y el Ordenamiento de los Servicios Financieros (2004)

2.2 The credit card contract

In Mexico, a typical credit card contract specifies the credit limit, the applicable fees, the benefits of the card, as well as the rules on how to compute the minimum payment and the finance charge on the outstanding balance. With the exception of the credit limit, all terms are set by the product, not the consumer. Cards differ from each other by the benefits they offer, by non-pecuniary features, by the fees and service charges, and by the interest rate (or more specifically, by the margin over the Prime Rate). Yet, there are several reasons why, in a given month, individuals with the same card may have different interest rates. We describe them in section 3.3.

The credit card activity is organized around the monthly billing cycle. Each month the cardholder receives an account statement specifying the purchases, cash advances and payments made during the previous cycle. The statement also includes the fees and interest accrued during the period as well as the outstanding balances, minimum payment due and current interest rate.

Figure 1 illustrates a typical billing cycle. This figure includes nearly all variables we will use in the analysis. Most variables are determined during the billing cycle, but observed at the end of the period. The exceptions are the Minimum Monthly Payment $PMin_{t+1}$ and the closing balances B_t, which are determined at the end of the cycle. During the billing cycle, the card holder can use the card for purchases or cash advances X_t up to the credit limit L_t assigned previously by the bank. If the credit limit is surpassed, an overlimit fee of approximately 100 pesos is charged. The cardholder can use up to 20% above the credit limit; beyond this point, the card is blocked.

Throughout the paper, we use the Average Daily Balances (ADB) as the key variable to measure interest-incurring debt. The ADB are calculated by the bank by taking the sum of the daily balances on each day of the period, including current purchases and cash advances, minus any payments received, divided by the number of days in the cycle. For a given period t, this variable is labeled D_t. The accrued interest is calculated by multiplying a pre-specified interest rate times the ADB in a given period. The card holder has a grace period of 20 days, counting from the closing date of the previous cycle, to make the minimum payment. This way, most monthly payments P_t correspond to those payments made to pay back the ending balance of the previous period B_{t-1} (or at least the associated Minimum Monthly Payment $PMin_t$). If the card holder pays the closing balance B_{t-1} in full, the ADB equal zero and no interest is accrued during the month. If the cardholder makes at least the minimum payment but does not pay the balance in full, the ADB are positive and an interest is accrued. If the cardholder fails to make even the minimum payment, the issuer charges a late payment fee of $230 pesos and extra interest, and the card is blocked until the time the payment is received.

3 Data

3.1 The data set

We constructed a novel panel date set of credit card usage in Mexico. The data collection involved the collaboration of the Credit Bureau and the three largest commercial banks in Mexico. The result is a unique proprietary data set that enables us to link consumers with all their credit card information. Our data contain the Credit Bureau reports as well as the entire credit card records observed by these three banks, including monthly balances, purchases, cash advances, payments, interest rates and some socio-demographic characteristics during part of 2004 and 2005.

To construct the data set, we asked the Credit Bureau to draw a random sample of 100,000 consumers who had at least one credit card in November 2004. This data set contains a unique identifier for each individual, all the account numbers for each one of them, the delinquency status of these accounts as well as some socio-demographic characteristics. Then, with the account identifiers at hand, the Ministry of Finance of the Mexican government asked the three largest banks to provide monthly information for every account in the sample. This data set comprises all the aggregate information contained in the accounts' monthly billing statements, including debt, purchases, cash advances, payments, minimum payments, credit limits and interest rates. The data also contain information about the 'type' of card (e.g. Classic, Gold, Platinum, etc.) and some additional demographics recorded at the time of solicitation. Jointly, the two data sets allow us to follow, for each individual, the monthly changes in their credit card accounts.

Due to the administrative costs involved in acquiring the data, we focused only on the three largest banks. This is not a problem for our analysis: since the industry is highly concentrated, we were able to match the entire credit card activity of 69% of the card holders in our original sample. The data cover the period 2004-2005 for two banks, and from October 2004 to September 2005 for the third one. Finally, to ensure confidentiality, all the information was provided in such a way that it is impossible to infer the identity of the account holders.

3.2 Sample construction

For the sake of simplicity, we focus on individuals holding 'exactly' two 'active' and 'comparable' credit cards issued by one of our three banks.[7] We consider a card as 'active' if the account remained open between January 1, 2004 and December 31, 2005, regardless of the payments or purchases made during this period. We define two credit cards as 'comparable' if their non-price attributes are similar. More specifically, we characterize two cards as 'comparable' if both cards are of the same 'type' (e.g. both cards are Gold) or if they have the same rewards programs. In general, cards

[7] we removed all the store cards (e.g. Wal-Mart Card) from the sample, as these cards can only be used in specific businesses and therefore are not good substitutes for the credit cards issued by banks.

of the same type issued by different banks share the same rewards structure. It is worth noting that there is little variation in the types of card held by most consumers; for example, around 65% of the cards are labeled as 'Classic'. In addition, during our sample period, rewards were just beginning to expand in Mexico.

We eliminated individuals who had less than six months of information in any of the variables required for our analysis. Due to administrative and technical issues while extracting the information, one of the banks provided data starting in late 2004, so our panel is unbalanced. This leaves us with a sample of 114,720 consumer-months. Out of these observations though, interest paying debt is positive in at least one of the cards for 103,343 consumer-months. It is only for these observations that the problem of which card to borrow from is really relevant. Therefore, our final sample contains 10,335 consumers and 103,343 consumer-months. The median number of periods by consumer in our final sample is 11, and the 95th percentile is 16.

We select the sample this way to simplify the allocation problem and to have a clean comparison between homogeneous cards. However, this selection precludes us from generalizing our results to the entire population of credit card holders. We acknowledge this limitation but point out two facts. First, approximately 17% of card holders in Mexico had exactly two active credit cards as of December 2005, making our analysis directly relevant for this population.[8] Second, there is no obvious reason why individuals with two cards allocate their debt less optimally than consumers with three or more cards. In fact, it is quite possible that individuals with more credit cards leave more money on the table. This could happen for two reasons. First, all else equal, it is likely the more cards individuals hold, the more complex their decision problems are, and therefore, the more likely it is that they misallocate their debt. Second, since the number of cards is positively correlated with total outstanding balances, it is possible that mistakes are more costly for individuals with more cards simply because the amounts at stake are larger.

3.3 Interest Rate differences

We exploit the variation provided by the difference in the interest rate across the credit cards already held by each consumer. In Mexico, interest rates are set by the product. Yet, there are three reasons why, in a given month, individuals may hold two cards with different interest rates. First, consumers may have simply chosen credit cards with different interest rates. Second, the interest rate in a given card could be lower because the bank granted a rebate as a complementary service for opening a checking or saving account or because borrowers have performed well during the lifetime of the card. Third, the bank could send a temporary reduction in the interest rate or teaser rate offer. These offers are targeted by banks to existing cardholders with a duration ranging

[8] In the original random sample derived from the Credit Bureau, 39% of consumers had one active credit card, 17% had two active credit cards, 10% had three credit cards, while 34% had more than three

from one to six months. Teaser rate offers are widespread in Mexico: in 2005, the average credit card had an active TRO for approximately 2.1 months.

The interest rate differences are quite significant. Figure 2 shows the distribution of the monthly interest rate differential among consumers' credit cards. The distribution has three modes and substantial variation. The median difference is 1.2%, and the overall standard deviation is 1.18% points. As expected, there is substantial variation within and among consumers. The standard deviations within consumers and between consumers are 0.77% and 0.92%, respectively.

3.4 Descriptive statistics

Table 1 presents summary statistics for the main variables we will use in our analysis. The top panel shows the mean of the variables that appear in the cardholder's balance statement once we combine the two credit cards held by each individual. Here, the unit of analysis is a consumer-month. Some facts stand out. First, the amount of interest paying debt is non-negligible. The average total debt equals $22,136 pesos, which corresponds to more than two times the median monthly income of the consumers in our sample. As a matter of fact, consumers use their credit lines extensively. The average utilization rate of the combined cards is 61%. Interestingly, consumers borrow on their credit cards even though the average monthly debt-weighted interest rate equals 2.54%. Second, most of the time, consumers pay interest and 42% of the time, individuals are charged a late payment or an overlimit fee. Finally, out of the consumers who borrow, 82% allocates their debt in both cards.

Monthly statistics can mask substantial cross-time heterogeneity for a given consumer. Panel B in Table 1 displays some summary statistics at the consumer level, that is, once we take the average across-months for each individual. Yet the picture looks quite similar. About 90% of consumers are charged interest half the time (in fact, 77% of consumers in our sample always incur interest). Similarly, 74% of consumers borrows in both cards at least half of the time. Taken together, these figures suggest that consumers in our sample seem to be heavily indebted and usually borrow in both cards.

The bottom panel presents some demographics. This information comes from the credit bureau reports or from the banks' records at the time of origination. Unfortunately, banks do not keep these records for long periods; consequently, the number of observations varies greatly. We deflate the monthly income using the Banco de Mexico consumer price index. The median monthly income in our data is $10,000 pesos. The average age is 44 years while the average tenure with the oldest credit card is 8 years.

4 Empirical analysis

4.1 Debt Allocation

In this section, we study whether individuals minimize their interest cost. Given the features of our sample, we would expect individuals to borrow as much as possible on the credit card with the lower interest rate, as long as they avoid paying over-limit fees. We start by analyzing how consumers allocate their debt among the credit cards they already hold, conditioning on their total credit card borrowing. There are two points deserving attention here. First, our focus is on how individuals pick which card to borrow on, among those available in their wallets. That is, we do not study how consumers decide which card to apply for, considering the available options in the market. Although relevant, this decision is made rarely and involves considerations such as the number of products available as well as search and switching costs. Instead, we focus on a simpler and recurrent decision, which allows us to abstract away from such factors and concentrate on the decision making process. Second, our focus is on the 'intra-temporal' instead of 'inter-temporal' decisions. That is, we do not study how much consumers borrow, or whether they borrow optimally given their inter-temporal constraints and time preferences. Instead, for every period, we take consumers' credit card debt as given, and investigate which cards individuals choose to borrow on. There are advantages and disadvantages to conditioning on total borrowing and abstracting from dynamic considerations. The main advantage is that we are able to focus on a much simpler problem. The main concern is that allocation may depend on the expected income and expectations about changes in the credit card features. For example, consumers may borrow from expensive credit cards today because they believe they will be able to pay the debt back in the future, although ex-post, they may not be able to do it [Ausubel (1991)]. Alternatively, if banks consider borrowing in a non-linear way when deciding to boost a credit line, consumers could anticipate such rule and borrow in all their credit cards, regardless of the interest rates, in order to relax their liquidity constraints in the future. We will address some of these concerns later on, when we study the allocation of payments and purchases as well as the persistence of the allocation decisions.

We begin by looking at the monthly allocation of debt across consumers' credit cards. The upper panel of Figure 3 shows the histogram of the share of interest-paying debt allocated to the cheaper card. The unit of analysis here is a consumer-month. The distribution is quite symmetric. Most of the time, consumers borrow on both credit cards. In fact, 48% of the time, consumers hold more than fifty percent of debt on their more expensive credit card. Consumers could borrow on the more expensive card because the credit limit of the cheaper card binds. To get a sense of the importance of the limit constraint, the lower panel of Figure 3 presents the distribution that would arise if consumers allocate debt up to the credit limit on the cheaper card, and the rest on the more expensive one. The histogram looks quite different from the previous figure. The results from the nonparametric Mann-Whitney test indicate that the difference between these distributions is statistically significant. In general, optimality implies a large share of debt to be put in the cheaper card. However, the striking difference between the upper and lower panels certainly suggests that feasibility or credit limit constraints do not fully explain why consumers allocate debt to the more expensive cards.

Studying debt may seem appropriate because interest is charged on outstanding balances. Nonetheless, focusing on debt is misleading if we want to study consumers' choices. On one hand, individuals decide on purchases, cash advances and payments; debt is simply the cumulative flow of these variables. On the other hand, individuals could exhibit a different behavior when shopping at stores than when handling the credit card payments (e.g. as in accountant-shopper models [Bertaut and Haliassos (2001)]). Because of these reasons, from now on, we concentrate on the allocation of new purchases, cash advances and payments. This approach has other advantages as it allows us to avoid double counting of one-time 'mistakes' and consider the role that liquidity constraints could play in consumers' choices. For example, if a consumer makes a purchase with an expensive card and lacks the money to pay off this debt afterward, this purchase will be carried forward as debt for many periods. Our approach allows us to evaluate this decision separately. On one hand, we are able to assess the allocation of purchases as a one-time decision. On the other hand, we can study how individuals regularly allocate their monthly payments to pay off this debt, conditional on the total amount repaid to both cards. Finally, this approach allows us to rule out balance transfer costs as an explanation for the observed allocation of debt.[9] If consumers minimize their financing costs when they decide which card to use and which card to pay off, then it is very likely that balance transfer costs explain the observed allocation of debt. If not, it is quite likely that the allocation of credit card debt is driven by some other factors.

4.2 Allocation of purchases and cash advances

In this section we describe how consumers allocate their monthly purchases and cash advances, conditional on the purchases and cash advances undertaken with both cards in a given month. Although purchases and cash advances are not exactly equal, for simplicity we will treat them alike and sum them up. Therefore, hereafter, we will refer to purchases and cash advances as simply purchases. This analysis presents some challenges. First, since banks usually put credit cards on hold (i.e. consumers are not able to use them) when individuals fail to make the minimum payment due, consumers who turn delinquent do not really have an allocation choice to make, as they cannot purchase with at least one credit card. Second, as cardholders may use their

[9]Balance transfer costs include not only the fees for transferring balances but more importantly, consumers' switching costs.

cards with the expectation to pay them off in the near future (for example, when receiving the billing statement), the allocation of purchases might be a dynamic decision. We acknowledge this, but argue an important point. Regardless of consumers' expectations, interest is charged immediately on purchases if consumers did not pay the previous balances in full. Therefore, dynamic considerations are only relevant when consumers start borrowing on any card. Finally, consumers may choose to purchase with an expensive card because the outstanding balances on the cheaper one are close to the credit limit. In the analysis below, we take these concerns into account.

