

Borrowing on the Wrong Credit Card? Evidence from Mexico

Online Appendix

Alejandro Ponce

Enrique Seira

Guillermo Zamarripa

This is the Online Appendix for “Borrowing on the Wrong Credit Card? Evidence from Mexico”. For readability the Appendix is organized following the structure of the paper.

I The Mexican credit card market

The credit market in Mexico has remained relatively underdeveloped in terms of credit penetration, lending practices, and product development. Mexico’s private credit to GDP ratio in 2015 was less than 30 percent, low even by Latin American standards. This has been partly due to the history of nationalization, privatization, and crises described in [Haber \(2005\)](#),¹ which have also affected the development of the credit card market. By 2004 — the beginning of our sample period — there were few credit card products available, little differentiation across suppliers, and shallow market penetration. At the time:

- There were basically four types of credit card products: Classic cards, Gold and Platinum cards, and store cards. Classic cards offered no rewards or benefits. Approximately 65 percent of the cards in circulation were of this type. Gold and Platinum cards offered rewards programs that had similar structures and comparable items for redemption across suppliers (at about 1 percent cashback equivalent).
- Rewards, fees, initial credit limits, and contractual interest rates were determined at the credit card-type level (e.g., Classic, Gold, etc.) and did not vary with the cardholder’s risk profile. This was partly because banks were unsophisticated in their statistical pricing techniques and had limited access to reliable information about potential clients. As of 2004, for example, the Credit Bureau was expanding its databases to include information from nonbank and unregulated institutions, and in 2006 it developed a credit score summarizing a person’s creditworthiness. Even today, Mexican banks are still relatively unsophisticated in their risk and behavioral based pricing.² Furthermore, due to regulatory constraints, only banks were able to issue credit cards, and all cards were linked to the Visa and MasterCard organizations that were universally accepted in all POS devices.
- There were only 0.13 credit cards per person in the country compared to 0.35 in Argentina, 0.38 in Brazil, and 2.53 in the U.S. ([US \(2008\)](#)). After 2004, the market began showing more dynamism, as shown by the growth in the number of credit cards from 9.8 million in 2004 to 13.5 million in 2005. This growth came largely from banks issuing new cards to existing cardholders, as reflected in the fact that between 2006 and 2008 the number of cards held by the average cardholder increased from 3.4 to 4.2 ([Banxico \(2009\)](#)).

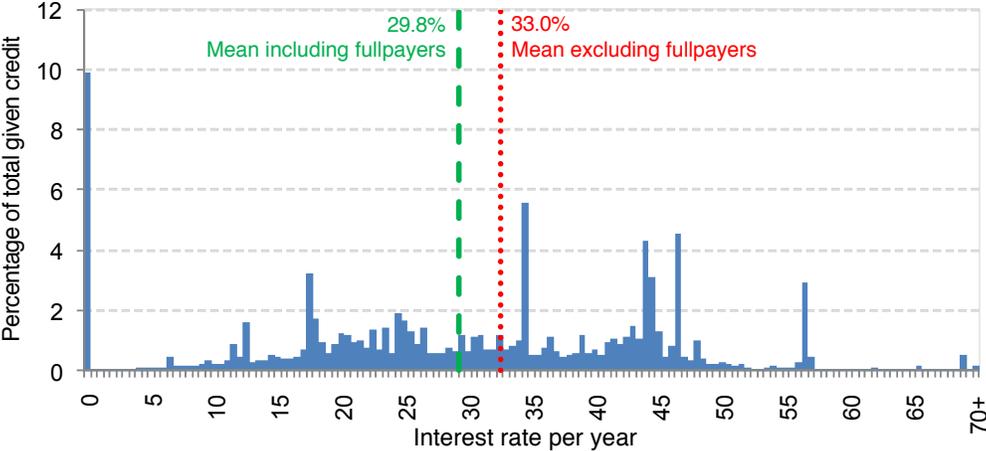
The banking industry in general and the credit card industry in particular have also been highly concentrated. From 2001 to 2012, the five largest banks held a steady market share of close to 90%. This, in turn, might explain the relatively high return on equity of Mexico’s

¹Banks were nationalized in 1982, privatized and sold to individuals with limited experience in 1990, and almost bankrupt in 1994 by a combination of bad lending practices and the Tequila crisis. The government bailed out the banking industry, after which banking underwent a slow process of capitalization and resale throughout the 2000s.

²As we describe in Section III, we obtained the credit card records of thousands of clients of one bank for about 64 months (the PT data set described below). This data set contains an estimate made by the bank of the probability of default in the next 12 months for each client. When we regress the interest rate on decile dummies of this probability, we get an R^2 of 0.12. In contrast, when we regress the interest rate on dummies for the type of card, we get an R^2 in the order of 0.45.

banks at about 25 percent in 2005 –above the ratios of Korea (19.6), the United Kingdom (17.5), Chile (16.9), the U.S. (14.1), and Spain (13.6), for example (Cofeco (2007))– as well as high fees and interest rates. As of 2005, overdraft fees ranged from \$50 to \$100 pesos and late payment fees ranged from \$180 to \$275 pesos³, depending on the bank. According to Avalos and Hernandez (2006), banking fees in Mexico ranked among the highest in Latin America. Similarly, as of 2005, the average credit card interest rate was close to 34 percent per year, while credit card loans in arrears amounted to a modest 3 percent and the government federal discount rate (TIIE) was about 7 percent (Banxico (2006)). Interestingly, though, high average prices were accompanied by substantial dispersion in interest rates across products and suppliers. Graph OA.1 comes from a 2009 report of Mexico’s Central Bank. It plots a histogram of the effective interest rate (interest paid/total monthly debt) pooling all banks and credit card products together. The Figure shows a staggering amount of dispersion in credit card interest costs across consumers. The graph, however, does not control for cardholders’ risk or any card or borrower characteristic, so we cannot directly ascertain whether there was the same price dispersion for comparable borrowers or credit card products. The Figure also shows that, in spite of the high interest rates, approximately 10 percent of the observations did not incur any interest (owing to a zero interest rate or a zero average balance). It is noteworthy that the average interest rates displayed in this Figure are very close to those in our sample.

Figure OA.1: Interest rate dispersion in the Mexican credit market in 2009



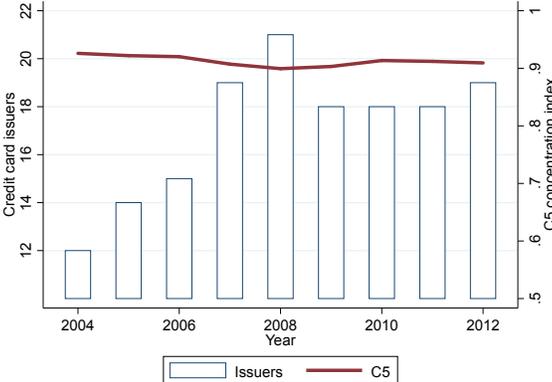
Notes: Data are from Banco de Mexico (2009a). This Figure plots a histogram of the effective interest rate (X-axis) against the percentage of total debt given in the market (Y-axis). For each credit card, the effective interest rate is the quotient between the total amount of interest incurred and the total amount of debt. This interest rate takes into consideration promotions and temporary teaser rate offers. The dashed line plots the mean effective interest rate including those consumers who usually do not pay interest. The dotted line plots the mean effective interest rate excluding consumers who do not pay interest.

In 2007, the competition authority issued an opinion stating that there was low competition

³The symbol for a Mexican peso is \$, and 10.6 Mexican pesos were equivalent to 1 U.S. Dollar as of December 2005.

in the credit card market. Faced with this situation Mexican authorities adopted a number of policies and regulations to foster competition, including promoting the entry of new banks, regulating fees, mandating banks to explain charges and interest rates more clearly,⁴ and requiring them to accept payments to credit card accounts via electronic transfers from other providers.⁵ According to the World Bank, Mexican banks now face fewer barriers to entry than countries like Canada, Spain, or Denmark, among many others [Banxico \(2013\)](#). In spite of these policies and the doubling in credit card issuers from 2004 to 2012, there has been almost no change in the market share of the largest five banks and only small reductions in prices (see [Figure OA.2](#) below), suggesting that in addition to supply side factors –such as substantial concentration, relatively low quality of credit bureau information, or regulatory constraints in assessing the riskiness of potential customers that gave banks with large branch networks and client bases informational advantages over smaller banks [Banxico \(2013\)](#) –, demand side factors may play an important role in shaping competition in Mexico.

Figure OA.2: Credit card issuers and C5 concentration index



Notes: The data presented here are from Banco de Mexico and CNBV. This chart shows the number of credit card issuers (bar graph) and the C5 concentration index (line graph), which is defined as the share of the five largest credit card providers in terms of all of the credit cards issued for the period of 2004-2012 in Mexico. The bar graph corresponds to the left axis and the line graph corresponds to the right axis.

⁴At the time, banks were required to send their clients their monthly credit card statements disclosing the applicable interest rate. The rate, however, was not displayed as transparently as in the Schumer Box used in the U.S.

⁵In December 2002, a new law, “Ley de Sistemas de Pagos”, was enacted. This law provided the general framework for balance transfers by mandating that banks accept payments to credit card accounts via electronic transfers from other providers. In January 2004, another law came into effect, “Ley para la Transparencia y el Ordenamiento de los Servicios Financieros”. This law regulated bank fees for the use of credit cards and ATMs and mandated that lenders disclose information about loan terms and financial transactions more clearly. This law has been improving and has seen at least three more modifications from 2004 to 2014.

II Data

C. Survey Data

(i) Survey Design

The survey used in the paper was designed primarily to elicit the reasons why cardholders misallocate their debt. The sample for this survey consisted of 1,003 adults living in Mexico City, who had at least one active credit card, and who were the main account holder. A private polling company, called *Data Opinión Pública y Mercados*, was hired to carry out the fieldwork. To explore approaches to construct the sample, we conducted a two-stage pilot survey using the Mexican Census to select the primary sampling units, and random route procedures to select the households. Unfortunately, due to the shallow credit card penetration and prevailing insecurity, this exercise resulted in a low hit rate. This precluded us from using a strictly stratified sampling technique and obliged us to employ intercept point sampling and snowball sampling techniques. The field work was carried out from September 13th, 2013 to January 15th, 2014. The initial interviews were conducted by 23 data collectors and 6 supervisors in the following public places of Mexico City: Bellas Artes, Alameda Central, Glorieta de los Insurgentes, Parque Mexico, Parque de los Venados, Parque de las Arboledas, Parque Viaducto, WalMart Universidad, Parque Tlacoquemecatl, Alameda Sur, Santo Domingo, Rectoría Explanada-UNAM, WalMart Perisur, Centro San Angel, Parque Cuitlahuac, Centro Iztapalapa, and Explanada San Cristobal Ecatepec. Starting in November, respondents were asked to provide referrals to other cardholders, who in turn were interviewed at their homes or workplaces. This snowball process continued until our target sample size was achieved. Sample quotas were established to over-represent consumers who had more than one card and who regularly carried interest paying debt. Early in the survey, respondents were asked three screening questions to determine if they fitted the profile needed. The final quotas were: individuals holding one card and outstanding balances (13%); individuals holding one card and no outstanding balances (paid off balances) (2%); individuals holding two cards and outstanding balances (40%); individuals holding two cards and no outstanding balances (5%); individuals holding more than two cards and outstanding balances (35%); and individuals holding more than two cards and no outstanding balances (5%). Respondents were interviewed in person using Google Nexus 7 tablets, and were given \$100 pesos in grocery coupons as an incentive for participating in the study. The interviews were done in Spanish and took between 30 and 50 minutes to complete. The questionnaire will be available at www.fundef.org.mx, both in English and in Spanish. In the paper and throughout this Online Appendix, we present summary statistics and results from this survey.

(ii) Measures of financial literacy and cognitive ability in the survey

The survey includes questions about financial literacy and cognitive ability. Some of the financial literacy questions were similar to the ones most used in the literature, while others were tailored to credit cards. We use 3 of these questions to build an index: **Q1**. Suppose

you had \$100 in a savings account and the interest rate was compounded at 2% per year. After 5 years, how much do you think you would have in the account if you left the money to grow: More than \$102 pesos; exactly \$102 pesos; less than \$102 pesos? **Q2:** Suppose you have an outstanding balance of \$1,000 pesos on your credit card. Which of the following options would generate the least interest on such a balance over a period of 6 months? If your card offered you: 3 months at 40% APR and 3 months at 0% APR; or 6 months at 30% APR. **Q3:** Suppose you have two credit cards: One has 30% APR and the other 40% APR. Imagine you are going to make a purchase worth \$5,000 pesos that would accrue interest over a period of one year. Approximately how much extra interest would you pay if you use your expensive card rather than your cheap card: Around \$200 pesos; around \$500 pesos; around \$1,000 pesos; around \$5,000 pesos? Our survey also includes 5 Raven matrices that we use to build a cognitive ability index. This summary measure is far from ideal. Unfortunately, due to possible respondent fatigue, we could not include the 16 or 64 item Raven Test.

D. Descriptive Statistics

Table OA.1 below presents month-level descriptive statistics for selected variables in the administrative and survey data sets for individuals holding two comparable cards, broken down by tertiles of the interest rate gap. In general, the variables considered show little variation across the three groups. Consumers holding cards with large price differentials display similar account and borrower characteristics as those holding two similarly priced cards.

Table OA.1: Summary statistics for selected variables in the administrative and survey data by tertiles of the interest rate gap

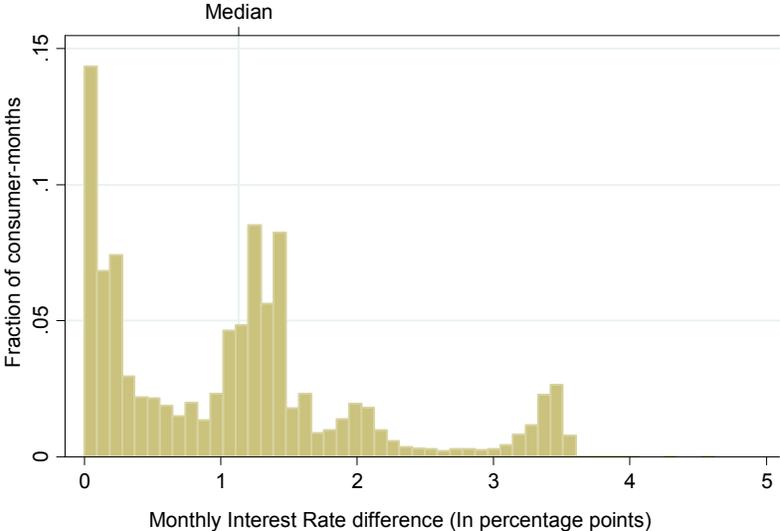
Range Gap tercile	Administrative Data			Survey data		
	[0.00, 0.73] (1)	[0.73, 1.32] (2)	[1.32, 3.49] (3)	[0.00, 0.25] (1)	[0.29, 0.79] (2)	[0.83, 3.83] (3)
Avg. monthly interest rate gap	0.31	1.09	1.92	0.11	0.49	1.37
Interest paying debt	23,968	24,132	18,307	17,791	22,109	19,790
Credit limit	48,459	48,221	54,379	35,117	38,077	54,050
Purchases	3,364	3,396	4,295	6,253	5,466	4,341
Payments	4,037	3,842	4,678	4,976	3,889	2,584
Fees	117	111	97	-	-	-
Pct. of months paying fees	44	43	39	-	-	-
Pct. of months paying interests	91	96	83	64	58	60
Debt weighted monthly interest rate	2.69	2.48	2.29	2.67	2.96	3.41
Tenure with the oldest card	8.57	7.94	8.87	5.04	5.74	6.04
Income	15,768	14,597	14,476	13,498	13,153	13,974

Notes: Data are from the administrative database (2-comparable cards sample) and from the survey data set (2-comparable cards sample). This table presents month-level means for selected variables for individuals holding two comparable cards, broken down by tertiles of the interest rate gap. .

E. Interest rate variation

Our empirical analysis exploits the variation provided by the difference in interest rates across the cards held by each consumer. Figure OA.3 below shows the distribution of this difference. The median difference is 1.1%, and the overall standard deviation is 1.18 percentage points, with 68% of it coming from variation across consumers and the rest from within-consumer time-series variation. An important source of this latter variation is the result of TROs that banks send to existing clients with a duration ranging from one to six months. On average, these short-term offers have a discount of 1.8 percentage points per month. In our sample, there are 3,979 individuals who received at least one such offer on at least one of their cards. In spite of these offers, 53% of the individuals in our data never experienced a change in the interest rank of their cards (i.e., a given card was always more expensive). The modes are generated by two levels of TROs, with some of them being zero interest rate offers and others being decreases of approximately 1 or 1.5 percentage points. These latter offers are similar to those we use in the paper to estimate elasticities.

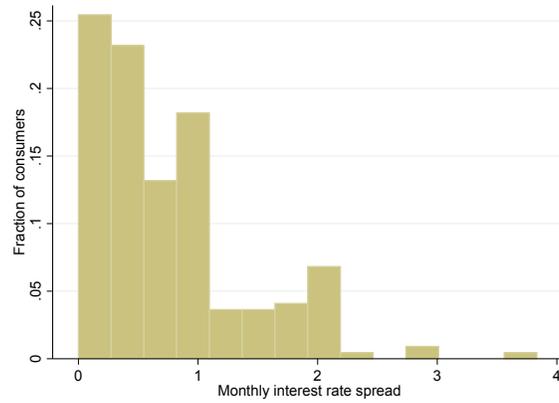
Figure OA.3: **Distribution of the monthly interest rate gap between credit cards (Administrative Data)**



Notes: Data are from the administrative database (2-card sample). An observation is a consumer-month. The figure shows the frequency distribution of the monthly interest rate gap (in percentage points) between the credit cards held by the individual.

Figure OA.4 below shows the distribution of the difference in interest rates across the cards held by survey respondents who reported having revolving debt.

Figure OA.4: **Distribution of the reported monthly interest rate gap between credit cards (Survey Data)**



Notes: Data are from the survey database. An observation is a consumer. This Figure shows the difference in interest rates for people who had two cards and who answered positively to the question asking if they regularly revolve debt.

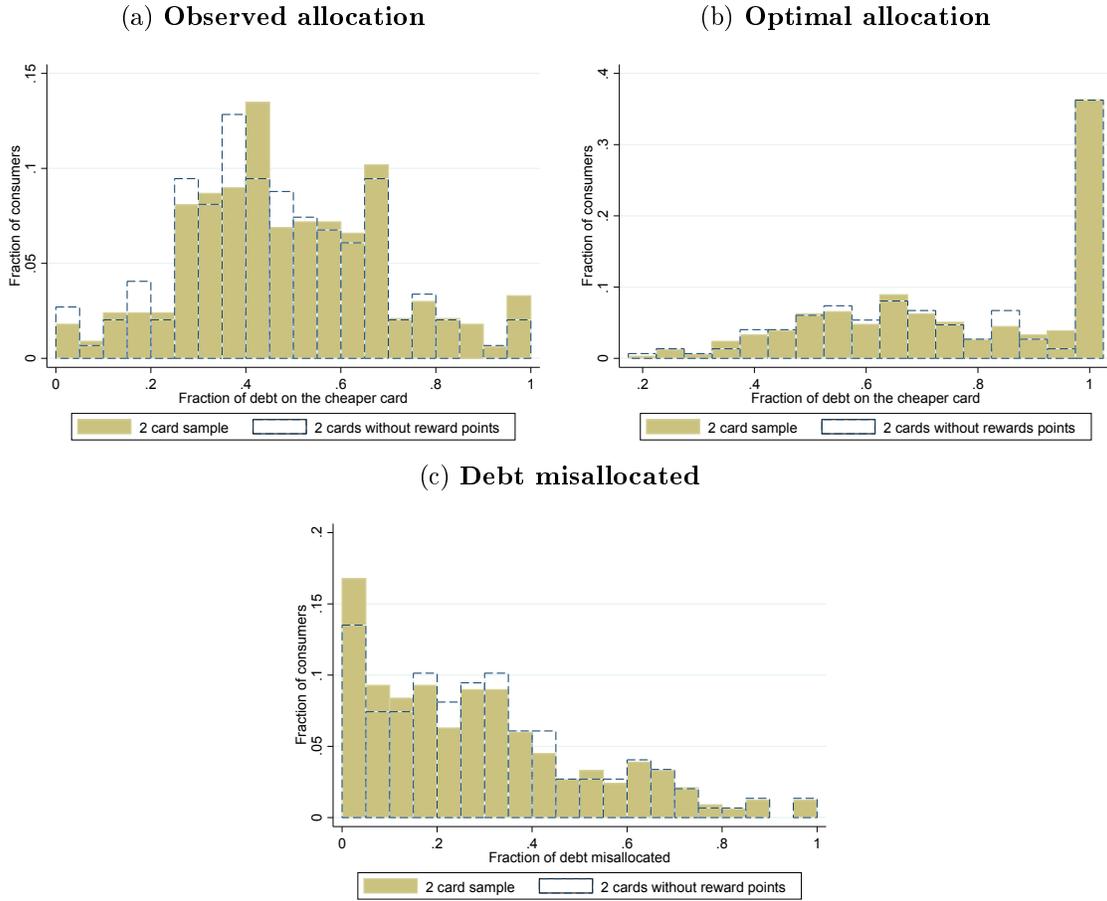
III Relative prices as determinants of debt, purchases, and payment allocations

A. Allocation of debt

(i) Robustness to various samples

Figure [OA.5](#) shows histograms of the fraction of credit card debt allocated to the low-interest card (Panel A), the fraction of debt on the low-interest card that would minimize interest cost (Panel B), and the fraction of misallocated debt for survey respondents holding two cards, or two cards without rewards. The fraction of misallocated debt is calculated as the high-interest debt that could be shifted onto a lower-interest card up to the credit limit divided by the total credit card debt reported by each respondent. These figures refer to the amounts reported as of the last month before the interview.

