Learning in the Limit: Income Inference from Credit Extension

Xiao Yin*

November, 2022

Abstract

Combining a randomized controlled trial with bank account and survey data, I show that credit-limit extensions significantly increase consumer expectations about their future income. A one-dollar increase in credit limit raises consumer income expectations over the next six months by 40 cents and total consumption by 34 cents. The expectation changes explain around 34% of the total spending responses to credit limit extensions. The results show that consumers infer information from lenders' credit-supply decisions, and this learning behavior impacts consumers' economic decision-making greatly.

Keywords: Consumption, MPC, MPB, Credit Supply, Field Experiments, Income Expectations.
JEL codes: D14, D15, D91, E21, E51, G21.

*Yin: Haas School of Business, University of California, Berkeley, yinxiao@berkeley.edu. I deeply appreciate the valuable help and comments from Ulrike Malmendier, and David Sraer, Yuriy Gorodnichenko, Michael Weber, and Matteo Benetton. I also appreciate the valuable comments from Stefano DellaVigna, Amir Kermani, Chen Lian, Peter Maxted, Emi Nakamura, and Annette Vissing-Jørgensen. I am thankful for valuable comments from UC Berkeley Macro Lunch Seminar, UC Berkeley Finance Seminar, and UC Berkeley Financial Economics Seminar. The field experiment was conducted within the guidelines of the UC Berkeley IRB-approved human subjects protocol (CPHS Protocol Number: 2019-11-12703), and is registered at the AEA RCT Registry (AEARCTR-0010315). All errors are my own.

I Introduction

Credit limit is one of the most important factors that affect household consumption-saving decisions, as it underpins how much consumers can borrow for consumption in the short run. Across many advanced economies, more than one-third of the consumers have positive consumption debt[1]. As predicted by the workhorse buffer-stock model, variations in credit limit should only have small impacts on total spending only except for those that are close to be borrowing constrained. However, existing literature documents a large spending response to changes in credit limit, and this is the case even for consumers that are far from their borrowing limits[2]. Hence, the micro-level mechanism of why credit limit extensions affect consumer spending remains an open question.

Standard estimation of spending responses to credit limit extension assumes that consumers view changes in credit limit as random shocks only to their budget constraints. But the nature of banks' decision-making processes indicates credit extensions are rarely a result of random decisions. In particular, banks use sophisticated statistical models, calibrated on large samples with rich cross-sectional variations, to infer borrowers' spending and repayment patterns. Such processes of systematic inference could generate additional information about consumers' characteristics that the consumers are unaware of.[3] Incorporating information generated by banks, limit changes serve as signals about consumers' future economic activities. In this paper, I study the effects of credit-limit changes on consumption through changing consumer expectations.

The data in this study comes from a large commercial bank that operates nationally in China and is among the top 10 commercial banks in the country, as ranked by total assets. I rely on a large-scale randomized controlled trial (RCT) run by the bank to

[1]See Gross and Souleles (2002), Zinman (2009), Fulford (2015) for examples in the US, Vihriälä (2020) for Finland, and Gathergood and Olafsson (2022) for Iceland. In addition, the calculation in this study shows that more than 40% of the consumers in China have positive credit card debt.

[2]See Gross and Souleles (2002), Agarwal et al. (2017), D'Acunto et al. (2020), and Aydin (2022) for some examples.

[3]Brunnermeier et al. (2021) analyze the context in which big data and AI allow insurers to infer statistical information and thereby reverse the information advantage from the insuree to the insurer. In Mulder (2022), risk reclassification due to better statistical models significantly changes insurees' behaviors through a possible information channel. In addition, using administrative data from the Chase Bank, Farrell et al. (2020) show that machine learning techniques could greatly improve the accuracy of predicting household income.

retrieve the causal responses of consumption and debt toward offers of credit-card-limit extensions. In the setting, the bank had planned to increase the credit card limits of a group of customers, with the amount of increase determined by its proprietary rule. The RCT then postponed the limit-change offers for a randomly-selected group of customers (the control group) for six months. The rest of the customers (the treatment group) received the pre-determined limit-change offers, and would then decide whether to accept the offers.

To identify the effects of the credit extensions on consumers' beliefs, I rely on two surveys: a pre-experiment survey and a post-experiment survey conducted one week before and one week after the experiment, respectively. Using the differences in the answers to the same questions in both the post- and pre-experiment surveys, I identify the effects of the experiment on consumer beliefs. The identification assumption is that during the week on both sides of the experiment, the credit-limit changes are the only systematic variations in the consumers' economic characteristics. The survey design has a high-frequency property that enables me to isolate the effects of the credit-limit extensions.

I first document that news about credit-limit extensions has a considerable positive effect on consumers' expectations of future income. Specifically, I find a \$1 increase in credit limit offered by the bank on average increases consumers' expectations about their income over the next six months by \$0.40. Given an average credit-limit increase of \$1,862, consumers on average revise up their income expectations by around \$744, which is equivalent to 3.88% of the participants' average annual income. The effects are significant for both those who accept and those who do not accept the limit changes. I find that for those who accept the limit increase, each \$1 increase in the credit limit increases six-month income expectations by \$0.48. Even for those who do not accept the offer, each \$1 increase in the credit limit increases income expectations by \$0.24. The findings show consumers infer future income growth from changes in credit limits, thus suggesting an income-inference channel through which credit extensions affect consumption.

Recent studies show that large variations in individual beliefs may have a limited relationship with their actions (Ameriks et al., 2020; Giglio et al., 2021). I continue to explore whether the updates in the beliefs regarding future income are simply "cheap talk"

or whether consumers indeed *act* on them. First, I study the effects of credit extensions on consumers' total borrowing, unconditional on the changes in their income expectations. I find a \$1 higher credit limit offered by the bank increases consumers' outstanding debt by \$0.127. For those who accept (do not accept) the offer, a \$1 higher credit limit increases consumers' outstanding debt by \$0.175 (\$0.042) six months after the experiment.

In the conventional buffer-stock model, credit-limit shocks affect total consumption only by relaxing consumers' precautionary motives, as a result of a lower possibility of binding borrowing constraints. In the presence of income inference, credit supply also affects spending by changing consumers' beliefs. I continue to decompose the changes in borrowing into an income-inference channel and a residual channel. Using a similar strategy as in Kling et al. (2007), Abdulkadiroğlu et al. (2014), and Kline and Walters (2016), I find that, conditional on the changes in income-growth expectations triggered by the random assignment, the effects of limit-extension news on borrowing decrease by around 35% for those who accept the offer. For those who do not accept the offer, the effects become insignificant. These results suggest that the income-inference channel is an economically important channel in explaining the borrowing responses to credit-limit extensions.

Using consumers' transaction-level data, I analyze the effects of relaxed credit limits on non-debt-financed spending. Similar to the findings on borrowing, I find that for a \$1 increase in the credit limit, non-debt-financed spending increases by \$0.276 (\$0.108) in total over a six-month horizon for those who accept (do not accept) the offer. After controlling for the changes in income-growth expectations, the estimate decreases to \$0.174 for those who accept the offers. For those who do not accept the offers, the estimate decreases to \$0.042 and becomes insignificant. Therefore, the weight of the income-inference channel is about the same for both debt-financed and non-debt-financed spending.

Although a change in consumer expectations after the experiment indicates consumers learn from credit-limit extensions, this finding does not provide direct evidence on whether the expectation changes are consistent with Bayesian learning. Credit-limit changes can affect consumer beliefs for two possible reasons. First, banks can extract

independent information about consumers' future earning potential. Consumers then rationally incorporate such ability and revise up their expectations in response to a credit extension. Second, credit supply is uncorrelated with future income growth; nonetheless, over-confident consumers believe credit extensions signal higher future income growth anyway. For credit-limit changes to be a signal of future income, bank credit supply should be correlated with consumer ex-post income changes. Consistent with credit limit being a signal of consumer future income, I find a \$1 higher credit limit is associated with \$0.309 higher realized income over a six-month horizon. However, the sensitivity of expectations to limit extension is much larger. Specifically, a \$1 increase in the credit limit increases participants' income expectations in the treatment group by \$0.396. Therefore, consumers tend to become over-optimistic about their future earnings after credit-expansion shocks.

Over-optimistic beliefs induce overspending and overborrowing. The resulting higher leverage increases default risks. Using 60-day delinquency as a default indicator, I find a \$1,000 higher credit limit offered by the bank increases the default rate of the consumers who accept the offers by 0.176 percentage points over the six months after the experiment. With a pre-experiment average default rate of 2.44 percentage points, this amount is equivalent to a 7.21% increase. After controlling for changes in income expectations, credit-limit extensions have marginally negative and insignificant effects on the default rate. Therefore, limit extensions do not seem to have significant effects on default risk beyond inducing over-optimistic income expectations.

Recent decades have witnessed a proliferation of artificial intelligence (AI) technologies that enable banks to have more accurate statistical inference and forecasting abilities. Incorporating the income-inference channel, more accurate statistical inferences of banks would potentially have large effects on the equilibrium credit supply and spending. At the end of the analyses, I structurally estimate a life-cycle model to study the equilibrium of credit supply and consumption when banks receive more precise signals about consumers' future income. I find that for a 25% increase in bank signal precision, equilibrium credit supply and consumption increase by 14.11% and 4.11%, respectively. Therefore, when advances in information technology enable banks to extract more precise signals about household future income, credit supply, and spending would increase significantly.

Related Literature This paper mainly contributes to three strands of literature. First, it contributes to the study of borrowing limits and consumption (Zeldes, 1989; Ludvigson, 1999; Gross and Souleles, 2002; Agarwal et al., 2017; Guerrieri and Lorenzoni, 2017; Chava et al., 2020; D'Acunto et al., 2020; Gross et al., 2020; Aydin, 2022). A recent major progress is Aydin (2022), which provides a clean empirical estimation of marginal propensity to borrow using an RCT in Turkey. Although previous literature relies on the mechanisms of credit limits affecting consumer budget constraints, the effect of credit expansions on consumer spending through changing consumer beliefs is still an open question. The lack of evidence lies in the difficulties of combining an RCT with both observational and expectations data. This paper aims to fill this gap by combining an RCT with bank-account data and high-frequency surveys, facilitating direct testing of the effects of an exogenous shock to credit constraints on consumers' beliefs and how the changes in beliefs affect households' consumption-debt decisions.

This paper also contributes to the rich literature on the marginal propensity to consume (MPC) out of a one-time wealth transfer (Parker et al., 2013; Fuster et al., 2020; Kueng, 2018; Olafsson and Pagel, 2018; Baugh et al., 2021; Fagereng et al., 2021).[4] As shown by Guerrieri and Lorenzoni (2017) and Aydin (2022), the estimates of MPB while holding beliefs fixed provide a lower bound of the MPC out of a one-time wealth transfer in the short run. The income-inference channel of credit expansion changes consumption by changing expectations about consumers' future income in the near future. Therefore, the estimation of MPC to a one-time wealth transfer in the short run based on credit-limit adjustment needs to consider the changes in consumers' beliefs. By directly controlling for the changes in consumers' expectations around a credit-expansion event, this paper identifies an exogenous shock that affects consumption and debt decisions through changing credit limits only, therefore providing a clean estimate of the lower bound of MPC to a one-time wealth shock.

Lastly, this paper contributes to the literature on using random or quasi-random variations for identification (See Deaton 2010, List 2011, Bouguen et al. 2019, and Duflo 2020 for some recent discussions). Recently, Chassang et al. (2015), Fudenberg and Levine

[4]See Attanasio and Weber (2010) and Jappelli and Pistaferri (2010) for a survey before 2010.

(2022), and Hennessy and Chemla (2022) propose that treatment effects estimated with random or quasi-random variations can be biased if subjects' beliefs are not sufficiently accounted for. I contribute to this literature by providing direct evidence showing that, in the markets with private information, random variations may not deliver the identification researchers expect. This is especially likely when the subjects in the experiments believe the variations in the forcing variables incorporate signaling effects.

II Methodology

A. Data and Institutional Environment

The data for this study comes from a large commercial bank in China. The bank operates nationally and is among the top 10 commercial banks in the country, as ranked by total assets. In 2020, the bank's total assets amounted to over $1 trillion.

The credit cards considered in this study are very similar to those in other countries. In general, each credit card is assigned a credit limit, and consumers can accumulate balances smaller than this limit every month and use the card as a payment method. Consumers earn different levels of discounts and cash back for purchasing certain types of goods or services, depending on the bank's current promotional strategy. At the end of each billing cycle, a minimum repayment is required on the credit cards (usually 10% of the current outstanding balance). Above this amount, consumers can choose to repay any proportion of the current outstanding balance. Consumers who repay all accumulated balances do not incur any interest costs and enjoy the rewards from the cash back and transaction discounts. For the unpaid amounts, the debt is carried over to the next billing cycle with a daily interest rate of five basis points[5].

A recent report shows credit card use in China has grown significantly since 2015.[6] The amount of credit card transactions in the top 14 commercial banks in China has grown from $2.61 trillion in 2015 to $5.63 trillion in 2019. At the same time, the total number of credit cards in the country has increased from 0.47 billion in 2015 to 0.78 billion in 2019.

[5]The daily interest rates on credit cards are five basis points for all the customers in the bank before 2021.

[6]See here for the report.

Among the different types of non-durable household loans, credit card debt accounts for the largest proportion in China: around 51.3% of non-durable household debt comes from credit cards. For comparison, credit card debt accounted for about 32% of non-durable household debt in the US in 2019.

B. Sample Restrictions

Consumers usually have multiple bank accounts. Therefore, single-provider transaction-level datasets raise concerns about the completeness of the data in covering the full extent of consumers' spending and cash savings. To alleviate these concerns, I follow recent work using single-provider transaction-level data (e.g., see Ganong and Noel (2019)) and impose two restrictions on the accounts in the empirical analysis to capture the consumers who are most likely to use the bank as their primary banking institution.

First, I include only those consumers in the sample whose bank accounts have at least 15 monthly outflow transactions on average during the sample period. An outflow is any debit from a checking account, including cash withdrawals, electronic payments, or debit card transactions. Imposing this criterion reduces the original sample by approximately 35%. The second restriction is that the consumers' income has to be identified by the bank by observing regular inflows to the checking accounts, which amounts to an additional drop of about 10% in the total observations. [7]

C. Measuring Income and Spending

The transaction-level data allow direct measurement of consumers' income inflows and spending outflows. In terms of income, I follow the steps the bank uses, which identify individual income according to a classification rule of regular inflows. The bank classifies income into three main categories: salary, business cash flows, and financial investment.

Salary is defined as the regular monthly income flow if the consumer declares that they work as an employee. The bank calculates this number in one of two ways. First, if

[7] I focus on the consumers with non-missing income information for the main analysis. In addition, the experiment also selects a small group of individuals whose income information is not observed by the bank. I use this group of individuals for robustness checks. See section IV.G for more details.

income is paid as a direct deposit from the consumer's employer to this bank, the number is directly labeled as salary in the bank's system. Otherwise, the bank can identify monthly income if the consumer's social security insurance is paid through this bank, which is usually a fixed portion of the consumer's income.[8]

Income from business operations is the difference between total inflow and total outflow when these transactions are categorized as business operations. This category is usually the main source of income for self-employed individuals. By contrast, for income from financial investment, the bank computes it as the difference between the total inflow and the total outflow from an investment account with the financial institutions.

When all the incomes in our sample are aggregated, the split of the three components come out to be 66.74% from salary, 23.37% from business operations, and 9.89% from financial investment. To verify these figures are accurately computed at the individual level, I match the income computed at the consumer-year level from the bank to the individual-level data from the tax administrative agency. The results of this comparison are shown in Panel A of Figure A.1. in the Online Appendix. A regression between the two measures gives an R^2 of 0.89. This finding indicates the income measure from the bank is of very high quality.