Our focus is on determining how the fraction of purchases allocated to a random card differs when the interest rate differential varies. In doing so, we examine whether this fraction differ for distinct situations or groups. For this, we select randomly one of consumers' credit cards and label it as card 1. We keep this label invariant throughout the paper. For such analysis, we also take the previous concerns into consideration and restrict the sample to include only observations in which consumers made at least the minimum due on both cards (to avoid cases in which one card was put "on hold"), had outstanding balances in their two cards, did not surpass the credit limit on any of them, and in which total purchases during that period could fit into any of the cards. The credit limit binds for 43% of the observations, while for the remaining 57%, it is feasible to allocate all purchases on the cheaper card. This leaves us with 32,267 observations.

We begin our analysis by providing a nonparametric description of the relationship between the allocation of purchases and the interest rate gap. Figure 4 displays the nonparametric kernel regressions of the fraction of purchases allocated to card 1 on the interest rate gap, defined as (r_1 - r_2). The figure also presents the scatter plot of a random sample of 15 percent of the observations. In this regression, we use the Gaussian kernel with the bandwidth chosen according to Silverman's rule of thumb for a normal density. The figure shows that the fraction of purchases allocated to card 1 diminishes slightly when card 1 becomes more expensive. Nonetheless, this fraction varies by less than 0.2 over the entire distribution of the interest rate gap. This figure suggests that interest rate differentials do not seem to determine the way consumers allocate their monthly purchases.[10] Still, we cannot be certain that this is the case for all consumers or all circumstances. To account for other factors that may affect the allocation of purchases, we follow a parametric approach and model the fraction of purchases allocated to a given card as a linear function of a time-varying indicator variable of the interest rate differential. Specifically, our estimation of interest is:

$$SH\ X1_{it} = \beta_0 + \beta_1 CHEAP1_{it} + \eta_{it} \tag{1}$$

[10]The distribution of the share of purchases allocated to the expensive card for this sample is bimodal with about 25% selecting only the cheaper card, 25% making all of their purchases with the more expensive one, and the rest mixing up their two cards. When we consider the distribution at the individual level the spikes disappear, suggesting that our results may not be only driven by heterogeneous consumers behaving always in the same way, but rather by heterogeneous consumers behaving differently over time.

where $SH\ X1_{it}$ denotes the fraction of purchases allocated to card 1 by consumer i at time t, $CHEAP1_{it}$ is a dummy variable equal to one if the interest rate on card 1 is lower than the rate on card 2, and η_{it} is the error term. When two cards have the same interest rate, we selected randomly which card (not card-month) had a larger value. We followed this approach to avoid dropping observations with potentially useful information. Our results however, are robust to excluding these cases.

In this regression, an observation is a consumer-month. Since the main covariate is a dummy variable, this specification is a comparison of means. The coefficient on the term $CHEAP1_{it}$ captures the average difference in $SH\ X1_{it}$ when card 1 has a lower interest rate than card 2. Under the null hypothesis, individuals minimize interest costs and allocate all their purchases to the cheaper card, so $\beta_0+\beta_1$ should be equal to one. The error term, η_{it}, captures all factors influencing the dependent variable other than differences in the interest rates. Since this variable is defined as the fraction of purchases made with card 1, and this label is assigned randomly, these factors embody other credit card features that individuals might pay attention to, for example, differences in credit limits, debt differentials, or logos, to mention just a few. It is worth noting that, in our specification, we do not include variables affecting the level of total purchases, for instance, time indicators, as they do not influence our dependent variable.

Given the large heterogeneity in our data, it seems reasonable to believe that each individual has a different intercept. In our analysis, we consider these intercepts as random and assume that the error term in equation (1) consists of two components, an individual-specific term, which does not vary over time, and a remainder component, which we assume to be uncorrelated over time. We assume this random-effects structure to exploit the information from the within and between dimensions of the data. This is appropriate as several individuals in our data show little variation in interest rate rankings over time. Moreover, if individuals respond to cues, it is likely that changes in the interest rate differential from one period to the other have different effects than changes from one person to another, suggesting that the fixed effect estimator may not be appropriate here. A remaining issue with using a random effects model is that the individual-specific error may contain elements that are correlated with the regressors. Following Chamberlain (1984), we address this concern and allow for this type of correlation by assuming that the individual-specific error is a linear function of the time-averages of the time-varying covariates. Under this assumption, we carry out the estimation and include the means of all time-varying covariates as additional regressors into the random effects model. Adding these terms as controls allows us not only to identify the influence of changes in the regressors over time, but also to quantify the differences due to individuals with different time-averages.

Table 2 presents the results. Each column shows a separate regression. In all but the second column, we use OLS with random effects and clustered standard errors within consumers. In all models, the Breusch-Pagan test rejects the null of a single intercept. Similarly, the Hausman test shows no significant difference between the fixed and random effect coefficients for any of these specifications. Column (1) reports the base regression. The fraction of purchases allocated to a credit card increases by 0.06 when this card becomes less expensive than the other one. Nonetheless, the sample mean of the fraction of purchases assigned to card 1 when this card is cheaper all the time is still 0.52. The Wald test rejects the hypothesis that, on average, individuals allocate all their purchases to the card with a lower interest rate. Not surprisingly, the interest rate differential explains only a tiny part of the variation in the allocation of purchases.

To check the robustness of our estimates, we consider a few variations on the previous regression. In Column (2), we use a Tobit regression to account for the fact that the dependent variable is censored at zero and one. The conditional expectation of the fraction of purchases allocated to card 1 when it has always a lower interest rate is 0.53.[11] Despite the double censoring, the estimates of this model correspond closely to those of the linear model, suggesting that reliance on the linear model is reasonable. Thus, for simplicity, we stick to the linear specification. Column (3) examines whether results depend on the gap between interest rates. We re-estimate our baseline model and include interactions of the covariate with an indicator equal to one if the interest rate differential is larger than the median (1.1% per month). The fraction of purchases assigned to card 1 when it has a lower interest rate and the gap is large is 0.52, leading us to reject the null hypothesis again. In addition, this result indicates that the possible selection bias associated with different individuals picking different contracts (i.e. individuals who picked two cheap cards, two expensive card, or one cheap and one expensive card) is not quantitatively important. Regression (4) is a modification of regression (1), in which a continuous variable is used instead of an indicator variable. The effects are qualitatively similar, indicating that results are not sensitive to what measures are actually used in the regression. While not shown, we re-run equation (1) and include bank dummies for card 1. Although the R^2 is higher, estimates are similar to the base case.

A possible objection to these results is that they are influenced by consumers who pay no attention to interest rates because of the low stakes involved in the allocation decision. We address this concern in two ways. In column (5), we re-estimate our basic model and add interactions of the main covariate with an indicator coded one if the total outstanding debt in both cards during that period is larger than the median ($20,000 pesos). There are no significant changes in the estimated effects. Next, we re-estimate our model and include interactions of the main covariate with an indicator equal to one if the total purchases during that period are in the highest quartile (monthly purchases above $3,840 pesos). This regression is particularly interesting as consumers tend to evaluate more carefully decisions involving significant amounts of money. If interest rates are

[11] In the Tobit model, the conditional expectation of the dependent variable is a nonlinear function of the covariates and the estimated parameters [See Greene (2006)]

important for consumers, it is likely that an effect will show up in these cases. Column (6) displays the estimated coefficients. The general pattern is similar to the previous regressions. When making large purchases, the average individual assigns 52% of them to the card with the lower interest rate.[12] Overall, the estimates in all specifications support the claim that consumers do not go by interest rates when deciding which card to purchase with.

4.3 Allocation of payments

In this section we study whether consumers minimize interests cost and allocate most of their payments to the more expensive card. As before, the analysis presents two challenges. First, since most consumers make at least the monthly minimum payment to avoid late payment fees, looking at the amount allocated to pay off a given card could be misleading, as it could simply reflect the constraints imposed by the monthly minimum payment instead of the choices made by the cardholders. This problem is more pronounced if individuals barely pay the minimum. We address this concern and focus predominantly on cases in which individuals make the minimum payment due on both cards and repay more than the minimum on at least one of them (72% of observations). Second, since interest charges are not incurred if consumers pay off the complete balance by the end of the billing term, we only consider those periods in which individuals had outstanding debt on their two cards, as only for these months the allocation of payments is relevant for minimizing the interest accrued. This sample includes 53,509 observations.

Our focus is on the relationship between the fraction of 'payments above the minimum' allocated to a random card and the interest rate differential. We define this fraction as the payment made to card 1 (we keep the same label as in the previous section), minus the minimum payment on that card, divided by the total payments made during that period minus the sum of the minimum payments on the two cards for that month. Figure 5 displays the nonparametric kernel regressions of the fraction of payments above the minimum allocated to card 1 on the interest rate gap, defined as (r_1 - r_2). The figure shows variation throughout the range of the independent variable. Nonetheless, the fraction of payments above the minimum allocated to card 1 moves by less than 0.25 across the distribution of the interest rate gap. [13] Next, we use linear regressions to allow for heterogeneous effects. As before, we model the fraction of payments above the minimum due allocated to card 1,

[12] This result also holds when we restrict the sample to observations in which the interest rate gap is larger than the median (not reported).

[13] The histogram of the fraction of 'payments above the minimum' allocated to the more expensive card is mostly uniform with two spikes at zero and one (not shown). The spikes stand for observations in which consumers only made the minimum payment due on one card and paid above the minimum on the other one. Only 24% of consumers allocates most of their payments to the more expensive card. The remaining 76% would save money by following a simple rule, that is, making the minimum payment on the cheaper card and allocating most repayments to the more expensive card. Although the spikes disappear, the general pattern of such figure remains when we plot the average at the individual level.

for a given consumer and time period, as a linear function of a time-varying indicator variable of the interest rate differential. Our basic specification is:

$$SH\ PAY\ 1_{it} = \beta_0 + \beta_1 CHEAP1_{it} + \epsilon_{it} \tag{2}$$

The dependent variable $SH\ PAY\ 1_{it}$ denotes the fraction of payments above the minimum due that are allocated to card 1 by consumer i at time t. The intuition behind $CHEAP1_{it}$ and ϵ_{it} is the same as in the previous sub-section. Under the null hypothesis, individuals minimize costs, make the minimum payment to both cards, and allocate all their remaining funds to the card with a higher interest rate, so $\beta_0+\beta_1$ should be equal to zero. We test this hypothesis and estimate the equation by OLS with random effects and clustered standard errors within individuals.

Table 3 summarizes the results for an array of specifications -each column is a separate regression. Column (1) reports the base regression. The sample mean of the fraction of payments over the minimum allocated to the expensive card equals 0.49. The Wald test for the hypothesis that $\beta_0+\beta_1$ are equal to zero is rejected at the 1 percent significance level. Column (2) reports the Tobit regression, to account for the fact that the dependent variable is censored at zero and one. The estimates of this model correspond closely to those in the OLS estimation. Column (3) studies whether the value of the difference in interest rates has an effect on the allocation of payments. This regression presents the estimates when we include an interaction term of our main covariate and an indicator variable that equals one if the interest rate gap is above the median (1.1% per month). Results do not change at all. Regression (4) is similar to regression (1), but a continuous variable is used instead of an indicator variable. The effects are qualitatively similar to the corresponding estimates in column (1), suggesting that our results do not depend on the definition of our covariates.

In columns (5) and (6), we probe the robustness of our estimates to cases in which mistakes are more costly. Specification (5) includes the interaction of $CHEAP1_{it}$ with an indicator variable for large outstanding debt. In our sample, the total debt of the average individual classifies as large 53% of the time. For such individual, the fraction of payments above the minimum allocated to the more expensive card when the amount of debt outstanding is large and the gap is above the median equals 0.50. Column (6) focuses on large payments. Again, the assumption behind this regression is that individuals tend to pay more attention to decisions involving significant amounts of money. We re-run our base specification and include interactions of our variables with an indicator coded one if the total payments over the minimum due are in the highest quartile of the distribution, i.e. the sum of the payments above the minimum due is larger than $8,400 pesos. Estimates are similar to the base case. When the average individual pays a large amount of money on her credit cards (which happens, on average, 23% of the time), she allocates only half of that amount (0.51) to pay back the more expensive card.

Another concern could be that, as a result of the teaser rate offers, consumers are not able to track the identity of the card with the lower interest rate. To investigate this possibility, regression (7) re-estimates our base model and includes interaction terms of the independent variable with an indicator equal to one if the cards of a given consumer never changed ranks; that is, the indicator is coded one for individuals whose cheaper card had always a lower interest rate than their more expensive one. Little changes in this regression. Another possibility is that cardholders are restricted by the closing dates of their cards. For instance, consumers may pay less to the more expensive card because the due date on this card is before payday, while the deadline on the other one is after that. This argument could explain why interest rates play no role in the allocation of payments. We run the basic model and incorporate interaction terms with an indicator equal to one for consumers whose credit cards have a closing date differential of eight days or less. The estimated parameters are shown in column (8). The coefficients on the regression remain unchanged. In regression (9) we explore the correlation between payments and purchases and include an interaction term of the independent variable with an indicator equal to one if the purchases made with card 1 were larger than the purchases made with card 2 during the previous period. We expect that individuals who are "savvy" and make most of their purchases with the cheaper card minimize interest costs and allocate most of their payments to the more expensive card. However, the estimates show that the effect for this group is not significantly different.