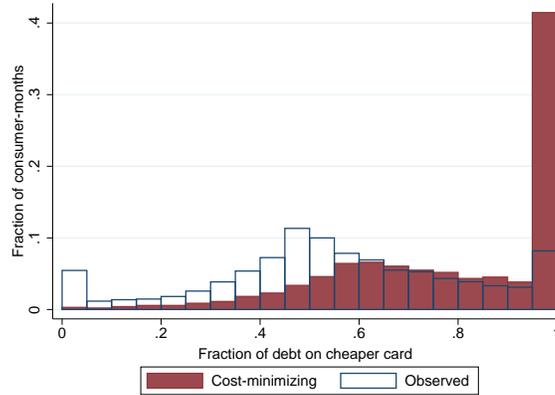
Figure OA.5: Allocation of debt to the low-interest card (Survey data)



Notes: Data are from the survey database (2-card sample). An observation is a survey respondent. The Figure shows a histogram of the fraction of credit card debt allocated to the low-interest card (Panel A), a histogram of the fraction of debt on the low-interest card that would minimize interest cost (Panel B), and the fraction of misallocated debt (Panel C). The solid line represents the sample of respondents holding two credit cards. The dotted line represents the sample of respondents holding two cards where neither of those cards have reward points.

Figure OA.6 replicates Figure 1 in the paper using the Purchase-Type data set. It shows the distribution of the fraction of credit card debt allocated to the low-interest card and the distribution of the fraction of debt on the low-interest card that would minimize interest cost.

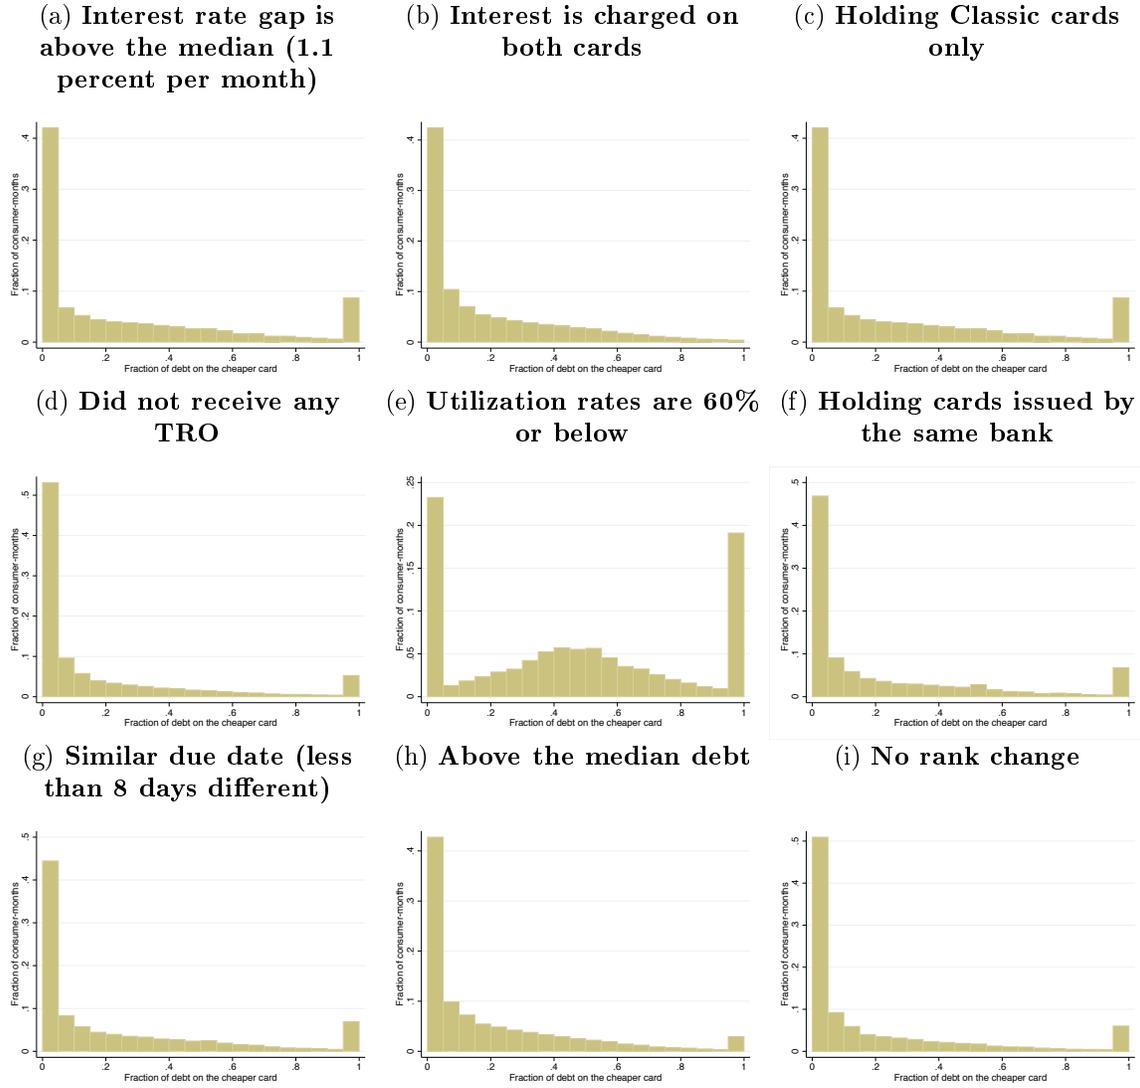
Figure OA.6: **Share of debt observed and misallocated on the Purchase-Type database**



Notes: Data are from the administrative database (purchase-type data). This figure shows the distribution of the fraction of debt allocated to the low-interest card (hollow blue) and the distribution of the fraction of debt on the low-interest card that would minimize interest cost (solid red).

Figure OA.7 below shows that debt misallocation is robust to 9 sub-samples of the administrative data where: (a) the difference in interest rates is above the median, (b) interest is charged on both cards, (c) both cards are classic cards and therefore have no reward points, (d) no card received TROs during the sample period, (e) both card have less than 60 percent line utilization, (f) both cards were issued by the same bank, (g) both cards have a due date differential of eight days or less, (h) total debt is above the median and, (i) there was no change in the price ranking of the two cards during the sample period.

Figure OA.7: Fraction of credit card debt that was misallocated to the high-interest card in various subsamples



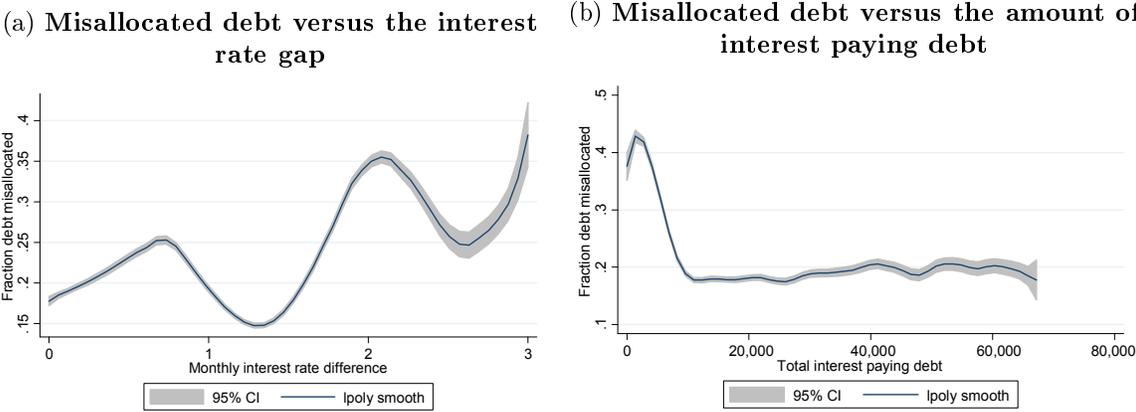
Notes: Data are from the administrative database (2-card sample). An observation is a consumer-month. The figure shows the frequency distribution of the fraction of credit card debt that was wrongly allocated (misallocated) to the high-interest card in several sub-samples. This fraction is calculated for each period as the high-interest debt that could be shifted onto the low-interest card up to its credit limit divided by the total credit card debt held by the consumer during that month.

(ii) Robustness to “stakes”

Individuals might borrow on their high-interest cards because the monetary stakes are not high enough. Figures OA.8a and OA.8b display non-parametric regressions of the fraction of

misallocated debt on the average monthly interest rate gap (Panel A) and on total credit card debt in pesos (Panel B). The interest rate gap and the total level of debt are not related to the fraction of misallocated debt.⁶ The fraction of misallocated debt decreases slightly to 20 percent for values of debt below \$10,000 pesos (about a third of the sample falls within this range), and remains flat for values above that threshold.

Figure OA.8: Fraction of debt misallocated, interest rate gap and total debt



Notes: Point estimates and 95-percent confidence intervals plotted. Data are from the administrative database (2-card sample). An observation is a consumer-month. The figure shows a kernel regression of the fraction of credit card debt that was misallocated to the high-interest card on the average monthly interest rate gap (in percentage points) on Panel A and on total credit card debt (in Mexican pesos) on Panel B, both using an Epanechnikov kernel. Total debt is trimmed at the 95th percentile. The interest rate gap is trimmed at 3 percentage points.

The table below shows a series of regression results of the fraction of debt allocated to the low-interest card on various measures of the interest rate gap. These regressions allow us to assess the percentage of variance explained by the interest rate gap, as well as the magnitude and statistical significance of the various coefficients. Column (1) of Table OA.2 shows the results from a linear-log specification with a full set of month dummies and individual fixed effects, which means that coefficients are identified using only within-individual time variation. The coefficient has the expected sign, but it is small in magnitude: a 100 percent increase in the interest rate gap is associated with a reallocation of 1 percent towards the low-interest card. We reach similar conclusions if we estimate the regression without individual fixed effects (not shown). Column (2) considers a more flexible specification and uses deciles of the interest rate spread and individual fixed effects. In this case, the variance explained, as reflected in the R^2 , is less than 0.01. Finally, Column (3) uses a double censored model to account for the fact that the share of debt on the low-interest card is a number between 0 and 1, and that the size of the credit line on the cheaper card may constrain this fraction below 1. In all of these cases the explained variance is close to zero and the estimated debt-interest elasticity is tiny

⁶The size of the interest gap is not random, of course. In the paper, we use random variation in the size of the interest rate gap and still find that consumers do not substitute balances across their cards.

at less than 2%.

Table OA.2: **Allocation of debt and the interest rate gap**

	Fraction of debt allocated to the low-interest card		
	(1)	(2)	(3)
$\log(1 + r_1 - r_2)$	-0.01** (0.00)		0.00 (0.00)
Range of β_i 's on interest rate spread deciles		[-0.01, 0.01]	
Censored regression			yes
Interest rate spread deciles		yes	
Monthly dummies	yes	yes	yes
Individual fixed-effects	yes	yes	
R-squared / Pseudo R-squared	0.00	0.00	0.00
Consumers	9,781	9,781	9,781
Observations	97,235	97,235	97,235

Notes: This table reports the results from regressions of the fraction of debt allocated to the low-interest card on the interest rate spread. Data are from the administrative database (2-card sample). Column (1) uses OLS and individual fixed-effects. Column (2) use OLS and individual fixed effects, but includes dummies for the deciles of the interest rate spread to allow for nonlinearities. The brackets show the range of coefficients for the decile dummies. Columns (3) use Tobit regression to account for the double censoring of the data. Standard errors clustered at the individual level are shown in parentheses. ***, **, and * indicate statistical significance at the 1, 5, and 10 percent confidence levels, respectively. We trimmed 1 percent of each tail of the interest gap distribution. R-squared is used for (1) and (2) and pseudo R-squared is used for (3).

(iii) Misallocation and demographics

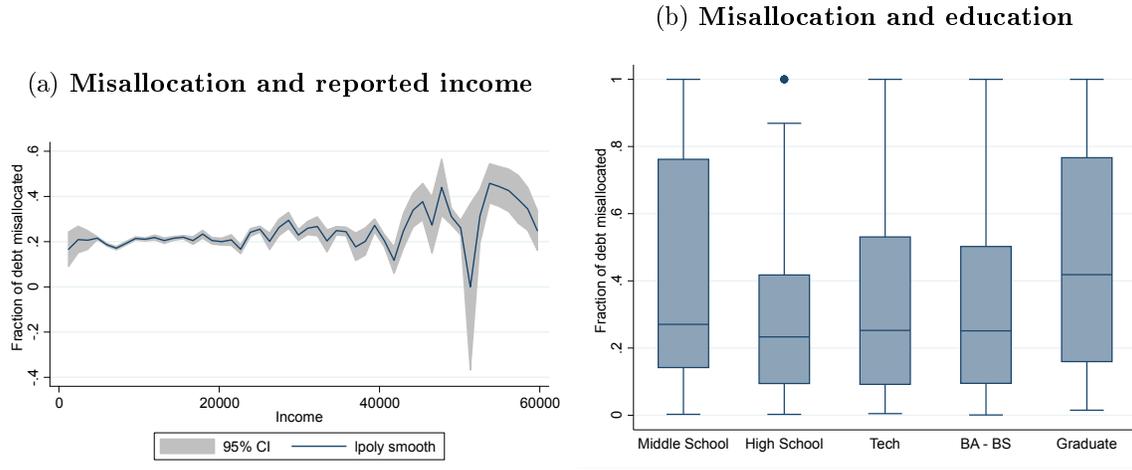
Next, we explore the relationship between misallocation and measured income and education using our administrative data set. This database contains information about educational attainment for individuals holding cards issued by one of our cooperating banks, and income for a subset of cardholders who reported it during the credit card application process. In Figure OA.9, we plot the relationship between the fraction of misallocated debt and these two variables.⁷ We use non-parametric regressions and box plots to allow for flexible relationships.

The relationships between the fraction of misallocated debt and income is flat. These relationship, however, should be interpreted with caution for two reasons, namely that income is self-reported during the credit card application process (which may have occurred a long time ago), and that the bank does not verify it (which implies that measurement error could be important).⁸ With regards to education, we find almost no change in the medians, although we observe a difference at the 75th percentile when we compare individuals with middle-school (about 1% of the sample) and graduate degrees (5% of the sample) against those with high-school degrees (about 6% of the sample).

⁷For readability, we right-trimmed income at the 97th percentile.

⁸In unreported graphs, we plotted the fraction of misallocated debt against income using our survey data and we also find no relationship

Figure OA.9: Misallocated debt and demographics



Notes: Point estimates and 95-percent confidence intervals plotted. Data are from the administrative database (2-card sample). An observation is a consumer. Panel A shows a kernel regression of the fraction of misallocated debt on the income reported by the individual using an Epanechnikov kernel. The figure is trimmed at the 97th percentile. Panel B shows box plots of the fraction of debt misallocated against educational attainment.

(iv) Cost of misallocation

In the paper, we define avoidable or extra pecuniary costs borne by consumers as the difference between the financing costs that a consumer actually incurs and the minimum costs necessary to finance her total debt. By this measure, we find that the average consumer leaves \$683 pesos per year on the table or 31% of her lowest-financing costs. To compute this cost, we take total debt, credit limits, interest rates, and overdraft fees as given (using a maximum overdraft of 20 percent, which is the policy that Mexican banks follow), and solve a linear cost-minimization problem treating the amount of debt allocated to the cheap card as the choice variable. In most cases, the solution to this problem is that individuals borrow on their low-interest card up to the credit limit and allocate the remaining debt to the high-interest card. We calculate all the costs on a monthly basis and convert them to annual figures. Figure OA.10 below shows the distribution of avoidable costs at the consumer level.

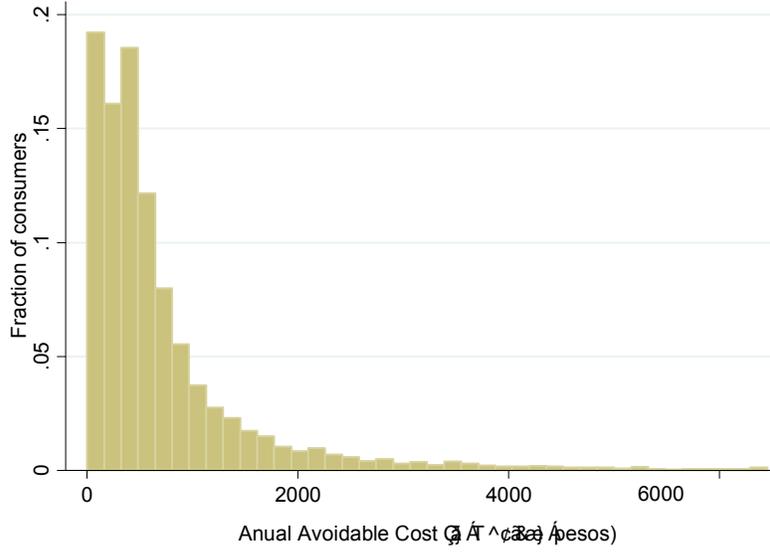


Figure OA.10: **Pecuniary cost of misallocation**

Notes: Data are from the administrative database (2-card sample). An observation is a consumer. The figure shows the annual interest costs in pesos of wrongly allocating debt to the high-interest card (in Mexican pesos). For each individual, this cost is calculated as the difference between the financing cost actually incurred and the minimum feasible interest cost necessary to finance her total debt. Costs are calculated on a monthly basis and converted to annual figures. The graph is trimmed at the 99th percentile.

As robustness, we also set up and solve a static cost-minimization model where consumers allocate their credit card purchases and payments (instead of debt) between their cards, taking into account interest rates as well as overdraft and late payment fees. To keep the model simple, we focus on the allocations on a period-by-period basis. As in the paper, we conceptualize the consumer's optimization problem in 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 payments across their two cards. Our model studies only the second stage. Specifically, consumers solve the following cost minimization problem:

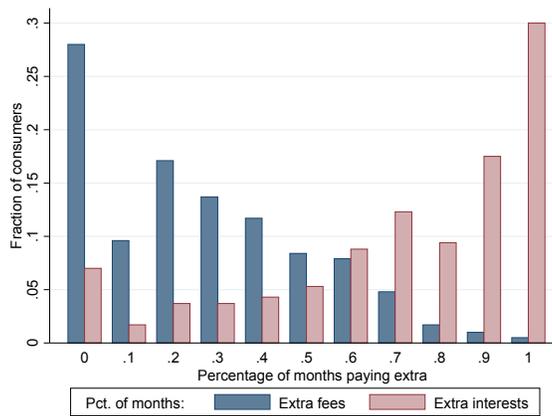
$$\begin{aligned}
 \mathcal{C}(P_{1t}, P_{2t}, X_{1t}, X_{2t}) \equiv & \\
 \min_{P_{1t}, P_{2t}, X_{1t}, X_{2t}} & \left\{ C_{1t} \cdot 1[P_{1t} < PMin_{it}] + F_{1t} \cdot 1[D_{1t-1} - P_{1t} + X_{1t} > L_{1t}] + \right. \\
 & [D_{1t-1} - P_{1t} + X_{1t}] \cdot 1[D_{1t-1} > P_{1t}] \cdot r_{1t} + \\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

subject to

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

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 overdraft fee charged if the outstanding balance exceeds the credit limit L_{it} , and D_{it-1} represents the outstanding balance at the end of the billing cycle. $1[\cdot]$ represents an indicator function. For each consumer, we search for the optimal purchases made with and payments made to card 1 in a {Purchase, Payment} grid of 10-peso increments. We use these allocations to estimate the optimal financing costs for each individual. We compare these costs with the actual financing costs faced by each cardholder and calculate the monetary costs of mistakes. This exercise yields similar results to those described above. We find that in 84 percent of consumer-months individuals pay extra costs due to their allocation of payments and purchases. On average, the extra interest and penalty fees incurred in this setting sum up to \$768 pesos per year, which is close to the figure we obtain when we use the stock rather than the flow of debt. We also find that most consumers pay extra interest frequently, but only occasionally incur in paying extra fees (Figure OA.11 below).

Figure OA.11: Fraction of months that consumers pay extra fees and interests



Notes: Data are from the administrative database (2-card sample). This figure shows the distribution of consumers by the frequency that they incur in extra interest and extra fees (x-axis). The figure is derived from the static model in which payments and purchases are allocated to minimize interest and fees. “Extra” means above the one that minimizes cost in our model.

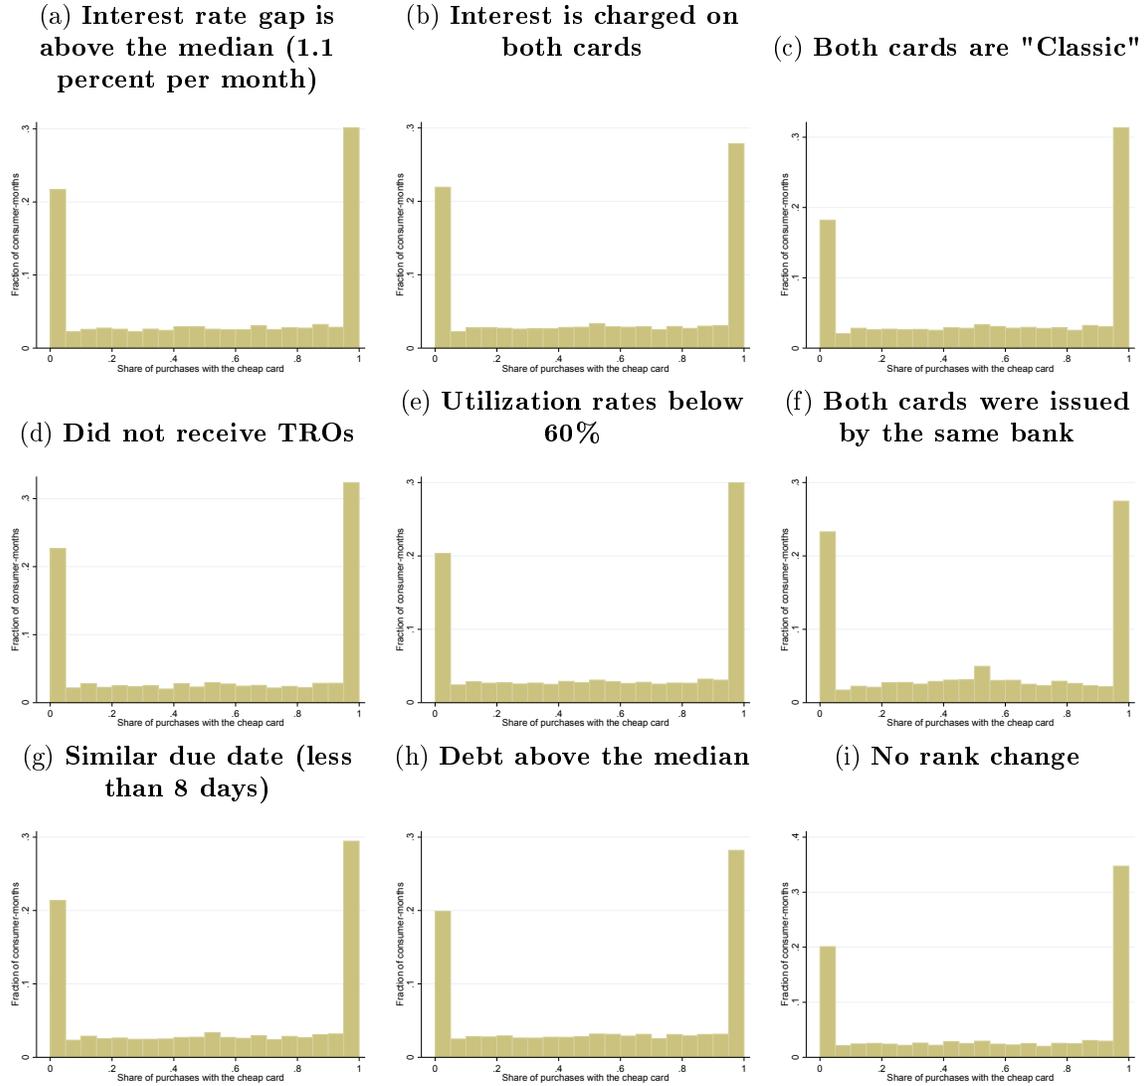
In the paper, we also present the results of a simulation of the distribution of avoidable costs that would arise in our data if consumers allocated their debt randomly. For this simulation, we compute the fraction of debt allocated to the low-interest card by drawing random values from a uniform distribution over $[0,1]$. In cases in which the drawn allocation is not feasible (because the debt on the low-interest card would exceed the credit limit), we impute an allocation in which the individual borrows up to the credit limit on the low-interest card and allocates the remaining debt to her high-interest card.

