The Free Installment Puzzle†

SungJin Cho
Seoul National University
John Rust
University Maryland

October, 2011

Abstract

This paper analyzes a new dataset on credit card spending decisions, with the unique feature that virtually any purchase made by its customers can be paid via *installment* over terms up to 12 months at an interest rate that is a function of the customer's credit score and the duration of the installment loan. We use these data to estimate the effect of interest rates on these consumers' demand for credit. We show that conventional econometric methods, including regression, instrumental variables, and matching estimators, predict that the demand for installment credit is an *increasing* function of the interest rate, an inference we dismiss as spurious due to the endogeneity of the interest rate and the effect of unobserved credit constraints that cause customers with worse credit scores to have higher demand for installment credit. To make more accurate inferences of the effect of the interest rate on customers' demand for credit, we exploit a novel feature in our data: customers are more or less randomly offered *free installments,* i.e. the opportunity to pay back a given purchase over a fixed term ranging from 2 to 12 months at an interest rate of *zero.* We treat these free installment opportunities as a *quasi-random experiment* and estimate a discrete choice model of installment choice that accounts for censoring (choice based sampling) in observed free installments. Despite the significant censoring, we show that it is possible to identify the probability of being offered a free installment option and consumers' probabilities of choosing various installment terms at positive interest rates, or to pay for the purchase in full at the next statement date. The *free installment puzzle* results from our finding that the average probability of being offered a free installment opportunity is 27% in our sample, while the fraction of purchases actually done under free installments is only 2.7%. Thus, we infer that the demand for credit of these customers, while downward sloping, is remarkably inelastic. In particular, our model predicts that there is only a 1 in 10 chance that a typical customer in our sample will choose a free installment option when it is offered to them.

†Preliminary draft, not for quotation or public distribution. We are grateful for feedback from presentations of earlier versions of this paper at the University of Maryland in fall 2010, at the CEAR Workshop in Denver, January 2011, and at the State University of New York at Albany and Penn State University in September 2011. Please send comments to SungJin Cho at Department of Economics, Seoul National University, Seoul, Korea, email: `sungcho@snu.ac.kr` or John Rust at Department of Economics, University of Maryland, College Park, MD 20742 e-mail: `jrust@gemini.econ.umd.edu.`

1 Introduction

This paper provides new evidence on the demand for credit from a new data set containing borrowing decisions made by a sample of customers of a major credit card company. We show that conventional reduced-form econometric approaches, including regression, instrumental variables, and matching estimators, all imply that the demand for credit is an *upward sloping* function of the interest rate charged to consumers. Of course, we believe this is a spurious finding, a likely result of unobserved factors that make consumers who have high need for credit to be charged higher interest rates than consumers who have better credit scores and other, lower cost borrowing opportunities or who are otherwise not "liquidity constrained."

To make more accurate inferences about the demand for credit, we estimate a structural model of consumer's choice over loan duration (i.e. the number of installments over which the amount borrowed is paid back) that accounts for a nonlinear increasing interest rate schedule (with higher interest rates for longer loan durations) faced by our sample of credit card customers. This model enables us to exploit quasi-random variability in the interest rates charged to consumers as a result of *interest-free installment opportunities* that arise from promotions offered by the credit card company, sometimes in conjunction with merchants. We show that our structural model provides remarkably good predictions of the borrowing decisions of our sample of consumers, and is able to control for the endogeneity of interest rates and result in a downward sloping demand for credit.

However we find that the demand for credit is remarkably inelastic and the take up rate for free installment offers is surprisingly low: we estimate that on average, the probability is less than 10% that the customers in our sample will accept a zero interest installment offer when it is presented to them. We view this as a surprising finding in light of the prevailing view that many consumers are liquidity constrained and face extremely high interest rates, and should therefore be highly responsive to zero interest borrowing opportunities. We refer to our finding as *the free installment puzzle* because our data do not enable us to provide a deeper explanation of why these consumers appear so reluctant to take zero interest installment opportunities when they are offered to them. It is also a puzzle why the credit card would place such heavy reliance on interest-free installments as a promotional device given that it appears to have such a weak impact on its customers.

While none of the customers in our sample whose credit cards are in good standing are "liquidity constrained" in the formal sense of the term, namely that the company does not impose an explicit borrowing limit on them, the conventional wisdom is that many consumers who hold credit cards at other credit card companies are liquidity constrained and are frequently spending near the maximum of their credit limits, and are therefore remarkably insensitive to the interest rates that they are charged. For example a recent paper by Alan et al. [2011] (ADL) analyze data from a randomized experiment undertaken by a British credit card company and find that "individuals who tend to utilize their credit limits fully do not reduce their demand for credit when subject to increases in interest rates as high as 3 percentage points." They interpret their finding as "evidence of binding liquidity constraints." (p. 1).

The fact that credit card borrowing is so high in most countries even though most credit card companies charge interest rates that are significantly higher than "traditional" sources of credit such as home mortgages or equity lines (for example in the U.S. the average household credit card balance is over $15,000 and the average credit card interest rate is 14.65% according to creditcard.com), could be regarded as evidence that many credit card holders are at least "credit constrained" in the sense that they either do not have access to, or have already exploited, other lower interest sources of credit and are therefore willing to borrow significant amounts on the margin at the much higher interest rates charged by credit card companies.

Thus, one possible explanation for ADL's results is that their credit card customers are liquidity constrained and "trapped" in a *corner solution* so that neither decreases nor even increases in interest rates have a measurable impact on their borrowing. However, what we find even more puzzling is that ADL found "no evidence of sensitivity to either a 1 or 3 percentage point increase (or the 3 percentage point decrease, cell 9) in our sample, even after conditioning on variables that are thought to be useful in characterizing *unconstrained* individuals." (p. 21, italics added). This suggests that demand for credit is inelastic even among individuals who are not facing binding borrowing constraints, and we regard this as a much more puzzling finding and one consistent with the new evidence we present in this paper.

The lack of sensitivity to interest rates may reflect some degree of "consumer inertia" either of the "rational inattention" variety (e.g. Sims [2003]) and or the impact of *switching and information costs* including the costs of becoming informed about other ways to borrow at lower interest rates, and switching balances to other credit cards in response to solicitations that offer consumers balance transfer opportunities at significantly lower interest rates.

This sort of inertia may explain additional types puzzling behavior observed in a different credit card data set analyzed by Ausubel and Shui [2005]. They analyzed an experiment conducted by a large U.S.

credit card company in 1995 that generated a mailing list of 600,000 consumers which was divided into six subsets with approximately 100,000 individuals each. Customers in each subset were offered (via a letter delivered by mail) the opportunity to apply for a "pre-approved" credit card from this company (including the opportunity to do balance transfers from other credit cards) at various low introductory rates for varying lengths of time. The most popular of these offers was the one offering the lowest interest rate, 4.9% for 6 months. However the response rate to these offers was uniformly small: only 1.07% of the customers offered the lowest interest rate offer actually responded and applied for the credit card, whereas the least popular offer, the one offering a 7.9% interest rate over a 12 month period, had a response rate of only 0.94% (a statistically significantly lower rate of acceptance).

Thus, while there is *prima facie* evidence of some level of consumer response to lower interest borrowing opportunities, the "take up rate" to the chance of a lower interest rate appears to be very small, and this is consistent with our findings. Ausubel and Shui describe several other puzzling aspects of the behavior of the consumers who responded to these offers. The first puzzle is one they call *rank reversal:* when they analyzed the *actual ex post* interest rate paid by customers for each of the six introductory offers over a 13 month period after the cards were adopted, the interest rate paid by customers who chose the least popular offer (7.5% for 12 months) was the *lowest* (just over 7.9%) whereas the interest rate paid by the customers who chose the most popular offer (4.9% for 6 months) was substantially higher (10.2%).

The explanation for the rank reversal that Ausubel and Shui found is that customers who chose the most popular lowest interest offer tended to behave too *optimistically* — they tended to transfer and spend more and acquire higher balances during the introductory period, but failed to pay down these balances or switch to another credit card after the 6 month introductory period ended. At that point the interest rates on their cards reverted back to the normal high annual rates the company charged customers with similar credit scores, ranging from 14 to 16%. Thus, it would appear that the individuals who responded to the most popular offer would have been better off *ex post* if they had taken the least popular offer, i.e. to have borrowed at 7.9% at 12 months rather than 4.9% for 6 months.

The rank reversal puzzle appears to be intimately connected with another puzzle, namely that once customers decided to adopt these cards and start spending on them, relatively few of these customers (40%) cancelled their accounts after the inroductory rates ended. As Shui and Ausubel note, it is puzzling why these customers were not motivated to reduce their balances or switch out of these cards when the low interest rates period expired, given that the low interest rates were evidently one of their primary motivations to switch into these cards in the first place. These results suggest that *switching costs* may be an important reason for the low response rates to the company's introductory low interest rate offers, and may explain the inertia that might be responsible for the relatively inelastic customer response to changes in interest rates overall.[1]

However the puzzle we uncover cannot be so easily ascribed to large switching costs since the ability to borrow on installment credit is an opportunity offered to customers *after* they have received their credit card and this opportunity is available for *every customer and for nearly every transaction.* Thus, there is no additional onerous "paperwork" that must be filled out to "apply" for the installment loan, and there is no issue about an installment loan being denied: these loans are essentially pre-approved and can be done by customers at the check out counter at very low marginal cost in terms of time and effort. Essentially, the customer just decides whether to pay a given balance in full, or to pay the balance in installments over a horizon the customer chooses just by telling the check out clerk or entering the installment duration on a keypad. Since installment transactions are designed to be "easy" and are not subject to credit limits (provided the customers is in good standing), our finding that customers are not very responsive to low interest rate installment opportunities (including "free installments") may be even more of a puzzle than the low response rates to low introductory interest rate opportunityies that Ausubel and Shui found in their study.

While the credit card data we have is of very high quality, it does present econometric challenges. First, unlike the ADR study which analyzes data from a *classical randomized experiment* our analysis is based on free installments as a *quasi random experiment* where we do not have a direct control and treatment group. While we can use customers as *self controls* (i.e. by comparing how much they spend on installments at positive interest rates relative to spending when they are offered free installments), we face an additional problem of *censoring*, namely, we only know when customers are offered free installment purchase opportunities when they actually choose them. How is it that we can infer the probability that customers are offered an free installments and the probability that customers will decline a free installment when it is offered? We show that these probabilities can be identified by the method of maximum likelihood using a *mixture likelihood* that results from treating installment purchase decisions as *individual choice experiments* where the presence of a free installment opportunity results in *unobserved choice*

[1] Shui and Ausubel argue that switching costs alone cannot fully explain the puzzles they find: they argue that the puzzling behavior of the customers they studied is best described by a hyperbolic discounting model than it is by a time-consistent dynamic programming model in the presence of switching costs.

sets. We show there is sufficient structure that it is possible to separately identify the probability of being offered a free installment opportunity from the probability of choosing it.

Section 2 provides an overview of the competitive environment that the credit card company whose data we study is operating in, and discuss the strategic motivations for why it offers free installment options to its customers. Section 3 describes the credit card data and documents the importance of merchant fees as a significant component of the profit that this company earns: we believe this is the main motivation for the company's frequent use of free installments. Section 5 introduces the econometric methods we employed to infer the demand for credit including regression-based and reduced-form treatment effect approaches. We also describe our "semi-structural" model of installment choice. The empirical findings from the reduced form methods are presented in section 6. We show that none of these methods result in plausible estimates of the demand for credit. In particular, all of the methods lead to the conclusion that the demand for credit is an *increasing* function of the interest rate. Section 7 derives the likelihood function for the semi-structural model, establishes the identification of the structural parameters, and presents the estimation results, including an evaluation of the goodness of fit of the model and the predicted installment credit demand function, as well as several counterfactual predictions of customer response to alternative installment credit policies. Section 8 presents our conclusions and speculative comments about the underlying reasons for the free installment puzzle, as well as suggestions for future research provided additional data and particularly new experimental data could be gathered.

2 The Strategic Role of Free Installments

We analyze a new data set containing credit card transactions and payments for a sample of 938 individuals who made purchases using their credit cards over a three year period from Janary 1, 2004 to May 1, 2007. These individuals had credit card accounts with a large bank that issued multiple credit cards. For confidentiality reasons we are unable to reveal the name of the credit card company or even the country in which it operates.

A unique feature of this credit card and therefore the data that we analyze is that *all customers who are in good standing have the option to make any purchase under installment credit.* That is, customers of this company make individual *transaction by transaction* decisions on whether to pay the amount of any purchase in full at the next statement date, or to spread the purchase out over multiple installments ranging from 2 months to 12 months, but at the cost (usually) of paying a relatively high interest rate for this borrowing opportunity. However, largely for "promotional" reasons, the company offers a significant fraction of its customers the opportunity to make *free installments,* that is, some customers are occasionally offered the chance to borrow on installment *over a fixed term at a zero percent interest rate.*

The company faces strong competition for customers from other credit card companies. Each of the firms in this market confronts difficult but very interesting strategic decisions relating to their choice of *credit policy* (i.e. which individuals to give credit card accounts to, which interest rates to charge them, which credit limits to set, and other policies such as the type of installment and revolving credit opportunities it provides), their *merchant policy* (i.e. the contracts it offers merchants including the level of transaction fees it charges to merchants who accept the company's credit card and any restrictions in the contracts it offers these merchants such as reimbursement in cases of fraud) and their *advertising policy* (i.e how much to spend on advertising to try to attract new customers, and how much to spend trying to convince new merchants to sign an agreement to accept payments from the company's credit card, the level of "rewards" it offers its customers and so forth).

In this paper we focus on the company's use of credit policy, and particularly on the company's use of *free installments* as a strategy for attracting new customers and encouraging existing customers to increase their use of the company's credit card. To large extent, the company is engaged in a race with its competitors to expand its market share. The larger the company's credit card market share, the more likely merchants are to to accept the company's credit card. The larger the rate of merchant acceptance, the more valuable the company's card is to consumers, which therefore tends to help the company add more customers and more merchants over time.

Thus, there is a strong element of "increasing returns" and "network externalities" to gaining a dominant market share. Ultimately, the firms that have the largest credit card market shares will be able to charge higher fees to merchants. As we show below these merchant fees are a major component of the company's revenues and profits, amounting to 36% of the revenues earned for the sample of customers we analyzed in this paper. That merchant fees are a major source of revenue for credit card companies is not a new finding, however. Evans and Schmalensee [2005] note that merchant fees account for nearly 65% of the revenues earned by American Express. We believe that the company's recognition of the importance of merchant fees as a source of revenue is the primary motivation for its "benevolence" in offering free installment opportunities to its customers. See also Rysman [2007] who analyzes Visa data that enables him to study credit card spending for consumers who hold multiple credit cards simultaneously. Rysman finds evidence "suggestive of the existence of a positive feedback loop between consumer usage and merchant acceptance." (p. 1) and this feedback loop is likely a main motivating factor for credit card companies' huge advertising spending to promote card use. Free installments is just one one of many ways that the company we study uses to try to encourage higher use of its credit cards.

The company we study is unique relative to many other U.S. based credit card companies in offering both *installment credit* and *revolving credit* options for its customers. Installment credit is made on a transaction by transaction basis: a customer who purchases an item using the company's credit card can decide to make the purchase as a *regular charge* which means the amount charged will be payable at the next statement date, or as an *installment charge* which means that the balance charged can be paid in multiple installments over periods ranging from two to twelve months. If a customer chooses the latter option, the customer can also choose the term of the installment (i.e. whether to pay back the amount charged in 2, 3 or up to twelve months). The customer does this at the time of the purchase, i.e. at the "check out counter" and usually at a significiant interest rate that is a function both of the customer's characteristics (including the customer's credit score) and the term of the installment loan. The average installment interest rate for the more than 6200 installment purchase transactions we observe for our sample of customers is 15%. Any customer who is in good standing (i.e. whose right to use their card is not suspended due to repeated refusal or inability to pay their balance due) has access to the installment credit option.

In mid 2005 the company offered a new *revolving credit* option to a subset of its "preferred customers" with the best credit scores and highest card usage. Normally, customers must pay their full credit card balance due at each statement date, otherwise a late fee is levied and their credit score is penalized. However for those customers who have the revolving credit option, they can choose to pay only part of the balance due and the remaining balance can be paid off at some subsequent statement date, along with accumulated interest. The company also offers related credit options such as *revolving cash advances* that can be paid back at the discretion of the customer. The interest rates for these revolving credit options are substantially higher than the installment credit options discussed above. For example, the mean interest rate on revolving cash advances for our sample is 27%.

Because revolving credit is a relatively new option for the customers of this company and was not universally offered to all customers, we chose to focus our analysis on the decisions customers make regarding installment credit. Since each customer faces this decision at essentially *every* purchase occasion (charges incurred at restaurants, hotels, airline ticket, doctor bills, consumer electronics, and even grocery bills can be paid by installment), we obtain a large number, *over 167,000,* of customer/purchase observations in our data set. Each of these purchases can be therefore be regarded as individual "micro borrowing/intertemporal choice decisions", i.e. in each one the customer is making choice whether or not to pay the charge in full at the next statement date, or to pay the amount off gradually over some number of monthly installments ranging from two to twelve, at an interest rate known to each customer that is a function of the customer's characteristics and the duration of the installment term. In particular, we show below that the interest rate schedule each customer faces is a *sharply monontonically rising function of the duration of the installment agreement.*

To understand the strategic role of free installment options, it is useful to provide some background on the competitive enviroment that the credit card company we are studying is operating under. The credit card market is highly competitive with nearly 20 banks and non-bank credit card companies offering competing credit cards. Competition for market share was particularly intense in the period just prior to our observations, from the late 1990s to 2002 when a combination of factors, including government policy favoring credit cards to improve tax collection (including paying taxes by credit card, and reducing tax evasion in the hard to monitor "cash economy") and the entry of new credit card issuers intent on capturing a larger share of the market lead to dramatic increase in the number of individuals using credit cards, and in overall credit card spending. At the peak of the credit card "boom" in 2002, the average credit card customer had more than 3 credit cards, average credit card balances were in excess of $2000 per capita, and aggregate credit card debt amounted to nearly 15% of GDP.

Much of the aggressive expansion of credit card accounts and unsecured lending by the new non-bank entrants proved unwise and in 2003, the year preceding most of our data, there was a significant "credit card bust" with default rates in the credit card industry as a whole exceeding 25%. This lead to massive losses in the financial sector, several near bankruptcies of major banks and major non-bank companies that entered the credit card market, and a government bailout to prevent a wider financial panic from ensuing. The move was largely successful and in combination with adoption of better risk-management policies at the major credit card companies, average credit card balances and default rates rates declined rapidly after 2003. By 2005 the credit card default rate had fallen by more than 50% to just over 10%, per capita credit card balances had fallen to less than $700, or about one third of their peak in 2002 just before the crisis, and credit card debt as a fraction of GDP had fallen to a much more reasonable level of approximately 4%. By 2007, the last year of our data, the default rate on credit cards had fallen to less than 4%, roughly comparable to credit card default rates in other OECD countries.

