Essays on Household Finance

by

Bruno Ferman

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Economics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2012

Author ....
Department of Economics
August 15, 2012

Certified by ....
Joshua Angrist
Ford Professor of Economics
Thesis Supervisor

Certified by ....
Esther Duflo
Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics
Thesis Supervisor

Certified by ....
Shawn Cole
Associate Professor of Graduate Business Administration, HBS
Thesis Supervisor

Certified by ....
Michael Greenstone
3M Professor of Environmental Economics
Thesis Supervisor

Accepted by ....
Michael Greenstone
3M Professor of Environmental Economics
Chairman, Department Committee on Graduate Theses

Essays on Household Finance

by

Bruno Ferman

Submitted to the Department of Economics
on August 15, 2012, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

This dissertation consists of three essays. The first chapter studies whether credit demand is sensitive to interest rates, to the prominence of interest rate disclosure, and to nudges. Consumer credit regulations usually require that lenders disclose interest rates. However, lenders can evade the spirit of these regulations by concealing rates in the fine print and highlighting low monthly payments. I explore the importance of such evasion in Brazil, where consumer credit for lower and middle income borrowers is expanding rapidly, despite particularly high interest rates. By randomizing contract interest rates and the degree of interest rate disclosure, I show that most borrowers are highly rate-sensitive, whether or not interest rates are prominently disclosed in marketing materials. An exception is high-risk borrowers, for whom rate disclosure matters. These clients are rate-sensitive only when disclosure is prominent. I also show that borrowers who choose this type of financing are responsive to nudges that favor longer-term plans. Despite this evidence, the financial consequences of information disclosure, even for high-risk borrowers, are relatively modest, and clients are less susceptible to nudges when the stakes are higher. Together, these results suggest that consumers in Brazil are surprisingly adept at decoding information even when lenders try to obfuscate the interest rate information, suggesting a fair amount of sophistication in this population.

The second chapter (co-authored with Leonardo Bursztyn, Florian Ederer, and Noam Yuchtman) studies the importance of peer effects in financial decisions. Using a field experiment conducted with a financial brokerage, we attempt to disentangle channels through which a person's financial decisions affect his peers'. When someone purchases an asset, his peers may also want to purchase it because they learn from his choice ("social learning") and because his possession of the asset directly affects others' utility of owning the same asset ("social utility"). We randomize whether one member of a peer pair who chose to purchase an asset has that choice implemented, thus randomizing possession of the asset. Then, we randomize whether the second member of the pair: 1) receives no information about his peer, or 2) is informed of his peer's desire to purchase the asset *and* the result of the randomization determining possession. We thus estimate the effects of: (a) learning plus possession, and (b) learning alone, relative to a control group. In the control group, 42% of individuals purchased the asset, increasing to 71% in the "social learning only" group, and to 93% in the "social learning and social utility" group. These results suggest that herding behavior in financial markets may result from social learning, and also from a desire to own the same assets as one's peers.

The third chapter (co-authored with Pedro Daniel Tavares) uses data on checking and savings accounts for a sample of clients from a large bank in Brazil to calculate the prevalence and cost of "borrowing high and lending low" behavior in a setting where the spread between the borrowing and saving rates is on the order of 150% per year. We find that most clients maintain an overdrawn account at least one day a year while having liquid assets. However, the yearly amount of avoidable financial charges would only correspond, on average, to less than 0.5% of clients' yearly earnings. We also show that consumers are less likely to engage in such behavior when the costs of doing so are higher. These results suggest that the spread between the borrowing and saving rates is a key determinant of this behavior.

Thesis Supervisor: Joshua Angrist
Title: Ford Professor of Economics

Thesis Supervisor: Esther Duflo
Title: Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics

Thesis Supervisor: Shawn Cole
Title: Associate Professor of Graduate Business Administration, HBS

Thesis Supervisor: Michael Greenstone
Title: 3M Professor of Environmental Economics

Acknowledgments

First, I thank my advisors, Joshua Angrist, Shawn Cole, Esther Duflo, and Michael Greenstone, for invaluable advice and feedback at all stages of my thesis research. They have taught me much, and it has been a great pleasure working with them.

I am very grateful to Juliano Assunção, Dana Chandler, Daniel Gottlieb, Horacio Larreguy, Simone Schaner, Heidi Williams, and numerous participants at the MIT Labor Lunch, MIT Development Lunch, MIT Finance Lunch, MIT PF/Labor Seminar, MOOD 2011 Doctoral Workshop, and CFS-EIEF Conference on Household Finance for helpful comments and suggestions.

It has been a great honor working with Leonardo Bursztyn, Florian Ederer, and Noam Yuchtman on the second chapter of thesis, and with Pedro Daniel Tavares on the third chapter of this thesis. I have learned a great deal from working with them, and I am very grateful for having this pleasant opportunity.

I would also like to especially thank the management and staff of the cooperating financial institutions during the implementation of the studies in chapters 1, 2 and 3. This thesis would not have been possible without the tireless assistance, hard work, and commitment of many employees in these firms.

Most of all, I thank my parents for the support they provided me through my entire life, and my wife and best friend, Marianna, for always cheering me up and supporting me through the bad and good times in this process. This thesis is dedicated to Marianna.

Contents

Chapter 1

Reading the Fine Print: Credit Demand and Information Disclosure in Brazil

1.1 Introduction

High-cost consumer credit has attracted significant academic and regulatory attention. The debate is intensified by claims that lenders conceal or misrepresent rates, which has made interest rate disclosure a major focus for regulators[1]. Even when information disclosure regulations require that lenders present interest rates in a standardized way, lenders can easily evade such regulations if clients have limited attention. As argued by Barr et al. (2008), lenders can use more salient terms (for example, "Low Monthly Payments!") to compete with the interest rate information for borrowers' attention, while still complying with disclosure regulations. These concerns are especially relevant in emerging markets, where consumer credit is novel and interest rates are particularly high[2]. In Brazil, the volume of credit card borrowing increased by a factor of 12 over the last 10 years, despite the fact that credit card revolving interest rates are usually higher than 10% *per month*[3].

[1]Typical lending practices before the enactment of the Truth-in-Lending-Act are described in National Commission on Consumer Finance (1972). There are also reports that microlenders stress weekly payments rather than long-term interest rates, and - when pressed to report their average annual rates - misrepresent rates by not taking into account declining balances (*BusinessWeek*, 2007).

[2]From 2000 to 2008, household debt as a proportion of GDP increased in all BRIC countries (Roxburgh (2010)).

[3]Brazilian Central Bank.

In this paper, I test whether credit demand is sensitive to interest rates, to the prominence of interest rate disclosure, and to nudges. In a randomized field experiment conducted with a large credit card company in Brazil, a sample of 19,690 clients were offered a menu of payment plans that allows them to pay down their balances over a 6-12 month period in fixed installments, as an alternative to the revolving credit line. When offering these contracts, credit card companies usually conceal rates in the fine print, while featuring a long-term payment plan with low monthly payments. In the experimental design, I randomly varied three features of the offer: (i) the offered interest rate: which ranged from 3.99% to 11.89% per month; (ii) the degree of interest rate disclosure: whether buried in a footnote, or prominently disclosed; and (iii) the featured payment plan: although all clients were offered 4 different payment plans, one plan was prominently featured in the advertisement.

The experimental results show that most clients are surprisingly adept at decoding loan terms, suggesting a fair amount of sophistication for this population. Clients are interest rate elastic, even when rates are not prominently disclosed. In addition, on average, prominent rate disclosure has only small and not statistically significant effects on take-up rates and interest rate elasticities. However, the degree of rate disclosure is relevant for an important subpopulation. High-risk clients are not rate-sensitive when the interest rate is concealed in the fine print, although they become interest rate elastic when this information is more prominently disclosed[4]. These results suggest that high-risk borrowers are less attentive to the details of the contract, so that prominent rate disclosure affects their borrowing decisions. However, even for the high-risk group, the financial consequences of information disclosure are modest.

Also, clients are no more likely to enroll in a payment plan when a longer-term plan (with lower monthly payments) is featured. This suggests that clients consider all available options. However, conditional on enrollment, there is evidence that "nudges" are relevant in determining which plans clients choose. Although clients revealed preferences for short-term plans, they can be nudged into enrolling in longer-term plans when a longer-term plan appears more prominently. However, clients are less susceptible to nudges when interest

[4]Risk categories are based on a borrower risk classification system used by the firm, which is based on information supplied by credit bureaus and on the credit card company's own data. Clients are more likely to be classified as high-risk when they have a lower credit score on credit bureaus, when they use the revolving credit line more often, when they make late payments, and when they use a higher proportion of their credit limit. According to this classification system, 10% of the clients are classified as high-risk.

rates are higher and when longer-term plans are featured.

The results reported here add to a recent body of experimental and quasi-experimental evidence in both developed and developing countries which suggests that consumers are responsive to prices in credit markets (for example, Karlan and Zinman (2008), Gross and Souleles (2002), Huang and Tan (2009), and Attanasio et al. (2008)). These results also fit into a small but growing literature on credit demand and information disclosure. In two recent randomized experiments, Bertrand et al. (2010) and Bertrand and Morse (2011) find that presenting the interest rate of loan contracts has no effect on credit demand[5]. Unlike designs in previous research, the experimental design in this paper allows the effects of interest rate disclosure to be estimated not only on average take-up rates, but also on the interest rate sensitivity of demand. While my results also provide evidence that, on average, interest rate disclosure has limited effect on take-up rates, I show that prominent rate disclosure matters for an important group of clients - namely, high-risk clients who are rate-sensitive only if the interest rate information is prominently disclosed. These results are consistent with the work by Stango and Zinman (2011), who studied the effect of weakening enforcement of APR disclosure in the Truth-in-Lending Act; they found that weak enforcement increased the disparity between interest rates paid by more and less sophisticated clients.

Finally, the results in this paper add to a large literature on nudges and default options in financial decisions. The estimated effects of nudges on the probability of choosing a specific plan are at the same order of magnitude as the effects of default options in 401(K) decisions in the US (Beshears et al. (2009)). Note, however, that in the experiment presented here clients had to make an active decision about in which payment plan to enroll. Still, the display of the options had a strong influence on payment plan choices, suggesting that consumers are susceptible to nudges even when they have to make an active decision. In addition, this paper provides novel evidence that the effectiveness of nudges is weaker when the stakes are higher.

The paper proceeds as follows. Section 1.2 describes the consumer credit market in Brazil, emphasizing the details of the payment plans offered by credit card companies. Section 1.3 describes the experimental design and the empirical strategy. The results of the

[5]Bertrand and Morse (2011) find that information that helps consumers aggregate the costs of payday loans over time has a significant effect in reducing take-up rates.

experiment are presented and discussed in Section 1.4, while section 1.5 concludes.

1.2 Economic Environment

1.2.1 Credit in Brazil

The rapid growth in consumer credit has been particularly pronounced in Brazil, where consumer credit doubled between 2000 and 2010. This increase in consumer credit in Brazil is in part explained by an increasing access to credit cards to lower and middle class consumers, especially through credit card companies associated with retail stores. In some cases these cards can be used only at the originating retail store; in other cases they can be used elsewhere. Between 2000 and 2010 the volume of credit card borrowing increased twelve-fold, despite the fact that credit card interest rates in Brazil are particularly high.

Credit card companies in Brazil usually offer two borrowing alternatives if clients do not pay their balance in full. The most common alternative is to use the credit card's revolving line of credit: clients pay an amount equal to or greater than the minimum required payment, but smaller than the credit card balance. The remaining balance plus interest accrued is carried over to the next billing period. For the clients in this study, the revolving rate ranges from 11.89% to 15.99% *per month*, and the minimum payment is equal to 15% of the credit card balance. For the consumer, the main advantage of this source of credit is that it is pre-approved, and clients have the flexibility to choose how much they want to pay, as long as it is at least the minimum payment.

Many credit card companies also offer a menu of installment plans to clients. Clients can choose a fixed period over which they can repay their entire balances with constant monthly payments. This type of credit is also pre-approved and easily accessible. To enroll in a payment plan, clients simply have to pay the exact amount of the monthly payment of the plan they have chosen. In doing so, they automatically enroll in the chosen plan, and they are charged the remaining installments on their credit card statements for the $m-1$ following months. For example, if the credit card statement presents a payment plan offer of 6 installments with a monthly payment of \$183.38, the client simply has to pay exactly \$183.38, and then \$183.38 will be added to his credit card balance every month for five months.

Given a balance (B), the number of installments (m), and the monthly interest rate

(r_m), the monthly payment (M_m) is[6]:

$$M_m = B \times \frac{(1+r_m)^{m-1} \times r_m}{(1+r_m)^m - 1}$$

The number of installments usually ranges from 4 to 24 months, and the interest rate may be equal to or lower than the revolving rate. When advertising a menu of payment plans, firms usually highlight low monthly payments (in some cases lower than the minimum payment), and advertise that these contracts have a "special interest rate" (if the payment plan rate is lower than the revolving rate), although the contract rate is usually concealed in the fine print. Firms offer these contracts in hopes of attracting clients that would not use the revolving credit line, and to lead clients to carry a larger outstanding balance[7].

Under a payment plan contract, clients need to pay the monthly installments in full in order to stay current. The credit card minimum payment will be equal to the monthly installment of the chosen payment plan, plus 15% of any new purchases made with the credit card. This implies that clients might face higher minimum payments when they enroll in a payment plan than if they used the revolving credit line, even if the revolving credit line has a higher interest rate.

Although there is no penalty for canceling a payment plan (if a client cancels the plan, then the present value of the unpaid installments is charged on his next billing cycle), many clients may not be aware of the possibility of canceling these payment plans, and may be afraid of facing additional fees and the hassle of dealing with a bureaucracy. If clients believe it is impossible or costly to cancel a payment plan, then a payment plan would imply a commitment to carry a balance for a longer period. These concerns might prevent some clients from choosing such longer-term credit offers.

Regardless of whether clients enrolled in a payment plan, in case of default or late payments with their credit cards, the interest rate and fees are substantial. If clients do not pay on time, then in addition to the revolving interest rate (which can be up to 15.99% per month) they are charged a late payment fee equal to 2% of their full balance, plus interest charges on future installments. If they remain in default for more than 70 days, their credit

[6]The amount of the monthly payments is defined such that, given the contract interest rate (r_m), the present value of the stream of payments is equal to the credit card balance (B).

[7]However, a separate test conducted by the same firm concluded that offering payment plans with interest rates equal to 6.39% or 9.59% would not be profitable for the firm, at least at those levels. The firm found that most of the clients who enrolled in a payment plan would have already been revolving large balances if they were not offered these plans, and that enrolling in a payment plan increased the probability of default.

cards are canceled. Their names are also reported to credit bureaus, making it harder for them to access other lines of credit.

Other borrowing alternatives for consumers include personal loans from banks (average monthly interest rate of 4.79%), checking account overdrafts (average monthly interest rate of 7.40%), personal loans from finance houses (average monthly interest rate of 9.87%), or informal loans. However, these alternatives might not be available for some clients and, even if available, may require additional applications, paperwork, and delays.

1.2.2 Borrowing Decisions, Limited Attention, and Information Disclosure

When deciding whether to enroll in a payment plan or to use a revolving line of credit, clients face a trade-off between lower interest rates and flexibility in the stream of payments: interest rates on payment plans are usually lower than revolving credit card rates, but payment plans have less flexibility in terms of the stream of payments clients must make. Clients are required to pay exactly the first monthly payment in order to enroll in a payment plan. For example, with a 6-month payment plan and an interest rate of 11.89%, this implies paying only around 22% of their credit card balances. In addition, because clients might be unaware of the option to cancel these contracts, they may believe they have an obligation to carry a balance over a longer period.

In both cases, that inflexibility implies that clients might have to distort their optimal consumption streams if they choose such a plan. It might instead be optimal for clients to use their revolving line of credit to pay off a higher fraction of their balances, thus paying off their debt in a shorter period, even if this means borrowing at a higher interest rate. Past data reveals that, in the absence of payment plan offers, clients pay more than 30% of their balances 93% of the time. Also, less than 10% of clients used the revolving line of credit for six consecutive months. In other words, even though payment plans usually have lower interest rates, enrolling in such plans might imply carrying a larger balance than usual for most clients.

The lower the interest rate, the more attractive the payment plans should be for the clients. However, figuring out whether a payment plan has an attractive interest rate may not be a straightforward task. While information disclosure regulations in Brazil require that credit card companies disclose information on the interest rates, this information is usually hidden, increasing the computational cost of comparing payment plans and alternative options. Although a "rational" consumer would be able to calculate the interest rate on payment plans given the number of installments and monthly payments, limited financial literacy and cognitive biases may prevent clients from correctly evaluating contract interest rates. In fact, there is no closed-form solution to calculate the interest rate of an installment plan given the number of installments and the monthly payments. There is a consistent body of evidence showing that consumers make mistakes when assessing interest rates[8], and Stango and Zinman (2009a) provide evidence that consumers systematically underestimate interest rates when given a other terms of the contract.

A second important decision in this setting is the maturity choice of the payment plan. Clients are usually presented with several payment plan options, with varying numbers of installments and monthly payments. In selecting a payment plan, clients must balance the size of their monthly payments with the duration over which they will have to repay their debt. The higher the interest rate, the more costly it is to choose a long-term relative to a short-term payment plan. Therefore, it would be expected that clients choose shorter-term contracts and make more careful decisions when the interest rates are higher.

1.3 Experimental Design and Implementation

The field experiment was carried out with a large credit card company in Brazil. The credit cards, issued by a major retailer in Brazil, are regular credit cards, accepted in most retail locations throughout the country. At the time of the experiment, this company had more than 5 million active clients, most of them lower- and middle-income. Therefore, the conclusions based on the sample analyzed in this study should be relevant for understanding an important group of consumers: lower- and middle-income consumers that are starting to gain access to credit in emerging markets. A sample of credit card clients was selected to receive a menu of payment plan offers, which were varied as described below.

1.3.1 Treatments

To understand the interest rate elasticity of credit demand, in a first treatment dimension clients were randomly assigned to groups in which the *monthly* interest rate of the payment

[8]For example, see Juster and Shay (1964), National Commission on Consumer Finance (1972), Day and Brandt (1974), and Parker and Shay (1974)

plans was 3.99%, 7.49%, or 11.89%. Given the payment plan's interest rate, clients had four different payment plan options, with the number of installments varying between 6, 8, 10, and 12 months. The Table below displays the monthly payments for each contract, assuming a balance of $1000.00. A client with an assigned interest rate of r_m would be able to choose among any of the contracts in the corresponding column.

Contract options for a client with a balance of $1000:

Assigned $i = 3.99\%$ Available choices:	Assigned $i = 7.49\%$ Available choices:	Assigned $i = 11.89\%$ Available choices:
6 × $183.38	6 × $198.14	6 × $216.70
8 × $142.77	8 × $158.77	8 × $179.22
10 × $118.50	10 × $135.47	10 × $157.47
12 × $102.40	12 × $120.21	12 × $143.55

Along with their credit card statement, clients received a one-page advertisement describing the payment plan offers.

In order to estimate how credit demand is affected by the prominence of the interest rate disclosure, in a second treatment dimension clients were randomly chosen to receive one of two different advertisement layouts. The standard advertisement states that the client could pay off his balance using payment plans with a *special interest rate*. The advertisement then displays an example of one of the payment plan options, saying "you can pay off your balance of B in m installments of M_m". Figure 1 presents the one-page advertisement for the payment plans, along with the credit card statement. In that advertisement layout, the interest rate of the contract is not prominently disclosed and is only present in a footnote. Clients were also shown a table with all four payment plan options at the top of their credit card statement. In this table, one of the payment plans (the same as in the one-page advertisement) appears more prominently, and the interest rate of the payment plans is presented in a small font size next to the table. Therefore, clients presented with this marketing material are encouraged to focus on the monthly payments (M_m) of the featured plan rather than on the interest rate.

The alternative layout for the one-page advertisement is exactly the same as the standard ad, except that the interest rate information is more prominently disclosed (Figure 2). If a client were assigned an interest rate of 11.89%, then this layout would state that he could pay off his balance using payment plans with a *special interest rate of 11.89%* (in large font size). The same applies for clients assigned an interest rate of 3.99% and 7.49%. Clients who received this layout also had the same table at the top of their credit card statement as those who received the standard layout. Even though both layouts present the same information, the different layouts could affect clients' decisions because of limited attention, since the value of the interest rate was more *salient* in the alternative layout (DellaVigna (2009)). Since the variation in the information disclosure is orthogonal to the changes in the interest rates, it is possible to estimate not only how information disclosure affects average take-up rates, but also how it affects the interest rate sensitivity of demand.

Finally, in order to test the hypothesis that consumers are more attracted by low monthly payments, in a third treatment dimension the featured plan (which is more prominently presented in the advertisement) was randomly assigned among clients. The 12-month contract is the standard plan featured by the firm, which expects clients to focus on the plan's monthly payments and thus be more attracted when a lower monthly payment is prominent. The featured plan can also be relevant in determining which plan clients actually choose among the menu of options. Having many options may create feelings of conflict and indecision (for example, Shafir et al. (1993), Bertrand et al. (2010), and Iyengar et al. (2004)). In this case, the featured plan can work as a nudge if clients use it as a "default option" in order to avoid making a decision about which payment plan to choose.

1.3.2 Description of Experimental Sample

In this experiment, 19,690 credit card clients received a one-time menu of payment plans, in either July or September 2010[9]. Using a borrower risk classification system used by the firm (based on credit bureaus and on the credit card company's own data), medium- and high-risk clients were oversampled in order to provide more precision on the estimates for these groups. Medium- and high-risk clients comprise, respectively, 7% and 10% of the population. All summary statistics and estimates are weighted by the inverse of the probability that the clients were selected so that they represent the original population. All results are similar if sampling weights are not used.

Table 1.1 shows the sample size in each treatment cell. Table 1.2 presents the baseline

[9]The credit card company only offers payment plans to clients with credit card balances greater than *R*\$100.00 who are not in default.

characteristics of the final sample. The average credit card balance was $R\$661$ (during the experiment, the exchange rate was US\$1 $\approx R\$1.75$). In around 27% of the cases, clients used the revolving credit line, and even conditional on using the revolving credit line, clients pay on average 60% of their balances. These numbers are lower than the proportion of households with outstanding credit card balances in the US, which might be because revolving interest rates are much higher in Brazil.

Columns 2 to 4 of Table 1.2 show baseline characteristics separately for each risk category group. According to this classification, around 10% of the clients are defined as high-risk. Using the revolving line of credit, making late payments, and using a higher proportion of the credit limit enter negatively in the risk assessment of the firm. Not surprisingly, these variables are all higher for high-risk clients. Table 1.2 also reports the 12-month probability of default for each category group[10]. As expected, low-risk clients have a lower probability of default (5%) than high-risk clients (22%).

Since clients were randomly assigned to each treatment group, averages for all baseline variables are well-balanced across the different treatment cells. Appendix table 1.A.1 presents the averages for baseline variables in the final sample in each interest rate x advertisement layout and suggested maturity cells, and p-values of the tests that each of these variables has the same mean across the different treatment cells. Appendix tables 1.A.2 to 1.A.4 present the same information when low-, middle-, and high-risk clients are analyzed separately.

1.3.3 Empirical Strategy

Given the fact that the interest rates on payment plans and the advertisement layouts were randomly assigned, the effect of the interest rate and information disclosure on credit demand can be estimated simply by comparing the mean take-up rate across cells. The following linear probability model (or logit) is used to estimate the interest rate sensitivity of demand:

$$E_i = \alpha + \beta_0 \cdot D_i + \beta_1 \cdot r_i + \beta_2 \cdot r_i \cdot D_i + \varepsilon_i \tag{1.1}$$

[10]This is the probability that, conditional on being current in the base month, a client does not make the minimum payment for 70 days at some point within the following 12 months. These calculations are based on an outside sample.

where E_i is equal to one if client i enrolled in a payment plan, r_i is the interest rate offered, and D_i is a dummy variable equal to one if client i received the alternative advertisement with the interest rate prominently disclosed. The omitted group is clients who received the standard advertisement layout (where the interest rate is not prominently disclosed). Because rates were randomly assigned, $\hat{\beta}_1$ yields consistent estimates of the interest rate sensitivity of the demand for payment plans when the lender conceals the interest rate information and encourages clients to focus on low monthly payments. Because the advertisement layout was also randomly assigned, $\hat{\beta}_2$ yields consistent estimates of the effects of information disclosure on the interest rate sensitivity.

The interest rate elasticities under the two different advertisement layouts may then be calculated as:

$$\hat{\eta}_{r,standard} = \hat{\beta}_1 \times \frac{\bar{r}_{standard}}{\bar{E}_{standard}} \text{ and } \hat{\eta}_{r,alternative} = (\hat{\beta}_1 + \hat{\beta}_2) \times \frac{\bar{r}_{alternative}}{\bar{E}_{alternative}} \tag{1.2}$$

where $\bar{r}_{layout}$ and $\bar{E}_{layout}$ are, respectively, average interest rate and take-up rates under the standard (interest rate presented only in the fine print) or the alternative (interest rate prominently disclosed) advertisement layout. The standard errors of $\hat{\eta}_{r,standard}$ and $\hat{\eta}_{r,alternative}$ are bootstrapped, which takes into account the fact that the average take-up rates also are estimated. Again, because of the random assignment of advertisement layouts, the differences between $\hat{\eta}_{r,alternative}$ and $\hat{\eta}_{r,standard}$ reveal the effect of prominent interest rate disclosure on interest rate elasticities.

Similarly, the importance of the featured plan can be estimated by comparing the mean take-up rate across the featured plan cells. Further, the following linear probability model (or logit) is estimated in order to test the hypothesis that demand is higher when a payment plan with lower monthly payments is more prominently presented (while holding the menu of options constant):

$$E_i = \alpha + \gamma \times m_i + \varepsilon_i \tag{1.3}$$

where m_i is the maturity of the featured plan offered to client i. The elasticity with respect to the maturity of the featured plan may be calculated as:

$$\hat{\eta}_m = \hat{\gamma} \times \frac{\bar{m}}{\bar{E}} \tag{1.4}$$

where $\bar{m}$ is the average maturity of the featured plan.

1.3.4 Experiment to Test the Effect of Payment Plan on Default

In addition to the field experiment described above, I will also rely on a larger scale field experiment conducted by the firm at the same time, with another group of clients. The credit card company carried out an experiment in July of 2010 with 103, 116 clients, randomly allocated in three groups. In the first group, 34, 743 clients were offered a menu of payment plans with interest rate equal to 6.39%, in a second group, 49, 573 clients were offered plans with interest rate equal to 9.59%, and finally a control group of 18, 800 clients did not receive any payment plan offer. This experiment was designed to measure the impact that enrolling in a payment plan has on future repayment. The results of this experiment will be combined with the results of my experiment in order to provide more precise estimates of the effects of information disclosure on probability of default in section 1.4.3.

1.4 Results

1.4.1 Interest Rate Elasticity and Information Disclosure

The first set of results shows the sensitivity of payment plans' demand to the interest rates when the interest rate is concealed in a footnote (as in the standard advertisement layout). Table 1.3, column 1, presents the payment plans' take-up rates for each interest rate offer when the interest rate is not prominently disclosed. The results show that payment plans' demand responds to interest rate changes. The average take-up rate is 2% when the payment plans' interest rate is 11.89%, and it doubles when the interest rate falls to 3.99%. I strongly reject that take-up rates are equal for all interest rate values. The implied interest rate elasticity, estimated from (1.2), is -0.713 (s.e. 0.137), negative and statistically different from zero[11].