B. Allocation of purchases

(i) Robustness to various samples

For ease of exposition, in Section III.B of the paper we restricted our sample to observations in which individuals carried outstanding debt and paid at least the minimum due on both cards, made positive purchases, and made purchases that could fit on *either* of their two cards. This left us with a sample of 24,267 consumer-months and 2,680 consumers. In this sub-section, we report robustness checks of our results to more restrictive and less restrictive samples. First, in Figure [OA.12](#) we take the (highly trimmed) sample used in that Section and check the robustness of our estimates to sub-samples of the administrative data in which: (a) the difference in interest rates is above the median, (b) interest is charged on both cards as a result of having outstanding debt in the previous period, (c) both cards are Classic cards, (d) none of the two cards received a TRO during the sample period, (e) the utilization rate of both cards is below 60%, (f) both cards were issued by the same bank, (g) the two cards have a closing date differential of eight days or less, (h) the total amount of debt is above the median, and (i) where the price ranking of the two cards did not change during the sample period.

Figure OA.12: **Fraction of misallocated purchases in various sub-samples (Restricted sample of Section III.B)**

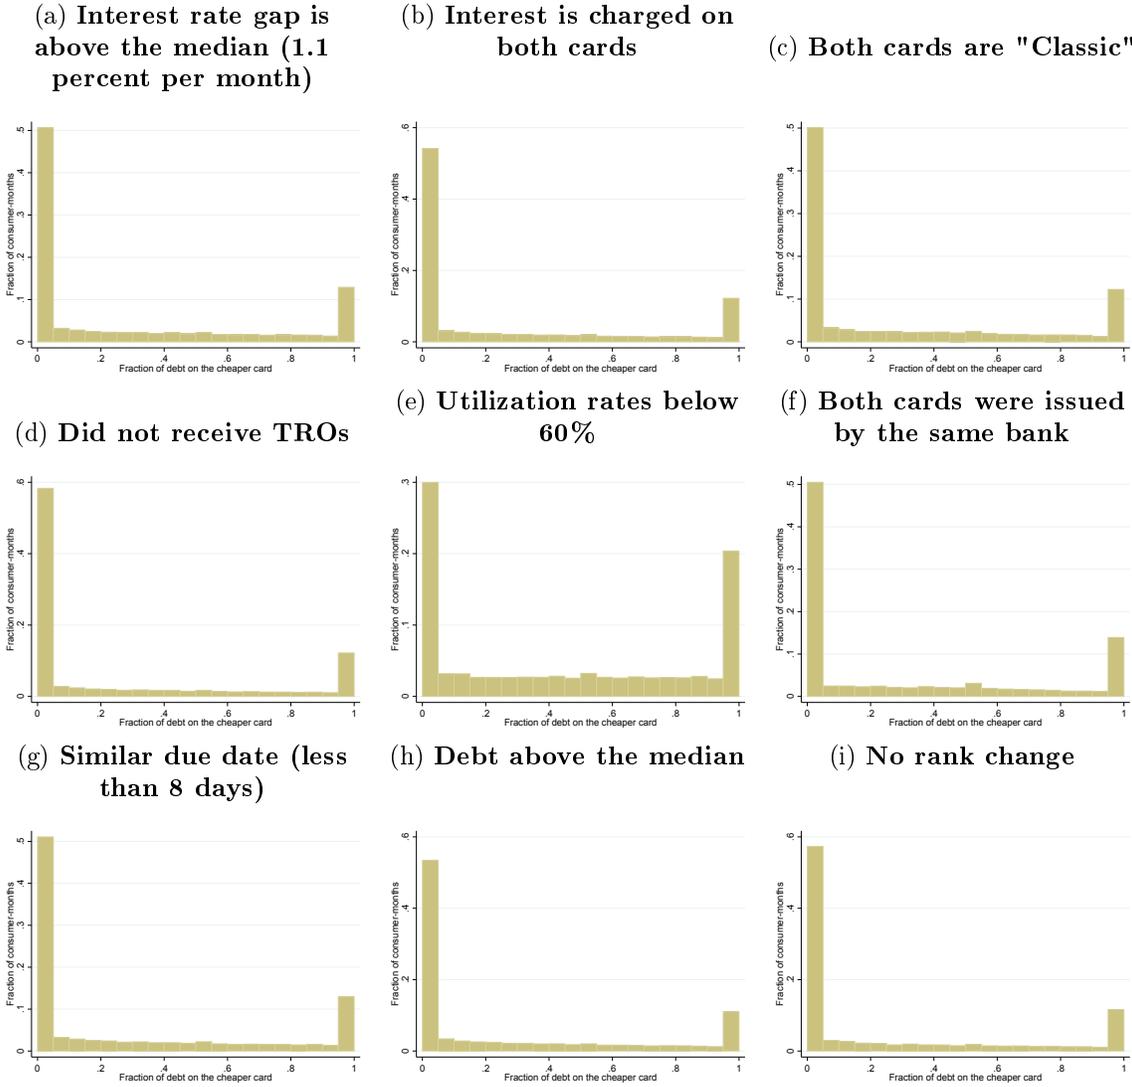


Notes: Data are from the administrative database (2-card sample), from the trimmed (24,267 consumer-months) sample used in the paper. An observation is a consumer-month. The Figure shows the frequency distribution of the fraction of credit card purchases that were allocated to the low-interest card in several sub-samples.

Next, we check whether our results hold when we use a less stringent criterion and analyze a sample in which the only restriction is that consumers made a purchase with one of their cards during that month (i.e. drop selection criteria (i) and (iii)). This leaves us with a sample of approximately 80,000 observations. Given that this sample includes observations for which the total amount of purchases may not fit on a card, we need to adjust our definition of misallocation. For this sample, we define the amount of misallocated purchases as the

purchases made with the high-interest card that could have been feasibly made with the low-interest card up to its credit limit. Figure OA.13 below shows the results of the fraction of misallocated purchases for several sub-samples within this expanded sample. We obtain distributions that are similar to those obtained when we examine the fraction of misallocated debt: on average, approximately one quarter of the total purchases are misallocated.

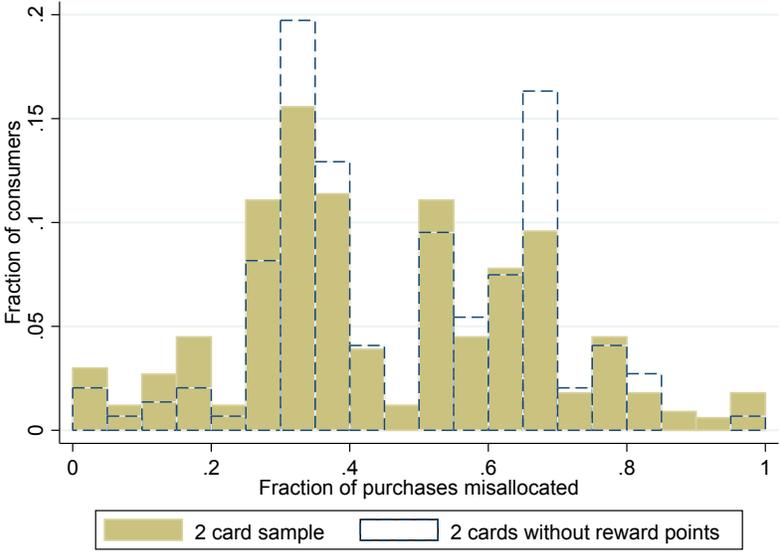
Figure OA.13: Fraction of credit card purchases that were MIS-allocated in various sub-samples (Complete sample)



Notes: Data are from the administrative database (2-card sample), from the untrimmed (80,000+) sample of consumer-months in which consumers made at least a purchase with one of their cards. The figure shows the frequency distribution of the fraction of misallocated credit card purchases in several sub-samples. In this case, the fraction of misallocated purchases is defined as the purchases made with the high-interest card that could have been feasibly made with the low-interest card up to its credit limit over the total amount of purchases made in that period.

Finally we show the distribution of the fraction of misallocated purchases in our survey. Figure OA.14 shows the results. The fraction of misallocated purchases here is calculated for each consumer as the high-interest purchases that could be shifted onto a lower-interest card up to the credit limit divided by the total credit card debt reported by the consumer. The amount of misallocated purchases is substantial in both the sample of cardholders with two cards, and in the sample of cardholders with two “comparable” cards (without rewards).

Figure OA.14: **Average fraction of misallocated purchases (by consumer) (Survey data)**

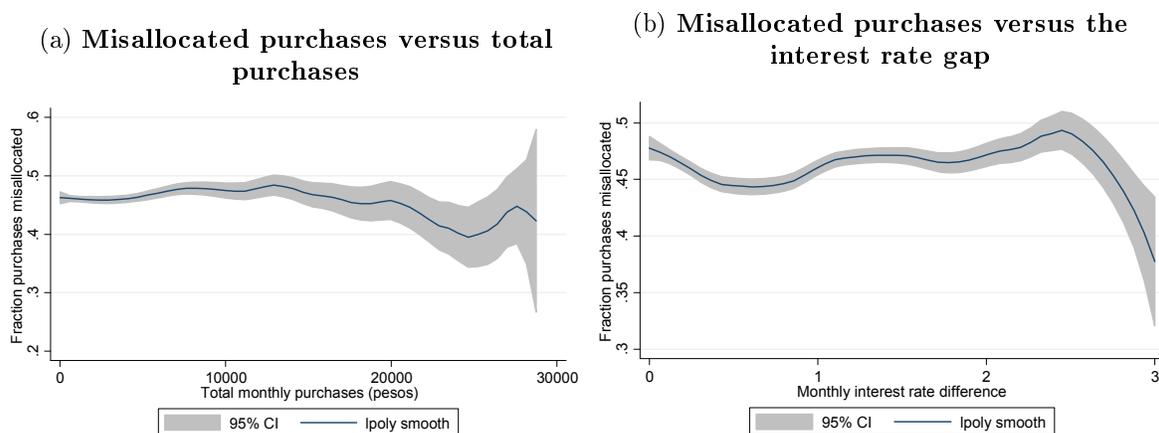


Notes: Data are from the survey database (2-card sample). An observation is a survey respondent. The figure shows the frequency distribution of the fraction of credit card purchases that were misallocated to the high-interest card. This fraction is calculated for each consumer as the high-interest purchases that could be shifted onto a lower-interest card up to the credit limit divided by the total credit card debt reported by the consumer. The solid line represents the sample of respondents with two credit cards (392 observations). The dotted line represents the sample of respondents with two credit cards without reward points (165 observations).

(ii) Robustness to “stakes”

As before, we examine graphically the relationship between the fraction of misallocated purchases and the “stakes” involved, measured by the difference in interest rates or by the total purchases made with both cards. In line with the results reported in the paper for debt, stakes matter little. Figure OA.15 plots non-parametric regressions of the fraction of misallocated purchases on the total monthly purchases in Mexican pesos (Panel A) and on the monthly interest rate difference (Panel B). The graphs are quite flat. Only for interest rate differentials above 3 percentage points per month does the fraction of misallocated purchases start to decline. The change, however, is small and there are few observations in this range.

Figure OA.15: **Fraction of purchases misallocated, total purchases, and interest rate gap (Kernel Regressions)**



Notes: Data are from the administrative database (2-card trimmed sample). An observation is a consumer-month. The Figure shows a non-parametric regression of the fraction of misallocated purchases on the average total monthly purchases in Mexican pesos (Panel A), and on the monthly interest rate difference, in percentage points using an Epanechnikov kernel. Point estimates and 95-percent confidence intervals plotted. Total purchases are trimmed at the 99th percentile. The interest rate gap is trimmed at 3 percentage points. The fraction of misallocated purchases is calculated as the purchases made with the high-interest card.

Regression results point to the same direction of no substitution across cards in the price margin. In Table OA.3, we replicate the analysis presented in Table OA.2 but use the fraction of misallocated purchases as the dependent variable and the less-restrictive sample of +80,000 observations (for consistency with the debt results). Column (1) shows that a doubling in the interest rate difference is associated with an increase of 5 points in the fraction of purchases allocated to the low-interest card. While the coefficient has the expected sign and is statistically significant, its magnitude is quite small. Column (2) and (3) uses deciles of the interest rate spread and individual fixed effects, and a double censored model, respectively. The results are qualitatively the same. In all cases, interest rate differences can explain at most 1 percent of the variance of the fraction of misallocated purchases.

Table OA.3: **Allocation of purchases and the interest rate gap**

	Fraction of purchases allocated to the low-interest card		
	(1)	(2)	(3)
$\log(1 + r_1 - r_2)$	0.05*** (0.00)		0.02*** (0.01)
Range of β_i 's on interest rate spread deciles		[-0.02, 0.07]	
Censored regression			yes
Interest rate spread deciles		yes	
Monthly dummies	yes	yes	yes
Individual fixed-effects	yes	yes	
R-squared	0.01	0.01	0.00
Consumers	9,574	9,574	9,574
Observations	80,298	80,298	80,298

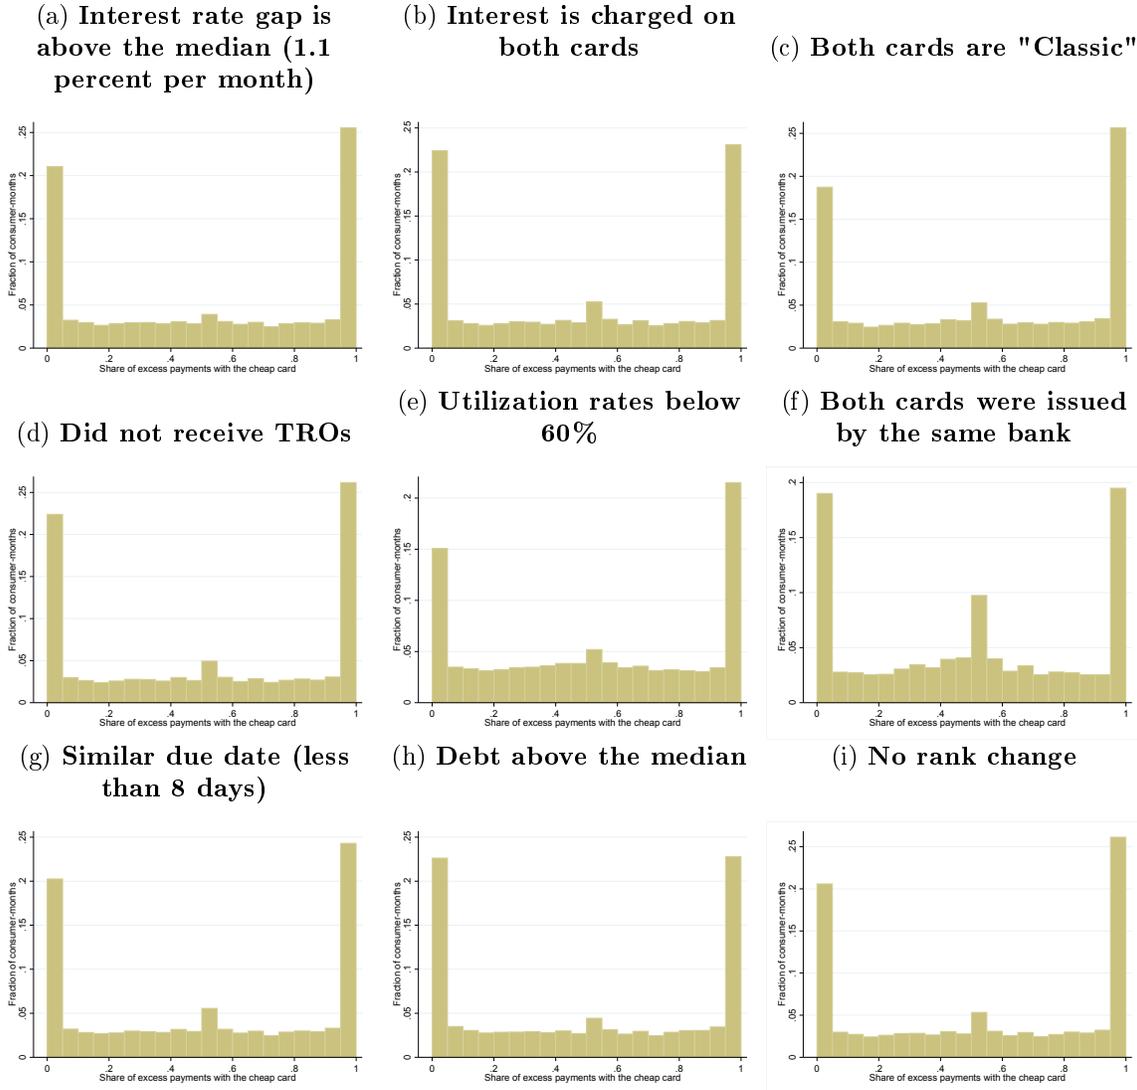
Notes: This table reports the results from regressions of the fraction of purchases allocated to the low-interest card on the interest rate spread. Data are from the administrative database (2-card less-restrictive sample (+80,000 observations)). Column (1) uses OLS and consumer fixed-effects. Column (2) uses OLS and consumer fixed effects, but includes dummies for the deciles of the interest rate spread to allow for nonlinearities. The brackets show the range of coefficients in the decile dummies. Columns (3) use Tobit regression to account for the double censoring of the data. Standard errors clustered at the individual level are shown in parentheses. ***, **, and * indicate statistical significance at the 1, 5, and 10 percent confidence levels, respectively. We trimmed 1 percent of each tail of the interest gap distribution. R-squared is used for (1) and (2) and pseudo R-squared is used for (3).

C. Allocation of payments

(i) Robustness to various samples

This sub-section presents robustness checks for our results on misallocated payments. Figure OA.16 below shows that our results in Section III.C of the paper are robust to 9 sub-samples of the administrative data, including samples in which the two cards had a due date differential of eight days or less, and in which both cards were issued by the same bank. The first of these sub-samples control for the fact that consumers may end up paying more to the low-interest card if it has an earlier due date and there is negative income shock in between the payment date of that card and the due date of the high-interest card. The second sub-sample controls for “ease to pay”, as in this case, the online payment platforms and the branch network are the same for the two cards.

Figure OA.16: Fraction of credit card payments that were wrongly allocated to the low-interest card in various sub-samples



Notes: Data are from the administrative database (2-card sample). The sample is the one used in Section III.C of the paper. An observation is a consumer-month. The figure shows the frequency distribution of the fraction of credit card payments that were misallocated to the low-interest card in several sub-samples. This fraction is calculated for each period, as the total payments above the required minimum allocated to the low-interest card divided by the total credit card payments made by the consumer during that month.

Next, we consider if our results continue to hold when we allow for the possibility that consumers minimize costs by missing the minimum due and paying a fee. This could happen, for example, if the outstanding balances, total payments, and interest rate differences are large enough that consumers prefer to be charged a late payment fee on their low-interest card to save on interest charges on their more expensive card. To estimate the cost minimizing allocations under this scenario, we solved the following problem, where consumers choose

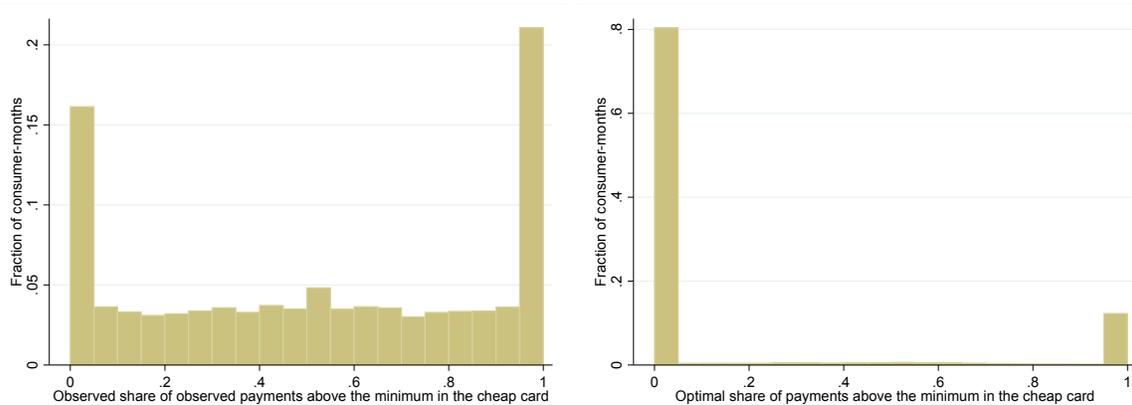
payments (P_1, P_2) such that:

$$\begin{aligned} \min_{P_1, P_2} \quad & r_1 \cdot \max\{(D_1 - P_1), 0\} + r_2 \cdot \max\{(D_2 - P_2), 0\} + \\ & \mathbb{1}(P_1 < MP_1) \cdot Fee_1 + \mathbb{1}(P_2 < MP_2) \cdot Fee_2 \\ \text{s.t.} \quad & 0 \leq P_i \leq \max\{D_i, MP_i\} \text{ for } i = 1, 2 \\ & P_1 + P_2 \leq P_{tot} \end{aligned}$$

where r_i represents the monthly interest rate of card i , D_i is the interest generating debt of card i , P_{tot} is the total amount of payments made during that period, MP_i is the required minimum payment for card i and Fee_i is the fee that the consumer has to pay if she fails to pay the minimum payment for card i . In this problem, we do not make any assumption about the relationship between D_i and MP_i and allow consumers to pay to each card more than the total amount of debt. Figure OA.17 shows the empirical distribution of the fraction of payments above the minimum allocated to the low-interest card together with the cost-minimizing distribution. As the Figure shows, the same qualitative results hold in this larger sample: actual payments are far from being cost minimizing. The figure is trimmed at $[0,1]$.