The credit card company that provided us the data we analyzed is a bank with relatively long experience in credit card lending, and it appears to have been much more conservative in its credit card issuance and lending policies than many of the new credit card entrants that contributed to the boom and bust cycle discussed above. In particular, to the extent we can measure in our sample of 938 customers, the credit card default rate for this company is much lower than for the industry as a whole, and did not change radically over the years 2004 to 2007.

Nevertheless, the company is keenly aware of the importance of increasing its customer base and spending market share, since as we will show below, merchant fees are a significant source of its revenues. As we noted above, credit card market share is a very important factor in merchant acceptance of credit cards, and also in the degree of market power the company has over its merchant fees.

In 2006 the company was among the top six largest credit card companies in the country in terms of credit card spending market share. The dominant firm had approximately a 25% market share in 2006 and 2007, and the combined market share of the six largest credit card companies was approximately 85% in each of these years. The importance of having a large market share is especially enhanced in due to the nature of electronic fund transfers between credit card customers and merchants which generates high returns to scale and "network externalities" for companies that have the largest market share.

Specifically, credit card companies need to have explicit agreements with individual merchants in order to have their credit card accepted, otherwise they need to have a partnership or agreement with another credit card company that has an agreement with the merchant. In the latter situation, the credit card company pays one of its competitors that does have an agreement with the merchant a fee to process a transaction over the competitor's network and under the competitor's merchant agreement. It is surprising that there is cooperation between competing firms where one firm allows a customer of competitor to carry out a transaction on its network when the merchant in question has an agreement with the company but not its competitor. One could imagine a different equilibrium where firms did not cooperate and excluded rivals from using their networks. However it appears that so far no single company has a sufficiently large share of agreements with all merchants to make exclusion an individually rational strategy. The collective benefit from pooling the set of merchant agreements, each of which is rather costly for any firm to acquire, seems to be sufficiently high that reciprocity rather than exclusion has emerged as the equilibrium in this market, perhaps much the same way as airlines have a reciprocal agreements to have their customers book flights on their competitors' plans in the event of cancellation due mechanical problems with airplanes.

However the level of reciprocity is informal and bilateral: there is no analog of the much more well developed *multi-party payment systems* that exist in the United States that began in the 1960s as cooperatives that "in effect, pool the merchants they had signed up so that any individual with a card from a member of the cooperative could use their card at any merchant also signed up by any member of the cooperative." (Evans and Schmalensee [2005]). As a result, in the U.S. market, each merchant does not need to make individual contracts with all credit card issuers (e.g. banks), instead merchants need only make contracts with a smaller number of credit card *brands*, primarily Visa, MasterCard, American Express, and Discover. This is not the situation in the country we are studying in this paper, and the presence of a MasterCard or Visa logo does not necessarily confer any advantage since merchants do not sign agreements with MasterCard or Visa directly, but only with each bank or credit card separately. The main advantage of a MasterCard or Visa logo is the ability to use the card for *international* transactions, but it does not confer any special advantage for domestic transactions.

In addition to the standard internationally branded logos there are also "private brand" credit cards and logos in this country. The private brand domestic credit cards gennerally charge lower merchant fees and have lower annual fees. However, some of the credit card companies issue credit cards which are a combination of the private brand/logo and a Visa or MasterCard logo. This happens especially on credit cards issued by banks or new entrants to credit card market, since they do not have existing payment network and part of the value of having the Visa or MasterCard logo is that it signals widespread acceptance of the card (including the ability to use the card internationally) and this attracts more customers and could help, in a positive feedback loop, to convince more merchants to accept cards with a Visa or MasterCard logos relative to cards that do not have these logos.

We do not directly observe the fees that the credit card company we are studying pays to its competitors for use of their networks in our data. However we do observe the frequency of these "out of network" transactions, since they are identified with a merchant code of -1 in our sales data. Of the 182,742 sales transactions in our data set, a total of 32,299 or 17.7% of the sales were 'out of network'. The average merchant fee recorded for these out of network transactions was 1.24% which is substantially less than the average 2.02% merchant fee for the company's "in network" transactions in our sales data set. A total of 45.5% of the out of network transactions in our data have a recorded merchant fee of 0, whereas the average merchant fee for the 17,601 out of network transactions for which a merchant fee was recorded was 2.28%, or about the same average merchant fee (as a percent of amount sold) as the company earns on its in-network transactions. It is not clear to us whether the large incidence of out of network transactions for which the company records no merchant fee constitute transactions where a competing firm that carries out the transaction charges a fee equal to the company's merchant fee, and whether on the out of network transactions where the company does record a merchant fee there is a separate payment to the competitor for the use of their network/merchant agreement to carry out the transaction that we do not observe in our data. However the general conclusion is that due to the way the fund transfer system works in this country, the credit card company has a very strong incentive to grow and increase its own network and set of agreements with merchants since this enables it to earn significantly higher merchant fees.

Like many other countries, merchant fees vary widely across different types of merchants. In a separate study, we will describe the nature of variation in merchant rates in more detail. However for the purposes of this study, we do show below that merchant fees are a significant component of both revenues and profits for this company, and it is our belief that a primary motivation for the use of free installments is as a promotional device to attract new customers and increase sales and usage of the company's cards by its existing customers. By doing this, the company expects to increase its market share, and thus merchant acceptance of its cards.

3 Credit Card Data

We were fortunate to obtain a data set from a credit card company that provided us information on the credit scores, purchases, payments, and credit decisions of 938 of its customers over the period 2004 to 2007. Our data consist of six data files: sales, billing, revolving and collection, credit rating, and a final fiile defining merchant the classification codes that appear in the sales data. For sales data, we should note that there are three types of sales 1) sales payable in full at the next statement date, 2) sales payable in installments over two or more statement dates, and 3) cash advances. Cash advances can either be paid in full at the next statement date, or paid by installment over multiple future statements. Generally purchases and cash advances that are paid by installment are done at relatively high interest rates, except when customers are offered free installment options.

The primary focus of this paper is to understand how customers decide whether to pay for individual purchases as a "regular purchase" (i.e. as payable at the next statement date to which the transaction is assigned) or as an installment purchase in which case the payment is spread out over 2 to 12 future statement dates. We are particularly focused on identifying the effect of the installment interest rate on the customer's choice of installment term. Although the availability of installment credit can potentially affect the customer's decision whether to purchase a given item or not, or to purchase via credit versus cash or some other credit card, as we discuss below, our data are of limited usefulness for studying these other, related effects on interest rates on spending and credit card usage decisions.

In our data we observe installment purchases of varying lengths, from 2 to 12 months. The most commonly chosen term is 3 months: 61.5% of all of the installment purchases we observe have a 3 month term. The maximum installment term we observe is 12 months, and is chosen in 1.7% of the cases. Other frequently chosen terms are 2 months (20.0% of cases), 5 months (5.0%), 6 months (4.9%), and 10 months (3.7%). There are no installment purchases with a term of 1 month, since this is equivalent to a regular charge, i.e. a payment due at the next billing statement. Thus, we define the "installment choice set" for a consumer as being $C = \{1, 2, \dots, 12\}$ where a choice of $c = 1$ is equivalent to a regular charge that will be due at the next billing statement, a choice of $c = 2$ corresponds to equal installments payable in the next two billing statements, and so forth, so that $c = 12$ denotes an installment contract that is payable over the next 12 billing statements (monthly) in 12 installments.

Customers typically pay off their installment purchases in equal installment amounts. For example, if a consumer purchases an amount P under an installment contract with a total of c installments payments, then the consumer will pay back the "principal" P in c equal installments of P/c over the next c billing periods. If the consumer is charged interest for this installment purchase, the credit card company levies additional interest charges that are due and payable along with the installment payment at each of the successive c statement dates. However in some cases there are unequal payments, sometimes as a result of late payments, or accelerated or pre-payment of installments. The installment agreement does not formally allow for a pre-payment option, so that if a consumer does pre-pay an installment contract, the credit card company still charges the interest at the successive c statement dates, as if the customer had not pre-paid.

We calculated the realized rates of internal rate of return on 8987 installment transactions in our credit card data set. The internal rate of return is the interest rate r that sets the net present value of the stream of cash flows involved in the installment transaction to 0, where the initial purchase is regarded as a cash outflow (from the credit card company) at time $t = 0$, and the successive payments (including interest) are treated as cash inflows at the successive statement dates $t_1, t_2, \ldots, t_c$. There were only 141 cases out of the 8987 installment transactions where the customer did not follow the original installment contract by paying in the c installments that the customer orginally agreed to pay. There were pre-payments in 127 cases, i.e. where the customer paid off the installment balance more quickly than necessary under the original installment agreement. Given that there is no direct benefit to the customer from pre-paying the installment (since the credit card company will continue to collect interest from the customer as if the installment loan had not been pre-paid), it seems hard to rationalize these cases under a standard model of a rational, well-informed consumer. In 31 of these cases, the customer was given the a 0% installment loan, and yet still pre-paid. One possible explanation is that these customers were not aware that they had what was in effect an interest-free loan, and not aware that there was no benefit to pre-paying. These customers might have believed (incorrectly) that by paying off their installment balance more quickly they were saving interest charges, or perhaps some other explanation such as "mental accounting" (e.g. the desire to be free of the mental burden of having a large outstanding installment balance to pay), might explain this behavior.

There were only 17 cases where the number of installment payments were greater than the number of installments originally agreed to in the oringal installment transactions. These do not appear to be "defaults" since the total amount collected in each of these cases equals the initial amount purchase. The delay in payment was typically only one billing cycle more than the orginally agreed number of installments. For this reason, we believe that these cases might reflect the effect of holidays (such as where a payment is allowed to be skipped since a statement falls on a special holiday) or some other reason (e.g. an agreed *ex post* modification in the installment agreement). Since there are so few of these cases, we basically ignore them in the analysis below.

In the data we observe most installment purchases have a positive internal rate of return, but in nearly half of all installment purchases we observed (47.7%) the internal rate of return was 0, so the customers were in effect given an interest-free loan by the credit card company. These "zero interest installments" are usually a result of special promotions that are provided either at the level of individual merchants (via agreement with the credit card company to help promote sales at particular merchants via the "free credit" aspect of an installment purchase with a 0% interest rate), or via "general offers" that the credit card compamy offers to selected customers during specific periods of time either to encourage more spending, increased customer loyalty, or as a promotion to attract new customers. Our data does not contain enough information for us to determine exactly which customers are offered 0% installment options, so we model it as occurring probabilitistically, depending on the characteristics of the customer, the merchant code at where the customer is making a purchase, and dummies for the partical time (since some of these promotions where offered at specific points in time). The vast majority of interest-free installment loans have a term of 6 months or less. If a customer wishes to have a longer term than the one being offered, the customer generally must pay a positive interest rate for longer term installments, according to the schedule described below. In our analysis below, we will assume that when a customer is offered a interest-free installment purchase option, the maximum term is exogenously specified according to a probability distribution that we will estimate from our data.

In order make customer-specific profit and rate of return calculations and analyze time patterns of credit card spending and installment usage, we had to assemble the data that were contained on customers in the sales, billing, and collections tables into a *longitudinal format* that would enable us to track the evolution of both credit card and installment balances on a *day by day basis.* We emphasize that the credit card company did not provide us with these latter data, rather we had to *construct the longitudinal data from the information we were provided.* While at first it may seem to be a relatively trivial exercise in stock/flow accounting to reconstruct these *balance histories* from the sales, billing and collection data, we faced a signficant *initial conditions problem.* That is, we were not given the outstanding installment and credit card balances at any initial date. Instead the collections table would tell us the *statement amount* and information on dates of collection and amounts received, but without knowing an initial balance, it was not always easy to determine if a customer had paid the initial statement or any previous statements in full, or had unpaid balances that needed to be carried over from previous statement dates. We could obtain some indirect evidence of the presence of such overdue balances from late fees charged, but without going into more detail, it proved to be a rather challenging accounting exercise to infer the initial balances of the customers in our sample accounting for the variable left and right censoring in the data.

In particular, not all sales records in the sales table could be matched with billing records in the billing table and vice versa. In some cases, we observed purchases that were at a date before any date in the billing table, and we also observed billing records for which we could not find a corresponding record in the sales table. Fortunately the billing table had redundant information on whether the transaction was on installment or not, so in most cases we could reconstruct an entire installment transaction even if we only observed a truncated series of installment payments in the billing record and no record of the initial sale in the sales table.

Similarly there were also problems of right censoring in our data, since in many cases we observe sales in April 2007 for which we had no corresponding billing records, or no collection records at the end of a balance history that would enable us to determine whether an outstanding balance would be fully paid at the next (yet to be observed) statement date that was missing in the collection table. In such cases after making the best inference on the value of the customer's initial balance at the start of the interval we observed the customer, we followed the customer for as long as possible where we could also match every sale with its corresponding record in the billing table and track payments received on balances due in the collections table. In some cases this required us to "back up" by one or more months on the full history of the customer and discard transactions in the last month that we could not provide the corresponding matches in the billing table and a record of payment in the collections table.

However, overall, our care in preparing the data paid off and we did not lose too many observations by doing this and the result is a considerably more accurate record for making profit/loss calculations on a customer by customer basis. If we did not do this, customers would be artificially classified as being in deficit if a balance due happened not to have been recorded for them in the collections table due to right censoring. Thus, we would end a record on a customer on a date where a balance due was received and for which all previous charges up to that date had been accounted for. Any subsequent charges that were made by the customer that would be billed and paid for in the future but which we could not yet observe in the billing or collections tables were discarded in our analyses of customer level profitability and returns.

Figure 1 plots our constructed longitudinal balance histories for one of the customers in our data set. We chose this example because the customer made only a single installment transaction and this makes it very easy to understand how the constructed balance histories behave. The top left panel of figure 13 is the overall creditcard balance for this customer. We start observing this customer making a charge of \$118.30 on December 12, 2003. However we did not know what the outstanding balance was for this customer at this date since the first statement date for the customers was on January 20, 2004. We were able to determine in this case that this customer had no outstanding unpaid balances and we were able to allocate all charges the customer made in the sales table to matching entries in the billing table and thus track this customer with an accurate determination of the customer's initial balance at the first installment date. Thus, the top right panel of figure 1 displays our inferred balance for this customer, \$427.24, on the

Figure 1 Balance and credit history of customer 125

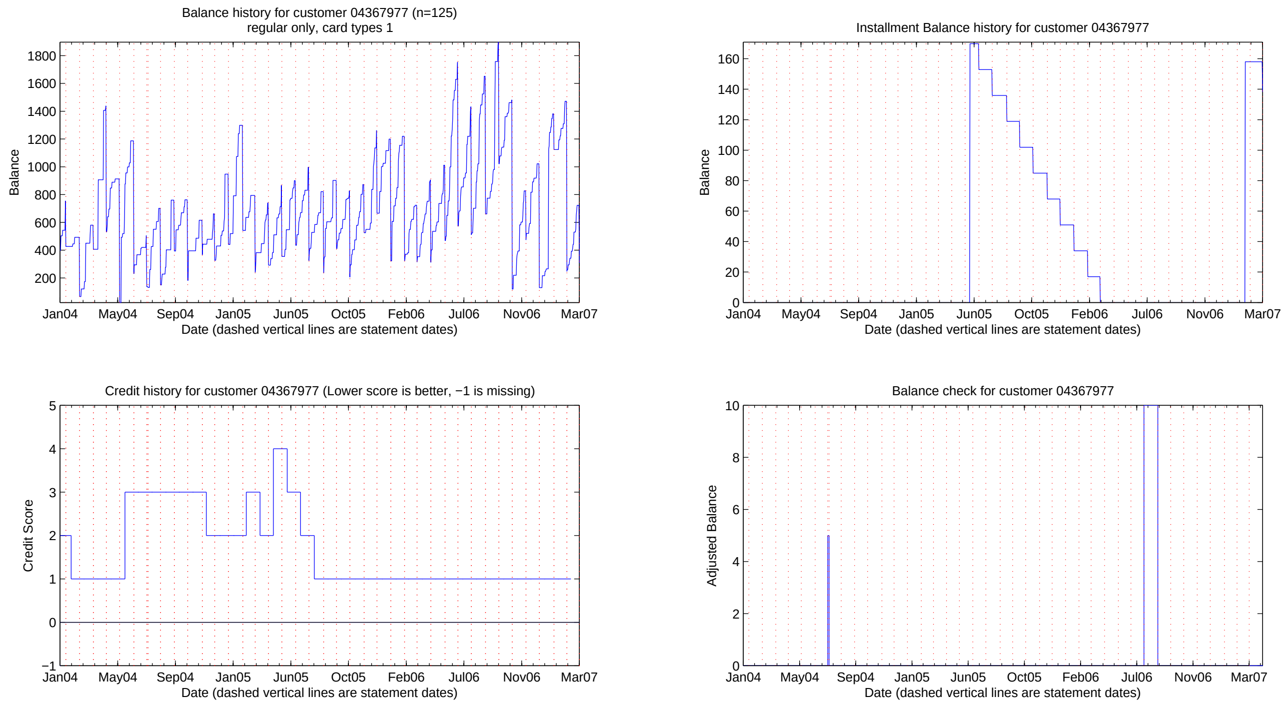

first statement date we observe for this customer, January 20, 2004.

The dashed vertical lines in the figures represent the statement dates. Because this company has links to its customers' bank accounts and auto-debits the statement due amount on each statement date, its customers almost always pay the full balance due *exactly* on each statement date, unlike for many American credit card companies where customers may mail in a check or pay online and the date paid may often be plus or minus the statement date by several days. Thus, this feature leads to the inverted sawtooth appearance of balances in the top right hand panel of figure 1: balances tend to grow monotonically (though stochastically) between successive statement dates representing the spending the customer is doing on their credit card, then it drops discontinuously on each statement date representing the payment of the balance due.

Note that the discontinous drops in the credit card balance at each statement date do not bring balances exactly to zero. The reason is that the credit card company assigns to each purchase a particular statement date at which that purchase will be due (unless it is an installment, which leads to a different treatment we will discuss shortly) and thus typically any purchases a customer makes that are sufficiently close to an upcoming statement date will be assigned as due and payable by the company to the *following* statement date. Thus, the level of credit card balances just after a statement date reflects the sum of all purchases made prior to that statement date that the company assigned to be due and payable at the next statement date. This implies that a person's credit card balance will almost never be exactly zero, even on a statement date — at least for customers who are sufficiently active users of their credit card.

Note the "balance check" in the lower right panel of figure 1. The balance check should be identically zero if we had correctly inferred the customer's initial balance and perfectly tracked all charges and fees. However there were some small charges and payments that we could not reconcile or ascribe to any late charge, annual fee or so forth. These appear as the spikes in the lower right panel of figure 1. In some cases the balance check will be non zero due to a pre-payment or some slightly mis-timed or out of sync payment but shortly after the balance check returns to zero showing that we have basically correctly calculated the full balance history for this customer.