Finally, we consider situations in which cardholders make at least the minimum payment on just one card and study which card consumers choose to make the minimum payment on (given that only one minimum payment was actually made). Specifically, we estimate the effect of interest rates on the probability of making the minimum payment due. The basic idea is simple. If individuals missed the minimum payment despite it was feasible for them to meet both of their minimum payment obligations, it must be because they lost the billing statement (or some other unlucky situation), missed the due date (or forgot about it), or preferred to allocate the funds to repay some other card. The first case is simply randomness. The second and third cases however, imply that individuals pay more attention to cards with certain characteristics or alternatively, prefer to repay some of them first. In either case, the features of those cards contain information about which factors matter for consumers when allocating their monthly payments.

Column (10) reports the marginal effects of the probit regressions using the same covariate described in the preceding analysis. Our sample consists of consumers who missed the minimum payment on one card, despite it was feasible for them to cover it up during that period (as the total amount paid to both cards is larger than the sum of the minimum payments). Our dependent variable is an indicator that equals one if the individual made the minimum payment to card 1 in that month. The unit of analysis is a consumer-month. The estimates are quite consistent with our previous results. As before, interest rates play no role at all in the decision to make the minimum payment.

In summary, the estimates for all specifications in this section support the claim that individuals do not allocate their monthly payments based on the interest rates. It is worth noting that, contrary to the case of purchases, where somebody could argue that allocations are influenced by unobservables not captured by our initial matching (for instance, discounts offered by banks when buying from certain retailers), the decision to pay one card or another is not affected by the discounts offered by banks to encourage credit card usage.

5 Possible explanations

The findings in the previous section seem to support the idea that individuals do not follow interest rates. Assuming (1) homogeneous goods, (2) no switching costs, and (3) perfect information, these results are difficult to reconcile with a cost minimizing model. Yet, there are two possible objections to this conclusion. The first one is that the previous assumptions may not be true in our setting. Although the first and second assumptions seem quite plausible considering the way we constructed the sample, the last assumption is more questionable. Namely, consumers may ignore the interest rate of their cards. In the first subsection, we investigate how likely this supposition is. The second objection is that consumers do not care about interest rate differentials due to the low stakes involved and therefore, find easier to stick to some other rules. In the second subsection, we address this concern and calculate the monetary cost of mistakes. Finally, in the third subsection, we argue that our results are more consistent with hypotheses related to mental accounting and financial unsophistication and provide further evidence supporting this statement.

5.1 Do consumers know the interest rates of their credit cards?

Consumers may exhibit the previous behavior because they do not know the interest rate of their credit cards. Testing for this hypothesis however, is complicated. An ideal test is to survey cardholders in our sample and ask them about the interest rate of their credit cards. This procedure however, is expensive and unfeasible due to confidentiality constraints. As a proxy, we elicit the proportion of consumers who know the interest rates of their cards using a survey of 200 randomly chosen respondents interviewed in public places in Mexico City with at least two credit cards. Among the questions of the survey, we asked individuals about the interest rates of their credit cards as well as their relative rank (in terms of interest rates) during the preceding month. The proportion of respondents who identified the exact rates (or a close rounded number) was 53%, while the percentage who knew how cards were ranked reached 67%. These figures give us a rough sense of consumers' acquaintance with interest rates. Nonetheless, as we cannot observe the actual allocations of respondents, these numbers should be interpreted as a purely descriptive exercise.

An alternative test is to analyze the behavior of consumers who recently received a reminder letter from banks concerning the interest rate of their credit cards. The basic idea here is that consumers are more likely to recognize interest rates after receiving this notice. Unfortunately for us, banks hardly ever send these announcements without other promotions. What banks do send, however, are temporary reductions in the interest rate or teaser rate offers (TROs). These offers are widespread in Mexico. In 2005, for example, the average credit card had an active TRO for approximately 2.1 months. TROs usually represent an important reduction in the interest rates. In our sample, the average interest rate during the teaser period was 1.3%. Banks send these offers by mail to selected customers. These letters include the description of the TROs as well as the otherwise applicable interest rates. We use these offers as proxies for interest rate reminders.

Our exercise has two steps. In the first one, we identify cardholders who received at least one TRO, and analyze how they react to it. This is to assess whether consumers are aware of the changes in the interest rates and consequently, of the rates that regularly apply. Although it is not the main focus of this exercise, this regression allows us to asses whether consumers substitute debt between cards. Next, we reproduce our previous analysis but restrict our sample to observations after the teaser period. This approach is the best available, given our data. Still, two possible concerns remain. First, limiting the analysis to individuals who received a teaser rate offer can introduce selection bias if the selection rule of the bank is correlated with allocations. For instance, banks could target TROs to consumers borrowing on other cards or who are not using the cards issued by them. Fortunately, this is not very likely here, as banks in Mexico cannot observe activity in other credit cards. Indeed, we find no statistical difference in the average allocation of payments and purchases between the full sample and the reduced sample, suggesting that the selection is unbiased for our purposes. A second concern is that even if consumers realize one of the interest rates after receiving a TRO, this does not imply they will be aware of the other one. To address this, we re-estimate our specification using data for consumers who got teaser rate offers on both cards, after they received the last one.

We begin by investigating whether consumers react to temporary changes in the interest rate. Our empirical strategy consists in regressing changes in debt, purchases or payments in a given card, on indicators of teaser rate offers, their forwards and lags, as well as a full set of controls. The base specification mirrors that used in previous research, e.g. Gross and Souleles (2002), except that these authors utilize changes in the interest rates as regressors. In contrast, we employ indicator variables of the TROs as covariates. This is to acknowledge that TROs are temporary. In Mexico, most of these offers last one or three months (82% in our sample). We concentrate on these two types. To properly account for these differences in duration, we introduce separate covariates for each kind of TRO. The unit of analysis is a consumer-month. The estimating equation of primary interest is:

$$\Delta Y_{it} = \alpha + \sum_{j=1}^{n} \beta_j TR1M1_{it+j} + \gamma_1 TR1M1_{it} + \sum_{k=1}^{n} \delta_k TR1M1_{it-k} + \sum_{l=1}^{n} \rho_l TR3M1_{it+l}$$
$$+ \theta_1 TR3M1_{it} + \theta_2 TR3M2_{it} + \theta_3 TR3M3_{it} + \sum_{m=1}^{n} \phi_m TR3M3_{it-m} + \lambda' X + \epsilon_{it} \quad (3)$$

where ΔY_{it} is either the change in actual debt, purchases or payments in card 1 or the change in debt on card 2. The term $TR1M1_{it}$ is an indicator equal to one if individual i received a one month TRO on card 1 during period t. The covariates $TR3M1_{it}$, $TR3M2_{it}$ and $TR3M3_{it}$ are indicator variables coded one for the first, second, or third month of the three month TRO received on card 1, by individual i at time t. The covariates in X are additional controls that include binary variables for other changes in interest rates, indicators for changes in credit limit and its lags and time dummies.[14] In addition, we include individual fixed effects.[15] Finally, since our panel is unbalanced and there is only one year of information available for one bank, we include two forwards and three lags for each one of the offers.

We estimate equation 3 by OLS, with clustered standard errors within accounts. Table 4 presents the results. Each column corresponds to a different dependent variable. Column (1) reports the estimates using the change in debt in card 1 as the outcome of interest. Consumers respond substantially and immediately to the two kinds of offers. In the case of the 1-month TROs, the average debt increases by \$346 pesos at the time of the announcement plus \$1,909 pesos during the first month of the offer. The corresponding figures for the 3-month TROs are \$1,023 and \$2,209. In addition, average debt either decreases or remains invariant during the month following the teaser period. Columns (2) and (3) present the estimates using the change in purchases and payments in card 1 as dependent variables, respectively. The regressions confirm the prior patterns. Purchases are estimated to increase at the beginning of the teaser period and decrease afterward. Conversely, payments decrease at the beginning but increase significantly when the TRO expires. These results confirm our supposition that consumers are aware of changes in the interest rates and suggest individuals must be familiarized with the rates typically charged.

Column (4) reports an additional regression using the change in debt in card 2 as dependent

[14]In unreported regressions, we include covariates to account for possible selection rule of the banks. These regressors include previous utilization rates and previous delinquencies. We include these covariates to control for an omitted variable bias that could arise due to the nonrandom credit supply decision of banks. For example, if banks extend teaser rate offers to consumers who are not using their card and who have a low elasticity with respect to temporary changes in the interest rate, then our estimates for the population will be biased downwards (although it would be correct for the treated population). Conversely, if banks target consumers who are increasing their debt, it is possible that we incorrectly attribute differences in debt (or the dependent variable) to TROs. Estimates are similar to the base case

[15]With fixed effects, the coefficients of interest are identified by departures from the average change in debt.

variable. Although not directly related to the goal of this section, we run this regression to investigate the sensitivity of debt in one card to changes in the interest rate of the other one. These findings are of potential interest. The estimate for the effect of teaser rate offers on debt in other cards is insignificantly different from zero.[16] This supports our previous findings and provides evidence against there being substitution of debt between credit cards. In fact, these estimates suggest that consumers pay attention to time-varying changes but not to time-invariant differences in interest rates pointing to the interpretation that consumers may categorize their expenditures in mental accounts or simply lack the financial sophistication to manage their personal finances. We explore this topic further in section 5.3.

So far, we have shown that consumers receiving TROs are actually aware of the applicable interest rates. We now investigate how these consumers allocate their monthly payments and purchases after receiving these "announcements". We re-estimate equations 1 and 2 restricting the sample to periods following the expiration of the first teaser offer (in case consumers received any). Columns (7) in Table 2 displays the results. In general, there does not appear to be an important difference with the base estimation.

Limiting the sample to observations following a teaser period slightly reduces the magnitude of the effect of previous purchases, but in general, results are consistent with the estimates obtained previously. Using the information on TROs, it is possible to construct a smaller sample of consumers that received at least one TRO on both of their cards. If consumers are aware of the interest rates after receiving a TRO, consumers in this sample should be presumably familiar with the terms of their two cards. The analysis for those observations is displayed in the last column of Table 2. Results do not change if we restrict the sample this way. In an unreported regression, we re-estimate these equations excluding observations for which the interest rate differential is smaller than the median. There are no significant changes in the estimated effects. The estimations for the payments specification using both samples are presented in columns (11) and (12) of Table 3. Limiting the sample to observations following a teaser period generates similar estimates to those obtained previously. Overall, these results seem to reject the hypothesis that individuals abide by interest rates after introducing reminders about them. While TRO may not be the best proxy for interest rate reminders, our findings cast doubts on the hypothesis that results in section 4 are mainly driven by unawareness about interest rates.

[16]It could be the case that the interest rate in card 2 is still lower than the one in card 1, despite the fact that card 1 has an outstanding TRO. In fact, this happens in 11% of the offers. Nonetheless, the results when we exclude these observations are quite similar to those using the full sample

5.2 Small stakes

Section 4 shows that consumers do not allocate their monthly purchases and payments based on the interest rate differential between their credit cards. Yet, this is not inconsistent with a cost minimization model, unless these deviations are systematic, costly and prevalent. We already showed that deviations are systematic. In this section, we focus on the last two of these features and calculate the monetary costs of mistakes. We set up a static model of optimal intra-temporal allocation of credit card payments and purchases, conditional on the purchases undertaken with both cards and the total amount repaid to them.[17] We use this model as a benchmark against which the actual financing costs can be compared. The main advantage of such a model is that it allows us to formally introduce into the decision problem contractual features such as the minimum payment due, the credit limit and fees.

The model is a simple period-by-period cost minimization problem. A formal description is presented in the Appendix. We solve the model and derive the optimal purchases and payments for each card, period and consumer. Next, we use these allocations to estimate the optimal financing costs for each individual. With this at hand, we compare these costs with the actual financing costs faced by each cardholder and calculate the monetary costs of mistakes. We refer to these extra expenditures as misallocation costs.

We measure financing costs as the sum of interest costs and fees. Although our main interest is in interest cost, we incorporate the cost of fees in our analysis. This is because fees could also be a consequence of allocation mistakes. For example, overlimit fees could be avoided by assigning part of the spending to another card. However, since the financing cost is additive, we are able to disaggregate the misallocation cost into extra interest and extra fees.[18] Table 5 presents the percentage of observations in which individuals incur in misallocation costs, disaggregated by its source. The table shows that mistakes are very frequent. In 84% of the cases, consumers pay extra costs due to the wrong allocation of payments or purchases. Most of these mistakes involve extra interest costs rather than additional fees, suggesting interest mistakes are likely to be persistent. To explore this hypothesis, Figure 6 presents the distribution of consumers by the frequency they pay extra interest and extra fees. Here, the unit of analysis is a cardholder. Consumers seem to incur in extra fees rarely. About 27% never pays extra fees and about 65% do not incur in these fees at least 30% of the time. On the contrary, interest mistakes seem to be quite persistent. About 54% pays extra interest at least 70% of the time and 75% incurs in these mistakes at least 50% of the time. This persistence suggests consumers have an objective different from minimizing financing

[17] As already mentioned, conceptually our approach is equivalent to separating the consumer problem into two stages. In the first stage, individuals decide on the total credit card spending as well as the total amount to be repaid at the end of the month. In the second stage, consumers allocate these purchases and payments between their cards to minimize the interest costs and fees.

[18] We define extra interest (fees) as the actual interest (fees) cost minus the optimal interest (fees) cost.

costs.

The previous table is silent about the magnitude of these mistakes. To address this, Figure 7 displays the distribution of the average misallocation cost. This cost is calculated for each consumer by taking the average of the monthly misallocation costs along time. The unit of analysis here is a cardholder. We focus on individuals rather than consumer-months to get a realistic measure of the cost faced by each person. To be more precise, using averages allows us to put a higher weight on costly but infrequent mistakes, such as those associated with large purchases, large payments and fees. The distribution in Figure 7 is skewed to the left. On average, consumers leave \$82 pesos on the table every month. Out of this, \$45 pesos correspond to extra interest cost. These results are not driven by outliers. The median of the distribution is \$57 pesos, which amounts to 0.006% of the median monthly income. These costs could be due however, to differences in the amount borrowed by distinct consumers. To examine the relative importance of these mistakes, we standardize the misallocation costs by dividing them by the average financing expenditure of each individual (interest accrued plus fees). The distribution of this measure is displayed in Figure 8. As before, the histogram is skewed to the left. The average individual leaves on the table a sum which amounts to 16% of her financing cost. The median individual leaves a sum which amounts to 10%. For the 90th percentile, more than 30% of the financing costs is due to misallocation mistakes.