Figure OA.17: Empirical and cost-minimizing distributions of the fraction of payments above the minimum when consumers are allowed to minimize costs by missing the minimum due

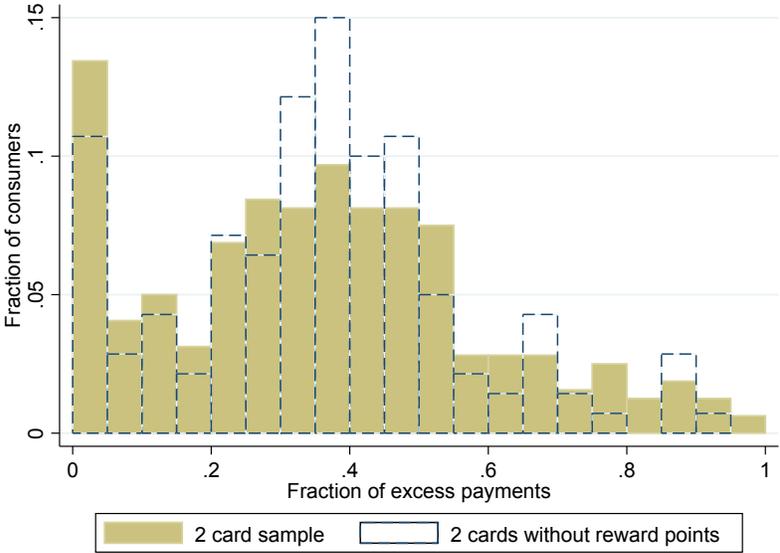
(a) **Observed allocation of payments above the minimum** (b) **Optimal allocation of payments above the minimum**



Notes Data are from the administrative database (2-card sample). The fraction of payments above the minimum (share of excess payments) is defined as the ratio between the payments made to the low-interest card above the required minimum and the difference between the payments made on both cards and the sum of the minimum due on both cards. Panel A plots the empirical distribution of the share of excess payments allocated to the low-interest card. Panel B plots the distribution of the cost-minimizing allocations of the share of excess payments. To have a cleaner looking graph, we trimmed 15 percent of the observations in panel A that were outside the $[0,1]$ range.

Figure OA.18 shows the the distribution of the fraction of misallocated payments above the minimum in our survey. Consistent with the other results, the average survey respondent allocates 36 percent of her payments above the minimum due to pay down the card with the lower interest rate.

Figure OA.18: **Average fraction of credit card payments above the minimum (by consumer) that were wrongly allocated to the low-interest card (Survey data)**

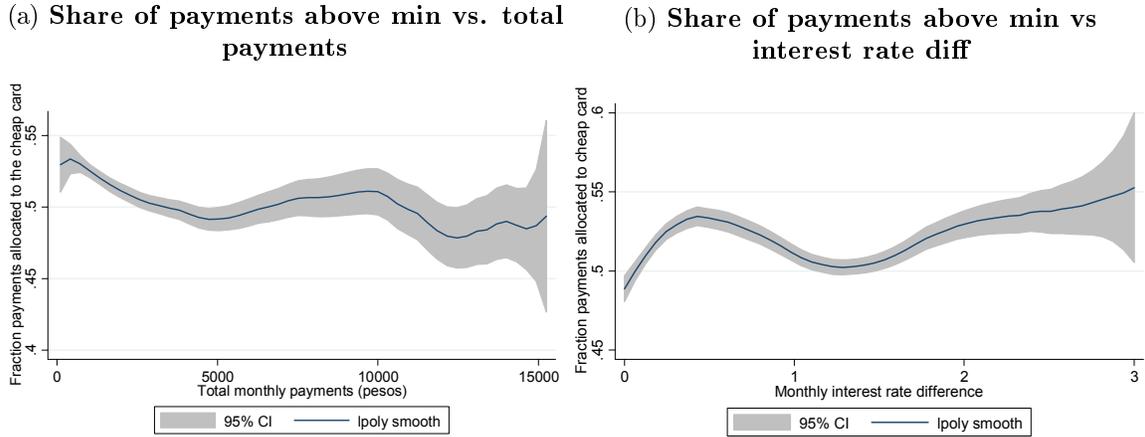


Notes: Data are from the survey database (2-card sample). An observation is a survey respondent. The figure shows the frequency distribution of the estimated fraction of credit card payments above the minimum due that were misallocated to the low-interest card. This fraction is calculated for each period, as the total payments above the required minimum allocated to the low-interest card divided by the difference between the reported payments made on both cards and the sum of the minimum due on both cards. The fraction is based on responses to questions asking consumers about the total payments made with each card during the month before the interview, the outstanding debt, the minimum payments, the standard interest rate on each card, and any outstanding offer. The solid line represents the sample of respondents with two credit cards. The dotted line represents the sample of respondents with two cards without reward points.

(ii) *Robustness to “stakes”*

Next, to evaluate whether consumers allocate their payments differently when the stakes involved are larger, we estimate non-parametric regressions of the fraction of misallocated payments above the minimum due on the interest rate gap and the amount of total payments. Figure OA.19b below shows the results. Consistent with our previous findings, the distributions of misallocated payments above the minimum do not vary with the interest rate gap or with the total amount paid.

Figure OA.19: **Fraction of payments misallocated above the minimum due, total payments, and interest rate gap (Kernel Regressions)**



Notes: Point estimates and 95-percent confidence intervals plotted. Data are from the administrative database (2-card sample). An observation is a consumer-month. The sample is the one used on Section III.C of the paper. The figure shows a kernel regression of the fraction of credit card payments above the minimum due that were misallocated to the low-interest card on the average total monthly payments (in Mexican pesos) (Panel A), and on the average monthly interest rate gap (in percentage points) (Panel B) using an Epanechnikov kernel. Total payments are trimmed at the 99th percentile. The interest rate gap is trimmed at 3 percentage points.

Following our analysis in the previous Sections, we present regression results of the relationship between the fraction of misallocated payments above the minimum due and different functions of the interest rate gap. We use the same specifications as before. Consistent with our results in this section, we find an economically small correlation between interest rate differences and the fraction of misallocated payments. Although the coefficients have the expected sign, the portion of explained variance is quite small.

Table OA.4: **Allocation of payments above the minimum due and the interest rate gap**

	Fraction of payments above the minimum to the low-interest card		
	(1)	(2)	(3)
$\log(1 + r_1 - r_2)$	-0.04*** (0.00)		-0.01** (0.00)
Range of β_i 's on interest rate spread deciles		[-0.09, 0.01]	
Censored regression			yes
Interest rate spread deciles		yes	
Monthly dummies	yes	yes	yes
Individual fixed-effects	yes	yes	
R-squared / Pseudo R-squared	0.00	0.01	0.00
Consumers	8,789	8,789	8,789
Observations	76,229	76,229	76,229

Notes: This table reports the results from regressions of the fraction of payments above the minimum due allocated to the low-interest card on the interest rate spread. Data are from the administrative database (2-card sample). Column (1) uses OLS and consumer fixed-effects. Column (2) uses OLS and consumer fixed effects, but includes dummies for the deciles of the interest rate spread to allow for nonlinearities. The brackets show the range of coefficients for the decile dummies. Columns (3) use Tobit regression to account for the double censoring of the data. Standard errors clustered at the individual level are shown in parentheses. ***, **, and * indicate statistical significance at the 1, 5, and 10 percent confidence levels, respectively. We trimmed 1 percent of each tail of the interest gap distribution. R-squared is used for (1) and (2) and pseudo R-squared is used for (3).

E. Own and cross price elasticities using random variation in interest rates

In Section III, we use experimental variation in interest rates to measure own and cross-price elasticities. For this, we rely on an overlap of our data with a large randomized experiment that one of our cooperating banks conducted in July 2005. The paper provides more details about the experiment and the results. Here, we focus on some properties of the experiment and the teaser rate offers, and check the robustness of our estimated elasticities to various sub-samples.

(ii) Targeting and summary statistics by treatment group

First, we show that (a) the experiment was indeed targeted to a broad (representative of the bank) non-delinquent population of cardholders, and (b) that in the intersection with our data, the treatment and control groups are statistically similar. These two features allow us to make causal inferences and cleanly estimate (short-term) interest rate elasticities.

According to bank officers, the experiment was targeted to the population of clients that were not delinquent at the time. To corroborate this, we estimate linear probability models

of the likelihood of receiving an experimental TRO on three monthly lags of utilization rates, purchases, payments, as well as dummy variables for whether the minimum payment was made, whether the outstanding balances were positive, and whether the total level of spending was positive. We use several specifications to allow for non-linearities. Table [OA.5](#) presents the results. Overall, we could explain 4% of the variance of the dependent variable, which suggests that the experimental teaser rate offers were targeted to a broad (unconditional) population of cardholders, as the bank claimed.

Table OA.5: **Determinants of Experimental TROs**

	1[Experimental TROs]
UR (log)	
L1.	-0.0005 (0.0007)
L2.	0.0005 (0.0009)
L3.	0.0001 (0.0007)
Purchases (log)	
L1.	0.003*** (0.0004)
L2.	-0.0006* (0.0003)
L3.	-0.0008 (0.0003)
Payments (log)	
L1.	-0.00002 (0.0001)
L2.	-0.00003 (0.0001)
L3.	-0.0006 (0.0001)
No interest incurred (dummy)	
L1.	-0.005*** (0.001)
L2.	-0.001 (0.001)
L3.	0.0007 (0.001)
Zero purchases (dummy)	
L1.	0.01*** (0.002)
L2.	-0.003* (0.002)
L3.	0.004* (0.002)
Zero payments (dummy)	
L1.	-0.0003 (0.0008)
L2.	-0.002*** (0.0006)
L3.	-0.001** (0.0006)
R-squared	0.04
Consumers	961
Observations	6,390

Notes: This table shows the estimates from linear probability models of indicators for receiving a randomized TRO (one-month or three-month) on time-varying card characteristics such as utilization rate, purchases, and payments; indicators for whether the card holder paid no interest, made no purchases, or made no payments; and their lags (L). Standard errors clustered at the individual level are given in parentheses. ***, **, and * indicate statistical significance at the 1, 5, and 10 percent confidence levels, respectively.

Next, we examine whether the main baseline characteristics are uncorrelated with treatment assignment. Table OA.6 below shows the means, for the treatment and control groups, of the most relevant variables during the three months prior to the experiment. It also reports the p-value of a t-test for differences in means. As the Table shows, we do not detect a sig-

nificant difference between treatment and control group in any of the six variables considered, which implies that the experiment seems valid in our sample. Another interesting aspect is that the means of these six variables for the cards that were part of the experiment are quite similar to the overall means for the cards in our main sample.

Table OA.6: **Randomization of teaser rate offers**

Variable at the time of the offer	3-month TRO	1-month TRO	Control	p-value (difference = 0)
	(1)	(2)	(3)	(4)
Credit limit	24,977	24,521	24,955	0.59
Monthly interest rate	0.028	0.028	0.028	0.60
Debt	9,974	9,956	10,061	0.56
Payments	2,383	2,428	2,431	0.57
Purchases	2,934	3,179	3,025	0.79
Fraction of misallocated debt	0.26	0.23	0.25	0.15
Observations	209	227	525	-

Notes: This table shows mean-comparison tests for the consumers that received the experimental teaser rate offers. All variables are measured at the time the offer was received. The table reports the p-values for the hypothesis that the difference between the mean of the treatment and the mean of the control is equal to zero.

(ii) Example of a teaser rate offer announcement

In this sub-section, we show the persuasive advertising accompanying a typical TRO, which makes the offers salient and therefore attention-drawing. Figure OA.20 below presents a letters that was mailed to cardholders congratulating them for their creditworthiness and informing them that their interest rate is reduced for a limited time period (i.e. that they had received a teaser rate offer). The TRO in these letters is recent. However, we were told by bank employees that the marketing materials announcing teaser rate offers have not changed much over the years, and were equally salient in the mid 2000s.

Figure OA.20: Example of teaser rate offer announcements



Notes: Shaded areas present information about the client. This is an example of a 0% offer for 12 months on purchases made between January 2015 and February 2015 from a Mexican bank. This offer is different from the offers we analyze in the paper. However, the saliency of the advertising accompanying the offer has not changed much over the years.

C. Robustness of our estimates

In Table OA.7 below, we examine the robustness of our results about the sensitivity to TROs to a different sub-sample and functional form. First, we address the concern that the *desired* own-response to a TRO may be understated because consumers may not be able to raise their debt above the credit limit and that the cross-response may be understated because debt on the (now) cheaper card may be already up to the limit. We estimate the same specification as in Table 2 of the paper, but conditioning on observations with less than 60% utilization in both cards in the month previous to the TRO. We obtain similar results: cross-price elasticities are zero and own price elasticities are large (although smaller than before) and statistically significant. Second, we redefine the dependent variable as the change in utilization rate in card 1 ($\Delta[debt/limit]_1$) instead of the amount of debt in pesos. We obtain qualitative similar results: large own-price elasticities (about an increase of 13 points in utilization for a 100 percent change in interest) and zero cross-price elasticities. Finally, to account for censoring in $\Delta[debt/limit]_1$, we estimate a Tobit model. Taking these three types of censoring into account does not change our qualitative conclusion: own-price elasticities

increase, but cross-price elasticities are still indistinguishable from zero.

Table OA.7: **Response to teaser rate offers (Robustness)**

	Δ Debt card 1 UR \leq 60% (OLS) (1)	Δ UR card 1 Full sample (OLS) (2)	Δ UR card 1 Full sample (Tobit) (3)
TRO $1_{t+2,Card1}$	137 (285)	-0.0038 (0.0103)	-0.0123 (0.0141)
TRO $1_{t+1,Card1}$	13 (260)	0.0130 (0.0144)	-0.0003 (0.0150)
TRO $1_{t,Card1}$	3,398*** (620)	0.1394*** (0.0166)	0.2082*** (0.0161)
TRO $1_{t-1,Card1}$	-2,991*** (722)	-0.0877*** (0.0164)	-0.0827*** (0.0166)
TRO $1_{t-2,Card1}$	-459 (459)	0.0029 (0.0112)	-0.0098 (0.0185)
TRO3 month 1 $t+2,Card1$	-8 (800)	0.0032 (0.0097)	0.0272 (0.0188)
TRO3 month 1 $t+1,Card1$	671 (533)	0.0160 (0.0127)	0.0400* (0.0240)
TRO3 month 1 $t,Card1$	1,002*** (302)	0.0480*** (0.0175)	0.0659*** (0.0141)
TRO3 month 2 $t,Card1$	973** (385)	0.0490*** (0.0117)	0.0748*** (0.0121)
TRO3 month 3 $t,Card1$	376 (476)	0.0361*** (0.0108)	0.0982*** (0.0174)
TRO3 month 3 $t-1,Card1$	216 (432)	-0.0094 (0.0098)	0.0412** (0.0176)
TRO3 month 3 $t-2,Card1$	-506 (354)	0.0065 (0.0108)	0.0338* (0.0193)
TRO $1_{t+2,Card2}$	-463 (324)	-0.0163* (0.0098)	0.0218 (0.0139)
TRO $1_{t+1,Card2}$	485* (285)	0.0149 (0.0101)	0.0260 (0.0198)
TRO $1_{t,Card2}$	90 (300)	0.0254** (0.0117)	0.0160 (0.0152)
TRO $1_{t-1,Card2}$	372 (253)	0.0286** (0.0113)	0.0179 (0.0152)
TRO $1_{t-2,Card2}$	-90 (333)	-0.0093 (0.0103)	0.0323 (0.0208)
TRO3 month 1 $t+2,Card2$	-237 (578)	0.0336** (0.0134)	0.0320 (0.0216)
TRO3 month 1 $t+1,Card2$	-336 (494)	-0.0024 (0.0109)	0.0350 (0.0259)
TRO3 month 1 $t,Card2$	298 (532)	-0.0047 (0.0104)	0.0358* (0.0200)
TRO3 month 2 $t,Card2$	306 (234)	0.0078 (0.0083)	0.0209 (0.0140)
TRO3 month 3 $t,Card2$	516 (349)	-0.0054 (0.0129)	0.0175 (0.0208)
TRO3 month 3 $t-1,Card2$	-266 (403)	0.0181 (0.0129)	0.0030 (0.0054)
TRO3 month 3 $t-2,Card2$	-348 (305)	-0.0262* (0.0138)	0.0144 (0.0171)
Adj / Pseudo R-squared	0.05	0.01	0.06
Consumers	380	961	961
Observations	2332	6391	6391

Notes: This table shows the effect of the one-month and three-month interest rate reductions on debt. Data are from the bank-run experiment (full sample of randomized offers). Each column represents a separate regression of the change in the interest-paying debt (Column 1) and the change in the utilization rate (Columns 2 and 3) on the treatment variable and its lags, controlling for changes in credit limit, indicator variables for non-random offers, its lags, and time dummies. Column (1) reports the results of the OLS regression for a sample of consumers with utilization rates of 60 percent or lower on both cards in the month before receiving the TRO. Column (2) shows the results of the OLS regression in which the dependent variable is the change in utilization instead of the change in debt. Column (3) takes into account censoring and presents the results of the Tobit regression of utilization on the same covariates as those in Column (2). Standard errors clustered at the individual level are given in parentheses.

F. Selection

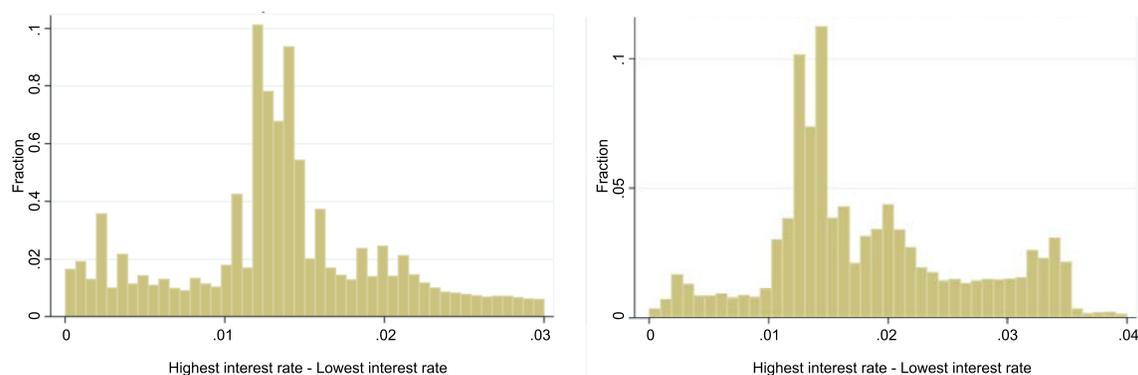
A possible concern with our analysis is that individuals holding two comparable cards might make systematically different allocation decisions from the rest of the cardholder population. In this sub-section, we examine the debt allocations and the responses to TROs of consumers holding three or four cards.

(i) Debt allocations

Before we start our analysis of debt allocations, we look at the distribution of the average difference in interest rates among the cards held by individuals holding three or four cards in our administrative data. As Figure OA.21 below shows, the distribution of the interest rate gap for individuals holding three or four cards is similar to the distribution shown in Figure OA.3 for individuals holding two comparable cards. The distribution for individuals holding more than two cards is smoother, probably because we are averaging over more cards, but the centrality of the distributions is alike.

Figure OA.21: **Monthly interest rate gap among credit cards for individuals holding more than two cardss**

(a) **Interest rate gap for individuals holding three credit cards** (b) **Interest rate gap for individuals holding four credit cards**

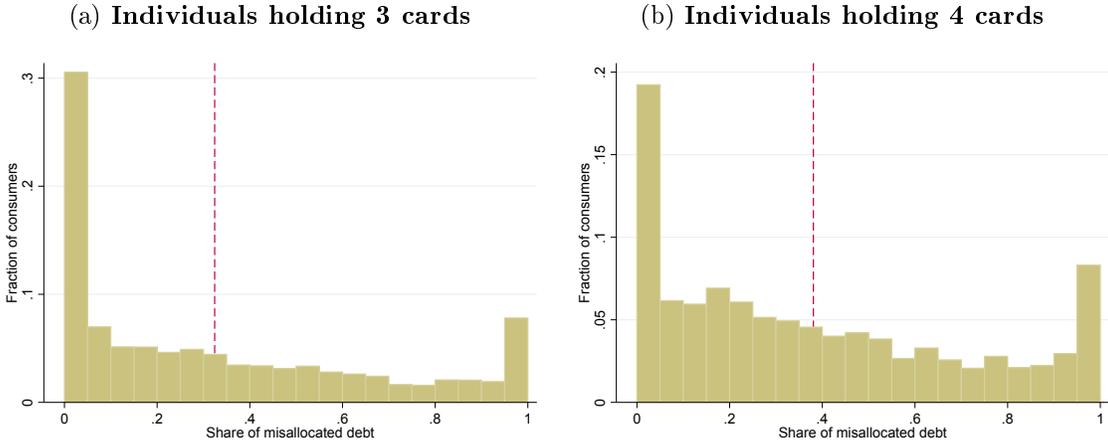


Notes: Data are from the administrative database (3 and 4-card samples). An observation is a consumer-month. The figure shows the frequency distribution of the average monthly interest rate gap (in percentage points) among the credit cards held by individuals with 3 (left-panel) or 4 credit cards (right-panel). For individuals holding 3 cards, the average is calculated as $(|r1 - r2| + |r1 - r3| + |r2 - r3|)/3$. For those holding 4 cards, the average is calculated as $(|r1 - r2| + |r1 - r3| + |r1 - r4| + |r2 - r3| + |r2 - r4| + |r3 - r4|)/6$.