Now consider the top right panel of figure 1, which shows the *installment balance history* for the customer. We keep two separate accounts for the customer, 1) the credit card balance and 2) the installment balance. In this case, we see that the customer did not charge anything on installment until May 31, 2005 when the customer made an installment purchase in the amount of $169.90. This is reflected by the discontinuous upward jump in the installment balance in the top right panel of figure 1. We can see from the graph that this balance was paid off in 10 equal installments of $16.99. This installment also happened to be an interest-free installment and so at each of the 10 succeeding statement dates after the item was purchased on May 31, 2005 the installment balance decreased by $16,99 until the balance was entirely paid off at the statement date of March 20, 2006. Note that on each such statement date, the amount currently due on the customer's installment balance *transfers* and is added to the customer's credit card balance.

The final, lower left panel of figure 1 plots the credit score that the company maintained on this customer. Credit scores are integers on a scale from 1 to 10 with 1 being the best possible credit score and 10 being the worst. This customer generally had excellent credit scores, though for reasons that are not entirely clear from figure 1, the customer had periods of time (particularly May to September 2004 and May to July 2005) where the credit score deteriorated for some reason. We see that the customer's worst credit scores appear to have coincided with the customer's installment purchase in May 2005.

We present another balance history for a more interesting customer, customer 809, in figure 2. This customer generally maintained larger credit card balances and also larger installment balances than cus-

Figure 2 Balance and credit history of customer 809

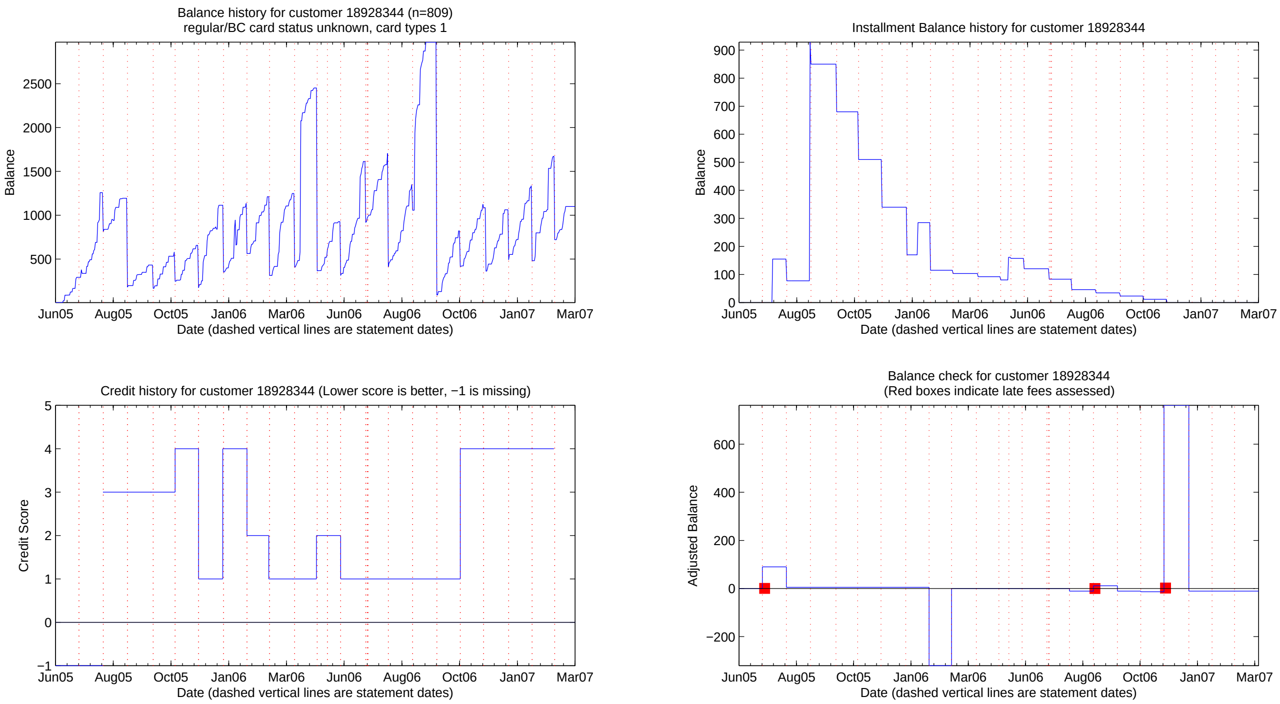

tomer 125, and we see that this customer also tends to have uniformly worse credit scores than customer 125 had. The red boxes in the lower right panel of figure 2 also indicate another behavior that is a big "no-no" for the credit card company: the customer was late in making payments and assessed late payments on three occasions. Because balances due are automatically debited from the customer's bank account, this means that on these three occasions the customer's bank account became *overdrawn* and the credit card company was unable to collect the full statement amount due. While the customer may have also been charged penalties for overdrawn account by his/her bank, the late payment penalties charged by this credit card on these three occasions was trivially small by American standards: \$0.18 in each case. The main penalty seems to be a degradation of the credit score, though the late fee of \$0.45 that the customer was assessed on September 4, 2006 did not seem to have any effect on the credit score around that time.

Now that we have shown how we were able to construct the spending and payment patterns and thus the balances histories of our sample of customers dynamically, we are now in a position to calculate returns and profitability on a *customer by customer basis.* In terms of profits, we can think of the primary cost of a customer is the company's *cost of credit,* i.e. the credit card company's borrowing cost or opportunity cost

Figure 3 Credit Card Balances, Average Balances and Observations in the Full Sample

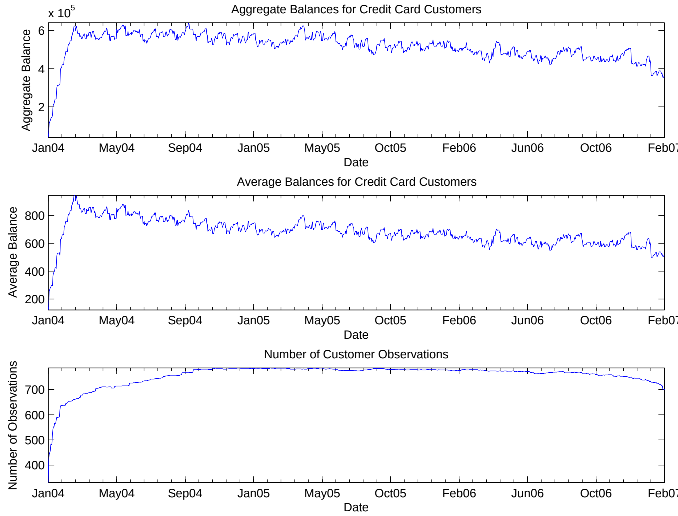

of capital. In the case of customers who default, the company also loses the unpaid balance of their load to the customer. The revenues include annual fees, late fees, interest and service charges, and merchant fees. We note that our measure is one of *gross profits*, i.e. we do not know the cost of things such as 1) rewards programs, 2) advertising costs, and 3) other fixed operating costs such as billing and collection costs and wages and salaries and payments to other credit card companies for out of network transactions.

Before we present these results, we start with figure 3 which gives an overview of credit card usage, total balances, and number of customer observatioins for our full sample. We see a slow downward trend in credit card balances over our sample period, declining from about \$800 at the beginning of our sample in 2004 and declining to \$600 at the end of our sample period in early 2007. Figure 4 shows the evolution of installment balances, which show an opposite, increasing trend over the sample period, with balances increasing from about \$400 per person at the start of our sample period to over \$700 by early 2007. The company data allow us to track installment and overall credit card balances separately. Note that overall credit card balances are *not* simply equal to installment balances plus balances for other non-installment balances for reasons we will explain shortly.

Installment credit predates the use of *revolving credit* which was introduced by this company in 2005 to a subset of its most profitable and creditworthy customers. Revolving credit differs from installment credit in that installment credit decisions are made on a *transaction by transaction basis.* Specifically, every time a customer holding this credit card uses their credit card to make a purchase, the cashier asks

Figure 4 Installment Balances, Average Balances and Observations in the Full Sample

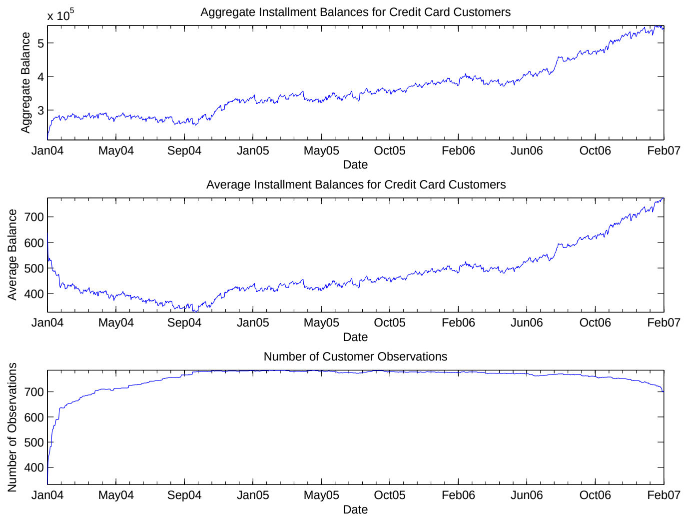

the customer whether the customer would like to pay by installment or to pay in full. The latter means that the amount purchased will be due and payable in full at the next *statement date* associated with the transaction. For each such transaction, the company assigns a specific statement date. If the transaction is done sufficiently in advance of the customer's next statement due date, the amount purchased will be assigned to that statement, otherwise if the purchase occurs too close to the next statement date, then the purchase will be assigned to the following statement date. It turns out that the majority of purchases, 63%, are due at a statement date that is more than 30 days from the date of purchase. From our data, the average delay between the date of purchase and the statement date at which the amount purchased is due is 50 days.

Credit card users who decide to purchase on installment also need to decide the term of the installment, which range from 2 to 12 months, although very infrequently installments over longer terms are possible. In our billing data, we have 11175 observations on installment purchases, but only 13 of these were for terms that were longer than 12 months. In each of these 13 cases, the number of installments was 24. We have decided to ignore these 13 cases as highly atypical of the installment purchases we observe, and because it simplifies econometric analysis in section 4 to consider installment durations as limited to between 2 and 12 months.

Each installment purchase specifies essentially equal repayments of the principal amount borrowed over the term of the loan, plus interest payable on the amount of the outstanding installment balance.

Thus, interest payments are largest at the first installment date and decline thereafter as the remaining installment balance declines. Our data provides sufficient level of detail including a unique transaction identifier called *NSS* (number of the sales slip) that we can track individual sales and installments in our data. There are separate billing records corresponding to each payment date of an installment purchase as identified by its NSS number. This enables us to compute the interest rate on the installment as the *internal rate of return*, i.e. the interest rate that sets the present value of future installment payments and interest payments equal to the amount initially purchased on the initial sales date.

As we noted above, installments are typically decided upon at the time of purchase, where the customer notifies the cashier of their intention to have the purchase be done on installment over a term decided on by the customer. The interest rate for the installment is typically not available at transaction time, though customers are informed of their installment interest rates on their statements and via their accounts on the company's web site. In situations where the customer is offered a free installment opportunity, the cashier will typically inform the customer at the time of purchase. The free installment term is almost always determined as part of the free installment offer, and thus is not a variable that the customer can choose (unlike the case of positive interest rate installments). The most common terms of free installments is 3 installments, though other common terms are 2, 5, 6 and 10 months.

Free installments are also made available to the company's customers for limited periods of time announced on the company's web site, or in flyers or ads that are included in paper bills that are mailed to its customers. Our presumption is that such offers are *universal* except for customers who are not in good standing, i.e. customers whose accounts have been classified as "in collection" for having unpaid balances for more than 6 months. In such circumstances, the customer must remember to tell the cashier to put the purchase on installment at checkout time, and specify that they are taking advantage of a free installment opportunity so they will not be listed as choosing to pay a positive installment rate for a term longer than the one offered under the announced free installment promotion.

Figure 5 plots the distributions of installment terms for 4700 installment transactions made by customers that chose installment with positive interest rates, and also the distribution of installment terms offered to 4287 customers who chose to exercise free installment options. The distributions are roughly similar except that the mean installment term chosen by customers under positive interest installments, 3.66 payments/months, is longer than the 3.42 payments/months offered to customers who chose free installment options. We see that when customers choose installments with a positive interest rate, they are

Figure 5 Durations of Free and Non-Free Installment Loans

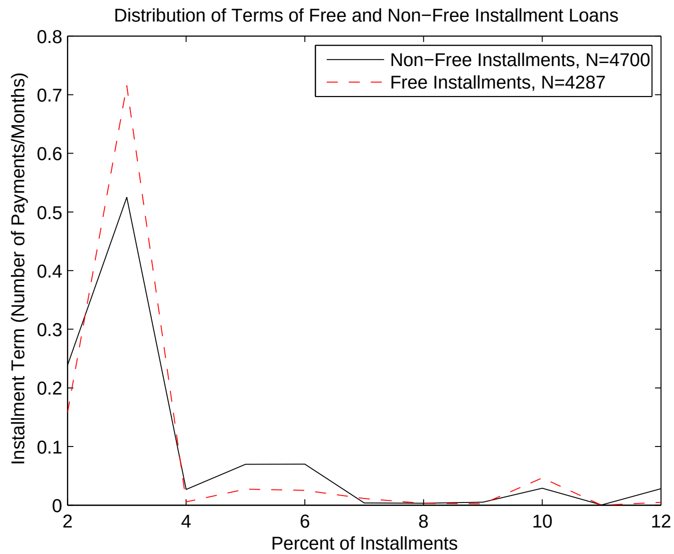

generally more likely to choose longer payment terms, though the difference in the two distributions is not particularly striking.

Note that due to censoring we are not always able to observe the full duration of installment transactions. For example we observe some installment NSS codes in our billing data for which the date of the initial installment purchase is not in our sales table. This is why, although we can identify 11175 installment transactions in our billing data, when we eliminate censored observations we obtain a smaller set of 8987 *uncensored* observations of installments where we can match the transaction NSS in the billing table to the NSS of the original sale in the sales table. The reason we want to make such matches is because the information on the merchant fee charged is only available in the sales table, not in the billing table. As we will show below, the merchant fee contributes a huge amount to the overall rate of return that the credit card company earns on installments. However the rates or return on installments quoted above are *net* of the merchant fee. That is, these are the effective rates of interest that the customer paid for the installment loan. The company earns a much larger rate of return when we also factor in the merchant fee it earns at the time of the installment transaction.

In addition to installments, the company allows its customers to borrow on *cash advances.* We observe 11,818 such transactions in our data. These are typically of shorter duration than installments: the average duration of a cash advance is 45 days. The interest rates for such loans is also typically higher than for installments: it averages 24% compared to an average of 15% for installment transactions that are done at a positive interest rate (i.e. excluding the free installment transactions). The average amount of cash advances, $734, is more than twice as high as the average installment purchase done at a positive interest rate, $352. However this ranking is reversed in the upper tails of the distributions of purchases and cash advances: the largest cash advance in our data was $8300 whereas the largest installment purchase done at a positive interest rate was $15,740.

Because the motives for cash advances are likely to be different than for installment purchases and because cash advance terms are shorter and because zero interest cash advance opportunities were not offered to the the company's customers (at least in our data for our sample of customers) we have chosen to limit our analysis to the choice of installment term and leave the analysis of cash advances to future work.

For all of sales (credit card purchase) observations we have (both for sales done as installment and not under installment) comprehensive purchasing data including customer IDs, types of credit cards (regular card, gold card, platinum card, debit card, check card, and etc), NSS (number of the sales slip, the unique identifier for each transaction discussed above), the type of sales (including whether the sale is a return or reversal or cancellation), the date of sale (both the date of the actual sale and the date it was "posted" to the credit card), the merchant fee earned by the credit card company, and a code for the merchant type, which will be -1 for merchants that are not "in network" (i.e. for which the credit card company does not have a formal merchant agreement but does the transaction via a competing credit card's network and merchant agreement as discussed above). The sales data also include the chosen term of the installment if the purchase was an installment sales transaction, and the up-front cash advance fees in case of cash advance transactions. Overall, we have a total of 182,742 observations for 884 customers. The average number of transactions per customer is therefore approximately 206. Figures 6, 7 and 8 below present the distribution of the transaction amounts or ordinary (non-installment) sales, installment purchases done at a zero interest rate, and installment purchases done at a positive interest rate.

We see that, as expected, the average installment purchases are sigificantly larger than the average non-installment purchase: on average interest-free installments are four times larger and positive interest installments are seven times larger than ordinary credit card purchases. However already we can see the *free installment puzzle* in figures 7 and 8: the average size of a positive interest rate installment is more than 75% larger than the average installment done under a zero interest rate. Economic intuition would suggest that installments done at a lower interest rate, and particularly at a *zero* interest rate should be

Figure 6 Distribution of non-installment credit card purchases

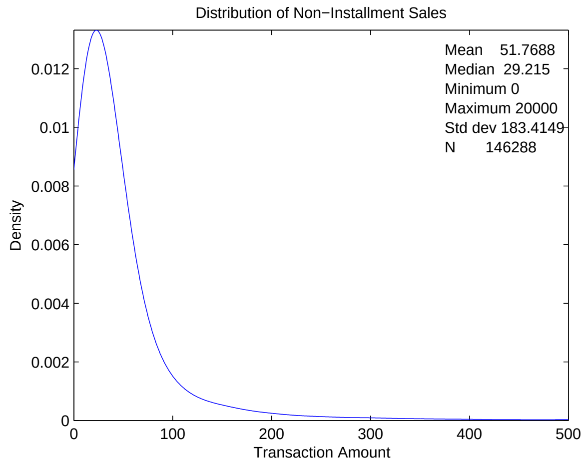

Figure 7 Distribution of positive interest installment purchases

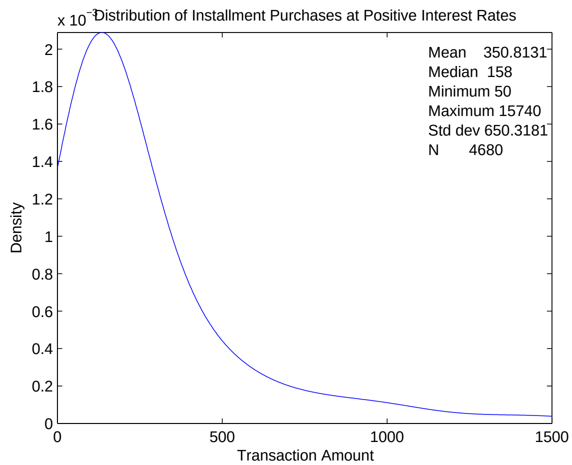

Figure 8 Distribution of zero interest installment purchases

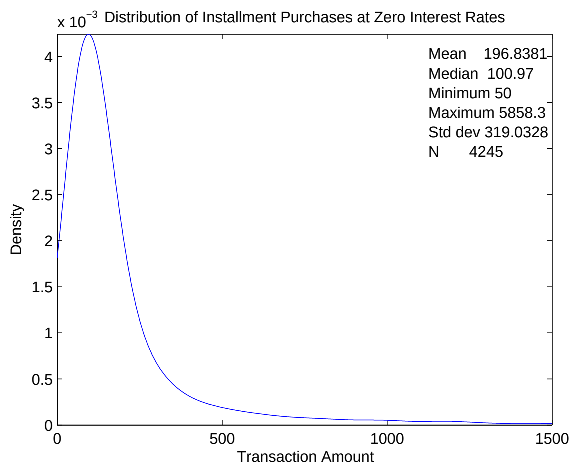

Figure 9 Cumulative Distributions of Credit Card Transaction Amounts

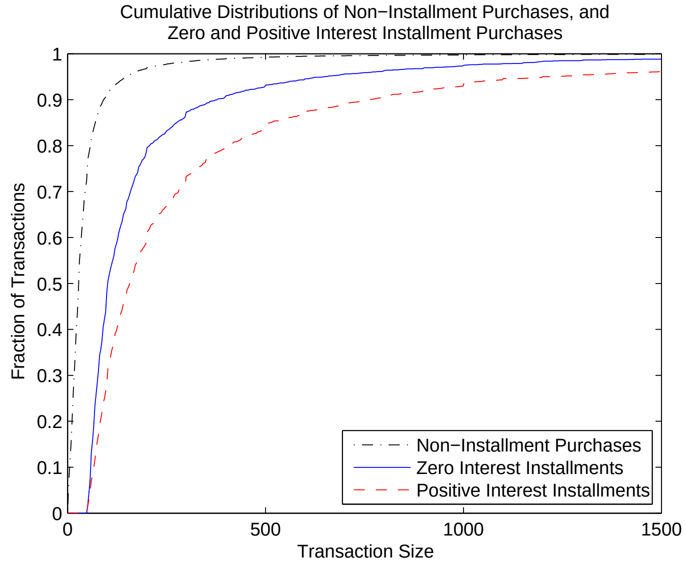

signficantly *larger* than those done at a positive interest rate.