Debt is the outstanding interest-incurring balance on credit cards. For the measurement of spending, I calculate the consumer's monthly total consumption as the sum of all purchasing transactions. When consumers make purchases from credit cards in the current billing cycle, they can either repay all or a proportion of their accumulated balances. Therefore, total consumption defined in this way also consists of newly accumulated debts that are not repaid at the end of the current billing cycle. In the analyses in this study, I also define non-debt-financed spending as the difference

[8]In China, social security payments have six components: five types of insurance and a housing provident fund. The types of insurance are paid with a fixed proportion of workers' monthly income. One such insurance is retirement savings insurance, which is similar to the retirement savings plan in the US. For a monthly income of 5,000 CNY, the monthly contribution is 8%. However, the income base for social security is usually capped at the two tails of the income distribution. The numbers differ for different geographic areas but are usually at 30% and 300% or 40% and 400% of the previous year's average income in that area. Therefore, for those who earn more than 300% of the last year's average income in the area, the total monthly payment is equal to $8\% \times 300\% \times \bar{Y}$, where $\bar{Y}$ is the previous year's average income in the area. However, the uncapped distribution is wide enough to cover most of the workers in China. In the analysis, I remove the consumers in the capped region from the final sample. Removing these customers drops the number of participants in the online sample by 9.6%.

between total consumption and newly accumulated debt. Thus, for the analysis of total spending, cumulative changes in non-debt-financed spending and those in debt are mutually exclusive.

D. Experimental Design

The procedure of the experiment was as follows:

1. At the end of March 2020, the bank selected a group of consumers and decided on a new credit limit that was higher than this group of customers' current credit card limit, based on the bank's own proprietary rules.[9]

   A random sample of approximately 16,000 consumers was selected from the group of consumers as the potential subjects in this study. Of the selected subjects, 7,000 consumers were randomly selected to form the control group, and the rest formed the treatment group.

2. From April 3 to April 7,[10] all the participants were given a survey asking about their expectations[11]. At the top of the front page of the survey, the participants were informed that the survey would be used to study the expectations and preferences of representative credit card holders in China and that the information would be used for only scientific research purposes. No one was informed about whether was an experiment.

   The survey was sent in two ways. First, it was designed in a Chinese survey app, and the link to the survey was sent to the participants using WeChat and text messages. Second, the questionnaire was printed out and handed out in person by the bank's staff to participants' companies if the companies paid their staff by depositing the salaries at the bank. Participants could participate through either method. After completing the survey, each participant received a gift worth approximately $2.50.

[9] Section D in Online Appendix describes the bank's credit supply rule.

[10] In China, COVID induced a nation-wide lockdown starting at the end of January. However, most areas turned to relatively normal conditions in early March. Wuhan was the latest for which the lockdown policy was removed, and the date was April 8, 2020. Therefore, COVID was expected to have a small effects on the study here.

[11] See Online Appendix A for the survey in English.

3. Within one week after step 2 (April 5–April 12), the participants in the treatment group were informed about the opportunity to increase their credit card limits to the amount offered by the bank. They could decide to either accept or ignore the offer within one week of the extension. The control-group offers were postponed to the beginning of October 2020.

4. Within one week after step 3 (April 12–April 19), all participants received a survey that was nearly the same as that in step 2 but with some slight changes.[12] A randomly chosen 15% of the participants were shown the following information at the top of the survey:

   To test their business strategies, banks often randomly select some people to have a change in their credit card limits and see how they change their spending.

   After answering the questionnaire, each participant received a gift worth approximately $2.

The setting is similar to that in Aydin (2022). In essence, the RCT temporarily pauses the internal underwriting process for a random subset of pre-selected customers for a lender-initiated credit limit increases. Therefore, the RCT only randomly prevents a group of consumers from receiving the pre-determined amout of limit increase. Within the control and treatment groups, the amount of limit change is not random and is based on the bank's reevaluation. Apart from the supplementation of the surveys, a key difference is that the change in credit limit is offered but not applied to these customers here. Therefore, after being noticed with the offers, consumers in this study could choose whether to accept the offers.

In the end, for those not shown the information treatment, 5,363 participants completed both surveys. Of those, 3,355 were from the treatment group. In addition, 1,983 participants completed the first survey but not the second. Of those, 1,032 were from the treatment group. The number of participants who completed the second survey but not the first was 483. Of those, 302 were from the treatment group. The rest did not complete either of the two surveys. Additionally, of those shown the information treatment, 1,045 completed both surveys.

[12] See the survey in Appendix for the changes.

E. Summary Statistics

Table 1 gives the summary statistics. Panel A summarizes the participants in the control group, and Panel B summarizes those in the treatment group. All level variables are converted to US dollars for ease of comparison with existing literature. Among the given offers, about 70% are accepted.[13] The average age of the participants is around 38 years old; their average annual income is around $19,000; and their average total savings is $23,000. The average credit limit from the bank is approximately $6,500, and the average total credit limit across all banks is around $13,000. The average outstanding interest-incurring debt is about $950 and approximately $2,400 conditional on holding a positive amount of debt before the experiment. A simple calculation indicates around 40% of the individuals in the sample hold positive credit card debt. This proportion is at the lower bound of the range of 40% to 80% found in the previous literature using US data (Gross and Souleles, 2002; Zinman, 2009; Fulford, 2015). Average increase in credit limit is around $1,850. This magnitude is economically significant. It is around 14.5% of the pre-experiment average total credit limit, and around 9.5% of the average pre-experiment annual income.

To check the quality of the surveys, I compare the survey answers with the information from the bank's database. Panel B of Figure A.1 from the Appendix presents the binned scatter plot of consumers' average monthly incomes over the six months before the experiment from the survey and that from the bank's database. The plot shows a clear linear relationship. A regression between the two measures gives an R^2 of 0.79. This finding indicates the high quality of the survey. Another requirement for the randomization to be effective is that the characteristics between the treatment and control groups are observationally indistinguishable before the experiment. As shown in Table 1, the consumers' characteristics are extremely similar between both groups. The differences

[13] Given that accepting the offer takes only a few seconds by clicking the accept button, the decision to not accept is unlikely to be a result of a high time cost. Online Appendix Figure A.6 shows approximately 80% of the participants who did not accept the offers indicated they were afraid of overspending. This finding suggests consumers might use a low credit limit as a commitment device against overspending, potentially due to present bias.

are all within 10% of the sample standard deviation, with large p-values. [14]

III Results

A. Learning from Credit Extension

This section examines the treatment effects of credit supply on consumer beliefs. I first present some motivating evidence of the effects of the experiment on consumer income expectations. I measure consumer expectations of future income change, $E_C[\Delta Y_i]$, using the following survey question:

Q1b: Your expected total income over the next 6 months is _____.

$E_C[\Delta Y_i]$ is the difference between the answers to Q1b and consumer average income six months before the experiment.

Figure 1 gives the scatter plots of consumer income-change expectations before and after the experiment. The red segments represent the answers for the treatment group; the blue segments represent the answers for the treatment group; the dashed segments represent the answers from the pre-experiment surveys; and the solid segments represent the answers from the post-experiment. As shown by the figure, consumer expectations about future income changes are in general positive before the experiment. After the experiment, there is no significant changes in the expectations of the control group. However, for the treatment group, the expectation about changes in the income over the next six months increases substantially. The difference-in-difference (DID) estimate gives the intent-to-treat (ITT) effect of the experiment on consumer expectations, which yields

[14] A frequent critique in using surveys to study patterns relates to the experimenter-demand effects (Orne, 1962; Zizzo, 2010; Mummolo and Peterson, 2019). I check whether the consumers in the two groups had different expectations about receiving changes in their future credit limits based on question 3b. The last row of Table 1 shows the average value of consumers' expectations about receiving future credit offers. Both averages are around 2.3. Therefore, consumers do not have heterogeneous expectations about being in the treatment or the control groups. In addition, de Quidt et al. (2018) recently show that in settings similar to those in this study, the experimenter-demand effects are plausibly small. In addition, Figure A.3 in the Online Appendix plots the evolution of spending and borrowing around the experiment for those who do not respond to the surveys. This pattern is similar to those in the main analysis.

an estimate of \$743.35. With an average increase in credit limit of \$1,866, consumers in the treatment group on average see a \$743.35 increase in income expectation over the next six months. This increase is equivalent to a 3.88% increase relative to the average annual income of the participants before the experiment.

I then estimate the treatment effects of the credit-limit offers on consumers' expectations regarding future income. The baseline specification is an instrumental variable (IV) regression with the following specification:

$$\begin{aligned} \Delta Limit_i &= \alpha_0 + \beta_0 Z_i + \gamma_0 X_i + e_i, \\ \Delta E_C[Y_i] &= \alpha_1 + \beta_1 \Delta \widehat{Limit}_i + \gamma_1 X_i + \epsilon_i. \end{aligned} \tag{1}$$

In the first-stage regression, Z_i is the treatment status and is equal to 1 if individual i is in the treatment group. $\Delta Limit_i$ refers to the changes in the credit limit the participants see on their offers. It is positive for those in the treatment group and 0 for those in the control group. Note that $\Delta Limit_i$ is equal to the realized changes in credit limit only for those who accept the offers. For those who do not accept the offer, $\Delta Limit_i$ is still positive, but the realized changes in credit limit are zero. Therefore, instead of a shock to the credit limit, $\Delta Limit_i$ is more closely tied with *news* in credit supply.

X_i are the province fixed effects.[15] The coefficient of interest is β_1, which measures the average causal response of consumers' expectations regarding future income to the credit limits offered to the consumers.[16]

The results are shown in Table 2, column (1). The estimated average effects of credit-limit offers on consumers' income expectations are both statistically and

[15]Given the randomization, the inclusion of the province fixed effects is not necessary. The inclusion of these fixed effects is for consistency with the analysis when decomposing the total effects of credit extensions. See section III.C for details.

[16]The analysis in the paper is preceded by a pilot study that happened in June 2019. The pilot study has the same design of the main analysis here. It involves more questions about expectations, no questions about preferences, and a much smaller sample size. In the pilot study, in addition to expectations about income, participants are also asked about their expectations of default probability, total saving, and total consumption. The results of other expectations are in Table A.1 of the Online Appendix. It shows that credit extensions increase expectations in consumption and income, but not wealth or default rate. Therefore, consumers tend to thing credit extensions are associated with larger spending, and the increases in spending is financed by higher income, but not higher default risk or lower savings. Therefore, in the main analysis, I only focus on expectations of income.

economically significant: a \$1 increase in the credit limit the bank offers increases consumers' expectations regarding their future income by around \$0.414 in total over the next six months. These findings show consumers infer future income growth from changes in credit limits and thus posit an income-inference channel through which credit extensions affect consumption. In column (2), I add province fixed effects, and the results hardly change.

A natural follow-up question is how the effects of the credit-limit offers differ for those who accept and do not accept the offers. Because the choice of accepting the offers is not randomized, respectively comparing the changes in the expectations of future income of those who accept and those who do not accept the offers with the control group would potentially yield selection biases. To cope with this problem, I split the control group into two sub-groups: those who would and would not accept the offers if they are given the offers. I perform an out-of-sample prediction with a LASSO logistic regression with consumers' pre-experiment characteristics as the predictors. [17] Specifically, I fit the LASSO logistic regression based on the treatment group using three-fold cross-validation and then predict who would accept the offers in the control group if they had received them. I label those who accept the offers in the treatment group and those who are predicted to accept the offers in the control group as the *acceptance* group, and label those who do not accept the offers in the treatment group and those who are predicted not to accept the offers in the control group as the *non-acceptance* group.

I check the effectiveness of the model in two ways. First, I randomly split the treatment group into a training sample and a test sample. Then, I fit the model using the training sample to predict who would accept the offer given the same covariates, and test the predictive power of the model by checking the error rates of the classifier using the test sample. The confusion matrix is shown in Online Appendix Table A.2. The results show the LASSO logistic model has a strong predictive ability. Out of a total of 1,678 observations, the LASSO logistic classifier is right in 1,457 cases. This finding implies an

[17] The predictors include gender, education, age, average income, average saving, average spending, average debt, average hours worked every week, credit score, changes in the credit score, number of credit cards owned, number of credit limit offers received before the experiment, subjective income volatility, city, short- and long-term discount rates, expectation about future income, and bank-proposed changes in the credit limits. All variables are from before the experiment.

error rate of only around 13%.

The classifier is akin to a propensity score matching without duplicates. Though the matching is not on the treatment status but a mediator variable after the treatment. A necessary condition for this strategy to be effective is to see if the distribution of the acceptance likelihood of the control and treatment groups overlap each other. Figure A.2 in the Online Appendix plots the distribution of the propensity scores of the treatment and the predicted control groups respectively for the acceptance and non-acceptance groups. The figure shows the distributions overlap each other nearly perfectly, indicating the matching procedure is successful.

I continue to provide some suggestive evidence on the effects of the experiment on consumer expectations separately for those in the acceptance and non-acceptance groups. The plots are in Figure 2. Regardless of being in the acceptance group and non-acceptance group, consumers in the control group have similar expectations about their future income before and after the experiment. For those in the treatment group, consumers in both the acceptance and non-acceptance groups have similar and marginally positive expectations before the experiment. This indicates that the pre-experiment income expectation is unlikely a determinant of if the consumers would accept the offers. After the experiment, both sub-groups in the treatment group see a significant change in their expectations. Besides, the expectation changes for those who accept the offers are much larger than the changes for those who do not accept the offers.

To estimate the average effects of credit-limit offers on those who accept and do not accept the offers, I re-fit (1) separately for the acceptance and non-acceptance groups. The identification assumption is that the prediction errors of the LASSO logistic identifiers are not correlated with the changes in consumers' expectations. This concern is mediated greatly, given that the matching process uses both the consumers' bank account information as well as their preferences and expectations before the experiment. In addition, the small error rates, as shown in Table A.1, indicate the selection issue is at most trivial. Columns (3) to (6) in Table 2 give the results. Consistent with Figure 2, the credit-limit offers significantly influence consumers' expectations regarding their future income for both those who accept and do not accept the offers: a $1 increase in the credit

limit offered by the bank increases consumers' expectations regarding their future income by around \$0.48 (\$0.24) over the next six months for those who accept (do not accept) the offers.

B. Heterogeneity in Income Inference

In this section, I explore the heterogeneity in the income-inference channel. I first study if credit-supply shocks affect consumers' expectations differentially, based on different levels of uncertainty. Presumably, if consumers infer information from the credit supply through Bayesian learning, the income-inference channel should be stronger for those with relatively more uncertain prior. Ideally, the heterogeneity analysis should be based on the signal-to-noise ratio of the learning process. However, because the survey design doesn't provide consumers' subjective uncertainty about the signals, I instead study the heterogeneity based on comparing consumers' subjective uncertainty about their future income growth prior to the experiment. The measure of subjective uncertainty in income growth is based on the surveys. Specifically, each participant is asked the following question:

Q1c: With a probability of 80%, your total income over the next 6 months will be between _____ and _____.

Given the answers to this question and the answers from Q1b, I calculate consumers' subjective income-growth uncertainty as the standard deviation from a normal distribution, assuming consumers' log income growth is normally distributed. This measure approximates the uncertainty in consumers' prior about future income growth when learning from the bank's actions.

To explore this heterogeneity, I first split the participants into deciles based on their pre-experiment subjective uncertainty, and then fit (1) respectively for the 10 subjective-uncertainty groups. Panel A of Figure (3) plots the 10 coefficients β_1 against the standardized average subjective uncertainty in each decile. As shown by the plot, there is a clear positive relationship between the sensitivity of income change to credit supply and consumer subjective uncertainty. This is consistent with a Bayesian learning

framework, in which consumers with more uncertain prior update more with respect to the signal.