The fact that clients are interest rate elastic even when the most salient information is the monthly payments is not surprising, since payment plans with higher interest rates have higher monthly payments. However, given the difficulties in assessing the interest rate of an installment plan based on the number of installments and the monthly payments,

[11]The estimated coefficients from model (1.1) using both a linear probability model and logit are presented in Appendix table 1.A.5.

it is harder for consumers to compare the payment plan to other borrowing alternatives based only on this information. Comparing the demand for payment plans under the two different advertisements will provide evidence on whether clients are able to assess the cost of credit based on the monthly payments (or able to look for the interest rate in the fine print). Column 2 of Table 1.3 presents the payment plans' take-up rates when the interest rate is prominently disclosed, while column 3 shows the difference in take-up rates when the interest rate is prominently versus when it is not prominently disclosed.

Displaying the interest rate prominently results in an increase in take-up rates when the interest rate is 3.99%, and a decrease in take-up rates when the interest rate is 11.89%. However, these differences are small and not statistically different from zero. Panel ii of Table 1.3 presents the interest rate elasticities. Clients are slightly more interest rate elastic when the interest rate information is more prominent (the elasticity goes from -0.713 to -0.880), but it is not possible to reject that the elasticities are equal under the two layouts[12].

These results suggest that, on average, clients are able to assess the cost of credit even when the credit card company conceals the interest rate information in the fine print, leading clients to focus on the monthly payments of the contracts. Therefore, as long as lenders are required to disclose interest rates, regulations that prevent lenders from concealing rates in the fine print should have, on average, small effects on consumers decisions.

1.4.2 Interest Rate Elasticity and Information Disclosure - Heterogeneity

The results for the full sample suggest that clients are sensitive to interest rate changes and that changing the salience of the contract interest rate has a limited effect on clients' behavior. However, these results hide an important heterogeneity when clients are classified according to their default risk.

Columns 1, 4, and 7 of Table 1.4 show payment plan take-up rates separately for low-, medium-, and high-risk clients. The demand for payment plans increases with the risk profile of the clients. More importantly, take-up rates are fairly constant across different interest rates for high-risk clients. A joint test fails to reject the hypothesis that that take-up rates are equal for all interest rates for these clients, with a p-value of 0.645. Panel ii reports the estimated interest rate elasticity for each risk group when the interest rate is

[12]Given the standard errors, it would be possible to detect an effect of prominent rate disclosure on the interest rate elasticity of the order of 0.3.

not prominently disclosed. Interest rate elasticity is equal to -0.884 (s.e. 0.173) for the low-risk clients and -0.526 (s.e. 0.270) for medium-risk clients. For the high-risk clients, though, the estimated interest rate elasticity is equal to -0.172 (s.e. 0.242), which is both small and statistically equal to zero.

Comparing take-up rates and interest rate elasticities under the two advertisement layouts for low- and medium-risk clients, we find that more prominent interest rate disclosure has no impact on the behavior of these clients. These results are reported in columns 3 and 6 of Table 1.4. For high-risk clients, though, prominent disclosure strongly reduces payment plans' demand when the interest rate is equal to 11.89% (from 5.7% to 2.9%), and increases payment plans' demand when the interest rate is equal to 3.99% or 7.49% (though these differences are not statistically significant). The interest rate elasticity for this high-risk group changes from -0.172 (s.e. 0.242) when the interest rate is concealed to -0.929 (s.e. 0.205) when it is prominently disclosed. The hypothesis that the interest rate elasticities are invariant to the advertisement layout for this group of clients can be rejected at the 1.7% level.

These results are consistent with low- and medium-risk clients being more careful when choosing among different borrowing options, suggesting that such clients look for the interest rate of contracts even when this information is not prominently displayed in the advertisement. These results also suggest, however, that high-risk borrowers are less attentive to the details of a contract, so that the salience of the interest rate actually impacts their borrowing decisions. The behavior of high-risk clients also suggests that the experimental manipulation produced a meaningful difference in the available information, although the more sophisticated clients were able to work around that.

1.4.3 Effects of Information Disclosure on Subsequent Financial Outcomes for High-Risk Clients

The results in sections 1.4.1 and 1.4.2 reveal that while the degree of interest rate disclosure has, on average, small and not statistically significant effects on payment plan enrollment decisions, it significantly affects high-risk clients decisions. In particular, high-risk clients are 2.8 percentage points less likely to enroll in a payment plan with a high interest rate when the interest rate information is prominently disclosed. But does this affect their welfare or even future financial position in any meaningful way?

While the scale of the experimental design in this paper is not sufficient to estimate how this change in behavior translates into subsequent financial outcomes, such as default probability[13], it is possible to use the larger scale experiment carried out by the same firm described in section 1.3.4. With this larger scale experiment, it is possible to estimate the causal effect of enrolling in a payment plan using the following specification:

$$Y_i = \alpha + \rho_{6.39} \times E_{6.39,i} + \rho_{9.59} \times E_{9.59,i} + \varepsilon_i \tag{1.5}$$

where Y_i is default in the 12 months after the offer, and $E_{r,i}$ is equal to 1 if client i enrolled in a payment plan with interest rate equal to r. Coefficient ρ_r is the causal effect of enrolling in a payment plan with interest rate r on outcome Y. Note that this causal effect can depend on the interest rate offered. In order to estimate $\rho_{6.39}$ and $\rho_{9.59}$, offers of payment plans with rates 6.39% and 9.59% are used as instruments for payment plan enrollment at these two rates.

The results presented in Table 1.5 indicate that, on average, enrolling in a payment plan at either of these two interest rates significantly increase the probability of default in the following 12 months. Considering only the high-risk clients, enrolling in a payment plan induces more clients to default when the interest rate is higher (9.59%).

Extrapolating these results, enrolling in a payment plan with an even higher interest rate should also increase the probability of default. Since prominent rate disclosure reduces the demand of high-risk clients for payment plans when the interest rate is equal to 11.89%, this information treatment likely reduces the probability of default for clients that change their payment plan enrollment decisions because of the treatment.

Although these estimates suggest that high-risk clients who did not enroll in a payment plan when offered the full information treatment are likely better off, another estimate of interest is the reduced form impact of information disclosure on default probability. Assuming that information disclosure only affects subsequent financial outcomes through the enrollment decision, consider the model:

[13]In a reduced form regression of information disclosure on default for high-risk clients offered a high interest rate payment plan, the standard error on the information disclosure coefficient is equal to 2.3 percentage points. Although the estimate is not statistically different from zero, it would not be possible to reject that information disclosure has large effects on default, given the large standard errors.

$$\begin{cases} E_{r,i} = \alpha_1 + \beta_r \times D_{r,i} + \varepsilon_{1,i} \\ Y_i = \alpha_2 + \rho_r \times E_{r,i} + \varepsilon_{2,i} \end{cases} \tag{1.6}$$

where β_r is the causal effect of information disclosure on payment plan enrollment, and ρ_r is the causal effect of payment plan enrollment on outcome Y_i. Note that these effects can vary with r. Combining these two equations:

$$Y_i = \tilde{\alpha} + \pi_r \times D_{r,i} + \tilde{\varepsilon}_i, \text{ where } \pi_r = (\beta_r \times \rho_r) \tag{1.7}$$

Therefore, the reduced form effect of prominent rate disclosure would be the product of the effect of prominent rate disclosure on enrollment (β_r) and the effect of enrollment on the outcome variable (ρ_r). This strategy is similar in spirit to Angrist and Krueger (1992) Two-Sample IV. However, instead of using estimates from the reduced form and from the first stage in order to back out the structural parameter, the strategy used here combines the first-stage estimates from my experiment, and the structural parameter estimates from the firm experiment in order to produce an estimate of the reduced form. That is, the overall effect of prominent rate disclosure on default[14].

The first experiment provides an estimate for $\beta_{11.89}$ ($\hat{\beta}_{11.89} = -0.028$, s.e. 0.012), while the second experiment provides an estimate for $\rho_{9.59}$ ($\hat{\rho}_{9.59} = 0.121$, s.e. 0.067). Assuming that $\rho_{9.59} \approx \rho_{11.89}$, combining these two estimates yields the reduced form effect of information disclosure on probability of default for high-risk clients $\hat{\pi} = (0.12) \times (-0.028) = -0.0034$ (s.e. 0.0025)[15]. Note that the standard error is around 10 times smaller than the standard error of the reduced form effects of prominent rate disclosure based on the main experiment (0.25 vs 2.3 percentage points), allowing even modest effects of information disclosure on probability of default to be ruled out.

Therefore, prominent rate disclosure has a significant effect in reducing default for clients that are induced to enroll in a payment plan with a high rate because the interest rate information is concealed in the fine print. However, since it affects only a small proportion of consumers, the aggregate effect of such policy would be small even for the population of

[14]In contrast to Angrist and Krueger (1992) Two-Sample IV strategy, my implementation relies on the assumption of homogeneous treatment effects. Otherwise, it might be that the estimated ρ_r is not the causal effect of enrolling in a payment plan for clients that change their enrollment decisions because of the information treatment.

[15]Standard error is calculated using the Delta Method.

high-risk clients.

1.4.4 Featured plan and Maturity Choice

Similar to the treatment where the interest rate is disclosed, changing which plan is selected to appear more prominently has no effect on the clients' choice set, and not even on the information content that is provided to them. One hypothesis in the firm, however, is that clients focus mostly on the monthly payments of the featured plan. In this case, take-up rates should be higher when a payment plan with lower monthly payments/longer maturity is featured.

Column 1 of Table 1.6 shows the take-up rates by featured plan. It is not possible to reject the hypothesis that take-up rates do not depend on which plan is featured. In particular, it is possible reject the hypothesis that clients are more likely to enroll in a payment plan when a longer-term plan is more prominently presented. When clients are analyzed separately by risk category groups, there is also no evidence that the changes in the featured plan affect take-up rates.

The lack of demand response to changes in the featured plan does not imply, however, that clients are indifferent with respect to the maturity of their payment plans, nor that the plan that is selected to appear more prominently has no effect on clients' decisions. Table 1.7 presents the distribution of payment plan choices by featured plan. Clients have strong preferences for short-term contracts. Of the clients who choose a payment plan, more then half of them choose the 6-month plan (the shortest available maturity choice), and when the 6-month plan is featured, more than 80% of the clients choose this plan. However, when a longer-term plan is featured, many clients continue to choose the 6-month plan, but a large fraction of them simply follow the featured plan.

These results suggest that even when one option is more prominently presented, clients are able to consider other alternatives if this option is not attractive. However, the payment plan that appears more prominently has a strong influence in determining which payment plan the clients choose. In particular, clients can be nudged into choosing longer-term plans.

Assuming that the featured plan affected which payment plan the clients chose, but that it had no effect their decision to enroll or not in a payment plan, it is also possible to determine if clients had worse subsequent financial outcomes because they were induced to choose a longer-term payment plan. Given that the featured plan had no effect on take-up rates, and that clients who enrolled in a payment plan in each featured plan treatment cell do not differ in terms of the baseline variables (see Appendix table 1.A.6), this seems to be a reasonable assumption. The following model is estimated:

$$Y_i = \alpha + \lambda \times m_i + \varepsilon_i \tag{1.8}$$

where m_i is maturity of the payment plan chosen by client i. This variable is instrumented by a set of dummy variables indicating whether the 8-, 10-, or 12-month plan was featured (the 6-month plan category was omitted). The results in Table 1.8 indicate that clients induced to enroll in a longer term payment plan were more likely to default. Therefore, the featured plan not only affects which plan clients choose, but it also has real effects on clients' financial outcomes.

However, clients are less likely to follow the featured plan when the cost of doing so is higher. The cost of following the featured plan relative to choosing the 6-month plan increases with the maturity of the featured plan and with the interest rate. Table 1.9 reports estimates on how the probability of following the featured plan correlates with the interest rate and with the maturity of the featured plan[16]. Clients are less likely to follow the featured plan when the interest rate is higher and when a longer-term plan is featured. Combining these two dimensions, when the interest rate is high and the 12-month plan is featured, only 20% of the clients who enroll in a payment plan follow the featured plan, while when the interest rate is low and the 8-month featured, almost 60% of the clients follow the featured plan. Therefore, although clients can be nudged into choosing a longer term payment plan, this effect is less relevant when the relative cost of following the featured plan is higher.

[16] It is important to note that all of the regressions in Table 1.9 are conditional on payment plan enrollment. Therefore, the identification assumptions in these models do not follow directly from the experimental design. For the regressions that estimate the relationship between probability of following the featured plan and the maturity of the featured plan, this selection problem should be less relevant. As reported above, the featured plan has no effect on average take-up rates. Also, there is no evidence that, conditional on payment plan take-up, clients differ in terms of the baseline variables presented in Table 1.2 (see Table A6). This problem can be more relevant for the interest rate regressions. Since clients are interest rate elastic, clients who enroll in a payment plan when the interest rate is high are different from clients who enroll when the interest rate is low. In order to mitigate this problem, I include all the baseline variables in Table 1.2 as controls variables.

1.5 Conclusion

In this paper, I test whether lenders are in fact able to exploit clients' limited attention through advertisement strategies that conceal the interest rate and induce clients to focus on low monthly payments. The results presented here indicate that, overall, these strategies have small and not statistically significant effects on clients' choices. On average, Brazilian credit card holders are sensitive to the interest rate even when interest rate information is not prominently disclosed, and making this information more salient has only a small and not statistically significant effect on take-up rates and interest rate sensitivity. Also, the firm is not able to attract more clients to enroll in payment plans by featuring a payment plan with lower monthly payments (and longer maturity), which again suggests that clients are sophisticated in considering all the available information.

The results in this paper also demonstrate the importance of considering the possibilities of heterogeneous effects of information disclosure by borrower risk. While the effects of interest rate disclosure are, on average, small and not statistically significant, it has a significant effect on credit demand decisions for an important population. High-risk clients are not sensitive to the interest rate when the information is concealed in the fine print, but they become interest rate elastic when the interest rate is prominently disclosed. However, even though information disclosure reduces the probability of default for high-risk clients who avoid a high interest rate payment plan when the interest rate information is prominent, the aggregate effects of information disclosure are small even for the pool of high-risk clients.

Overall, these results suggest that, as long as lenders are required to present interest rates, most consumers are adept at decoding this information, even when lenders try to obfuscate the interest rate information by making it less salient. Furthermore, even though the degree of information disclosure affects the decisions of high-risk clients, regulating how lenders must disclose the information would have only small effects on consumers' subsequent financial positions, even if such regulation is targeted for this group of clients. Therefore, these results suggest that the benefits of regulating the way that interest rate information must be disclosed to consumers are unlikely to outweigh its compliance and enforcement costs.

Finally, conditional on enrollment, there is evidence that consumers can be nudged into enrolling in longer-term payment plans. Although clients are apparently able to consider all their options, a large fraction of them simply follow the plan that is more prominently presented. This is consistent with clients following the featured plan in order to avoid making a decision about which payment plan to choose, suggesting that clients might face a high cost of deciding between alternatives. In this case, default options and nudges can have a strong influence in shaping consumers decisions, and therefore should be on the radar of consumer credit regulators. However, consumers are less susceptible to such nudges when the decision involves higher stakes, which implies that the effects of nudges are limited.

1.A Appendix: Figures and Tables

Figure 1-1: **Standard Layout (Interest Rate Concealed in the Fine Print)**

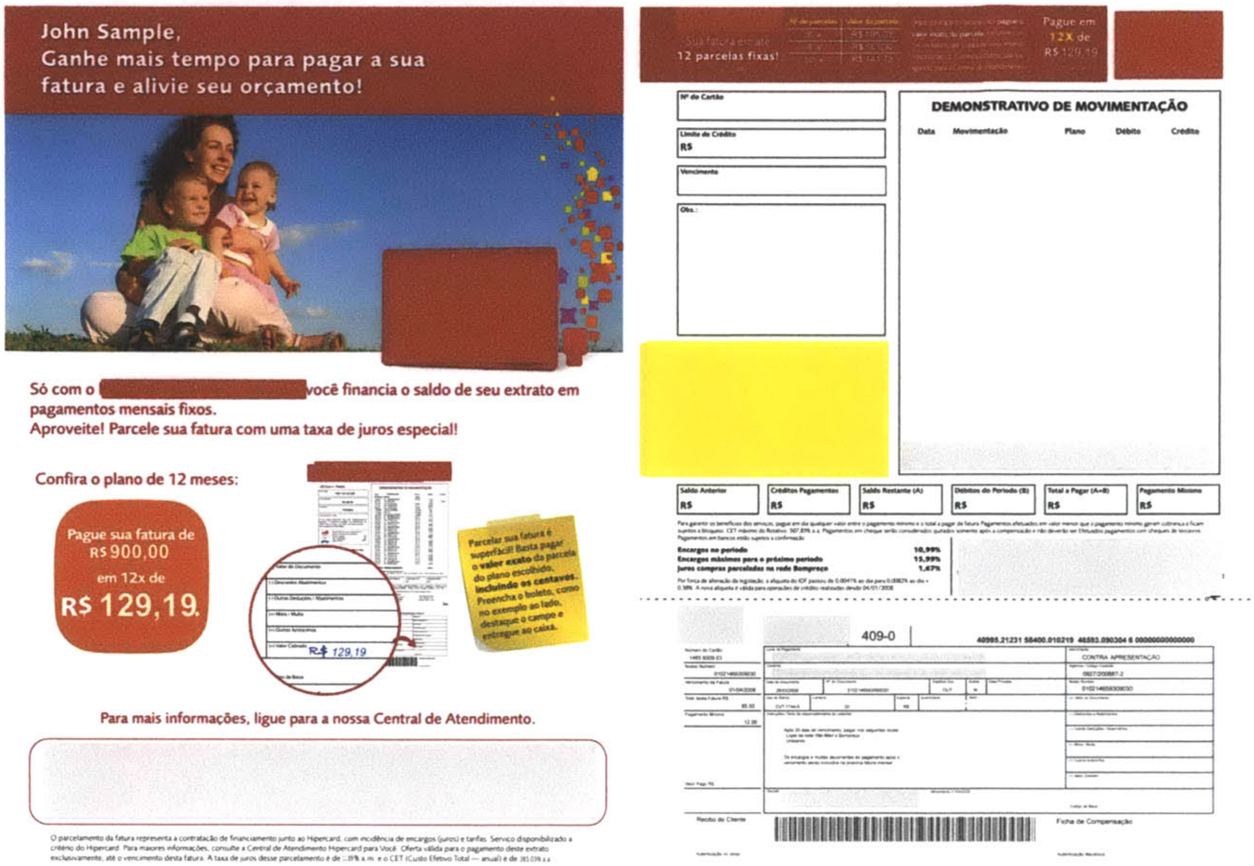

At the top of the page on the left, the advertisement says that the client can have more time to pay off his balance. The text below the pictures says that, with this contract, the client could finance his credit card balance in fixed monthly payments with a special interest rate. The orange box features the 12-month plan, saying that the client could pay off his balance of R900.00 in 12 installments of R129.19. The footnote states that the interest rate of this contract is 11.89% per month.

The page on the right is the credit card statement. On the top of this page, there is a red box stating that the client can pay off his balance in up to 12 fixed installments. In this box, the client is presented with all 4 payment plan options, although one plan is featured. The interest rate of the contract is presented in a small font size.

Figure 1-2: **Alternative Layout (Interest Rate Prominently Disclosed)**

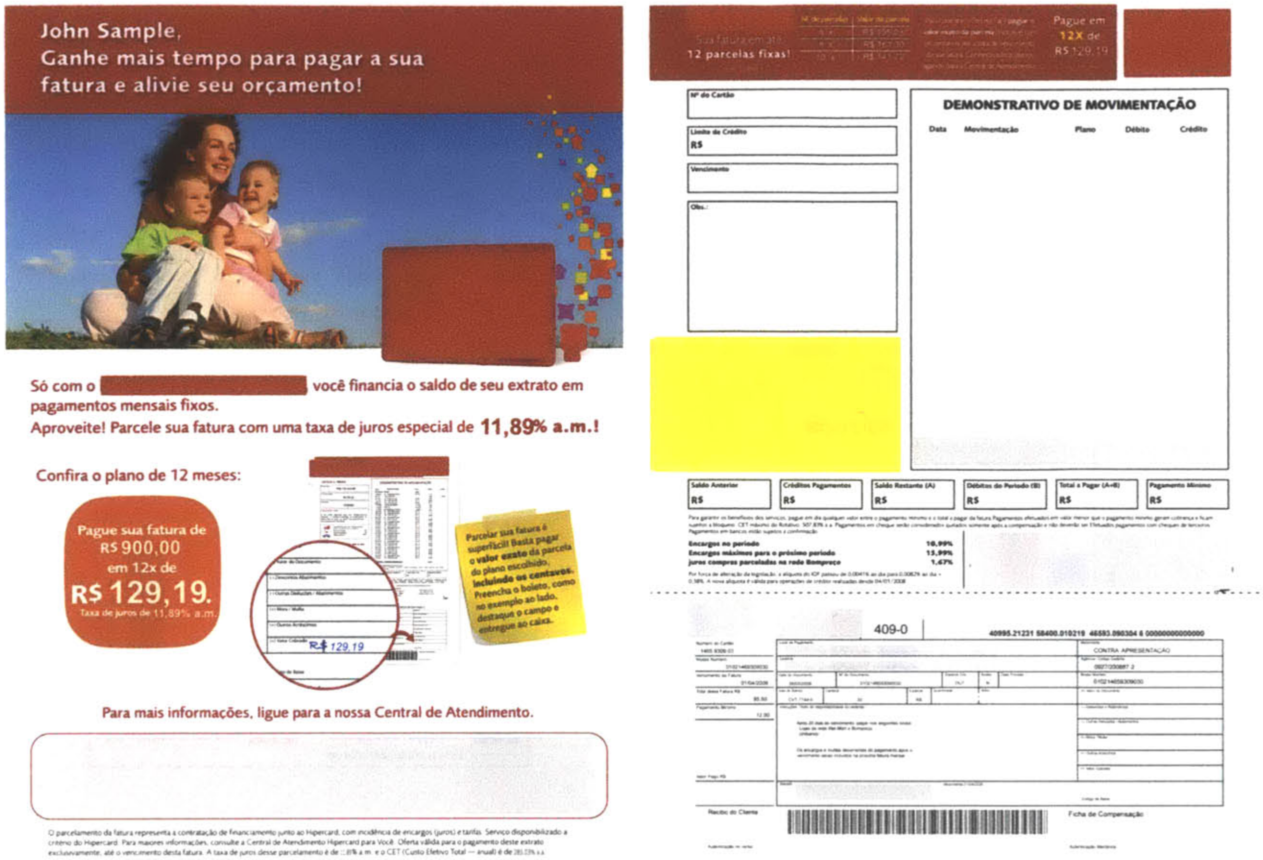

The one-page advertisement is the same as the one in Figure 1, except that the main text says that the client could finance his credit card balance in fixed monthly payments with a special interest rate of 11.89%.

Figure 1-3: **Take-up Rates by Interest Rate x Advertisement Layout**

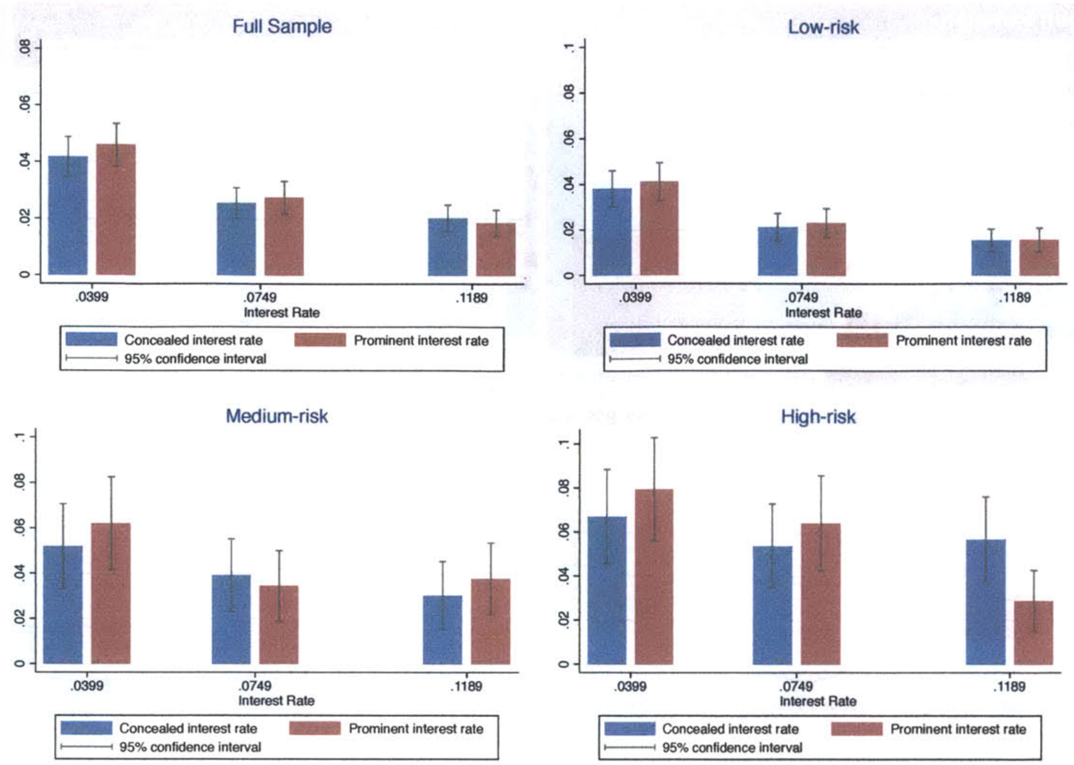

Figure 1-4: **Take-up Rates by Featured Payment Plan)**

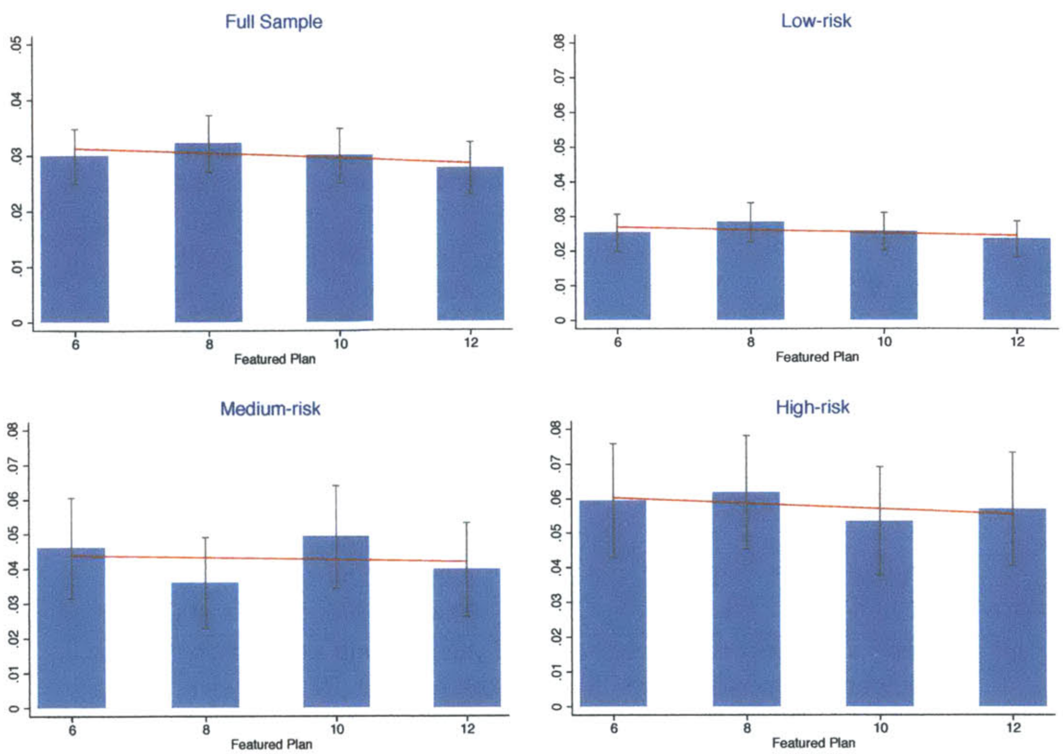

Table 1.1: **Sample Size by Treatment Cells**

Interest rate	Prominent rate disclosure?	Featured plan 6	8	10	12	Total
	No	842	826	801	824	3293
3.99%	Yes	789	817	785	802	3193
	Total	1631	1643	1586	1626	6486
	No	833	855	828	796	3312
7.49%	Yes	800	803	807	819	3229
	Total	1633	1658	1635	1615	6541
	No	797	833	864	828	3322
11.89%	Yes	867	802	834	838	3341
	Total	1664	1635	1698	1666	6663
	No	2472	2514	2493	2448	9927
Total	Yes	2456	2422	2426	2459	9763
	Total	4928	4936	4919	4907	19690

Notes: includes clients who received the payment plan offers (that is, clients with credit card balances greater than R$100.00 who are not in default), and received either the standard or the alternative advertisement layouts presented in Figures 1 and 2.