Figure 9 plots the cumulative distribution of non-installment purchases, as well as zero and positive-interest installments. We see a striking pattern: the distribution of positive-interest installments *stochastically dominates* the distribution of zero-interest installments, and this in turn stochastically dominates the distribution of non-installment purchases. Again the latter is to be expected: we would expect consumers to put mainly their larger expenditures on installment and the remaining smaller charges are regular, non-installment credit card charges. However the surprising result is that installments done at a positive rate of interest are substantially larger than installments done at a zero interest rate, at *every quantile* of the respective distributions. For example, the median installment at positive interest rates is nearly 60% larger than the median installment done at a zero interest rate.

In summary, the vast majority of transactions in our sales dataset, 87%, are regular (non-installment) credit card purchase transactions. These tend to be smaller in size with an average size of $50. The remaining transactions consist of cash advances (7% of the transactions) and installments (6% of the transactions). The installments we observe are roughly equally divided between zero interest and positive interest transactions. Specifically, for the subset of installment transactions that we are able to match to the billing table (which enables us to determine the interest rates actually paid, which are not contained in the sales table), approximately 47% of the installments are at zero interest and the remaining ones are done at a positive rate of interest.

Figure 10 Distribution of Rates of Return on Installments, Net of Merchant Fee

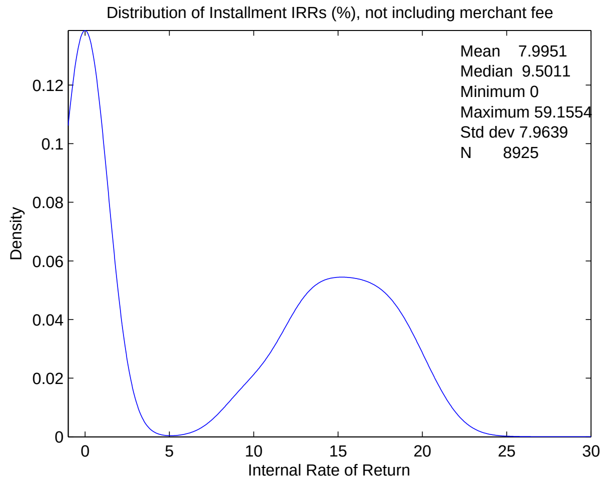

Figures 10 and 11 below show the distribution of internal rates of return that the credit card company earns on these installment sales, before and after accounting for the merchant fee. Recall that the internal rate of return is the (continuous time) rate of interest that sets the net present value of the cash flow stream associated with an installment purchase to zero. The credit card company experiences a cash outflow (to the merchant for the amount of the purchase) on the date the customer makes the purchase which we normalize as "day 0". At the same time the firm received a cash inflow equal to the merchant fee received, which is actually an amount discounted from the amount paid to the merchant (if the merchant is not in-network, then the discounted payment is made to the credit card company that handles the transaction). Then at the next n statement dates the credit card company receives cash inflows equal to the repayment of "principal" plus interest on the installment loan.

Figure 10 shows the distribution of internal rates of returns when the merchant fee is not accounted for. This distribution is effectively the distribution of interest rates charged to the company's customers. We see the pronounced bi-modal distribution reflecting the fact that roughly 50% of installment purchases are done at a zero percent interest rate and the other half is done at a positive interest rate. As noted above, the mean interest rate for positive interest rate installments is 15.25%.

However figure 11 shows that when we add the merchant fee, which provides the distribution of gross returns that the credit card company earns on its installment loans, we see the disrtribution of returns is shifted significantly to the right. Even with the "free installments" included, the company is earning an average rate of return of 23% on its installment loans, and for the positive interest installment loans the

Figure 11 Distribution of Rates of Return on Installments, Including Merchant Fee

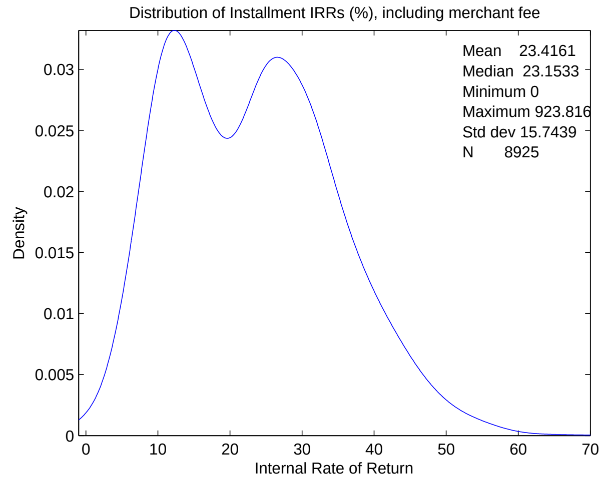

average internal return inclusive of the merchant fee is 31.4%. Of course, these calculations do not include *defaults.* However fortunately for the credit card company we studied, there were only 23 individuals of the 938 in our sample who defaulted and whose credit card accounts were sent to collection. We cannot determine the amount of the unpaid balances that the company was ultimately able to recover from these 23 individuals, however even if all 23 were declared complete losses, including the losses into the distributions in figures 10 and 11 would not significantly diminish the estimated rates of returns that the company earns on its installment loans. Overall, we conclude that at least for this company, the installment loan business is a very good one: it pays very high rates of return with relatively low risk of default.

The high rates of return from installments point to the profitability of the company's non-installment credit card purchases as well. Ae we noted above, the average duration between a purchase and repayment of a non-installment purchase transaction is about 50 days. The average merchant fee that the company earns on its purchase is 2%. Therefore the internal rate of return on the typical sales transaction is the solution to

$$0 = -0.98A + exp\{-r(50/365)\}A, \tag{1}$$

or $r = -365 * \log(.98)/50 = 0.1474$. *Thus, the firm earns an average gross return of 15% even on its regular credit card transactions even when it is giving all of its customers on average a 50 day interest-free loan.*! We conclude that the credit card business appears to be a highly profitable one, and perhaps this should not be a big surprise. However our calculations point out the importance of extending the network of merchants and raising the merchant fee that the company can charge them. If the company were able to raise its average merchant fee to 4%, then the rate of return it earns on ordinary purchases more than doubles, to 29.8% (assuming the same average delay between purchase and repayment on non-installment purchases).

Of course the transaction by transaction analysis ignores the other fees that the company earns, including 1) annual fees, 2) late fees and penalty payments, and 3) interest on revolving credit. Figures 12, 13 and 14 below calculate the gross profits and rates of return that the company earns on a *per customer* basis. The profits are gross profits since we do not know the cost of rewards program benefits provided to customers and are not deducting advertising/marketing and other fixed "front office costs". The internal rates of return are calculated based on treating each full customer record as a stream of cash inflows and outflows, with outflows corresponding to purchases made by the customer and inflows being payments received from merchant and from the customer when the customer pays balances due and other fees on the statement dates.

We see from figure 12 that though the average gross daily profits that the company earns per customer is about 60 cents per day, there is huge variability, and the company can incur large losses (amounting to as much as $14 per day) for the customers who default, but balanced by profits as high as $19 per day for some of the most profitable customers (note that we calculated the daily profits only for a subsample of customers for which we could observe at least several hundred transactions over an account duration of at least 3 months, so we do not believe the maximum and minimum gross profit values are likely to be results of sampling noise from customers who made only a few transactions and were observed only over short periods of time).

Figures 13 and 14 express the profitability of customers in terms of their rate of return. Here again we see the large degree of variability in return, reflecting that the credit card business does represent an "investment" that has both a high risk and return. However the most important conclusion to take away from these figures is the huge effect merchant fees have on the overall profitability of this firm. Without merchant fees, the company is already earning a respectable 29% rate of return on its customers, however when we include merchant fee the mean return increases to 89%! Thus, merchant fees account for more than a third of total revenues in our sample, and they account for an even greater share of total profits and the overall high rates of return earned by the firm. The reason for this, of course, is that the merchant fee is a cash flow the company receives right away at the time of each transaction, and there is virtually no risk associated with this stream of revenues. This is why even modest merchant fees equal to 2 to 4% of the

Figure 12 Distribution of Daily Profits per Customer

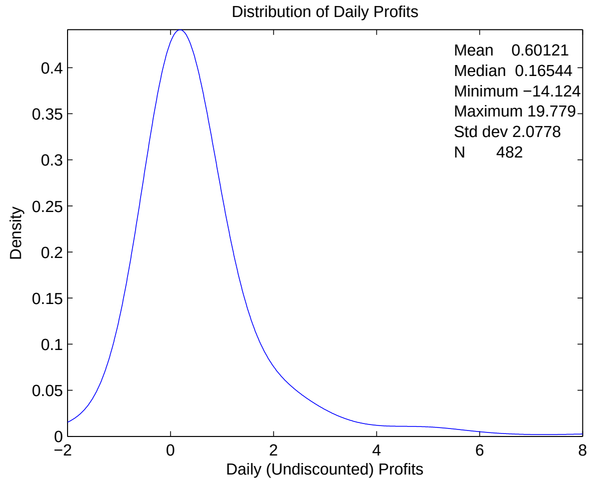

Figure 13 Distribution of Customer-specific Internal Rates of Return, Excluding Merchant Fee

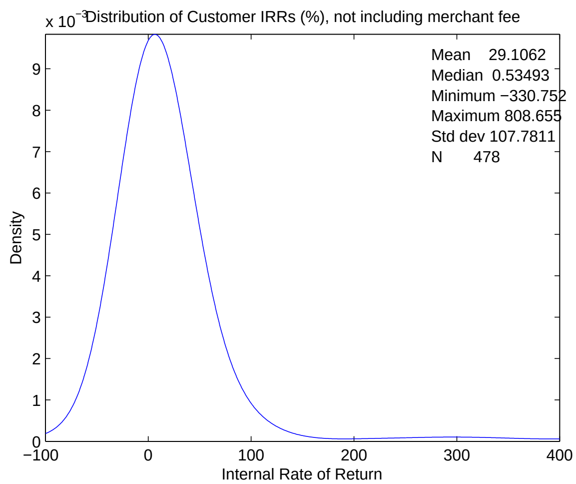

Figure 14 Distribution of Customer-specific Internal Rates of Return, Including Merchant Fee

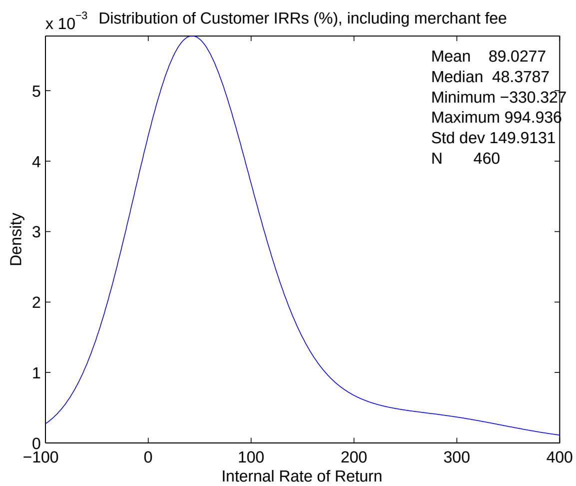

transaction price contribute so importantly to the bottom line of this company.

Nevertheless we think calculations of the gross returns and profits on individual customers is an interesting and meaningful calculation to make. It can help to identify characteristics of the most profitable customers, and it also provides further illustration of the importance of merchant fees in the overall profitability of the company. Although we defer an analysis of merchant fees to a separate paper, we already showed that customer profitability is highly linked to the magnitude of the merchant fee that the company earns on the customer's transactions. We will show in subsequent work that there is a wide variation in the merchant fees (as a percent of the amount charged) with merchant fees ranging from as low as zero to has high as 6% of some purchases. Thus, it is not only the *volume* of spending but also the *distribution of spending* (i.e. whether the customer tends to use their care at merchants that pay high fees or not) that can determine the overall profitability of a customer. Clearly, the customer's interest in borrowing and willingness to pay high interest charges also affects their profitability as well.

Overall, our preliminary analysis of the credit card data leads to a number of key conclusions. First, we already see the "free installment puzzle" emerging by comparing the distributions of expenditures for zero interest installments to the corresponding distribution of positive interest installments. We showed that the latter distribution stochastically dominates the former distribution, so that at every quantile in the distribution, these customers are spending more on installments that come with a large interest rate than for installments that are offered at an interest rate of zero. Secondly, we showed that the company is highly profitable and that merchant fees contribute in an important way to the overall profitability of the firm.

Specifically, when we computed the (undiscounted) revenues of the firm for the 938 customers we analyzed, we found that merchant fees amounted to 36% of the total revenues received from these customers. We have argued that the company sees merchant fees as a major component of its profits and due to the structure of payments in this country, the company places great importance on rapid growth, both in absolute and in terms of its market share, as the key to its future success. For the reasons we discussed above, there is a strong element of increasing returns to scale and network externalities that lead to the cards offered by the dominant firms being accepted by more merchants and this in turn enables these firms to charge higher merchant fees.

We have argued that this provides a strong incentive for the firms to try to attract new customers and to stimulate the credit card spending of its existing customers by offering free installment opportunities to its customers. However this only heightens the basic puzzle: if consumers appear to be spending *less* per transaction on the free installment opportunities they are offered in comparison to their average transaction sizes when they pay the full interest rate, what evidence is there that free installments are really stimulating spending or enabling the company to attract a significant number of new customers?

Our analysis is limited by an important *sample selection bias:* our data do not allow us to observe all transactions a customer makes using the various possible means of payment at their disposal, including paying in cash, and paying using an alternative credit card. As a result, our data are not informative as to whether the possibility of free installments induces customers to make a *greater number of transactions* even if the average size of a free installment is less than installments done at a positive interest rate. We do not observe whether customers are aware of the free installment opportunity *prior* to undertaking any purchase using their credit card, and thus, whether this option caused them to make an extra purchase, or switch from paying in cash or using another credit card to making the purchase using the company's credit card under the free installment option.

However what we can learn from our data is the likelihood, conditional on being presented with a free installment opportunity, a customer is willing to take this opportunity when it is offered to them. We think that all customers are aware that they can make purchases under installment at a positive interest rate, since this information (and the interest rate schedule they are facing) is part of their monthly statements. Further, customers are usually informed about the option to do an installment on an interest-free basis at the checkout counter, though there are some interest-free installment options that are offered by any merchant during a specific interval of time, and the creditcard company usually heavily advertises these special promotional periods, including in flyers included in customers' monthly statements. Thus, we believe it is highly plausible that the customers we are studying are fully aware of the various options that they have for making a purchase, including to purchase under a free installment option when it is available.

Thus, our data allows us to ask and hopefully answer questions such as, "conditional on deciding to make the purchase, how does the magnitude of the interest rate affect the likelihood the customer will pay for the item on installment?" More specifically, we can use our data to try to answer the question, "conditional on deciding to make a given purchase and being offered a free installment purchase option, what is the likelihood that a customer will choose the free installment option?" While our data do not allow us to identify the complete *demand curve for credit* if we are able to provide answers to the questions raised above, we can at least gain new insights into the *conditional demand for credit,* i.e. how interest rates affect the probability that the consumer will borrow (via deciding to pay the amount on installment) conditional on their having made a decision to buy a given item (or spend a given amount on a bundle of goods).

4 Exploratory, Multivariate Analysis of Installment Transactions

Before we go into a more focused empirical analysis directed at the specific issue of attempting to estimate the "demand for credit" we find it useful to present some additional scatterplots that reveal some important facts and features and correlates of installment purchase decisions. In particular, we are interested in understanding what types of purchases are made via installment credit, and which types of individuals are the most likely users of installment credit.

Figures 15 and 16 show the distribution of the number of credit card *transactions* and the *share of all credit card spending* done as installment purchases. We see that while installments are less than 9% of all credit card transactions, they account for more than 25% of all credit card spending.

Of course, this is due to the fact that the average credit card purchase is $74 while the average installment purchase is $364, with the full distributions of the average purchase and installment transaction sizes over the consumers in our sample plotted in figures 17 and 18. Thus, consumers generally pay for much larger items (or more expensive baskets) on installment, but choose to pay smaller amounts in full at the next statement date. We are also struck by the much greater skewness of the distribution of installment purchases relative to that of credit card purchases as a whole.

Our analysis reveals a substantial degree of heterogeneity across credit card customers in their propensity to make use of installments to pay for their credit card purchases. Overall our analysis suggests that the best single measure of the propensity to use installments is not the mean fraction of transactions done via installment, but rather the mean share of credit card purchases paid for by installment. Hereafter we will refer to the latter measure as the *installment share.* Now we will turn to a series of scatterplots that relate the installment share to other covariates we observe in our credit card data set.