Next, we look at the distribution of the fraction of misallocated debt for individuals holding three or four cards. To calculate the interest minimizing allocation, we assign debt sequentially to the lowest interest card up to its credit limit (first to the lowest card, then to the second lowest, and so on). Then, we define for each consumer-month, the amount of misallocated debt as the debt that pays a higher interest than this optimal scenario. To obtain the fraction

of misallocated debt we divide this figure by the total credit card debt held by a consumer that month. We calculate a consumer-level measure by taking the average over months. Figure OA.22 shows the two distributions at the consumer level. The median fraction of misallocated debt for individuals carrying three cards is 33%, and 38% for those carrying four cards.

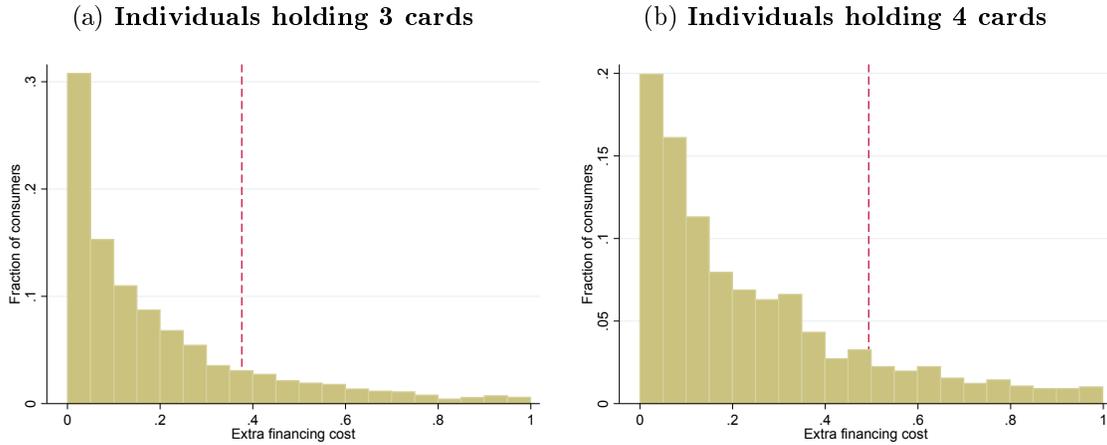
Figure OA.22: **Average fraction of credit card debt (by consumer) that was wrongly allocated to the high-interest card by individuals with 3 or 4 cards**



Notes: Data are from the administrative database (3 and 4-card sample). An observation is a consumer. The figure shows the frequency distribution of the average fraction of credit card debt that was wrongly allocated ("misallocated") to the high-interest card by individuals holding 3 (Panel A) and 4 cards (Panel B). The average is taken over the months with information for each consumer. Histograms in Panel A and B are trimmed at the 95th percentile.

Next, we look at the extra (or avoidable) costs borne by consumers holding three and four cards. This cost is calculated as the difference between the financing costs (interest costs) that a consumer actually incurs and the minimum feasible costs necessary to finance her total debt. As Figure OA.23 reveals, the magnitudes of the extra financing cost are higher when we look at consumers who have more than two cards. The average extra cost incurred as a percentage of the minimum cost is 31% for consumer holding two cards, 38% for those holding three cards (mean = \$1,339 pesos), and 49% for those holding four cards (mean = \$2,406 pesos).

Figure OA.23: Cost of misallocation for individuals holding 3 or 4 cards



Notes: Data are from the administrative database (3, and 4-card sample). An observation is a consumer. The figure shows the annual interest costs of wrongly allocating debt to the high-interest cards (in Mexican pesos) as percentage of the minimum cost. For each individual, the avoidable cost is calculated as the difference between the financing cost actually incurred and the minimum interest cost necessary to finance the individual's total debt. Costs are calculated on a monthly basis and converted to annual figures. Histograms are trimmed at 100 percent of the extra financing cost.

(ii) Responses to TROs

In this sub-section, we carry out the same exercise as in Section III.E of the paper and estimate own-card elasticities for consumers holding one, two, three or four cards. To estimate these elasticities, we first drew randomly four groups of 2,500 consumers holding 1, 2, 3, or 4 cards (for a total of 10,000 borrowers) and selected one card at random (for a total of 147,776 consumer-months). Using this sample, we then estimate a regression of the change in debt on an interaction of the interest rate of the card picked at random and the number of cards held by the consumer, time dummies, and consumer fixed-effects (to control for time-invariant unobserved differences across accounts). Specifically, we estimate the following specification: $\ln(debt_{i,t} + 1) = \alpha_i + \gamma_t + \beta \ln(r_{i,t-1} + 1) * NumberCC + \nu_{i,t}$ using OLS. Admittedly, since we are not using an experiment, our estimates could be biased, which hinders our ability to interpret β as an elasticity. In this exercise, however, we are not particularly interested in estimating the elasticities themselves, but rather the differences in these elasticities among consumers holding different number of cards. As long as the bias is not different across these groups, our exercise is informative. Table OA.8 presents the results. The estimated elasticities for individuals holding 1, 2, 3, and 4 cards are -0.41 , -0.43 , -0.43 , and -0.38 respectively, which are close to those we estimate in the paper using random variation. The F-test of the hypothesis that all elasticities are equal cannot be rejected, suggesting that consumers who hold more than two cards are not less sensitive to changes in interest rates than those who hold one, three, or four cards.

Table OA.8: **Debt-interest elasticity estimates by the number of credit cards**

Dep Var: $\ln(\text{Debt} + 1)_t$	Number of credit cards			
	One (1)	Two (2)	Three (3)	Four (4)
$\ln(\text{interest rate} + 1)_{t-1}$	-0.41 (0.14)	-0.43 (0.15)	-0.43 (0.14)	-0.38 (0.16)
Month dummies		Yes		
Consumer dummies		Yes		
R-squared		0.04		

Notes: Data are from the administrative database. The sample was constructed by drawing randomly four groups of 2,500 consumers holding 1, 2, 3, or 4 cards (for a total of 10,000 borrowers) and selecting one of their cards at random (for a total of 147,776 consumer-months). An observation is a consumer-month. The table presents estimates of debt-interest elasticities $\hat{\beta}$ using the following fixed effects regression: $\ln(\text{debt}_{i,t} + 1) = \alpha_i + \gamma_t + \beta \ln(r_{i,t-1} + 1) * \text{NumberCC} + \nu_{i,t}$. The regression was estimated using OLS. To interpret this as debt responses we have to assume that the previous month interest rate is exogenous conditional on the card fixed effect and the time dummies. Standard errors clustered at the individual level are given in parentheses.

IV Potential explanations

In this section, we consider some explanations for our five stylized facts and bring new data and modeling. Our intention is not to develop an all-encompassing model, but a series of models and model approaches that provide important information about the main neoclassical explanations.

IV.A Switching costs and stochastic interest rates

Consumers may not transfer balances to their low-interest card if this is costly and there is uncertainty about the future interest rate on each card. Depending on the stochastic properties of the interest rates and the size of the switching cost, these two features of the environment could potentially explain why consumers allocate their balances to their more expensive card. In this sub-section, we calibrate a dynamic stochastic model that captures this tradeoff parsimoniously.

In our model –described in equation (2)– consumers observe current interest rates r_{1t}, r_{2t} , as well as total debt and credit limits (\bar{D}, L_1, L_2) , and choose how to allocate their current debt between their two cards (D_{1t}, D_{2t}) in order to minimize their current and *expected* future financing cost, inclusive of the debt switching cost θ .⁹ Consumers incur the switching cost only when they change the allocation of their debt. This makes the problem dynamic and generates an optimal inactivity range that may explain low substitution.¹⁰

$$\begin{aligned}
 V_t(D_{1,t-1}, D_{2,t-1}) = & \min_{D_{1,t}, D_{2,t}} r_{1t}D_{1,t} + r_{2t}D_{2,t} + \theta * \mathbb{1}_{\{D_{1,t-1} \neq D_{1,t}\}} + \beta \mathbb{E}[V_{t+1}(D_{1,t}, D_{2,t})] \\
 \text{s.t. } & D_{1,t} + D_{2,t} = \bar{D}, \quad D_{1,t} \leq L_1, \quad D_{2,t} \leq L_2
 \end{aligned} \tag{2}$$

We take the following information directly from the data: total debt (\bar{D}_i) , the initial time zero allocation of debt across cards D_{i0} , the credit limits at time zero $(L_{1,0}, L_{2,0})$, and the realized interest rates (r_{1t}, r_{2t}) each month.¹¹ We also take expectations with respect

⁹We think of this cost as a reduced form of all the costs to reallocate debt across cards. It includes, for instance, not only the direct balance transfer fee, but also the time and effort cost of thinking which card to use or pay. For simplicity, we ignore the variable cost component in this decision. In our view, the variable cost is small. Nevertheless, we experimented with a variable cost of 5 % and 10% of the transferred debt and obtained similar results.

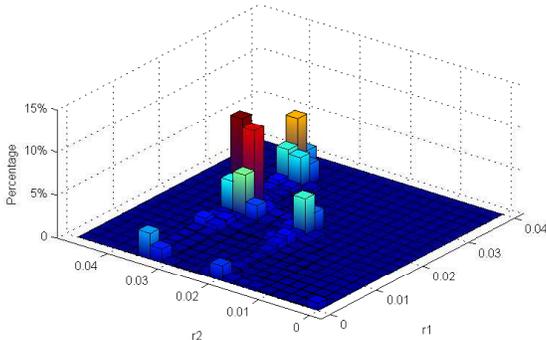
¹⁰The objective function assumes risk neutrality. We believe this is a good approximation for credit card debt as what is at stake is not a large part of consumers' wealth. However, we experimented with a convex function $g(x) = -x^2$ –where x represent the current part of the value function– and got similar results.

¹¹To save on computational time, we assume that total debt and credit limits are constant over time. An alternative is to use actual total debt *each* period and calculate the optimal allocation: i.e. for each month t take \bar{D}_{it} from the data and solve the problem with $T-j$ periods remaining. This has two complications: the first is that it is more intensive in computing time because the max total debt in the grid of debt changes period by period. The second is that we need to make an assumption about how agents form expectations regarding the evolution of \bar{D}_{it} , for instance: assuming myopic no debt growth, perfect foresight of debt, or an expected value on a distribution of debt increases calibrated from the population distribution. We ran an exercise for 500 randomly chosen individuals and 10 months assuming perfect foresight and we find that results do not change significantly.

to the empirical distribution of interest rates in our data, and assume a Markov process for the distribution of the interest rate pair r_{1t}, r_{2t} , using a 9×9 transition matrix (described below). Figure OA.3 above shows that the distribution of interest rate gap has three salient modes. We use these modes to discretize the transition matrix $T(r_1, r_2)$ into 9 states from the 3×3 combinations of (r_1, r_2) . The combinations for the model were taken from $(r_1, r_2) \in \{[0, 0.66], (0.66, 1.99), (1.99, 3.33]\}^2$. The transition matrix is shown in Figure OA.24.¹² For robustness, we experimented with finer transition matrices using 20 equally spaced states and a transition matrix that depended on the type of card or bank and got similar results.

Figure OA.24: **Parameters for the switching cost model with stochastic interest rates**

(a) **Density function plot**



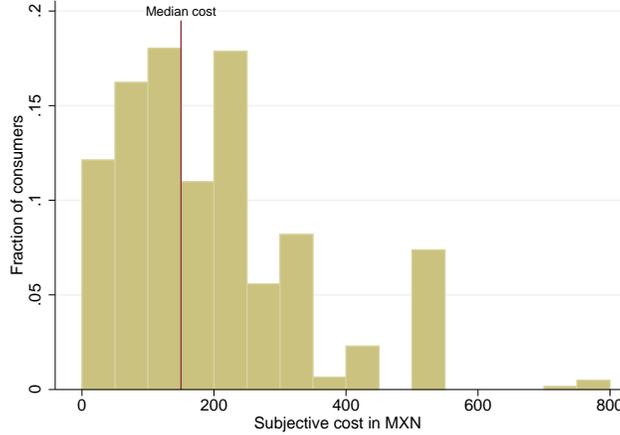
(b) **Transition matrix**

	(0.66, 0.66)	(0.66, 1.99)	(0.66, 3.33)	(1.99, 0.66)	(1.99, 1.99)	(1.99, 3.33)	(3.33, 0.66)	(3.33, 1.99)	(3.33, 3.33)
(0.66, 0.66)	0.71	0.05	0.06	0.05	0.03	0.01	0.05	0.01	0.02
(0.66, 1.99)	0.02	0.71	0.05	0.00	0.11	0.01	0.01	0.08	0.01
(0.66, 3.33)	0.02	0.03	0.74	0.00	0.00	0.10	0.01	0.00	0.09
(1.99, 0.66)	0.02	0.00	0.00	0.70	0.11	0.09	0.06	0.01	0.01
(1.99, 1.99)	0.00	0.03	0.00	0.02	0.78	0.08	0.00	0.07	0.01
(1.99, 3.33)	0.00	0.00	0.02	0.01	0.03	0.88	0.00	0.00	0.05
(3.33, 0.66)	0.02	0.00	0.01	0.03	0.00	0.00	0.74	0.11	0.09
(3.33, 1.99)	0.00	0.01	0.00	0.00	0.04	0.00	0.02	0.87	0.05
(3.33, 3.33)	0.00	0.00	0.02	0.00	0.00	0.04	0.02	0.04	0.88

We note that if θ is sufficiently small, or if r_1 and r_2 are i.i.d, or if interest rates have enough inertia, then the problem above is close to the static version we calculated in Section III of the paper. Since we do not have data on the balance transfer cost, we experimented with $\theta \in \{100, 200, 300, 400, 500, 1000, 2000, 4000\}$. This cost range is on the high end since, in our survey, the objective median value of θ is 150 pesos including non-pecuniary costs, and the subjective estimate for those who have not transferred balances is approximately \$200 pesos (see Figure OA.25 below). Obviously, for a sufficiently high θ , there will be no reallocation at all.

¹²Note that the diagonal shows substantial inertia with values between 70% and 80%. Interest rates in our sample are pretty stable across time, with a correlation of 0.84, and a median of $|r_t - r_{t-1}| = 0.1\%$.

Figure OA.25: **Objective and subjective costs of transferring balances across cards as reported by survey respondents**



Notes: Data are from the survey database (2-card sample). The histogram shows the expected subjective cost for respondents who had not transferred balances across cards and had at least two credit cards. It computes the answer for the question: “Taking into consideration the bank fees and the cost of both your time and effort, how much do you think would it cost to transfer debt across cards?” The subjective expected cost histogram is computed with 610 observations and it is trimmed at the 95th percentile. The red line represents the realized cost for people who *did* transfer balances across their cards. This line shows the median for the realized cost, which is computed using the sum of the answers from the questions: *How much did it cost to transfer balances across cards?* and *If you had to assign a monetary value to the time and effort that you invested in transferring balances, how much would that be?* The median cost is computed with 35 observations of respondents who had transferred balances across their cards.

We simulate the distribution of the optimal allocation of debt for 30 months by backward induction, and compare it with the empirical distribution of debt allocations. The top panels of Figure OA.26 plot the distributions of the fraction of debt allocated to the low-interest card, as given by the model and the actual data, using switching costs of \$200 pesos and of \$2,000 pesos.¹³ In both cases, the Kolmogorov-Smirnov test rejects the hypothesis of equality of distributions with a ($p < 0.001$). When the switching cost is \$200 pesos –in line with our survey results– our model predicts that on average, consumers allocate a larger fraction of their debt to the low interest card than in the actual data. In the model about 33% of observations allocate *all* of their debt to their cheap card, while in the actual data only 11% of observations do. In fact, the model implies a distribution that is quite close to that in Figure 1 from the paper, suggesting that the simple static model is a good approximation to this more complex dynamic stochastic model.

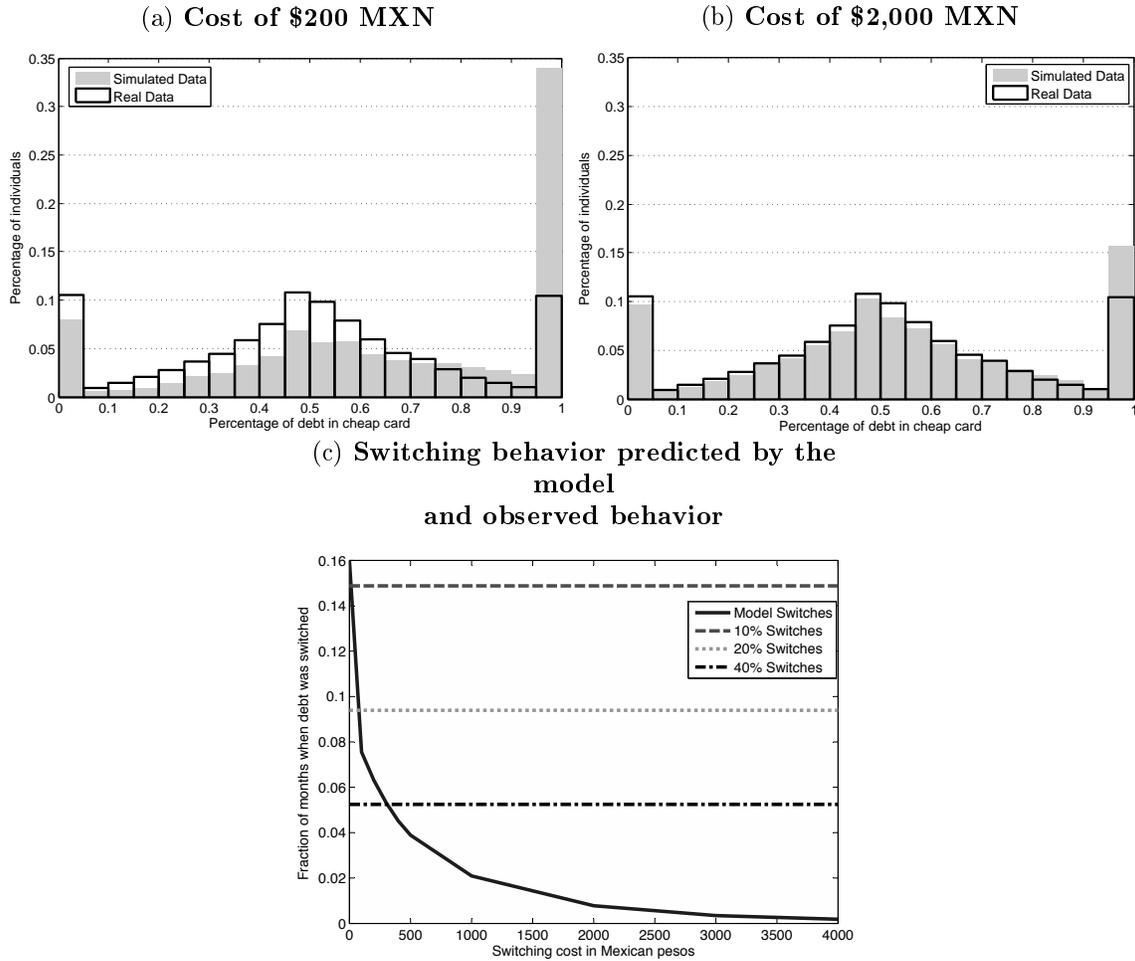
According to the model, to rationalize the actual distribution of the fraction of debt allocated to the low-interest card we would need a fixed cost close to or above \$2,000 pesos. This amounts to 10% of the average total debt on both cards. In our view, this cost is too

¹³Since we do not have data on this cost, we experiment with values of \$100, \$200, \$300, \$400, \$500, \$1000, \$2000, and \$4000 pesos. We pick \$200 pesos and \$2,000 because these values are close to the median and the 99th percentile of the distribution of the subjective expected costs of transferring debt across cards –inclusive of effort and pecuniary costs– reported by survey respondents (shown above).

large and inconsistent with the survey answers. More importantly, a switching cost of \$2,000 pesos implies that cardholders should change their preferred allocation of debt rarely, which cannot be reconciled with the amount of switching –changes in the share of debt allocated to each card– we observe in the data. Panel C of Figure 7 plots the number of times per year that the average consumer switches her debt from one card to another in the model and in the data for different switching costs and different thresholds for what we consider a switch.¹⁴ In the data, consumers change the allocation of their debt (“switch”) much more than in the simulated data, and they switch to the low-interest card as often as to the high-interest card. This is hard to square with an environment with high switching costs. For a cost of \$2,000 pesos, the model substantially underestimates the amount of switching in the data by a factor of about 20 times. We conclude that stochastic interest rates and switching costs, as modeled, do not explain the low substitution in our data.

¹⁴In the model it is very clear when there is a switch, but in the data it is trickier to identify switching, since there might be changes in the share of debt allocated to a particular card due to changes in total debt, fees, etc. that are not literally debt transfers. We experiment with several thresholds of what constitutes a switch. We define a change in the fraction of debt allocated to a given card as a switch when it differs from the share of the previous month by 10%, 20%, or 40%; these correspond to the horizontal lines in the Figure.

Figure OA.26: Allocations predicted by the switching cost model



Notes: Data are from the administrative database (2-card sample). An observation is a consumer-month. The figure compares the empirical distribution of the fraction of credit card debt allocated to the low-interest card (red line) to the frequency distribution predicted by the switching cost model (blue line). Panel A assumes a switching cost of \$200. Panel B uses a switching cost of \$2,000 pesos. Panel C shows the predicted number of times per year that a consumer would transfer her credit card debt from one card to another for different switching costs (θ , in pesos) (solid line). For comparison, Panel C also shows the average number of times per year that consumers *switch* their debt from one card to another for different definitions of switching (dotted lines). In these figures, a *switch* is defined as a change of 10 percent or more (20 percent or more; 40 percent or more) in the fraction of debt allocated to a random card (labeled card 1) from one month to the next. The lines from the observe data are horizontal since we don't observe switching cost.