Table 1.2: **Sample Characteristics**

	Full Sample	Risk categories		
		Low-risk	Medium-risk	High-risk
	(1)	(2)	(3)	(4)
Credit card limit (R$)	1,514.0	1,668.2	766.3	675.9
	[2,184.8]	[2,305.9]	[1,189.4]	[1,009.6]
Credit card balance (R$)	661.0	678.0	603.0	551.5
	[964.4]	[993.4]	[872.0]	[729.4]
Probability of using the revolving credit line	0.304	0.264	0.432	0.572
	[0.460]	[0.441]	[0.495]	[0.495]
Average revolving balance (proportion of current balance)	0.123	0.097	0.192	0.306
	[0.241]	[0.213]	[0.286]	[0.332]
Average monthly interest and fees charges (R$)	23.1	19.4	35.5	47.1
	[73.7]	[66.4]	[96.5]	[104.3]
Probability of making a late payment	0.188	0.168	0.249	0.316
	[0.390]	[0.374]	[0.433]	[0.465]
Time with the credit card (years)	4.65	4.82	3.81	3.72
	[4.02]	[4.15]	[3.14]	[3.16]
12-month probability of default	0.078	0.051	0.136	0.220
	[0.268]	[0.220]	[0.343]	[0.414]
Proportion of population	-	0.839	0.067	0.095
Sample size	19690	13304	3207	3179

Notes: This table presents summary statistics for the final sample of 19,690 clients offered payment plans with the standard or alternative layouts. All summary statistics are weighted by the inverse of the probability that the client was selected so that they represent the original population. Standard deviations in brackets. During the experiment, the exchange rate was US$1≈ R$1.75. Probability of default is the probability that, conditional on being current in the base month, a client does not make the minimum payment for 70 days at some point within the following 12 months. These calculations are based on an outside sample.

Table 1.3: **Effects of Interest Rate and Interest Rate Disclosure on Payment Plan Demand**

	Prominent rate disclosure?		
	No	Yes	Difference
	(1)	(2)	(3)
	i. Take-up rates by interest rate		
All	0.029	0.031	0.001
	(0.002)	(0.002)	(0.002)
Interest rate = 3.99%	0.042	0.046	0.004
	(0.004)	(0.004)	(0.005)
Interest rate = 7.49%	0.026	0.028	0.002
	(0.003)	(0.003)	(0.004)
Interest rate = 11.89%	0.020	0.019	-0.002
	(0.002)	(0.002)	(0.003)
p-value (equal for all interest rates)	0.000	0.000	
	ii. Interest rate elasticities		
Interest rate elasticity	-0.713***	-0.880***	
	(0.137)	(0.132)	
p-value (elasticities are equal)	0.381		
N	9927	9763	

Notes: panel i presents the take-up rates for each interest rate cell and a p-value of an F-test that average take-up rates are equal for all interest rate cells when the interest rate is hidden (column 1) and when the interest rate is emphasized (column 2). Column 3 presents the differences in take-up rates when the interest rate was emphasized and when it was hidden, and for each interest rate groups. Panel ii reports the interest rate elasticities when the interest rate is hidden and when it is emphasized, along with the p-value of a test that these two elasticities are equal. The interest rate elasticities are calculated based on a Linear Probability Model, using equation 2. All estimates are weighted by the inverse of the probability that the client was selected so that they represent the original population. Robust standard errors in parentheses. For column 3 and panel ii: * significant at 10%; ** significant at 5%; *** significant at 1%

Table 1.4: **Effects of Interest Rate and Interest Rate Disclosure on Payment Plan Demand - Risk Category Heterogeneity**

	Low-risk			Medium-risk			High-risk		
	Prominent rate disclosure?			Prominent rate disclosure?			Prominent rate disclosure?		
	No	Yes	Diff	No	Yes	Diff	No	Yes	Diff
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
i. Take-up rates									
All	0.025	0.026	0.002	0.041	0.045	0.004	0.059	0.057	-0.002
	(0.002)	(0.002)	(0.003)	(0.005)	(0.005)	(0.007)	(0.006)	(0.006)	(0.008)
Interest rate = 3.99%	0.038	0.041	0.003	0.052	0.062	0.010	0.067	0.079	0.012
	(0.004)	(0.004)	(0.006)	(0.010)	(0.010)	(0.014)	(0.011)	(0.012)	(0.016)
Interest rate = 7.49%	0.021	0.023	0.002	0.039	0.034	-0.005	0.054	0.064	0.010
	(0.003)	(0.003)	(0.004)	(0.008)	(0.008)	(0.011)	(0.010)	(0.011)	(0.015)
Interest rate = 11.89%	0.015	0.016	0.000	0.030	0.038	0.007	0.057	0.029	-0.028**
	(0.003)	(0.003)	(0.004)	(0.008)	(0.008)	(0.011)	(0.010)	(0.007)	(0.012)
p-value (equal for all interest rates)	0.000	0.000		0.210	0.090		0.645	0.000	
ii. Interest rate elasticities									
Interest rate elasticity	-0.884***	-0.927***		-0.526*	-0.514*		-0.172	-0.929***	
	(0.173)	(0.170)		(0.270)	(0.293)		(0.242)	(0.205)	
p-value (elasticities are equal)	0.860			0.976			0.017		
N	6723	6581		1596	1611		1608	1571	

Notes: panel i presents the take-up rates for each interest rate cell and a p-value of an F-test that average take-up rates are equal for all interest rate cells when the interest rate is hidden and when the interest rate is emphasized for each risk category group. Panel ii reports the interest rate elasticities when the interest rate is hidden and when it is emphasized, along with the p-value of a test that these two elasticities are equal for each category group. The interest rate elasticities are calculated based on a Linear Probability Model, using equation 2. All estimates are weighted by the inverse of the probability that the client was selected so that they represent the original population. Robust standard errors in parentheses. For columns 3, 6, 9, and panel ii: * significant at 10%; ** significant at 5%; *** significant at 1%

Table 1.5: **Causal Effects of Enrolling in a Payment Plan on Probability of Default**

	Full Sample		Low-risk		Medium-risk		High-risk	
	Reduced Form	2SLS	Reduced Form	2SLS	Reduced Form	2SLS	Reduced Form	2SLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Offered i=6.39%	0.004		0.006*		0.007		-0.003	
	(0.003)		(0.003)		(0.007)		(0.007)	
Offered i=9.59%	0.006**		0.006**		-0.001		0.012*	
	(0.003)		(0.003)		(0.007)		(0.007)	
Accepted i=6.39%		0.065		0.101*		0.071		-0.025
		(0.040)		(0.057)		(0.078)		(0.058)
Accepted i=9.59%		0.114**		0.140**		-0.014		0.121*
		(0.048)		(0.071)		(0.094)		(0.067)
Control mean	0.085		0.057		0.132		0.208	
	(0.002)		(0.003)		(0.006)		(0.006)	
N	103114		57934		18930		26250	

Notes: this table presents the reduced form and 2SLS estimates of the main effect of enrolling in a payment plan on the probability of default in the following 12 months, using the larger scale experiment carried out by the firm. All estimates are weighted by the inverse of the probability that the client was selected so that they represent the original population. Robust standard errors in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%

Table 1.6: **Payment Plan Take-up Rates by Featured Plan**

		Risk categories		
	Full Sample	Low-risk	Medium-risk	High-risk
	(1)	(2)	(3)	(4)
	i. Take-up rates by featured plan			
All	0.030	0.026	0.043	0.058
	(0.001)	(0.001)	(0.004)	(0.004)
Featured plan = 6-month	0.030	0.025	0.046	0.059
	(0.002)	(0.003)	(0.007)	(0.008)
Featured plan = 8-month	0.032	0.028	0.036	0.062
	(0.003)	(0.003)	(0.007)	(0.008)
Featured plan = 10-month	0.030	0.026	0.049	0.053
	(0.002)	(0.003)	(0.008)	(0.008)
Featured plan = 12-month	0.028	0.023	0.040	0.057
	(0.002)	(0.003)	(0.007)	(0.008)
p-value (equal for all featured plans)	0.667	0.675	0.543	0.902
	ii. Elasticities with respect to the maturity of the featured plan			
Maturity elasticity	-0.130	-0.138	-0.048	-0.127
	(0.168)	(0.201)	(0.350)	(0.286)

Notes: column 1 presents take-up rates by featured plan and elasticity with respect to the maturity of the featured plan for the full sample. Elasticity is calculated based on equation 4. The p-value presented in column 1 is of an F-test that take-up rates are the same for all featured plans. Columns 2 to 4 present the same information separately for the low-, medium-, and high-risk clients. All estimates are weighted by the inverse of the probability that the client was selected so that they represent the original population. Robust standard errors in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%

Table 1.7: **Payment Plan Take-up Rates by Featured Plan**

		Featured Payment Plan			
		6	8	10	12
	6	83.47%	37.45%	41.79%	47.91%
Payment Plan	8	4.85%	50.32%	12.20%	9.04%
Choice	10	5.18%	6.20%	37.64%	8.72%
	12	6.50%	6.03%	8.37%	34.33%

Notes: this table presents the distribution of payment plan choices for each featured plan group. Sample is conditioned on clients who enrolled in a payment plan.

Table 1.8: **Payment Plan Take-up Rates by Featured Plan**

	First Stage	Reduced Form	2SLS
	(1)	(2)	(3)
Featured plan = 8	0.863***	-0.004	
	(0.172)	(0.042)	
Featured plan = 10	1.487***	0.098**	
	(0.211)	(0.047)	
Featured plan = 12	1.945***	0.048	
	(0.263)	(0.046)	
Maturity of payment plan chosen			0.040*
			(0.022)
Mean (featured plan=6)		0.160	
		(0.031)	
N		662	

Notes: this table presents the 2SLS estimates of the effects of enrolling in a longer-term payment plan on the 12-month probability of default. Sample is restricted to clients that enrolled in a payment plan. Observations are weighted by the inverse of the probability that the client was sampled. Robust standard errors in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%

Table 1.9: **When are clients More Likely to Follow the Featured Plan?**

Dependent variable:	Followed the Featured Plan							
Independent variable:	Maturity of the Featured Plan				Interest rate			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mean (featured plan = 8)		0.503 (0.041)						
Featured plan = 10	-0.127** (0.058)	-0.128** (0.059)						
Featured plan = 12	-0.160*** (0.059)	-0.155*** (0.059)						
Maturity of the featured plan			-0.041*** (0.015)	-0.039*** (0.015)				
Mean (interest rate = 3.99%)						0.564 (0.030)		
Interest rate = 7.49%					-0.081* (0.049)	-0.081 (0.050)		
Interest rate = 11.89%					-0.101* (0.054)	-0.100* (0.055)		
Interest rate							-1.341** (0.672)	-1.345* (0.688)
Include controls	No	Yes	No	Yes	No	Yes	No	Yes
N	494	494	494	494	662	662	662	662

Notes: In all regressions, sample is restricted for clients who enrolled in a payment plan. Columns 1 to 4 report the correlations between probability of following the featured plan and the maturity of the featured plan. Columns 1 and 2 reports coefficients of a LPM with dummies for featured plan equal to 10 and 12 (omitted category is featured plan equal to 8, clients with featured plan equal to 6 were excluded). In the regressions 3 and 4, the maturity of the featured plan enters linearly in the LPM. Columns 5 to 8 report the correlations between probability of following the featured plan and the interest rate. Control variables include all variables in Table 2. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 1.A.1: **Randomization Balance**

	Credit card limit	Credit card balance	Prob. of using revolving credit	Average revolved balance	Monthly interest and fees charges	Prob. of making a late payment	Time with the credit card	Medium-risk	High-risk
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
				i. Interest rate x advertisement layout					
Concealed interest rate									
3.99%	1,471.3	644.2	0.265	0.122	17.6	0.162	4.60	0.067	0.093
	(37.6)	(17.7)	(0.005)	(0.003)	(0.8)	(0.003)	(0.07)	(0.003)	(0.004)
7.49%	1,444.2	634.6	0.270	0.128	17.6	0.160	4.57	0.070	0.096
	(40.7)	(19.1)	(0.005)	(0.003)	(0.7)	(0.003)	(0.07)	(0.003)	(0.004)
11.89%	1,577.0	691.7	0.269	0.124	17.5	0.160	4.66	0.061	0.096
	(43.4)	(18.2)	(0.005)	(0.003)	(0.6)	(0.003)	(0.07)	(0.003)	(0.004)
Prominent interest rate									
3.99%	1,511.4	654.5	0.259	0.121	17.1	0.156	4.69	0.068	0.095
	(41.6)	(15.6)	(0.005)	(0.003)	(0.7)	(0.003)	(0.08)	(0.003)	(0.004)
7.49%	1,570.4	668.4	0.260	0.122	17.4	0.160	4.69	0.066	0.090
	(44.0)	(20.7)	(0.005)	(0.003)	(0.7)	(0.003)	(0.08)	(0.003)	(0.004)
11.89%	1,509.0	671.7	0.266	0.124	17.7	0.162	4.71	0.069	0.098
	(42.6)	(17.1)	(0.005)	(0.003)	(0.7)	(0.003)	(0.08)	(0.003)	(0.004)
p-value (all averages are equal)	0.160	0.274	0.622	0.598	0.990	0.840	0.731	0.240	0.792
				ii. Featured plan					
6-month	1,491.1	665.7	0.263	0.123	17.7	0.165	4.59	0.067	0.094
	(34.0)	(15.3)	(0.004)	(0.002)	(0.6)	(0.003)	(0.06)	(0.002)	(0.003)
8-month	1,543.8	667.8	0.260	0.124	17.5	0.161	4.65	0.065	0.098
	(35.6)	(16.8)	(0.004)	(0.002)	(0.6)	(0.003)	(0.06)	(0.002)	(0.003)
10-month	1,533.0	644.4	0.270	0.125	17.7	0.161	4.68	0.068	0.094
	(33.6)	(12.4)	(0.004)	(0.002)	(0.5)	(0.003)	(0.06)	(0.002)	(0.003)
12-month	1,487.9	666.0	0.267	0.122	17.0	0.154	4.68	0.067	0.092
	(32.8)	(14.4)	(0.004)	(0.002)	(0.6)	(0.003)	(0.06)	(0.002)	(0.003)
p-value (all averages are equal)	0.552	0.553	0.332	0.775	0.828	0.046	0.729	0.819	0.665
N	19690	19690	19690	19690	19690	19690	19690	19690	19690

Notes: each column presents averages of the corresponding variable for each treatment cell. Robust standard errors in parentheses. For each column, the p-value of an F-test that the mean the corresponding variable is the same for all treatment groups is presented. All estimates are weighted by the inverse of the probability that the client was selected so that they represent the original population.

Table 1.A.2: **Randomization Balance - Low-risk Clients**

	Credit card limit	Credit card balance	Prob. of using revolving credit	Average revolved balance	Monthly interest and fees charges	Prob. of making a late payment	Time with the credit card
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			i. Interest rate x advertisement layout				
Concealed interest rate							
3.99%	1,613.6	657.5	0.227	0.094	13.9	0.141	4.77
	(44.0)	(20.6)	(0.006)	(0.003)	(0.9)	(0.004)	(0.08)
7.49%	1,595.5	651.5	0.234	0.103	14.4	0.140	4.75
	(48.1)	(22.4)	(0.006)	(0.003)	(0.7)	(0.004)	(0.09)
11.89%	1,729.6	708.0	0.234	0.097	13.7	0.139	4.82
	(50.7)	(20.9)	(0.006)	(0.003)	(0.6)	(0.004)	(0.09)
Prominent interest rate							
3.99%	1,662.5	672.4	0.221	0.093	13.3	0.131	4.84
	(48.9)	(18.1)	(0.006)	(0.003)	(0.7)	(0.004)	(0.09)
7.49%	1,736.4	684.7	0.226	0.096	14.1	0.140	4.87
	(51.5)	(24.0)	(0.006)	(0.003)	(0.7)	(0.004)	(0.09)
11.89%	1,670.0	692.8	0.230	0.096	14.2	0.140	4.89
	(50.3)	(19.9)	(0.006)	(0.003)	(0.7)	(0.004)	(0.09)
p-value (all averages are equal)	0.219	0.395	0.636	0.277	0.892	0.399	0.837
			ii. Featured plan				
6-month	1,605.8	669.9	0.224	0.092	13.0	0.141	4.67
	(54.4)	(25.8)	(0.007)	(0.003)	(0.7)	(0.004)	(0.10)
8-month	1,640.4	673.7	0.220	0.097	13.8	0.143	4.69
	(56.9)	(26.8)	(0.007)	(0.003)	(1.0)	(0.004)	(0.10)
10-month	1,652.3	644.4	0.241	0.102	14.3	0.144	4.78
	(52.1)	(20.0)	(0.007)	(0.003)	(0.7)	(0.004)	(0.10)
12-month	1,690.6	703.6	0.243	0.101	15.0	0.131	4.99
	(56.9)	(25.3)	(0.007)	(0.004)	(0.9)	(0.004)	(0.11)
p-value (all averages are equal)	0.757	0.332	0.038	0.153	0.340	0.123	0.105
N	13304	13304	13304	13304	13304	13304	13304

Notes: sample conditioned on low-risk clients. Each column presents averages of the corresponding variable for each treatment cell. Robust standard errors in parentheses. For each column, the p-value of an F-test that the mean the corresponding variable is the same for all treatment groups is presented. All estimates are weighted by the inverse of the probability that the client was selected so that they represent the original population.

Table 1.A.3: **Randomization Balance - Medium-risk Clients**

	Credit card limit	Credit card balance	Prob. of using revolving credit	Average revolved balance	Monthly interest and fees charges	Prob. of making a late payment	Time with the credit card
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			i. Interest rate x advertisement layout				
Concealed interest rate							
3.99%	730.0	585.9	0.385	0.194	25.5	0.239	3.78
	(42.0)	(32.9)	(0.013)	(0.008)	(2.2)	(0.009)	(0.13)
7.49%	713.1	600.2	0.376	0.182	24.4	0.230	3.68
	(40.6)	(38.7)	(0.013)	(0.008)	(1.8)	(0.009)	(0.12)
11.89%	837.4	629.2	0.379	0.194	25.8	0.241	3.84
	(59.5)	(39.5)	(0.014)	(0.008)	(2.0)	(0.009)	(0.14)
Prominent interest rate							
3.99%	789.2	576.7	0.379	0.187	27.0	0.273	3.93
	(60.3)	(31.6)	(0.014)	(0.008)	(2.3)	(0.010)	(0.15)
7.49%	749.8	590.3	0.357	0.185	25.0	0.247	3.69
	(49.8)	(37.5)	(0.013)	(0.008)	(1.8)	(0.009)	(0.13)
11.89%	785.6	636.3	0.379	0.196	27.0	0.236	3.91
	(54.9)	(44.2)	(0.013)	(0.008)	(2.3)	(0.009)	(0.13)
p-value (all averages are equal)	0.552	0.848	0.746	0.770	0.936	0.036	0.674
			ii. Featured plan				
6-month	669.2	572.2	0.398	0.197	25.9	0.245	3.76
	(44.8)	(36.5)	(0.016)	(0.009)	(2.2)	(0.011)	(0.14)
8-month	789.5	578.1	0.377	0.190	24.2	0.229	3.62
	(54.6)	(35.0)	(0.016)	(0.010)	(2.1)	(0.010)	(0.15)
10-month	827.2	684.8	0.375	0.192	26.6	0.240	3.84
	(67.0)	(56.8)	(0.015)	(0.009)	(2.6)	(0.010)	(0.16)
12-month	745.0	581.4	0.371	0.179	24.3	0.232	3.84
	(49.7)	(38.6)	(0.015)	(0.009)	(2.4)	(0.010)	(0.15)
p-value (all averages are equal)	0.171	0.364	0.597	0.567	0.862	0.674	0.700
N	3207	3207	3207	3207	3207	3207	3207

Notes: sample conditioned on medium-risk clients. Each column presents averages of the corresponding variable for each treatment cell. Robust standard errors in parentheses. For each column, the p-value of an F-test that the mean the corresponding variable is the same for all treatment groups is presented. All estimates are weighted by the inverse of the probability that the client was selected so that they represent the original population.

Table 1.A.4: **Randomization Balance - High-risk Clients**

	Credit card limit	Credit card balance	Prob. of using revolving credit	Average revolved balance	Monthly interest and fees charges	Prob. of making a late payment	Time with the credit card
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			i. Interest rate x advertisement layout				
Concealed interest rate							
3.99%	719.3	566.1	0.523	0.325	44.5	0.296	3.64
	(51.8)	(32.8)	(0.014)	(0.010)	(3.6)	(0.010)	(0.13)
7.49%	663.2	513.8	0.505	0.310	40.7	0.289	3.63
	(47.4)	(31.3)	(0.013)	(0.010)	(2.7)	(0.009)	(0.13)
11.89%	701.4	587.7	0.516	0.319	45.8	0.296	3.78
	(45.5)	(36.3)	(0.013)	(0.010)	(3.1)	(0.009)	(0.14)
Prominent interest rate							
3.99%	703.7	553.4	0.506	0.320	43.4	0.292	3.82
	(42.7)	(29.6)	(0.013)	(0.010)	(2.8)	(0.010)	(0.15)
7.49%	616.4	572.6	0.512	0.321	42.2	0.283	3.72
	(32.1)	(36.0)	(0.014)	(0.010)	(3.0)	(0.010)	(0.14)
11.89%	650.0	517.9	0.492	0.308	41.3	0.295	3.71
	(39.9)	(23.0)	(0.014)	(0.010)	(3.2)	(0.009)	(0.13)
p-value (all averages are equal)	0.417	0.436	0.699	0.807	0.837	0.920	0.911
			ii. Featured plan				
6-month	741.2	605.6	0.504	0.313	46.9	0.290	3.59
	(65.0)	(46.3)	(0.015)	(0.011)	(4.7)	(0.010)	(0.16)
8-month	737.3	547.7	0.523	0.327	43.0	0.295	3.76
	(56.4)	(29.5)	(0.015)	(0.011)	(2.6)	(0.011)	(0.15)
10-month	662.6	537.7	0.504	0.317	43.3	0.289	3.69
	(53.3)	(37.7)	(0.016)	(0.012)	(3.3)	(0.011)	(0.16)
12-month	635.6	534.5	0.526	0.314	41.7	0.299	3.70
	(46.9)	(41.1)	(0.015)	(0.011)	(3.8)	(0.011)	(0.15)
p-value (all averages are equal)	0.408	0.640	0.615	0.781	0.855	0.892	0.896
N	3179	3179	3179	3179	3179	3179	3179

Notes: sample conditioned on high-risk clients. Each column presents averages of the corresponding variable for each treatment cell. Robust standard errors in parentheses. For each column, the p-value of an F-test that the mean the corresponding variable is the same for all treatment groups is presented. All estimates are weighted by the inverse of the probability that the client was selected so that they represent the original population.

Table 1.A.5: **Price Sensitivity of Payment Plans' Demand**

	Full Sample	Risk categories Low-risk	Medium-risk	High-risk
	(1)	(2)	(3)	(4)
Average take-up	0.029	0.025	0.041	0.059
(concealed interest rate)	(0.002)	(0.002)	(0.005)	(0.006)
Panel i: linear probability model				
Prominent interest rate	0.001	0.002	0.004	-0.002
	(0.002)	(0.003)	(0.007)	(0.008)
Interest rate	-0.266***	-0.281***	-0.279*	-0.130
	(0.054)	(0.060)	(0.153)	(0.186)
Interest rate x Prominent interest rate	-0.077	-0.032	-0.013	-0.537**
	(0.078)	(0.086)	(0.225)	(0.253)
N	19690	13304	3207	3179
Panel ii: Logit marginal effects				
Prominent interest rate	0.001	0.001	0.004	-0.006
	(0.002)	(0.003)	(0.007)	(0.008)
Interest rate	-0.267***	-0.285***	-0.287*	-0.122
	(0.054)	(0.060)	(0.160)	(0.175)
Interest rate x Prominent interest rate	-0.066	-0.014	0.016	-0.564**
	(0.078)	(0.088)	(0.224)	(0.248)
N	19690	13304	3207	3179

Notes: panel i presents coefficients of a linear probability model of take-up on interest rate, a dummy for emphasized advertisement, and the interaction of these two variables. All models include payment day fixed effects. Panel ii presents logit marginal effects. Column 1 presents estimates for the full sample, while columns 2 to 4 present estimates separately for low-, medium- . and high-risk clients. All estimates are weighted by the inverse of the probability that the client was selected so that they represent the original population. Robust standard errors in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%

Table 1.A.6: **Covariates Balance Across Featured Plans Conditional on Take-up**

	Credit card limit	Credit card balance	Prob. of using revolving credit	Average revolved balance	Monthly interest and fees charges	Prob. of making a late payment	Time with the credit card	Medium-risk	High-risk
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Featured Plan									
6-month	1,249.9	1,057.2	0.479	0.279	46.6	0.159	4.21	0.103	0.188
	(138.1)	(97.0)	(0.027)	(0.018)	(5.0)	(0.014)	(0.34)	(0.017)	(0.027)
8-month	1,690.6	1,315.0	0.480	0.283	45.0	0.152	4.70	0.073	0.189
	(175.2)	(245.8)	(0.028)	(0.019)	(4.9)	(0.014)	(0.33)	(0.014)	(0.026)
10-month	1,686.7	1,211.7	0.516	0.282	48.1	0.165	4.86	0.112	0.168
	(207.1)	(144.2)	(0.027)	(0.017)	(5.5)	(0.014)	(0.34)	(0.018)	(0.026)
12-month	1,394.3	930.0	0.499	0.280	41.5	0.183	4.48	0.097	0.190
	(171.7)	(83.9)	(0.028)	(0.019)	(4.7)	(0.015)	(0.29)	(0.018)	(0.028)
p-value (all averages are equal)	0.147	0.221	0.731	0.999	0.810	0.481	0.553	0.314	0.921
N	662	662	662	662	662	662	662	662	662

Notes: each column presents averages of the corresponding variable for each treatment cell, conditional on payment plan enrollment. Robust standard errors in parentheses. For each column, the p-value of an F-test that the mean the corresponding variable is the same for all treatment groups is presented. All estimates are weighted by the inverse of the probability that the client was selected so that they represent the original population.