Additionally, the sensitivity of income changes to credit extension is likely to be stronger when there is a larger cross-sectional variation in consumer income growth that can be observed by the bank but not the consumers. Specifically, if the bank could observe a larger amount of information on individuals who are similar to a specific consumer, the bank is more likely to extract more high-dimensional information that this consumer is lacking. Based on this logic, I construct a measure of income-growth uncertainty that is more likely to capture the cross-sectional uncertainty that the consumers face. I define $SD(\Delta \log Income)$ as the within-industry standard deviation of consumers' annual income growth in the year prior to the experiment. Specifically, I select all of those consumers at the bank who have two years of income data preceding the experiment and group those consumers into 18 industries. Then, I calculate the income-growth rates at the individual level and residualize them by log age, gender, highest degree earned, and log savings. I then take the standard deviation of the residualized income growth at the industry level to form $SD(\Delta \log Income)$. $SD(\Delta \log Income)$ gives a measure of the variability of individual income growth for each industry and serves as a proxy for the cross-sectional uncertainty of consumers' income growth. When $SD(\Delta \log Income)$ is high, the bank is more likely to observe the income-growth rates of some consumers who are similar to the participants. The heterogeneity analysis of income expectation sensitivity to credit supply with respect to $SD(\Delta \log Income)$ is in Panel B of Figure 3. Similar to subjective uncertainty, there is a clear positive relationship between the sensitivity of income change to credit supply and industry uncertainty. This supports the conjecture that a consumer would infer more from bank supply when the bank can observe a larger cross-sectional variation in the information that are more relevant to this consumer.

Recent studies document that the effects of credit supply on consumer spending are also large for those with high liquidity buffers (D'Acunto et al., 2020; Aydin, 2022). Explaining the finding with the standard buffer-stock model is hard. I continue to study whether consumers with different levels of liquidity buffers infer different levels of information from credit supply. To do so, I repeat the heterogeneity analysis with the

consumer deposit-to-income ratio and utilization ratio. The latter is defined as the ratio of outstanding interest-incurring debt over the total credit limit. The results are in panels C and D of Figure 3. Both plots show a negative relationship between the sensitivity of income change to credit supply and the degree of borrowing constraints. This helps to explain the large consumption responses to credit limit extensions for consumers that are far from their borrowing limits.

One potential explanation of the positive relationship between expectation changes and liquidity is the degree of overconfidence in interpreting a positive signal. Previous literature mostly documents a positive relationship between overconfidence and one's social economic class. For example, Bénabou and Tirole (2002) analyze how overconfidence could induce higher ex-post outcomes by overcoming present bias. On the empirical side, Bhandari and Deaves (2006) and Belmi et al. (2020) document that the degree of overconfidence increases with one's social economic class including income and education. Given that liquidity including the wealth-to-income ratio and credit availability in general increases with income and education, it is expected that the degree of overconfidence also increases with liquidity. Therefore, individuals with higher liquidity will overestimate how much credit extension as a positive signal tells about their income growth, consequently having a larger expectation change after seeing the limit-extension shock.

C. Spending Responses to Credit-Supply Shocks

The previous section shows that credit-supply shocks have a considerable influence on consumers' expectations regarding their future incomes. However, recent studies have documented that consumers' expectations sometimes have limited effects on their actions (Ameriks et al., 2020; Giglio et al., 2021). To see if consumers indeed act on the changes in their beliefs, I further analyze the effects of credit expansions on consumers' spending and debt-taking behaviors after the experiment. I first show the monthly evolution of borrowing scaled by the proposed credit limit increases around the time of the experiment. Then, I calculate the changes in debt as the difference between total outstanding interest-incurring debt at the end of a month and that right before

the experiment. Total borrowing is residualized by quarter and city fixed effects.

Panel A of Figure 4 plots the evolution of debt changes scaled by the changes in the proposed credit limit around the experiment for the participants in the treatment and control groups. The x-axis denotes the number of months away from the experiment. In each of the two subplots, the solid red line represents the treatment group, and the dotted blue line represents the control group. As shown, the sharp increase in borrowing right after the experiment for those in the treatment group indicates the effectiveness of the experiment. Additionally, borrowing for the control group starts to increase in the seventh month after the experiment, the time when the control group receives the postponed offers. Panels C and E plot the evolution of the debt separately for those in the acceptance and non-acceptance groups. Similar to Panel A, a sharp increase can be seen in borrowing right after the experiment for those who are in the treatment group and accept the offer. In addition, even for those who do not accept the offer, borrowing starts to increase in the experimental period.

The empirical strategy for assessing the statistical behavior of debt after the experiment is the same as that in (1), with changes in interest-incurring outstanding debt as the left-hand-side variables. Table 3 presents the results. Column (1) gives the marginal effects of the credit limits on total interest-incurring debts using all observations in the treatment and control groups. The estimate gives the marginal propensity to borrow out of *news* in limit changes (MPB^{N}). It measures the total newly-accumulated debt for each dollar higher qualified credit limit. The statistic is different from the MPB estimated in the previous literature (Gross and Souleles, 2002; Agarwal et al., 2017; Gross et al., 2020; Aydin, 2022), which measures the marginal propensity to borrow out of *realized* changes in credit limit.

Column (1) shows that for a \$1 increase in the credit-supply news, consumers' debts increase by \$0.127 on average six months after the experiment. Columns (2) and (3) give the results respectively for those who accept and do not accept the offers. Column (2) shows that for a \$1 increase in the credit limit, consumers who accept the offer increase their debt by \$0.175 on average six months after the experiment, whereas, as column (3) shows, even for those who do not accept the offer, a \$1 increase in the credit limit

increases their debt by \$0.042. This finding is indicative of credit-supply shocks affecting consumer borrowing in addition to a relaxed borrowing limit.[18]

The estimated MPB^{N} in column (1) is similar to the MPB found in the previous literature. For example, the MPB is 0.11 at a 12-month horizon in Gross and Souleles (2002), between 0.08 and 0.3 in Agarwal et al. (2017), and approximately 0.16 at a nine-month horizon in Aydin (2022). However, a difference in the setting here is that I have an estimate of the effects of the credit-limit offers on borrowing for all the consumers, regardless of whether the offers are accepted. However, in previous studies, analyses are usually based on the changes in credit limits that have to be accepted. In my setting, for those who accept the offer, MPB^{N} is the same as the MPB estimated in the previous literature, whereas for those who do not, MPB^{N} should be the conventional MPB minus the additional effects on borrowing through relaxed borrowing limits. In this case, MPB^{N} should be weakly smaller than the conventional MPB. Consequently, the conventional MPB under the scenario where credit-limit changes are directly applied to everyone should be a number between the estimate from column (1) and the estimate focusing on those who accept the number, which is between \$0.127 and \$0.175.

Credit-limit shocks have great effects on consumers' borrowing. Conventionally, the mechanism is based on the buffer-stock model. Specifically, a relaxed credit limit reduces consumers' precautionary motives, thereby increasing their current consumption. However, previous results show credit-supply shocks also change consumers' expectations regarding their future income. Given the rosier beliefs about their earning ability, consumers increase spending even if their borrowing limit does not change. This mechanism is suggested by column (3) of Table 3. A direct test of the effects of the changes in expectations on borrowing is to control for the expectation changes. Because the changes in expectations are also a result of the experiment, I treat $\Delta E_C[Y_i]$ as a second endogenous variable and employ the location-by-treatment IVs to separately identify the effects of the offered credit-limit changes and the changes in expectations on consumers'

[18]Given that the data is from a single bank, the changes in debt could be merely a result of transfers from other banks. In Online Appendix Figure A.6, I plot the changes in the net transfers between this bank and other banks. The results show that the changes in the net transfer are statistically indifferent from zero for both the control and treatment groups. This indicates that the change in debt is not a results of the participants moving funds around different bank accounts.

borrowings (Kling et al., 2007; Abdulkadiroğlu et al., 2014; Kline and Walters, 2016). Specifically, I use the interaction between the province dummies and the treatment group assignment as the IVs for $\Delta Limit_i$ and $\Delta E_C[Y_i]$, controlling for the province fixed effects. The specification is

$$\Delta B_i = \alpha_2 + \beta_2 \Delta Limit_i + \omega_2 \Delta E_C[Y_i] + \gamma_2 X_i + u_i. \tag{2}$$

The site-by-treatment strategy requires that the consumer's degree of learning varies across provinces. This cross-province heterogeneity generates additional variation in the weights of the income-inference channel in affecting consumer spending. This requirement is satisfied if people in different provinces have different beliefs about how well the bank could predict their future income.[19] In addition, there are several additional assumptions to satisfy for the unbiased estimation of (2) (see Reardon and Raudenbush (2013) for details). A stronger sufficient assumption is that, conditional on X_i, β_2 and ω_2 are homogeneous across the provinces (see Hull (2015), Kirkeboen et al. (2016), and Kline and Walters (2016) for related results). This assumption can be examined with the overidentification tests.

Columns (4) and (5) in Table 3 present the results respectively for those who accept and do not accept the offers using the province-by-treatment interactions as IVs. The first-stage partial F statistics are large and are respectively 74.27 and 19.26. This observation implies using the province-by-treatment interactions as instruments yields significant independent variations in the two channels. In both columns, a \$1 increase in the expectations of future income over the next six months increases debt by around \$0.12. Conditional on the changes in expectations, the main effects of credit-limit offers are significantly smaller. As column (4) shows, after controlling for $\Delta E_C[Y_i]$, a \$1 higher offered credit limit on average increases the debt of those who accept the offer by \$0.118 six months after the experiment. The estimate decreases by around 32.57%. Column (5) shows that for those who do not accept the offer, after controlling for $\Delta E_C[Y_i]$, the

[19]For example. the bank could have different market shares and numbers of customers in different provinces. Therefore, consumers in a market where the bank's market share is larger could have a stronger impression of the bank's information-processing ability in that province.

effects of a credit-limit shock on borrowing are no longer significant. This finding suggests the income-inference channel is the only significant channel through which credit-supply shocks increase the spending of the consumers who do not accept the offers.

The assumption that β_2 and ω_2 do not vary at the province-level can be tested with the overidentification tests. From Table 3, the overidentification tests all have large p-values, indicating the data are consistent with a constant-effects framework. That is, the variation in treatment effects is not coming from a heterogeneous-effects model, in which, in some provinces, the same degree of changes in expectations triggers a larger change in total debt. Instead, the findings support a dose-response relationship. In provinces where the reaction in expectations is larger (larger dose), the effects on outcomes are larger (larger response).

For a robustness check, I decompose the two channels, with the additional sample receiving the information treatment (see Step 4 in section II.D.). The results are in Table A.3 in the Online Appendix. Columns (1) and (2) focus only on the additional sample. The information treatment induces a smaller change in income expectations. In columns (3) and (4), the decomposition using the information treatment and randomization as the IVs yields similar results, thereby providing additional supportive evidence of the relative weights of the two channels. However, the first-stage F-statistics are much smaller. This indicates the information treatment has relatively weaker relevance, cautioning the possible problems of weak instruments. Because of this concern, I use the location-by-treatment IVs as my strategy for the decomposition.

Consumers' precautionary motives are usually larger in the case of larger uncertainty or a smaller liquidity buffer. Table 4 gives the estimates of MPB for the consumers in different uncertainty and liquidity groups. Considering the insignificant responses after controlling for changes in expectations for those who don't accept the offers, I focus on those who accept the offers. Panel A splits the sample by the participants' ex-ante subjective uncertainty; Panel B splits the sample by the within-industry volatility of income growth; Panel C splits the sample by the participants' wealth-to-income ratio right before the experiment; and Panel D splits the sample by the participants' utilization rate before the experiment. The results show income expectations have similar effects on

consumer borrowing; a \$1 increase in the expectations of future income over the next six months increases debt by \$0.11 after six months. Given that credit-limit shocks have considerable effects on expectations for those with higher uncertainty, the heterogeneity in MPB is smaller across the uncertainty group after controlling for the change in expectations. For example, without conditioning on the changes in expectations, a \$1 increase in the offered credit limit on average increases the debt of those who accept the offer and are in the high- (low-) subjective-uncertainty group by \$0.208 (\$0.133) six months after the experiment. After controlling for the changes in expectations, this number decreases to \$0.141 (\$0.102).

As for the heterogeneity by liquidity, because the income-inference channel is stronger for those with higher liquidity buffers, the heterogeneity in MPB is larger across the wealth-to-income group after controlling for the change in expectations. As Panel C shows, unconditional on the changes in expectations, a \$1 increase in the offered credit limit on average increases the debt of those who accept the offer and are in the high (low) wealth-to-income group by \$0.142 (\$0.213) six months after the experiment. After controlling for the changes in uncertainty, this number decreases to \$0.068 (\$0.175). The results are similar for utilization rate.

Given that I observe consumer transaction-level data, I can study the effects of credit-limit shocks on consumers' non-debt-financed spending. I define non-debt-financed spending as consumers' cumulative spending over the six months after the experiment minus the newly accumulated debt. Panels B, D, and F of Figure 4 plot the evolution of the cumulative changes in non-debt-financed spending, scaled by the changes in the proposed credit limits around the experiment for the participants in different groups. Similar to the debt responses, a sharp increase can be observed in the spending right after the experiment for both the acceptance and non-acceptance groups.

Table 5 presents the estimates of the average effects of the proposed limited changes on non-debt-financed spending. The results are qualitatively similar to those for debt. From column (1), the consumers on average increase non-debt-financed spending by \$0.216 six months after the experiment. Combined with the estimates from Table 3, this finding is equivalent to total spending of \$0.34 over a six-month horizon. Columns (2) and (3)

show the results respectively for those who accept and do not accept the offers. As shown in the columns, for a \$1 increase in the offered credit limit, the consumers who accept (do not accept) the offer on average increase non-debt-financed spending by \$0.276 (\$0.108) six months after the experiment. Combined with the estimates from Table 3 give the estimates of the MPC out of news in limit change (MPC^{LN}). It measures the total amount of spending for each dollar higher credit limit qualified. From Table 3 and Table 5, MPC^{LN} for the accepters and non-accepters are respectively \$0.451 and \$0.150 over a six-month horizon.

The estimates in front of $\Delta E_C[Y_i]$ from Table 5 show a \$1 increase in the expectations of total income over the next six months increases total non-debt-financed spending by around \$0.25. Combined with debt-financed spending, a \$1 increase in the expectations of total income over the next six months increases total non-debt-financed spending by around \$0.37. If the changes in income expectations are with respect to the permanent component of consumers' lifetime income and no adjustment costs are incurred, the coefficient in front of $\Delta E_C[Y_i]$ in column (3) should be close to 1. Therefore, the estimate, which is smaller than 1 but still economically significant, indicates changes in credit supply serve as a signal of a temporary though persistent shock, to consumer beliefs about their future income[20].

Controlling for the changes in expectations, consumers who accept the offers on average increase non-debt-financed spending by \$0.174 six months after the experiment. For those who do not accept the offer, the increase in non-debt-financed spending is a statistically insignificant amount of 4.2 cents. The identified effects of credit constraints on spending through the income-inference channel can explain several findings in the previous empirical literature. Many studies find relaxed credit limits have considerable effects on both spending and borrowing for high-liquidity consumers. This finding is hard to reconcile with a conventional buffer-stock model with rational expectations. Table II shows the income-inference channel is stronger for high-liquidity consumers. The findings therefore help explain the high MPB among high-liquidity consumers. Additionally, the

[20]Given that the spending response seems to level off after six months, the smaller coefficient is unlikely a result of habit formation(Campbell and Cochrane, 1999; Dynan, 2000; Christiano et al., 2005; Havranek et al., 2017).

finding that the income-inference channel is stronger when uncertainty is high could also explain the findings in Gross et al. (2020), who document that MPB is likely to be 20% – 30% higher during a recession in additional to the presence of liquidity constraints.

D. Weights of the Income-Inference Channel

The estimates in Table 3 and Table 5 suggest the weight of the income-inference channel in MPC^{LN}. That is, how much do expectation changes affect total spending when consumers are qualified for one dollar higher borrowing limit. In this cases, the consumers, like those in this study, could choose whether to realize the one-dollar higher credit limit or not. Summing up the estimates in column (1) in Table 5 and Table 3, total spending for an average consumer increases by 34.3 cents for each dollar higher offered credit limit. Based on column (4), total spending increases by 20.5 cents after controlling for expectation changes. Therefore the weight for the income-inference channel in MPC^{LN} is 1-20.5/34.3 $\approx$ 40.23%.