A troubling objection to the previous results is the possibility that misallocation costs are highly influenced by overlimit and late payment fees. To further assess the relative importance of interest against fees, we combine Table 5 and Figure 7 and calculate the misallocation cost by source and deciles. Figure 9 shows the average misallocation cost within each decile, decomposed in extra interest costs and extra fees. Such desegregation allows us to assess the relative importance of both sources across the entire distribution of misallocation costs. With the exception of the top decile, interest and fees costs are spread fairly equally across the distribution. For the top decile, interest costs are significantly higher and add up to \$2,228 pesos. These results suggest that interest costs do constitute an important fraction of the money consumers leave on the table.

The cost of mistakes is not trivial, particularly when we compare it with other figures in the literature. Zinman (2007) uses the Survey of Consumer Finances for the United States to study foregone arbitrage between credit card debt and demand deposit accounts. He finds that “fewer than 10% of credit card holders could save as much as \$10 per month by managing their liquidity more aggressively”. We look at missed arbitrage between much closer substitutes and with more accurate data. In spite of the presumably higher substitutability, we find larger sums left on the table. What is more, since mistakes are persistent, costs accumulate over time. These conclusions however, should be interpreted with caution as they may be influenced by factors such as interest rate differentials, fees and outstanding balances.

In summary, mistakes are very frequent. The cost of these mistakes is not trivial, but it is not so large either, and it is spread almost equally between extra interest and extra fees. What is more interesting is the fact that mistakes involving extra interests are quite persistent over time, suggesting consumers do not pay attention to interest rate differentials. Still, it is possible to argue that potential savings are not large enough to overcome the costs involved in performing the required calculations. We cannot prove this statement wrong. However, this explanation cannot fully rationalize the patterns documented in the next section. More specifically, if consumers do not really mind about which card to purchase with or to pay off, we should observe individuals choosing randomly between their two credit cards. Although there is substantial heterogeneity in consumers' choices, this implication contradicts our findings in the next section. In addition, the fact that consumers respond to temporary changes in the interest rates (as documented in the previous section) suggest they *actually care* about small potential savings.

5.3 Heuristics and financial unsophistication

A final explanation to our results is related to mental accounting and financial unsophistication (or lack of financial literacy). Mental accounting refers to the cognitive operations used by individuals to organize, evaluate and keep track of their financial activities [Thaler (1985), Thaler (1999)]. Financial unsophistication denotes the inability of consumers to understand how credit cards work or how interest is accrued.[19] These explanations seem plausible as credit card features are complex. Testing for these stories, however, is complicated. Ideally, we would like to survey cardholders in our sample and study the reasons for their behavior, yet this is not feasible due to the confidential nature of the data. As an alternative, we investigate whether individuals adhere to heuristics and analyze how individuals allocate their monthly purchases and payments when certain features of the environment, such as differences in balances or in previous purchases, change.

Looking for heuristics and rules of thumb in the data is not easy. In order to guide our search, we designed and applied a small survey in several public places of Mexico City during July, 2006.[20] We interviewed people until we gathered 200 individuals who had at least two comparable credit cards, carried interest paying debt on at least one of them and whose average utilization rate in the month before the interview was below 0.6.[21] We asked for respondents with low utilization rates to avoid situations in which individuals make their decisions purely on the basis of credit limit constraints. We asked respondents about the identity of the credit cards in which they carried

[19] In a recent paper, Hastings and Tejeda-Ashton (2008) present evidence supporting this story and show that financially illiterate workers in Mexico pay much more attention to fees when fees are presented in pesos instead of APRs, when choosing between retirement investment funds (Afores).

[20] Due to budgetary and time constraints, the sample size is small and the range of questions is limited. Although a larger and more rigorous survey would be required to properly evaluate the usage of heuristics, our survey can still provide a guide to behavioral patterns to look for in the data.

[21] To calculate the total utilization rate, we asked individuals about the sum of interest paying debt and credit limit on their cards.

outstanding balances on and then, inquired about the interest rate on them. We asked respondents who carried debt on a more expensive card about the reasons to do so. We showed them several answers including an alternative for respondents to talk about not listed options. It is worth noting that we asked about the reasons to borrow on a certain card, not about the reasons to use it. Therefore, our answers combine motives such as preferences, usage patterns as well expectations at the time of spending.

Of the 200 consumers, 67% (134) claimed they knew for certain which card was more expensive (although many of them did not know the exact interest rates). Out of this, 29 respondents sustained the interest rate in all their credit cards was equal. From the 105 individuals remaining, 76% (80) holds interest paying debt on the more expensive card; their reasons were as follows: a) 3% (3) answered that they thought they would be able to pay the expensive card but in the end, they could not; b) 3% (3) said they were concerned about one of their cards to be stolen or incorrectly rejected at a store, so having available limit in both cards was an insurance against these kinds of events; c) 43% (34) replied that they did not like to have "too much" debt on a single credit card, even though they had available revolving credit on the cheaper one; d) 8% (6) said that they preferred to have available credit limit on both cards; e) 31% (25) answered they used different cards for different purposes and; f) 11% (9) said that they had not thought about it.

Answers a) and b) are some of the standard arguments found in the literature, however only 6% of respondents chose them. Interestingly, only 3% of the surveyed consumers selected option a), which features prominently in Ausubel (1991). This provides support to our assumption that dynamic considerations play a smaller role in the allocation problem. Options c) and d) are consistent with the idea that consumers aim to keep a constant utilization rate in their credit cards. This motive is consistent with the results obtained by Gross and Souleles (2002). Finally, a considerable amount of consumers selected option e), which is closely tied to categorization of purchases and mental accounting [Thaler (1985), Thaler (1999), Prelec and Loewenstein (1998) and Ranyard et al. (2006)]. Taken together, these answers support the idea that consumers follow rules other than interest rates when deciding which credit card to borrow on. In what follows, we use these answers to guide our data analysis.

In this paper, we do not attempt to unravel all the heuristics in the data. Instead, our goal is to demonstrate that consumers tend to pay attention to features of the environment other than interest rates, which would suggest that individuals are prone to guide their choices by rules of thumb rather than by cost minimization motives. Based on the responses to our survey, we examine three heuristics, which in one way or another, could be rationalized by the theory of mental accounting. First, we test the hypothesis that individuals are reluctant to accumulate too much debt into a single card, so they tend to purchase with the card which has less outstanding balances and pay back the card with more debt. We refer to this behavior as the "**debt equalization**" heuristic.

Second, we test whether individuals tend to consolidate debt into a single card. This is consistent with the responses c) and d) in our survey and seems reasonable considering the 'advice' provided by some popular financial planners. For example, Dave Ramsey's website[22] proposes the following strategy to get out of debt fast:

> *"**Myth**: I should pay off the debt with the highest interest rate first to get out of debt quickly.*
>
> ***Truth**: You should pay off the smallest debt first to create the greatest momentum in your debt snowball."*

This strategy, commonly known as "**snowball**", has gained popularity and is considered the primary debt-reduction method by many financial experts in the United States. Finally, we test the idea of "**inertia**" (inclination to repeat-purchase a brand over successive purchases), which is closely tied to the idea of categorization and is often mentioned in the credit card industry (the so-called "front of the wallet" behavior).[23]

5.3.1 Debt equalization effect

We begin by analyzing how differences in the outstanding balances on each credit card affect the allocations of purchases. We estimate the following regression:

$$SH\ X1_{it} = \gamma_0 + \gamma_1 CHEAP1_{it} + \gamma_2 DEBT1_{it} + \gamma_3 LIMIT1_{it} + \eta_{it} \tag{4}$$

where $SH\ X1_{it}$ and $CHEAP1_{it}$ are defined as in the preceding section. The term $DEBT1_{it}$ is an indicator coded one if the outstanding balances on card 1 at the beginning of period t are greater than those on the other card, and $LIMIT1_{it}$ is an indicator coded one if card 1 has more available limit (limit minus outstanding balances) than card 2. As before, we include the means of all time all time-varying covariates as additional regressors and carry out the estimation by OLS with random effects and clustered errors within individuals.

The coefficient on $L1MIT1_{it}$ measures how credit limit constraints affect the allocation of purchases. Although limit constraints do not bind in our sample, we include this regressor to account for consumers who may worry about getting close to their credit limits. If these constraints are important, this coefficient should be positive. The coefficient on $DEBT1_{it}$ measures how outstanding balances impact the allocation of purchases. One might expect this coefficient to be negative if consumers tend to equalize debt across cards (debt equalization heuristic).

[22] Source: http://www.daveramsey.com/etc/cms/index.cfm?intContentID=4055

[23] Ideally, we would like to also test the hypothesis of categorization, in which individuals tend to put their purchases into categories or mental accounts, and associate a different card to each one of them [Thaler (1999)]. However, since the data we use do not disaggregate into expenditure categories, this is not empirically feasible.

The first column of Table 6 reports OLS estimates for the specification shown in equation 4. The first point to note is that the difference in credit limits plays a minor role in the allocation of purchases when the limit constraints do not bind. The identity of the card which has more debt, however, seems to play a more important role. The estimated coefficient on the debt differential dummy implies that consumers tend to allocate 0.09 points less in purchases with the card with more outstanding balances, regardless of the differences in credit limits or interest rates. For instance, consumers who typically borrow *less* on card 1, i.e. those for which $\overline{DEBT1}_i < 0.5$, tend to make only 0.41 of their purchases with this card when such card has *more* debt and is less expensive. Since total purchases in our final sample are roughly \$4,000 pesos, this number implies that consumers tend to purchase \$800 pesos less with the card which has more debt, even when this card has a lower interest rate. Column (2) displays the estimates when we restrict the sample and include only observations in which total purchases are in the highest quartile. The results are similar, although the estimated coefficient on the debt differential indicator is slightly larger. These figures suggest that the effect of the so-called debt equalization heuristic is more important than the one induced by the differences in interest rates.

What about the allocations of consumers who tend to borrow more on card 1 when such card has more debt? Unfortunately, we cannot say much about them. While it is possible that some of these individuals show a similar behavior to the one previously documented, in practice it is not possible to distinguish such effect because it is likely to be confounded with other factors affecting both, the sum purchased with this card and the amount of debt kept on it. More precisely, the fact that card 1 has often more debt might be a proxy for unobserved preferences. Therefore, while some consumers might tend to equalize debt and purchase less with the card which has higher outstanding balances, others might purchase more with such card simply because it represents the card they frequently use and not pay in full.

In columns (3) and (4), we re-estimate equation 4 using as dependent variable the fraction of payments above the minimum due that are allocated to card 1. We omit the covariate of differences in the available credit limit. The estimates in column (3) indicate that, on average, individuals pay 0.05 points *more* to the card with *more* debt. This effect increases to 12 percentage points when individuals make large payments. As previously noted, this is particularly interesting, as people tend to pay more attention to decisions involving larger amounts of money. When making large payments our estimates imply that, once the minimum payment due is covered, individuals pay approximately \$2,000 pesos more to the card with more debt. To get a sense of the size of this effect, under the same assumptions, the difference between the payments made to the more expensive card and those made to the less expensive one is only \$600 pesos. These findings all indicate that consumers seem to pay more attention to heuristics than to prices. There are several explanations for this. If individuals categorize their expenditures into mental accounts, they might prefer to repay cards that have been spent up to its limits and 'mentally' reduce the stress that a single large debt put on them. Alternatively, if consumers are unsophisticated, they may pay more attention to the *interest accrued*, rather than the *interest rate*. What is more, it is quite possible that the effect of the "debt equalization" heuristic might be stronger that the previous statistics show because of two contradictory trends - those who tend to purchase less and pay more to the card with larger balances (debt equalization effect) as opposed to those who try to pay off the lower debt card and consolidate their debt over, say, a period of 3 or 6 months (snowball effect). Unfortunately, it is quite difficult to distinguish these two effects in the data. In the next subsection, we present evidence that suggest that at least some consumers abide to the second heuristic.

5.3.2 Snowball effect

Another possibility is that some people look for paying off the card with the smallest balance first. This kind of behavior would be justified if consumers value the psychological benefits of quickly crossing off the debt on as many cards as possible, which would make sense if individuals categorize and treat each credit card as a separate mental account. Looking for evidence of this behavior, however, is not easy. When consumers carry debt on their two cards, regression analysis is not useful to identify the percentage of consumers making larger payments on the card with the lowest balance. This is due to the presence of two opposing trends - people making larger payments on the credit card with the lowest balance, and consumers who tend to equalize debt and pay off the largest debt first. Here we adopt an alternative approach. First, we estimate the fraction of consumers paying off the card with the smallest balance in two scenarios: (1) when, after subtracting the minimum payment on the card with the largest debt, the payment amount remaining is enough to pay off the card with the smallest balance, and (2) when, after subtracting the minimum payment on either card, the payment amount remaining is enough to pay off the other card. If consumers follow the snowball method, it is very likely they will pay off the card with the lowest balance when they have enough money (as measured by the total payment amount) to do so. Next, we examine the allocations leading consumers to pay off the card with the lowest balance.