IV.B Indivisible purchases

Consumers could minimize costs by making purchases with their high-interest card today in order to keep available credit to accommodate an indivisible large expense on the low-

interest card in the future (explaining Facts 1 and 2). In reality consumers in Mexico can split a purchase over multiple cards at the counter, which means this explanation is not really operational. The raw data provides evidence suggesting indivisibilities are not important in explaining Facts 1 and 2. In our data large purchases are uncommon, the 90th percentile of purchases made *with both cards* is 8,267 pesos and the 75th is 5,140 pesos. Most purchases are small relative to the available credit limit: the median purchase is 1,180 pesos and the median available credit after purchases is 6,180 pesos. Furthermore, only 7% of observations that report purchases with the expensive card also display large purchases (>90th percentile) in the next 3 months on the cheap card, making it unlikely that consumers were making room on such a card for a large purchase. Direct survey responses also point to the irrelevance of indivisibilities. The results in Table 3 of the paper show some interesting statistics in this regard. When asked, “So far this year, how many times have you used other credit cards instead of your [name of the cheapest card] card because you wanted to leave enough room on it to make a major purchase?”, the mean answer is 0.4 times, with more than 80 percent of respondents reporting zero instances.

To be more formal, set up a simple two-period model to illustrate the basic trade-off of purchase indivisibilities. The model has one period divided in two parts. At the beginning of the period, interest rates are observed and consumers allocate their debt among their two cards to accommodate a potential future indivisible purchase (we set up the model assuming without loss of generality that card 1 is the low-interest card). At the end of the period, consumers face the necessity to make an indivisible expense of size C^I with probability p .¹⁵ The cost of that good follows a parametric probability distribution $f(C^I|c > \bar{C})$, estimated from the data.

Every period t , the consumer faces a trade-off between purchasing today with the high-interest card to free space on the low-interest card in case they have to make an indivisible expense in the future with positive probability. The model assumes that, if consumers had to make an indivisible purchase C^I of a size that would fit on the low-interest card, they would automatically allocate it to that card. Otherwise, they allocate it to the more expensive card.¹⁶ Specifically, consumers solve the following problem:

$$\begin{aligned} \min_{D_1, D_2} \quad & r_1 D_1 + r_2 D_2 + r_1 \mathbb{E}[C^I | C \leq L_1 - D_1] \mathbb{P}(C^I \leq L_1 - D_1) \\ & + r_2 \mathbb{E}[C^I | C > L_1 - D_1] \mathbb{P}(C^I > L_1 - D_1) \\ \text{s.t.} \quad & D_1 + D_2 = \bar{D}, \quad D_1 \leq L_1, \end{aligned} \tag{3}$$

The expectation in equation (3) is taken with respect to the mixed distribution $f_{C^I, p}(C) = (1 - p) * 0 + p * f(C^I | C^I > \bar{C})$. To simulate the model, we take (r_{1t}, L_{1t}, r_{2t}) from the data.¹⁷ Since we only observe the sum of monthly purchases, we have to take a stand on what counts as an indivisible purchase, and how its distribution $f(C^I | c > \bar{C})$ looks like. We approximate the distribution of indivisible purchases by the empirical distribution of the total monthly

¹⁵Thinking of p in the model as the probability of making an indivisible expense in the future helps bridge the gap between the multi-period and the one period models.

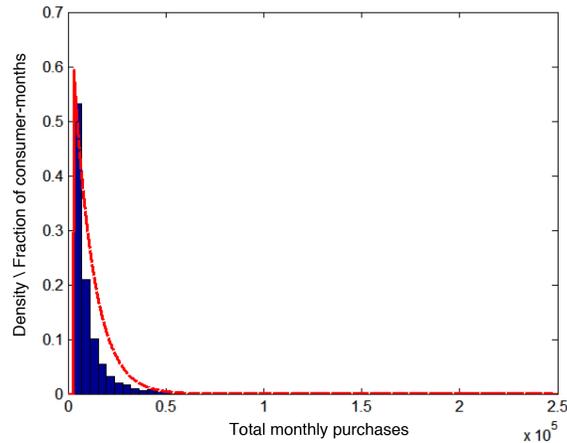
¹⁶By forcing consumers to make the expense, we may be exaggerating the financing cost of indivisibilities. The same is true about our assumption that C^I can always be allocated to the more expensive card.

¹⁷We take credit limits and interest rates from the data as given. Since interest rates do not seem to depend on changes in debt (Figure OA.32), we think it is reasonable to assume that interest rates are exogenous to debt allocations.

purchases truncated from the left at \bar{C} , and experiment with several cutoffs of \bar{C} . The idea is to capture “large” purchases.

To simulate the model, we need to compute the expectations in equation 3. For this, we obtain the truncated empirical distribution from the data and approximate it with a parametric distribution, selecting it as the one that maximized the log-likelihood among a large family of distributions.¹⁸ We obtained the best fit by using the Generalized Pareto distribution. Figure OA.27 shows the fit graphically.

Figure OA.27: **Fit of the Generalized Pareto Distribution**



Notes: Data are from the administrative database (2-card sample). The bars show the observed distribution of total monthly purchases in the administrative database. The red (dotted) line shows the Generalized Pareto distribution that was fitted for the optimization problem.

Intuitively and very informally — given (p, r_1, r_2, L_1) — indivisible purchases would matter if they are small enough to fit on the cheap card, whose credit limit is L_1 , but big enough to outweigh the cost of incurring in r_2 instead of r_1 today. The model shows that this scenario is not common in the data: Figure OA.27 below shows the simulated and empirical distributions of the fraction of debt allocated to the low-interest card for a probability of one and an minimum indivisible purchase of \$3,000 pesos. The difference between the empirical distribution and the simulated distribution is noteworthy, suggesting that this simple model is unable to rationalize Fact 1. As a test of robustness, we simulated the model using different thresholds of truncation $\bar{C} \in \{3000, 5000, 10000, 20000\}$ and different probabilities $p \in \{1/3, 1/4, 1/6, 1\}$, or assuming heterogeneity in the size of the indivisible expense $f(C^I | C^I > 0)$. In all cases, we obtained similar results.

¹⁸We tried the log-normal, beta, gamma, exponential, logistic, Weibull, and Generalized Pareto, among others.

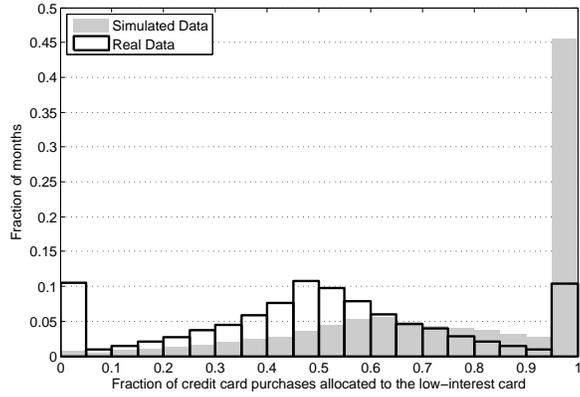


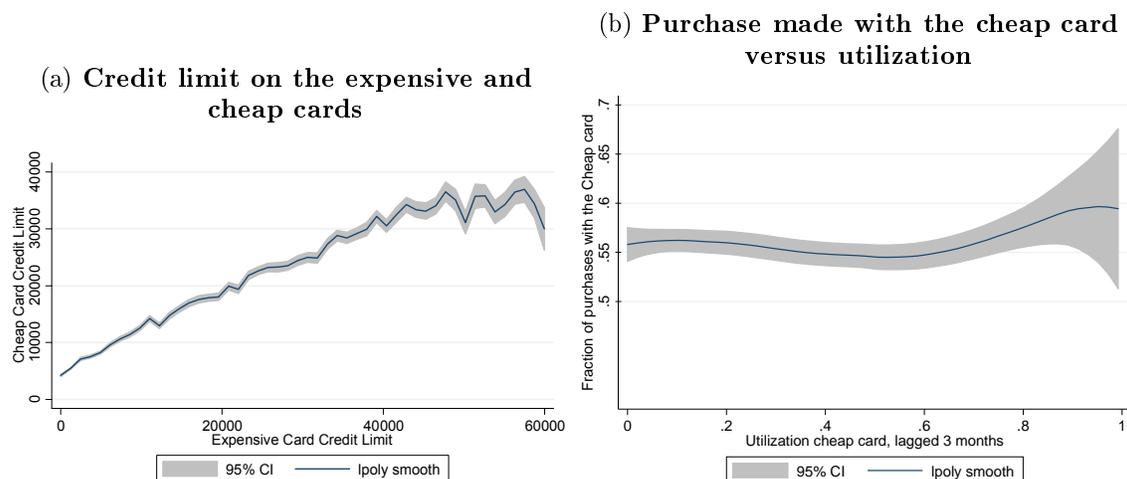
Figure OA.28: Allocation of Debt Predicted by the Indivisible Purchases Model

Notes: Data are from the administrative database (2-card sample). An observation is a consumer-month. The figure shows the frequency distribution of the fraction of debt allocated to the low-interest card predicted by the indivisible purchases model when $p = 1$ and $C = 3,000$.

IV.C Uncertainty in available credit

Rational consumers may make purchases with their high-interest card if they do not know how much available credit they have on the low-interest card and want to avoid overdraft fees (explaining Facts 1 and 2). This explanation is more likely if individuals have less available credit on their low-interest cards or if these cards have larger overdraft fees. In the data, however, credit limits are similar across low-interest and high-interest cards (see Panel A of Figure OA.29), and so are overdraft fees, \$84 vs. \$71 pesos on average, respectively. Similarly, according to this explanation, there should be a negative relationship between the fraction of purchases allocated to the low-interest card and the utilization rate of such a card. Nonetheless, Panel B of Figure OA.29 below shows that there is virtually no relationship between these two variables.

Figure OA.29: **Allocation of Debt and Available Credit (Kernel regressions)**



Notes: Point estimates and 95-percent confidence intervals plotted. Data are from the administrative database (2-card sample). An observation is a consumer-month. Panel A shows a kernel regression of the credit limit on the low-interest card against the credit limit of the high-interest card, using an Epanechnikov kernel. Panel B displays a kernel regression of the fraction of purchases made with the low-interest credit vs the utilization rate of that card lagged 3 months.

To examine the issue more formally, we also simulate a parsimonious static model where consumers are uncertain about the amount of available credit and tradeoff the expected overdraft fees against lower interest costs when allocating their purchases.

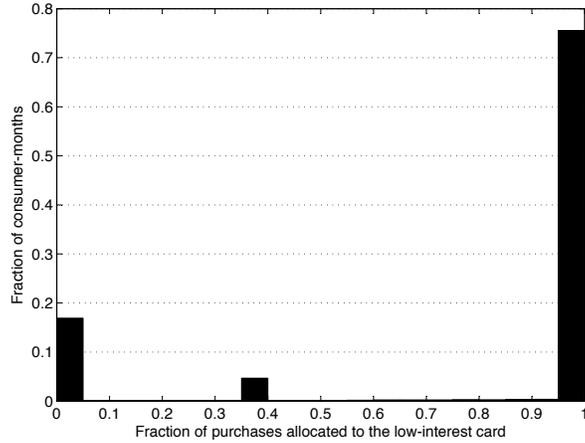
For each month, we take from the data the interest rates, overdraft fees, credit limits, outstanding balances on each card, and the total amount of purchases ($r_{1t}, r_{2t}, ODF_1, ODF_2, L_{1t}, L_{2t}, D_{1t}, D_{2t}, \bar{C}$) and estimate the allocation of purchases (C_1, C_2) that minimizes the *expected* cost, considering overdraft fees and uncertainty about the available credit on each card. Specifically, the optimization problem the consumer faces is given by:

$$\begin{aligned} \min_{C_1, C_2} & r_1 C_1 + r_2 C_2 + ODF_1 \cdot \Pr[L_1 - D_1 < C_1] + ODF_2 \cdot \Pr[L_2 - D_2 < C_2] \\ \text{s.t.} & C_1 + C_2 = \bar{C}, C_{it} \leq 1.2L_{it} - D_{it} \text{ for } i = 1, 2 \end{aligned}$$

In this exercise, we need to calibrate uncertainty, ideally at the consumer level to allow for heterogeneity. Given that the object of interest — over which there is uncertainty — is the available limit on each card, and that we have a limited number of observations for each individual, we introduce uncertainty in the amount of available credit on each card by drawing values from a uniform distribution with mean equal to the available credit at the beginning of the period and variance equal to the — within card — variance of the available limit on each card over time.

Figure OA.30 plots the distribution of optimal allocations that results from this model. The predicted distribution of the share of purchases allocated to the low-interest card from this model is significantly different from the allocations observed in the data.

Figure OA.30: **Purchases Allocation with Uncertain Limits**



Notes: Data are from the administrative database (2-card sample). An observation is a consumer-month. The figure shows the frequency distribution of the fraction of purchases allocated to the low-interest card predicted by the uncertainty and overdraft model.

IV.D Card heterogeneity

Another possibility is that consumers borrowed on their expensive cards because these cards had unobserved characteristics that made them valuable or more convenient for cardholders to use.¹⁹ To examine whether heterogeneity could account for Facts 1-4 we perform several exercises. First, we replicate our analysis using different samples. Second, we use our ancillary administrative data to assess the prevalence of online transactions and automatic payments, which could be considered differentiators. Finally, we use our survey to evaluate whether a rich set of card characteristics could explain misallocation and find that these characteristics explain only 5 percent of its variation.

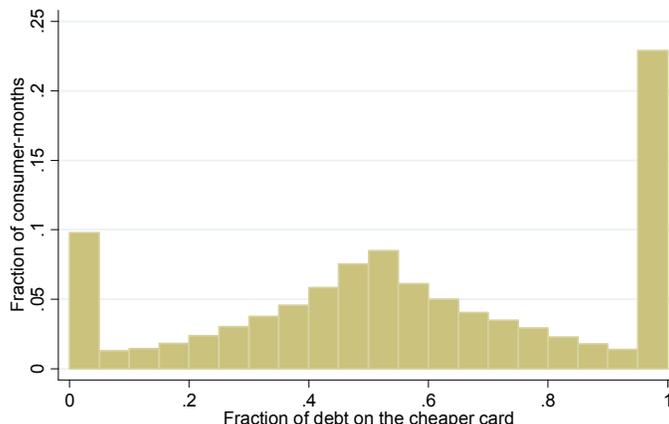
A. Robustness to different samples

First, we show that our results do not change when we examine *more homogeneous* samples. Figures OA.7, OA.12, and OA.16 above show that our results do not change when we restrict the sample to individuals holding two Classic cards, two cards with similar due dates, or two cards issued by the same bank. Similarly, Panel C in Figure OA.5 shows that the distributions of misallocated debt in our survey data for individuals holding two cards or two cards without reward points, as reported by respondents, are quite similar.

¹⁹If cards are heterogeneous, the extra financing cost we document could be viewed as a compensating differential for desirable product attributes. Even if this were the case, our Facts would still be interesting as they would show that small unobserved differences (to the econometrician) in card characteristics would be enough to eliminate substitution.

Second, we check the robustness of our results to the use of a *more heterogeneous* sample. Figure OA.31 below plots the frequency distribution of the fraction of misallocated purchases for consumers holding any pair of cards (i.e. without restricting the sample to include only comparable cards). The distribution looks very similar to the distribution presented in Panels E and F of Figure 2 in the paper.

Figure OA.31: **Histogram with no restriction to comparable cards**



Notes: Data are from the administrative database (2-card sample). An observation is a consumer-month. This figure plots the histogram of the share of misallocated purchases without restricting the sample to include only comparable cards.

B. Online transactions and automatic payments

Next, we examine whether online transactions and automatic payments are prevalent enough to account for the observed cardholders' allocations. To this end, we use a data set that includes the records of all the transactions listed on the credit card billing statements of 1,500 clients of one of our cooperating banks during the sample period. We classify each transaction as online or offline, and identify those that could be associated with bills paid automatically every month (such as internet services, cell phone, home phone, cable/satellite TV, newspaper/magazines, health clubs, utility services, car loans, departmental store loans, and insurance). Table OA.9 below summarizes the incidence and value of the periodic/automatic and online purchases made by credit card holders. Overall, only 0.2% of transactions correspond to online purchases and only 6% to bills paid automatically.²⁰

²⁰In our survey only 3% of transactions correspond to online purchases.

Table OA.9: **Periodic/automatic charges and online purchases**

	Percentage
<i>Panel A: Periodic/automatic purchases</i>	
Total number of automatic charges	6
Individuals with at least three automatic charges of the same kind during the sample period	23
<i>Panel B: Online purchases</i>	
Total number of online purchases	0.2
Individuals with at least one online purchase during the sample period	5
<i>Panel C: Survey information</i>	
Percentage of direct debit purchases out of total monthly purchases	5
Percentage of online purchases out of total monthly purchases	3

Notes: Data are from a subset of the administrative database (Panel A and B) and from the survey database (Panel C). This table summarizes the incidence and value of the periodic/automatic and online purchases made by credit card holders. The data set consists of twelve months of transaction data (with each charge displayed as in the monthly statement of the individual) from 1,500 random accounts in 2005 from one of our cooperating banks.

C. Card characteristics and misallocation

Finally, we use our survey — which contains detailed information on card and consumer’s characteristics — to evaluate whether these characteristics as well as attributes perceived by consumers have explanatory power for misallocation. Table OA.10 shows the results of OLS regressions of the fraction of misallocated debt (column 1), misallocated purchases (column 2), and misallocated payments (column 3) on an indicator for whether the high interest card has reward points, but not the other card, an indicator for whether the expensive card is perceived as more prestigious; an indicator for whether both cards were issued by the same bank; indicators for the type of credit card; account tenure; the interest rate gap; an indicator for consumers who perceived balance transfers as costly, defined as those consumers who assigned a value above the median in the survey; and the ratio of distances in minutes (as reported by respondents) to the most frequented branch of the bank that issued the high-rate and low-rate cards. The Table shows that only 4 out of 27 coefficients are statistically different from zero at the 5% significance level. An interesting exception is the indicator for “prestige”. Cardholders who indicated that their expensive card is the most valuable misallocate 12 percentage points more of their purchases and 9 percentage points less of their payments. Overall, however, these characteristics explain a small portion of the variation in misallocations. The explained variance is 4%, 11% and 7% for the fraction of misallocated debt, purchases, and payments, respectively.

Table OA.10: **Allocation of debt, purchases, and payments versus possible determinants of misallocation**

	Share misallocated		
	Debt (1)	Purchases (2)	Payments (3)
Expensive card has points and cheap does not	-0.024 (0.036)	0.071** (0.032)	-0.047 (0.040)
Expensive card is considered to be more prestigious	-0.020 (0.027)	0.122*** (0.024)	-0.094*** (0.030)
1 = Classic	0.093** (0.045)	0.011 (0.039)	-0.043 (0.049)
1 = Gold	0.114* (0.064)	-0.011 (0.055)	-0.063 (0.069)
1 = Platinum	0.097 (0.061)	-0.013 (0.053)	-0.064 (0.067)
Both cards issued by the same bank	0.001 (0.029)	0.029 (0.025)	0.049 (0.031)
Years with a credit card	-0.001 (0.003)	-0.001 (0.003)	0.002 (0.003)
$ r_1 - r_2 $	-0.000 (0.000)	-0.000* (0.000)	0.001*** (0.000)
Perceives balance transfers as costly	-0.089*** (0.026)	0.019 (0.022)	0.006 (0.028)
Relative bank distance (distance cheap/distance expensive)	0.011 (0.012)	0.002 (0.011)	0.016 (0.013)
Constant	0.259*** (0.053)	0.390*** (0.046)	0.461*** (0.058)
Observations	391	391	388
R-squared	0.04	0.11	0.07
max(VIF)	2.37	2.36	2.38

Notes: Data are from the survey database (2-card sample). An observation is a survey respondent. This table shows the results of OLS regressions of the fraction of misallocated debt (column 1), misallocated purchases (column 2), and misallocated payments (column 3) on several potential misallocation predictors. The first is a dummy variable equal to one when a consumer has a high interest card with points whereas the low interest card does not. The second is a dummy variable equal to one for those consumers who stated that the card that they value as the most prestigious one coincides with their expensive card. The Classic, Gold and Platinum dummies refer to the type of the low interest card; the omitted variable is the basic credit card type. We also included a dummy variable indicating those consumers that perceived balance transfers as costly — defined as those consumers who assigned a value above the median value in the survey. Finally, the relative distance is the ratio of distances in minutes (as reported by respondents) to the most used branch of the bank(s) that issued the high-rate and low-rate cards. ***, ** and * indicate statistical significance at the 1, 5, and 10 percent confidence levels, respectively. For each column, we also present the maximum variance inflation factor (ie. $\max(VIF) = \max_i \{VIF_i\}$).

Table OA.11 estimates the analogous regressions using elicited subjective beliefs as explanatory variables. Again, explained variance is at most 5% and coefficient estimates are small.

Table OA.11: **Allocation of debt, purchases, and payments versus subjective beliefs**

	Share misallocated		
	Debt (1)	Purchases (2)	Payments (3)
<i>Probability of in the following 12 months...</i>			
getting a card cloned or stolen	-0.000 (0.000)	0.001** (0.000)	-0.001 (0.001)
receiving a TRO	0.001** (0.001)	-0.001 (0.001)	-0.001 (0.001)
card being denied	0.002*** (0.001)	-0.001*** (0.001)	0.000 (0.001)
exceeding the credit limit	-0.000 (0.001)	-0.000 (0.000)	0.000 (0.001)
having to incur in an unexpected expense	0.000 (0.000)	-0.000 (0.000)	0.000 (0.000)
Constant	0.234*** (0.028)	0.495*** (0.025)	0.469*** (0.031)
Observations	392	392	388
R-squared	0.05	0.05	0.01
max(VIF)	1.60	1.61	1.60

Notes: Data are from the survey database (2-card sample). An observation is a survey respondent. This table shows the results of OLS regressions of the fraction of misallocated debt (column 1), misallocated purchases (column 2), and misallocated payments (column 3) on several variables indicating the probability that each consumer assigned on each event. Schooling is measured in years the individual attended school. Robust standard errors are given in parentheses. ***, ** and * indicate statistical significance at the 1, 5, and 10 percent confidence levels, respectively. For each column, we also present the maximum variance inflation factor (ie. $\max(VIF) = \max_i\{VIF_i\}$).