Another statistic of interests is MPC out of realized changes in borrowing limit (MPC^{L}). This statistic measures the spending responses when limit offers are forced to be accepted, and is usually the focus of the previous literature (Gross and Souleles, 2002; Agarwal et al., 2017; Chava et al., 2020; D'Acunto et al., 2020; Gross et al., 2020; Aydin, 2022). In this case, the weight of the income-inference channel is the same as that in MPC^{LN} for the accepters, as there is no difference on these consumers regardless of whether the limit change is offered or applied. Calculating the weight for MPC^{L} requires the hypothetical value of the total spending when the non-accepters have to accept the offers. To provide a bound of this number, first assume that total spending is always larger for the non-accepters if they had to accept the offers. In addition, assume the effects of news in limit changes on expectations stay the same. Then there are two extreme cases. First, if total spending does not change after limit offers are accepted, then the weight of the income-inference channel in MPC^{L} is the same as that in MPC^{LN}, which is 40.23%. The other extreme is when total spending increases to infinity for the non-accepters when they have to accept the offers. Then the weight of the income-inference channel for these people is zero. Therefore, a lower bound of the weight in MPC^{L} for an average consumer

is the weighted average of that for the accepters and zero, which is around 24.88%. In sum, the weight of the income-inference channel in MPC^{L} should be a number between 24.88% and 40.23%.

A more straightforward way to calculate the weight of the income-inference channel in MPC^{L} that requires a much stronger assumption is directly comparing total spending for accepters and non-accepters. If we assume that the only difference between accepters and non-accepters is the realized change in credit limit, then the weight of the income-inference channel in MPC^{L} is just 0.150/0.451, which is roughly 33.26%. To sum up, in general, the decomposition exercises suggest that the weight of the income-inference channel in MPC^{LN} is around 40.23%, and is between 24.88% and 40.23% for MPC^{L}.

E. Overreaction to Credit Supply

The previous sections show consumers make significant inferences from credit supply decisions. There are two possibilities that contribute to this finding. First, by employing advanced statistical inference methods, banks can extract additional information about consumers' earnings ability in the future. Consumers then rationally incorporate such information from credit supply. Second, credit extensions could be completely uncorrelated with future income growth. Nonetheless, over-confident consumers believe credit extensions signal higher future income growth anyway. In this section, I study which mechanism is more likely in the data by comparing consumer post-experiment expectations and the realized income changes.

Figure 5 gives the binned-scatter plots of consumer income-change expectations and the realized income changes versus the pre-determined limit changes. All variables are residualized by age, degree, gender, income, saving, total spending, industry fixed effects, city fixed effects. In all four panels, the x-axes are the limit changes as proposed before the random assignment. These numbers are positive before residualization for all participants. In Panel A, the y-axis is consumer expectations about income changes over the next periods before the experiment. As shown, pre-experiment expectations are not significantly correlated with proposed limit changes, and this is the case for both the control and treatment groups. From Panel B, realized income changes are positively

correlated with proposed limit changes for both control and treatment groups, and the associations are similar for the two groups. Panels A and B indicates that, when banks actively offer to increase the credit limits of the consumers, banks are more likely to have more information about consumer income in the near future.

Panel C plots consumer post-experiment expectations. Since the control group never receives the offer, there is no change in their expectations. While for the treatment group, there is a positive relationship between expectation changes and proposed limit changes. This finding confirms the previous results. Besides, comparing Panel B and Panel C, an interesting finding is that the association between post-experiment expectations and proposed limit changes is stronger than the association between realized income changes and proposed limit changes. In other words, consumers in the treatment group get over-optimistic about their earnings ability after receiving credit supply shocks. In Panel D, I plot consumer expectation errors after the experiment. Expectation errors are defined as the difference between post-experiment expectations and realized income. Confirming the results in panels B and C, expectations errors are negatively correlated with proposed limit changes for the control group, and positively correlated with proposed limit changes for the treatment group.

From a Bayesian learning perspective, Figure 5 suggests that, for a credit-supply event, credit extension as a signal for consumer future income has a positive correlation with consumer realized income. At the same time, as compared with the signal, consumer prior is less correlated with realized income. Consumers then learn from the signal. As a result, the posterior moves in the direction of the signal. However, the over-optimism after the credit-supply event indicates that the consumers overreact to the signals.

Aydin (2022) explains that credit-limit increases do not have informational content for individuals with rational expectations, by showing the realized income between those in the control and the treatment groups does not differ after the experiment. However, this finding is necessary as long as the experiment does not change the labor supply[21].

[21] There is mixed evidence on the findings about how access to higher credit limits affects labor supply. For example, Aydin (2022) shows that there are no significant changes in labor supply after given higher credit limits. While in Herkenhoff et al. (2021), self-employment increases monotonically with increases in credit limits.

Otherwise, the experimental design would fail due to a lack of effective randomization. Similarly, as shown in Figure 5, I do not find any significant difference in the realized income growth between the control and the treatment groups. But this lack of a finding does not indicate the **expectations** of future income change remain unchanged. When the bank has superior information about consumers' future income growth, credit supply is somewhat more tightly correlated with true future income growth. With effective randomization, this true future income growth, as well as the ex-ante expectation regarding future income growth, is the same for the participants in both the control and the treatment groups. However, consumers' ex-ante expectations are more *wrong* to begin with, possibly because consumers' information set is noisier with respect to the variables that the bank has better information on, at least during the period of active credit expansion. Therefore, as long as the changes in the income perspective do not change labor supply, the realized income growth will stay the same, while the expectations will be corrected to some extent, which would also change consumption and borrowing.

In addition, even if consumer expectations changes after credit extension are due to mistakes in guessing the bank's supply function, the assumption that consumers with rational expectations would learn to correct these mistakes are challengeable. Specifically, because of the many types of shocks to income throughout their lifetime and the sparsity of bank-offered credit-limit extensions, consumers with recency effects would have difficulty associating the lower ex-post income with inaccurate learning from credit-limit extensions. This is especially likely in the consumer credit market. As shown by Agarwal et al. (2013), consumers usually learn about their mistakes slowly in the consumption credit market but forget them quickly. Altogether, the assumption on rational expectations in the long run to correct consumer misbeliefs does not seem to hold in the consumer credit market.

F. Delinquency Rates

Previous literature has found mixed evidence on the effects of relaxed credit constraints on consumers' default probability. For example, in Agarwal et al. (2017), relaxed credit constraints increase the delinquency rate of all borrowers except those with relatively high FICO scores. In Aydin (2022), however, a credit-supply shock doesn't have a significant

effect on borrower delinquency rates. Here, I provide some additional evidence on the effects of credit-supply shocks on borrower delinquency rates, especially when accounting for the income-inference channel.

Table 6 presents the results. I define default as a dummy variable equal to 1 if the participants have a 60-day delinquency over the debt taken during the experimental period, and 0 otherwise. $\Delta Limit_i$ and $\Delta E_C[Y]$ are in thousands of dollars. All coefficients are multiplied by 100 for easier interpretation. As shown in column (1), for all treated participants, a \$1,000 increase in the credit limit increases the default probability by 0.108 percentage points. With a sample average of 2.44%, this increase is equivalent to a 4.43% increase. For the participants in the treatment group who accept (do not accept) the offer, a \$1,000 increase in the credit limit increases the default probability by 0.176 (0.073) percentage points. However, the estimated coefficient for those who do not accept the offer is statistically insignificant due to large standard errors. Columns (4) and (5) show updates in the expectations of future income have considerable effects on default rates: for a \$1,000 higher expectation about future income, the default rate increases by around 0.2 percentage points. After controlling for the updates in the income-growth expectations, the estimated effects of credit supply on the default rate become insignificant. Therefore, the extension of credit limits does not seem to have any effects on consumers' default rates beyond inducing a too-rosy expectation.

G. Beliefs about Bank Supply

The results show consumers make inferences about their future earning ability from bank credit supply. In addition, consumers tend to get over-optimistic about their income growth after the news about credit supply. This section provides some suggestive evidence about the source of such overreactions.

Consider a simplified version of the learning process in this setting. Suppose the true income change for the representative consumer is

$$\Delta Y = \rho.$$

Before seeing the credit supply, the consumer has a prior of their income changes in the next period, which follows

$$\rho \sim N(\rho_0, \sigma_0^2).$$

At the same time, the bank receives a signal about the consumer's future income and decides on a level of limit change that follows

$$\Delta l = f(\rho_B + \epsilon_B),$$

where ρ_B is the level of income changes the bank expects about the consumer, and $\epsilon_B \sim N(0, \sigma_B^2)$ is the error in bank belief. $f(\cdot)$ is the credit supply rule. The consumer would try to infer ρ_B by inverting the credit supply rule when seeing the credit limit change, Δl. That is, from the consumer's perspective,

$$E_C[\rho_B] = \eta \cdot f^{-1}(\Delta l) = \eta \cdot \rho_B, \tag{3}$$

where $E_C[\rho_B]$ is the expectation of bank belief from the perspective of the consumers. η is a misperception parameter. When $\eta = 1$, the consumer on average has rational expectations. The consumer is assumed to update the income expectation through Bayesian learning with the subjective belief of the signal. After observing credit supply, the consumer's posterior about the income change is

$$\hat{\rho} = (1 - K)\rho_0 + K \cdot \eta \cdot \rho_B, \tag{4}$$

where $K = \sigma_0^2/(\sigma_0^2 + \sigma_B^2)$ is the learning rate.

The second part of the right-hand-side of (4) indicates over-optimism could come from three components. First, consumer learning rate K is different from the true value. This is possible if the consumers have diagnostic expectation (Bordalo et al., 2019), self-attribution bias (Daniel et al., 1998), or other mechanisms of under- or over-reaction to bank belief shock ϵ_B in the signal extraction process. Second, consumers overestimate the true sensitivity between bank belief and credit supply, therefore having $\eta > 1$. Third, the bank might systematically overestimate consumer income changes. In this case, ρ_B is

much larger than the true value ρ. Empirically testing if the learning rate K as perceived by the consumers differs from the true value is difficult, since this requires the simultaneous measurement of the true distribution of σ_B^2 and that as perceived by the consumers. Instead, I use the survey questions to study consumer subjective beliefs about the inverse credit supply rule, $E_C[\rho_B]$. Specifically, I measure consumer subjective sensitivity of bank beliefs to credit supply as

$$\lambda = \frac{\partial E_C[\rho_B]}{\partial \Delta l} = \eta \cdot \frac{\partial \rho_B}{\partial \Delta l}. \tag{5}$$

From a Bayesian inference perspective, λ calculates the likelihood of income changes signaled by a marginal change in credit limit, as perceived by the consumers.

I rely on the following questions from the survey to measure λ[22]:

Q4a: Suppose your bank increases your credit card limit by 5,000 CNY this month. This would mean that the bank expects your total income to change by _____ over the next 6 months.

Q4b: Suppose your bank increases your credit card limit by 10,000 CNY this month. This would mean that the bank expects your total income to be changed by _____ over the next 6 months.

Suppose the answers from the two questions are respectively x_1 and x_2, I calculate the consumers' subjective beliefs about the sensitivity of credit supply to the bank-perceived income growth, λ_i, as

$$\lambda_i = \frac{x_2 - x_1}{10,000 - 5,000}. \tag{6}$$

Panel A of Figure 6 plots the distribution of λ_i. It shows a large heterogeneity in consumers' subjective beliefs about the sensitivity of credit supply to bank-perceived income growth. Around 28.73% (6.53%) of the consumers believe $\lambda = 0$ ($\lambda < 0$). However, most of the participants believe credit-limit extensions signal a large degree of income growth in the future. The economic significance of λ is large. Its average value is 0.86, and the median is 0.50. Thus, for a \$1 increase in the credit limit, consumers on average believe the bank expects their income to increase by \$0.86 over the next six months. This number is much larger than the relationship between limit changes and realized income

[22]These survey questions are sent to a random 50% of the participants.

change, which is around 0.3 as shown by Figure 5. In addition, the distribution of λ is relatively smooth, indicating the answers are not based on any rules of thumb (e.g., a one-for-one increase or always zero).

From (5), λ_i can be decomposed into two components: (1) the true marginal relationship between credit supply and the bank-perceived income growth of the consumers, $\partial\rho_B/\partial\Delta l$, and (2) the degree of misperception about bank credit supply rule, η. To measure η for the participants, I use a random 5% of all the data available from the bank during 2017 – 2019 and out-of-sample predict the income growth of the participants over the experimental periods.[23] I then split the participants into 200 groups based on their cumulative changes in credit scores six months before the experiment, and regress the predicted income growth on the proposed changes in the credit limit. The coefficients in front of the proposed limit changes give $\partial\rho_B/\partial\Delta l$. Then η_i is extracted by taking the ratio of λ_i and $\partial\rho_B/\partial\Delta l$.

The estimated $\partial\rho_B/\partial\Delta l$ has a mean value of 0.245 with a standard deviation of 0.07. This is to say, a one-dollar increase in credit limit is associated with 24.5 cents higher consumer future income as predicted by the bank. The mean value is smaller than the sensitivity of realized income changes to credit supply of 30.9 cents. Therefore, the relationship between bank supply and bank belief is close to but smaller than the relationship between bank supply and the realized income changes. Panel B of Figure 6 plots the misperception parameter η_i. There is a large variation in η_i. Besides, consumers on average considerably overestimate the relationship between credit supply and bank belief. The average value of η_i is 3.50, with a median of 3.04. This is to say, the perception of the sensitivity of bank belief to the credit supply of the average consumer is around 2.5 times larger than the true relationship. The findings suggest that consumers' overreaction behaviors are consistent with them being over-optimistic about how much bank credit extensions could signal their future income changes.

[23]The model is similar to the one the bank uses to predict consumer default risks.

H. Robustness of the Results

A concern about combining surveys with RCT is that the surveys might provide some cues that can induce the participants to change their behaviors. For example, the surveys in this study ask consumers about their expectations regarding their future income. Additionally, the hypothetical questions in section 4 of the survey could lead the participants to think about the relationship between the credit supply and banks' perceived income growth of the consumers. If these concerns exist, the income-inference channel documented here would be exaggerated. To provide evidence of the robustness of the results, I first focus on the 50% of consumers to whom the hypothetical questions are not sent. The results are shown in Table A.4 of the Online Appendix. The results are close to the estimates using all observations, indicating the hypothetical questions have little impact on consumers' beliefs.

To further test if the survey has any significant influence on the participants, I plot the evolution of debt and non-debt-finance spending again using the participants who do not fill out the survey, only fill out the first survey, and only fill out the second survey. The results are shown in figures A.3 to A.5 in the Online Appendix. If the surveys have any framing effects, the spending and debt responses should have been smaller than the responses using the participants who filled out all surveys. However, from figures A.3 to A.5, the evolution of debt and spending is similar in both magnitude and pattern to those in Figure 4. Therefore, the framing effects or the experimenter-demand effects because of the surveys are likely to be trivial.

In addition, the analysis here focuses on consumers whose income information is observable by the banks. In this case, the bank may have greater predictive power regarding these participants' future income, thereby inducing a larger income-inference channel than the general population. As presented in Table A.5 of the Online Appendix, I study the effects of the experiment on consumers' expectations and spending behaviors using the 1,383 observations for which the bank could not observe the income information. In general, these results are close to the results obtained using consumers with observable income information. However, the strength of the income-inference channel is slightly smaller when income information is not observable; controlling for the changes in

expectations, for a \$1 increase in the offered credit limit, the total spending decreases from \$0.495 to \$0.364 for those who accept the offers. This finding is equivalent to a 26.46% decrease, compared with around 32% decrease for the consumers whose income is observable. Thus, the weight of the income-inference channel is likely to be smaller, but still economically significant, for the consumers whose income information is not known by the bank.