Chapter 2

Understanding Peer Effects in Financial Decisions: Evidence from a Field Experiment*

2.1 Introduction

People's choices often look like the choices made by those around them: we wear what is fashionable, we "have what they're having," and we try to "keep up with the Joneses." Such *peer effects* have been analyzed across several fields of economics and social psychology.[1] An especially active area of research has examined the role of peers in financial decisions; beyond studying *whether* peers affect financial decisions, different *channels* through which peer effects work have generated their own literatures linking peer effects to investment decisions, and to financial market instability. Models of herding and asset-price bubbles,

*This chapter is co-authored with Leonardo Bursztyn, Florian Ederer, and Noam Yuchtman.

[1]In the economics literature, theoretical models of herding and social learning include Banerjee (1992) and Bikhchandani et al. (1992). Becker (1991) studies markets in which a consumer's demand for a product depends on the aggregate demand for the good. Empirical work has studied a wide range of contexts: the impact of classmates or friends on education, compensation, and other outcomes Sacerdote (2001); Duflo et al. (2011); Card and Giuliano (2011); Shue (2012); the impact of one's peers and community on social indicators Bertrand et al. (2000); Kling et al. (2007); Bobonis and Finan (2009); Dahl et al. (2012); and, the impact of coworkers on workplace performance Guryan et al. (2009); Mas and Moretti (2009). Durlauf (2004) surveys the literature on neighborhood effects. Within the social psychology literature, Asch (1951) studied individuals' conformity to group norms; Festinger (1954) posited that one resolves uncertainty by learning from others; Burnkrant and Cousineau (1975) distinguished between "informational" and "normative" reasons for conformity.

potentially based on very little information, focus on *learning* from peers' choices.[2] Models in which individuals' relative income or consumption concerns drive their choice of asset holdings, and artificially drive up some assets' prices, focus on peers' *possession* of an asset (and consequent income or consumption).[3]

Identifying the causal effect of one's peers' behavior on one's own is notoriously difficult.[4] Correlations in the investment or consumption choices of socially-related individuals might arise without any causal peer effect: for example, peers select into social groups, and this might generate correlated choices; peers share common environments (and changes in those environments), and this, too, might generate correlated choices. In the context of financial decisions, experimental work surmounting these challenges has been done by Duflo and Saez (2003) and Beshears et al. (2011), among others.

Equally difficult is identifying *why* one's consumption or investment choices have a social component.[5] Broadly, there are two reasons why a peer's act of purchasing an asset (or product, more generally) would affect one's own choice:

1. One infers that assets (or products) purchased by others are of higher quality; we refer to this as *social learning*.[6]

2. One's utility from possessing an asset (or product) depends directly on the possession of that asset (or product) by another individual; we call this *social utility*.

Suppose an investor i considers purchasing a financial asset under uncertainty. In canonical models of herding based on social learning (e.g., Banerjee (1992) and Bikhchandani et al. (1992)), information that a peer, investor j, purchased the asset will provide favorable information about the asset to investor i: investor j would only have purchased the asset if

[2]See for instance Bikhchandani and Sharma (2000) and Chari and Kehoe (2004). Social learning has also been studied experimentally in a laboratory setting by Celen and Kariv (2004).

[3]Preferences over relative consumption can arise from the (exogenous) presence of other individuals' consumption decisions in one's utility function, (e.g. Abel (1990), Gali (1994), Campbell and Cochrane (1999)) or can arise endogenously when one consumes scarce consumption goods, the prices of which depend on the incomes (and consumption and investment decisions) of other individuals (DeMarzo et al. (2004), DeMarzo et al. (2008)).

[4]An early, thorough discussion of the challenge of identifying causal peer effects is found in Manski (1993). Unsurprisingly, the empirical evidence of peer effects in financial markets has typically been correlational (e.g., Hong et al. (2005), Ivkovic and Weisbenner (2007), and Li (2009)).

[5]Banerjee et al. (2011) study the diffusion of microfinance through social networks, and structurally estimate the importance of different potential channels linking peers' decisions. Cooper and Rege (2011) attempt to distinguish among peer effect channels in the lab.

[6]Identifying social learning alone is the focus of Cai et al. (2009) and Moretti (2011); they try to rule out the importance of peer effects through other channels (e.g., joint consumption) in the contexts they study.

he observed a relatively good signal of the asset's return.[7] The favorable information conveyed by investor j's revealed preference increases the probability that investor i purchases the asset, relative to making a purchase decision without observing his peer's decision.[8]

In our study, we focus on social learning arising only from the information one acquires from the fact that one's peer purchased a financial product. We abstract away from the additional information one might acquire *after* a peer's purchase (e.g., by talking to the peer and learning about the quality of a product) and from any change in behavior due to increased salience of a product when consumed by one's peers.[9] The impact of learning from a peer's purchase decision is the social learning channel.

Typically, investor j's decision to purchase the asset will also imply that investor j possesses the asset. If investor j's possession of the asset directly affects investor i's utility of owning the same asset, then observing investor j purchasing the asset (implying investor j's possession of the asset) can increase the likelihood that investor i purchases the asset for reasons other than social learning.

In this work, we pool together individuals' concerns about relative income or consumption ("keeping up with the Joneses"), greater utility from joint consumption of a good, etc.[10] The impact of a peer's possession of an asset on an individual's utility derived from owning the same asset (for multiple reasons) is the social utility channel.

A comparison of investor i's investment when no peer effect is present to the case in which he observes investor j purchasing an asset will generally identify the combined social learning and social utility channels. To disentangle social learning from social utility, one needs to identify, or create, a context in which investor j's decision to purchase an asset is decoupled from investor j's possession of the asset (in Appendix 2.B we present a formal model of peer effects in financial decisions that features both the social learning and social

[7]We assume that investor j here made a choice in isolation, i.e., without peer effects.

[8]Avery and Zemsky (1998) present a model in which information based herds do not occur, due to price adjustments; however, in our setting there is no asset price adjustment (see also Chari and Kehoe (2004)).

[9]The asset sold in our field experiment was designed to make this abstraction possible. This is discussed in detail in Section 2.1.

[10]Note that even in the absence of truly "social" preferences, one might observe greater demand for an asset simply because a peer holds it: for example, this might arise as a result of competition over scarce consumption goods. Because we do not wish to abuse the term, "social preferences," we prefer the broader term, "social utility," which will include social preferences, as well as general equilibrium-induced differences in demand. Note also that in principle, social utility might lead to *negative* correlations between peers' choices (see Clark and Oswald (1998)), but we focus here on the case of positive social utility effects, as these are predominant in the literature on peer effects in financial decisions. Evidence consistent with individuals caring about relative outcomes has been presented by Luttmer (2005), Fliessbach et al. (2007), and Card et al. (2010), among others.

utility channels).

Disentangling the channels through which peer effects impact financial decisions is of more than academic interest. Identifying causal effects of the individual channels can provide important evidence on the sources of herding behavior in financial markets. For example, finding a causal effect of the social learning channel suggests that information provision can limit herding based on little actual information. On the other hand, if there is evidence for herding driven by factors other than social learning, improved information provision will not eliminate correlated choices among peers.

Our experimental design represents an attempt to surmount both the challenge of identifying a causal peer effect, and the challenge of separately identifying the effects of social learning and social utility. Working closely with us, a large financial brokerage in Brazil offered a new financial asset, designed exclusively for our experiment, to pairs of clients who share a social relationship (either friends or family members).[11] The stakes were high: minimum investments were R$2,000 (over $1,000 U.S. dollars at the time of the study), around 50% of the median investor's monthly income in our sample; furthermore, investors were not allowed to convert existing investments with the brokerage to purchase the asset and thus were required to allocate new resources.

To identify any sort of peer effect on investment decisions, we randomly informed one member of the peer pair (investor 2) of the investment made by the other member of the pair (investor 1).[12] To disentangle the effect of investor 1's possession from the effect of the information conveyed by investor 1's revealed preference (his decision to purchase the asset), we exploit a novel aspect of our experimental design. The financial brokerage with which we worked implemented a lottery to determine whether individuals who chose to purchase the asset would actually be allowed to make the investment (see Figure 2-1 for a graphical depiction of the experimental design).[13] Thus, half of the investor 1's who chose to purchase the asset revealed a preference for the financial asset, but *did not* possess it.

Among investor 1's who chose to purchase the asset, we implemented a second, indepen-

[11]The experimental design is discussed in more detail in Section 2. In particular, we describe the sample of investors, the characteristics of the financial asset, as well as the details of the experimental treatments.

[12]The assignment to the roles of investor 1 and investor 2 was random.

[13]Individuals understood that their decisions to purchase the asset might not be implemented. We do not believe that the investors made their decisions lightly: take-up rates in the brokerage's pilot sale of the asset (where no lottery was used) were similar to those in the control (no information about one's peer) condition in this study. Furthermore, purchase decisions that were confirmed by the lottery were actually implemented.

dent randomization to determine the information received by the associated investor 2's: we randomly assigned investor 2 to receive either *no information* about investor 1's investment decision, or to receive information about *both* the investment decision *and* the outcome of the lottery determining possession. Thus, among investor 1's who chose to purchase the asset, the associated investor 2's were randomly assigned to one of three conditions: (1) no information about investor 1's decision[14]; (2) information that investor 1 made a decision to purchase the asset, but was not able to consummate the purchase (so learning occurred *without possession*[15]); and, (3) information that investor 1 made a decision to purchase the asset, and was able to consummate the purchase (so learning occurred, along with possession).[16] A comparison of choices made by investor 2 in conditions (1) and (2) reveals the effect of social learning; a comparison of (3) and (2) reveals the impact of investor 1's possession of the asset over and above the information conveyed by his purchase[17]; a comparison of (3) and (1) reveals the total effect of these two channels. This design allows us to cleanly identify contributions of social learning and social utility in generating the overall peer effects we observe.[18]

Our experimental evidence suggests that *both* channels through which peer effects work are important. Among investor 2's whose peer chose to purchase the asset we find the following: individuals in the no information control condition chose to purchase the asset at a rate of around 42%; among those informed that investor 1 wanted the asset, but was unable to purchase it, the rate increased to 71%; finally, among those informed that investor 1 wanted the asset, and was able to purchase it, the rate increased to 93%.[19] There are large, statistically significant peer effects; in addition, we find that each channel – social utility and social learning – is individually economically and statistically significant. We

[14]We attempted to ensure that no information spread independently of the brokers' phone calls by requiring that calls be made to both investors on the same day. Only 6 out of 150 investor 2's had communicated with their associated investor 1's about the asset prior to the phone call from the brokerage; dropping these 6 observations does not affect our results.

[15]In Section 2.3.4, we discuss various other channels through which peer effects work that might be active in this condition.

[16]Investor 2 also had his decision to purchase confirmed or rejected by lottery, so he was made aware of this randomization in all of the conditions. We have no reason to expect that investor 2 viewed investor 1's decision as anything other than a revealed preference.

[17]It is important to stress that we cannot identify *why* one's peer's possession of the asset matters, but the finding that possession matters, regardless of the channel, is of interest.

[18]It is difficult to quantitatively estimate the effect of possession above learning, as the purchase rate in condition (3) is very close to 100%. Because the purchase rate is bounded above, we estimate what may roughly be thought of as a lower bound of the effect of possession beyond learning.

[19]The take-up rate in the control group is relatively high by design: the firm offered an appealing asset that was not available outside of the experiment. We discuss the the asset in detail in Appendix 2.C.

find that individuals learn from their peers, but also that there is an effect of possession beyond learning. This is true not only for the purchase decision, but also for the decision of whether to invest more than the minimum investment amount, and how much to invest in the asset.

Our design allows us to examine another aspect of peer effects: the role played by selection into a socially-related pair in generating correlated choices among peers. In particular, when investor 1 chose *not* to purchase the asset, his associated investor 2 was assigned to a "negative selection" condition, in which no information was received about about the peer.[20] We refer to these investor 2's as the "negative selection" condition, because the information they receive is identical to that received by investor 2's in the control condition (1); however, the investor 2's in the "negative selection" condition are those whose peers specifically chose *not* to purchase the asset. Interestingly, we do not find large selection effects: the take-up rate in the "negative selection group was 39%, not very different from the take-up rate in condition (1).

If investors exhibit different levels of financial sophistication, a natural extension of our analysis of social learning is to explore heterogeneity in that dimension. One would expect that an investor with a high degree of financial sophistication is less likely to follow the revealed preference decisions of other investors, because he will rely more on his own assessment of the quality of an asset. On the other hand, an investor with a low degree of financial sophistication will be more likely to be influenced by the revealed preference decisions of other investors. Analogously, financially sophisticated investors may be more likely than unsophisticated investors to influence their peers through the social learning channel, given their greater knowledge of financial investments.

Using occupational categories to indicate financial sophistication[21] ("sophisticated" investors have occupations in engineering, finance, accounting, etc.), we find that there is *no* significant effect of social learning among sophisticated investor 2's, and a large, significant social learning effect among unsophisticated investor 2's. The difference in the treatment effects across these groups is statistically significant as well. We find weaker evidence suggesting that investor 2's exhibit larger social learning effects when their associated investor 1

[20]We did not reveal their peers' choices because the brokerage did not want to include experimental conditions in which individuals learned that their peer *did not* want the asset.

[21]Goetzmann and Kumar (2008), Calvet et al. (2009), and Abreu and Mendes (2010) find that investors' occupations are correlated with measures of their financial sophistication.

is sophisticated than when he is unsophisticated. These results suggest that learning about the asset plays the predominant role in the social learning effects we observe, and that competing stories do not seem to be driving our findings in the social learning condition.[22]

A final important concern with our design, common to many experimental studies, regards the external validity of the findings. In the discussion of our empirical results, in Section 2.3.4, we present evidence suggesting that our results do capture relevant effects, albeit not representative of all peer effects in financial decisions. Importantly, we show: (i) the set of investors in our study who chose to purchase the asset are similar to those who do not want to make the purchase – this is important because our treatment effects of interest condition on investor 1's wanting to purchase the asset; (ii) the characteristics of our sample of investors – selected because they shared a social connection with another client – are roughly similar to those of the full set of clients of the main office of the brokerage with which we worked; (iii) sales calls similar to those used in the study are common, accounting for a large fraction of the brokerage's sales – this indicates the importance of communication from brokers to clients in financial decision making.[23]

The paper proceeds as follows: in Section 2, we describe in detail our experimental design, which attempts to separately identify the channels through which peer effects work; in Section 3, we present our empirical specification and the results of our experiment, and discuss our findings; finally, in Section 4, we offer concluding thoughts.

2.2 Experimental Design

The primary goal of our design was to decouple the decision to purchase the asset from possession of the asset. To this end, we generated experimental conditions in which individuals would make decisions: 1) uninformed about any choices made by their peer; 2) informed of their peer's revealed preference to purchase an asset, but the (randomly determined) inability of the peer to make the investment; and, 3) informed of their peer's revealed preference to purchase an asset, and the peer's (randomly determined) successful investment. We now describe the design in detail; in particular, the structuring of a financial asset that possessed particular characteristics, and the implementation of multiple stages of randomization in

[22]We assess some alternative interpretations of our findings and some potential confounding factors in the discussion of our empirical results in Section 2.3.4.

[23]While brokers generally do not provide information about specific clients' purchases, discussions with the brokerage indicate that brokers do discuss the behavior of other investors in making their sales calls.

the process of selling the asset.

2.2.1 Designing the Asset

The asset being offered needed to satisfy several requirements. Most fundamentally, there needed to be a possibility of learning from one's peers' decisions; and, the asset needed to be sufficiently desirable that *some* individuals would choose to purchase it, even in the absence of peer effects.[24] To satisfy these requirements, the brokerage created a desirable, new, risky asset specifically for this study. The asset is a combination of an actively-managed mutual fund and a real estate note (the asset is described in detail in Appendix 2.C). The brokerage piloted the sale of the asset (without using a lottery to determine possession), to clients other than those in the current study, in order to ensure a purchase rate of around 50%.

Another requirement was that there be no secondary market for the asset, for several reasons. First, we hoped to identify the impact of learning from peers' decisions to purchase the asset, rather than learning from peers based on their experience possessing the asset. Investor 2 may have chosen not to purchase the asset immediately, in order to talk with investor 1, then purchase the asset from another investor. We wished to rule out this possibility. In addition, we did not want peer pairs to jointly make decisions about selling the asset. Finally, we did not want investor 2 to purchase the asset in hopes of selling it to investor 1 when investor 1's investment choice was not implemented by the lottery. In response to these concerns, the brokerage offered the asset only at the time of their initial phone call to the client – there was a single opportunity to invest – and structured the asset as having a fixed term with no resale – once the investment decision was made, the investor would simply wait until the asset matured and then collect his returns.

A final requirement, given our desire to decouple the purchase decision from possession, was that there must be limited entry into the fund to justify the lottery to implement purchase decisions. The brokerage was willing to implement the lottery design required, justified by the supply constraint for the asset they created for the study.[25]

[24]Many of our comparisons of interest are among those investor 2's whose associated investor 1's chose to purchase the asset, and investor 1 never receives any information about his peer (i.e., about investor 2).

[25]There was a maximum allowable investment of R$10,000 per client on the real estate note component of the investment; one investor invested the maximum amount.

2.2.2 Selling the Asset

To implement the study, we designed (in consultation with the financial brokerage) a script for sales calls that incorporated the randomization necessary for our experimental design (the translated script is available in Appendix 2.D).[26] The sales calls made by brokers possessed several important characteristics. First, and most importantly, the calls were extremely natural: sales calls had frequently been made by the brokerage in the past; investments resulting from brokers' calls are thus in no way unusual.[27] We believe that no client suspected that the calls were being made as part of an experiment.[28]

Second, the experimental calls were made by the individual brokers who were accustomed to working with the clients they called as part of the study; and, the calls only deviated from brokers' typical sales calls as required to implement the experimental design. Thus, clients would have trusted the broker's claims about their peer's choices, and would have believed that the lottery would be implemented as promised. Third, because brokers were compensated based on the assets they sold, they were simply incentivized to sell the asset in each condition (rather than to confirm any particular hypothesis).[29]

Between January 26, and April 3, 2012, brokers called 150 pairs of clients whom the brokerage had previously identified as having a social connection (48% are members of the same family, and 52% are friends; see Table 2-1).[30] It is important to note that, although

[26]We created the script using Qualtrics, a web-based platform that brokers would access and use to structure their call; although brokers were not able to use the web-based script in all of their calls, two-thirds of the calls used the Qualtrics script (it was occasionally abandoned when the website was not accessible). The brokers were made very familiar with the script, however, and used Excel to generate the randomization needed to execute the experimental design when unable to use Qualtrics. The treatment effects are not significantly different if we restrict ourselves to the pairs whose calls followed the Qualtrics script (results available upon request). The brokers entered the results of the randomization and the purchase decisions in an Excel spreadsheet, which they then delivered to the authors.

[27]The brokerage communicated to us that approximately 70% of their sales were achieved through such sales calls.

[28]Thus, our study falls into the category "natural field experiment", according to the classification in Harrison and List (2004).

[29]Thus, brokers would have used the available information in each experimental treatment as effectively as possible. Any treatment effects measured can be thought of as the effects of information about a peer's choice (or choice plus possession) when that information is "optimally" used by a salesperson. Of course, in reality, information about peers' choices may be received from the peer (rather than from a salesperson), or may not be observed at all; the magnitudes of our effects should be interpreted with this in mind.

[30]In Table 2-1, we generally present means of the various investor characteristics. The exception is the earnings variable, the median of which is shown to mitigate the influence of outliers: while the median monthly income in our sample is R$4,500, one investor had monthly earnings of R$200,000. In addition to the brokerage's record of a pair's social relationship, investor 2 was directly asked about the nature of his relationship with investor 1, after investor 2 had decided whether to purchase the asset (investor 1 was not asked about investor 2 at all). Finally, note that one of the authors (Bursztyn) was present for some of the phone calls. In addition, for those and many of the other calls, we had a research assistant present. We discuss the implications of broker monitoring in Section 2.3.4.

the investors in this sample are not a random sample of the brokerage's clients, we find that their observable characteristics are roughly similar to the full set of clients of the brokerage's main office (see Table 2-1, columns 1 and 8).

Information on these clients' social relationships was available for reasons independent of the experiment: the firm had made note of referrals made by clients in the past.[31] In the context of our experiment, this is particularly important because clients' social relationships would not have been salient to those whose sales call did not include any mention of their peer. If the second member of the peer pair had thought about his peer's potential offer and purchase of the same asset, our measured effects would be attenuated. However, data from the pilot calls made by the brokers prior to the experiment suggest that these potential peer effects in the control condition are not of great importance: the brokerage called clients who were *not* socially connected to other clients of the brokerage, and their purchase rate of the asset was very similar to that of our control condition. We thus believe that without any mention of the offer being made to the other member of the peer pair, there will be no peer effect, though of course this may not be exactly true in reality.[32]

One member of the pair was randomly assigned to the role of "investor 1," while the second member was assigned to the role of "investor 2."[33] Investor 1 was called by the brokerage and given the opportunity to invest in the asset without any mention of their peer. The calls proceeded as follows. The asset was first described in detail to investor 1. After describing the investment strategy underlying the asset, the investor was told that the asset was in limited supply; in order to be fair to the brokerage's clients, any purchase decision would be confirmed or rejected by computerized lottery.[34] If the investor chose to purchase the asset, he was asked to specify a purchase amount. Then, a computer would generate a random number from 1 to 100 (during the phone call), and if the number was greater than 50, the investment would be authorized; the investor was informed of

[31]Some clients were known to have links to more than one other client. We only called a single investor 1 and a single investor 2 from these "networks" containing more than one individual to ensure that investors did not learn about the asset outside of the brokers' sales calls.

[32]We also asked the brokerage if any client mentioned their peer in the sales call, and the brokerage indicated that this never occurred.

[33]A comparison of the characteristics of investor 1's and investor 2's can be seen in Appendix table 2.A.1, columns 1 and 2. The randomization resulted in a reasonable degree of balance across groups: 5 of 6 tests of equality of mean characteristics across groups have p-values above 0.10. One characteristic, gender, is significantly different across groups.

[34]This is not as unusual as it may appear at first glance: for example, Instefjord et al. (2007) describe the use of lotteries to allocate shares when IPO's are oversubscribed.

the details of the lottery before making his purchase decision.[35] One might naturally be concerned that knowledge of the lottery would affect the decision to invest. This would, of course, be of greatest concern to us if any effect of the lottery interacted with treatment status. It is reassuring to know, however, that in the brokerage's initial calls to calibrate the asset's purchase rate, which did not mention the lottery, the purchase rate was 12 of 25, or 48% – very similar to what we observe among investors in our study receiving no information about their peers.

Following the call to investor 1, the brokerage called the associated investor 2. The brokers were told that, for each pair, both investors had to be contacted *on the same day* to avoid any communication about the asset that might contaminate the experimental design.[36] If the broker did not succeed in reaching investor 2 on the same day as the associated investor 1, the broker would not attempt to contact him again; this outcome occurred for 12 investor 1's, who are not included in our empirical analysis.[37] When the broker reached investor 2, he began the script just as he did for investor 1: describing the asset, including the lottery to determine whether a purchase decision would be implemented. Next, the broker implemented the experimental randomization and attempted to sell the asset under the experimentally-prescribed conditions (described next). If investor 2 chose to purchase the asset, a random number was generated to determine whether the purchase decision would be implemented, just as was the case for investor 1.

2.2.3 Randomization into Experimental Conditions

The experimental conditions were determined as follows. Among the group of investor 1's who chose to purchase the asset, their associated investor 2's were randomly assigned to receive information about investor 1's choice and the lottery outcome, or to receive no information. There was thus a "double randomization" – first, the lottery determining

[35]Among investor 1's who wanted to purchase the asset, a comparison of the characteristics of investor 1's whose purchase decision was authorized and investor 1's whose purchase decision was not authorized can be seen in Appendix table 2.A.1.The randomization resulted in a reasonable degree of balance across groups: 6 of 7 tests of equality of mean characteristics across groups have p-values above 0.10. One characteristic, gender, is significantly different across groups.

[36]This restriction also limited the ability of investor pairs to coordinate their behavior (for example, organizing side payments or transfers of the asset between friends). As noted above, 6 investor 2's had communicated with their associated investor 1's about the asset prior to the phone call from the brokerage, but dropping these 6 observations does not affect our results.

[37]Thus, to be clear, brokers called 162 investor 1's in order to attain our sample size of 150 pairs successfully reached.

whether investor 1 was able to make the investment, and second, the randomization determining whether investor 2 would be informed about investor 1's investment choice and the outcome of the first lottery.

This process assigns investor 2's whose associated investor 1's chose to purchase the asset into one of three conditions (Figure 2-1 presents a graphical depiction of the randomization); investor characteristics across the three experimental conditions can be seen in Table 2-2. One-third were assigned to the "no information," control condition, condition (1). Half of these come from the pool of investor 2's paired with investor 1's who wanted the asset but were not authorized to make the investment, and half from those paired with investor 1's who wanted to make the investment and were authorized to make it. Investor 2's in condition (1) were offered the asset just as was investor 1, with no mention of an offer made to their peer.[38]

Two-thirds received information about their peer's decision to purchase the asset, as well as the outcome of the lottery that determined whether the peer was allowed to invest in the asset.[39] The randomization resulted in approximately one-third of investor 2's in condition (2), in which they were told that their peer purchased the asset, but had that choice rejected by the lottery. The final third of investor 2's were in condition (3), in which they were told that their peer purchased the asset, and had that choice implemented by the lottery.

The three conditions of investor 2's whose associated investor 1's wanted to purchase the asset are the focus of our analysis. Importantly, the investor 1's who chose to purchase the asset were not an unusual subset of the clients in the study. When comparing investor 1's who chose to purchase the asset to those who chose not to purchase it, the means of observables are very similar, as can be seen in Table 2-1, columns 3 and 4. This suggests that the selection of investor 1's who invest are not very different from the original pool of clients who were reached.[40]

Given the double randomization in our experimental design, investor 2's in conditions

[38]We can think of these investor 2's as "positively selected" relative to the set of investor 1's, as the latter were a random sample of investors, while the former are specifically those whose peer chose to purchase the asset.

[39]Investor 2's were not informed of investor 1's desired investment size, only whether investor 1 wished to purchase the asset.

[40]In addition, Table 2-1, columns 6 and 7, also suggests that investor 2's associated with investor 1's who chose to purchase the asset are similar to investor 2's associated with investor 1's who chose not to purchase the asset.

(1), (2), and (3) should have similar observable characteristics, and would differ only in the information they received. As a check of the randomization, we present in Table 2-2 the individual investors' characteristics for each of the three groups, as well as tests of equality of the characteristics across groups. As expected from the random assignment into each group, the sample is well balanced across the baseline variables.