Finally, Table OA.12 below displays other statistics from our survey data that show that online purchases and direct debits were minor even in 2013. The table also shows that events like getting the card stolen or cloned are not prevalent. The table also includes other interesting statistics.

Table OA.12: **Statistics associated with credit card activity**

	# of credit cards	
	Two (1)	Three + (2)
<i>Panel A. Online purchases and direct debits</i>		
Percentage of respondents who...		
used a credit card to make purchases online during the previous month	9	17
Online purchases as a percentage of total purchases	3 (11)	3 (9)
Percentage of respondents who...		
have direct debit set up on the more expensive card	12	15
Direct debit purchases as a percentage of total purchases	5 (11)	6 (12)
<i>Panel B. Awareness</i>		
Percentage of cards for whom consumers...		
claim to know exact interest rate	76	83
claim to know their credit limit	98	99
<i>Panel C. Other selected questions</i>		
Percentage of respondents who...		
lost one of their credit cards or had it cloned or stolen during 2013 ¹	6	11
received a teaser rate offer in one of their credit cards during 2013 ¹	19	19
have transferred balances from a high-interest credit card to a low-interest card	1	3

Notes: This table shows summary statistics for selected variables for respondents in our survey who responded positively to the question "do you normally pay interest on your credit card(s)?" Each column represents a sub-population of interest. Standard deviations for each variable within each population are shown in parenthesis. The balance transfer cost in pesos includes a self-reported cost of the time and effort expended ¹. The survey was done by the 10th month of 2013, hence answers were standardized to represent this period of time.

IV.E Strategic manipulation of interest rates and credit limits

Consumers may borrow on their high-interest cards to try to influence their contract terms and obtain interest rate discounts or higher credit limits in the future through the rules that banks use to assign these terms (explaining Facts 1-4). In this sub-section, we estimate various parametric and non-parametric regressions of the likelihood of receiving a TRO or a change in credit limit on a large number of time-varying account characteristics

A. Determinants of TROs

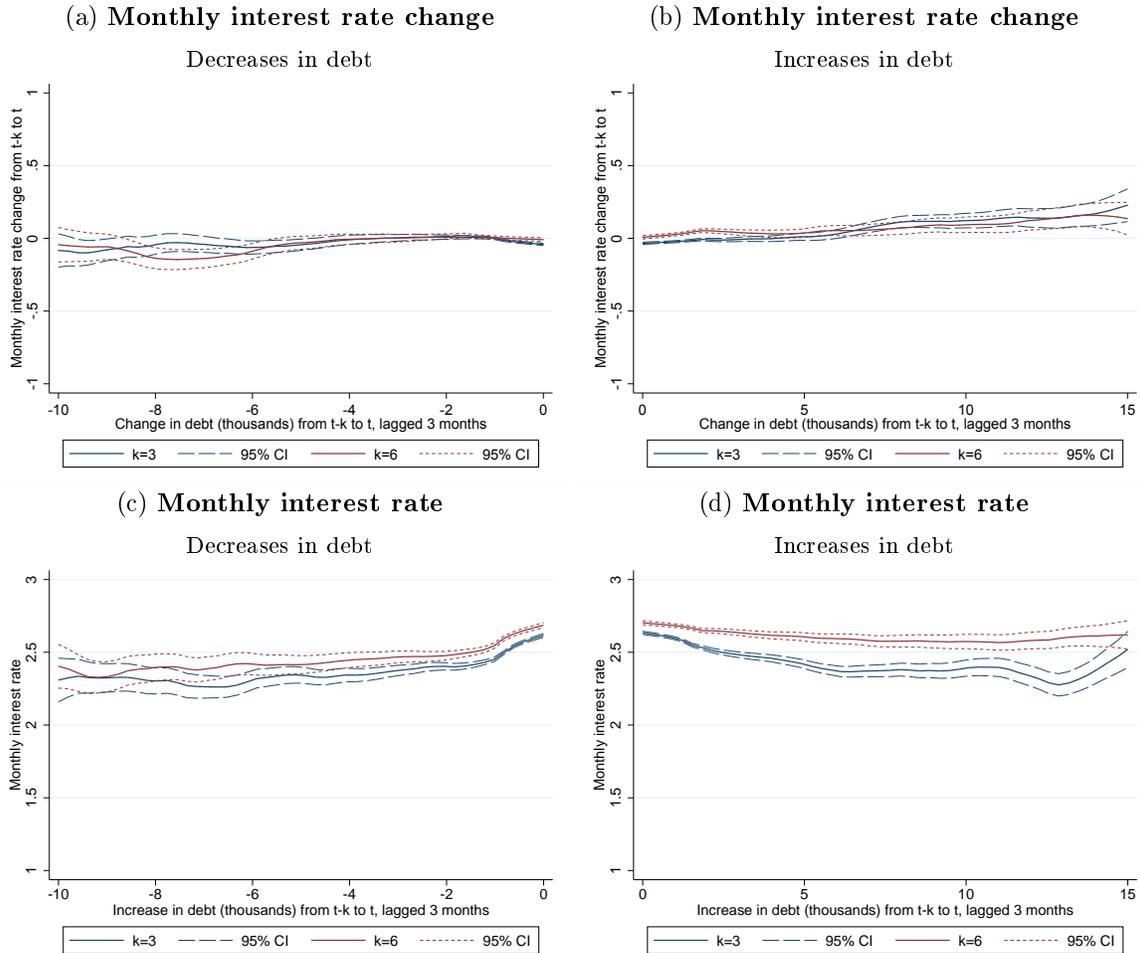
In the paper, we use a randomized experiment to estimate debt-interest rate elasticities consistently. We are comfortable with this approach, as the focus of the paper is on the behavior of cardholders. Nonetheless, there are a couple of issues that require attention if interest rates are set as a function of the allocation of debt. The first is that sophisticated consumers might hold debt on their high-interest card to obtain interest rate discounts or

higher credit limits in the future. The second issue is that our estimates of the extra costs may be biased upward as they do not take into account this future discounted benefit of borrowing on the more expensive card.²¹ In this sub-section, we present some results that argue against this possibility.

We begin by looking at non-parametric relationships between interest rates and past debt. Figure OA.32 presents kernel regressions of the increase (or decrease) in interest paying debt (in thousand pesos) between period $t - 3 - k$ and period $t - 3$ for $k = 3$ and $k = 6$ on the change in the monthly interest rate in t in percentage points (Panels A and B) and the level of interest rate in t (Panels C and D). The Figure shows that an increase (or decrease) in interest paying debt between period $t - 3 - k$ and period $t - 3$ for $k = 3$ and $k = 6$ bears no relation to the change in the monthly interest rate in t or the level of interest rate in t .

²¹To the extent that the decrease in the expensive card's interest rate occurs in the time-frame of our sample, we are taking *some* of these benefits into account as it is reflected in lower extra cost in later periods.

Figure OA.32: Relationship between the change in debt and interest rates



Notes: The figure shows a kernel regression of the increase in interest paying debt (in thousand pesos) between period $t-3-k$ and period $t-3$ for $k=3$ and $k=6$ on the change in the monthly interest rate in t in percentage points (Panels A and B) and the level of interest rate in t (Panels C and D) using an Epanechnikov kernel. Point estimates and 95-percent confidence intervals are plotted. Data are from the administrative database (2-card sample). An observation is a consumer-month. The change in debt is trimmed at the 95th percentile.

To supplement this analysis, we examine the relationship between past borrower behavior and TROs in a regression framework using our administrative sample. We construct two indicator variables for short-term offers. The first indicator equals one if, during that period, the individual received an offer that had an interest rate decreases of 50 percent or more (low interest offers), and zero otherwise.²² The second indicator is coded as one if, during that period, the individual received a zero interest offer. For comparison, we built an indicator for whether the consumer received an experimental TRO during that period. In Table OA.13, we estimate a linear regression of receiving an offer on past purchases, payments, debt, and other indicators of inactivity (i.e. zero purchases, zero debt, and zero payments). Overall, we

²²In the case of 3-month offers, the indicator equals one only for the first month of the offer

can explain at most 5 percent of the variance. We also tried specifications with longer lags, interactions, and quadratic terms. In all cases, the adjusted R-squared was as small as in our original specification, at 0.05. We also tried separate regressions by bank and duration of the teaser offer and obtained similar results. Overall, it is hard to predict who is more likely to receive a teaser rate offer.

Table OA.13: **Determinants of TROs**

	1[Low-interest offer]	1[Zero-rate offer]	1[Random offer]
	(1)	(2)	(3)
UR (log)			
L1.	-0.02*** (0.003)	0.02*** (0.003)	-0.0005 (0.0007)
L2.	-0.001 (0.002)	0.01*** (0.003)	0.0005 (0.0009)
L3.	0.002 (0.002)	0.01*** (0.002)	0.0001 (0.0007)
Purchases (log)			
L1.	0.006*** (0.001)	-0.001 (0.001)	0.003*** (0.0004)
L2.	-0.0004 (0.0008)	0.0006 (0.0009)	-0.0006* (0.0003)
L3.	-0.001 (0.0008)	0.002** (0.0009)	-0.0008 (0.0003)
Payments (log)			
L1.	0.001*** (0.0003)	0.008*** (0.0006)	-0.00002 (0.0001)
L2.	-0.0006 (0.0003)	0.006*** (0.0005)	-0.00003 (-0.35)
L3.	0.0003 (0.0003)	0.002*** (0.0005)	-0.0006 (0.0001)
No interest incurred (dummy)			
L1.	-0.05*** (0.006)	0.06*** (0.006)	-0.005*** (0.001)
L2.	0.02*** (0.006)	0.07*** (0.007)	-0.001 (0.001)
L3.	0.03*** (0.005)	0.05*** (0.009)	0.0007 (0.001)
Zero purchases (dummy)			
L1.	0.03*** (0.005)	-0.001 (0.006)	0.01*** (0.002)
L2.	0.0008 (0.005)	0.005 (0.006)	-0.003* (0.002)
L3.	0.007 (0.005)	0.01** (0.006)	0.004* (0.002)
Zero payments (dummy)			
L1.	0.001 (0.002)	0.04*** (0.003)	-0.0003 (0.0008)
L2.	-0.0007 (0.002)	0.02*** (0.003)	-0.002*** (0.0006)
L3.	0.001 (0.002)	0.009*** (0.003)	-0.001** (0.0006)
R-squared	0.04	0.04	0.04
Consumers	9,916	9,916	961
Observations	70,597	70,597	6,390

Notes: This table shows the estimates from linear probability models of indicators for receiving a low-interest offer on time-varying card characteristics such as utilization rate, purchases, and payments; indicators for whether the card holder paid no interest, made no purchases, or made no payments; and their lags (L). Column (1) reports the results using all low-interest offers as the dependent variable. We define a low interest offer as one where the interest rate is decreased by at least 50% from one month to the next. Column (2) focuses on zero-rate offers. Column (3) displays the regression results using the randomized TRO as the dependent variable. Standard errors clustered at the individual level are given in parentheses. ***, **, and * indicate statistical significance at the 1, 5, and 10 percent confidence levels, respectively.

Finally, in the table below we estimate the effect of lagged utilization on the likelihood of receiving a TRO using a linear regression with fixed effects [columns (3) and (4)]. In the two regressions, the coefficient estimates are not statistically significant. In particular, a 50 percentage point increase in utilization three months before is associated with a change on the likelihood of receiving a TRO by 1pp, 0.5pp or 0.03pp depending on the specification.

Table OA.14: **Low interest offers vs utilization**

	I(TRO_t)			
	Mean dep. = 0.025			
	(1)	(2)	(3)	(4)
Utilization_t-3	-0.0002*** (0 .00009)		0.0001 (0.0001)	
Utilization_t-6		-0.0003*** (0.00009)		-0.00006 (0.00004)
Time effects	Yes	Yes	Yes	Yes
Fixed effects	No	No	Yes	Yes
R2	0.02	0.02	0.05	0.05
Num Obs	76407	53028	76407	53028
Consumers	10300	10157	10300	10157
Mean of utilization	49.0	50.0	49.0	50.0
Within s.d. of utilization X	13.9	13.2	13.9	13.2
Prob points of 1 s.d. increase in X	-0.003	-0.004	0.001	-0.001
%50 increase in utilization is associated with Y prob points	-0.01	-0.015	0.005	-0.003

In a framework of rational expectations in which consumers know and anticipate banks' rules, these results would be enough to rule out the notion that consumers strategically manipulate the supply rules. But if we allow for the possibility that consumers may have wrong or misinformed expectations, one could still claim that consumers may *think* that they know the rule the bank uses to determine TROs, even when these rules do not exist or when it is practically impossible for consumers to infer such policies. As reported in the paper, we asked respondents in our survey questions that speak closely to this issue. First, we asked respondents whether they thought banks followed a rule to determine if and when to send them a TRO based on their credit card use, and if they had an idea of what that rule was. Next, we asked respondents whether they actually acted upon this rule and managed their credit card debt to try to get better terms. Only 11 percent of survey respondents claimed to know the rule that banks use to decide whether to send a TRO, and a mere 2 percent claimed to allocate their balances to influence that rule and obtain a low interest rate offer. In our view, this evidence argues against the interpretation that consumers borrow on their high-interest cards to try to obtain interest rate discounts in the future through the rules that banks use to assign TROs, and suggests that not modeling the supply side is a reasonable simplification that does not affect the paper's main conclusions.

B. Determinants of changes in credit limits

We conducted analogous regressions for the likelihood that the bank changes the credit limit of a card. Table OA.15 below show the results. Columns (1) and (2) show the results

for credit limit increases without using card fixed effects.²³ Columns (3) and (4) include card fixed effects, which means that the parameters are identified only by within-consumer variation, bolstering the case for a causal interpretation of the results. The first thing to notice is that, in all cases, the correlations are tiny. Take for instance column 1: an increase of one standard deviation in past utilization is associated with just a 0.001 percentage point decrease in the likelihood of getting a credit limit increase, or 0.7 percent of the mean likelihood. Once we control for card fixed effects, the coefficient becomes *positive* and still tiny in column (3), and negative and statistically insignificant in column (4). In column (5), we use dummies for various levels of past utilization to account for the possibility that banks might reduce the likelihood of offering credit limit increases only after cardholders reach a certain threshold of past utilization. We obtain qualitatively the same results to those shown in the other four columns. In columns (6)-(10), we use credit limit decreases as the dependent variable and find comparable results. Taken together, these findings argue against the idea that consumers allocate their balances to obtain higher credit limits.

²³We do this for comparison purposes only. We think that using card fixed effects gets closer to the causal effect of interest as the simple OLS regressions compare across consumers with different utilization rates which may be very different from one another undermining a causal interpretation of the estimated coefficient

Table OA.15: Utilization rate and likelihood of credit limit increases

	I(D limit_t > 0)				I(D limit_t < 0)					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
		Mean dep. = 0.2				Mean dep. = 0.0006				
Utilization_t-3	-0.0001*** (.00003)		0.0003** (0.0001)			-4.54e-06** (2.2e-06)		-4.1e-06 (5.9e-06)		
Utilization_t-6		-0.0002*** (0.00006)		-0.0002 (0.0002)				-3.5e-06 (2.94e-06)	3.19e-06 (5.6e-06)	
Utilization_t-3 Q2					-0.001 (0.006)					-0.0005 (0.0005)
Utilization_t-3 Q3					0.015** (0.007)					-0.0002 (0.0004)
Utilization_t-3 Q4					0.043*** (6.12)					(0.0006)
Time effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects	No	No	Yes	Yes	Yes	No	No	Yes	Yes	Yes
R2	0.17	0.19	0.18	0.21	0.18	0.004	0.005	0.005	0.006	0.005
Num Obs	76407	53028	76407	53028	76407	76407	53028	76407	53028	76407
Consumers	10300	10157	10300	10157	10300	10300	10157	10300	10157	10300
Mean of utilization	49.0	50.0	49.0	50.0	49.0	49.0	50.0	49.0	50.0	49.0
Within s.d. of utilization X	13.9	13.2	13.9	13.2	13.9	13.9	13.2	13.9	13.2	13.9
Prob points of 1 s.d. increase in X	-0.001	-0.003	0.004	-0.003	0.000	0.000	0.000	0.000	0.000	0.000

IV.F Limited attention and salience

Even if consumers know their interest rates, they could still misallocate their debt if they do not pay enough attention to prices at the time of making their purchases and payments. In this section, we display copies of monthly credit card statements as well as summary statistics for variables related to the reasons for using certain cards and the attention paid to interest rates by survey respondents.

A. Monthly billing statements

Figure [OA.33](#) and [OA.34](#) below present copies of two monthly billing statements issued during our sample period by two banks that do not necessarily correspond those in our main data set. As the Figures show, the monthly billing statement did disclose the annual interest rate, but it was shrouded (not-salient) in a mist of other quantities and therefore easily lost.

Figure OA.33: Sample of a monthly billing statement 1



ESTADO DE CUENTA

CLIENTE NUMERO	12042610
TARJETA NUMERO	X
FAVOR DE PAGAR ANTES DE	16 DIC, 2005
SALDO NUEVO	\$ 2,244.70
PAGO MINIMO	\$ 200.00
PAGO PARA NO GENERAR INTERESES	\$ 1,176.47
	X



ESTADO DE CUENTA

12271 102

X

SI TIENES MAS DE UNA TARJETA BANCOMER AHORA SOLO PAGAS LA CUOTA ANUAL DE LA DE MAYORES BENEFICIOS (APLICAN RESTRICCIONES).

CLIENTE NUMERO	12042610
TARJETA NUMERO	4931-6180-0002-3791
FAVOR DE PAGAR ANTES DE	16 DIC, 2005
PAGO MINIMO	\$ 200.00
PAGO PARA NO GENERAR INTERESES	\$ 1,176.47

INFORMACION GENERAL

LIMITE DE CREDITO	\$ 22,000.00
CREDITO DISPONIBLE	\$ 19,755.50
FECHA DE CORTE	27 NOV, 2005
DIAS TRANSCURRIDOS EN EL CICLO	31
SALDO PROMEDIO	\$ 822.70
TASA MENSUAL	3.45%
TASA ANUAL	41.40%
MONTO NO EXIGIBLE	\$ 0.00
COMPRAS DIFERIDAS	\$ 0.00

RESUMEN DE SALDOS Y MOVIMIENTOS

SALDO ANTERIOR	\$ 2,078.87
PAGOS	- \$ 857.00
OTROS ABONOS	- \$ 0.00
RENDIMIENTO	- \$ 0.00
COMPRAS	+ \$ 952.65
OTROS CARGOS	+ \$ 0.00
DISPOSICIONES DE EFECTIVO	+ \$ 0.00
INTERESES (SIN IVA)	+ \$ 28.42
I.V.A.	+ \$ 41.76
SALDO NUEVO	= \$ 2,244.70

DETALLE DE TRANSACCIONES TARJETA TITULAR 4931-6180-0002-3791

FECHA	CONCEPTO	IMPORTE \$
17 - NOV	*** SU PAGO GRACIAS ***	\$ 857.00-
17-11	74931615521950172181932	
17 - NOV	COMISION PAGO TARDIO	\$ 250.00
17-11	74931615521970172007125	
20 - NOV	IUSACELL	\$ 702.65
22-11	74555465325850060857617 5968	
27 - NOV	CARGO 05 DE 12 AOL 12 H SIN IN	\$ 154.00
28-11		

BBVA BANCOMER, S.A. INSTITUCION DE BANCA MULTIPLE, GRUPO FINANCIERO
25,056 140

Notes: The figure shows a monthly billing statements issued during our sample period by a bank in Mexico.

Figure OA.34: Sample of a monthly billing statement 2

 					
SALDO ACTUAL 10,030.55 Saldo del mes 9,069.22		PAGO MINIMO 1,079.97		FECHA LIMITE DE PAGO CLIENT Feb 4 2006	
HOJA 1 de 3					
FECHA REG	CONCEPTO	FECHA TRAN.	REFERENCIA	IMPORTE \$	
TRANSACCIONES DEL -					
12/16	LUAU MEXICO DF	12/16	74941335350001908191514	191.00	
12/20	SANBORN S PERISUR CIUDAD DE ME	12/18	74540615352605021035974	63.50	
12/20	COSTCO SAN ANTONIO CIUDAD DE ME	12/19	74540615353237070237415	434.99	
12/21	ITALIANNI S ZONA ROSA MEXICO DF	12/20	7455485355955568245778	204.00	
12/23	ZARA ZONA ROSA CIUDAD DE ME	12/22	74540015357625216204950	1,470.00	
12/23	ITALIANNI S INSURGENTE MEXICO DF	12/22	7455485357955567255909	403.00	
12/26	MAYAN COUNTRY CLUB MEXICO DF	12/21	7455485358951426237334	1,829.64	
12/27	ZARA ZONA ROSA CIUDAD DE ME	12/26	74540615361625218844367	80.00	
12/27	ZARA ZONA ROSA CIUDAD DE ME	12/26	74540615361625218844656	517.00	
12/28	VIPS HAMBURGO MEXICO DF	12/27	7455485362909702286835	154.50	
12/29	FARM SAN PABLO MEXICO DF	12/28	74912845362025494540429	732.70	
12/30	ITALIANNI S INSURGENTE MEXICO DF	12/29	7455485364955567253061	436.00	
12/30	BABY OUTLET NAUCALPAN NAUCALPAN EM	12/30	74912595364001285951595	35.00	
01/02	ZARA ZONA ROSA CIUDAD DE ME	12/29	74540615364625218264431	119.00	
SALDO ANTERIOR		LIMITE DE CREDITO		20,000.00	
+ DEBITOS, COMPRAS Y COMISIONES		FECHA DE CORTE		Ene 14 2006	
+ DISPOSICIONES EN EFECTIVO		DIAS DE PERIODO		32	
- PAGOS		TASA MENSUAL		1.50%	
- OTROS ABONOS		TASA ANUAL		18.00%	
+ INTERESES		TASA ANUAL DE INVERSION		0.00%	
+ I.V.A.		TASA MORATORIA ANUAL		36.00%	
= SALDO ACTUAL		CREDITO DISPONIBLE		9,069.45	
		SALDO PROMEDIO DIARIO		4,877.17	
		NUMERO DE CUENTA			

Notes: The figure shows a monthly billing statements issued during our sample period by a bank in Mexico.