The estimation strategy in this study is based on an RCT that randomly increases consumer credit limits for a treatment group. In the existing literature, the effects of credit limits are often estimated with quasi-random variations induced by consumers having FICO scores passing some thresholds (Agarwal et al., 2017; D'Acunto et al., 2020). In such a setting, if consumers receive a change in their credit limit because of this quasi-random variation, credit-limit changes shouldn't have any effects on consumer expectations if the consumers also believe such changes in the credit limit are random. However, in the field, consumers might have limited knowledge about such mechanisms and treat any changes in credit limits as a result of banks re-evaluating their customers' future economic activities. In this case, credit extensions due to quasi-random variation in credit scores would also have effects on consumer expectations. I test this conjecture using a similar fuzzy regression discontinuity (RD) design. Specifically, the bank assigns each participant a credit score that ranges from 577 to 800. It then groups each consumer into five bigger groups for the high-level evaluation of risk composition. The four thresholds are 661, 678, 717, and 722 using the four thresholds as discontinuities. Table A.6 in the Online Appendix gives the results with the fuzzy RD design. The results are based only on the consumers in the treatment group. I define consumers above and below the threshold with the middle point in each of the three intervals. In the odd columns, I control for a second-order polynomial of credit score; in the even columns, I control for a third-order polynomial of credit score. The results mostly confirm the findings in the main analysis, indicating that even if variations in credit score sometimes induce quasi-random changes in credit limits, consumers in general are unlikely to treat such changes as completely random.

I. Implications from US Data

To shed light on the external validity of the income-inference channel, I present some survey results based on US data collected through SurveyMonkey, an online survey platform.[24] Recently, Bentley et al. (2017) and Haaland et al. (2020) note data from online survey platforms such as SurveyMonkey and MTurk can have high quality when compared with traditional, larger surveys and field experiments. To verify the survey's representativeness of the population, I confirm the basic demographic patterns by looking at the association between individual income level and age. The patterns are consistent with that generated from a conventional life-cycle model (Gourinchas and Parker, 2002) and are close to that from the 2020 US census.

Without an experiment and bank account data, I cannot determine the causal effects of credit expansion on consumers' expectations and spending behaviors through the income-inference channel in the US. However, I use two separate surveys to show consumers in the US also believe banks' credit-extension decisions signal changes in their future income growth. [25] The survey questions and results are in section B of the Online Appendix. Basically, when banks increase a consumer's credit limit, the common belief is that this person would see higher spending and higher income over the next year, but not lower savings or higher defaults. In addition, the λ elicited based on (3) using the same questions shows a distribution similar to that using the Chinese data. Specifically, for a \$1 increase in the credit limit, the average consumer in the US believes the bank expects their income to increase by \$1.60 over the next year.

A natural question is which types of income-related information in the future do consumers believe that the banks can have a better predictive ability over. To provide some stylized evidence, question 4 in the second US survey lists a set of income-related options and asks the participants if banks have a better predictive ability about the near-future growth rate of these options than the participants themselves. The options include the participants' income, the macro economy, the local economy in which the

[24] The results, a brief description of SurveyMonkey, and the collection method are provided in the Online Appendix.

[25] This strategy is similar to the *reported preference* approach in estimating MPCs. See Shapiro and Slemrod (2003), Jappelli and Pistaferri (2014), Graziani et al. (2016), Parker and Souleles (2019), Coibion et al. (2020), Fuster et al. (2020), and Jappelli and Pistaferri (2020) for some examples.

participants live, the industry/sector the participants work in, and their income compared with their peers. The participants can also select an option of *None of the above* if they don't believe that banks have better predictability over any of these options. The frequency of selection is in Online Appendix Figure B.4. As shown, only around 18% of the participants believe that the participants themselves have a better predictive ability over all these components. Among the listed components, the most selected is the near-future growth rate of the local economy where the participants live (42%).

The results from the US survey indicate consumers believe a higher credit limit is associated with higher future spending, and this future spending is financed by higher incomes but not lower savings or higher defaults. Therefore, the online surveys show consumers believe income growth is correlated with bank credit extensions, thus providing supporting evidence for the external validity in the US context.

IV Structural Estimation

Recent decades have seen a proliferation of AI and big data. This increase is expected to significantly improve banks' statistical models. As a result, banks would expect to receive more precise signals about consumers' future economic activities. In this section, I structurally estimate a life-cycle model incorporating the income-inference channel to study how shocks to banks' beliefs affect their credit extensions and consumers' subsequent spending behaviors, especially when the signals banks receive have different levels of precision.

A. Setup

A.1 Income Process

The income growth of the consumers follows:

$$\begin{aligned} \log y_{i,t} &= \alpha + z_{i,t} + \epsilon_{i,t} \\ z_{i,t} &= \rho z_{i,t-1} + \nu_{i,t}, \end{aligned} \tag{7}$$

where $\epsilon_{i,t}$ and $\nu_{i,t}$ are i.i.d. normal shocks with $\mathbb{E}[e^{\epsilon_{i,t}}] = 1$ and $\mathbb{E}[e^{\nu_{i,t}}] = 1$. The variances of $\epsilon_{i,t}$ and $\nu_{i,t}$ are σ_ϵ^2 and σ_ν^2, respectively. α is the life-cycle component. It is assumed to be a constant and is known to everyone. The consumers do not know the true value of $z_{i,t}$ and need to make inference based on Bayesian learning. The Kalman-filtering problem with respect to the persistent component of $\log y_{i,t}$ here is a simplified version of Guvenen (2007). At time 0, i's prior of $z_{i,t}$ follows $N\left(z_{i,0}, \sigma_{z,0}^2\right)$. In each period, consumers observe $y_{i,t}$ and update their beliefs accordingly. Therefore, the posterior of $z_{i,t}$ is

$$\begin{aligned}
\hat{z}_{0,i,t+1} &= \rho\left[\hat{z}_{i,t} + \frac{\hat{\sigma}_{z,t}^2}{\hat{\sigma}_{z,t}^2 + \sigma_\epsilon^2}(y_{i,t} - \alpha - \hat{z}_{i,t})\right] \\
\hat{\sigma}_{0,z,t+1}^2 &= \sigma_\nu^2 + \rho^2 \frac{\hat{\sigma}_{z,t}^2 \sigma_\epsilon^2}{\hat{\sigma}_{z,t}^2 + \sigma_\epsilon^2},
\end{aligned} \tag{8}$$

where the subscript 0 in $\hat{z}_{0,i,t+1}$ and $\hat{\sigma}_{0,z,t}^2$ captures the posterior before receiving credit shocks. The posterior distribution of income is then $\log y_{0,i,t+1} \sim N(\alpha + \hat{z}_{0,i,t+1}, \hat{\sigma}_{0,z,t+1}^2 + \sigma_\epsilon^2)$.

B. Consumer Preferences

Household preferences are similar to those studying consumer credit and default (e.g. Chatterjee et al. (2007) and Livshits et al. (2007)). Consumers at time 0 maximize their expected life-time utility as

$$U_{i,0} = E_0\left[\sum_{t=0}^{T-1} \delta^t \frac{c_{i,t}^{1-\gamma}}{1-\gamma} + \delta^T \frac{w_{i,T}^{1-\gamma}}{1-\gamma} | I_{i,0}\right],$$

where for each period t, $w_{i,t}$ is the consumers' total saving. The frequency t is set to be one year to be consistent with the average frequency of changes in the credit limit. The budget constraint in each period t is

$$\begin{aligned}
w_{i,t+1} &= \begin{cases} (1 + r_{i,t})(w_{i,t} - c_{i,t}) + y_{i,t+1} & \text{if } d_{i,t} = 0 \\ (1 - \chi) y_{i,t+1} & \text{if } d_{i,t} = 1 \end{cases} \\
w_{i,t} &\geq -(1 - d_{i,t-1}) l_{i,t},
\end{aligned} \tag{9}$$

where $d_{i,t}$ is a default indicator, $l_{i,t}$ is the credit limit, and $\chi \in [0,1]$ is the marginal rate of garnishment. (9) states that when consumers do not default, their wealth in the next period is the sum of their income and gross saving. When the consumers default, their saving becomes zero; at the same time, they need to pay for a garnishment cost equal to χ times their income in the next period. In addition, after defaulting, their credit limit becomes zero, and they can no longer borrow. The interest rate is different for saving and borrowing and takes the value

$$r_{i,t} = \begin{cases} r_b & \text{if } w_{i,t} < 0 \\ r_s & \text{if } w_{i,t} \geq 0. \end{cases}$$

C. Bank

This economy contains a monopolistic bank. The bank takes interest rates as given.[26] At the beginning of each period, the bank chooses the credit limit for each consumer to maximize expected profit.

The bank doesn't observe consumers' total saving. Instead, it perceives their total saving at time t as

$$\tilde{w}_{i,t} = \omega \cdot w_{i,t},$$
$$\omega \sim N(1, \sigma_\omega^2).$$

At the beginning of each period t, after the realization of $\log y_{i,t}$ and before consumers make decisions, the bank receives a noisy signal about $\log y_{i,t+1}$. The signal is

$$\log y_{B,i,t+1} = \mathbb{E}[\log y_{i,t+1}] + \xi_{i,t+1},$$

[26]During the experimental period, the interest rates of the credit card debt at the bank are five basis points daily for everyone.

where $\xi_{i,t+1}$ is normally distributed with $\mathbb{E}[e^{\xi_{i,t+1}}] = 1$ and variance σ_ξ^2.[27]

Given the bank's signal and consumers' policy functions after observing the bank's decision, the bank's problem follows:

$$\max_{l_{i,t}} \Pi = \mathbb{E}\left[\kappa \cdot \tilde{c}_{i,t}(l_{i,t}; \tilde{\theta}_{i,t}) + (1 - d_{i,t}(l_{i,t}; \tilde{\theta}_{i,t})) \cdot r_b \cdot b_{i,t}(l_{i,t}; \tilde{\theta}_{i,t}) - \phi_0 \cdot l_{i,t}^{\phi_1} | I_B\right], \tag{10}$$

where κ is the income from each dollar of transaction using credit cards. $\tilde{c}$ is the total spending the bank could earn transaction income from. I assume $\tilde{c} = q \cdot c_{i,t} + (1 - q) \cdot \min\{c_{i,t}, l_{i,t}\}$ so that, for a proportion q of the time, all transactions are from credit cards. But for the rest of the time, total consumption is bounded by the credit limit, and the excess proportion is transacted using other payment methods for which the bank does not earn income.[28] $\phi_0 \cdot l_{i,t}^{\phi_1}$ is the cost of supplying credit limits of $l_{i,t}$ to the consumers, $\tilde{\theta}_{i,t} = \{\tilde{w}_{i,t}, t, \hat{y}_{i,t+1}\}$ is the collection of consumer state variables from the perspective of the bank, and $\hat{y}_{i,t+1}$ is the consumers' posterior after observing the bank's credit supply.

Optimization yields the supply function

$$\log l_{i,t} = f(\log y_{B,i,t+1}; \theta_{i,t}).$$

I assume $f(\log y_{B,i,t+1}; \theta_{i,t})$ is monotonic in $\log y_{B,i,t+1}$.

D. Learning from Bank Supply

The consumers treat the credit-expansion decision as a signal to $z_{i,t+1}$. From the consumers' perspective,

$$\mathbb{E}_C[\log y_{B,i,t+1}] = \eta \cdot f^{-1}(\log l_{i,t}; \theta_{i,t}) \equiv \eta \cdot g(\log l_{i,t}; \theta_{i,t}), \tag{11}$$

[27] A difference between the model and the experimental setting is that in the experiment, consumers could choose to accept the offer or not. However, for simplicity, I assume the agents in the model have to accept the credit-limit changes. When matching the model with the data on the relationship between credit supply and consumers' expectations, I focus on the average treatment effects on the treatment group as a whole, instead of on either those who accept or do not accept the offers.

[28] This functional form is to capture the scenario in which some large transaction is bounded by the credit limit. Consumers in this case would shift from using credit cards to other payment methods. In addition, I use this functional form to match bank total profits from transactions to total spending.

where η is a misperception parameter. When $\eta = 1$, consumers have the right perception of the bank's inverse supply function. As a result, $\mathbb{E}_C[\log y_{B,i,t+1}] = \log y_{B,i,t+1}$, and consumers have rational expectations. Then the problem is characterized by a signal-jamming equilibrium *a la* Stein (1989). (11) says that when receiving a credit limit of $\log l_{i,t}$, consumers do not know the true belief of the bank; rather, they believe the bank's belief about their income in the next period is η times the bank's true belief. Therefore, η captures all perception errors in the inference process.

After seeing the bank's offer, the consumers Bayesian update their knowledge and form the following belief:

$$\begin{aligned}
\hat{z}_{i,t+1} &= \hat{z}_{0,i,t+1} + \frac{\hat{\sigma}^2_{0,z,t+1}}{\hat{\sigma}^2_{0,z,t+1} + \sigma^2_\xi}(\eta \log y_{B,i,t+1} - \mathbb{E}[\log y_{0,i,t+1}]), \\
\hat{\sigma}^2_{z,t+1} &= \frac{\hat{\sigma}^2_{0,z,t+1}\sigma^2_\xi}{\hat{\sigma}^2_{0,z,t+1} + \sigma^2_\xi}.
\end{aligned} \tag{12}$$

The posterior distribution of the next-period income is then $\log y_{i,t+1} \sim N(\alpha + \hat{z}_{i,t+1}, \hat{\sigma}^2_{z,t+1} + \sigma^2_\epsilon)$.

D.1 Consumer Problems

Given an overall state $\theta_{i,t} = (w_{i,t}, t, \hat{y}_{i,t})$, the consumer value function at t is

$$V(\theta_{i,t}) = \max\left\{V_D(\theta_{i,t}), V_N(\theta_{i,t})\right\}. \tag{13}$$

The continuation value from defaulting is

$$V_D(\theta_{i,t}) = \max_{c_{i,t}} \frac{c_{i,t}^{1-\gamma}}{1-\gamma} + \delta\,\mathbb{E}[V((1-\chi) \times y_{i,t+1}, 0, t+1)|I_{i,t}], \tag{14}$$

and the continuation value from not defaulting is

$$V_N(\theta_{i,t}) = \max_{c_{i,t}} \frac{c_{i,t}^{1-\gamma}}{1-\gamma} + \delta\,\mathbb{E}[V(w_{i,t+1}, l_{i,t+1}, t+1)|I_{i,t}]. \tag{15}$$

E. Results

I study consumer behavior in the model at age 38 so that the age matches that of the average participant in the data. Table 7 gives the estimated parameters.[29] Panel A presents the parameters estimated in the first stage. Estimated η is on average 3.5, indicating consumers' perception about the relationship between bank belief and credit supply is around 3.5 times of the true relationship. The interest rates are set to the averages in the data, which is 19.7% for borrowing and 2% for saving. $\kappa = 0.025$. Thus, for \$1 of transaction from credit cards, the bank earns \$0.025.

Panel B gives the parameters estimated in the second stage, and Panel C shows the matched moments. As shown in Panel C, SMM is successful in matching the moments. Consumers in the data have a risk aversion of 2.46. The marginal rate of garnishment is 0.375. Therefore, in case of default, consumers begin with zero saving and incur a cost that is equal to 37.5% of their income in the next period. $\phi_1 = 0.71$ implies the cost function of the credit limit is slightly concave.[30]

F. Counterfactual Analysis

For the counterfactual analysis, I first study the credit supply and MPC^{L} at age 38 when incorporating the income-inference mechanism. Specifically, I induce a one-time shock of \$666 dollars to bank beliefs about consumer income at age 39 and study the credit supply and spending at age 38. I select the magnitude to match the average changes in a credit limit of around \$1,850 in the experiment. The first row of Panel A in Table 8 gives the results. Columns (1) and (2) give the resulting changes in credit-limit spending. A \$666 shock to bank beliefs increases the credit limit by \$1,833, increasing total spending by roughly \$647. This increase implies an MPC^{L} of around 0.353.

The estimates in the experiment do not directly provide an empirical counterpart of

[29] A detailed description of the model solution and estimation is in Online Appendix section C.

[30] Figure A.8 in the Online Appendix plots the average policy function as a function of the consumer wealth-to-income ratio at age 38. In general, spending increases with wealth-to-income ratio, and debt and default rate decrease with wealth-to-income ratio. The credit limit also decreases in liquidity, because high liquidity is associated with low borrowing. Even though the default probability decreases with liquidity, the low default rate in general and the high interest rate imply profitability is higher due to high debt net of default when liquidity is low. Therefore, the credit supply is high when liquidity is low.