Along with the three conditions of interest, in some analyses we will consider those investor 2's whose associated investor 1 chose not to invest in the asset (the characteristics of these investor 2's can be seen in Table 2-1, column 7). We assign these investor 2's to their own "negative selection" condition, in which they receive no information about their peer. We did not reveal their peers' choices because the brokerage did not want to include experimental conditions in which individuals learned that their peer *did not* want the asset. These individuals were offered the asset in exactly the same manner as were investor 1's and investor 2's in condition (1). We refer to these investor 2's as those in the "negative selection" condition, because the information they receive is identical to that received by investor 2's in condition (1); however, the investor 2's in the "negative selection" condition are those whose peers specifically chose *not* to purchase the asset.

2.2.4 Treatment Effects of Interest

Our experimental design allows us to make several interesting comparisons across groups of investors. First, we can estimate overall peer effects, working through both social learning and social utility channels. Consider the set of investor 2's whose peers had chosen to purchase the asset (whether the investment was implemented or not). Among these investor 2's, a comparison of those in conditions (1) and (3) reveal the standard peer effect: in condition (1), there is no peer effect active, as no mention was made of any offer being made to the other member of the peer pair; in condition (3), the investor 2 is told that investor 1 successfully invested in the asset (so both social utility and social learning channels are active).

Second, we can disentangle the channels through which peers' purchases affect investment decisions. Comparing investor 2's (again those whose peer chose to purchase the asset) in conditions (1) and (2) will allow us to estimate the impact of social learning resulting from a peer's decision *but without possession.* Comparing investor 2's purchase decisions in conditions (2) and (3) will then allow us to estimate the impact of a peer's *possession* *alone*, over and above learning from a peer's decision, on the decision to invest.[41]

In addition to identifying these peer effects, we will examine the role of "selection", by comparing investor 2's in condition (1) to those in the "negative selection" condition; we will also test for heterogeneous effects of social learning according to the financial sophistication (proxied by occupation) of investor 2's and investor 1's.

2.3 Empirical Analysis

Before more formally estimating the effects of the experimental treatments, we present the take-up rates in the raw data across categories of investor 2's (see Figure 2-2). Among those investor 2's in the "no information" condition (1), the take up rate was around 42%; in the "social learning alone" condition (2), the take-up rate was around 70%; in the "social learning plus social utility" condition (3), the take-up rate was just over 90%. These differences represent economically and statistically significant peer effects: the difference of around 50 percentage points in take-up rates between conditions (1) and (3) is large (and statistically significant at 1%), indicating the relevance of peer effects in financial decisions; moreover, we observe sizable and statistically significant effects of learning alone, and of possession above learning as well.[42] Interestingly, we do not see a big "selection" effect: there was a 42% take-up rate in condition (1) and 39% in the "negative selection" group (the difference is not statistically significant).

2.3.1 Regression Specification

To identify the effects of our experimental treatments, we will estimate regression models of the following form:

$$Y_i = \alpha + \sum_c \beta_c I_{c,i} + \gamma' \mathbf{X}_i + \epsilon_i. \tag{2.1}$$

[41]It is important to emphasize that our estimated effect of possession is conditional on investor 2 having learned about the asset from the revealed preference of investor 1 to purchase the asset. One might imagine that the effect of possession of the asset by investor 1 *without* any revealed preference to purchase the asset could be different. It is also important to point out that the estimated effect of "possession" is difficult to interpret quantitatively: the measured effect is bounded above by 1 minus the take-up in Condition 2, working against finding any statistically significant peer effects beyond social learning.

[42]The p-value from a test of equality between take-up rates in conditions (1) and (3) – the overall peer effect – is 0.000. The p-value from a test that (2) equals (1) – social learning alone – is 0.040. The p-value from a test that (3) equals (2) – possession's effect above social learning – is 0.047.

Y_i is an investment decision made by investor i: in much of our analysis it will be a dummy variable indicating whether investor i wanted to purchase the asset, but we will also consider the quantity invested, as well as an indicator that the investment amount was greater than the minimum required. The variables $I_{c,i}$ are indicators for investor i being in category c, where c indicates the experimental condition to which investor i was assigned. In all of our regressions, the omitted category of investors to which the others are compared will be investor 2's in condition (1); that is, those investor 2's associated with a peer who wanted to purchase the asset, but who received no information about their peer. In most of our analyses, we will focus on investor 2's, so $c \in \{\text{condition (2)}, \text{condition (3)}, \text{"negative selection"}\}$. In some cases, we will also include investor 1's in our analysis, and they will be assigned their own category c. Finally, in some specifications we will include control variables: X_i is a vector that includes broker fixed effects and investors' baseline characteristics.

2.3.2 Empirical Estimates of Social Utility and Social Learning

We first present the estimated treatment effects of interest using an indicator of the investor's purchase decision as the outcome variable, and various specifications, in Table 2-3. We begin by estimating a model using only investor 2's and not including any controls, in Table 2-3, column 1. These results match the raw data presented in Figure 2-2: treatment effects are estimated relative to the omitted category, investor 2's in condition (1), who had a take-up rate of 42%. As can be seen in the table, the overall peer effect – the coefficient on the indicator for condition (3) – is over 50 percentage points, and is highly significant.

In addition, social learning and social utility are individually significant. Investor 2's in condition (2) purchased the asset at a rate nearly 29 percentage points higher than those in condition (1), and the difference is significant. This indicates that learning without possession affects the investment decision. The difference between the coefficient on condition (3) and the coefficient on condition (2) indicates the importance of possession *beyond* social learning. Indeed, as can be seen in Table 2-3, the 22 percentage point difference between these conditions is statistically significant.[43]

Finally, the coefficient on the "negative selection" variable, in column 1 of Table 2-3,

[43]The possession effect could also capture a demand for joint insurance: closely-related peers might want to diversify the risk in their investments. Joint insurance would imply a reduction of the measured effects of possession above learning and thus attenuate our findings relating to that channel. Interestingly, we find that friends choose assets similar to their peers', rather than trying to diversify their investments.

gives us the difference in take-up rates between investor 2's whose peers did not want to purchase the asset and investor 2's whose peers wanted to purchase the asset, with neither group receiving information about their peers. As suggested by Figure 2-2, the estimated selection effect is economically small, and it is not statistically significant.

We next present regression results including broker fixed effects (Table 2-3, column 2) and including both broker fixed effects and baseline covariates (Table 2-3, column 3); then, we estimate a regression including these controls and using the combined sample of investor 1's and investor 2's in order to have more precision (Table 2-3, column 4). The overall peer effect, as well as the individual social learning and social utility channels, estimated using these alternative regression models are very similar across specifications, providing evidence that our randomization across conditions was successful.

We also consider two alternative outcome variables: the amount invested in the asset, and a dummy variable indicating whether the investment amount was greater than the minimum required. In Tables 2-4 and 2-5, we replicate the specifications presented in Table 2-3, but use the alternative outcome variables. As can be seen in Tables 2-4 and 2-5, the results using these alternative outcomes closely parallel those using take-up rates. We observe significant peer effects, significant social learning, and significant effects of possession beyond learning.[44]

The results across specifications in Table 2-3 through 2-5 indicate sizable peer effects in financial decision making. Moreover, they suggest that both channels through which peer effects work are important. Though our context is not perfectly general, the results lend support both to models of peer effects emphasizing learning from others as well as to those emphasizing keeping up with the Joneses and other "social utility" channels. We discuss in more detail the external validity of the magnitudes of our coefficients in Section 2.3.4.

2.3.3 Treatment Effect Heterogeneity

While we estimate sizable average treatment effects of social learning, as discussed above, the importance of the social learning channel is likely to depend on investors' financial sophistication. Financially sophisticated investors should put more weight on their own signal of the quality of the asset relative to information derived from their peers' revealed

[44]The only difference worth noting is that we observe marginally significant differences between investor 1's and investor 2's in condition (1) using these two outcomes.

preferences. Therefore, the learning channel should be less important for more financially sophisticated investor 2's. Similarly, the information revealed by the action of one's peer should be more influential if this peer is more financially sophisticated, and is thus likely to have received a more precise signal of the asset's quality (see Appendix 2.B for a formal treatment of this argument). Exploiting this source of heterogeneity is not only interesting from a theoretical standpoint – since it is a natural extension of a social learning framework – but it also helps rule out other alternative explanations driving our measured social learning effects (as we discuss in Section 2.3.4).

Although we do not observe direct measures of financial sophistication for the investors in our sample, we can use information on their occupations as a proxy.[45] Investors who have technical occupations (for example, professions related to engineering, finance, accounting, etc.) are likely to be more financially sophisticated than those who do not (see the list of occupations in our sample and their coding in Appendix table 2.A.2). Therefore, we examine the heterogeneity in the social learning effect according to whether investor 2 or investor 1 has a technical occupation.[46]

To identify the social learning effect for different groups of investors, we compare the take-up rates of investor 2's with differing degrees of financial sophistication in conditions (1) and (2).[47] In Figure 2-3.1, we present raw take-up rates in conditions (1) and (2) for investor 2's who have, and who do not have, technical occupations, respectively. One sees no significant effect of social learning among investor 2's with technical occupations; on the other hand, the impact of social learning is very large among those without technical occupations. In Figure Figure 2-3.2, we present raw take-up rates in conditions (1) and (2) for investor 2's whose associated investor 1's have, and who do not have, technical occupations, respectively. One sees moderate social learning effects among both groups of investor 2's. The effect appears larger for those learning from investor 1's with technical occupations, but the cell sizes are small, and confidence intervals are correspondingly large.

We present regression estimates in Table 2-6. Panel A presents comparisons of means for

[45] Goetzmann and Kumar (2008), Calvet et al. (2009), and Abreu and Mendes (2010) find that investors' occupations are correlated with measures of their financial sophistication.

[46] Since our experimental design allows us to quantify the importance of the social learning channel, but it only allows us to test for a qualitative effect of possession over and above learning (because our estimates of the latter are more likely to be affected by the upper bound of 100% take-up), we focus on the social learning channel when contrasting the magnitude of peer effects for the different groups of investors.

[47] We drop observations with "undetermined occupations", i.e., occupations that cannot be coded as technical or non-technical, such as broad categories like "teacher" and "retired."

investors who have, and who do not have, technical occupations (columns 1-2); and, whose associated investor 1's have, and do not have, technical occupations (columns 4-5). Panel B, columns 1-2, reports effects estimated using the full specification of Table 2-3, column 4, but adding an interaction of the condition (2) dummy with a dummy indicating that the investor has a technical occupation, and also an interaction of the condition (2) dummy with with a dummy indicating that the investor does not have a technical occupation. Panel B, columns 3-4, reports an analogous specification, but including an interaction of the condition (2) dummy with a dummy indicating that the associated investor 1 has a technical occupation, and also an interaction of a condition (2) dummy with a dummy indicating that the associated investor 1 does not have a technical occupation.[48]

In columns 1 and 2 of Table 2-6, one sees that the social learning effect is large and statistically significant when investor 2 does not have a technical occupation, while it is not statistically significant when investor 2 has a technical occupation. While sophisticated and unsophisticated investor 2's have similar take-up rates in the control condition (1), we can reject at 5% or better that the social learning effect is the same for the two groups (see Table 2-6, column 3).[49] In Appendix tables 2.A.3 and 2.A.4, we show that social learning's effects on investment amounts, and on an indicator that the investment amount was greater than the minimum, are also significantly larger when investor 2 does not have a technical occupation.

In Table 2-6, columns 4-6, one sees that social learning from investor 1's with technical occupations is not significantly greater than social learning from investor 1's with non-technical occupations. The point estimate of the effect of social learning is larger when learning is from an investor 1 who had a technical occupation, but this difference should be interpreted cautiously.[50]

[48]Note that we do not combine into a single analysis the study of social learning *by* sophisticated and unsophisticated investors with the study of social learning *from* sophisticated and unsophisticated investors. One might expect that the greatest effect of social learning would occur when unsophisticated investors learn from sophisticated ones; however, our sample size prevents us from running this sort of test. Finally, it is important to note that sophisticated investors (both investor 1's and 2's) are somewhat more likely to have sophisticated peers, but this positive assortative matching is not statistically significant. Data are available from the authors upon request.

[49]We also find that the social utility effect beyond learning is positive for investor 2's who have a technical occupation, and zero for those who do not. However, this difference is most likely driven by the fact that the take-up rate for the latter group is already close to one with only the learning channel active, leaving little space for possession beyond learning to have a measurable effect. These results are available upon request.

[50]In Appendix tables 2.A.3 and 2.A.4, we show that social learning's effects on investment amounts, and on an indicator that the investment amount was greater than the minimum, are not clearly related to the occupation of the associated investor 1.

2.3.4 Discussion

Our results thus far present evidence strongly suggesting an important role played by social learning alone and by possession above learning in driving individuals' financial decisions. However, experiments in the field are typically imperfect, and ours is no exception. Thus, we now discuss, in turn, alternative hypotheses (or confounding factors) that might contaminate our experimental treatments; whether supply-side (broker) behavior played an important role in generating the observed treatment effects; and, the external validity of our findings.

Alternative Hypotheses and Confounding Factors

One important concern in interpreting our results is that factors other than learning from investor 1's revealed preference are present in the social learning condition (condition (2)). Investor 2's might have had different behavioral responses to their peer's loss of the lottery. For instance, they might have felt guilt purchasing an asset that their peer wanted but could not purchase; alternatively, they might have wanted the asset even more, to "get ahead of the Joneses." These two behavioral responses would have opposite implications in terms of changes in the magnitudes of the treatment effects observed in the social learning condition. In evaluating this concern, we are reassured by our results (in Table 2-6) showing heterogeneous effects of social learning depending on financial sophistication (proxied by occupation). They suggest that the alternative stories do not drive our findings, unless guilt or a desire to get "ahead of the Joneses" is systematically stronger among individuals with more- or less-technical occupations.

Another potential concern within the social learning condition relates to side payments. Since investor 1's wanted, but could not get, the asset, some investor 2's may have felt tempted to make the investment and pass it along, or sell it, to their peer. Our design reduces the impact of this type of concern for several reasons. First, investor 1's who lost the lottery did not know that investor 2's would receive the offer, so investor 1's would likely not have initiated this strategy following their sales call. Second, even had they suspected that their friend would receive the offer, there was limited time between calls to investor 1's and investor 2's – indeed, only 6 out of 150 investor 2's reported that they heard about the asset from their associated investor 1. Finally, once investor 2 received the offer, he was unable to communicate with investor 1 prior to making his investment decision in order to facilitate coordination.

We can also address this concern to some extent with our experimental data. One might expect side payments, if present and important, to be most prominent among peers who are family members, as family members would have an easier time coordinating such payments than would mere friends or coworkers. This would tend to drive up the estimated social learning effect for pairs who are family members. In column 1 of Table 2-7, we consider the full specification from column 4 of Table 2-3, but also including interactions of each of our treatment (i.e., condition c) dummy variables with a dummy variable equal to one if investor 1 and investor 2 are family members. The results suggest that the treatment effects from social learning are not stronger among family members – the point estimate of the interaction is almost exactly zero.

One might also think that knowing that a peer desired to purchase an asset (even if he was unable to make the purchase) provides an indication of that peer's portfolio (or future asset purchases outside of the study). As a result, the social learning condition could potentially also contain some (anticipated or approximate) possession effect. This inference, however, is a sophisticated one, and we would expect financially sophisticated clients to be more likely to make it. However, as mentioned above, the effects of social learning are actually very *small* for clients with technical occupations, arguing against this hypothesis.

There are also some concerns that could affect both the "learning alone" and the "learning plus possession" conditions. When individuals observe their peers desiring to purchase an asset, they might update their priors about the existing demand for the asset and thus about future asset prices. However, in our experiment, the asset is only sold during the broker's call, and resale is not possible; thus, this concern does not seem to be relevant. Another potential issue relates to trust in the information provided by the brokers during the phone calls. However, we have no reason to think that clients would mistrust brokers with whom they have had an ongoing relationship (especially regarding easily verifiable claims).

Changes in Supply Side Behavior

Our study manipulates information received by agents on the demand side of a financial market. Of course, there is a supply side in this market as well, and one could be concerned that supply-side factors interact with our measured treatment effects. Brokers could exert differential effort toward selling the asset under different experimental conditions; they could try to strategically sort subjects across conditions, overturning our randomization; more generally, the experiment was not double blind, which might affect the implementation of the design.

Fortunately, we believe that the impact of these various concerns was likely small, for several reasons. First, as mentioned above, because brokers were compensated based on the assets they sold, they were incentivized to sell the asset in all conditions (rather than to confirm any particular hypothesis). Brokers would have used the available information in each experimental treatment as effectively as possible.

As a more direct check of supply-side effects, we examine the impact of broker experience on the treatment effects we estimate. Within the experiment, we view broker experience as a measure of broker knowledge of our study, and hence, the "double-blindness" of a phone call. We thus estimate the full specification of column 4 from Table 2-3, but including an interaction of the treatment dummies with a measure of broker experience: for each date of the study, we calculate the number of calls each broker had made before that date. The estimates of these interactions, presented in column 2 of Table 2-7, show that broker experience does not significantly affect the estimated treatment effects.

Brokers adjusting their effort across conditions or trying to overturn our randomization would also likely be correlated with broker experience within the study: both of these would be more profitably executed with some knowledge of investors' responses to the information provided in different treatments.[51] The lack of sizable effects from the interactions of broker experience with the treatment dummies suggests that these concerns do not drive our findings.

Finally, our research assistant randomly visited the brokerage on 6 out of the 12 dates on which sales calls were made, and monitored the brokers to check that they were following the script. In column 3 of Table 2-7, we present the results of estimating estimate the full specification of column 4 from Table 2-3, but including an interaction of the treatment dummies with a dummy variable indicating whether the research assistant was present at the brokerage. The results indicate that research assistant presence does not significantly affect

[51]For example, if brokers knew that investors were more likely to purchase the asset in the "learning and possession" condition (3), they may have been willing to exert more costly effort to make a sale in that condition. If certain types of investor 2's were more responsive to the information provided in a particular experimental condition, brokers might have tried to shift them into the relevant condition.

our treatment effects; this provides further evidence that the experiment was implemented as designed, even when brokers were not closely monitored.

External Validity

A final important concern with our design regards the external validity of the findings. There are several reasons to question the generality of the treatment effects we estimate. First, our comparisons among investor 2's in conditions (1), (2), and (3) are conditional on investor 1 wishing to purchase the asset. If this was an unusual sample of investor 1's, perhaps the associated investor 2's were unusual as well, and thus reveal peer effects that cannot be viewed as representative even within our experimental sample.

In fact, when comparing investor 1's who chose to purchase the asset to those who chose not to purchase it, one sees that their observable characteristics are very similar (see Table 2-1, columns 3 and 4). The investor 2's associated with investor 1's who chose to purchase the asset are also similar to those associated with investor 1's who chose not to purchase it (see Table 2-1, columns 6 and 7).[52] We also find that in the "no information" condition, investor 2's associated with investor 1's who chose to purchase the asset have a similar purchase rate to investor 2's associated with investor 1's who chose not to purchase the asset (Table 2-3). This suggests that conditioning on investor 1's wanting to purchase the asset does not produce an unusual subsample from which we estimate treatment effects.

A second question is how different our sample of investors is from other clients of the brokerage – perhaps individuals who had referred (or had been referred by) other clients in the past (and who were thus selected into our study) are a highly atypical sample. In Table 2-1, column 8, we present characteristics of the full set of the brokerage's clients from the firm's main office.[53] Although the clients in our study's sample are not a random sample of the brokerage's clients, we find that their observable characteristics are roughly similar to the full set of clients of the main office.

One may question the representativeness of the form of communication studied in our experiment. Certainly, peers often communicate among themselves, rather than being informed about each other's activity by a broker trying to make a sale. While our study

[52]This is a comparison of investor 2's in conditions (1), (2), and (3) to investor 2's in the "negative selection" condition.

[53]The majority of the individuals in our experimental sample were selected from this office's clients, though some were selected from other offices.

focuses on only one type of communication generating peer effects, we believe it is important: as noted above, approximately 70% of the brokerage's sales were derived from sales calls, so information conveyed through such calls is extremely relevant to financial decisions in this setting. Of course, in interpreting the magnitude of our effect, one might wish to consider the likelihood of information transfer in the real world; our design estimates the impact of information about one's peers conditional on receiving information – the endogenous acquisition of information is not studied here.[54]

Finally, it is important to note that the type of social learning we focus on is that of classic models, such as Banerjee (1992) and Bikhchandani et al. (1992): learning that occurs upon observation of the revealed decision of purchase by a peer. Of course, there are other types of social learning that might occur in finance, such as learning about the existence of an asset, or learning occurring after the purchase (when to sell the asset, when to buy more); these channels are shut down in our study because of the design of the financial asset. However, these other forms of social learning are likely important as well.

2.4 Conclusion

Peer effects are an important, and often confounding, topic of study across the social sciences. In addition, in many settings – particularly in finance – identifying *why* a person's choices are affected by his peers' is extremely important, beyond identifying peer effects overall. Our experimental design not only allows us to identify peer effects in investment decisions, it also decouples revealed preference from possession, allowing us to provide evidence that learning from one's peer's purchase decision and changing behavior due to a peer's possession of an asset *both* affect investment decisions. Perhaps the most important implication of these findings is support for models of social-learning based herding in financial markets, and also for herding based purely on a desire to own assets possessed by one's peers.

Our findings should be extended in several directions. Most fundamentally, it is important to determine their external validity. We are limited to studying a single asset; a single mechanism through which peers' choices were communicated; and, pairs of socially-related

[54]Duflo and Saez (2003)) present evidence that information about investment opportunities does flow naturally to peers; our goal of disentangling separate channels of peers' influence required *control* over these information flows.

individuals. One might be interested in whether our findings extend to assets with different expected returns or different exposures to risk; or, to investment decisions made from a larger choice set. One might also wish to study whether information transmitted directly among peers has a different effect from information transmitted through brokers. Of course, the selection of information transmitted by brokers and by peers will be endogenous, and studying the process determining *which* information gets transmitted, and to whom, would be of great interest. Finally, studying information transmission through a broader network of socially-related individuals is important as well.

To the extent that our results do shed light on financial decision making beyond our experimental context, it is important to understand what they imply for asset pricing and for policies that attempt to limit financial market instability. For example, herds based on social learning could be mitigated if more, and better, information is made available to investors. On the other hand, information provision will be less successful in limiting correlated choices among peers if those correlated choices are driven by social utility. Our findings suggest that information provision could reduce market instability, but only to a certain point, because a significant component of the peer effects we find was generated by social utility.[55]

In addition to the context of financial decision making, our experimental design could be used in other settings to identify the channels through which peer effects work. In marketing, various social media rely on different peer effect channels: Facebook "likes", Groupon sales, and product give-aways all rely on some combination of the channels studied here. Future work can compare the effectiveness of these strategies, and their impact through different channels, using designs similar to ours. One could also apply our experimental design to the study of technology adoption: one might wish to distinguish between the importance of learning from a peer's purchase decision and the desire to adopt technologies used by people nearby.[56] Finally, health-promoting behavior often is affected both by learning from peers' purchases and by peers' actual use of health care technology (e.g., vaccination or smoking cessation).[57] In these settings and others, separately identifying the roles of social learning

[55]Understanding *why* a peer's possession of an asset affects his decisions (beyond learning) is an important next step in clarifying the implications of our findings here.

[56]Foster and Rosenzweig (1995) and Conley and Udry (2010) identify the important role played by social learning in technology adoption. Social utility might exist in this setting because using a technology might be easier or less expensive when others nearby use it (network externalities); because one wishes not to fall behind those living nearby; etc.

[57]One can also learn about health care products from a peer's experience of the product – a type of

and social utility might be of interest to policymakers.

learning we do not consider in this study. Kremer and Miguel (2007) study the transmission of knowledge about de-worming medication through social networks.

2.A Figures and Tables

Figure 2-1: **Experimental design "roadmap"**

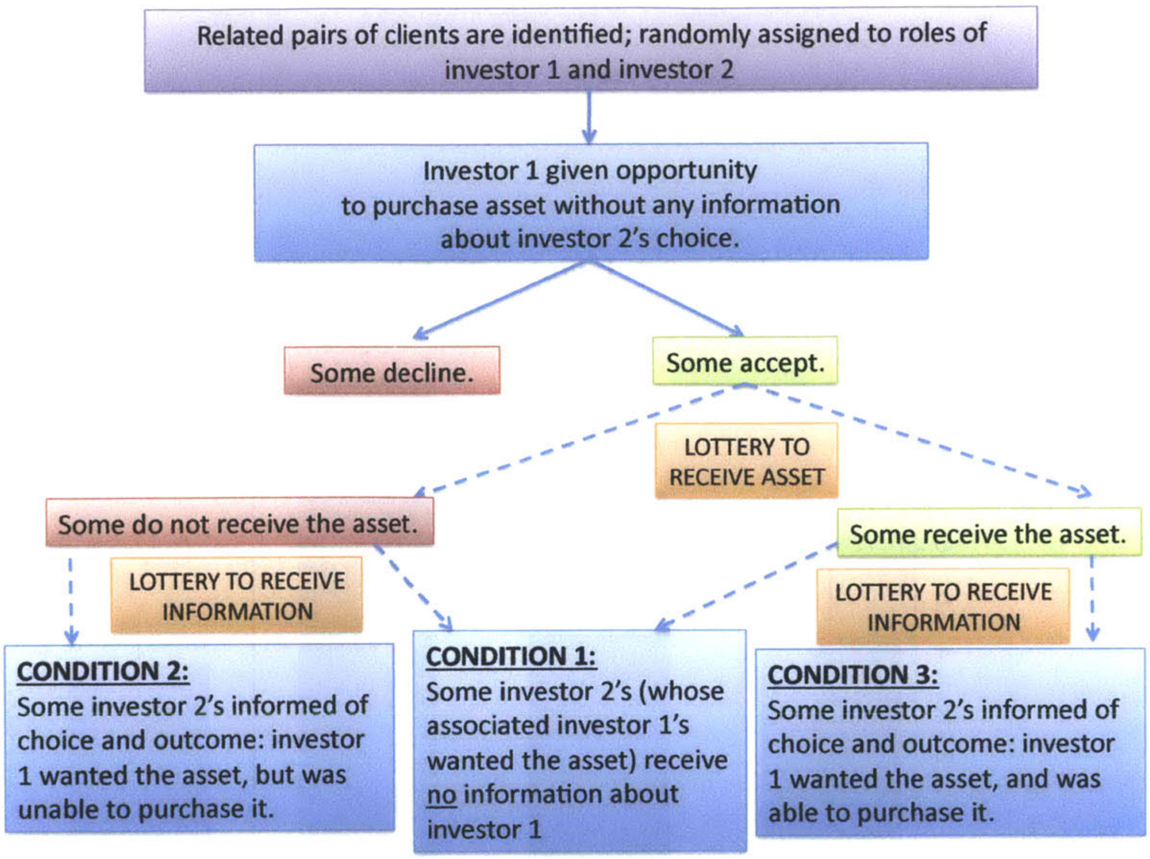

Figure 2-2: **Investor 2's take-up rates**

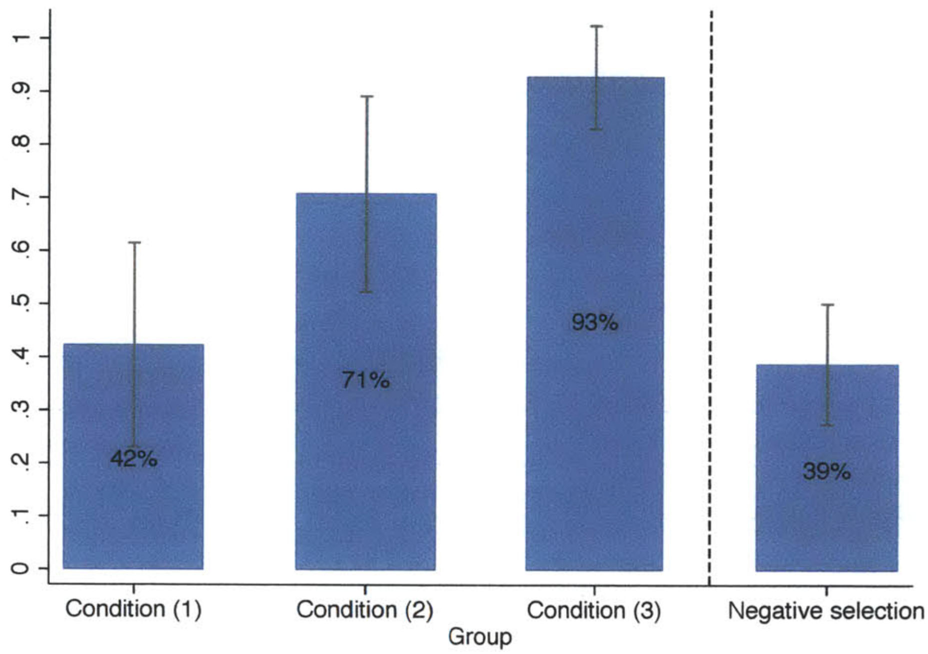

Note: This figure presents the mean (and 95% confidence intervals) of take-up rates for each group of investor 2's. Investors in conditions (1) to (3) have peers who wanted the asset. These investors were randomly allocated to one of these 3 groups. Those in condition (1) had no information about their peers. Those in condition (2) had information that their peers wanted to purchase the asset but had that choice rejected by the lottery. Those in condition (3) had information that their peers wanted and received the asset. Investors in the negative selection group have peers who did not want to purchase the asset (and received no information about their peer).