Figures 19, 20 and 21 present scatterplots (with the central tendency of the data indicated by a local linear regression fit to the data) of how the installment share relates to various measures of creditworthiness. Figure 19 plots the installment share against customer credit scores, using the company's internal (proprietary) credit scoring system where a score of 1 represents the best possible creditworthiness and 12 is the worst. Customers who have credit scores in this range are still allowed to borrow on installment and face no credit limits. However consumers who are in the process of collection will have their credit card

Figure 15 Distribution of the Fraction of Credit Card Transactions done as Installments

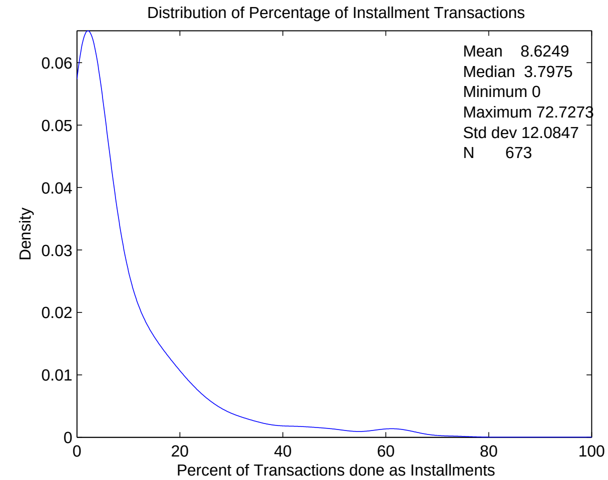

Figure 16 Distribution of the Share of all Credit Card Spending done as Installments

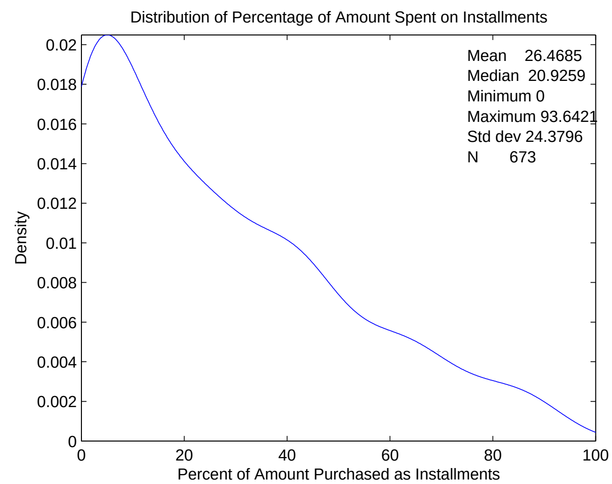

Figure 17 Distribution of the Average Amount of a Credit Card Purchase across Customers

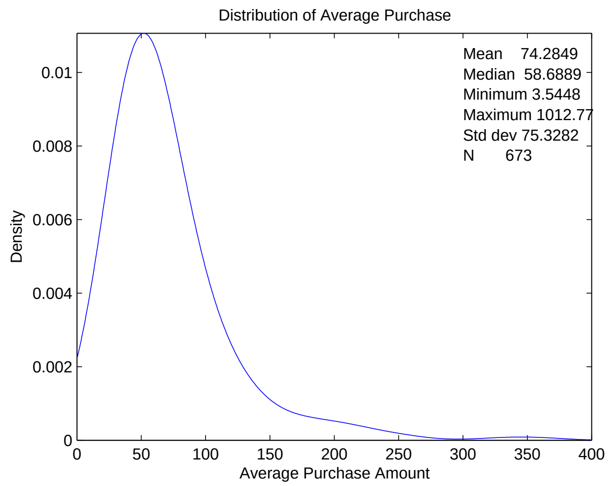

Figure 18 Distribution of the Average Amount of an Installment Purchase across Customers

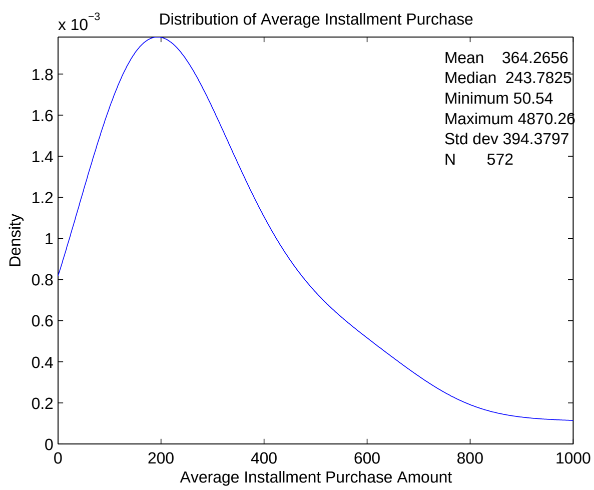

borrowing and spending privileges suspended and they show up in our data set as having a credit score of 0. We see generally negative correlation between the credit score and the installment share (remember that higher credit scores indicate worse credit, so the relationship in figure 19 is actually positively sloped).

We see figure 19 as a potential first indication of possible credit constraints, or at least *high demand for credit* among the customers that are heavy installment spenders. Perhaps their poor credit score indicates that they are also regarded as poor credit risks to other lenders, and as a result of this, they are forced to make heavier use of the installment credit facility of this credit card company at relatively high rates. On the other hard, the customers with the best credit scores also generally are the least heavy users of installment, which could be an indication that they are not liquidity constrained, or have other lower cost sources of access to credit elsewhere.

Figures 20 and 21 illustrate the incidence of late payments. Figure 20 shows that the average number of late payments per customer is positively correlated with the installment share, and figure 21 shows that the number of *seriously late* payments (i.e. payments that are 90 or more days past due, or at about the threshold where the company suspends credit card charging privileges) is also positively correlated with the installment share. These figures confirm the conclusion we obtained in figure 18 *that customers who are heavy users of installment spending are also worse credit risks.*

Figures 22 and 23 relate the installment share to three separate indicators of the type of installment spending that customers do. Figure 22 presents a scatterplot of the ratio of the size of a typical installment purchase to the typical credit card purchase. As we noted previously, credit card customers generally pay for only relatively large purchases on installment, and pay for the smaller transactions in full at the next statement date. We see that as a function of the installment share, the low intensity installment users tend to buy items on installment that are between 4 and 6 times as large at their typical credit card purchase. However for the heaviest users of installment spending this ratio falls to less than 3, which potentially indicates a more “desperate” individuals who are more likely to pay for smaller “everyday” items by installment.

Figure 23 shows a scatterplot of the ratio of the installment balance to the average statement balance as a function of the installment share. Of course, that this ratio is positively correlated with the installment share is almost definitional, but the figure does show that the heaviest installment users carry installment balances that are on average 10 times larger than their typical monthly credit card balances (statement amounts).

Figure 19 Customer-Specific Credit Scores by Installment Share

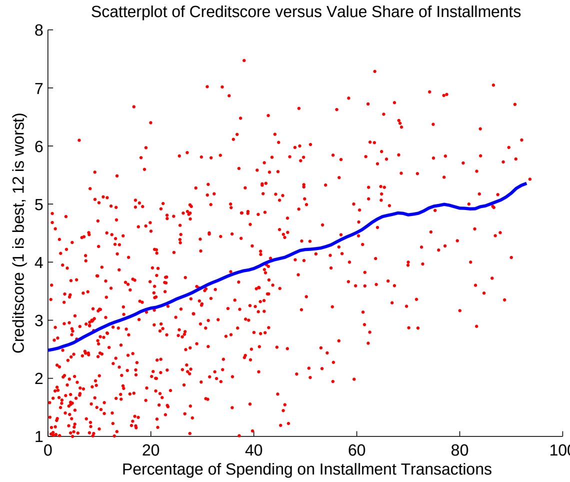

Figure 20 Number of Late Payments by Installment Share

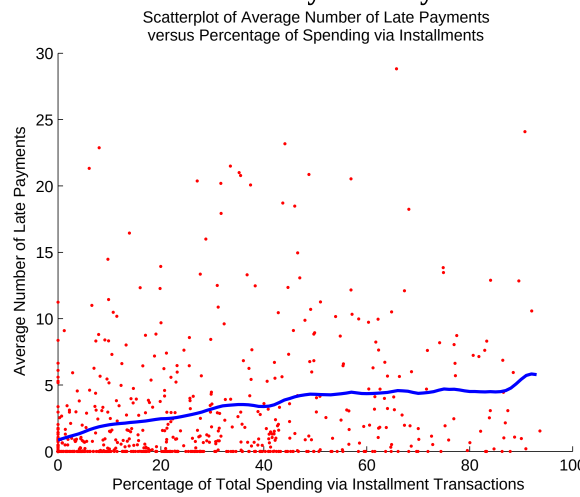

Figure 21 Number of Seriously Late Payments (over 90 days) by Installment Share

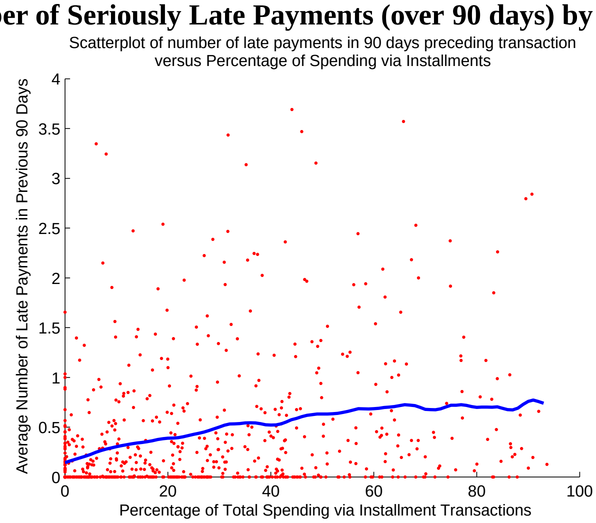

Figure 22 Ratio of Installment Size to Typical Purchase Size by Installment Share

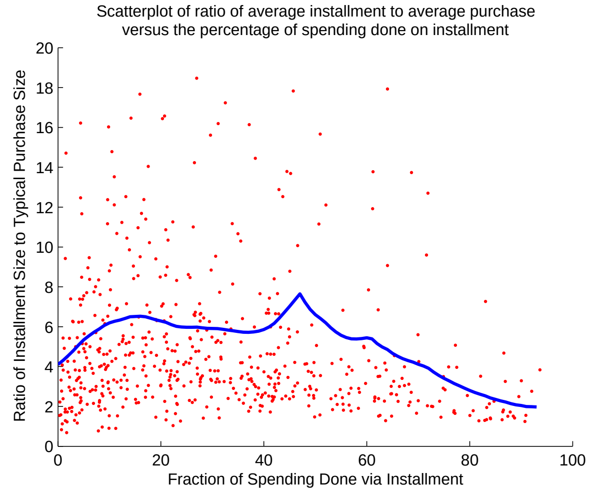

Figure 23 Ratio of Installment Balance to Average Statement Balance by Installment Share

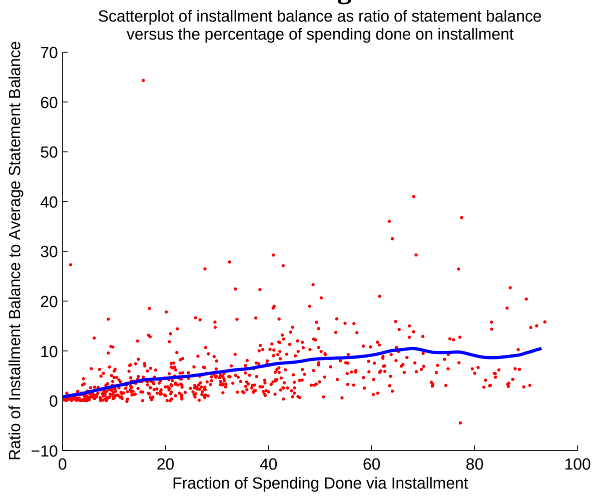

Figure 24 Fraction of Installment Transactions done as Free Installments by Installment Share

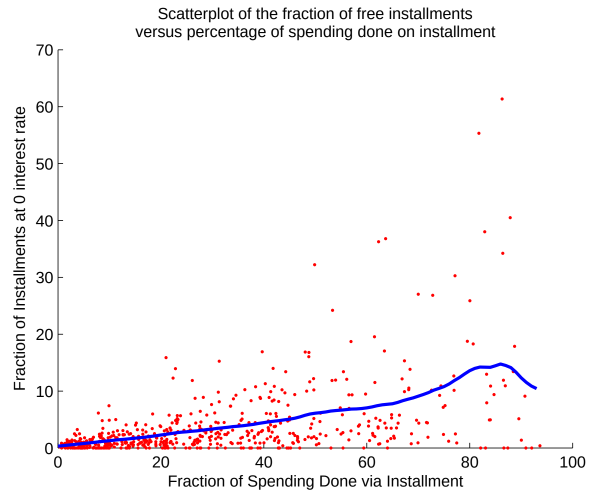

Figure 25 Share of Installment Spending Done as Free Installments by Installment Share

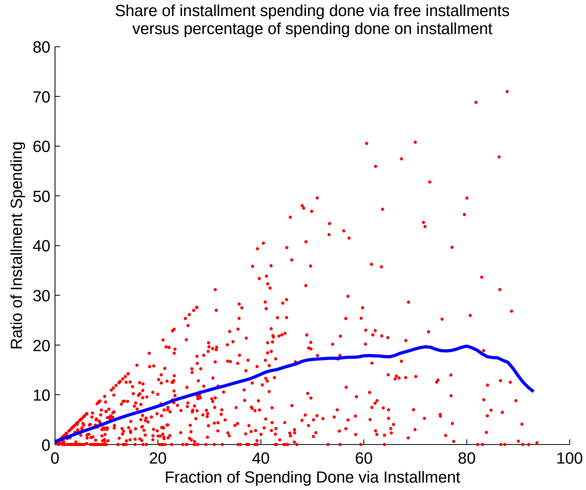

Figures 24 and 25 relate the usage of free installments to the installment share. In figure 24 we see that the fraction of installment transactions done as free installments is positively correlated with the installment share. The previous figures in this section lead to an impression that the heavy installment spenders are relatively desperate for credit, and thus, it would seem logical that they are the ones who would be most likely to take the greatest advantage of free installment opportunities when they encounter them. The upward sloping relationship in figure 24 is consistent with this interpretation, and shows that the heaviest installment users are doing as much as 20% of their installment purchase transactions as free installments (i.e. at 0% interest rate).

Figure 25 shows a similar relationship but instead of plotting the fraction of installment transactions that are done as free installments it shows the share of installment spending that is done via free installments. Both of these graphs show a similar pattern, namely that the customers with the highest installment shares are doing about 20% of all of their installment transactions and 20% of all installment spending via free installment offers. If we assume that these most high intensity installment "addicts" will take almost any free installment offer that is presented to them (provided the transaction is of a sufficiently large size to make an installment purchase worthwhile), these graphs suggest that approximately 25% of all relevant installment purchase opportunities should be at a 0% interest rate (i.e. a free installment). The reason is that if the high intensity installment addicts are doing approximately 80% of their spending on installment, and free installments are being offered about 25% of the time, then this would imply the observed 20% fraction of free installments for these individuals, since 20% = 80% $\times$ 25%.

Finally, we conclude this section with figures 26, 27 and 28 that give us some insight into the profitability of the "free installment marketing strategy" used by this firm. Recall from section 2 that we suggested that the company's use of free installment offers seems motivated by a desire to increase its customers' use of its credit cards in an attempt to increase its credit card market share, since doing this increases its leverage in setting merchant fees, which we showed in section 3 are a major component of the high profitability of this company. However we have also shown in this section that the customers that are most likely to act on the free installment offers are those with worse credit scores and higher incidence of late payments. As such, the use of free installments as a promotional device may have the perverse effect of offering free credit to the company's least creditworthy customers, and this group may be the most likely to default. This creates the possibility that free installments might be a relatively ineffective and/or highly costly means of increasing credit card usage.

Figure 26 plots the average internal rate of return on all installment transactions (including free installments) against the installment share. We see that this curve is upward sloping, which indicates that even though the "installment addicts" are the ones most likely to be taking up the free installment opportunities, the interest rates that they pay on their positive interest installment transactions are rising sufficiently fast with the installment share that it counteracts the "free installment effect" so that overall average installment interest rates paid by its customers increase monotonically as a function of the installment share. Of course the reason for this is likely to be related to the fact that the customers with high installment shares have significantly worse credit scores, and as we will show in section 5, the interest rates that customers pay is a monotonically increasing function of their credit score (i.e. customers with higher scores, which indicate worse credit risks, pay higher interest rates).

Figure 27 plots the overall internal rate of return for each customer as a whole, for *all* transctions, installment and non-installment and revolving, and including all other fees such as late charges, annual fees and so forth, and also inclusive of merchant fees. We already presented the distribution of these rates of returns in figure 14 in the previous section. In figure 27 we have trimmed the sample of customers to those whose internal rates of return are between -20% and +50%, and since the distribution of returns is so skewed to the right, this trimming has the effect of substantially reducing mean returns from the 89% presented in table 14 to approximately 20-30% in figure 27. Here, we see that in part due to the large variability in customer-specific rates of returns, there is no obvious upward sloping relationship between the installment share and the customer-specific return. If anything, figure 27 suggests that there is a

Figure 26 Average Internal Rates of Return on Installments by Installment Share

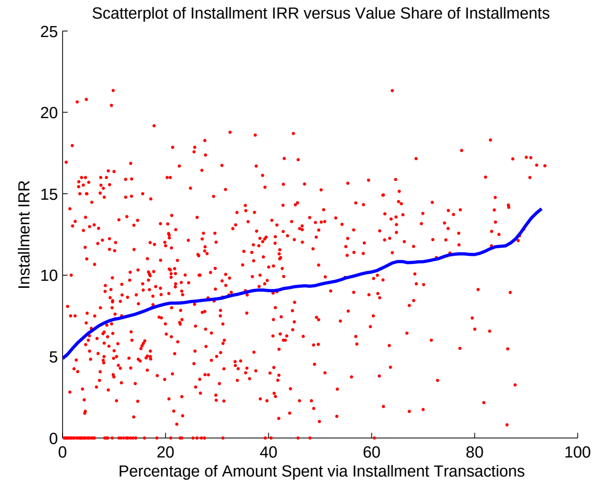

Figure 27 Customer-Specific Overall Internal Rates of Return by Installment Share

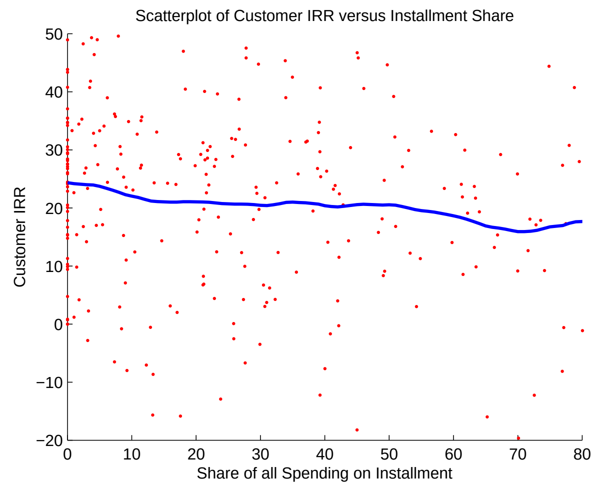

Figure 28 Customer-level Daily Profits by Installment Share

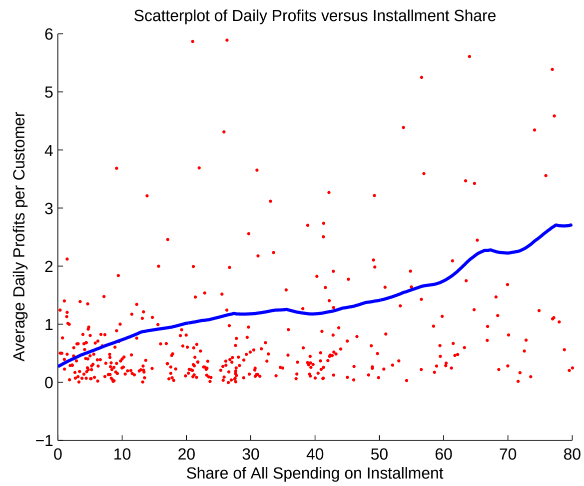

negative correlation between the customer-specific rates of return and the installment share.