We begin by estimating the fraction of people paying off the card with the lowest balance via OLS regressions.[24] Since all the covariates are binaries, our analysis is a comparison of means. The

[24] We estimate the following specification:

$$PAY\ OFF1_{it} = \delta_0 + \delta_1 CHEAP1_{it} + \delta_2 ABLE\ TO\ PAY\ OFF1_{it} + \delta_3 INT_{it} + \eta_{it} \tag{5}$$

where $PAY\ OFF1_{it}$ is an indicator coded one if consumer i paid off just card 1 in period t, and had outstanding balances on that card during the previous period. The variable $ABLE\ TO\ PAY\ OFF1_{it}$ is an indicator measuring whether individual i is able to pay off the card 1 in period t. We define this variable in two ways, according to the two cases outlined above. In the first case, this variable equals one if all the conditions on the top row of column (6) of Table 6 are satisfied. In the second case, the variable is codified using the conditions in column (7).

results are reported in columns (6) - (7) of Table 6. As reported in column (6), when individuals have the option of paying off only one card, 31% decide to do so, regardless of the difference in the interest rate on their cards. The equivalent figure when individuals are able to pay either card in full is 20% (Column (7)). Interestingly, the percentage of people paying off the card with the *largest* amount of debt equals 34% (not shown). Although these figures are correct, they do not prove that individuals follow the snowball method. Further analysis is needed to evaluate whether individuals were actually paying a larger amount to the card which was later paid off. In fact, this does not seem to be the case. Out of the 791 cases in which the card with the smallest balance was paid off (column (6)), only 277 (35%) allocated 75% or more of the payments over the minimum to the card with the lowest balances in the preceding period. The equivalent figure in column (7) is only 28% percent. Although these numbers are relatively small, they still show that, relative to the amount of savvy consumers, many still abide to heuristics. For instance, in the case portrayed in column (6), the percentage of people allocating 75% or more of the payments over the minimum to the more expensive card equals 20 percent. Nonetheless, our results suggest that, in spite of certain people following it, this method is not really popular in Mexico. This makes sense, as the market for financial advisers in Mexico is less developed. Moreover, since most of the monthly payments in our data are relatively small, consumers might get discouraged and abandon the strategy as they might not see quick results.

5.3.3 Inertia

So far, we have shown that some individuals use rules of thumb to manage their credit card debt before situations such as large outstanding balances on one card. A natural question arises concerning how individuals get into such situations in the first place. In this section, we explore this issue and provide evidence that individuals tend to stick to their initial choices and repeat purchases with the same card. This behavior is related to the concept of inertia in the marketing literature and is generally referred to as "front of the wallet" behavior in the credit card jargon.[25] There are at least two reasons why individuals might be inertial. First, since credit cards are homogeneous products, individuals might not have the necessity to seek for variety.[26] Second, individuals might use different cards for different purposes [for example, one for recurrent purchases (i.e. groceries) and another one for large buys (i.e. durables, traveling, etc.)]. In fact, this is an implication of the mental accounting theory. More precisely, an important component of the mental processes involves

The covariate $CHEAP1_{it}$ is defined as before. Finally, the term INT_{it} is the interaction between $CHEAP1_{it}$ and $ABLE\,TO\,PAY\,OFF1_{it}$.

[25]The concept of inertia is not new in financial markets. For instance, Madrian and Shea (2001) study the impact of automatic enrollment on 401(k) savings behavior and find that a substantial fraction of 401(k) participants hired under automatic enrollment retain both the default contribution rate and fund allocation.

[26]This is particularly true in our sample as we select only cards of the same 'type' or with the same rewards structure

the assignment of expenditures to specific accounts which, in our framework, are represented by credit cards. If individuals categorize their purchases, they are likely to stick to a certain card and repeat purchases. In this sense, inertial tendencies are consistent with the theories of categorization and mental accounting.

We first investigate whether making more purchases on a given card increases the percentage of purchases made with that card in the next period. We model the fraction of purchases allocated to card 1 as a linear function of indicator variables of interest rate differentials, differences in debt, differences in available limits, differences in the purchases made during the previous month, and differences in the purchases made during the first month with available information. Specifically, we estimate the following linear model:

$$SH\ X1_{it} = \beta_0+\beta_1 CHEAP1_{it}+\beta_2 DEBT1_{it}+\beta_3 LIMIT1_{it}+\beta_5 PURCH1_{it-1}+\beta_6 PURCH1_{i0}+\eta_{it} \tag{6}$$

where the covariates $SH\ X1_{it}$, $CHEAP1_{it}$, $DEBT1_{it}$, and $LIMIT1_{it}$ are defined as before, $PURCH1_{it-1}$ takes on the value of one if consumer i carried more purchases on card 1 at time $t-1$, and $PURCH1_{i0}$ equals one if consumer i made more purchases on card 1 during the first period with available data. We include the variable $PURCH1_{i0}$ to address the so-called initial conditions problem and distinguish econometrically between state dependence and unobserved heterogeneity [See Wooldridge (2002)].[27] We are interested in β_5. If individuals tend to repeat purchases once we control for unobserved characteristics, we would expect this coefficient to be positive. As before, we include the means of all time-varying covariates as additional variables and carry out the estimation by OLS with random effects and clustered errors within individuals.

Column (8) of Table 6 presents estimates of the parameters in equation 6. The estimate of β_5 suggests that consumers are prone to repeat purchases with the same credit card. After controlling for time-invariant unobserved heterogeneity, consumers purchase 20 percent more with the card which had larger purchases during the preceding period. This is four times the effect of interest rates. Interestingly, the fraction of purchases allocated to card 1 is 12 points larger if that card had larger purchases during the first period, suggesting that unobserved brand preferences play an important role in allocation. These results do not change when we use a Tobit specification (not shown). A possible objection is that these results are explained by small and recurrent purchases and automatic deductions. While not reported, we re-run regression 6 using a sub-sample of purchases in the highest quartile. Overall, the estimated coefficients do not change substantially, although the point estimate of the indicator for previous purchases drops to 0.14 when restricting the sample this way. Finally, although not shown either, we repeat the previous exercise for the

[27] Even though we are not including a lag of the dependent variable in the regression, the identification problem still exists because the term $PURCH1_{it-1}$ is a non-linear function of that regressor.

payments specification. Not surprisingly, we find no evidence that individuals pay larger amounts to the card which was paid relatively more during the former period.

Finding that individuals purchase constantly with the same card does not necessarily mean that they categorize purchases into mental accounts. As a simple test, we look at the behavior of people using and paying off at least one credit card. In particular, we study how prevalent consumers using one card for transacting and another one for borrowing are. Although not perfect, this exercise shows that at least some people follow one particular kind of categorization that is not explained by interest rates. Here we also include in the sample individuals not borrowing or not making purchases on any card. Of the 14,926 cases in which consumers paid off at least one card, 4,258 (29%) carried debt in one card and paid off the other one. Out of this figure, 48% paid off the card with the lowest interest rate. To show that this behavior is persistent, column (9) in Table 6 present the estimates of a linear probability model with random effects of the probability of paying off card 1 as a function of interest rate differentials, its own lag, and an indicator coded one if card 1 was paid off in the first period with available information.[28] The point estimate of the marginal effect indicates that paying card 1 in full in the previous period increases the probability of settling the debt again by 0.20. Taken together, these results provide statistical support to the theories of categorization and mental accounting.

Taken as a whole, mental accounting could explain why individuals do not pay attention to differences in interest rates, react to teaser rate offers, systematically use one card for recurrent purchases and another one for borrowing, or tend to equalize debt. Mental accounting, however, is not the only explanation. An alternative story is that individuals do not understand how credit cards work. This is plausible because credit card features are complex. For example, individuals may pay the credit card with more debt, with the 'wrong' idea that it is more expensive because it accrues more interest. Unfortunately, there is no way to disentangle these hypotheses in the data. In any case, consumers do not seem to make random mistakes. A more reasonable explanation is that financially unsophisticated consumers stick to other rules. In fact, it is quite plausible for unsophisticated consumers to rely on mental accounts if these processes serve as substitutes for otherwise complex financial decisions. In this regard, lack of financial literacy, mental accounting, and the use of heuristics are not mutually exclusive.

In summary, out of the explanations above, the most successful ones, in our view, involve mental accounting and lack of financial sophistication. Nonetheless, we acknowledge that small costs, unawareness of interest rates and even the cost-minimizing model of allocation are also plausible explanations for some consumers or time periods, particularly when we take into account the enormous heterogeneity in the data.

[28] In this regression, we exclude here cases in which individuals paid off both cards

6 Supply side implications

In this section we briefly discuss the supply side implications of our results. If consumers do not pay attention or do not understand interest rates, high and disperse prices can prevail in this market. This result mirrors the well understood effects of search and switching costs in the literature. Here, these costs include the cognitive efforts of performing mental calculations, the learning costs associated with these or alternatively, the mental costs of using a different card. Regardless of the motives, the implication is that banks will compete hard to attract 'new' customers in the first period, anticipating that if these clients borrow later on, they will do it on any card with available limit, regardless of the interest rate. In other words, since consumers can hold multiple credit cards, banks have incentives to assure clients keep the cards they issue in their wallets. This is particularly true when the market exhibits high inertia, as we documented before.

This up-front competition can take many forms, depending on the situation of the industry and the strategy used by each issuer. For instance, some banks could exploit the fact that consumers react to changes in the interest rates and compete on prices through introductory teaser rate offers. Likewise, other issuers could compete on the product space and introduce credit cards with rewards, mileages or logos. There is, in fact, substantial evidence that such strategies exist in the market.[29]

A natural question arises concerning why, if consumers do not pay attention to price differentials, some banks still compete up-front with low interest rate products (as we observe in our data). There are at least two reasons for this. First, certain consumers may indeed pay attention to prices and hold only low interest rate cards or alternatively, hold low and high interest rate cards but borrow mostly on the cheaper one. Second, as previously documented, many consumers take into account interest rates when making large purchases. In the end, whether a bank provides a certain kind of product depends on their assessment of the market demand for it.

Finally, our findings have another interesting implication. Namely, the fact that consumers allocate a larger fraction of their monthly purchases to their oldest card serves as an effective barrier to potential entrants. The extent to which these are relevant in practice requires further research.

7 Conclusions

Most consumers have more than one credit card. Yet, it is not evident how consumers manage these accounts. Here, we drew attention to the intra-temporal allocation of debt, payments and purchases among the credit cards consumers already hold. We constructed a novel data set that includes information on all the credit cards held by more than 10,000 consumers in Mexico. We find

[29] In Mexico, although introductory TROs are recent (mainly due to informational barriers), there is considerable evidence about such strategies on the product space.

that the difference in the interest rates between homogeneous cards is not an important determinant of allocations. Namely, individuals do not borrow on their cheapest card despite it is feasible for them to do so. On average, these cardholders forego potential savings for a sum that amounts to 16% of their financing cost. These results are difficult to reconcile with a standard cost minimizing model. We discuss potential explanations for these findings. Although we cannot fully disregard some hypotheses, we believe the most compelling explanation is mental accounting and financial unsophistication. Nonetheless, further research is necessary to disentangle the importance of the different factors affecting the allocation decisions of consumers. This is particularly relevant from a policy perspective, as the public policy implications of our study rely largely on the drivers of such behavior.

Overall, our paper speaks to the issue of financial decision making and financial mistakes in real world environments. In addition, this paper also adds to the so-called credit card puzzle and the literature on market implications of heuristics. Although the literature on these topics is still scarce, we hope that the results documented in this paper will generate further research.

Appendix

In this appendix we derive a simple period-by-period cost minimization problem. To compute the minimum cost we divide the consumer's optimization problem into two stages. In the first stage, for each period, consumers select the total amount of purchases and payments to be made with their credit cards on that month, and as a result, the total amount of debt to carry on. In the second stage, once individuals have selected the totals, they decide on the allocation of those purchases and repayments between the cards they already hold. Our model studies only the second stage. Therefore, we take total payments and purchases as given and solve for the allocation that minimizes interest and fees costs.

We assume that *rational* consumers decide the allocation of payments and purchases between cards to minimize the expenditure on fees and interests. That is, for the two card case, we solve for the optimal purchases (X^R_{it}) and payments (P^R_{it}) in each card ($i = 1, 2$) which solve the following cost minimization problem:

$$
\begin{aligned}
&\mathcal{C}(P_{1t}, P_{2t}, X_{1t}, X_{2t}) \equiv \\
&\min_{P_{1t},P_{2t},X_{1t},X_{2t}} \Big\{ C_{1t} \cdot 1[P_{1t} < PMin_{it}] + F_{1t} \cdot 1[D_{1t-1} - P_{1t} + X_{1t} > L_{1t}] + \\
&\qquad [D_{1t-1} - P_{1t} + X_{1t}] \cdot 1[D_{1t-1} > P_{1t}] \cdot r_{1t} + \\
&\qquad C_{2t} \cdot 1[P_{2t} < PMin_{2t}] + F_{2t} \cdot 1[D_{2t-1} - P_{2t} + X_{2t} > L_{2t}] + \\
&\qquad [D_{2t-1} - P_{2t} + X_{2t}] \cdot 1[D_{2t-1} > P_{2t}] \cdot r_{2t} \Big\}
\end{aligned} \tag{7}
$$

subject to

$$
\begin{aligned}
&PT_t = P_{1t} + P_{2t} \\
&P_{1t} \geq 0 \ , \ P_{2t} \geq 0 \\
&XT_t = X_{1t} + X_{2t} \\
&X_{1t} \geq 0 \ , \ X_{2t} \geq 0 \\
&D_{1t-1} - P_{1t} + X_{1t} \leq L_{1t} \ , \ D_{2t-1} - P_{2t} + X_{2t} \leq L_{2t} \\
&r_{it} \ , \ C_{it} \ , \ F_{it} \ , \ L_{it} \ , \ PMin_{it} \ , \ X_t \ , \ D_{it-1} \text{ and } P_t \text{ given for i=1,2.}
\end{aligned}
$$

where P_{it} is the payment made to card i in period t, X_{it} are the monthly purchases made with card i in period t, $PMin_{it}$ is the minimum monthly payment, r_{it} the monthly interest rate and L_{it} the credit limit. C_{it} is the late payment fee charged by the bank if the minimum monthly payment $PMin_{it}$ is not made during the billing cycle, F_{it} is the overlimit fee charged if the outstanding balance exceeds the credit limit L_{it} and D_{it} represents the outstanding balance at the end of the billing cycle. $1[\cdot]$ represents an indicator function.