Figure 2-3: **Heterogeneity of Social Learning Effects**

Figure 2-3.1: Investor 2 has a technical occupation

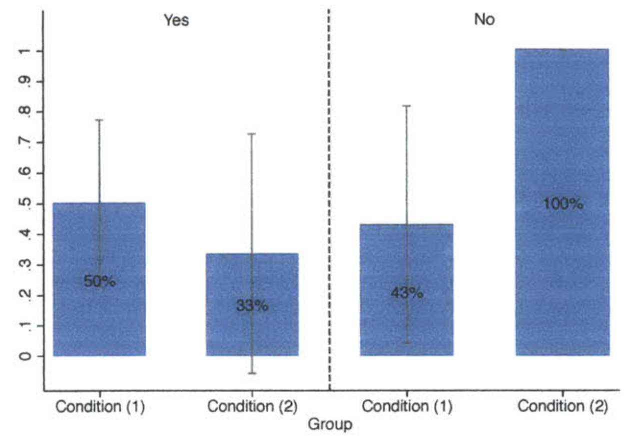

Figure 2-3.2: Associated investor 1 has a technical occupation

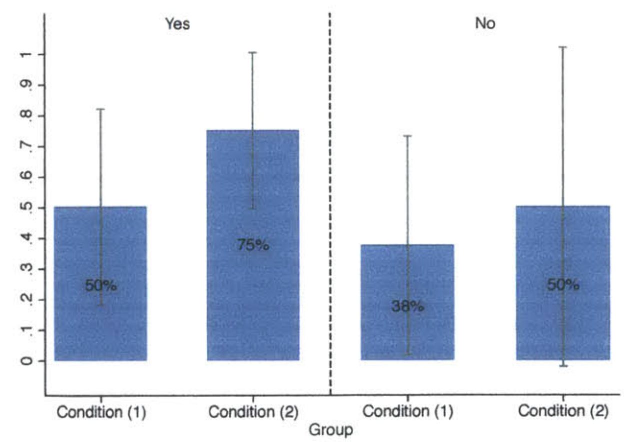

Note: Panel A presents figures with the mean (and 95% confidence intervals) of take-up rates for investor 2's in conditions (1) and (2), separately for those who have and who do not have a technical occupation. Panel B presents these figures separately for those whose associated investor 1's have and who do not have a technical occupation. Investors in conditions (1) and (2) have peers who wanted the asset. Those in condition (1) had no information about their peers. Those in condition (2) had information that their peers wanted to purchase the asset but had that choice rejected by the lottery.

Table 2-1: **Characteristics of the Experimental Sample**

	Experimental Sample							Universe
		Investor 1			Investor 2			
	Full Sample	All	Wanted the asset?		All	Peer wanted the asset?		
			Yes	No		Yes	No	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Age	38.15	39.12	39.60	38.60	37.18	36.45	37.97	34.14
	(0.80)	(1.14)	(1.60)	(1.62)	(1.12)	(1.50)	(1.68)	(0.16)
Gender (=1 If male)	0.680	0.747	0.769	0.722	0.613	0.641	0.583	0.729
	(0.027)	(0.036)	(0.048)	(0.053)	(0.040)	(0.055)	(0.059)	(0.006)
Married	0.413	0.440	0.436	0.444	0.387	0.333	0.444	0.340
	(0.028)	(0.041)	(0.057)	(0.059)	(0.040)	(0.054)	(0.059)	(0.006)
Single	0.557	0.527	0.513	0.542	0.587	0.628	0.542	0.647
	(0.029)	(0.041)	(0.057)	(0.059)	(0.040)	(0.055)	(0.059)	(0.006)
Earnings	4,500	5,000	5,000	5,000	4,000	4,000	3,500	3,200
	(256)	(499)	(501)	(775)	(507)	(504)	(650)	(126)
Technical Occupation	0.62	0.64	0.69	0.58	0.59	0.55	0.64	0.59
	(0.03)	(0.05)	(0.06)	(0.07)	(0.05)	(0.07)	(0.07)	(0.01)
Relationship with associated investor (=1 if family)	0.48	0.48	0.53	0.43	0.48	0.53	0.43	-
	(0.03)	(0.04)	(0.06)	(0.06)	(0.04)	(0.06)	(0.06)	
N	300	150	78	72	150	78	72	5506

Notes: Column 1 presents the characteristics of the experimental sample, combining investor 1's and investor 2's. Column 2 presents the sample characteristics of investor 1's in the experimental sample, while columns 3 and 4 present the information for investor 1's who wanted and who did not want the asset, respectively. Column 5 presents the characteristics of investors 2's in the experimental sample, while columns 6 and 7 present the information for investor 2's whose peers wanted and did not want the asset, respectively. Column 8 presents the characteristics of the universe of investors in the main office of the brokerage. Each line presents averages of the corresponding variable. For earnings, we present the median value instead of the mean due to large outliers. The sample size for the earnings variable is smaller due to missing values. The sample size for the technical occupation dummy is smaller because we considered as missing values professions that are indeterminate. The omitted value for "Relationship with associated investor" is "friends". This variable is not defined for investors outside the experiment's sample.

Table 2-2: **Covariates balance**

	Investor 2 conditional on investor 1 wanted to purchase the asset							
	Control	Learning only	Learning + possession	p-value of test:				N
				(1)=(2)=(3)	(1)=(2)	(1)=(3)	(2)=(3)	
	N=26	N=24	N=28					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Age	37.92	34.50	36.75	0.59	0.31	0.75	0.57	78
	(2.16)	(2.55)	(2.98)					
Gender (=1 If male)	0.654	0.667	0.607	0.90	0.93	0.73	0.66	78
	(0.095)	(0.098)	(0.094)					
Married	0.385	0.250	0.357	0.56	0.31	0.84	0.41	78
	(0.097)	(0.090)	(0.092)					
Single	0.538	0.708	0.643	0.47	0.22	0.44	0.62	78
	(0.100)	(0.095)	(0.092)					
Earnings	4,000	4,000	4,500	0.81	0.79	0.68	0.52	67
	(782)	(534)	(1,941)					
Technical Occupation	0.67	0.40	0.53	0.29	0.12	0.44	0.48	51
	(0.11)	(0.13)	(0.13)					
Relationship with investor 1 (=1 if family)	0.46	0.67	0.46	0.24	0.15	0.98	0.14	78
	(0.10)	(0.10)	(0.10)					
Joint test				0.15	0.26	0.86	0.29	

Notes: The sample is conditioned on investor 2's whose associated investor 1's wanted to purchase the asset. Each line presents averages of the corresponding variable for each treatment group. Robust standard errors in parentheses. For each variable, the p-value of an F-test that the mean of the corresponding variable is the same for all treatment groups is presented in column 4. The p-value of a joint test of equality for all variables is also presented. The p-values of F-tests on pairwise treatment group comparisons are presented in columns 5 to 7. For earnings, we present the median and the p-value of a test that the median of this variable is the same for all treatment groups. The sample size for the earnings variable is smaller due to missing values. The sample size for the technical occupation dummy is smaller because we considered as missing values professions that are indeterminate.

Table 2-3: **Peer Effects, Social Learning, Social Utility, and Selection: Take-up Rates**

Dependent variable	Wanted to purchase the asset			
	(1)	(2)	(3)	(4)
Learning alone	0.285**	0.298**	0.328**	0.278**
(Condition (2) - Condition (1))	(0.136)	(0.140)	(0.134)	(0.127)
Learning and possession	0.505***	0.540***	0.552***	0.500***
(Condition (3) - Condition (1))	(0.110)	(0.122)	(0.123)	(0.111)
Negative selection	-0.034	0.011	-0.005	0.042
	(0.114)	(0.124)	(0.118)	(0.117)
Investor 1				0.128
				(0.106)
Possession alone	0.220**	0.242**	0.224*	0.222**
(Condition (3) - Condition (2))	(0.106)	(0.109)	(0.124)	(0.108)
Mean (no information; peer chose the asset)	0.423			
(Condition (1))	(0.099)			
Broker fixed effects	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
N	150	150	150	300
R^2	0.186	0.228	0.283	0.219

Notes: Column 1 presents the results of a regression of a dummy variable equal to one if the investor wanted to purchase the asset on a dummy for condition (3), a dummy for condition (2), and a dummy indicating whether the associated investor 1 did not want to purchase the asset ("Negative selection"). Investor 2's in condition (1) is the omitted group. This regression uses only the sample of investor 2's. The regression presented in column 2 includes broker fixed effects. The regression presented in column 3 includes the covariates presented in Table 2-2. We did not include earnings or the technical occupation dummy as this would reduce our sample size (results including these variables are similar). The regression presented in column 4 combines the sample of investors 1 and 2, and includes an indicator variable for investor 1. Standard errors are clustered at the pair level. In all columns, "Possession alone" gives the difference between the coefficient on "Learning and possession" and the coefficient on "Learning alone." * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 2-4: **Peer Effects, Social Learning, Social Utility, and Selection: Amount Invested**

Dependent variable	Amount invested			
	(1)	(2)	(3)	(4)
Learning alone	948.7***	861.3**	825.3*	715.2*
(Condition (2) - Condition (1))	(357.7)	(379.5)	(421.2)	(394.5)
Learning and possession	2,633.2***	2,556.5***	2,564.8***	2,521.4***
(Condition (3) - Condition (1))	(702.9)	(633.0)	(613.4)	(611.9)
Negative selection	-106.8	-3.0	-6.4	123.9
	(239.0)	(272.7)	(305.3)	(308.6)
Investor 1				503.8*
				(300.1)
Possession alone	1,684.5**	1,695.1**	1,739.6**	1,806.1**
(Condition (3) - Condition (2))	(731.4)	(721.2)	(743.4)	(727.0)
Mean (no information; peer chose the asset)	884.6			
(Condition (1))	(210.0)			
Broker fixed effects	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
N	150	150	150	300
R^2	0.251	0.277	0.316	0.264

Notes: Column 1 presents the results of a regression of the amount invested in the asset on a dummy for condition (3), a dummy for condition (2), and a dummy indicating whether the associated investor 1 did not want to purchase the asset ("Negative selection"). Investor 2's in condition (1) is the omitted group. This regression uses only the sample of investor 2's. The regression presented in column 2 includes broker fixed effects. The regression presented in column 3 includes the covariates presented in Table 2-2. We did not include earnings or the technical occupation dummy as this would reduce our sample size (results including these variables are similar). The regression presented in column 4 combines the sample of investors 1 and 2, and includes an indicator variable for investor 1. Standard errors are clustered at the pair level. In all columns, "Possession alone" gives the difference between the coefficient on "Learning and possession" and the coefficient on "Learning alone." * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 2-5: **Peer Effects, Social Learning, Social Utility, and Selection: Invested More than Minimum**

Dependent variable	Invested more than minimum			
	(1)	(2)	(3)	(4)
Learning alone	0.212**	0.203**	0.186*	0.173*
(Condition (2) - Condition (1))	(0.097)	(0.093)	(0.098)	(0.095)
Learning and possession	0.497***	0.475***	0.481***	0.485***
(Condition (3) - Condition (1))	(0.103)	(0.102)	(0.104)	(0.101)
Negative selection	-0.038	-0.029	-0.033	-0.016
	(0.038)	(0.042)	(0.046)	(0.049)
Investor 1				0.097*
				(0.053)
Possession alone	0.286**	0.271**	0.295**	0.311**
(Condition (3) - Condition (2))	(0.131)	(0.130)	(0.132)	(0.128)
Mean (no information; peer chose the asset)	0.038			
(Condition (1))	(0.038)			
Broker fixed effects	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
N	150	150	150	300
R2	0.338	0.366	0.402	0.295

Notes: Column 1 presents the results of a regression of a dummy variable equal to one if the investor invested more than the minimum amount on a dummy for condition (3), a dummy for condition (2), and a dummy indicating whether the associated investor 1 did not want to purchase the asset ("Negative selection"). Investor 2's in condition (1) is the omitted group. This regression uses only the sample of investor 2's. The regression presented in column 2 includes broker fixed effects. The regression presented in column 3 includes the covariates presented in Table 2-2. We did not include earnings or the technical occupation dummy as this would reduce our sample size (results including these variables are similar). The regression presented in column 4 combines the sample of investors 1 and 2, and includes an indicator variable for investor 1. Standard errors are clustered at the pair level. In all columns, "Possession alone" gives the difference between the coefficient on "Learning and possession" and the coefficient on "Learning alone." * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 2-6: **Heterogeneity of Social Learning Effects**

	Investor 2 has a technical occupation			Associated investor 1 has a technical occupation		
	Yes	No	p-value of test (1)=(2)	Yes	No	p-value of test (4)=(5)
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: no controls						
Learning alone	-0.167	0.571***	0.027	0.250	0.125	0.749
(Condition (2) - Condition (1))	(0.249)	(0.198)		(0.215)	(0.323)	
Panel B: full specification						
Learning alone	-0.153	0.543***	0.003	0.208	0.077	0.659
(Condition (2) - Condition (1))	(0.227)	(0.139)		(0.177)	(0.286)	
Mean (no information; peer chose the asset)	0.500	0.429	0.771	0.500	0.375	0.620
(Condition (1))	(0.140)	(0.197)		(0.168)	(0.182)	

Notes: Panel A reports comparisons of means for each group; panel B reports coefficients using the full specification of column 4 from Table 2-3 with the condition 2 dummy interacted with a dummy indicating that the relevant investor (either 1 or 2) has a technical occupation, and also interacted with a dummy indicating that the relevant investor does not have a technical occupation. We also include a technical occupation dummy. Columns 1 and 2 present the learning effects (investor 2's in condition (2)) and the take-up rates for investor 2's in the control group (condition (1)) for investor 2's who have a technical occupation and for investor 2's who do not have a technical occupation, respectively. Column 3 reports the p-value of a test that the learning effect and the baseline take-up rates are the same for these two groups. Investor 2's for whom it is not possible to determine whether they have a technical occupation or not are excluded from this analysis. Columns 4 to 6 are analogous, but examine heterogeneity associated with investor 2' s peer's (i.e., investor 1's) having a technical occupation. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 2-7: **Robustness Tests**

Interaction of the treatment effects with:	Relationship with investor 1 (=1 if family) (1)	Broker experience within the experiment (2)	RA was present during the call (3)
Learning alone	0.077	-0.001	0.140
	(0.305)	(0.008)	(0.268)
Learning and possession	0.417*	-0.003	0.248
	(0.232)	(0.008)	(0.234)
Possession alone	0.340	-0.001	0.108
	(0.220)	(0.007)	(0.203)

Notes: This table presents coefficients on the interactions of the variables at the column heading with the treatment effects of interest. These results are based on the regressions used in the full specification of column 4 from Table 2-3, including interactions of the group dummies ($I_{c,i}$, where $c \in \{\text{condition (2)}, \text{condition (3)}, \text{“negative selection”}, \text{investor 1}\}$) with the corresponding variables. We also include the main effect of the corresponding variable. In column 1, we interact the treatment effects with a dummy variable equal to one if the investors 1 and 2 are family members. The omitted category is “friends”. In column 2, we interact the treatment effects with a variable indicating the number of calls that the broker made before the day of the call. In column 3, we interact the treatment effects with a variable indicating whether our research assistant was present during the call. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 2.A.1: **Covariates Balance - Other Randomizations**

	Assignment to investor 1 or investor 2				Lottery for investor 1's who wanted the asset			
	Investor 1	Investor 2	p-value of test (1)=(2)	N	Won	Lost	p-value of test (5)=(6)	N
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Age	39.12	37.18	0.22	300	39.47	39.71	0.94	78
	(1.14)	(1.12)			(2.34)	(2.23)		
Gender (=1 If male)	0.747	0.613	0.01	300	0.861	0.690	0.07	78
	(0.036)	(0.040)			(0.058)	(0.072)		
Married	0.440	0.387	0.35	300	0.472	0.405	0.56	78
	(0.041)	(0.040)			(0.084)	(0.077)		
Single	0.527	0.587	0.30	300	0.528	0.500	0.81	78
	(0.041)	(0.040)			(0.084)	(0.078)		
Earnings	5,000	4,000	0.22	270	5,000	5,000	0.59	74
	(499)	(507)			(925)	(754)		
Technically skilled	0.64	0.59	0.47	195	0.73	0.66	0.55	55
	(0.05)	(0.05)			(0.09)	(0.09)		
Relationship with peer (=1 if family)	-	-	-	-	0.44	0.60	0.19	78
					(0.08)	(0.08)		
Joint test			0.15				0.23	

Notes: Columns 1 and 2 present the averages of the corresponding variable, respectively, for investors assigned to be in the role of investor 1 and for those assigned to be in the role of investor 2. Robust standard errors in parentheses. Relationship with peer is not considered in this comparison since this variable is equal for both groups by construction. Column 3 presents the p-value of an F-test that the mean of the corresponding variable is the same for these two groups. The p-value of a joint test of equality for all variables is also presented. Column 5 presents the averages for investor 1's who wanted the asset and won the lottery, while column 6 presents the averages for investor 1's who wanted the asset but did not win the lottery. Column 7 presents the p-value of an F-test that the mean of the corresponding variable is the same for these two groups. For earnings, we present the median and the p-value of a test that the median of this variable is the same for the corresponding groups. The sample size for the earnings variable is smaller due to missing values. The sample size for the technical occupation dummy is smaller as we considered as missing value professions that are indeterminate.

Table 2.A.2: **List of occupations**

Occupation	Technical occupation?	Number of investors in the sample
Engineer	Yes	39
Business administrator	Yes	34
Bank employee or clerk	Yes	8
Accountant	Yes	6
Real estate, insurance or securities broker	Yes	5
Economist	Yes	5
Systems analyst	Yes	4
Capitalist receiving income from invested capital	Yes	2
Commercial establishment owner	Yes	2
Insurance professional	Yes	2
Actuary or mathematician	Yes	1
Industrial establishment owner	Yes	1
Lawyer	No	15
Medical doctor	No	13
Architect	No	5
Journalist	No	5
Retail or wholesale salesperson	No	4
Administrative agent	No	3
Physical therapist or occupational therapist	No	3
Aircraft pilot	No	3
Advertising agent	No	3
Office assistant or similar occupation	No	2
Nurse or nutritionist	No	2
Pharmacist	No	2
Dentist	No	2
Social worker	No	1
Professional athlete or sports coach	No	1
Liibrarian, archivist, museologist or archeologist	No	1
Communications specialist	No	1
Interior designer	No	1
Designer	No	1
Maid	No	1
Speech therapy	No	1
Member of the judiciary: federal supreme court justice	No	1
Professional driver	No	1
Psychologist	No	1
Public relations	No	1
Student	Undetermined	31
Other	Undetermined	29
Retired (exception: civil servants)	Undetermined	14
Entrepreneur	Undetermined	11
Teacher	Undetermined	7
Professor	Undetermined	6
Manager	Undetermined	2
State public servant	Undetermined	2
Federal public servant	Undetermined	2
Technical specialist	Undetermined	2
Coordinator or supervisor	Undetermined	1
Retired	Undetermined	1
Electricity, electronics or telecommunications technician	Undetermined	1
Laboratory or X-ray technician	Undetermined	1
Chemical technician	Undetermined	1

Table 2.A.3: **Heterogeneity of Social Learning Effects - Amount Invested**

	Investor 2 has a technical occupation			Associated investor 1 has a technical occupation		
	Yes	No	p-value of test (1)=(2)	Yes	No	p-value of test (4)=(5)
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: no controls						
Learning alone	-404.8	1,920.6***	0.004	650.0	750.0	0.927
(Condition (2) - Condition (1))	(513.6)	(538.7)		(523.5)	(955.0)	
Panel B: full specification						
Learning alone	-807.4	1,797.9***	0.004	929.7	-635.2	0.326
(Condition (2) - Condition (1))	(764.1)	(480.3)		(660.7)	(1,314.4)	
Mean (no information; peer chose the asset)	1,071.4	857.1	0.673	1,100.0	750.0	0.515
(Condition (1))	(309.0)	(393.3)		(381.0)	(363.1)	

Notes: Panel A reports comparisons of means for each group; panel B reports coefficients using the full specification of column 4 from Table 2-3 with the condition 2 dummy interacted with a dummy indicating that the relevant investor (either 1 or 2) has a technical occupation, and also interacted with a dummy indicating that the relevant investor does not have a technical occupation. We also include a technical occupation dummy. Columns 1 and 2 present the learning effects (investor 2's in condition (2)) and the take-up rates for investor 2's in the control group (condition (1)) for investor 2's who have a technical occupation and for investor 2's who do not have a technical occupation, respectively. Column 3 reports the p-value of a test that the learning effect and the baseline take-up rates are the same for these two groups. Investor 2's for whom it is not possible to determine whether they have a technical occupation or not are excluded from this analysis. Columns 4 to 6 are analogous, but examine heterogeneity associated with investor 2' s peer's (i.e., investor 1's) having a technical occupation. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 2.A.4: **Heterogeneity of Social Learning Effects - Invested More than Minimum**

	Investor 2 has a technical occupation			Associated investor 1 has a technical occupation		
	Yes	No	p-value of test (1)=(2)	Yes	No	p-value of test (4)=(5)
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: no controls						
Learning alone	-0.071	0.444**	0.011	0.067	0.250	0.512
(Condition (2) - Condition (1))	(0.073)	(0.176)		(0.153)	(0.230)	
Panel B: full specification						
Learning alone	-0.042	0.371**	0.022	0.040	0.097	0.842
(Condition (2) - Condition (1))	(0.088)	(0.158)		(0.139)	(0.258)	
Mean (no information; peer chose the asset)	0.071	0.000	0.336	0.100	0.000	0.335
(Condition (1))	(0.072)	(0.000)		(0.101)	(0.000)	

Notes: Panel A reports comparisons of means for each group; panel B reports coefficients using the full specification of column 4 from Table 2-3 with the condition 2 dummy interacted with a dummy indicating that the relevant investor (either 1 or 2) has a technical occupation, and also interacted with a dummy indicating that the relevant investor does not have a technical occupation. We also include a technical occupation dummy. Columns 1 and 2 present the learning effects (investor 2's in condition (2)) and the take-up rates for investor 2's in the control group (condition (1)) for investor 2's who have a technical occupation and for investor 2's who do not have a technical occupation, respectively. Column 3 reports the p-value of a test that the learning effect and the baseline take-up rates are the same for these two groups. Investor 2's for whom it is not possible to determine whether they have a technical occupation or not are excluded from this analysis. Columns 4 to 6 are analogous, but examine heterogeneity associated with investor 2' s peer's (i.e., investor 1's) having a technical occupation. * significant at 10%; ** significant at 5%; *** significant at 1%.

2.B Appendix: A Simple Model of Financial Decisions Under Social Influence

Our model studies an investment decision made by an individual under several conditions. First, we present the investment decision under uncertainty, but with no social influence. Second, we present the investment decision with social learning present, using the ingredients of a canonical social learning model: a peer makes an investment acting on a private signal, and this action can be used by another investor to make an informational inference before taking his own action. Third, we allow the ownership of an asset to affect a socially-related investor's utility of owning the asset, aside from any learning – that is, we allow for a social utility effect. A peer's purchase decision typically will produce both social learning and social utility effects; we consider a case in which both effects are active (the full "peer effect") and a case in which the revealed preference purchase decision is de-coupled from possession. This de-coupling allows one to observe each channel through which peer effects work, and motivates our experimental design.

Investment without Peer Effects

Consider an investor i's decision to invest in a risky asset.[58] The asset's return is given by x, with probability density function $f(x)$, and investor i's utility is $u_i(x) = u(x)$ for all i. In our field experiment, investors received calls from brokers who offered them a financial asset for purchase. The brokers attempted to convey the same information about the asset in every call using a pre-specified script; thus, the information they provided can be thought of as a signal, s_i, coming from a single distribution, with probability density function $g(s_i)$. Importantly, not every investor would have received exactly the same information: calls evolve in different ways, investors ask different questions about the asset, etc., meaning that each investor received a different signal realization, s_i, from the common distribution of signals.

For expositional simplicity, assume that the conditional density $f(x|s_i)$ satisfies the monotone likelihood ratio property (MLRP) such that, intuitively, higher values of s_i are indicative of higher values of x. Under these conditions, investor i is willing to invest if and

[58]Note that we implicitly assume that when investing in isolation, investor i does not take into consideration any investor j ($j \neq i$) at all – he is "unaware." In the context of our experiment, we believe that this assumption is reasonable, as we discuss in the text.

only if

$$\int u(x)f(x|s_i)dx \geq \bar{u}, \tag{2.2}$$

where $\bar{u}$ denotes the outside option of the investor. Given that $f(x|s_i)$ satisfies MLRP and given mild monotonicity assumptions on the utility function $u(\cdot)$ of the investor, there exists a unique threshold $\bar{s}_1$ such that for any $s_i \geq \bar{s}_1$ investor i is willing to invest. Denote the decision to buy the asset made by investor i by $b_i = \{0, 1\}$. Hence, for an investor making a purchase decision in isolation, we have

$$b_i = 1 \Leftrightarrow s_i \geq \bar{s}_1. \tag{2.3}$$

Investment with Social Learning Alone

Suppose that instead of making his investment choice in isolation, before making his own decision, investor i observes the investment decision of investor j which is given by b_j. Assume that investor j made his choice $b_j = 1$ in isolation and hence his decision rule is given by (2.3).[59] Thus, when investor i observes $b_j = 1$ he correctly infers that $s_j \geq \bar{s}_1$ and he is willing to invest if and only if

$$\int u(x)f(x|s_i; s_j \geq \bar{s}_1)dx \geq \bar{u}. \tag{2.4}$$

Furthermore, given that $f(x|s_i; s_j)$ satisfies MLRP we have

$$\int u(x)f(x|s_i; s_j \geq \bar{s}_1)dx \geq \int u(x)f(x|s_i)dx, \tag{2.5}$$

for all s_i. It is straightforward to show by comparing (2.4) and (2.2) that the signal realization threshold for investor i that is necessary to induce purchase of the asset is lower when $b_j = 1$ is observed than when investor i makes his choice in isolation. This is because in the former case, regardless of his own private information summarized by s_i, investor i has additional favorable information about the asset from observing the purchase of investor j. This is the pure social learning effect.