However we obtain a different perspective on this in figure 28, which plots the average daily profits for each consumer against the installment share. This figure indicates a pronounced upward sloping relationship between the installment share and the profitability of customers. If we believe this is the relevant figure to focus on, then the company's free installment marketing policy seems rational and well targetted: it appears to be succeeding in having the biggest impact on the most profitable customers, but these customers also happen to have worse credit scores and present higher credit risks.

However given the relatively small number of observations and the relatively large number of outliers, we think it is hazardous to come to any definite conclusion one way or the other about the wisdom of free installments at this point. As we noted in the previous section, we cannot address with our data a crucial missing piece of information that would be needed to provide a fuller answer to this question: to what extent does the knowledge of free installments cause customers to increase their spending? Recall that we are doing our analysis *conditional* on the decision to purchase a given item. We would need different data to determine whether the existence and knowledge of free installment opportunities causes the company's customers to go to stores more often, purchase more at a given store than they otherwise would, or increase their likelihood of using the company's credit instead of paying for the item using a competing credit card or cash.

5 Inferring the Demand for Installment Credit

The data we have would appear ideal for empirically modeling *the conditional demand for credit* — at least as it pertains to relatively smaller scale short term borrowing decisions. As we noted above, we define the conditional demand for credit as the demand to finance a given credit card purchase through borrowing rather than to pay the amount purchased in full at the next purchase date. It is conditional on having made a decision to make a given purchase of a given size in the first place. As we noted above, we do not have the appropriate data that would enable us to model how access to borrowing and how the interest rate schedule that a customer can borrow at also affects the frequency and amounts of purchases. We would need additional sources of data, then, to attempt to estimate the fuller *unconditional demand for credit.*

To make this a bit more precise, we introduce a bit of notation. Let c denote the decision by the consumer to pay using the company's credit card (as opposed to paying by cash, or using some other credit card). Let r be the interest rate charged to a customer with observed characteristics x for purchasing via installment credit. As we show in more detail below, we should interpret r as an entire *interest rate schedule* since the customer can ordinarily choose the term of the installment loan and thus faces a consumer-specific "term structure" of interest rates. Consider the demand for credit via the company's credit card c over a specific interval of time, say one month. The (unconditional) expected demand for credit by a single customer with characterics x, $ED(r,x,c)$ (where x includes variables such as the customer's credit score, spending history, and might also include information on interest rates offered by competing credit cards or interest rates for other sources of credit) can be written as follows

$$ED(r,x,c) = \left[\int_0^\infty a[1-P(1|a,r,x,c)]\,f(a|x,r,c)da\right]\pi(c|r,x)EN(x,r). \tag{2}$$

where $P(1|a,r,x,c)$ is the probability that a customer will choose to pay for a purchase amount a in full at the next statement date given the interest schedule r, the consumer charateristics x and the decision to use the company's credit card c to carry out the transaction. We let $\pi(c|r,x)$ denote the customer's decision to use the credit card company c's credit card to pay for the transaction, and $f(a|x,r,c)$ denotes the density of the amount purchased using the company's credit card during any given shopping trip. Finally $EN(x,r)$ denotes the expected number of shopping trips that the customer makes during the specified interval of time. The overall expected demand for credit from the customers of credit card company c is then just the sum over the customer-specific expected demand curves $ED(r,x,c)$.

The data we have are not sufficient to estimate the objects $\pi(c|r,x)$ or $EN(x,r)$. Separate survey data would have to be collected that would enable us to study the purchase habits of a sample of the company's customers, and how something like free installment offers during a given period of time might affect the number of shopping trips they make (thus enabling us to estimate $EN(x,r)$), or the likeilihood that they will use the company's credit card c to pay for the purchase (thus enabling us to estimate $\pi(c|a,r,x)$).

However since we do observe all of the purchase amounts that a given consumer makes during any given shopping trip where the customer uses the company's credit card, we can potentially estimate $f(a|x,r,c)$. Further, since we also observe customers' choices of whether to purchase on installment or whether to pay the amount a in full at the next statement date conditional on having decided to use the company's credit card, we can potentially estimate the *installment choice probability* $P(d|a,r,x,c)$, where the option $d=1$ indicates a choice to pay the purchase amount a in full at the next statement date. If so, then by segregating customers' purchases into those that are paid in full at the next statement date and those that are paid on installment, we can estimate two conditional densities, $f_0(a|x,r,c)$ (i.e. the distribution of purchase amounts that are paid in full at the next statement date) and $f_1(a|x,r,c)$ (the distribution of purchase amounts that are paid for by installment). We have already presented the unconditional analogs of f_0 and f_1 in figures 17 and 18 of section 4, where we showed in particular that the average size of an installment purchase was nearly 5 times larger than the average size of a non-installment transaction. Since f_0 and f_1 are conditional distributions, we can write them according to the usual formulas of probability theory

$$\begin{aligned} f_0(a|x,r,c) &= \frac{P(1|a,r,x,c)f(a|x,r,c)}{\int_0^\infty P(1|a,r,x,c)f(a|x,r,c)da} \\ f_1(a|x,r,c) &= \frac{[1-P(1|a,r,x,c)]f(a|x,r,c)}{\int_0^\infty [1-P(1|a,r,x,c)]f(a|x,r,c)da}. \end{aligned} \tag{3}$$

Thus, we can at least use our data to estimate the *conditional expected demand for credit* $ED_1(r,x,c)$ which we define as

$$ED_1(r,x,c) = \int_0^\infty a f_1(a|x,r,c)da. \tag{4}$$

Just as we expect the unconditional demand curve to be a downward sloping function of r, we also expect the conditional demand for credit to be downward sloping in r because we expect that customers to borrow larger amounts on installment when the interest rate is lower. Even if the distribution of purchase sizes was unaffected by r (i.e. if $f(a|x,r,c)$ was not a function of r), a downward sloping demand would still follow if the probability that a customer chooses to pay the purchase amount a in full at the next statement date is an increasing function of r (in which case the customer's credit demand is nothing beyond that inherent in the typical "float" i.e. the lag between buying an item with a credit card and paying for it at the next statement date).

It follows that if we restrict attention to the subset of transactions that a customer purchases on installment credit, we have the regresion equation

$$\tilde{a}_i = ED_1(r,x,c) + \tilde{\varepsilon}_i \tag{5}$$

where $\tilde{a}_i$ is the amount borrowed in the i^{th} installment transaction made by the customer, and $\tilde{\varepsilon}_i$ is a residual satisfying $E\{\tilde{\varepsilon}_i|r,x,c\} = 0$. We refer to the regression equation (5) as the conditional demand curve for credit, and it seems like a natural place to start is to estimate this regression by ordinary least squares. However rather than attempt to specify parametric functional forms for the underlying components of the regression function $ED_1(r,x,c)$, i.e. the probability $P(1|a,r,x,c)$ and the density $f(a|x,r,c)$ which would result in a specification that is nonlinear in the underlying parameters, it is also natural to start by estimating a flexible linear-in-parameters approximation to the regression function $ED_1(r,x,c)$.

However, perhaps not surprisingly, we find that when we do these ordinary least squares regressions for every specification we tried where the dependent variable is the amount of an installment purchase and for different combinations of right hand side (x,r) variables, we always found that the regression predicted a strong, and statistically significant *positive relationship* between the expected amount of installment borrowing and the interest rate r. to have a *positive and statistically significant coefficient.* That is, the regressions are suggesting that the *conditional (expected) demand for credit is upward sloping!*

Of course, the ordinary least squares regression results are likely to be spurious due to the *endogeneity of the interest rate.* That is, we can imagine that there are *unobserved characteristics* of consumers that affect both their willingness/desire to make purchases on credit and the interest rate they are charged. In particular, we would imagine that customers who are *liquidity constrained* and who might exhibit *bad characteristics* that can lead them to simultaneously wish to borrow more but at the same time constitute a *higher credit risk* will have worse credit score and therefore face a higher rate of interest, but will still have a higher propensity to borrow due to their liquidity constraints and a dearth of alternative, better borrowing options. Indeed, we show in section 3 that there is a strong correlation between the fraction of spending on installment credit and the credit score: indiividuals with worse credit scores tend to do a higher fraction of their credit card purchases on installment. Given the monotonic relationship between credit scores and installment interest rates, it is not hard to see why the regression estimate of the installment interest rate is positive and statistically significant.

We attempted to deal with the endogeneity problem using the standard arsenal of "reduced form" econometric techniques, including *instrumental variables.* In particular, we have access to daily interest rates that measure the "cost of credit" to the bank for the loans it makes to its customers, including 1) *the certificate of deposit CD rate* and 2) *the call rate.* The latter is an interbank lending rate for "one day loans." Both the CD rate and the call rate change on a daily basis. We use these rates as instrumental variables on the theory that in a competitive banking market, no single bank can affect the CD or call rates, and thus changes in these rates can be regarded as exogenous changes in the cost of credit that the credit card companies ultimately "pass on" to their credit card customers. However the instrumental variables (two stage least squares) estimate of the coefficient of the interest rates the company charges its customers becomes *statistically insignificant* as you can see in table 1 below. The coefficient estimates of the interest rate r are highly sensitive to whether we include all installment transactions (including those with $r = 0$ or just those with $r > 0$. We obtain a highly negative but statistically insignificant point estimate in the former case, and positive and statistically insignficant estimate in the latter.

We define the average treatment effect (ATE) as our "parameter of interest" even though our actual interest is to estimate the conditional demand curve for credit. Given the poor results from instrumental variables estimation, we are now willing to settle for a much less ambitious goal: can we even show that people will borrow more when offered 0% interest compared to when they must pay high positive rates of interest? The ATE is simply an estimate of the difference between mean borrowing for the treatment group who were offered zero interest

We do not really believe the inferences from our instrumental variables regressions, or the suggestion that we have a unique finding that the demand for credit is some sort of *Giffen good.* After all, if the firm believed that charging higher interest rates causes its customers to spend *more,* why would it offer free installment opportunities? Instead we believe that the reduced-form results are spurious, and in particular both the CD and call rate are *weak instruments.* Indeed, not only are they weakly correlated wtih consumer interest rates, we find that the CD and call rates are *negatively correlated* with the interest rates the firm charges to its customers. We view this as evidence that the credit card market is not "competitive" and the are substantial "markups" in the interest rates charged to customers over the cost of credit to the banks, and this markup is driven more by customer specific risk factors and by competitive trends within the credit card market itself than by the the much smaller day to day fluctuations in the CD and call rate. The latter have hovered in a fairly narrow band between 3 or 4 percent over the period of our analysis whereas installment interest rates vary much more widely across customers and over time as their credit scores change, ranging from as low as 5% to 25% or higher.

The next approach we considered in order to try to infer the "causal effect" of interest rates on the demand for credit was *matching estimators.* The idea behind these estimators is to compare the average amount purchased by individuals who were offered free installments (the "treatment group") with a corresponding and "similar" set of individuals who took out installment loans when purchasing from similar merchants at similar periods of time but at a positive interest rate (the "control group"). Since there are many individuals in our sample for which we observe a large number of installment transactions (these are the heavy installment "addicts" that we discussed in the previous section who have installment shares in excess of 50%), we can even use a number of individuals as "self-controls" — that is we can compare the

Table 1: Instrumental Variables-Fixed Effects Regressions of Conditional Demand for Credit
Dependent variable: $log(a)$ where a =amount borrowed. Amounts in parentheses are P-values for tests of the hypothesis that the coefficient/statistic is zero.

Item	Specification 1	Specification 2	Specification 3	Specification 4
Instruments	CD rate	CD rate	CD rate	CD rate
	credit score	credit score		
Fixed effects	yes	yes	yes	yes
Free Installments	yes	no	yes	no
Variable	Estimate	Estimate	Estimate	Estimate
r	0.965	-72.903	0.739	-102.20
	(0.249)	(0.591)	(0.382)	(0.628)
credit score	0.001	-0.002		
	(0.442)	(0.835)		
$d = 2$	0.314	2.733	0.317	3.695
	(0.000)	(0.531)	(0.000)	(0.588)
$d = 3$	0.896	4.9644	0.912	6.543
	(0.000)	(0.690)	(0.000)	(0.561)
$d = 4$	1.028	5.434	1.042	7.099
	(0.000)	(0.474)	(0.000)	(0.549)
$d = 5$	1.06	6.623	1.061	8.668
	(0.000)	(0.472)	(0.000)	(0.547)
$d = 6$	1.828	7.172	1.840	9.243
	(0.000)	(0.450)	(0.000)	(0.533)
constant	11.500	20.559	11.519	24.104
	(0.000)	(0.216)	(0.000)	(0.349)
σ_u	0.651	1.276	0.652	1.708
σ_ε	0.656	1.788	0.657	2.420
ρ	0.495	0.337	0.496	0.331
Sample size	8183	4109	8078	4049
F-test $(u_i = 0)$	$F(613) = 8.03$	$F(474) = 0.97$	$F(598) = 8.08$	$F(464) = 0.53$
	(0.00)	(0.687)	(0.00)	(1.00)
Hausman test	$H(8) = 6.54$	$H(8) = 1.96$	$H(6) = 4.45$	$H(6) = 0.23$
	(0.59)	(0.96)	(0.61)	(0.99)

Table 2: Effect of Free Installments: Results from Matching Estimators

Matching Criteria	Estimated ATE	Standard Error	P-value for H_o : ATE $= 0$
customer, credit score CD rate, merchant code	-\$56.60	\$15.20	0.000
customer, credit score merchant code	-\$69.51	\$16.45	0.000
customer, merchant code	-\$79.33	\$19.93	0.000
customer	-\$76.72	\$18.75	0.000
merchant code	-\$61.07	\$16.00	0.000

average size of free installments with the average size of installments done at positive interest rates for the same individual, where we do additional matching by selecting a set of free installments and positive interest rate installments that were done at approximately the same intervals of time and from approximately the same set of merchants.

Specifically, we focuses on attempting to estimate the "average treatment effect " (ATE) where the "treatment" in question is offering a customer a free installment borrowing opportunity, which we denote as $r = 0$. The ATE is defined as the difference in the expected borrowing between the treatment group $r = 0$ and control group $r > 0$

$$\text{ATE} = E\{a|r = 0\} - E\{a|r > 0\}, \tag{6}$$

where a is the amount borrowed and r is the interest rate. The idea behind the matching estimator is that if we are able to match a sufficiently large number of customers in the treatment and control groups on a sufficiently narrow set of criteria X such that we can plausibly assume that the "assignment" of the "treatment" $r = 0$ is essentially random for the matched individuals/transactions, then we can infer what the installment spending for a treated person would be by taking the mean installment spending for the matched individuals in in the control group (and vice versa) and essentially estimate the ATE as if it were a result of a classical controlled randomized experiment for subsets of matched individuals and transactions and averaging these match-specific treatment effects across all matched groups in the sample. The validity of this approaches depends on a conditional independence assumption known by the (unfortunate) name, "the unconfoundedness assumption" (or also, the "strong ignorability assumption"). The table below presents our estimates of the ATE, which we would expect to be positive if the demand for credit were downward sloping.

We can see from table 2 that regardless of how we do the matching of individuals/transactions the estimated treatment effects are all estimated to be of the *wrong sign and highly statistically significant.* The estimated treatment effects become increasingly negative as we use increasingly relaxed criteria for matching individuals, but overall given the magnitude of the estimated standard errors for the estimated ATE's, there is no strong evidence that the various estimates are statistically significantly different from each other. However we can strongly reject the hypothesis that the ATE is zero. Thus, we are left with the paradox that the matching estimator predicts that free installment opportunities cause customers to *reduce the amount of their borrowing* and therefore, the matching estimators imply an *upward sloping demand for credit.*

6 Exploiting the Quasi-Random Nature of Free Installment Offers

In view of the failure of the various reduced form methods that we tried in the previous section (regression, instrumental variables estimators, and matching estimators) to produce plausible estimates of the slope of the conditional demand for credit (with regression estimates resulting in significantly positive estimates of the coefficient of r, IV estimates being insignificantly different from 0, and matching estimators implying that free installments have a statistically significant *negative* average treatment effect), we had to start to think "outside the box" about other ways to provide more credible and econometrically valid estimates of the conditional demand for credit.

Our next approach was to see if there is some other way to exploit the company's use of free installment offers as a *quasi random experiment.* We already tried to do this in the previous section, where we applied one of the standard approaches in the "treatment effects" literature, namely the use of matching methods. Unfortunately the matching estimators were all strongly statistically and economically significant but of the wrong sign. Although the quasi-random nature of the way the credit card company offers free installment offers to its customers does provide a strong degree of *prima facie* plausibility for the validity of the key conditional independence assumption that justifies the use of matching estimators, the fact that there is a great deal of *self-selection* in which individuals choose to take free installment offers suggests that there could be an important problem of *selection on unobservables* that could invalidate the conditional independence assumption and cause the matching estimators to result in spurious estimates. So we began to explore whether there is an alternative way to exploit the quasi random nature of free installment offers that could be robust to the possibility of selection on unobservables.