MPC^{L}. However, the estimates of MPC^{LN} for the accepters and non-accepters provide a range of MPC^{L}. Suppose that non-accepters' consumption would be larger if they have to accept the offers. In addition, assume that the accepters accept the offers because they want to consume more. Therefore, the lower bound of MPC^{L} is the MPC^{LN} for the whole sample, and the upper bound is that of the accepters. Therefore, a guess of the range of MPC^{L} is 0.343 and 0.451, which contains the estimates of around 0.4 in the previous literature (Agarwal et al., 2017; Aydin, 2022), and also the 0.353 from the structural estimation. Therefore, the model is successful in generating the average spending responses to limit extensions.

I then consider the case in which consumers' action of learning from the bank is shut down. To do so, I increase the credit limit by the same amount as when there is a shock of \$666 to bank beliefs at age 39. However, I assume the consumers do not update their expectations based on bank action but treat limit changes as exogenous changes. The second row of Panel A in Table 8 gives the results. When the income-inference channel is shut down, the same increase in the credit limit increases total spending by roughly \$425. This finding implies an MPC^{L} of around 0.232. Therefore, when the income-inference channel is shut down, MPC^{L} decreases by around 34%, roughly matching the range of weights from the empirical estimates.

For the second counterfactual, I study the scenario when the variance of the signals the bank receives decreases by 25% from 0.077 to 0.096. The latter value is similar to the estimated value based on the consumers whose income data are not observed by the bank. This analysis seeks to provide insights into the equilibrium credit supply and consumption when banks have more accurate information-processing technology, and also shed lights on a policy that shares consumer income data from the administrative agency to the banks. Panel B of Table 8 presents the results. When σ_ξ^2 increases by 25% from 0.077 to 0.096, equilibrium credit supply and spending at age 38 respectively decrease by 14.11% and 4.11%. Therefore, when advancement of information technology enables banks to extract more precise signals about household future income, credit supply and spending tend to increase. The elasticity of spending with respect to bank-signal precision is around 0.16.

V Conclusion

This study provides the first set of results about the effects of credit-limit expansions on consumers' beliefs. The combination of bank account data, survey data, and an RCT provides a clean identification of the effects. The results show credit expansions strongly increase consumers' expectations regarding their future income growth. At the same time, consumers become more optimistic about their future income growth, increase their spending and borrowing, and have a higher default rate. The identified income-inference channel accounts for around 35% of MPC^{L} and MPB. The findings have important implications for the micro-level mechanism about why consumers have large spending and debt responses to credit-limit extensions. In addition, this study provides accurate estimation of MPB, MPC^{L}, and MPC out of one-time wealth shocks, thereby suggesting the effective design of monetary and fiscal polices.

Future analysis could investigate several avenues. First, the study here is based on an RCT with only a cross section of study. Given that the income-inference channel has different weights during periods with different levels of uncertainty, a potential topic for future research is the strength of the income-inference channel across business cycles. Moreover, the experiment involved only credit expansions and not contractions. Consumers' learning from credit supply might involve other behavioral biases that can cause asymmetric responses. For example, suppose consumers have motivated beliefs about their earning potential (Bénabou and Tirole, 2002; Kőszegi, 2006; Zimmermann, 2020). Then they would overweight the income component in banks' credit-supply rules when facing increases in the credit limit but underweight the income component when facing decreases. This mechanism is similar to holding some form of optimal expectation (Brunnermeier and Parker, 2005; Oster et al., 2013) or having self-attribution bias (Kahneman and Tversky, 2000; Gilovich et al., 2002; Baker et al., 2004) about future abilities. Therefore, a potential direction for future research is to study the asymmetric responses of consumers' beliefs respective to both credit expansions and contractions.

This paper also sheds light on how we should think about information asymmetry in the credit market. Conventionally, the credit market is characterized with adverse

selection from the borrower side. That is, lenders always suffer from an informational disadvantage. However, the findings in this paper show that lenders could extract independent signals about borrower future economic activities. This could potentially invert the adverse problems from the borrowers to the lenders. This is closely related to the situation of inverse selection, as recently dubbed by Brunnermeier et al. (2021). What's more, I find that consumers are over-optimistic about how positive signals from lenders could signal about their future income growth. This is to say, even if lenders may not be able to extract independent information about the borrowers, over-confident borrowers could still believe that positive lender decisions signal good information about themselves. Therefore, featuring borrower learning from lender decisions would greatly alter the optimal contract design in the credit market.

References

Abdulkadiroğlu, A., J. Angrist, and P. Pathak (2014). The elite illusion: Achievement effects at boston and new york exam schools. *Econometrica 82*(1), 137–196.

Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2017, 07). Do Banks Pass through Credit Expansions to Consumers Who want to Borrow? *The Quarterly Journal of Economics 133*(1), 129–190.

Agarwal, S., J. C. Driscoll, X. Gabaix, and D. Laibson (2013, April). Learning in the credit card market. Technical report.

Ameriks, J., G. Kézdi, M. Lee, and M. D. Shapiro (2020). Heterogeneity in expectations, risk tolerance, and household stock shares: The attenuation puzzle. *Journal of Business & Economic Statistics 38*(3), 633–646.

Attanasio, O. P. and G. Weber (2010, September). Consumption and saving: Models of intertemporal allocation and their implications for public policy. *Journal of Economic Literature 48*(3), 693–751.

Aydin, D. (2022, January). Consumption response to credit expansions: Evidence from experimental assignment of 45,307 credit lines. *American Economic Review 112*(1), 1–40.

Baker, M., R. S. Ruback, and J. Wurgler (2004, November). Behavioral corporate finance: A survey. Working Paper 10863, National Bureau of Economic Research.

Baugh, B., I. Ben-David, H. Park, and J. A. Parker (2021, January). Asymmetric consumption smoothing. *American Economic Review 111*(1), 192–230.

Belmi, P., M. A. Neale, D. Reiff, and R. Ulfe (2020). The Social Advantage of Miscalibrated Individuals: The Relationship between Social Class and Overconfidence and Its Implications for Class-based Inequality. *Journal of Personality and Social Psychology 118*(2), 254â€“282.

Bénabou, R. and J. Tirole (2002, 08). Self-Confidence and Personal Motivation. *The Quarterly Journal of Economics 117*(3), 871–915.

Bentley, F., N. Daskalova, and B. White (2017, 05). Comparing the reliability of amazon mechanical turk and survey monkey to traditional market research surveys. *CHI '17: CHI Conference on Human Factors in Computing Systems*, 1092–1099.

Bhandari, G. and R. Deaves (2006). The demographics of overconfidence. *Journal of Behavioral Finance 7*(1), 5–11.

Bordalo, P., N. Gennaioli, R. L. Porta, and A. Shleifer (2019). Diagnostic expectations and stock returns. *The Journal of Finance 74*(6), 2839–2874.

Bouguen, A., Y. Huang, M. Kremer, and E. Miguel (2019). Using randomized controlled trials to estimate long-run impacts in development economics. *Annual Review of Economics 11*(1), 523–561.

Brunnermeier, M. K., R. Lamba, and C. Segura-Rodriguez (2021, December). Inverse selection. Available at SSRN: https://ssrn.com/abstract=3584331.

Brunnermeier, M. K. and J. A. Parker (2005, September). Optimal expectations. *American Economic Review 95*(4), 1092–1118.

Campbell, J. Y. and J. H. Cochrane (1999). By force of habit: A consumptionâ€based

explanation of aggregate stock market behavior. *Journal of Political Economy 107*(2), 205–251.

Chassang, S., E. Snowberg, B. Seymour, and C. Bowles (2015, June). Accounting for behavior in treatment effects: New applications for blind trials. *PLOS One*.

Chatterjee, S., D. Corbae, M. Nakajima, and J.-V. RÃos-Rull (2007). A quantitative theory of unsecured consumer credit with risk of default. *Econometrica 75*(6), 1525–1589.

Chava, S., R. Ganduri, N. Paradkar, and L. Zeng (2020, June). Shocked by bank funding shocks: Evidence from consumer credit cards. Technical report.

Christiano, L., M. Eichenbaum, and C. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy 113*(1), 1–45.

Coibion, O., Y. Gorodnichenko, and M. Weber (2020, May). The cost of the covid-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending. Working Paper 27141, National Bureau of Economic Research.

D'Acunto, F., T. Rauter, C. Scheuch, and M. Weber (2020, December). Perceived precautionary savings motives: Evidence from digital banking. Technical report.

Daniel, K., D. Hirshleifer, and A. Subrahmanyam (1998). Investor psychology and security market under- and overreactions. *The Journal of Finance 53*(6), 1839–1885.

de Quidt, J., J. Haushofer, and C. Roth (2018, November). Measuring and bounding experimenter demand. *American Economic Review 108*(11), 3266–3302.

Deaton, A. (2010, June). Instruments, randomization, and learning about development. *Journal of Economic Literature 48*(2), 424–55.

Duflo, E. (2020, July). Field experiments and the practice of policy. *American Economic Review 110*(7), 1952–73.

Dynan, K. E. (2000, June). Habit formation in consumer preferences: Evidence from panel data. *American Economic Review 90*(3), 391–406.

Fagereng, A., M. B. Holm, and G. J. Natvik (2021, October). Mpc heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics 13*(4), 1–54.

Farrell, D., F. Greig, and E. Deadman (2020, May). Estimating family income from administrative banking data: A machine learning approach. *AEA Papers and Proceedings 110*, 36–41.

Fudenberg, D. and D. K. Levine (2022, August). Learning in games and the interpretation of natural experiments. *American Economic Journal: Microeconomics 14*(3), 353–77.

Fulford, S. L. (2015). How important is variability in consumer credit limits? *Journal of Monetary Economics 72*, 42–63.

Fuster, A., G. Kaplan, and B. Zafar (2020, 11). What Would You Do with $500? Spending Responses to Gains, Losses, News, and Loans. *The Review of Economic Studies 88*(4), 1760–1795.

Ganong, P. and P. Noel (2019, July). Consumer spending during unemployment: Positive and normative implications. *American Economic Review 109*(7), 2383–2424.

Gathergood, J. and A. Olafsson (2022). The co-holding puzzle: New evidence from transaction-level data. Available at SSRN: https://ssrn.com/abstract=3607560.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2021, May). Five facts about beliefs

and portfolios. *American Economic Review 111*(5), 1481–1522.

Gilovich, T., D. Griffin, and D. Kahneman (2002). Heuristics and biases. the psychology of intuitive judgment.

Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the life cycle. *Econometrica 70*(1), 47–89.

Graziani, G., W. van der Klaauw, and B. Zafar (2016, November). Workers' spending response to the 2011 payroll tax cuts. *American Economic Journal: Economic Policy 8*(4), 124–59.

Gross, D. B. and N. S. Souleles (2002, 02). Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data. *The Quarterly Journal of Economics 117*(1), 149–185.

Gross, T., M. J. Notowidigdo, and J. Wang (2020, April). The marginal propensity to consume over the business cycle. *American Economic Journal: Macroeconomics 12*(2), 351–84.

Guerrieri, V. and G. Lorenzoni (2017, 03). Credit Crises, Precautionary Savings, and the Liquidity Trap. *The Quarterly Journal of Economics 132*(3), 1427–1467.

Guvenen, F. (2007, June). Learning your earning: Are labor income shocks really very persistent? *American Economic Review 97*(3), 687–712.

Haaland, I., C. Roth, and J. Wohlfart (2020). Designing information provision experiments. Available at SSRN: https://ssrn.com/abstract=3644820.

Havranek, T., M. Rusnak, and A. Sokolova (2017). Habit formation in consumption: A meta-analysis. *European Economic Review 95*, 142–167.

Hennessy, C. A. and G. Chemla (2022). Signaling, instrumentation, and cfo decision-making. *Journal of Financial Economics 144*(3), 849–863.

Herkenhoff, K., G. M. Phillips, and E. Cohen-Cole (2021). The impact of consumer credit access on self-employment and entrepreneurship. *Journal of Financial Economics 141*(1), 345–371.

Hull, P. (2015). Isolateing: Identifying counterfactual-specific treatment effects by stratified comparisons. *Working Paper*.

Jappelli, T. and L. Pistaferri (2010). The consumption response to income changes. *Annual Review of Economics 2*(1), 479–506.

Jappelli, T. and L. Pistaferri (2014, October). Fiscal policy and mpc heterogeneity. *American Economic Journal: Macroeconomics 6*(4), 107–36.

Jappelli, T. and L. Pistaferri (2020, November). Reported mpc and unobserved heterogeneity. *American Economic Journal: Economic Policy 12*(4), 275–97.

Kahneman, D. and A. Tversky (2000). Choices, values, and frames.

Kőszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association 4*(4), 673–707.

Kirkeboen, L. J., E. Leuven, and M. Mogstad (2016, 05). Field of Study, Earnings, and Self-Selection. *The Quarterly Journal of Economics 131*(3), 1057–1111.

Kline, P. and C. R. Walters (2016, 07). Evaluating Public Programs with Close Substitutes: The Case of Head Start. *The Quarterly Journal of Economics 131*(4), 1795–1848.

Kling, J. R., J. B. Liebman, and L. F. Katz (2007). Experimental analysis of neighborhood effects. *Econometrica 75*(1), 83–119.

Kueng, L. (2018, 06). Excess Sensitivity of High-Income Consumers. *The Quarterly Journal of Economics 133*(4), 1693–1751.

List, J. A. (2011, September). Why economists should conduct field experiments and 14 tips for pulling one off. *Journal of Economic Perspectives 25*(3), 3–16.

Livshits, I., J. MacGee, and M. Tertilt (2007, March). Consumer bankruptcy: A fresh start. *American Economic Review 97*(1), 402–418.

Ludvigson, S. (1999, 08). Consumption and Credit: A Model of Time-Varying Liquidity Constraints. *The Review of Economics and Statistics 81*(3), 434–447.

Mulder, P. (2022, March). Mismeasuring risk: The welfare effects of flood risk information. Technical report.

Mummolo, J. and E. Peterson (2019). Demand effects in survey experiments: An empirical assessment. *American Political Science Review 113*(2), 517–529.

Olafsson, A. and M. Pagel (2018, 04). The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software. *The Review of Financial Studies 31*(11), 4398–4446.

Orne, T. M. (1962). On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist 17*(11), 776–783.

Oster, E., I. Shoulson, and E. R. Dorsey (2013, April). Optimal expectations and limited medical testing: Evidence from huntington disease. *American Economic Review 103*(2), 804–30.

Parker, J. A. and N. S. Souleles (2019, December). Reported effects versus revealed-preference estimates: Evidence from the propensity to spend tax rebates. *American Economic Review: Insights 1*(3), 273–90.

Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland (2013, October). Consumer spending and the economic stimulus payments of 2008. *American Economic Review 103*(6), 2530–53.

Reardon, S. F. and S. W. Raudenbush (2013). Under what assumptions do site-by-treatment instruments identify average causal effects? *Sociological Methods & Research 42*(2), 143–163.

Shapiro, M. D. and J. Slemrod (2003, March). Consumer response to tax rebates. *American Economic Review 93*(1), 381–396.

Stein, J. C. (1989, 11). Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior. *The Quarterly Journal of Economics 104*(4), 655–669.

Vihriälä, E. (2020). Intra-household frictions, anchoring, and the credit card debt puzzle. Available at SSRN: https://ssrn.com/abstract=3710505.

Zeldes, S. P. (1989). Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy 97*(2), 305–346.

Zimmermann, F. (2020, February). The dynamics of motivated beliefs. *American Economic Review 110*(2), 337–61.

Zinman, J. (2009). Debit or credit? *Journal of Banking Finance 33*(2), 358–366.

Zizzo, D. J. (2010). Experimenter demand effects in economic experiments. *Experimental Economics 13*, 75–98.