Denote the threshold for s_i when investor i observes $b_j = 1$ by $\bar{s}_2$ and note that $\bar{s}_2 \leq \bar{s}_1$.

[59] We focus on the case of investor i observing that investor j chose to purchase the asset (rather than choosing *not* to purchase it) because in the experimental design, we were not allowed to inform investors that their peer chose not to purchase the asset.

In particular, after observing a purchase decision made by investor j, the decision rule of investor i is given by

$$b_i = 1 \Leftrightarrow s_i \geq \bar{s}_2. \tag{2.6}$$

Social Utility and Social Learning

We now consider the situation in which both social utility and social learning effects are present. Our focus (following much of the literature on peer effects in financial decisions) is on social utility effects that result in a *positive* effect of a peer's possession of an asset (denoted by $p_j = \{0, 1\}$) on one's own utility.[60] In particular, when investor i considers purchasing the asset, we assume that $u(x|p_j = 1) \geq u(x|p_j = 0)$ for all x. That is, investor i's utility is higher for all asset return realizations if the asset is also possessed by an investor j who is a peer of investor i. Using the notation of our model, an investor j's purchase of an asset, $b_j = 1$, typically implies both that investor i infers favorable information about the asset, $s_j \geq \bar{s}_1$, and that investor j now possesses the asset, $p_j = 1$, which might affect investor i's utility of owning the asset (due to a taste for joint consumption, "keeping-up-with-the-Joneses" preferences).

When investor i observes that investor j expressed an intention to invest, $b_j = 1$, and was allowed to invest, $p_j = 1$, both investor i's utility $u(x|p_j = 1)$ *and* his information about the asset $f(x|s_i; s_j \geq \bar{s}_1)$ are affected, relative to his choice in isolation (that is, relative to $u(x) = u(x|p_j = 0)$ and $f(x|s_i)$).[61] In this case, one observes the "full" peer effect, and investor i invests if and only if

$$\int u(x|p_j = 1) f(x|s_i; s_j \geq \bar{s}_1) dx \geq \bar{u}. \tag{2.7}$$

Denote the threshold for s_i above which investor i is willing to invest when exposed to both peer effects channels by $\bar{s}_3$. Then, the decision rule for investor i is given by

$$b_i = 1 \Leftrightarrow s_i \geq \bar{s}_3. \tag{2.8}$$

To separate the effects of social learning and social utility, we need to decouple will-

[60] One could also imagine a *negative* correlation, for example, out of a desire to insure one's peers, or to differentiate oneself. See Clark and Oswald (1998).

[61] We are assuming here that the utility function discussed above, $u(x)$, is the same as $u(x|p_j = 0)$ here. In addition, we are assuming that investor j made his decision in isolation.

ingness to purchase (and the informative signal of the purchase decision) from possession. Consider the situation where investor i observes that investor j expressed a revealed preference to invest, but was not allowed to do so (perhaps due to capacity constraints). In this case, investor i infers that $s_j \geq \bar{s}_1$, but also knows that investor j did not obtain the asset, so $p_j = 0$. This condition is equivalent to the "social learning alone" problem discussed above: there is no direct effect of possession on investor i's utility from the asset, but there is social learning. Thus, investor i purchases the asset if and only if (2.4) is satisfied (since $u(x) = u(x|p_j = 0)$) and this leads to the same decision rule as (2.6) with the threshold $\bar{s}_2$.

The following proposition summarizes investor i's purchase decisions across conditions.

Proposition 1. *The threshold for the signal s_i above which investor i is willing to purchase the asset (and, the likelihood of a purchase of the asset by investor i) is highest (lowest) when the investor makes his decision in isolation, lower (higher) when he observes that investor j intended to purchase the asset but did not obtain it, and lowest (highest) when investor j intended to purchase the asset, and obtained it: $\bar{s}_1 \geq \bar{s}_2 \geq \bar{s}_3$ (and $\Pr(s_i \geq \bar{s}_3) \geq \Pr(s_i \geq \bar{s}_2) \geq \Pr(s_i \geq \bar{s}_1)$).*

Proof. The relationship between $\bar{s}_1$ and $\bar{s}_2$ follows immediately from comparing the inequalities (2.2) and (2.4) and the monotone likelihood ratio property of $f(x|s_i; s_j)$. Similarly, comparison of the inequalities (2.4) and (2.7) and $u(x) = u(x|p_j = 0) \leq u(x|p_j = 1)$ establishes that $\bar{s}_2 \geq \bar{s}_3$. Finally, $\Pr(s_i \geq \bar{s}_3) \geq \Pr(s_i \geq \bar{s}_2) \geq \Pr(s_i \geq \bar{s}_1)$ follows from the ranking of the thresholds. □

The difference between $\bar{s_2}$ and $\bar{s_3}$ is the result of a difference in investor j's possession of the asset.[62] In one situation investor j received favorable information and expressed an intent to purchase the asset, but was unable to execute the purchase due to supply restrictions. In the other situation investor j received a favorable signal and was also able to obtain the asset. Thus, in the two cases investor i infers the same information (via investor j's choice) about the potential returns of asset x. However, only in the latter case is investor i's utility directly influenced by the investment *outcome* (and not just the purchase *intention*) of investor j. This is the social utility effect that raises the expected

[62] Note that the difference between $\bar{s_2}$ and $\bar{s_3}$ measures the impact of possession conditional on the presence of social learning. This is consistent with our experimental design, in which we are not able to measure the impact of possession in the absence of social learning.

utility of purchasing the asset for investor i over and above the social learning effect. In the inequalities in Proposition 1, the effect of social learning is captured by the difference between $\Pr(s_i \geq \bar{s}_2)$ and $\Pr(s_i \geq \bar{s}_1)$, and the effect of social utility is the difference between $\Pr(s_i \geq \bar{s}_3)$ and $\Pr(s_i \geq \bar{s}_2)$. The total peer effect is the difference between $\Pr(s_i \geq \bar{s}_3)$ and $\Pr(s_i \geq \bar{s}_1)$.

Our analysis readily extends to the case in which investor i's investment choice is continuous rather than limited to a binary decision. In particular, since $f(x|s_i; s_j)$ satisfies MLRP, the optimal investment in the asset is increasing in s_i and s_j and the expected equilibrium investment amounts will follow exactly the prediction regarding purchase rates in Proposition 1. Suppose individual i chooses an investment magnitude q_i^*, rather than making a binary investment decision. Since $f(x|s_i; s_j)$ satisfies MLRP, the optimal investment in the asset is increasing in s_i and s_j and we can rank the expected equilibrium investment amounts.

Proposition 2. *The expected equilibrium investment amount q_i^* of investor i is lowest when the investor makes his decision in isolation, higher when he observes that investor j intended to purchase the asset but did not obtain it, and highest when investor j intended to purchase, and obtained, the asset.*

Proof. The inference problem of investor i is the same as in Proposition 1. Thus, for a given signal s_i the described relationship holds for the actual equilibrium investment amount and follows immediately from comparing the expression for the utilities on the left-hand side of the inequalities (2.2), (2.4) and (2.7) and by noting that the optimal investment amount is increasing in s_i and s_j. Finally, taking expectations over the signal realizations s_i yields the ranking in expected investment amounts. □

Heterogeneous Investors

In practice, some investors are more financially sophisticated than others, and one would expect that this variation will affect the peer effects we study here – especially the impact of social learning. In particular, an unsophisticated investor may have much more to learn about an asset from the purchase decision of their peer than does a sophisticated investor, as the sophisticated investor likely has a very good sense of the asset's quality from his signal alone. Differing financial sophistication can be captured in our model by allowing the signals s_i and s_j to be drawn from distributions with differing precision. For simplicity, we make the assumption that, in contrast to unsophisticated investors, sophisticated investors receive perfectly informative signals. This assumption generates the following prediction of heterogeneous effects of social learning.

Proposition 3. *The thresholds $\bar{s}_1$ and $\bar{s}_2$ for the signal s_i above which investor i is willing to purchase the asset (and hence the likelihood of investor i purchasing the asset) are identical if investor i is financially sophisticated (i.e., signal s_i is perfectly informative). If investor j is sophisticated, then investor i follows the choice of investor j when observing the decision of investor j.*

Proof. If s_i is perfectly informative (i.e., investor i is sophisticated), then s_i is a sufficient statistic for x. As a result, s_j, and hence the purchase decision of investor j, has no informational value for sophisticated investor i and does not influence the threshold $\bar{s}_1$. Hence, $\bar{s}_1 = \bar{s}_2$. If s_j is perfectly informative, then investor j knows the value of x and makes a perfectly informed investment decision. As a result, investor i follows investor j's choice. □

Proposition 3 suggests that social learning will be limited (in fact, given the simplifying assumptions made, will be nonexistent) for sophisticated investors. These investors are sufficiently well-informed that they are not influenced by the revealed preference of another investor. The proposition further shows that social learning will have relatively strong effects on investment choices if the investor whose choice is observed is sophisticated.[63]

[63] We have assumed that sophisticated investors receive perfectly informative signals. Our results can be extended to the case in which sophisticated investors receive more informative, but still imperfectly informative, signals. While results for general distributions of x, s_i and s_j that satisfy MLRP do not exist, it is straightforward to show that for binary signal structures, the impact of social learning will be relatively small when the observing investor is sophisticated and relatively large when the observed investor is sophisticated. Finally, it is worth noting that, another investor's *possession* of the asset could still affect financially sophisticated investors' choices; similarly a financially unsophisticated investor's purchase decision – when accompanied by possession – could influence a peer's choice. Both of these effects would work through the social utility channel. Thus, we emphasize that these predictions of heterogeneous treatment effects apply to social learning effects alone, but not necessarily the overall peer effect.

2.C Appendix: The Financial Asset

The asset offered to clients in the experiment was a combination of an actively-managed, open-ended long/short mutual fund and a real estate note (LCI, *Letra de Créditto Imobiliário*), for a term of 1 year. A client was required to make a minimum investment of at least R$1,000 (approximately US$550) in each individual asset if he chose to purchase the combination asset (for a minimum total investment of approximately US$1,100). The long/short fund seeks to outperform the interbank deposit rate (CDI, *Certificado de Depósito Interbancário*) by allocating assets in fixed-income assets, equity securities and derivatives. The LCI is a low-risk asset which is attractive to personal investors, because it is exempt from personal income tax. The LCI offered in this particular combination had better terms than the real estate notes that were usually offered to clients of the brokerage with which we worked. First, the return of the LCI offered in the experiment was 98% of the CDI, while the best LCI offered to clients outside of the experiment had a return of 97% of the CDI. In addition, the brokerage firm usually required a minimum investment of R$10,000 to invest in an LCI, while the offer in the experiment reduced the minimum investment threshold to R$1,000. As noted in the text, the maximum investment in the LCI component was R$10,000.

2.D Appendix: Experimental Documentation

We enclose here English versions of the Qualtrics scripts used by the brokers in the phone calls, first to investor 1's and then to investor 2's.

Client number

Name of broker making phone call

Client number

Introduction

Description of asset

Combination of two investments:
- Fundo Long-Short multi-mercado (read brochure)
- LCI de 98% do CDI (read brochure)

Minimum investment:
- R$1,000 in LCI and R$1,000 in Fundo Long-Short

Maximum investment:
-R$10,000 in LCI and no limits in Fundo Long-Short

Observations to be told to client:
1) Special LCI usually not avaliable to clients. LCI typically available to clients has return of 97% of CDI and minimum investment of R$10,000

2) Emphasize that product can only be purchased during this call (take it or leave it): will not be sold on other occasions

3) Remind that LCI is exempt from income tax

4) Explain that only new resources will be accepted (and not resources already invested with the brokerage)

Limited supply

This is a special asset, only available in limited supply, and only to special clients like you.

As so, unfortunately, some of the clients that want the asset will not be able to actually purchase it.

Since we are a company that always wants to be as fair as possible, we want to give a chance to all the special clients we are calling and who are interested in the product. In addition to that, we would like to give the same chance to everyone.

Because of that, we will use a lottery to determine which clients will actually be able to implement the purchase, among those that chose to purchase the asset.

In this lottery half (50%) of the clients that choose to purchase the asset will have their choices authorized and implemented.

The lottery consists in drawing a random integer number between 1 and 100. If the number is 50 or less, the lottery will not authorize the investment. If the number is greater than 50, the lottery will authorize and make the investment.

It is important that you know that the decision you will make now if final. If you decide to purchase the asset, you will be authorizing the purchase. Therefore, if the lottery authorizes the purchase, the investment will be made.

Take advantage of this great opportunity to buy this exclusive product!

Investment decision

Ask the client what their decision is

- Wants to invest
- Does not want to invest

How much does he want to invest in the Fundo Long-Short multi-mercado?

How much does he want to invest in the LCI?

Investment authorization

A random number will now be drawn to determine whether or not you will be able to actually make the investment. The random number is ${e://Field/random}

Due to the outcome of the lottery, your investment was not authorized.

Due to the outcome of the lottery, your investment was authorized.

Was the investment authorized?

- Yes
- No

End of Call and Summary

Finish the call

This is the summary of the call. Please put the following information in your Excel spreadsheet of communication with the clients:

Client number: ${q://QID20/ChoiceTextEntryValue}

Did this client want to invest in the product? Yes

Amount invested in the Fundo Long-Short: ${q://QID18/ChoiceTextEntryValue}

Amount invested in the LCI: ${q://QID26/ChoiceTextEntryValue}

Was this client authorized to make the investment? Yes

This is the summary of the call. Please put the following information in your Excel spreadsheet of communication with the clients:

Client number: ${q://QID20/ChoiceTextEntryValue}

Did this client want to invest in the product? Yes

Amount invested in the Fundo Long-Short: ${q://QID18/ChoiceTextEntryValue}

Amount invested in the LCI: ${q://QID26/ChoiceTextEntryValue}

Was this client authorized to make the investment? No

This is the summary of the call. Please put the following information in your Excel spreadsheet of communication with the clients:

Client number: ${q://QID20/ChoiceTextEntryValue}

Did this client want to invest in the product? No

Amount invested in the Fundo Long-Short: 0

Amount invested in the LCI: 0

Was this client authorized to make the investment? N/A

Client number

Name of broker making phone call

Client number

Number of client of the (first) friend of this investor

Previous Choice by FRIEND 1

Did the first friend of this investor want to invest in this asset?

- Yes
- No

Was the first friend of this investor authorized to make the investment?

- Yes
- No

Introduction

Description of Asset

Combination of two investments:
- Fundo Long-Short multi-mercado (read brochure)
- LCI de 98% do CDI (read brochure)

Minimum investment:
- R$1,000 in LCI and R$1,000 in Fundo Long-Short
Maximum investment:
-R$10,000 in LCI and no limits in Fundo Long-Short

Observations to be told to client:
1) Special LCI usually not avaliable to clients. LCI typically available to clients has return of 97% of CDI and minimum
investment of R$10,000

2) Emphasize that product can only be purchased during this call (take it or leave it): will not be sold on other occasions

3) Remind that LCI is exempt from income tax

4) Explain that only new resources will be accepted (and not resources already invested with the brokerage)

Limited Supply

This is a special asset, only available in limited supply, and only to special clients l ke you.

As so, unfortunately, some of the clients that want the asset will not be able to actually purchase it.

Since we are a company that always wants to be as fair as possible, we want to give a chance to all the special clients we are calling and who are interested in the product. In addition to that, we would like to give the same chance to everyone.

Because of that, we will use a lottery to determine which clients will actually be able to implement the purchase, among those that chose to purchase the asset.

In this lottery half (50%) of the clients that choose to purchase the asset will have their choices authorized and implemented.

The lottery consists in drawing a random integer number between 1 and 100. If the number is 50 or less, the lottery will
not authorize the investment. If the number is greater than 50, the lottery will authorize and make the investment.

It is important that you know that the decision you will make now if final. If you decide to purchase the asset, you will be authorizing the purchase. Therefore, if the lottery authorizes the purchase, the investment will be made.

Take advantage of this great opportunity to buy this exclusive product!

Only Learning Treatment

Before asking whether or not the client wants to purchase the asset, tell him the information associated with the choice of the first friend and the outcome of the lottery for the first friend:

"We would like to inform you, before you make your decision, that [FIRST FRIEND'S NAME], your [RELATIONSHIP TO THIS CLIENT], received the same offer today. He/she chose to purchase the product. However, the lottery did not authorize him/her to make the purchase, so he/she will not make the investment."

SUMMARIZING: He/she wanted to make the investment but was not able to invest.

Possession and Learning Treatment

Before asking whether or not the client wants to purchase the asset, tell him the information associated with the choice of the first friend and the outcome of the lottery for the first friend:

"We would like to inform you, before you make your decision, that [FIRST FRIEND'S NAME], your [RELATIONSHIP TO THIS CLIENT], received the same offer today. He/she chose to purchase the product. The lottery authorized him/her to make the purchase, so he/she will make the investment."

SUMMARIZING: He/she wanted to make the investment and was able to invest.

Investment Decision

Ask the client what their decision is

- ○ Wants to invest
- ○ Does not want to invest

How much does he want to invest in the Fundo Long-Short multi-mercado?

How much does he want to invest in the LCI?

Investment Authorization

A random number will now be drawn to determine whether or not you will be able to actually make the investment.
The random number is ${e://Field/random}

Due to the outcome of the lottery, your investment was not authorized.

Due to the outcome of the lottery, your investment was authorized.

Was the investment authorized?

○ Yes

○ No

Relationship with First Investor, End of Call, and Summary

Had you previously heard about this offer/this product from [FIRST FRIEND'S NAME]?

○ Yes

○ No

What is your degree of relationship with [FIRST FRIEND'S NAME]? Examples: sibling, parent, friend, co-worker, etc.

Finish the phone call

This is the summary of the call. Please put the following information in your Excel spreadsheet of communication with the clients:

Client number: ${q://QID27/ChoiceTextEntryValue}

First friend's client number: ${q://QID30/ChoiceTextEntryValue}

Did the first friend want to invest in the product? ${q://QID21/ChoiceGroup/SelectedChoices}

Was the first friend authorized to make the investment? ${q://QID25/ChoiceGroup/SelectedChoices}

Did this client (second friend) want to invest in the product? Yes

Amount invested in the Fundo Long-Short: ${q://QID28/ChoiceTextEntryValue}

Amount invested in the LCI: ${q://QID38/ChoiceTextEntryValue}

Was this client authorized to make the investment? Yes

This is the summary of the call. Please put the following information in your Excel spreadsheet of communication with the clients:

Client number: ${q://QID27/ChoiceTextEntryValue}

First friend's client number: ${q://QID30/ChoiceTextEntryValue}

Did the first friend want to invest in the product? ${q://QID21/ChoiceGroup/SelectedChoices}

Was the first friend authorized to make the investment? ${q://QID25/ChoiceGroup/SelectedChoices}

Did this client (second friend) want to invest in the product? Yes

Amount invested in the Fundo Long-Short: ${q://QID28/ChoiceTextEntryValue}

Amount invested in the LCI: ${q://QID38/ChoiceTextEntryValue}

Was this client authorized to make the investment?: No

This is the summary of the call. Please put the following information in your Excel spreadsheet of communication with the clients:

Client number: ${q://QID27/ChoiceTextEntryValue}

First friend's client number: ${q://QID30/ChoiceTextEntryValue}

Did the first friend want to invest in the product? ${q://QID21/ChoiceGroup/SelectedChoices}

Was the first friend authorized to make the investment? ${q://QID25/ChoiceGroup/SelectedChoices}

Did this client (second friend) want to invest in the product? No

Amount invested in the Fundo Long-Short: 0

Amount invested in the LCI: 0

Was this client authorized to make the investment?: N/A

Chapter 3

Paying a High Price to Save?*

3.1 Introduction

It is well documented that consumers in the US could pay less in financial charges by changing the way they manage different accounts. In particular, many consumers borrow at high interest rates from their credit cards while having enough liquid assets to cover their credit card debt (Gross and Souleles (2002) and Stango and Zinman (2009b)). While this borrowing high and lending low (BHLL) behavior has usually been associated with consumer mistakes and behavioral biases[1], two recent papers suggest that the amount of financial charges that households face due to this behavior is not that high, and could be rationalized by demand for liquidity (Zinman (2007) and Telyukova (2011)). However, this evidence does not directly imply that consumers would avoid BHLL when costs are higher There is still limited evidence on whether households engage in such behavior in settings where interest rates are higher, and on whether the spread between borrowing and savings rates is a key determinant of these decisions.

This paper examines this phenomenon in a setting where borrowing costs are much higher. Specifically, we look at this phenomenon in Brazil, where checking account overdraft interest rates are, on average, equal to 8.33% per month, which corresponds to 161% per year[2], while interest rates on savings accounts are no higher than 0.6% per month.

*This chapter is co-authored with Pedro Daniel Tavares.

[1]Bertaut and Haliassos (2006), Reiter and Haliassos (2007), and Bertaut et al. (2009) present an "accountant-shopper" model, where mental accounting would explain the decisions to simultaneously hold credit card debt and liquid assets. Laibson et al. (2007) present a model where demand for a commitment device would explain credit card borrowing and investments in illiquid assets.

[2]February of 2012. Source: National Association of Executives in Finance, Administration and Account-

Therefore, the interest rate spread that households pay to borrow further from their overdrawn accounts instead of using funds from their savings accounts is around 7.7% per month (144% per year), which is about 10 times greater than the usual spread between credit card revolving rates and the return on liquid assets in the US.

Using information on checking and savings accounts for a sample of clients from a large bank in Brazil, we document the prevalence of this BHLL behavior, and calculate the amount of avoidable financial charges that households face in such a high interest rate setting. Also, by exploring variations in the overdraft interest rates, we test whether consumers are less likely to engage in such behavior when the associated costs are higher.

Our results suggest that the decision to borrow at high rates while holding low-yield liquid assets is prevalent even in a setting where the costs associated with this behavior are much higher. Around 70% of clients maintain an overdrawn account at least one day a year while having liquid assets in their savings accounts to cover at least part of this debt. However, the amount of financial charges that clients incur due to this behavior is not excessively high. On average, clients would be able to avoid approximately R\$100 (approximately US\$60) in financial charges by transferring funds from their savings to their checking accounts, which would represent less than 0.5% of their yearly earnings. The BHLL cost faced by Brazilian consumers is not much larger than the one faced by US consumers, despite the fact that the spread between borrowing and savings rates in Brazil is around 10 times larger[3]. This suggests that the high interest rates in Brazil prevent consumers from excessively engaging in this behavior.

There is also evidence that clients with lower overdraft rates are more likely to engage in this behavior. Bank employees face an overdraft interest rate on the order of 4%, compared to an interest rate on the order of 9% for non-employees. We show that employees borrow from their overdrawn checking accounts while having liquid funds more frequently than non-employees. After controlling for income, employees end up facing approximately the same in financial charges that could be avoidable by using liquid assets to cover their overdraft debt, despite the fact that they face an interest rate more than 50% lower than

ing (Anefac). Interest rates are usually stated in monthly terms in Brazil.

[3]Zinman (2007) estimates that the average U.S. household with a credit card pays US\$100 per year in financial charges to hold bank account balances instead of using them to pay down credit card debt. The estimated BHLL cost faced by Brazilian consumers is smaller than that for US consumers. Considering that the per capita GDP in the US is 4 times larger than the per capita GDP in Brazil, the BHLL cost relative to the per capita GDP in Brazil would still be far less than 10 greater than that in the U.S.

non-employees. These results suggest that although consumers engage in this behavior even when the interest rate spread is high, the interest rate spread is a key determinant in these decisions.

3.2 Economic Environment

This paper focuses on the decision to simultaneously borrow from an overdrawn checking account while holding liquid assets in a savings account. Checking account overdraft represents 32% of the volume of consumer credit borrowing in Brazil[4]. Bank clients in Brazil usually have predetermined overdraft limits on their checking accounts. If a client's balance is below zero, but still within the overdraft limit, he will be charged the overdraft interest rate on the average overdraft balance for a given month. The average overdraft interest rate is equal to 8.33% per month[5]. If he crosses the overdraft limit, then he will also be charged a penalty fee, and face a higher interest rate on the amount that exceeded the overdraft limit.

These higher overdraft interest rates make it more costly for consumers in Brazil to borrow while still possessing liquid assets to cover their debts than for consumers in the United States and in other developed countries. While, on average, the return on savings accounts is also higher than in developed countries, it is still on the order of 0.6% per month, which implies that consumers in Brazil are essentially paying around 7.7% per month on their overdraft balances in order to keep money in their savings accounts. This implies that a client who keeps a $100 balance in his overdrawn checking account while having liquid funds in his savings account would have been able to save in a year $144 in financial charges in Brazil. In contrast, considering a revolving credit card APR of 15%, a client in the US would pay $16 in financial charges to revolve $100 with his credit card while having liquid funds to cover this debt. Therefore, the BHLL cost in Brazil is about 10 times greater than that in the US and in other developed countries.

In addition to being more expensive, this BHLL behavior cannot be justified by the usual explanations for borrowing at high rates while having liquid assets. First, checking and savings accounts are closer substitutes in terms of means of payments. A client could

[4]Source: Brazilian Central Bank

[5]February of 2012. Source: National Association of Executives in Finance, Administration and Accounting (Anefac)

convert funds from his savings account to cash as easily as he can with his checking account. In contrast, credit cards and liquid assets would not be perfect substitutes in terms of means of payments, as suggested by Zinman (2007). Also, when a client transfers funds from his savings account to cover his overdraft balance, he would restore his available overdraft limit, meaning his available liquidity would not be negatively impacted. Finally, the risk of overdraft fees (Fusaro (2003)) or minimum balance requirements (Stavins (1999)) would not apply to this setting. On the contrary, a client would face a *lower* risk by transferring funds from savings account to his checking account, as this would reduce the likelihood of exceeding his overdraft limit, which would result in fees and higher interest rates.

Still, it is possible that some consumers keep their funds in savings accounts while paying checking account overdraft financial charges for rational reasons. For instance, a client could worry that the bank would reduce his overdraft limit once he has paid off his overdraft debt. In this case, covering the overdraft debt with funds from his savings account would reduce his available liquidity. In fact, the overdraft terms are renewed every month, and the bank is allowed to reduce the overdraft limits for clients at that point. In practice, the bank policy is to reduce the overdraft credit limit when a client exceeds his overdraft limit for more than 90 days. Another important potential rationale for this behavior would be strategic default (as in Lehnert and Maki (2002)). While both the checking and savings accounts are held by the same bank, the savings account is not used as collateral for the overdraft debt. Therefore, a client would be able to default on his overdraft debt, but still access the funds from his savings account.