Consider first an idealized *randomized controlled experiment* (RCE). Though the company we are studying has not done this to our knowledge, one could imagine that the company could be convinced to undertake such a study to get better estimates its customers' demand for installment credit (the ADL 2011 study (Alan et al. [2011]) discussed in in the introduction is an example where an enlightened credit card company did choose to undertake a large scale RCE to better understand its customers' demand for credit). In a classical RCE the company would randomly assign a subset if its customers to a control group and a treatment group. Individuals in the control group would continue to receive the same interest rates for installments that they receive under the *status quo* however individuals in the treatment group would be offered randomly assigned alternative installment interest rates. The alternative interest rates could be either higher or lower, or even zero, and by comparing the demand for installment loans for the treatment and control groups, we could essentially use the random assignment as a valid "instrument" to help solve the problem of endogeneity in the interest rate, and make valid inferences about the conditional demand for credit.[2]

In this paper we exploit the quasi-random nature of the free installment promotions as a *quasi random experiment* (QRE) to help us estimate customer responses to the free installment opportunity. However in order to fully exploit the opportunity provided by the free installment offers, we do have to make some additional assumptions. In particular, the self-selected nature of customers' decisions to take advantage of free installment offers is compounded by another potentially serious measurement issue, namely *censoring.* That is, *our data only allows us to observe free installment offers when customers actually choose them, however for all other non-free installment transactions, we cannot observe whether the customer was not offered a free installment opportunity, or if the customer was offered a free installment opportunity but the customer chose not to take it.* Since we are willing to make some reasonable assumptions and put some additional structure on the credit choice problem, the next section presents our "solution" for inferring the choice of installment term and the conditional demand for credit.

[2]Note that Ausubel and Shui [2005] analyzed data from a randomized experiment, but it was not a RCE since there were no "controls" corresponding to the subjects who were offered the "treatments" (i.e. the six introductory offers). However to a certain extent the individuals who were offered different introductory offers could be regarded as controls. For example the individuals who were offered a 7.9% 12 month introductory offer could serve as controls for the individuals who were offered the 4.9% 6 month introductory offer, but doing this only allows us to test how customers respond to one of these offers relative to the other one. They cannot tell us how the customers who accepted either of these introductory offers behaved relative to customers who were not offered either introductory offer: the company would have have to have included an explicit control group to do this —- i.e. a 7th group of customers who decided to sign up for the credit card without being offered any special introductory offer.

7 A structural model of the conditional demand for credit

In view of these problems with the reduced-form approaches, we adopted an alternative *structural approach* to analyzing the effect of the interest rate on a customer's choice of installment term. This approach does require us to make some assumptions and develop a simple but we think relatively flexible *model* of how consumers make installment choices. We assume that a customer, with characteristics x, evaluates each transaction in terms of the *net utility* of postponing the payment of the purchase over a term of d months. The customer faces an interest rate $r(x,d)$ for borrowing over a term of d months, except that $r(x,1) = 0$, i.e. all customers get an "interest free loan" if they choose to pay the purchase amount a in full on the next statement date. We normalize the net utility of this "pay in full" option, $d = 1$, to 0. However for the installment purchase options $d = 2,3,\ldots,12$ we assume that the net utility is of the form $v(a,x,r,d) = ov(a,x,d) - c(a,r,d)$ where $ov(a,x,d)$ is the *option value* to a customer with characteristics x of paying for the purchase amount a over d months rather than paying the amount in full a the next statement date (which has an option value normalized to 0 as indicated above, $ov(a,x,1) = 0$).

The function $c(a,r,d)$ is the *cost of credit* equal to the (undiscounted) interest that the customer pays for an installment loan of amount a over duration d at the interest rate r. The net utility

$$v(a,x,r,d) = ov(a,x,d) - c(a,r,d) \tag{7}$$

can therefore be regarded as capturing an elementary cost/benefit calculation that the customer makes each time he/she makes a transaction with their credit card.

We add onto each of the net utilities $v(a,x,r,d)$, $d = 1,2,\ldots,12$ an additional Type III extreme value error component $\varepsilon(d)$ that represent the effect of "other idiosyncratic factors" that affect an individual's choice of installment term that are independent across successive purchase occasions, so that the overall net utility of choosing to purchase an amount a on an installment of duration d months is $v(a,x,r,d) + \sigma\varepsilon(d)$, where $\sigma > 0$ is a scale parameter that determines the relative impact of the "idiosyncratic factors" $\varepsilon(d)$ relative to the "systematic factors" affecting decisions as is captured by $v(a,x,r,d) = ov(a,x,d) - c(a,r,d)$. Examples of factors affecting a person's choice that might be in the $\varepsilon(d)$ term is whether there is a long line at checkout (so the customer feels uncomfortable weighing the options $d = 2,\ldots,12$ relative to doing the "default" and choosing $d = 1$), or if a customer has time-varying but uncorrelated psychological uncertainty about what other bills or payments may be due at various upcoming months $d = 2,\ldots,12$.

As is well known, when we "integrate out" these unobserved components of the net utilities we obtain a *mulitinomial logit model* that represents a consumer's choice of the installment term. For consumers who are not offered any free installment purchase opportunity, their choice set is the full set of 12 alternatives $d \in \{1,2,\dots,12\}$. However for a consumer that is offered a free installment opportunity, which is an offer to pay the purchase amount a over a term of δ months, we make a *dominance assumption* that the customer will strictly prefer a free installment opporunity of duration δ over any positive interest rate installment of *shorter* duration, $d = 2,3,\dots,\delta-1$. This implies that when a customer is offered a free installment opportunity over a duration of δ months, the customer's choice set is $\{1,\delta,\delta+1,\dots,12\}$, and the corresponding choice is a multinomial logit model over this reduced choice set. Notice also that $r(x,\delta) = 0$ for a free installment purchase opportunity, but a consumer might rationally choose a *longer installment horizon at a positive interest rate* if the option value of having a loan for a longer duration is sufficiently high to outweigh the added interest cost.

If we observed whether consumers had a free installment option *regardless of whether or not they choose the free installment option* our life would be much simpler. Then we could write a *full information likelihood function* that is the product of the probablity of whether or not the customer is offered a free installment option or not on any specific purchase occasion times the probability of their choice of installment term (where the choice probability is conditional on whether they are offered a free installment option or not). This would result in a relatively easy estimation exercise, where we could use a flexibly parameterization for the option value function and estimate the model no differently than most static discrete choice models are estimated.

However we face a more difficult task since our observations of free installment options are *censored* in a way that is very similar to *choice based sampling:* that is, we only observe whether a consumer is offered a free installment option for those purchases where the consumer actually chose the free installment option. In such a situation, how is it possible to infer the probability that customers are offered free installment options? More importantly, how can we estimate the probability that customers do not choose the free installment option when it is offered to them? We show that we can solve the problem by forming a likelihood function that accounts for the censoring. The likelihood function takes the form of a *mixture model* where the probability of being offered a free installment option is a key part of the *mixing probabilities* (there are additional component corresponding to a probability distribution over the duration δ offered to customers who are offered free installment options).

Though there are well know econometric diffficulties involved in identifying mixture models, and the degree of censoring in our application is very high (we only observe free installments being chosen in 2.7% of the 167,946 customer-purchase observations used in our econometric analysis), we show that under reasonable but *parametric* assumptions about the forms of the probability function governing free installment options and for flexibly parameterized functional forms for customers' option value functions $ov(a,x,d)$, we are able to separately identify the probability of being offered a free installment, $\pi(z)$ which depends on a set of variables z including time dummies and merchant class code dummies and consumers' conditional choice probabilities for installments $P(d|a,r,d,x)$.

We find that our model fits the data well, but implies a highly inelestic demand for credit. In particular, we find a relatively limited degree of consumer responsiveness to free installment options: the probability of turning down these options is relatively high even though we estimate that for our sample customers are offered free installments approximately 27% of the time, Thus, these customers are taking free installments in only about 10% of the times that they are offered them. We refer to this low take-up rate of what would appear to be a "costless" option for an interest-free loan as the *free installment puzzle.*

Our data are not sufficiently detailed to enable us to delve a great deal further and uncover a more detailed explanation for the reasons *why* customers appear so unwilling to take up free installments and their demand for credit is so inelastic. Our model attributes the reasons for this low takeup rate to a combination of a relatively low option value of credit relative to the cost of credit and to relatively high fixed transactions costs associated in undertaking each installment purchase transaction. However these "transactions costs" could also be interpreted as capturing *stigma* associated with installment transactions, and the low option value may be associated with a fear that installment credit balances could undermine one's credit rating, or that there are some unspecified hidden future fees associated with installments beyond the interest rate (e.g. an unfounded belief that there are pre-payment penalties, or a concern that an installment balance could lead to a higher risk of missed future payments and thus late fees). Unfortunately, we are unable to delve further to determine which of these various more subtle psychological explanations is the dominant explanation of the free installment puzzle.

Customers who were not offered interest-free installment purchase options, or who desire a greater number of installment payments than they are offered under an interest-free installment opportunity can borrow (with no explicit borrowing limits) to pay for a credit card purchase for durations ranging from 2 to 12 months (future statement dates) according to a *nonlinear, customer-specific interest rate schedule.* These schedules are determined according to a rather complex function of a) the consumer's credit score

Figure 29 Interest Premium for Installment Purchases as a function of the Installment Term

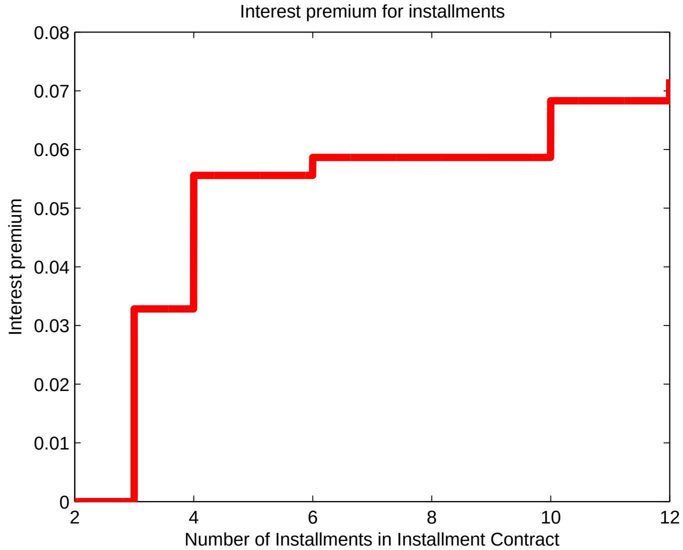

and payment history (including the number of recent late payments), b) the number of installment payments, and c) the current economic environment, including the level of overall interest rates and dummy variables capturing current economic conditions. The most important factors are a) and b). It is a good approximation that the consumer characteristics a) determine the "base interest rate" for an installment loan with $d = 2$ payments, and there is a step-wise increasing schedule (common to all consumers) that determines successive increases in the interest rate offered for longer installment terms $d > 2$. Figure 29 graphs the interest "premiums" customers must pay for successively longer installment terms d.

Below we will let $\overline{r}(d,x)$ denote the *installment interest rate schedule* offered to a consumer with characteristics x who desires to finance an installment purchase with d installments. By our discussion above, this schedule has the form

$$\overline{r}(d,x) = \rho_0(x) + \rho_1(d), \tag{8}$$

where the characteristics of the particular consumer x only enter via the "intercept" term $\rho_0(x)$, and $\rho_1(d)$ represents the *interest premiums* for installments longer than $d = 2$ months. Thus $\rho_2(d) = 0$ for $d \leq 2$ and $\rho_2(d) > 0$ is given by the function graphed in figure 29 for $d \geq 2$. Note that our regression analysis of actual interest rates charged to customers confirms that the ρ_2 function is, to a first approximation, common for all consumers x.

Consider a consumer with characteristics x who is interested in purchasing a given item that costs an amount a. We take as a given that the consumer is going to make the purchase and has decided to use their credit card to make this purchase, but we do model here the customer's choice of whether to pay the balance a in full at the next statement ($d = 1$), or request an installment purchase option with $d > 2$ installments at an interest rate of $r = \overline{r}(d,x)$. Let $v(d,x,a,r)$ represent the net gain in utility the consumer obtains from choosing installment option d (where again, we have normalized the net gain for paying in full, $d = 1$ to $v(1,x,a,r) = 0$. Since we do not expect to be able to perfectly predict every consumer's choice of installment period d, we introduce to commonly used device of extreme value distributed unobservable (to the econometrician) component of utility $\varepsilon(c)$ that also affects the choice of installment period. Thus, the consumer chooses installment period $d \in D = \{1,2,\ldots,12\}$ if and only if

$$v(d,x,a,\overline{r}(c,x)) + \varepsilon(d) \geq \max_{d' \in D} \left[v(d',x,a,\overline{r}(d',a)) + \varepsilon(d') \right]. \tag{9}$$

The extreme value assumption implies that the conditional probability of observing the consumer choose installment period d is (after integrating out the unobserved components of utility $\{\varepsilon(d')|d' \in D\}$ is given by the standard multinomial logit model

$$P(d|a,a,r_+) = \frac{\exp\{v(d,x,a,\overline{r}(d,X))\}}{\sum_{d' \in D} \exp\{v(d',x,a,\overline{r}(d',x))\}}, \tag{10}$$

where r_+ denotes a choice situation where the consumer can only choose among installment periods that have positive interest rates.

The consumer's choice problem is slightly more complicated when the consumer is offered an interest-free installment option. Recall from the discussion above, that when the consumer is offered such an option, the term of the installment is generally fixed at some value $d > 1$ but d is generally less than or equal to 6 months (although in certain circumstances interest-free installment options are available for as long as $d = 12$ months. In this case, the consumer must choose the best of the following options, a) $d = 1$, b) the d offered under the interest-free installment option, which we denote by d_0, and c) the possibility of choosing an installment period d longer than d_0 at an interest rate of $\overline{r}(d,X)$. The consumer will choose the interest-free installment option d_0 if

$$v(d_0,x,a,0) + \varepsilon(d_0) \geq \max \left[v(1,x,a,0) + \varepsilon(1), \max_{d' > d_0} \left[v(d',x,a,\overline{r}(d',a)) + \varepsilon(d') \right] \right], \tag{11}$$

and the probability of this happening is

$$P(d_0|x,a,d_0) = \frac{\exp\{v(d_0,x,a,0)\}}{\exp\{v(1,x,a,0)\} + \exp\{v(d_0,x,a,0)\} + \sum_{d' > d_0} \exp\{v(d',x,a,\overline{r}(d',X))\}}, \tag{12}$$

where the conditioning argument d_0 in $P(d_0|x,a,d_0)$ denotes the situation where a customer is offered an interest-free installment purchase option with term off d_0. Similarly, the probability $P(1|x,a,d_0)$ denotes the probability the customer will choose to pay in full at the next statement date even though he/she was offered an interest-free installment option with term d_0, and for $d > d_0$, $P(d|x,a,d_0)$ denotes the probability that a customer chooses to pay by installment, but with a longer term d than offered under the interest-free installment option, at the cost of paying a positive interest rate $r = \overline{r}(d,x)$.

Now consider the likelihood function for a given customer who makes purchases at a set of times $T = \{t_1,\ldots,t_N\}$. Of these times, there is a subset $T_I \subset T$ where the customer purchased under installment, i.e. where $d > 1$. The complement T/T_I consist of times where the customer purchased without installment, i.e. where $d = 1$. We face a censoring problem that in many cases where $d = 1$, we do not know if the consumer was eligible for an interest-free installment purchase option or not. Even when $d > 1$, we only know if the consumer was offered an interest-free installment purchase option when the customer actually chose that alternative. However it is possible that in some cases customers may have been offered an interest-free installment purchase option with term c_0 but decided to choose a longer term option at a positive interest rate. Our likelihood must be adjusted to account for these possiblities and to "integrate out" the various possible interest-free installment options that the consumer could have been offered but which we did not observe.

Let T_0 be the subset of purchase dates T where the customer did choose the installment option and we observe that this was an interest-free installment option (we can determine this by observing that the consumer never made interest payments on the installments as described above). For this subset, the component of the likelihood is

$$L_0(\theta) = \prod_{t \in T_0} P(d_0|x_t,a_t,d_0) \tag{13}$$

where for each transaction in the set of times T_0, the chosen installment term d_t is equal to the term offered to the customer under the interest-free installment option, d_0. The vector θ denotes parameters entering the function $v(d,x,a,r)$ function to be estimated, and this will be discussed in more detail below. Thus, when we observe a customer choosing the interest-free installment option, it is clear that they were actually offered it. However in the other cases, $t \in T/T_0$, we do not know for sure if the customer was offered the interest-free installment option or not. There are two possibilities here: a) the consumer chose not to purchase under installment, b) the consumer chose to purchase under installment but paid a positive interest rate.

Let $\Pi(z_t)$ be the probability that a consumer who purchases an item at time t is offered an interest-free installment purchase option. The vector z_t includes dummies indicating the time of the purchase and the type merchant the customer is purchasing the item from, since as we noted above the main determinants of the interest-free installment option are a) the time of year, and b) the type of merchant (since different merchants can negotiate interest-free installment deals with the credit card company as a way of increasing their sales). Conditional on being offered an interest-free installment purchase option, let $f(c_0|Z)$ be the conditional distribution of the installment term that is associated with the interest-free installment option. Note that $f(1|z) = 0$: by definition an installment payment plan must have 2 or more future payment dates. Equivalently, every consumer has the option to pay in a single installment, and they get what amounts to an interest free loan covering the duration between the date of purchase until the next billing date.

Now consider the probability that a consumer chooses to purchase the item at a cost of a without installment, as observed by the econometrician who does not know whether this customer was offered an interest-free installment purchase option or not. Let $P(1|x,z,a)$ denote this probability, which is given by

$$P(1|x,z,a) = \Pi(z)\left[\sum_{d_0>1} f(d_0|z)P(1|x,a,d_0)\right] + [1-\Pi(z)]P(1|x,a,r_+). \tag{14}$$

Similarly consider the probability that a customer chooses to purchase the item at a cost of a under installment with d payments with a positive interest rate, again as observed by the econometrician who does not know whether or not the customer was offered an interest-free installment purchase offer or not. Let $P(d|x,z,a)$ denote this probability, which is given by

$$P(d|x,z,a) = \Pi(z)\left[\sum_{d_0<d} f(d_0|z,d_0<d)P(d|x,a,d_0)\right] + [1-\Pi(z)]P(d|x,a,r_+), \tag{15}$$

where $f(d_0|z,d_0 < d)$ is the conditional probability that a customer who is offered an interest-free installment option is offered one with a term less than the term c that the customer actually chose, given by

$$f(d_0|z,d_0<d) = \frac{f(d_0|z)}{\sum_{d'<d} f(d'|z)}, \quad d_0 \in \{2,\dots,d-1\}. \tag{16}$$

The reasoning is that an interest-free installment option with a term of d would always strictly dominate an installment position with positive interest rates and a term of d. Thus, we assume that the consumer would always choose the former over the latter, so that if we observe that a customer had chosen an installment term d with a positive interest rate, we infer that this customer could not have been offered an interest-free installment option with a term greater than or eqaul to d. Note that if $d = 2$, then we deduce that the customer could not have been offered any interest-free installment option, so $P(2|x,z,a) = P(2|x,a,r_+)$.