Figure 1. Expectations of Income Changes

This figure gives the averages of participants' expected income changes over the next six months. Expectation changes are defined as the differences between the answers to survey question 1b and consumer average income over the six months before the experiment. The dots are the averages and the segments are the averages ± two times the standard errors.

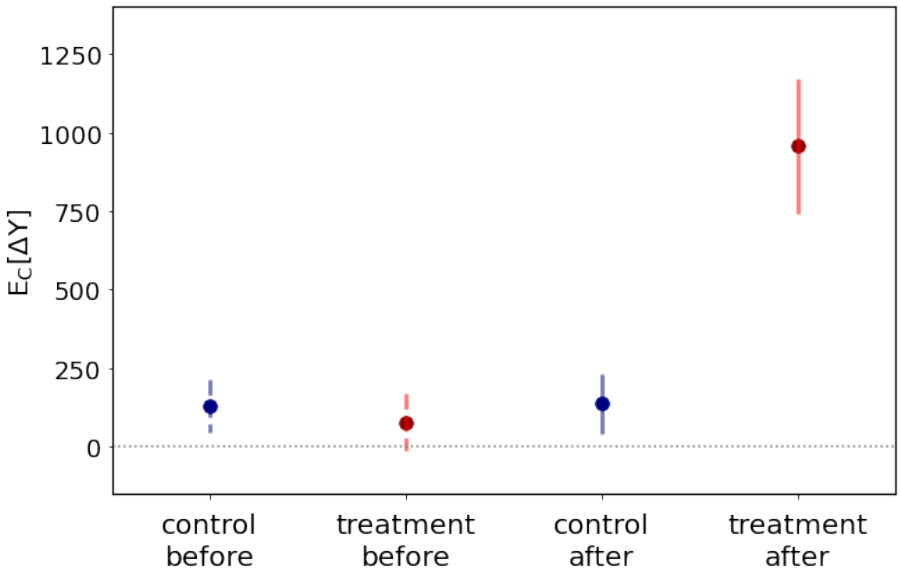

Figure 2. Income Expectation and Pre-determined Credit Limit Changes – By Acceptance

This figure gives the averages of participants' expected income changes over the next six months, separately for the acceptance and non-acceptance group. Expectation changes are defined as the differences between the answers to survey question 1b and consumer average income over the six months before the experiment. The dots are the averages and the segments are the averages ± two times the standard errors. The classification of acceptance group are in Section III.A.

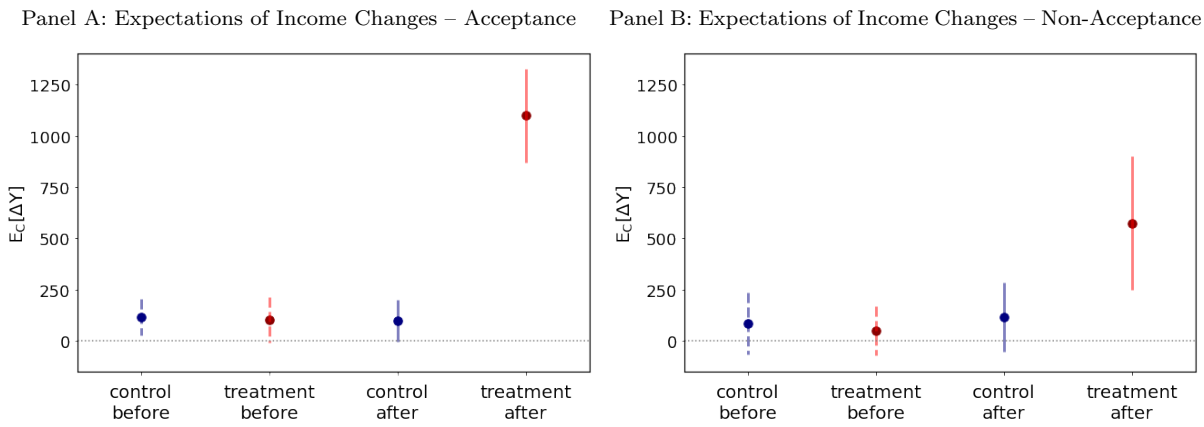

Figure 3. Sensitivity of Income Changes to Credit Limit Extensions – Heterogeneity

This figure assesses the effects of credit expansion on the expectation of future income growth by different groups of consumers' characteristics. For each sorting variables x, I first split the participants into n groups by x, and then fit (1) respectively for the n x-sorted groups. I then plot the coefficients β_1 from fitting the n IV regressions (1) with the average of x in each decile. The four panels sort the participants respectively by their subjective uncertainty, industry uncertainty, deposit-to-income ratio, and utilization ratio. $n = 10$ for deciles for panels A, C, and D, and $n = 18$ for 18 industries for panel B.

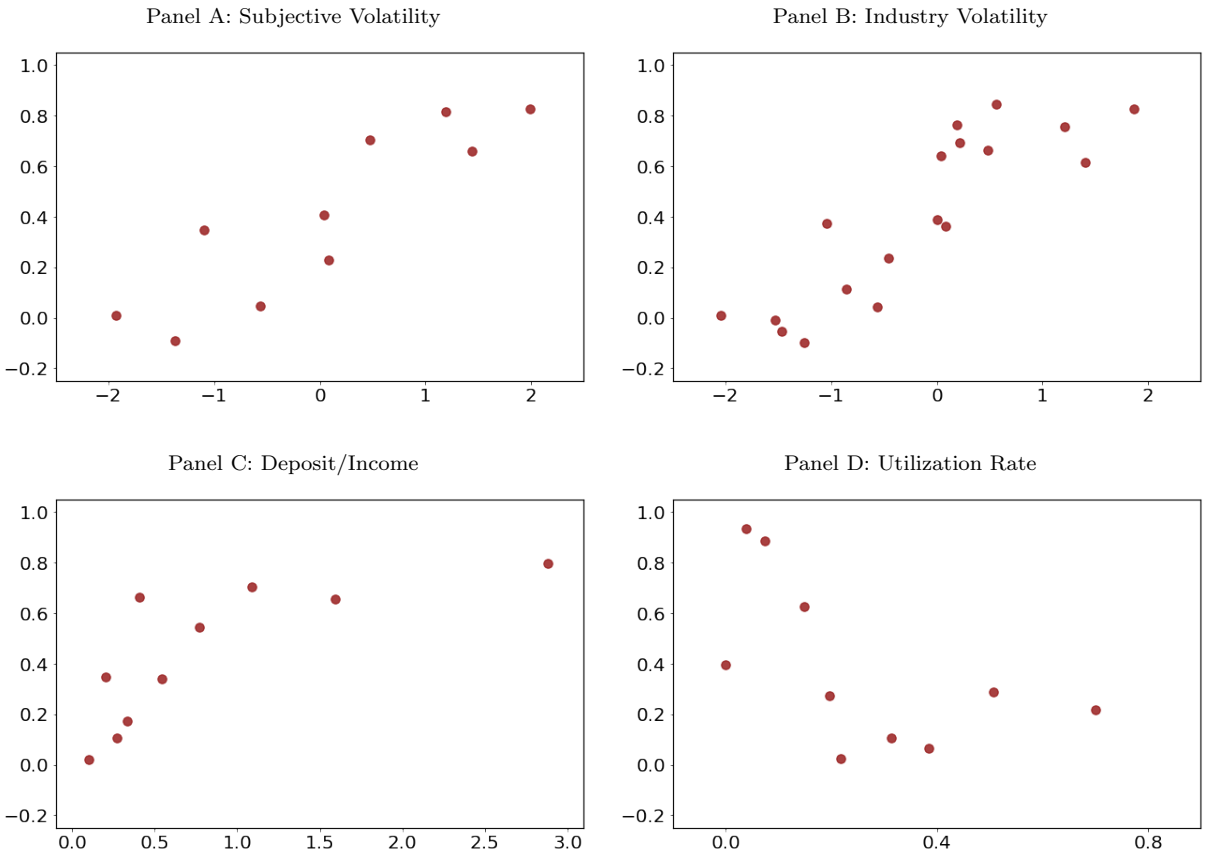

Figure 4. Evolution of Spending and Borrowing

This figure plots the evolution of total interest-incurring debt and cumulative non-debt-financed spending on both sides of the experimental period. In each panel, the x-axis is the number of months away from the experiment. The solid red line shows the evolution of the treatment group, and the blue dotted line shows the evolution of the control group. All lines are vertically shifted so that the value for the control group at time zero is 0.

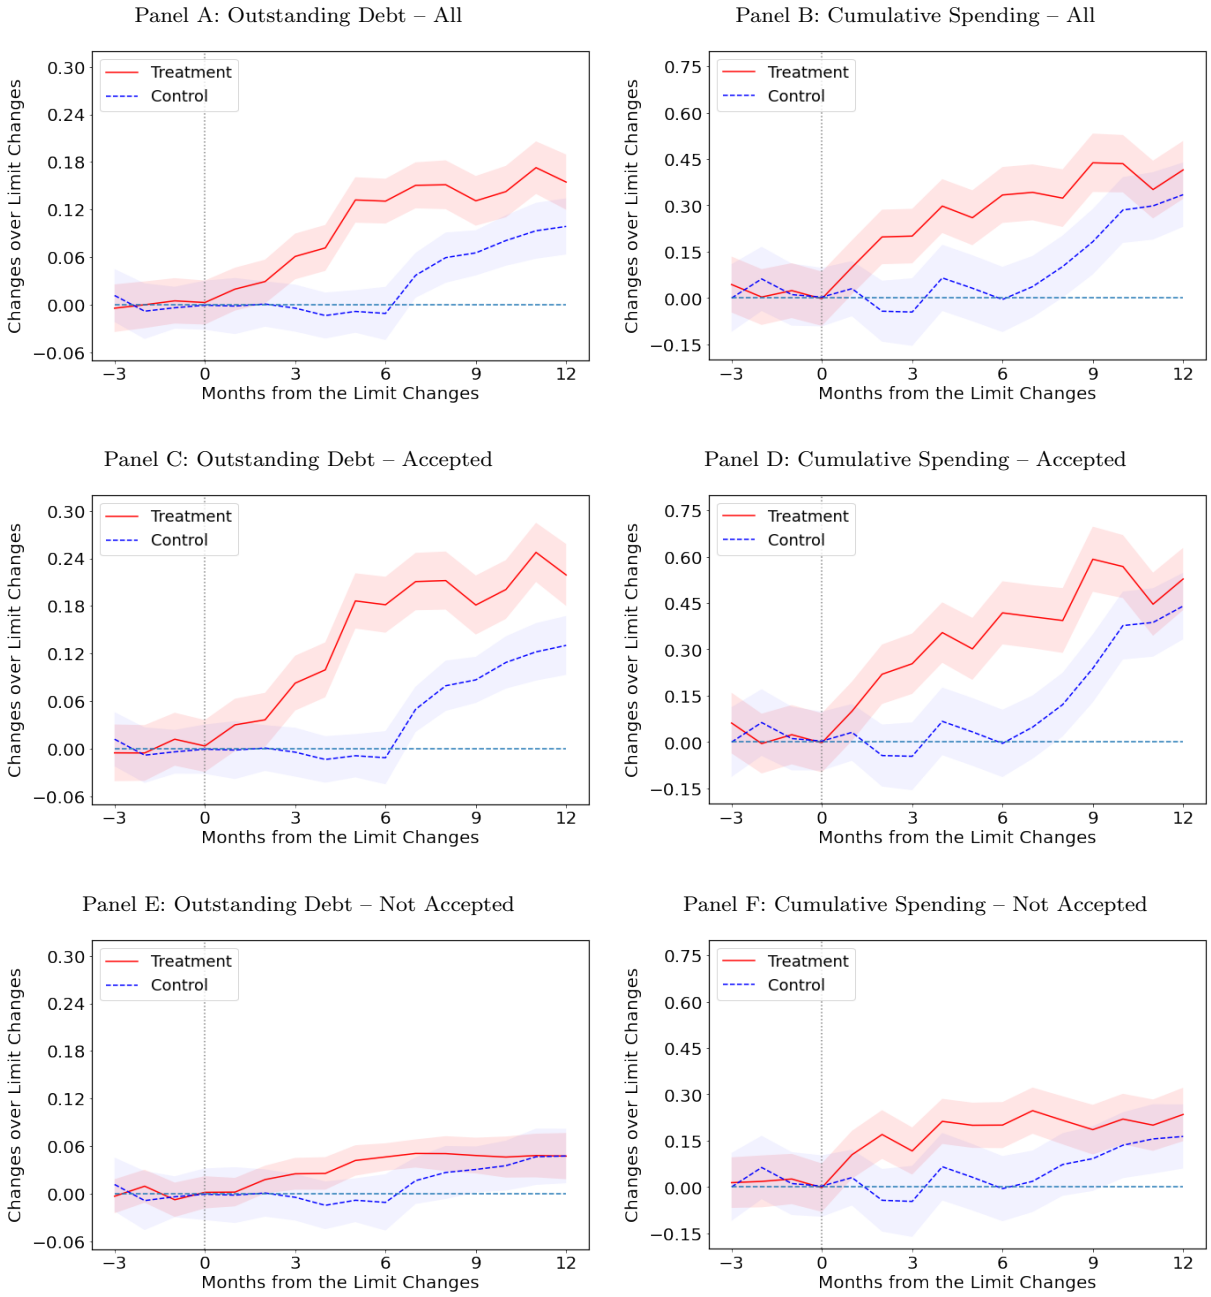

Figure 5. Expectations and Realizations of Income Changes

This figure plots consumer expectations and realized income changes of versus the pre-determined limit changes. The x-axis is the limit changes as proposed before the random assignment. the y-axis of the four panels is consumer pre-experiment expected income changes, realized income changes six months around the experiment, post-experiment expected income changes, and expectation errors after the experiment, respectively. expectation errors are defined as the differences between post-experiment expectation and income realization. All variables are residualized by age, degree, gender, income, saving, total spending, industry fixed effects, city fixed effects.

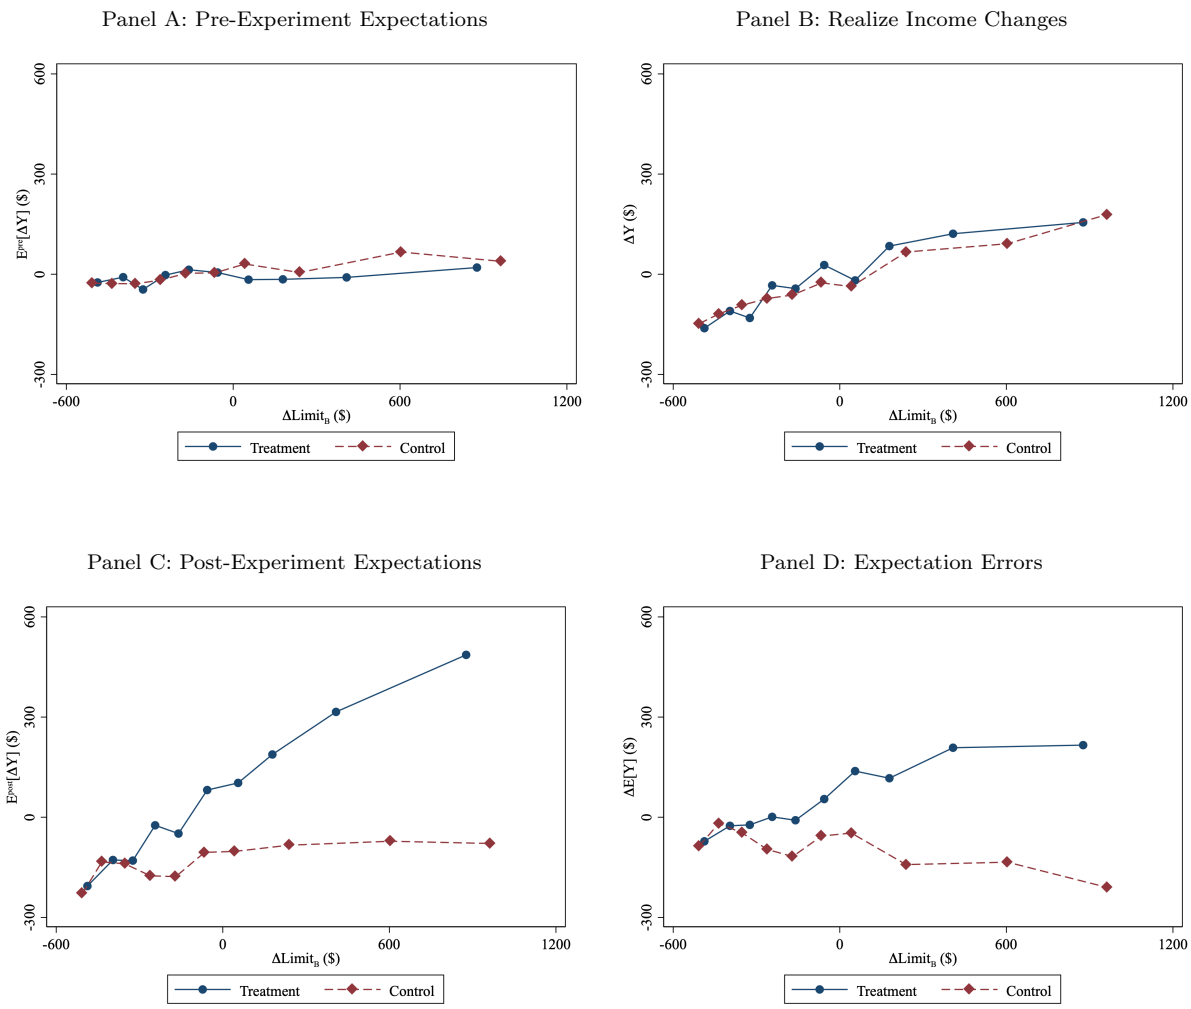

Figure 6. Distribution of λ and η

This figure plots the distribution of consumer subjective sensitivity of income growth to credit extension. Panel A plots λ, and Panel B plots η. λ and η are calculated using (5) and (6). Both variables are winsorized at the 1% level.