Other reason for BHLL could be transaction costs. Presumably, transferring funds from savings accounts to checking accounts should be easy and costless. The transfer's monetary cost is zero, and a client may use any bank channel: online banking, automated teller machines, branches, and call centers. However, there might still be some non-monetary transaction costs. Finally, clients might keep an overdrawn checking account due to limited attention, as they might not keep track of their checking account balances (Grubb (2012)).

3.3 Data and Methodology

In this paper, we examine information taken from the checking and savings accounts of a sample of clients from a major Brazilian bank. Since our focus is the decision to simultaneously hold overdraft debt and liquid assets, we restrict our sample to clients who have predetermined limits on their checking accounts. We constructed a sample with $6,000$ observations, half of which consists of clients with the standard overdraft interest rate of approximately 9% per month, while the other half consists of employees of the bank, who pay the reduced overdraft interest rate of approximately 4% per month. Our information for each client includes daily checking account balances, overdraft limits, and overdraft interest rates, as well as their savings account balances from the end of each month. We have information for one entire year[6].

Using the above information, we are able to calculate the overdraft financial charges for each consumer, and estimate the amount of financial charges that each consumer could avoid by using liquid assets from his savings account to cover his overdraft balance. The actual financial charges a consumer faces in a given month (C) are defined by:

$$C = \sum_{t=1}^{T} D_t \times \frac{r}{T} + \sum_{t=1}^{T} (D_t - \bar{L}\}) \times \frac{(\bar{r} - r)}{T} + F \times 1\{B_t > \bar{L} \text{ for some t}\} \tag{3.1}$$

where D_t is the overdraft balance at day t (equal to zero if the checking account has a positive balance), T is the number of days of the specific month, r is the overdraft interest rate, $\bar{L}$ is the overdraft limit, $\bar{r}$ is the interest rate the consumer faces when he exceeds the overdraft limit, and F is the fee that the consumer faces when he exceeds the overdraft limit.

In order to estimate the total amount of avoidable financial charges, we calculate the charges that a client would face if he used funds from his savings account to cover this balance ($C^{\star}$). Unfortunately we do not have information on daily savings account balances. Instead, we consider that a client would be able to use at least the minimum between his savings balance at the beginning and at the end of the month ($B = min\{B_m, B_{m+1}\}$). In this case:

$$\begin{aligned} C^{\star} = \sum_{t=1}^{T} max\{0, D_t - B\} \times \frac{r}{T} + \sum_{t=1}^{T} (max\{0, D_t - B\} - \bar{L}\}) \times \frac{(\bar{r} - r)}{T} + \\ + F \times 1\{max\{0, D_t - B\} > \bar{L} \text{ for some t}\} \end{aligned} \tag{3.2}$$

[6]We are not allowed to disclose the year of data we are using due to the disclosure agreement with the cooperating bank.

Therefore, the amount of avoidable financial charges would be equal to $C - C^{\star}$. The opportunity cost of using funds from the savings account to cover the overdraft debt and avoid financial charges would be the foregone return on the savings account, which is around 0.6% of the amount used to cover the debt.

In order to study how the prevalence and cost of this behavior is affected by the overdraft interest rate, we compare bank clients and bank employees. The overdraft interest rate for bank employees is approximately 4% per month. While this rate is still much higher than revolving credit card interest rates in the United States, it is less than half the rates generally faced by clients. Of course, there are many inherent differences between bank clients and bank employees. In particular, bank employees have, on average, higher income and higher credit limit. To mitigate this obvious omitted variable bias, we control flexibly for income, credit limit, and number of months since the checking account was opened. In some regressions, we also condition our sample of bank clients to those who receive their salaries through the cooperating bank. Even after these controls, it might still be possible that bank employees have a better understanding of financial decisions. Also, bank employees might also face lower transaction costs to transfer funds from their savings to their checking accounts. These omitted variables would bias our estimates towards finding no effect of interest rates on the prevalence of BHLL.

3.4 Results

3.4.1 Prevalence and Cost of BHLL

We start by presenting a descriptive analysis of the prevalence and cost of this BHLL behavior in a setting where borrowing rates are approximately equal to 9% per month. In column 1 of Table 3-1, we present the proportion of clients that borrowed from their overdrawn checking accounts while having liquid funds in their savings accounts at some point in one year. This behavior is highly prevalent, and around 70% of the clients were in this situation at least one day in a period of one year. Still, it is important to understand whether this is a recurrent behavior, and whether clients end up facing significant costs by engaging in this behavior. In columns 2 and 3 of Table 3-1, we present, respectively, the mean and percentiles of the number of days that clients were in a BHLL state and the monetary cost due to this behavior. On average, clients stay around 84 days per year in a

BHLL state, which leads to around $R\$100$ (approximately US\$60) in financial charges that could be avoided. This number is actually not higher than the BHLL costs that Zinman (2007) calculates for US consumers, despite the fact that consumers in Brazil face much higher interest rates. As presented in column 4, the BHLL cost represents, on average, only around 0.5% of clients' yearly earnings. Looking at the distribution, the avoidable financial charges per year is smaller than 1.3% of clients' yearly earnings for 90% of the clients.

3.4.2 Prevalence and Cost of BHLL - Clients vs. Employees

In addition to documenting BHLL prevalence and cost among bank clients in Brazil, we are able to use variation in the overdraft interest rate between bank employees and bank clients to estimate how the decision to borrow at high interest rates while holding liquid funds in the savings account is affected by the spread between borrowing and savings rates.

Table 3-2 presents the prevalence and cost of BHLL behavior separately for bank clients and bank employees. Panel i of Table 3-3 presents the comparison in means of the relevant dependent variables, using OLS regressions without controls. Bank employees are 15 percentage points more likely to borrow from their overdrawn checking accounts while having liquid funds on their savings account at least one day per year, and stay in this position for almost 40 days more per year than bank clients. The annual BHLL cost for bank employees is significantly higher than that of bank clients, despite the fact that bank employees face much lower interest rates.

These numbers should be taken with caution, since bank employees and non-employees might differ in a variety of dimensions. In order to mitigate the problem of omitted variables, the regressions presented in panel ii of Table 3-3 include as control variables 5 order polynomials of income, overdraft limit, and number of months since checking account was opened. Although the difference in prevalence of BHLL between bank employees and bank clients is smaller after including these controls, bank employees are still more likely to borrow from their overdrawn checking accounts while having liquid funds on their savings account at least one day per year, and stay in this situation for almost 20 days more per year than bank clients. After including these controls, however, the difference in annual BHLL costs between bank clients and employees is close to zero, and not statistically significant.

In order to take into account the fact that bank employees have lower interest rates, in Table 3-4 we calculate the BHLL cost calculated as if both bank clients and bank employees had the higher overdraft interest rate of approximately 9%. This is a measure of the effect of the overdraft interest rate on the "quantity" of BHLL. As expected, even after controlling for income and credit limit, employees would face a much higher BHLL cost, and this would represent a higher fraction of their monthly earnings, if they were charged the same interest rate as bank clients. Given these estimates, the elasticity of BHLL with respect to the overdraft interest rate would be around -2.9[7].

In Table 3-5, we restrict the sample of bank clients to those who receive their salaries through the cooperating bank. This way we consider only bank clients who have a stable stream of earnings and, therefore, should be more comparable to bank employees. All results remain similar[8].

Although there are presumably other omitted variables that might bias these results, the comparison between bank clients and bank employees provides suggestive evidence that the spread between borrowing and savings rates is a key determinant of the decisions to borrow at high interest rates while having liquid funds yielding a lower return to cover this debt.

3.5 Concluding Remarks

Borrowing high and lending low (BHLL) is prevalent for many Brazilian consumers. Around 70% of the consumers face financial charges that could be avoided by transferring funds from their savings to their checking accounts. However, the high interest rates faced by Brazilian consumers seem to prevent them from engaging in this behavior recurrently and excessively. The BHLL cost faced by Brazilian consumers is no higher than what is documented in developed countries, where the spread between borrowing and saving rates is much lower. In addition, a comparison between bank clients (who have an overdraft interest rate of approximately 9%) and bank employees (who have a lower overdraft interest rate of approximately 4%) suggests that the spread between borrowing and savings rates is indeed a key determinant of these decisions. Bank employees borrow more from their overdrawn checking accounts while having liquid funds in their savings accounts, even after controlling

[7]The elasticity is calculated as $\epsilon = \frac{\beta}{\Delta r} \times \frac{\bar{r}}{\bar{Y}}$, where β is the coefficient on the employee dummy, Δr is the difference in the overdraft interest rate for employees and clients, and $\bar{Y}$ and $\bar{r}$ are the BHLL and interest rate averages.

[8]The only qualitative difference in the restricted sample regressions is the effect of bank employee on BHLL cost. In these regressions, there is a small positive effect of the bank employee dummy on the BHLL cost (significant at 5%), even after controlling for income, credit limit, and time since account was opened.

for income, credit limit, and time since the account was opened. Taken together, these results suggest that consumers are less likely to engage in this behavior when costs are higher, which implies in an upper bound on the financial problems that this behavior might entail. Still, these results would be consistent with rational models with demand for liquidity and strategic default, behavioral biases, or consumer mistakes. Consumers might see the spread between borrowing and saving rates as a "price" to keep their funds in their savings accounts. Alternatively, consumers might be more attentive and correct their decisions when the stakes are higher. Understanding the rationale behind these decisions is still an open question.

3.A Appendix: Figures and Tables

Table 3-1: **Prevalence and Cost of Borrowing High and Lending Low Behavior**

Variable:	Probability of BHLL (1)	Number of of BHLL days (2)	Yearly BHLL cost (3)	$\frac{\text{Yearly BHLL cost}}{\text{Yearly earnings}}$ (4)
Mean	0.702	83.8	103.8	0.005
25th percentile		0.0	0.0	0.000
50th percentile		42.0	2.9	0.000
75th percentile		142.0	56.4	0.003
90th percentile		248.0	233.8	0.013
95th percentile		295.0	426.6	0.024
99th percentile		348.0	1,539.4	0.086
# of observations	6000	6000	6000	5147

Notes: Column 1 shows the proportion of clients that were negative on their checking accounts while holding a positive balance on their savings accounts at least one day during a period of one year. Column 2 shows the mean and percentiles of the number of days that clients were in a BHLL position. Column 3 shows the mean and percentiles of the avoidable financial charges due to BHLL. Column 4 shows the mean and percentiles of the avoidable financial charges calculate in column 3 divided by client's yearly earnings. These results are calculating using the pooled sample of 3,000 bank clients and 3,000 bank employees. Observations are weighted by the inverse of the probability that the client was selected so that they represent the original population.

Table 3-2: **Prevalence and Cost of BHLL, for Clients ($i \approx 9\%$ per month) and Employees ($i \approx 4\%$)**

Variable:	Probability of BHLL		Number of of BHLL days		Yearly BHLL cost		$\frac{\text{Yearly BHLL cost}}{\text{Yearly earnings}}$	
	Client	Employee	Client	Employee	Client	Employee	Client	Employee
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mean	0.700	0.851	83.2	120.9	102.9	162.7	0.005	0.003
25th percentile			0.0	13.5	0.0	0.0	0.000	0.000
50th percentile			41.0	92.0	2.8	31.5	0.000	0.001
75th percentile			142.0	210.0	55.6	174.1	0.003	0.004
90th percentile			247.5	297.0	228.0	476.7	0.013	0.010
95th percentile			294.5	326.0	422.7	769.6	0.025	0.016
99th percentile			348.0	360.5	1,505.4	1,539.9	0.091	0.030
# of observations			3,000	3,000	3,000	3,000	2,702	2,431

Notes: This table shows the prevalence and cost that consumers faced for being negative on their checking account while holding positive balances on their savings accounts, separately for clients (who have an overdraft interest rate of approximately 9% per month) and bank employees (who have an overdraft interest rate of approximately 4%). See footnote of Table 3-1 for the definition of the corresponding variables.

Table 3-3: **Effect of Employee Status on BHLL - Regression Results**

Variable:	Probability of BHLL (1)	Number of of BHLL days (2)	Yearly BHLL cost (3)	$\frac{\text{Yearly BHLL cost}}{\text{Yearly earnings}}$ (4)
		i. No controls		
Employee dummy	0.151***	37.7***	59.8***	-0.002***
	(0.011)	(2.7)	(12.1)	(0.000)
Number of observations	6000	6000	6000	5133
		ii. Include controls		
Employee dummy	0.076***	19.0***	-28.1	-0.001***
	(0.012)	(3.0)	(17.9)	(0.000)
Number of observations	5873	5873	5873	5133
Mean for clients	0.700	83.2	102.9	0.005

Notes: This table shows regression estimates comparing the BHLL prevalence and cost for clients and bank employees. Omitted category is non-employees, who had an overdraft interest rate of approximately 9%. Regressions in panel i do not include controls. Regressions in panel ii include a 5 order polynomials on income, overdraft limit, and number of months since checking account was open as controls. Robust standard errors in parenthesis. See footnote of table 3-1 for the definition of the corresponding variables. * significant at 10%; ** significant at 5%; *** significant at 1%

Table 3-4: **Effect of Employee Status on BHLL - Normalized**

Variable:	Normalized yearly BHLL cost	$\frac{\text{Normalized yearly BHLL cost}}{\text{Yearly earnings}}$
	(1)	(2)
	i. No controls	
Employee dummy	316.1***	0.003***
	(18.7)	(0.000)
Number of observations	6000	5133
	ii. Include controls	
Employee dummy	162.5***	0.004***
	(20.3)	(0.000)
Number of observations	5873	5133
Mean for clients	101.5	0.005

Notes: This table shows regression estimates comparing the BHLL costs and relevance for clients and bank employees. BHLL costs are normalized to take into account that bank clients and bank employees have different interest rates. BHLL normalized costs are calculated considering that everyone had an interest rate of 8.75%. Omitted category is non-employees. Regressions in panel i do not include controls. Regressions in panel ii include a 5 order polynomials on income, overdraft limit, and number of months since checking account was open as controls. Robust standard errors in parenthesis. * significant at 10%; ** significant at 5%; *** significant at 1%

Table 3-5: **Effect of Employee Status on BHLL - Restricted Sample**

Variable:	Probability of BHLL	Number of of BHLL days	Yearly BHLL cost	Yearly BHLL cost / Yearly earnings	Normalized BHLL cost	Normalized BHLL cost / Yearly earnings
	(1)	(2)	(3)	(4)	(5)	(6)
			i. No controls			
Employee dummy	0.065***	27.7***	90.2***	-0.001*	352.9***	0.005***
	(0.013)	(3.4)	(8.9)	(0.000)	(17.3)	(0.000)
N	4192	4192	4192	3639	4192	3639
			ii. Include controls			
Employee dummy	0.045***	17.1***	23.9**	0.000	196.1***	0.005***
	(0.014)	(3.8)	(10.3)	(0.000)	(14.8)	(0.000)
N	4192	4192	4192	3639	4192	3639
Mean for clients	0.802	95.9	75.8	0.004	74.5	0.004
	(0.000)	(0.0)	(0.0)	(0.000)	(0.0)	(0.000)

Notes: This table shows regression estimates from Tables 3-3 and 3-4 using a restricted sample of clients who receive their salaries through the cooperating bank. See footnote of Tables 3-1 and 3-4 for the definition of the corresponding variables. * significant at 10%; ** significant at 5%; *** significant at 1%

Bibliography

Abel, Andrew B., "Asset Prices under Habit Formation and Catching Up with the Joneses," *American Economic Review*, 1990, *80* (2), 38–42.

Abreu, Margarida and Victor Mendes, "Financial Literacy and Portfolio Diversification," *Quantitative Finance*, 2010, *10* (5), 515–528.

Angrist, Joshua D. and Alan B. Krueger, "The Effect of Age at School Entry on Educational Attainment: An Application of Instrumental Variables with Moments from Two Samples," *Journal of the American Statistical Association*, 1992, *87* (418), 328–336.

Asch, Salomon E., "Effects of Group Pressure on the Modification and Distortion of Judgments," in Harold Guetzkow, ed., *Groups, Leadership and Men*, Pittsburgh, PA: Carnegie Press, 1951, pp. 177–190.

Attanasio, Orazio P., Pinelopi Koujianou Goldberg, and Ekaterini Kyriazidou, "Credit Constraints In The Market For Consumer Durables: Evidence From Micro Data On Car Loans," *International Economic Review*, 05 2008, *49* (2), 401–436.

Avery, Christopher and Peter Zemsky, "Multidimensional Uncertainty and Herd Behavior in Financial Markets," *American Economic Review*, 1998, *88* (4), 724–748.

Banerjee, Abhijit V., "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 1992, *107* (3), 797–817.

_ , Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson, "The Diffusion of Microfinance," August 2011. MIT Department of Economics Working Paper.

Barr, M.S., S. Mullainathan, and E. Shafir, "Behaviorally Informed Financial Services Regulation," *New America Foundation White Paper*, 2008.

Becker, Gary S., "A Note on Restaurant Pricing and Other Examples of Social Influences on Price," *Journal of Political Economy*, October 1991, *99* (5), 1109–1116.

Bertaut, Carol C. and Michael Haliassos, "Credit Cards: Facts and Theories," CFS Working Paper Series 2006/19, Center for Financial Studies September 2006.

_ , _ , and Michael Reiter, "Credit Card Debt Puzzles and Debt Revolvers for Self Control," *Review of Finance, Forthcoming*, 2009.

Bertrand, Marianne and Adair Morse, "Information Disclosure, Cognitive Biases and Payday Borrowing," *Journal of Finance, forthcoming*, 2011.

—, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman, "What's Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment," *Quarterly Journal of Economics*, 2010, *125* (1), 263–305.

—, Erzo F. P. Luttmer, and Sendhil Mullainathan, "Network Effects and Welfare Cultures," *Quarterly Journal of Economics*, August 2000, *115* (3), 1019–1055.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian, *The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States*, University of Chicago Press, June

—, —, —, —, and Katherine L. Milkman, "The Effect of Providing Peer Information on Retirement Savings Decisions," August 2011. NBER Working Paper 17345.

Bikhchandani, Sushil and Sunil Sharma, "Herd Behavior in Financial Markets: A Review," 2000. IMF Working Paper No. 00/48.

—, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 1992, *100* (5), 992–1026.

Bobonis, Gustavo and Frederico Finan, "Neighborhood Peer Effects in Secondary School Enrollment Decisions," *Review of Economics and Statistics*, 2009, *91* (4), 695–716.

Burnkrant, Robert E. and Alain Cousineau, "Informational and Normative Social Influence in Buyer Behavior," *Journal of Consumer Research*, 1975, *2* (3), 206–215.

Cai, Hongbin, Yuyu Chen, and Hanming Fang, "Observational Learning: Evidence from a Randomized Natural Field Experiment," *American Economic Review*, 2009, *99* (3), 864–882.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini, "Measuring the Financial Sophistication of Households," *American Economic Review*, May 2009, *99* (2), 393–398.

Campbell, John Y. and John H. Cochrane, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 1999, *107* (2), 205–251.

Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez, "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction," 2010. NBER Working Paper 16396.

— and Laura Giuliano, "Peer Effects and Multiple Equilibria in the Risky Behavior of Friends," May 2011. NBER Working Paper 17088.

Celen, Bogachan and Shachar Kariv, "Distinguishing Informational Cascades from Herd Behavior in the Laboratory," *American Economic Review*, 2004, *94* (3), 484–497.

Chari, Varadarajan V. and Patrick J. Kehoe, "Financial crises as herds: overturning the critiques," *Journal of Economic Theory*, 2004, *119* (1), 128–150.

Clark, Andrew E. and Andrew J. Oswald, "Comparison-concave Utility and Following Behaviour in Social and Economic Settings," *Journal of Public Economics*, 1998, *70* (1), 133–155.

Conley, Timothy G. and Christopher R. Udry, "Learning about a New Technology: Pineapple in Ghana," *American Economic Review*, March 2010, *100* (1), 35–69.

Cooper, David J. and Mari Rege, "Misery loves company: Social regret and social interaction effects in choices under risk and uncertainty," *Games and Economic Behavior*, September 2011, *73* (1), 91–110.

Dahl, Gordon B., Katrine V. Loken, and Magne Mogstad, "Peer Effects in Program Participation," June 2012. UC-San Diego Working Paper.

Day, George S. and William K. Brandt, "Consumer Research and the Evaluation of Information Disclosure Requirements: The Case of Truth in Lending," *The Journal of Consumer Research*, 1974, *1* (1), 21–32.

DellaVigna, Stefano, "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, June 2009, *47* (2), 315–72.

DeMarzo, Peter M., Ron Kaniel, and Ilan Kremer, "Diversification as a Public Good: Community Effects in Portfolio Choice," *The Journal of Finance*, 2004, *59* (4), 1677–1716.

_ , _ , and _ , "Relative Wealth Concerns and Financial Bubbles," *Review of Financial Studies*, 2008, *21* (1), 19–50.

Duflo, Esther and Emmanuel Saez, "The Role of Information and Social Interactions in Retirement Plans Decisions: Evidence from a Randomized Experiment," *Quarterly Journal of Economics*, 2003, *118* (3), 815–842.

_ , Pascaline Dupas, and Michael Kremer, "Peer Effects, Teacher Incentives, and the Impact of Tracking: Evidence from a Randomized Evaluation in Kenya," *American Economic Review*, 2011, *101* (5), 1739–1774.

Durlauf, Steven N., "Neighborhood Effects," *Handbook of Regional and Urban Economics*, 2004, *4*.

Festinger, Leon, "A Theory of Social Comparison Processes," *Journal: Human Relations*, 1954, *7*, 117–140.

Fliessbach, Klaus, Bernd Weber, Peter Trautner, Thomas J. Dohmen, Uwe Sunde, Christian E. Elger, and Armin Falk, "Social Comparison Affects Reward-Related Brain Activity in the Human Ventral Striatum," *Science*, 2007, *318* (5854), 1305–1308.

Foster, Andrew D. and Mark R. Rosenzweig, "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy*, December 1995, *103* (6), 1176–1209.

Fusaro, Marc, "Consumers' Bank Choice and Overdraft Volume: An EmpiricalStudy of Bounce Protection Programs," *Working Paper*, 2003.

Gali, Jordi, "Keeping Up with the Joneses: Consumption Externalities, Portfolio Choice, and Asset Prices," *Journal of Money, Credit and Banking*, 1994, *26* (1), 1–8.

Goetzmann, William N. and Alok Kumar, "Equity Portfolio Diversification," *Review of Finance*, 2008, *12*, 433–463.

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

Grubb, Michael D., "Consumer Inattention and Bill-Shock Regulation," *Available at SSRN: http://ssrn.com/abstract=1983518*, 2012.

Guryan, Jonathan, Kory Kroft, and Matthew J. Notowidigdo, "Peer Effects in the Workplace: Evidence from Random Groupings in Professional Golf Tournaments," *American Economic Journal: Applied Economics*, 2009, *1* (4), 34–68.

Harrison, Glenn W. and John A. List, "Field Experiments," *Journal of Economic Literature*, December 2004, *42* (4), 1009–1055.

Hong, Harrison, Jeffrey Kubik, and Jeremy Stein, "Thy Neighbor's Portfolio: Word-of-Mouth Effects in the Holdings and Trades of Money Managers," *The Journal of Finance*, 2005, *60*, 2801–2824.

Huang, Dandan and Wei Tan, "Estimating the Demand for Credit Card: A Regression Discontinuity Approach," *SSRN eLibrary*, 2009.

Instefjord, Norvald, Jerry Coakley, and Zhe Shen, "The Winner's Curse and Lottery-Allocated IPOs in China," February 2007. University of Essex Working Paper.

Ivkovic, Zoran and Scott Weisbenner, "Information Diffusion Effects in Individual Investors' Common Stock Purchases: Covet Thy Neighbors' Investment Choices," *Review of Financial Studies*, 2007, *20* (4), 1327–1357.

Iyengar, S., G. Huberman, and W. Jiang, *How much choice is too much? Contributions to 401(k) retirement plans*, Oxford University Press, 2004.

Juster, F. Thomas and Robert P. Shay, "Consumer Sensitivity to Finance Rates: An Empirical and Analytical Investigation," June 1964, (just64-2).

Karlan, Dean S. and Jonathan Zinman, "Credit Elasticities in Less-Developed Economies: Implications for Microfinance," *American Economic Review*, June 2008, *98* (3), 1040–68.

Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz, "Experimental Analysis of Neighborhood Effects," *Econometrica*, 2007, *75*, 83–119.

Kremer, Michael and Edward Miguel, "The Illusion of Sustainability," *Quarterly Journal of Economics*, August 2007, *122* (3), 1007–1065.

Laibson, David, Andrea Repetto, and Jeremy Tobacman, "Estimating Discount Functions with Consumption Choices over the Lifecycle," Technical Report 2007.

Lehnert, Andreas and Dean M. Maki, "Consumption, debt and portfolio choice: testing the effect of bankruptcy law," Finance and Economics Discussion Series 2002-14, Board of Governors of the Federal Reserve System (U.S.) 2002.

Li, Geng, "Information sharing and stock market participation: evidence from extended families," 2009. Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, 2009-47.

Luttmer, Erzo F. P., "Neighbors as Negatives: Relative Earnings and Well-Being," *Quarterly Journal of Economics*, 2005, *120* (3), 963–1002.

Manski, Charles F., "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies*, 1993, *60* (3), 531–542.

Mas, Alexandre and Enrico Moretti, "Peers at Work," *American Economic Review*, 2009, *99* (1), 112–145.

Moretti, Enrico, "Social Learning and Peer Effects in Consumption: Evidence from Movie Sales," *Review of Economic Studies*, 2011, *78* (1), 356–393.

Parker, George G C and Robert P Shay, "Some Factors Affecting Awareness of Annual Percentage Rates in Consumer Installment Credit Transactions," *Journal of Finance*, March 1974, *29* (1), 217–25.

Reiter, Michael and Michael Haliassos, "Credit Card Debt Puzzles," *SSRN eLibrary*, 2007.

Roxburgh, Charles, "Debt and deleveraging : The global credit bubble and its economic consequences," *McKinsey Global Institute*, 2010, (January).

Sacerdote, Bruce, "Peer Effects with Random Assignment: Results for Dartmouth Roommates," *Quarterly Journal of Economics*, May 2001, *116* (2), 681–704.

Shafir, E., I. Simonson, and A. Tversky, "Reason-based choice," *Cognition*, 1993, *49*, 11–36.

Shue, Kelly, "Executive Networks and Firm Policies: Evidence from the Random Assignment of MBA Peers," April 2012. University of Chicago, Booth School of Business.

Stango, Victor and Jonathan Zinman, "Exponential Growth Bias and Household Finance," *Journal of Finance*, December 2009, *64* (6), 2807–2849.

_ and _ , "What Do Consumers Really Pay on Their Checking and Credit Card Accounts? Explicit, Implicit, and Avoidable Costs," *American Economic Review*, 2009, *99* (2), 424–29.

_ and _ , "Fuzzy Math, Disclosure Regulation, and Market Outcomes: Evidence from Truth-in-Lending Reform," *Review of Financial Studies*, 2011, *24* (2), 506–534.

Stavins, Joanna, "Checking accounts: what do banks offer and what do consumers value?," *New England Economic Review*, 1999, (Mar), 3–14.

Telyukova, Irina, "Household Need for Liquidity and the Credit Card Debt Puzzle," *SSRN eLibrary*, 2011.

Zinman, Jonathan, "Household Borrowing High and Lending Low Under No-Arbitrage." 2007.