B. Summary statistics for selected variables (survey data)

Table OA.16 presents the frequency distribution of responses for selected questions related to attention and price shopping. As the table shows, most consumers do not take interest rates into consideration when making a credit card purchase or a credit card payment. Along the same lines, most individuals do not compare the interest rates on their cards when making a credit card purchase. However, most individuals compare prices or brands when buying groceries or take interest into account when getting a new card. The Table also shows that most consumers are passive when it comes to getting a new card: only 30 percent of respondents claimed to have actively searched for their last card, while the rest answered that they simply got it because the bank offered it to them. Furthermore, 83 percent of respondents stated that they did not compare prices across providers when applying for a new card.

Table OA.16: **Attention to interest rates (survey data)**

	Pct. of cardholders with 2 cards (1)	Pct. of cardholders with 2 cards w/o points (2)
Overall, how much do you take interest rates into account when taking out a new card?		
A great deal	31	45
To a certain extent	25	34
Not much/Not at all	44	21
Overall, how much do you take interest rates into account when making a credit card purchase?		
A great deal	17	25
To a certain extent	27	35
Not much/Not at all	56	40
Overall, how much do you take interest rates into account when making your credit card payments?		
A great deal	19	29
To a certain extent	30	39
Not much/Not at all	51	32
Do you compare prices of different brands when purchasing your groceries?		
Almost always	48	45
Occasionally	26	24
Seldom / Never	26	31
Do you compare the interest rates of your cards when making a credit card purchase?		
Almost always	13	12
Occasionally	25	28
Seldom / Never	62	60
Thinking about your last card, did you apply for it or was it offered to you?		
I looked for it and applied for it myself	29	
I was offered it	70	
Does not know	1	
Regarding the last card obtained...] How many banks did you compare?		
I didn't compare any banks	83	
I compared 2 banks	12	
I compared 3 or more banks	4	
Does not know	1	

Notes: This table shows summary statistics for selected variables for respondents in our survey who responded positively to the question "do you normally pay interest on your credit card(s)?" Each column represents a sub-population of interest.

IV.G Mental accounting

Consumers could misallocate their debt, purchases, and payments if they associate a particular mental budget with a particular credit card, allocating specific types of purchase to such a card or establishing a maximum level of leverage for each particular card separately. Mental accounting can act as a psychological switching cost and as such can explain Facts 1, 2, 3, and 5b; however, it is silent about Facts 4 and 5a. In this section, we evaluate the performance of this theory in explaining Facts 1 and 2 using two pieces of evidence. First, we use our purchase-type data set to examine whether consumers earmark a specific card for certain spending categories (for example, a card for “vacations” and another for “electronics”), assuming that these categories correspond to the consumers’ mental accounts.²⁴ Next, we present evidence on individual responses to some questions related to categorization and mental accounting in our survey.

A. Earmarking specific cards for certain spending categories (*Purchase-type data*)

In Table OA.17 below, we show the means and standard deviations of the total monthly purchases in each category, conditional on them being positive, as well as the number of observations (consumer-months) in which cardholders made a purchase in each particular category with either of their cards. The last four columns present a measure of earmarking of a given card for each category. This measure is calculated for each individual i , card j , and spending category z as the average across months of the fraction of purchases made with card j by individual i that belongs to the spending category z . The card with the biggest share of purchases in category z is the preferred-for- z card ($\mathcal{P}z_i$).²⁵ Overall, we find that consumers tend to use *the same* card for most of their purchases in a given category. Column 4 shows that the average purchase share of the $\mathcal{P}z_i$ card by type of purchase ranges between 72% and 97% and has a mean value across purchase types of 87%. For example, out of 100 pesos spent on lodging, 86 are purchased with the preferred-for-lodging card. If there were no specialization, we would expect this figure to be close to 50 pesos.²⁶ Consumers do not always use the same card, but rather use a particular card for certain types of expenses. We define for each consumer their most preferred card *overall* (\mathcal{P}_i) as the card that they prefer for the majority of the 38 purchase categories. We find that the average consumer uses their \mathcal{P}_i card for 27 of the 38 spending categories and use the other for the remaining 11. This means there is *card switching* depending on the type of purchase. We also find that individuals who use one card for half of their spending categories and the other card for the other half incur a

²⁴We recognize that the assumption that these administrative spending categories represent the consumers’ mental accounts is ad hoc. In fact, a complex problem that the empirical literature on mental accounting faces – and our approach is no exception – is that the econometrician does not observe the actual mental accounts that consumers might construct on the basis of purpose, source, or timeframe. Using administrative categories that do not match perfectly with these accounts may make it harder for us to detect that they matter for misallocation, should a relationship truly exist. In spite of this, we find card/purchase-type specialization and a correlation of this specialization with extra interest costs.

²⁵This measure ranges between 50% (half of the purchases in that category are made with one card and half with the other) and 100% (all purchases in that spending category are made always with the same card).

²⁶The other columns show that the numbers are similar when we condition on cards with at least 5 type- z purchases and on consumers holding exactly the same type of card (e.g., Visa Gold). This latter fact is important, since these cards have the same reward point structure.

financing cost 28 percent larger compared with the cost incurred by those who specialize on their cheapest card, so that specialization may have cost implications.

Table OA.17: **Specialization**

	Mean	Std. dev.	# obs.	Degree of specialization			
				All	$\geq 5P$	SCCT	Both
				(4)	(5)	(6)	(7)
Recreational purchases	708	1,490	10,962	97	95	97	96
Home	3,491	6,532	508	96	75	98	81
Car Rental	3,140	4,052	657	94	77	94	80
Art	1,647	6,779	498	93	81	92	87
Computers	2,983	6,876	1,401	93	76	93	70
Pets	829	1,395	634	92	75	93	80
Door to door sales	1,865	3,029	4	92		75	
Children and toys	1,272	3,849	1,083	92	76	90	77
TV and cable	664	421	16,385	92	91	92	91
Uncategorized	750	5,948	7,073	92	87	92	87
Books and press	781	1,335	2,231	91	75	92	71
Jewel shops	2,556	5,302	1,818	91	74	91	73
Office material	1,380	3,739	1,560	91	72	91	73
Government	3,685	7,984	2,449	90	79	90	78
Wine and liquor	1,528	4,568	1,382	90	76	90	78
Financial, insurance	542	2,130	57,513	90	88	89	88
Sports	1,638	3,531	2,809	89	72	89	71
Travels and entertainment	2,005	5,844	5,290	89	78	87	76
Education	5,202	8,185	2,703	88	76	88	72
Various services	794	11,188	35,002	88	87	88	88
Cleaning	463	1,121	404	88	70	85	60
Airlines	6,228	13,297	4,910	87	74	88	71
Lodgings	2,702	5,102	6,262	86	75	86	73
Telephone	1,314	2,617	27,994	86	83	86	82
Construction	3,949	8,968	7,043	86	73	87	73
Furniture and equipment	3,115	8,523	9,133	86	74	86	74
General entertainment	831	1,870	8,499	85	75	84	75
Food and groceries	792	5,364	8,833	84	74	84	72
Telemarketing	1,923	5,764	14,045	84	74	84	73
Automobile	3,306	8,891	9,282	84	72	84	71
Medicines	1,538	5,987	15,600	82	73	82	72
International purchases	5,595	12,867	18,642	82	76	82	76
Clothing, footwear, accessories	1,867	7,055	21,473	80	72	79	71
Gasoline and chemistry related	1,085	4,043	27,891	78	72	77	71
Restaurants and bars	1,029	1,923	29,314	77	70	78	69
Department store	1,630	3,371	29,457	77	71	78	70
Cash withdrawal	4,804	10,066	41,657	75	71	74	71
Supermarkets	1,824	3,838	67,547	72	69	71	68

Data are from the administrative database purchase-type data. This table shows summary statistics of the monthly purchases made by individuals holding two cards in one of our cooperating banks by spending category. One of their two cards was sampled at random. Column (1) presents mean purchase size conditional on purchase, while column (2) shows standard deviations. Column (3) reports the number of purchase transactions. Columns (4), (5), (6) and (7) display the measure of specialization described in the paper. This measure ranges between 50 percent when half of the purchases in that category are made with one card and half with the other, and 100 percent when all purchases in that spending category are made always with the same card. Column (4) uses the full sample. Column (5) restricts the sample to cases with at least 5 purchases in each spending category. Column (6) restricts the sample to individuals holding cards of the same type. Column (7) restricts the sample to cases with at least 5 purchases in each spending category for individuals holding two cards of the same type.

B. Survey responses related to mental accounting

In this sub-section, we look for evidence of mental accounting in our survey. Table [OA.18](#) presents the reasons that our survey respondents reported for using their different credit cards. Table

Table OA.18: Reasons for using certain cards (Survey)

	Percentage (1)
Percentage of respondents that claimed to prefer borrowing on different cards	55
Stated reasons for using the card with the higher balance (%):	
The card has a higher credit limit available	30
The card has a lower interest rate	25
The card had a low-interest rate offer	13
The card has a good rewards program	8
Had too much debt on the other card and I did not want to overdraft	8
Had too much debt on the other card and wanted to spread it	6
Other	12
Stated reasons for using other cards besides the card with the higher balance (%):	
Prefer to distribute debt over several cards	23
Did not want to go over the limit on the other card	22
Prefer to use specific cards for certain purchases	20
The card has a lower interest rate	11
The card had a low-interest rate offer	10
Have direct debit set up on the card	6
Other	8
Stated reasons for paying a larger amount to this card (%):	
Wanted to pay off some of my recent purchases as soon as possible	28
Have more debt on this card	24
The due date coincided with the day I had more cash	16
The minimum was higher than the minimum on my other card	12
The card has a higher interest rate	10
Other	10
(Regarding purchases), do you prefer to make them with the same card or to use different cards?	
Use different cards	61
Use the same card	39
Number of times over the past year that the respondent used the high-interest card (instead of the low-interest card) because:	
He/she wanted to avoid going over the limit on the low-interest card	0.9
He/she wanted to leave enough room on the low-interest card for a major purchase	0.4
The low-interest card was not accepted	0.4
He/she was afraid that the low-interest card might be cloned	0.2
Percentage of consumers who...	
Claim to know the rule that the bank follows to determine interest rates	11
Claim to allocate their debt to try to obtain lower interest rates	2

Notes: This table shows summary statistics for selected variables for respondents in our survey with 2 credit cards who responded positively to the question "do you normally pay interest on your credit card(s)?" (386 observations). Direct responses were constructed from direct questions asked in the survey to people who had at least two cards. The questions used for this table, in the respective order, are 3.39 When you accrue credit card debt, do you prefer to spread it over various cards or consolidate it on a single card?, 3.41 You previously mentioned that last month you spent more on your [NAME OF CARD ON WHICH MOST WAS SPENT] card than on your other credit cards. Could you tell me what the main reason was for your using that card rather than a different one?, 3.42 Now let's talk about your other cards. Based on your previous answers, although you didn't use your other cards to make most of your purchases and payments, you did, however, use them. Could you tell me the main reason you used the other credit cards and not just your [NAME OF CARD ON WHICH MOST WAS SPENT] card?, 3.43 Now let's talk about your credit card payments. What was the main reason for your paying off more on your [NAME OF CARD ON WHICH MOST WAS PAID] card than on your other credit cards?, 3.37 In the case of the credit card payments and purchases you made: Do you prefer to use various cards or always use the same one?, 3.29 So far in 2013, how many times have you had to use other credit cards instead of your [NAME OF CHEAP CARD] card because you were afraid you would go over the credit limit on that?, 3.32 So far this year, how many times have you used other credit cards instead of your [NAME OF CHEAPEST CARD] card because you wanted to leave enough room on that to be able to make a major purchase?, 3.33 So far this year, how many times have you had to use other credit cards instead of your [NAME OF CHEAP CARD] card because the place where you wanted to use it didn't take it?, 3.34 So far this year, how many times have you had to use other credit cards instead of your [NAME OF THE CHEAP CARD] card because you were afraid it might get cloned?, 3.62 In general, do you think you have a good idea of what that rule is? and 3.63 Do you try to influence the bank's decision whether to lower the rate you pay by choosing which card you use to buy things or which you choose to pay the balance on, yes or no?.

Table OA.19: Survey questions related to mental accounting

	Percentage (1)
<i>Panel A. Mental accounting</i>	
Percentage of people who would buy 2 items in the same category with the same card. . . [<i>Stated preference exercise</i>]	
vacations	76
restaurant expenses	68
household appliances	72
groceries	82
all of them	37
Percentage of people who always use the same card	12
Given that they do not use the same card for every purchase ¹ . . .	
Percentage of people who would buy 2 items in the same category with the same card. . . [<i>Stated preference exercise</i>]	
vacations	73
restaurant expenses	63
household appliances	69
groceries	79
Some people tend to use one card for certain type of purchases and another card for certain others. For example, some people use their Bancomer card for vacations, while they use the Banamex for the supermarket. Do you tend to organize your expenses in this way and assign only certain purchases for an specific card? [<i>Direct question</i>]	
Almos always	25
Most of the times	28
Rarely	19
Never	27
Does not know	1
<i>Panel B. Expectations</i>	
Number of times over the past year that respondents accumulated debt on the high-interest card thinking that they would be able to pay the entire balance off, but weren't able to.	0.9

This table shows summary statistics for selected variables for respondents in our survey who responded positively to the question “do you normally pay interest on your credit card(s)?” For each category on the hypothetical choices (Panel B) (except the one called *all of them*) we had two items listed in the survey. For vacations, the items were a vacation package and a plane ticket; for restaurant expenses, the items were a family meal at a restaurant and a romantic dinner; for household appliances, the items were a refrigerator and a television; for groceries, the items were supermarket shopping and grocery (*abarrotes* in spanish) shopping. It is important to note that items were listed in a random order. For each item listed, we asked the survey respondents who had at least two cards “*From your credit cards, which one would you use to buy x?*” where x represents each item.

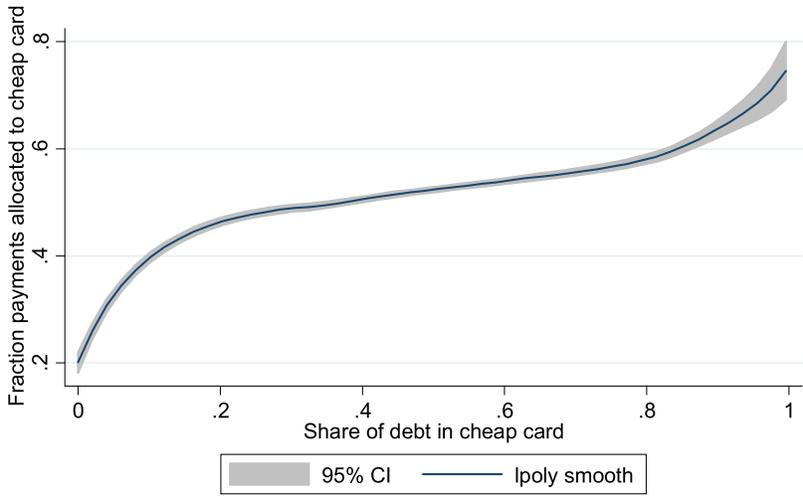
When asked directly about whether they use different cards for different types of purchases, about 53% of cardholders reported using one of their cards for certain expenses and the other for other purchases “almost always” or “most of the time” (Panel A of Table OA.19). We also run a hypothetical choice experiment that asked respondents holding two cards to indicate

which of their actual credit cards they would use to purchase various items.²⁷ If consumers engage in mental accounting, they will tend to use the same card for purchases belonging to the same spending group or mental account. We find that approximately three-quarters of the respondents use the same card for purchases belonging to the same category. This is not driven by consumers who responded that they use the same card for all purchases (12% of them).

Finally, when asked about the reasons for making payments on the card on which cardholders paid the most, the two most cited reasons were that they wanted to pay off some of their recent purchases as soon as possible (28%) and because they had more debt on that card (24%). We find a similar pattern in our administrative data: the amount of outstanding debt on a particular card is a powerful predictor of the allocation of payments. This suggests that consumers use debt instead of interest rates as a reference to allocate their balances. A regression of the fraction of misallocated payments against a five-degree polynomial of the fraction of outstanding balance on the card explains 10% of the variance in the allocation of payments. This is almost ten times more than the percentage of variation explained by a polynomial of the interest rate difference. We estimate an elasticity of 35% (t-stat=44.36) of the fraction of debt on a particular card and the fraction of payment to that card. Figure OA.35 shows the explanatory power of this relationship graphically. The graph plots the relationship between the fraction of total debt allocated to the low-interest card on the fraction of payments allocated to the low-interest card. As the Figure shows, the amount of outstanding debt is a powerful predictor of the allocation of payments.

²⁷The pairs of items are: (1) a refrigerator and (4) a TV; (2) a vacation package and (5) an airplane ticket; (3) a romantic meal at a restaurant and (7) a family meal; and (6) assorted groceries and (8) supermarket shopping.

Figure OA.35: **Share of payments above the minimum allocated to the low-interest card (Kernel regression)**



Notes: Data are from the administrative dataset (2-card sample). This figure shows a kernel regression of the share of total debt allocated to the low-interest card on the fraction of payments allocated to the low-interest card using an Epanechnikov kernel. The share of debt in the cheap card is trimmed at the 95th percentile. For this figure, the sample is restricted to consumers paying interests in the expensive card.

V Possible supply-side responses and market implications

This section presents evidence consistent with the idea that, in response to the low levels of substitution and the high sensitivity to salient own-card rate changes, banks have incentives to temporarily reduce account specific interest rates (through TROs) to entice consumers to borrow, and then raise them to exploit locked-in consumers (in the spirit of the switching cost literature).

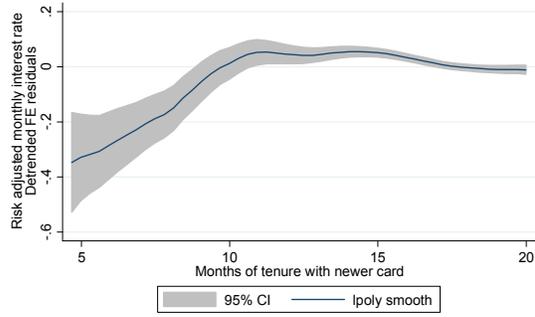
The left panel of Figure OA.36 plots the predicted OLS residual $\epsilon_{i,t}$ of the regression $IntRate_{i,t} = \alpha_i + \gamma_t + \theta X_i + \delta * LaggedRiskVars_{i,t} + \epsilon_{i,t}$ as a function of time. As the Figure shows, there is a positive relationship between account tenure and interest rates, controlling for risk,²⁸ demographics, as well as month, bank, and type of card indicators: in a span of one year, *monthly* risk-adjusted interest rates increase by 0.4 percentage points after consumers open their accounts, after which it becomes flat.

Next, to assess whether banks issuing new cards set prices assuming that consumers will not switch their debt back to their old cards after the teaser period ends, we estimate the same specification using the percentage difference in interest rates between the card the consumer got last (new) and the card that she got first (old) ($[r_{new}/r_{old}] - 1$) as the dependent variable, and plot the residuals of this regression against the tenure of the new card. The right-hand side panel of Figure OA.36 shows that for consumers holding two cards, the interest rate on the card they opened last (new) is initially 10 percent lower than the rate on the card they opened first (old), but 15 months later, it is 5 percent *higher* than that on the old card. This is consistent with the notion that banks issuing new cards set prices assuming that consumers will not transfer their balances away towards their older cards.

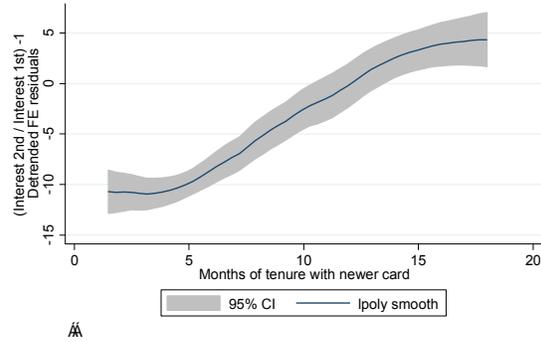
²⁸Risk variables include four lags of delinquency (failure to pay the minimum payment), credit utilization rate, and payments as a fraction of the minimum. We include these variables because they are the most predictive of default in Logit regressions in our data and in the Mexican market in general (see the methodology used by the Mexican Banking Commission (CNBV) to calculate credit card reserves).

Figure OA.36: Interest rates and Tenure

(a) Risk Adjusted monthly interest rates versus account tenure



(b) Interest rates of the old and new cards



Notes: Data are from the administrative database (2-card sample). Both figures plot Kernel regressions using and Epanechnikov kernel. Panel A plots the predicted OLS residual $\hat{\epsilon}_{i,t}$ of the regression $IntRate_{i,t} = \alpha_i + \gamma_t + \theta X_i + \delta * LaggedRiskVars_{i,t} + \epsilon_{i,t}$. It shows that ‘risk-adjusted’ prices increase with time with the card. Panel B compares relative prices of the newer versus the older card by estimating the same equation as above but using instead $([r_{new}/r_{old}] - 1)$ as the dependent variable; it plots the residuals of this regression against the tenure of the new card.

References

- Avalos, Marcos and Fausto Hernandez**, “Competencia Bancaria en Mexico,” 2006. Serie Estudios y Perspectivas Núm. 62 CEPAL.
- Banxico**, “Reporte sobre el Sistema Financiero,” Technical Report, Banco de Mexico 2006.
- , “Reporte sobre el sistema financiero,” Technical Report, Banco de Mexico 2009.
- , “Reporte sobre las condiciones de competencia en el mercado de emision de tarjetas de credito,” *Reporte Banxico*, 2013.
- Cofeco**, “Opinion, Oficio PRES-10-096-2007-110,” April 2007.
- Haber, Stephen**, “Mexico’s experiments with bank privatization and liberalization, 1991–2003,” *Journal of Banking and Finance*, 2005, (29), 2325–53.
- US**, “Credit Card Market: Economic Benefits and Industry Trends,” *International Trade Administration Report, Department of Commerce*, 2008.