Let $L_1(\theta)$ denote the component of the likelihood corresponding to purchases that the consumer makes in the subset T/T_0, i.e. purchases either that were not done under installment, or which were done under installment but at a positive interest rate. This is given by

$$L_1(\theta) = \prod_{t \in T/T_0} P(d_t|x_t, z_t, a_t). \tag{17}$$

where $d_t = 1$ if the customer chose to purchase an item at time t without installment, and $d_t > 1$ if the customer chose to purchase via installment, but with a positive interest rate. The overall likelihood is then $L(\theta) = L_0(\theta)L_1(\theta)$. We maximize this with respect to θ for various "flexible functional forms" for $v(d,x,a,r)$ that are designed to capture the net "option value" to the customer of purchasing an item under installment. We assume that $v(d,x,a,r)$ has the additively separable representation given in equation (7) above. Thus, we can view consumers as making "cost-benefit" calculations where they compare the benefit or option value $ov(a,x,d)$ of paying a purchase amount over $d > 1$ installments with the interest costs $c(a,r,d)$. For free installments, we have $c(a,r,d) = 0$, but this does not necessarily imply that customers will necessarily always take every free installment option. One reason is due to the randomly distributed *IID* extreme value shocks $\varepsilon(d)$ representing unboserved idiosyncratic factors that affect a consumer's choice of the installment term. In some cases these shocks will be sufficiently negative to cause a consumer not to take a free installment offer even if $ov(a,x,d)$ is positive (and thus higher than the utility of paying the purchase in full at the next statement date, which is normalized to 0). Another reason is that we specify the option value function as follows

$$ov(a,x,d) = a\rho(x,d) - \lambda(x) \tag{18}$$

where we can think of $\rho(x,d)$ as the customer x's *subjective term structure* (i.e. the subjective interest rate that represents the customer's maximal willingness to pay to obtain a loan of duration d months) and $\lambda(x)$ represents the *fixed transaction costs of deciding and undertaking an installment transaction at the checkout counter.* Note that this component is assumed not to be a function of the amount purchased a whereas the other component of the option value, $a\rho(x,d)$ is a linear function of the amount purchased. This implies that *consumers will not want to pay for sufficiently small credit card purchases on installment since the benefit of doing this, $a\rho(x,d)$, is lower than the transactions cost $\lambda(x)$.*

Notice that we assume the option value of having the benefit of extended payment does not depend on the interest rate the credit card company charges the customer, and that the interest rate schedule $\overline{r}(d,x)$ only enters via the cost function $c(a,r,d)$. This is an important identifying assumption. Furthermore we assume that the financial cost that a customer percieves due to purchasing an item under installment equals the excess of the total payments that the customer makes over the term of the agreeement less the current cost a of the item. That is, we assume c equals the difference between the total payments the customer makes under the installment agreement *cumulated with interest to the time the installment agreement ends* less the amount the customer purchased, a, discounted back to the date t when the customer purchased the item. This value can be shown to be

$$c(a,r,d) = a\,(1 - \exp\{-rt_d/365\})\,, \tag{19}$$

where t_d is the elapsed time (in days) between the next statement date after the item was purchased and the statement date when the final installment payment is due. The interest rate r is the internal rate of return on the installment loan, and is given by $r = \overline{r}(d,x)$. Recall that this is the positive interest rate that company offers to the customer for an installment purchase with term d. Notice that $d = 0$ if the consumer chooses not to do an installment. In this case $t_d = 0$, and the consumer automatically gets an interest-free loan from the credit card company from the date of purchase until the next statement date. Notice also that for any interest-free installment opportunity, $r = 0$ and so $c = 0$ in this case as well. To a first approximation (via a Taylor series approximation of the exponential function) we have $c(a,r,d) = \overline{r}(d,x)at_d/365$, so the cost of the installment loan is proportional to the duration of the loan, the amount of the loan, and the interest rate offered to the consumer.

Notice that the $c(a,r,d)$ function has no unknown parameters to be estimated. The parameters to be estimated are the parameters ϕ entering the option value function, $ov(a,x,d,\phi)$, the scale parameter σ of the Type III extreme value distributions for the unobserved components of the $v(a,x,r,d,\phi)$ functions, and α, the parameters entering the probability of being offered a free installment, $\Pi(z)$ and the conditional distributions of the term of the free installments that are offered to consumers, $f(d_0|z,d_0 < d)$. So we let $\theta = (\sigma,\phi,\alpha)$ as the full set of parameters to be estimated. Table **??** presents the maximum likelihood estimates of (σ,ϕ) and appendix 1 presents the maximum likelihood estimates of the 26 α parameters. Clearly, the parameters of interest are (σ,ϕ). We are interested in the α parameters only to the extent that we are interested in learning the conditional probability $\Pi(z,\alpha)$ that credit card customers are offered free installment options when shopping at different merchants at different periods of time.

To understand the parameter estimates, note that we have specified $ov(a,x,d) = a\rho(x,d)$ where

$$\rho(x,d) = \frac{1}{1+\exp\{h(x,d,\phi)\}} \tag{20}$$

where

$$h(x,d,\phi) = \phi_0 + \phi_1 \text{ib} + \phi_2 \text{installlshare} + \phi_3 I\{r=0\} + \phi_4 \text{creditscore} + \phi_5 \text{nlate} + \sum_{j=6}^{12} \phi_j I\{d=j\} + \tag{21}$$

Further, the fixed transaction cost of choosing an installment term at the checkout counter, $\lambda(x)$, is assumed to be given by

$$\lambda(x) = \phi_{13} \exp\{\phi_{14} \text{installshare} + \phi_{15} I\{r=0\}\}. \tag{22}$$

The variable *creditscore* is the interpolated credit score for the customer at the date of the transactions (the company only periodically updates its credit scores so we only observed them at monthly intervals), and *nlate* is the number of late payments that the customer had on his/her record at the time the transaction was undertaken, and *ib* is the customer's installment balance at the time of the transaction. Note that we normalized both *a* and *ib* to represent them as *ratios over the customer-specific average statement amount.*

The most important variable of the *x* variables turned out to be *installshare,* the share of creditcard spending that the customer does under installment. Though one could accuse of us as making the mistake of including *installshare* in our model due to the fact that it is "endogenous" our explanation for doing so is that our analysis in section 4 revealed that *installshare* carries important information about which consumers are most likely to be liquidity constrained. We found that neither *creditscore* nor *nlate* has as powerful impact on enabling our model to fit the data as *installshare* has. We agree that it would be preferable to replace *installshare* by a random parameter τ representing *unobserved heterogeeity* with the interpretation that lower values of τ indicate customers who are more desperate for liquidity and thus have a higher subjective willingness to pay for loans of various durations, $\rho(x,d,\tau,\phi)$. However, we have had considerable difficulty so far in estimating specifications with unobserved heterogeneity due to the fact that we have an unbalanced panel where for some consumers we observe many hundreds of transactions. Conditioning on τ, the likelihood for these hundreds of conditionally independent choices of installment duration is typically a *very very small number.* Unobserved heterogeneity specifications require us to take averages (i.e. integrate over the distribution of τ) of these very small numbers and we often found that when we tried to take the logarithm of the resulting *mixture probability* it was sufficiently small to be below the "machine epsilon" i.e. the lowest positive number a computer is capable of representing, even

Table 3: Maximum Likelihood Parameter Estimates
Dependent variable: choice of installment term, d. Amounts in parentheses are estimated standard errors.

Variable	ρ function parameter estimate
σ	0.0340
ϕ_0 (constant)	3.537
ϕ_1 (ib)	-0.959
ϕ_2 (installshare)	-3.878
ϕ_3 $(I\{r=0\})$	1.568
ϕ_4 (creditscore)	1.99^{-5}
ϕ_5 (nlate)	-0.009
ϕ_6 $(I\{d=2\})$	-2.306
ϕ_7 $(I\{d=3\})$	-2.576
ϕ_8 $(I\{d=4\})$	-0.861
ϕ_9 $(I\{d=5\})$	-1.422
ϕ_{10} $(I\{d=6\})$	-1.590
ϕ_{11} $(I\{d=7\})$	1.162
ϕ_{12} $(I\{d=10\})$	-0.159
Variable	λ function parameter estimate
ϕ_{13} (constant)	-0.225
ϕ_{14} (installshare)	-0.614
ϕ_{15} $(I\{r=0\})$	0.101
$\log(L(\hat{\theta}))$	-54836.6
Number of observations	167,946

on 64-bit machines. So we have found *installshare* to be extremely convenient as an "observed indicator" of the underlying unobserved heterogenity τ. We conjecture that if we can somehow resolve the problem of "underflow" in computing the mixing probabilities, the estimation results (particularly the overall fit of the model) of a specification with a sufficiently rich specification of unobserved heterogeneity but omitting *installshare* will be quite similar to the results presented below with τ omitted and *installshare* included.

The estimation results in table 3 show the result, that is not surprising in light of our analysis in section 5, that the "subjective discount factor" $\rho(x,d)$ is an increasing function of *installshare.* Thus, the maximum likelihood estimation attempts to best-fit the data by making sure that the "installment addicts" have the highest option value for credit. In addition, the negative coefficient of *installshare* in the $\lambda(x)$ function means that it is possible to further improve its fit of the model to the data by having the transactions costs of doing an purchase via installment be lower for customers who are more installment-prone as measured by *installshare.*

The variables *creditscore* and *nlate* variables have very negligible impacts on the option value (and they are not important determinants of transactions costs either). We believe that this is largely because the *installshare* variable does such a much better job of capturing the underlying unobserved heterogeneity in our sample of customers. The estimated installment dummies reflect a rather strange pattern in the estimated subjective discount factor $\rho(x,d)$ to help the model match the fact that installment loans of relatively shorter term, particularly 3 month installment loans, are substantially more popular than the longer term loans. Of course the sharply increasing interest rate schedule for longer term installment loans that we plotted in figure 29 does a large of the work in helping the model to fit the data: the significantly higher cost $c(a,r,d)$ of the longer duration installment loans does dissuade many customers from selecting them, however the disincentives to borrowing for longer durations inherent in the $c(a,r,d)$ function is evidently not enough to fit the model well. The maximum likelihood algorithm estimates values to the installment duration dummies that provide an extra boost (or extra penalty by sharply lowering the implied subjective discount factor for longer term loans) to enable the model to better explain the popularity of 3 month installment loans.

Figures 30 and 31 make the estimation results easier to understand by plotting the implied discount factors $\rho(x,d)$ and the implied continuous time "subjective interest rates" $r_s(x,d)$ defined by

$$r_s(x,d) = -\log(\rho(x,d))/d, \tag{23}$$

i.e. $r(x,d)$ represents a sort of "subjective term structure of interest rates" implied by the model estimation results. $r_s(x,d)$ can be regarded as the maximum interest rate that a consumer of type x is willing to pay for a loan of duration d. We see that both the subjective discount factors $\rho(x,d)$ and the estimated subjective term structure are generally strongly declining in the duration of the installment loan. Figure 31 indicates that the subjective interest rates start out as high as 10% for 2 or 3 month loans but drop rapidly to nearly zero for loans that have a duration of 7 months or longer. The model also implies that consumers regard that there is some extra value for interest free installments: they assign a significantly higher subjective option value to interest-free loans than they do to loans done at positive interest rates. Of course, the maximum likelihood estimates the parameters this way in order to best fit the take-up of free installments. The figures also illustrate how the *installshare* increases the subjective option values and interest rates.

We conclude this section with figures 32 an 33 which summarize the ability of the structural model to fit the credit card data. Of course the predominant choice by consumers is to pay their credit card purchases in full by the next installment date: this is the choice made in 93.57% of the customer/purchase transactions in our data set. When we simulate the estimated model of installment choice, taking the x and

Figure 30 Predicted versus Actual Installment Choices

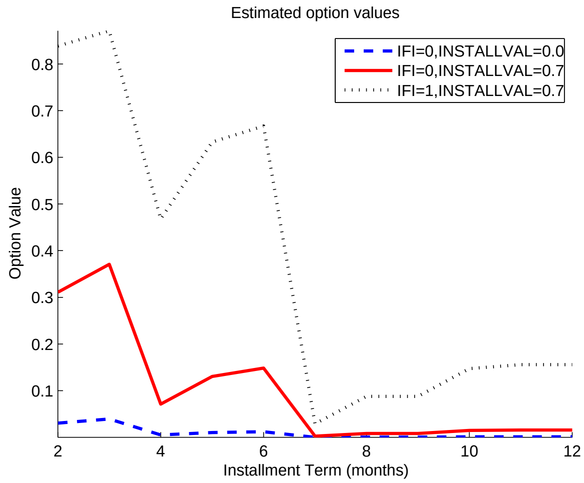

Figure 31 Predicted versus Actual Free Installment Choices

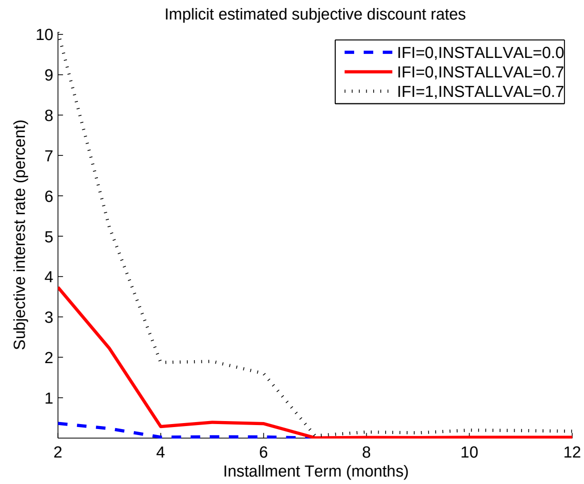

purchase amounts a as given for the 167,946 observations in our data set, we obtain a predicted (simulated) choice of paying in full at the next statement (i.e. to choose $d = 1$) of 93.06%.

Of more interest is to judge the extent to which our model can predict the installment choices made by the customers in our sample, i.e. to predict the incidence of choices $d > 1$. Figure 32 plots the predicted versus actual set of *all* installment choices made the customers in our sample. We see that the model fits the actual choices well overall, but underpredicts the choice of 3 month installment terms. Figure 33 shows that the model does much better in predicting the number of free installments that consumers chose. In the actual data, 2.7% of our customer-purchase transactions were done as free installments. Our model simulations predict that 2.4% of the simulated customer-purchase transactions would be done as free installments.

As we noted in the introduction and elsewhere, our simulations also predict something that we could not otherwise learn from our data without having a structural model: the model predicts that in 27% of 167.946 simulated customer-purchase transactions, the company would offer the customer a free installment opportunity. As we noted in section 4, this estimate strikes us as quite reasonable. Recall that we found that the most installment prone "addicts" were doing about 20% of all of their purchases as free installments. Howevever these individuals are doing roughly 70-80% of all their transactions via installment but the remaining 20-30% are smaller purchases that even the "installment addicts" do not bother to do as installments (perhaps even for them the transaction costs are too large for the smallest purchases). So assume that these installment addicts are encountering free installment opportunities 27% of the time, and whenever the amount they are purchasing is sufficiently large, the installment addicts will always to choose to take the free installment option. Then since 20% = 75% $\times$ 27%, our maximum likelihood estimates are consistent with our observation that approximately 20% of all of the installment transactions by the "installment addicts" are done as free installments.

8 Conclusions

The main contribution of our paper is to show, despite our highly censored data, that for whatever reason, our sample of customers appear to be highly reluctant to spending on installment, to the extent that even free installment offers fail to motivate most customers to take advantage of them for most transactions. Instead, we find that the customers who are most likely to take advantage of free installlment offers are the

Figure 32 Predicted versus Actual Installment Choices

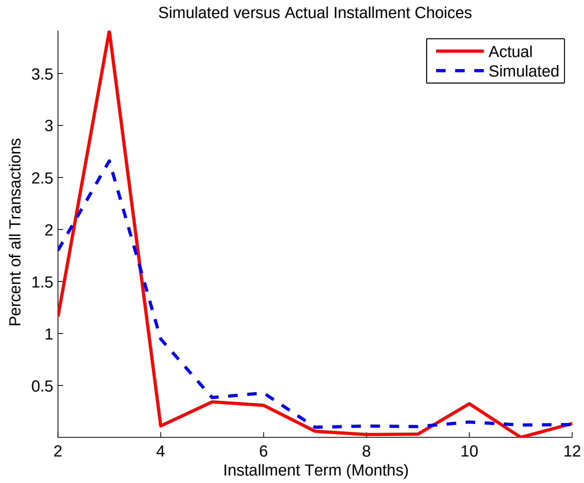

Figure 33 Predicted versus Actual Free Installment Choices

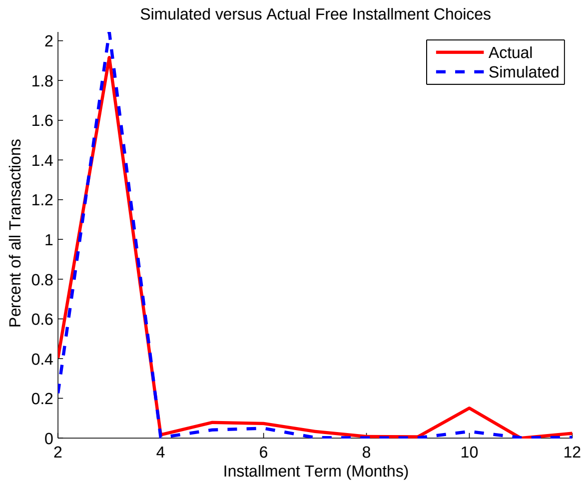

"installment addicts" that we identified in our exploratory data analysis in section 4.

Otherwise the structural approach to estimation of demand for installment credit produces reasonable estimates of the demand for installment credit, and we showed the model fits the observed patterns of spending quite well. We can use our model to make out of sample predictions of the effects of a wide range of policy changes, including the impact of abolishing the free installment option on the use of installments, and the impact of raising or lowering the installment interest rate. However we require a more complete dynamic model, and more data, to study a larger range of credit policy questions, including the impact of interest free installments on credit card adoption decisions by individuals who do not already have this credit card, and how credit policy affects the spending/purchase decision itself.

Ultimately, we need to compliment the use of econometric models and forecasts with *experiments* that enable us to test to accuracy of the out of sample forecasts of our econometric models. We hope this study will be a stepping stone that will convince the credit card company to provide more data, and undertake key experiments that would be necessary to better model and identify key aspects of the credit card usage and spending decisions that affect the company's market share and ultimately its profitability.

References

Sule Alan, Ruxandra Dumitrescu, and Gyongyi Loranth. Subprime consumer credit demand – evidence from a lender's pricing experiment, 2011.

Lawrence Ausubel and Haiyan Shui. Time inconsistency in the credit card market, 2005.

David S. Evans and Richard Schmalensee. The economics of interchange fees and their regulation: An overview, 2005.

Marc Rysman. Empirical analysis of payment card usage. *Journal of Industrial Economics*, 60:1–36, 2007.

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