The first line of the objective function corresponds to credit card 1. The first term represents the late payment fee C_{it} times an indicator variable that turns on when the payment to card 1 is less than the minimum payment[30]. The second term is the overlimit fee, F_{it} charged when the outstanding balance on card 1, including current payments and purchases, surpasses the credit limit. The third term corresponds to the interest rate paid. This expression is composed of three terms. The first one consists on the balance on which interests are charged. This balance is composed by the previous debt, minus the payments, plus the current purchases. Notice that, for debt revolvers, new purchases start accruing interests immediately. That is, wrong purchases allocations are costly regardless of the expectation about future payments when individuals hold debt in the more expensive card. The second term represents whether the individual pays the closing balance in full or not; if she does not do so, interests are accrued. The third term refers to the monthly interest rate of card 1. The line for the credit card 2 is analogous.

This is a linear problem that involves corner solutions. In general, these corner solutions are going to be defined by the values of payments and/or purchases for which the indicator functions turn on. Because of the interactions between payments and purchases, we search for the optimal purchases and payments in a {Purchase, Payment} grid of $1 dollar increments. We solve the model and derive the optimal purchases and payments for each card, period and consumer. Next, we use these allocations to estimate the optimal financing costs for each individual. With this at hand, we compare these costs with the actual financing costs faced by each cardholder and calculate the monetary costs of mistakes.

[30]Since we do not know the day the payment was made, we assume that they payment was within the grace period, thus we underestimate the cost of mistakes

References

Agarwal, Sumit, John C. Driscoll, Xavier Gabaix, and David Laibson, "Learning in the Credit Card Market," Working Paper 13822, National Bureau of Economic Research February 2008.

_ , Souphala Chomsisengphet, Chunlin Liu, and Nicholas S. Souleles, "Do Consumers Choose the Right Credit Contracts?," *CFS Working Paper Series*, 2005.

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

Avalos, Marcos and Fausto Hernandez, "Competencia Bancaria en Mexico," 2006. Serie Estudios y Perspectivas N 62 CEPAL.

Banxico, "Reporte sobre el Sistema Financiero," Technical Report, Banco de Mexico 2006.

Benartzi, Shlomo and Richard H. Thaler, "Naive Diversification Strategies in Defined Contribution Saving Plans," *American Economic Review*, March 2001, *91* (1), 79–98.

Bernheim, B. Douglas and Daniel M. Garrett, "The Effects of Financial Education in the Workplace: Evidence from a Survey of Households," *Journal of Public Economics*, August 2003, *87*, 1487–1519.

Bertaut, Carol C. and Michael Haliassos, "Debt Revolvers for Self Control," University of Cyprus Working Papers in Economics 0208, University of Cyprus Department of Economics May 2001.

_ and _ , "Credit cards: Facts and theories," in G. Bertola, R. Disney, and C. Grant, eds., *The Economics of Consumer Credit Demand and Supply*, MIT Press 2006.

Campbell, John Y, "Household Finance," *Journal of Finance*, 2006, *61*, 1553–1604.

Chamberlain, G., "Multivariate Regression Models for Panel Data," *Journal of Econometrics*, 1982, *18*, 5–46.

_ , "Panel Data," in "Handbook of Econometrics," Amsterdam: North Holland, 1984.

Choi, James J., David Laibson, and Brigitte C. Madrian, "Mental Accounting in Portfolio Choice: Evidence from a Flypaper Effect," *American Economic Review*, forthcoming.

DellaVigna, Stefano and Ulrike Malmendier, "Contract Design and Self-control: Theory and Evidence," *The Quarterly Journal of Economics*, May 2004, *119* (2), 353–402.

_ and _ , "Paying Not to Go to the Gym," *American Economic Review*, June 2006, *96* (3), 694–719.

Ellison, Glenn, "Bounded Rationality in Industrial Organization," in Whitney Newey Richard Burnell and Torsten Persson, eds., *Advances in Economics and Econometrics: Theory and Applications*, Cambridge University Press, forthcoming.

Greene, William H., *Econometric Analysis*, 6th edition ed., Prentice Hall, 2006.

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

Hastings, Justine and Lydia Tejeda-Ashton, "Financial Literacy, Information, and Demand Elasticity: Survey and Experimental Evidence from Mexico," February 2008. Unpublished manuscript, Yale University.

Lusardi, Annamaria and Peter Tufano, "Debt Literacy, Financial Experience and Overindebtedness," September 2008. Unpublished manuscript.

Madrian, Brigitte C. and Dennis F. Shea, "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior," *The Quarterly Journal of Economics*, nov 2001, *116* (4), 1149–1187.

Miravete, Eugenio J., "Choosing the Wrong Calling Plan? Ignorance and Learning," *American Economic Review*, 2003, *93*, 297–310.

Prelec, Drazen and George Loewenstein, "The Red and the Black: Mental Accounting of Savings and Debt," *Marketing Science*, 1998, *17*, 4–28.

Ranyard, Rob, Lisa Hinkley, Janis Williamson, and Sandie McHugh, "The role of mental accounting in consumer credit decision processes," *Journal of Economic Psychology*, August 2006, *27* (4), 571–588.

Shui, Haiyan and Lawrence M. Ausubel, "Time Inconsistency in the Credit Card Market," January 2005, (176).

Thaler, Richard, "Mental Accounting and Consumer Choice," *Marketing Science*, 1985, *4*, 199–214.

_ , "Mental Accounting Matters," *Journal of Behavioral Decision Making*, 1999, *12*, 183–206.

Wooldridge, Jeffrey M., *Econometric Analysis of Cross Section and Panel Data*, The MIT Press, 2002.

Zinman, Jonathan, "Household Borrowing High and Lending Low Under No-Arbitrage," April 2007. Dartmouth College Working Paper.

Figure 1: **The billing cycle**

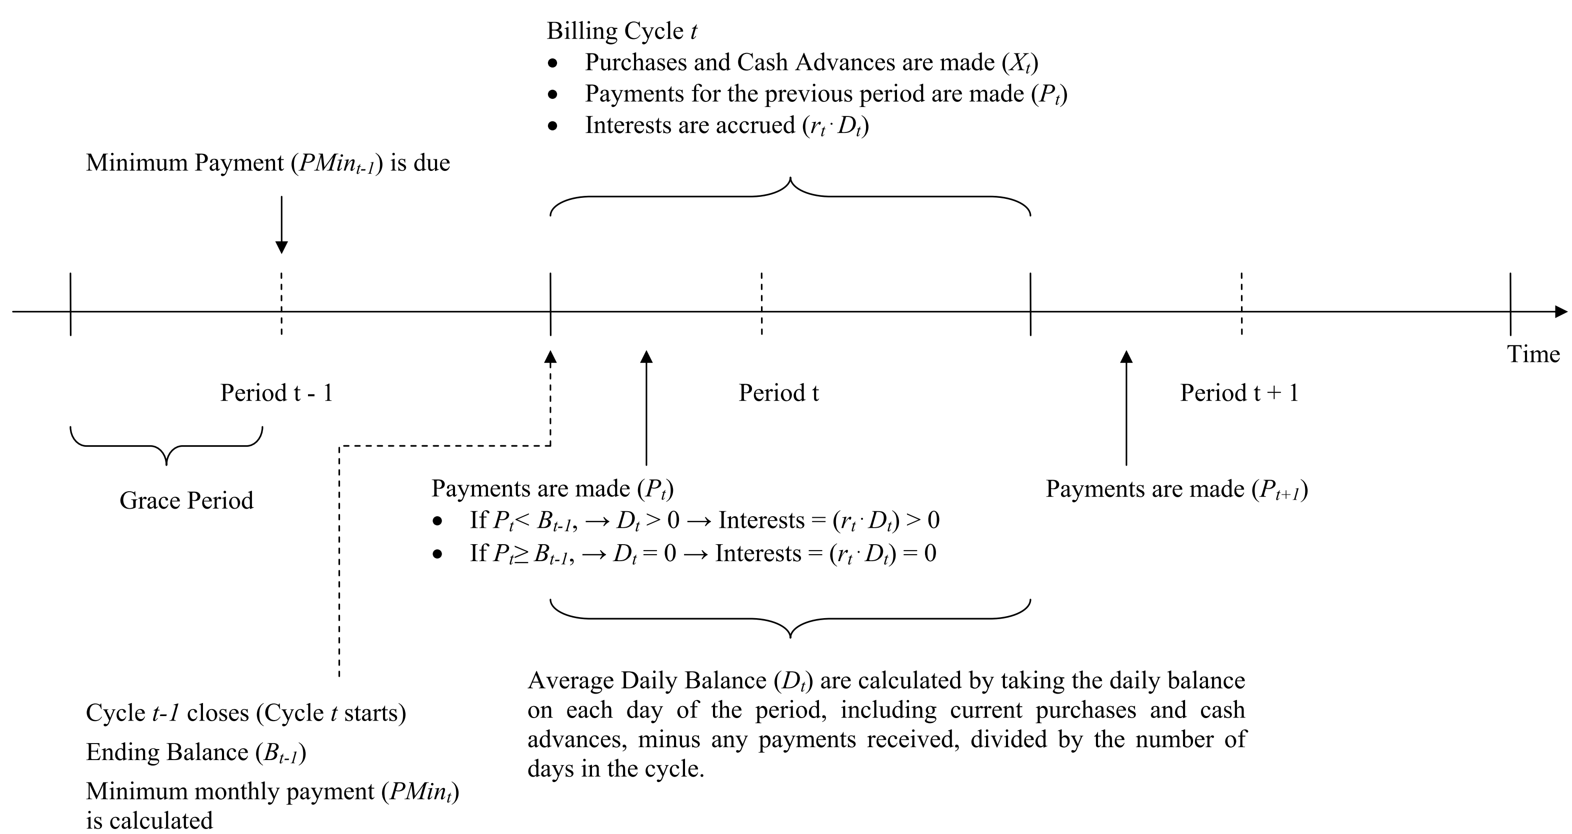

Figure 2: **Monthly interest rate differential across credit cards**

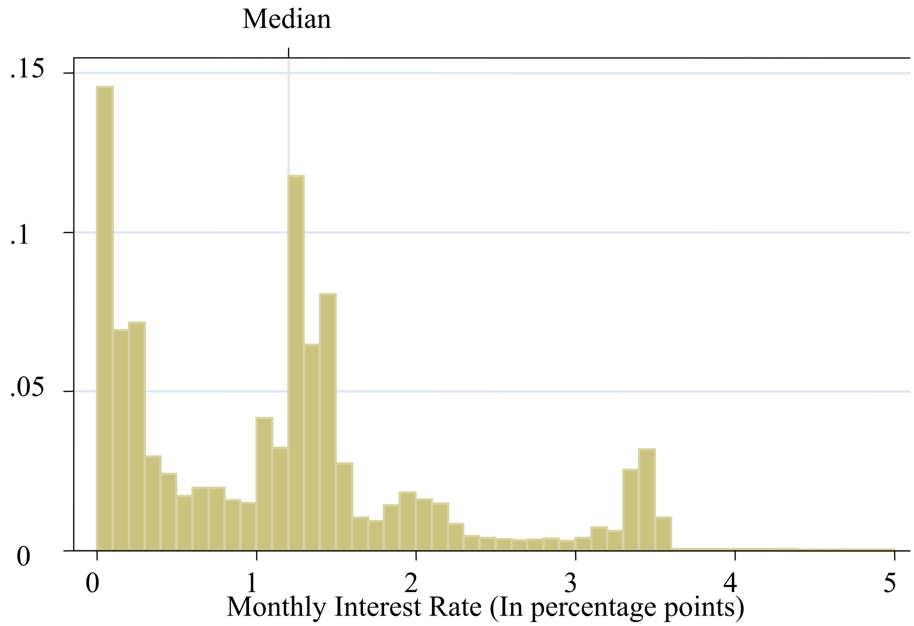

Based on 103,343 consumer-months

Figure 3: **Allocation of Interest Paying Debt**

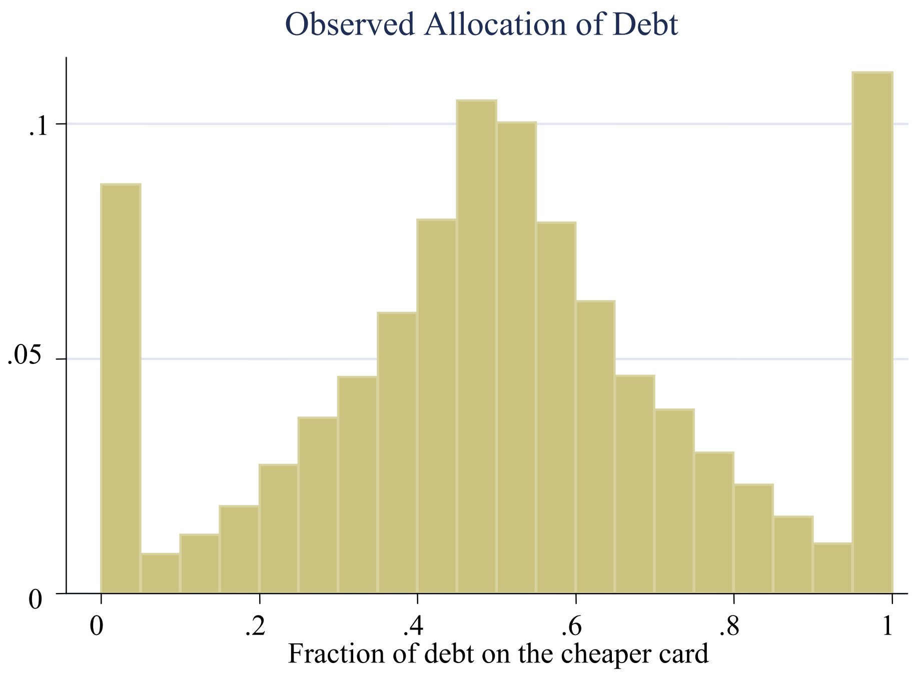

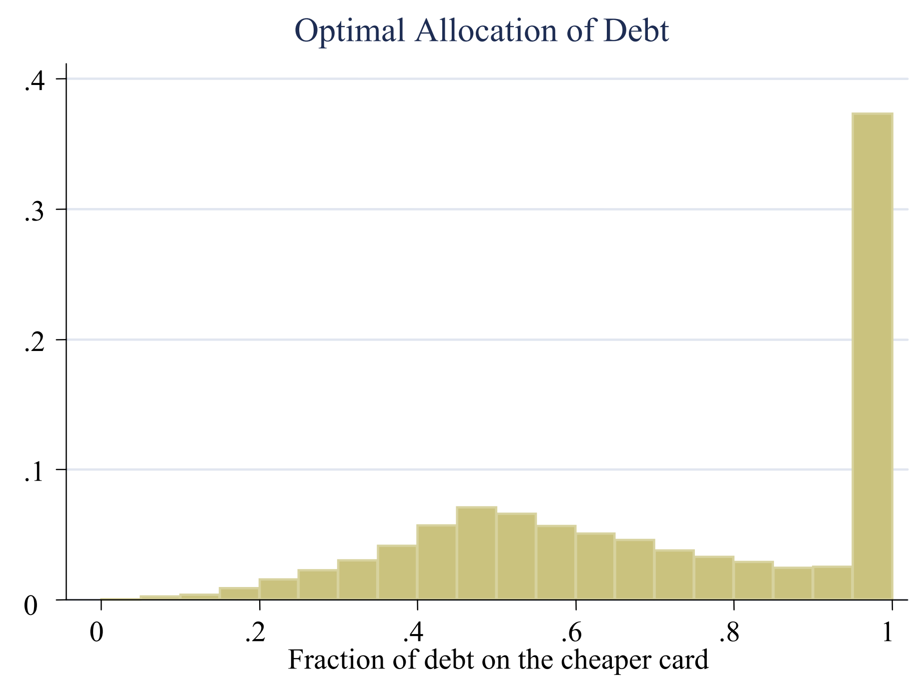

Figure 4: **Kernel Regression of the share of monthly purchases made with card 1 on the interest rate gap (r_1 - r_2)**

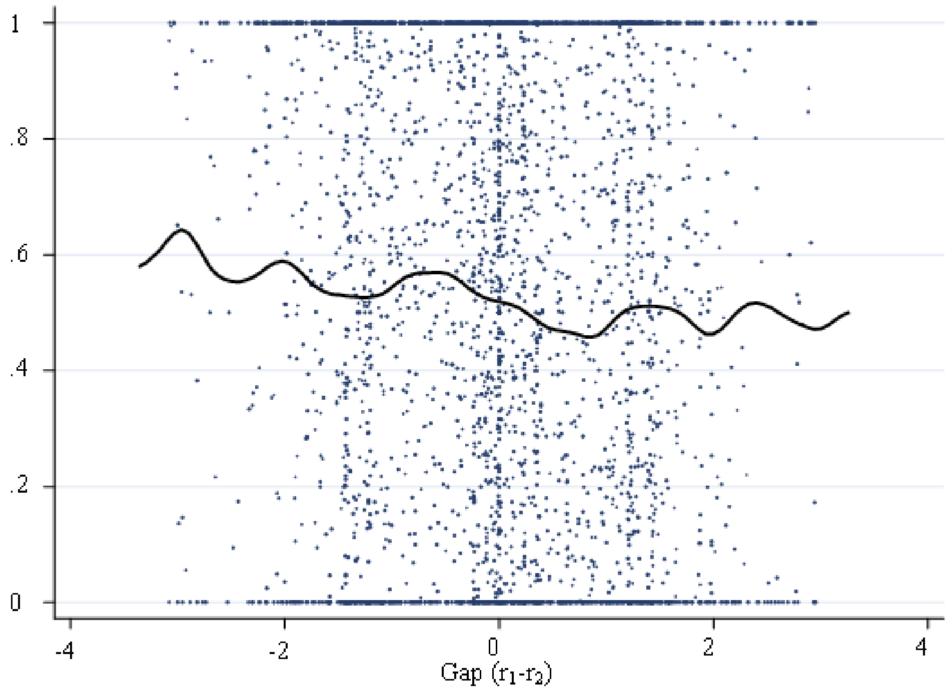

Figure 5: **Kernel Regression of the share of monthly payments above the minimum allocated to card 1 on the interest rate gap** (r_1 - r_2)

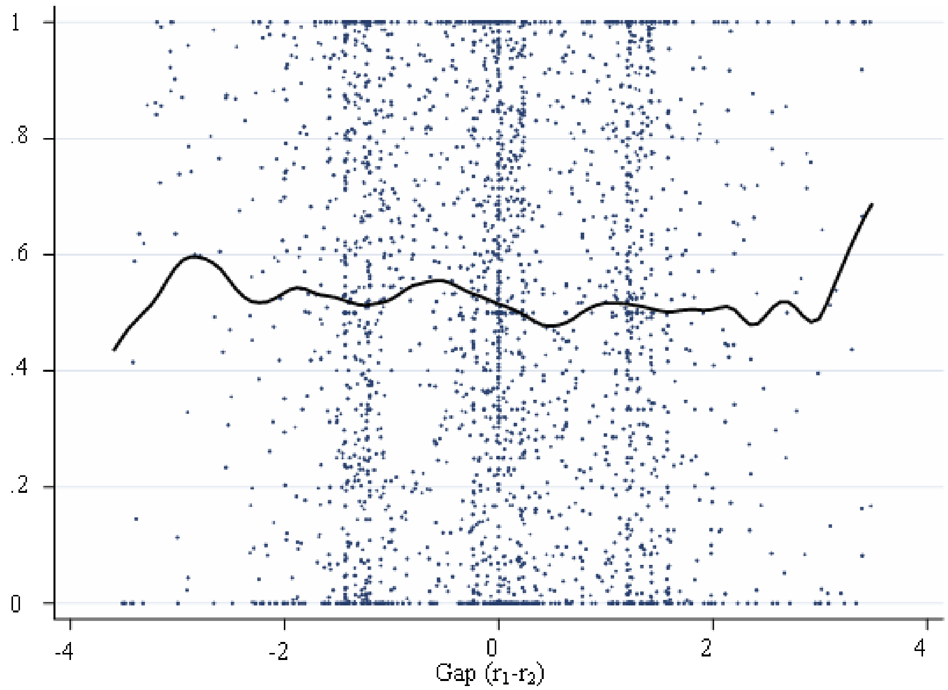

Figure 6: **Fraction of time consumers pay extra interest or extra fees**

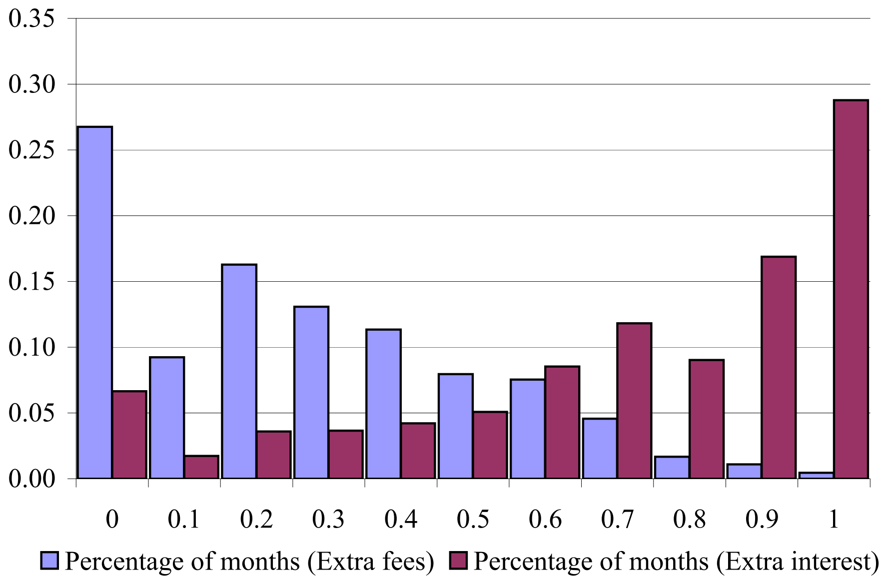

Figure 7: **Average misallocation cost in pesos**

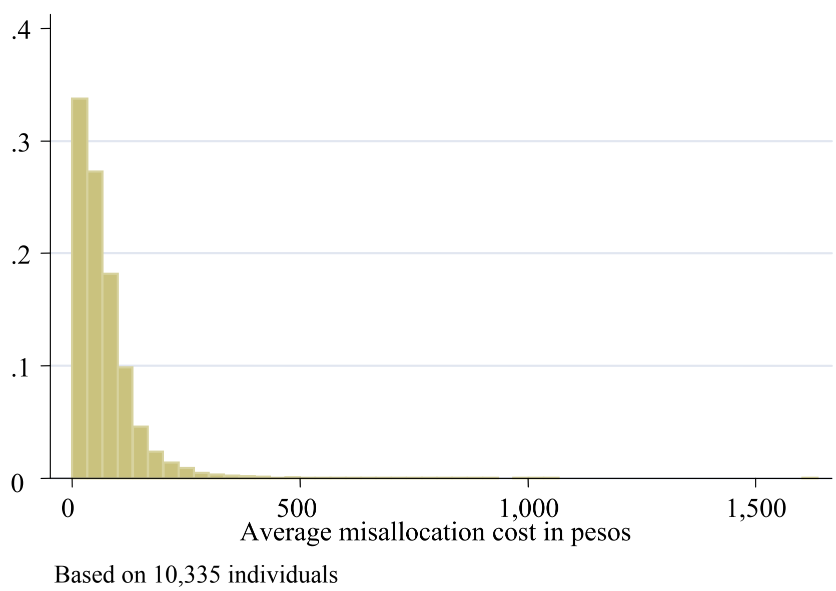

Figure 8: **Total misallocation cost as a fraction of total financing cost**

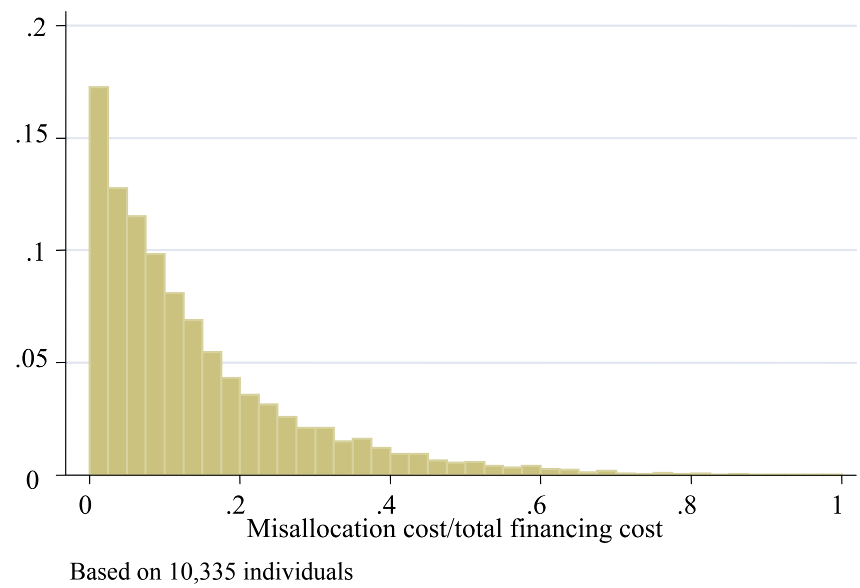

Figure 9: **Average misallocation cost by decile and source**

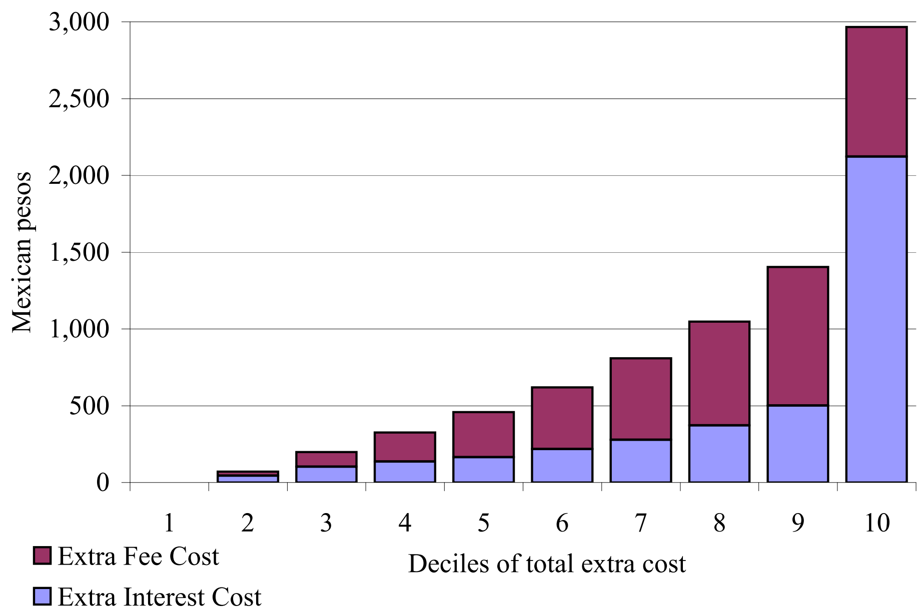

Table 1: **Summary statistics for consumers with two comparable credit cards**

The following table summarizes the mean of consumers' monthly flows and stocks for their combined credit cards. In Panel A an observation is a consumer-month; in Panel B and Panel C, the unit of analysis is a consumer. Figures are in Mexican Pesos. The raw sample includes 114,720 consumer-months. Once we consider only the periods in which debt is positive, the sample reduces to 103,343 consumer-months and 10,335 consumers. The statistics in Panel B are based only on the non-missing observations. The utilization rate was calculated as total debt over total credit limit. Standard errors are in parenthesis. In Panel C, the number of observations are in parenthesis.

	Raw (1)	Given > 0 (2)
Panel A: Monthly Statistics		
Interest paying debt	22,136	24,573
	(25,826)	(26,087)
Credit Limit	50,351	-
	(51,411)	
Purchases and cash advances	3,685	4,749
	(6,987)	(7,607)
Payments	4,185	4,465
	(6,159)	(6,263)
Fees	108	257
	(168)	(169)
Percentage of months paying fees	42	-
	(49)	
Percentage of months paying interests	90	-
	(30)	
Percentage of months borrowing in both cards	73	82
	(44)	(38)
Average Utilization Rate	61	65
	(40)	(37)
Debt weighted monthly interest rate	2.54	-
	(0.9)	
Panel B: Statistics by consumer (Percentage of consumers)		
Incur fees at least one month	81	-
Incur interests at least one month	97	-
Incur interests at least half the time	90	-
Borrow in both cards at least half of the time	74	-
Panel C: Demographics		
Monthly Income (Median)	10,000	-
	(5,012)	
Male	0.61	-
	(7,748)	
Age	44	-
	(4,999)	
Avg. years of tenure with a credit card	8	-
	(10,335)	

Table 2: **Allocation of purchases**

This table reports the effect of the interest rate differential on the fraction of purchases allocated to a randomly selected card. The columns labeled (1) through (8) each report separate regressions. The sample in these regressions includes observations for which total spending can fit into both of consumers' credit cards. The term labeled X in regressions (3), (5), and (6) is described in the top row of these columns. In all the specifications, we follow Chamberlain (1982) and include the means of all time-varying covariates as regressors. Standard errors, are given in parentheses. Letters a, b, and c indicate statistical significance at the 1, 5 and 10 percent confidence level, respectively.

	Base	Base	X = Large gap	Continuous	X = Large debt (large gap)	X = Large purchases	After TRO	After TRO both
	OLS	Tobit	OLS	OLS	OLS	OLS	OLS	OLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$CHEAP1_{it}$	0.06^a	0.10^a	0.04^a		0.06^a	0.06^a	0.05^b	0.09^a
	(0.01)	(0.01)	(0.01)		(0.04)	(0.01)	(0.02)	(0.03)
X_{it}			-0.02		0.01	-0.02^c		
			(0.01)		(0.02)	(0.01)		
$CHEAP1_{it} \cdot X_{it}$			0.05^a		-0.02	0.01		
			(0.02)		(0.03)	(0.02)		
$\overline{CHEAP1}_i$	-0.02	-0.01	0.02		-0.04	-0.02	0.00	0.03
	(0.01)	(0.03)	(0.02)		(0.04)	(0.02)	(0.03)	(0.07)
$\overline{X}_i$			0.04^c		-0.06^b	-0.02		
			(0.02)		(0.03)	(0.02)		
$\overline{CHEAP1 \cdot X_i}$			-0.08^a		0.08^c	0.04		
			(0.03)		(0.04)	(0.03)		
$r2_{it}$-$r1_{it}$				0.03^a				
				(0.00)				
$\overline{r2-r1}_i$				-0.01^b				
				(0.01)				
CONST	0.48^a	0.47^a	0.47^a	0.51^a	0.51^a	0.49^a	0.48^a	0.47^a
	(0.00)	(0.01)	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)	(0.03)
Random Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj/Pseduo R^2	0.005	-	0.006	0.004	0.003	0.005	0.006	0.030
Observations	19,773	19,773	19,773	19,773	9,042	19,773	3,982	422
Consumers	4,860	4,860	4,860	4,860	3,189	4,860	1,435	189

Table 3: **Allocation of payments above the minimum due between cards**

This table reports the effect of interest rate differentials on the fraction of payments above the minimum due allocated to a randomly selected card. The columns labeled (1) through (12) each report separate regressions. The term labeled X in regressions (3) and (5)-(9) is described in the top row of these columns. Column (10) presents the effect of interest rates on the probability of making the minimum payment due. This sample includes only observations in which individuals miss the minimum payment on one card despite it was feasible for them not to do so. In columns (1)-(5), (8)-(9), and (11)-(12) we follow Chamberlain (1982) and include the means of all time-varying covariates as regressors. Standard errors, clustered at the account level, are given in parentheses. Letters a, b, and c indicate statistical significance at the 1, 5 and 10 percent confidence level, respectively.

	Base	Base	X = Large gap	Continuous	X = Large debt (large gap)	X = No Δ Ranks*	X = Close due date*	X = Large Excess Payments	X = Previous Purchases (large gap)	Miss minimum payment	After TRO	After TRO both
	OLS	Tobit	OLS	OLS	OLS	OLS	OLS	OLS	OLS	Probit	OLS	OLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$CHEAP1_{it}$	0.01	0.01	0.01^a		0.03	0.01	0.00	0.02^a	-0.00	0.00	0.00	0.01
	(0.00)	(0.01)	(0.00)		(0.02)	(0.00)	(0.01)	(0.00)	(0.02)	(0.01)	(0.01)	(0.04)
X_{it}			0.02^b		-0.00	-0.01^c	-0.00	0.03^a	0.03^a			
			(0.01)		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)			
$CHEAP1_{it} \cdot X_{it}$			-0.03^a		-0.01	0.03^a	0.02^b	-0.05^a	0.01			
			(0.01)		(0.02)	(0.01)	(0.01)	(0.01)	(0.01)			
$\overline{CHEAP1}_i$	0.02^a	0.03^a	0.03^b		0.01			0.04^a	0.02		-0.03	-0.01
	(0.01)	(0.01)	(0.01)		(0.02)			(0.01)	(0.02)		(0.02)	(0.06)
$\overline{INT}_i$			0.00		0.03^c			0.05^a	0.16^a			
			(0.01)		(0.02)			(0.01)	(0.02)			
$\overline{CHEAP1 \cdot X}_i$			-0.01		-0.07^a			-0.12^a	-0.04			
			(0.02)		(0.02)			(0.02)	(0.01)			
$r1_t$-$r2_t$				0.001								
				(0.00)								
$\overline{r2 - r1}_i$				0.004								
				(0.00)								
CONST	0.49^a	0.48^a	0.48^a	0.50^a	0.49^a	0.50^a	0.50^a	0.47^a	0.40^a	-	0.51^a	0.51^a
	(0.00)	(0.01)	(0.01)	(0.00)	(0.01)	(0.00)	(0.00)	(0.01)	(0.01)	-	(0.01)	(0.03)
Random Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
Adj/Pseduo R^2	0.001	-	0.002	0.001	0.003	0.001	0.001	0.007	0.041	0.000	0.001	0.000
Observations	53,509	53,509	53,509	53,509	25,124	53,509	53,509	53,509	22,890	6,709	6,848	611
Consumers	7,807	7,807	7,807	7,807	5,452	7,807	7,807	7,807	5,262	3,683	2,103	230

Table 4: **Response to teaser rate offers**

This table reports the response to teaser rate offers. Each column corresponds to a different dependent variable. Standard errors clustered at the consumer level, are given in parentheses. Letters *a*, *b*, and *c* indicate statistical significance at the 1, 5 and 10 percent confidence level, respectively.

	Δ Debt1$_t$ OLS (1)	Δ Purchases1$_t$ OLS (2)	Δ Payments1$_t$ OLS (3)	Δ Debt2$_t$ OLS (4)
TR1M1$_{t+2}$	33	811^a	54	-257
	(208)	(281)	(238)	(241)
TR1M1$_{t+1}$	346^b	1,190^a	232	169
	(163)	(342)	(230)	(205)
TR1M1$_t$	1,909^a	-349	858^a	69
	(227)	(350)	(249)	(198)
TR1M1$_{t-1}$	-1,693^a	-1,281^a	123	-133
	(279)	(302)	(310)	(184)
TR1M1$_{t-2}$	-122	335	-1,229^a	2
	(193)	(345)	(263)	(199)
TR3M1$_{t+2}$	241	-93	-265	-212
	(193)	(306)	(290)	(286)
TR3M1$_{t+1}$	1,023^a	2,277^a	-406^c	-123
	(236)	(317)	(232)	(265)
TR3M1$_t$	2,209^a	-1,301^a	-38	-42
	(210)	(372)	(215)	(230)
TR3M2$_t$	1,112^a	-206	77	145
	(188)	(310)	(184)	(247)
TR3M3$_t$	1,289^a	121	-175	254
	(218)	(428)	(172)	(471)
TR3M3$_{t-1}$	269	-632	484^b	10
	(248)	(419)	(229)	(226)
TR3M3$_{t-2}$	686^c	813	635	-266
	(361)	(557)	(395)	(441)
Int. Rate Offers and its lags	Yes	Yes	Yes	Yes
Δ Limit and its lags	Yes	Yes	Yes	Yes
Time indicators	Yes	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes	Yes
Adj/Pseduo R^2	0.009	0.004	0.004	0.002
Observations	97,905	97,905	97,905	97,905
Consumers	10,329	10,329	10,329	10,329

Table 5: **Percentage of observations in which individuals make allocation mistakes**

The following table presents the percentage of observations in which individuals make allocation mistakes. The rows show the percentage of observations in which individuals pay less, equal or more interest relative to the optimal allocation. Similarly, the columns correspond to the percentage of observations in which consumers pay less, equal or more fees compared with the optimum. The table is based on 103,343 consumer-months.

		Fees		
		Less or Equal	More	Total
Interest	Less	0%	8%	8%
	Equal	16%	1%	17%
	More	58%	16%	74%
	Total	74%	26%	100%

Table 6: **Patterns of behavior in the allocation of payments and purchases**

This table reports the effect of proxies for heuristics on the allocation of payments and purchases. The dependent variable in column (1), (2), and (9) is the fraction of payments above the minimum due allocated to a randomly selected card (labeled card 1). The dependent variable in columns (3), (4), and (10) is the fraction of purchases allocated to a randomly selected card. In column (5), the dependent variable is an indicator that equals one if the individual made the minimum payment to card 1 in that month. The sample in that regression consists of consumers who missed the minimum payment on one card, despite it was feasible for them to cover it up during that period. The dependent variable in columns (6)-(7) is an indicator defined in the top row of these columns. Standard errors, clustered at the account level, are given in parentheses. Letters a, b, and c indicate statistical significance at the 1, 5 and 10 percent confidence level, respectively.

Hypotheses:	**Debt equalization (1) - (5)**					**Snowball (6) - (7)**		**Inertia (8) - (9)**	
Dependent Variable:	Share of purchases in card 1 (1) - (2)		Share of payments in card 1 (3) - (4)		Miss minimum in card 2	Paid card 1 in full (6) - (8)		Sh. of purchases in card 1	Paid in full card 1
	Base	Large Purchases	Base	Large Payments	Minimum payment	X = $1[D_1<\text{TP}<D_2]$	X = $1[D_2<D_1<\text{TP}]$	Previous Purchases	Base
	OLS	OLS	OLS	OLS	Probit	OLS	OLS	OLS	Probit
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$CHEAP1_{it}$	0.06^{a}	0.07^{a}	0.01	-0.00	0.00	-0.06^{a}	-0.00	0.05^{a}	0.00^{c}
	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.01)	(0.00)
$DEBT1_{it}$	-0.09^{a}	-0.10^{a}	0.05^{a}	0.12^{a}	0.23^{a}			-0.08^{a}	
	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)			(0.01)	
$LIMIT1_{it}$	0.02^{b}	0.04^{b}						0.04^{a}	
	(0.01)	(0.02)						(0.01)	
$PURCH1_{it-1}$ / $PFULL_{it-1}$								0.20^{a}	0.20^{a}
								(0.01)	(0.01)
$PURCH1t0_i$ / $PFULL0_i$								0.12^{a}	0.09^{a}
								(0.01)	(0.01)
X_{it}						0.06^{a}	-0.04		
						(0.02)	(0.02)		
$CHEAP1_{it} \cdot X_{it}$						0.08^{a}	-0.05		
						(0.03)	(0.03)		
$\overline{CHEAP1}_i$	-0.02	0.01	0.02^{b}	-0.05^{a}				0.02	-0.00^{b}
	(0.01)	(0.02)	(0.01)	(0.01)				(0.02)	(0.00)
$\overline{DEBT1}_i$	0.22^{a}	0.21^{a}	0.02^{a}	0.12^{a}				0.09^{a}	
	(0.01)	(0.02)	(0.01)	(0.01)				(0.02)	
$\overline{LIMIT1}_i$	0.01	0.03						0.01	
	(0.01)	(0.02)						(0.02)	
CONST	0.41^{a}	0.36^{a}	0.45^{a}	0.41^{a}		0.44^{a}	0.48^{a}	0.28^{a}	
	(0.01)	(0.02)	(0.01)	(0.01)		(0.02)	(0.01)	(0.01)	
Random Effects	Yes	Yes	Yes	Yes	No	No	No	Yes	Yes
Adj/Pseduo R^2	0.029	0.049	0.011	0.108	0.038	0.013	0.003	0.165	
Observations	19,773	4,935	53,509	12,690	6,709	9,609	9,609	8,561	67,624
Consumers	4,860	2,013	7,807	4,166	3,683	3,367	3,367	1,710	7,574