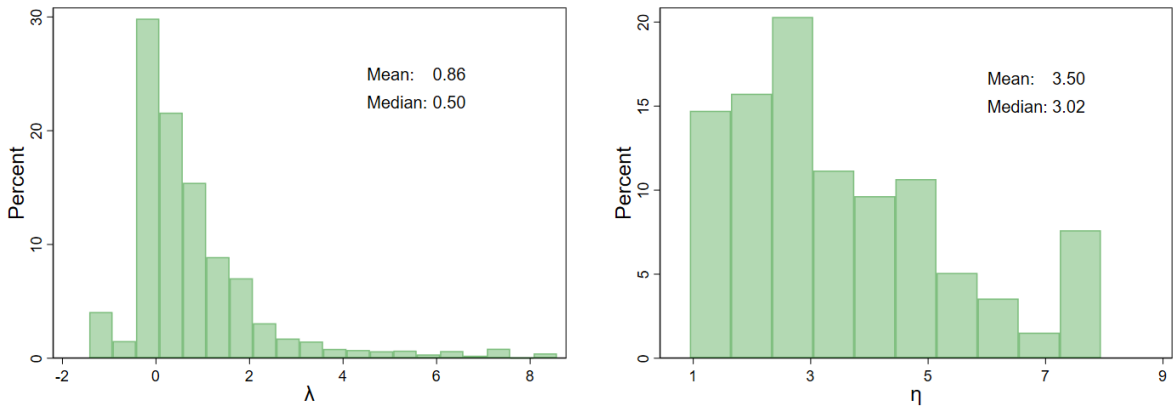

A. Tables

TABLE 1. Summary Statistics

Age is the consumer age immediately before the experimental period. Female is a dummy variable that is 1 if the participant is female. Degree is a categorical variable taking a value from 1 to 5 that labels the participants' highest educational attainment. Hours is the product of the answers to survey questions 2a and 2b; it gives the number of hours the subjects usually work in a week. Saving, Income, Spending, and Debt are the average saving, annual income, annual spending, and interest-incurring outstanding debt, respectively, over the year before the experiment. E[Δ log Income] is the log difference between the answer from survey question 1b and Income, which records the expected 6-month income growth before the experiment. Debt|Debt > 0 gives the summary statistics of Debt conditional on those with positive debt. Limit is the credit card limit at the bank immediately before the experiment. Limit (All) is the answer to survey question 3d. δ is the measure of the discount rate based on survey question 5b. p(Offer) is based on survey question 3b before the experiment. Δ Limit is the proposed change in the credit limit. All level variables are converted to dollars and are winsorized at the 1% level.

	Panel A: Control			Panel B: Treatment			Panel C: Differences	
	Mean	SD	N	Mean	SD	N	Diff/SD	*p*-value
Age	38.65	9.91	2,008	39.25	9.68	3,355	6.02%	0.14
Female	0.49	0.50	2,008	0.52	0.50	3,355	6.00%	0.15
Degree	3.42	1.20	2,008	3.35	1.28	3,355	5.56%	0.84
Hours	41.15	9.72	2,008	40.77	9.60	3,355	-3.95%	0.75
Saving	23904.31	48197.15	2,008	22419.18	40202.07	3,355	-3.24%	0.72
Income	19328.64	17247.00	2,008	19181.16	14445.58	3,355	-0.90%	0.56
Spending	12693.36	21975.24	2,008	13949.28	23578.44	3,355	5.44%	0.17
Debt	959.12	1206.03	2,008	968.91	1472.20	3,355	0.70%	0.45
Debt|Debt>0	2230.74	1378.47	837	2324.46	1431.28	1,379	6.62%	0.23
Limit	6731.76	8413.57	2,008	6416.16	9004.12	3,355	-3.62%	0.74
Limit (All)	13131.76	14161.42	2,008	13216.16	17166.32	3,355	0.54%	0.46
Δ Limit	1874.72	935.89	2,008	1853.32	867.85	3,355	-2.38%	0.66
δ	0.95	0.05	2,008	0.95	0.05	3,355	-5.62%	0.84
p(Offer)	2.32	1.97	605	2.38	2.12	1,001	-2.89%	0.30
Electronically Collected			89.50%			88.11%		

TABLE 2. Income-Inference of Credit Extension

This table assesses the effects of credit expansion on the expectation of future income growth. The specification is based on (1). $\Delta E_C[Y]$ is the difference between the answer of Q1b of the post-experiment survey and that of the pre-experiment survey. $\Delta Limit$ is the proposed changes in credit limit consumers see on the offers. It is positive for those in the treatment group, and zero for those in the control group. In the treatment group, Acceptance (Non-Acceptance) group contains those who accept (do not accept) the offers. In the control group, the classification of acceptance and non-acceptance is explained in section III. A. All variables are winsorized at the 1% level.

	$\Delta E_C[Y]$ All (1)	$\Delta E_C[Y]$ All (2)	$\Delta E_C[Y]$ Acceptance (3)	$\Delta E_C[Y]$ Acceptance (4)	$\Delta E_C[Y]$ Non-Acceptance (5)	$\Delta E_C[Y]$ Non-Acceptance (6)
ΔLimit	0.414***	0.396***	0.492***	0.480***	0.240***	0.240***
	(0.033)	(0.048)	(0.034)	(0.060)	(0.037)	(0.055)
Province FE	N	Y	N	Y	N	Y
First-Stage F	2403.73	1293.45	1040.27	2403.73	1293.45	1040.27
N	5,363	5,363	3,786	3,786	1,577	1,577

Standard errors clustered at city level in parentheses
* p <0.10 ** p <0.05 *** p <0.01

TABLE 3. Borrowing Responses after Credit-Extension Shocks

This table assesses the effects of credit extensions on borrowing. The specification is based on (1) and (2). ΔB is the difference between the total outstanding interest-incurring debt six months after the experiment and that at the beginning of the experiment. $\Delta E_C[Y]$ is the difference between the answer to Q1b in the post-experiment survey and that in the pre-experiment survey. $\Delta Limit$ is the proposed changes in credit limit consumers see on the offers. It is positive for those in the treatment group, and zero for those in the control group. In the treatment group, Acceptance (Non-Acceptance) group contains those who accept (do not accept) the offers. In the control group, the classification of acceptance and non-acceptance is explained in section III. A. All variables are winsorized at the 1% level.

	ΔB All (1)	ΔB Acceptance (2)	ΔB Non-Acceptance (3)	ΔB All (4)	ΔB Acceptance (5)	ΔB Non-Acceptance (6)
ΔLimit	0.127***	0.175***	0.042***	0.073***	0.118***	0.011
	(0.009)	(0.009)	(0.011)	(0.008)	(0.008)	(0.010)
$\Delta E_C[Y]$				0.124***	0.121***	0.129***
				(0.013)	(0.014)	(0.022)
Province FE	Y	Y	Y	Y	Y	Y
First-Stage F	2403.73	1293.45	1040.27	101.33	70.49	19.79
Overid. p-value				0.13	0.30	0.65
N	5,363	3,786	1,577	5,363	3,786	1,577

Standard errors clustered at city level in parentheses
* p <0.10 ** p <0.05 *** p <0.01

TABLE 4. Borrowing Responses after Credit-Extension Shocks – Heterogeneity

This table assesses the effects of credit extensions on borrowing by consumer characteristics before the experiment, focusing on the participants in the acceptance group. The specification is based on (1) and (2). ΔB is the difference between the total outstanding interest-incurring debt six months after the experiment and that at the beginning of the experiment. $\Delta E_C[Y]$ is the differences between the answer of Q1b of the post-experiment survey and that of the pre-experiment survey. $\Delta Limit$ is the proposed changes in credit limit consumers see on the offers. It is positive for those in the treatment group, and zero for those in the control group. The four panels splits the consumers into high and low groups based on their pre-experiment subjective beliefs of income growth uncertainty, within-industry standard deviation of income growth, wealth-to-income, and utilization rate, respectively. In the treatment group, Acceptance group contains those who accept the offers. In the control group, the classification of acceptance is explained in section III. A. All variables are winsorized at 1% level.

	ΔB Low (1)	ΔB High (2)	ΔB Low (3)	ΔB High (4)
	Panel A:Subjective Volatility			
ΔLimit	0.133***	0.208***	0.102***	0.141***
	(0.010)	(0.012)	(0.011)	(0.015)
$\Delta E_C[Y]$			0.105***	0.096***
			(0.019)	(0.021)
First-Stage F	2021.18	741.79	36.60	45.05
Overid. p-value			0.24	0.22
N	1,893	1,893	1,893	1,893
	Panel B: Industry-Level Volatility			
ΔLimit	0.148***	0.214***	0.114***	0.139***
	(0.010)	(0.013)	(0.011)	(0.013)
$\Delta E_C[Y]$			0.104***	0.111***
			(0.022)	(0.015)
First-Stage F	1618.68	504.67	35.01	47.41
Overid. p-value			0.23	0.29
N	2,117	1,669	2,117	1,669
	Panel C: Wealth/Income			
ΔLimit	0.213***	0.142***	0.175***	0.068***
	(0.014)	(0.009)	(0.012)	(0.011)
$\Delta E_C[Y]$			0.117***	0.118***
			(0.020)	(0.018)
First-Stage F	1056.36	980.63	55.16	45.84
Overid. p-value			0.86	0.18
N	1,893	1,893	1,893	1,893
	Panel D: Utilization Rate			
ΔLimit	0.252***	0.098***	0.215***	0.031***
	(0.014)	(0.009)	(0.015)	(0.013)
$\Delta E_C[Y]$			0.126***	0.109***
			(0.023)	(0.024)
First-Stage F	1113.43	693.25	61.12	50.63
Overid. p-value			0.77	0.15
N	1,893	1,893	1,893	1,893

Standard errors clustered at city level in parentheses
* p <0.10 ** p <0.05 *** p <0.01

TABLE 5. Spending Responses after Credit-Extension Shocks

This table assesses the effects of credit extension on non-debt-financed spending. The specification is based on (1) and (2). ΔC is the difference between the total spending minus newly accumulated interest-incurring debt six months after the experiment and that at the beginning of the experiment. $\Delta E_C[Y]$ is the differences between the answer to Q1b in the post-experiment survey and that in the pre-experiment survey. $\Delta Limit$ is the proposed changes in credit limit consumers see on the offers. It is positive for those in the treatment group, and zero for those in the control group. In the treatment group, Acceptance (Non-Acceptance) group contains those who accept (do not accept) the offers. In the control group, the classification of acceptance and non-acceptance is explained in section III. A. All variables are winsorized at the 1% level.

	ΔC All (1)	ΔC Acceptance (2)	ΔC Non-Acceptance (3)	ΔC All (4)	ΔC Acceptance (5)	ΔC Non-Acceptance (6)
ΔLimit	0.216***	0.276***	0.108**	0.132***	0.174***	0.042
	(0.026)	(0.031)	(0.052)	(0.031)	(0.036)	(0.067)
$\Delta E_C[Y]$				0.243***	0.209***	0.275**
				(0.076)	(0.051)	(0.116)
Province FE	Y	Y	Y	Y	Y	Y
First-Stage F	2403.73	1293.45	1040.27	101.33	74.27	19.26
Overid. p-value				0.14	0.20	0.26
N	5,363	3,786	1,577	5,363	3,786	1,577

Standard errors clustered at city level in parentheses
* p <0.10 ** p <0.05 *** p <0.01

TABLE 6. Default Probability after Credit-Extension Shocks

This table assesses the effects of credit extensions on consumers' default probability. The specification is based on (1) and (2). $Default$ is equal to 100 if there is a 60-day delinquency for the debt taken over the six months after the experiment, and 0 otherwise. $\Delta E_C[Y]$ is the difference between the answer to Q1b in the post-experiment survey and that in the pre-experiment survey. $\Delta Limit$ is the proposed changes in credit limit consumers see on the offers. It is positive for those in the treatment group, and zero for those in the control group. $\Delta Limit$ and $\Delta E_C[Y]$ are multiplied by 1,000. In the treatment group, Acceptance (Non-Acceptance) group contains those who accept (do not accept) the offers. In the control group, the classification of acceptance and non-acceptance is explained in section III. A. All variables are winsorized at the 1% level.

	Default All (1)	Default Acceptance (2)	Default Non-Acceptance (3)	Default All (4)	Default Acceptance (5)	Default Non-Acceptance (6)
ΔLimit	0.108*	0.176**	0.073	-0.052	-0.074	-0.007
	(0.062)	(0.076)	(0.070)	(0.053)	(0.064)	(0.054)
$\Delta E_C[Y]$				0.212***	0.245**	0.167*
				(0.084)	(0.119)	(0.099)
Province FE	Y	Y	Y	Y	Y	Y
First-Stage F	2403.73	1293.45	1040.27	101.33	19.26	19.26
Overid. p-value				0.21	0.43	0.65
N	5,363	3,786	1,577	5,363	1,577	1,577

Standard errors clustered at city level in parentheses
* p <0.10 ** p <0.05 *** p <0.01

TABLE 7. Estimated Parameters

This table gives the estimated parameters of the structural model. Panel A presents the parameters estimated in the first stage, Panel B gives the parameters estimated in the second stage, and Panel C gives the matched moments.

Panel A: First-Stage Parameters		Panel B: Second-Stage Prameters			Panel C: Matched Moments		
	Estimates (1)		Estimates (2)	S.E. (3)		Data (4)	Model (5)
α	-0.180	γ	2.460	(0.137)	w/c	0.866	0.866
ρ	0.940	χ	0.375	(0.006)	$p(default)$	0.023	0.023
σ_ν^2	0.002	ϕ_0	0.043	(0.001)	$\bar{b}/y$	0.650	0.650
σ_ϵ^2	0.088	ϕ_1	0.710	(0.001)	$\partial l/\partial \log y_B$	2.106	2.106
σ_ξ^2	0.077						
η	3.500						
κ	0.025						
q	0.087						
r_b	0.197						
r_s	0.025						
σ_ω^2	0.384						

TABLE 8. Counterfactual Analysis

This table gives the equilibrium condition under different scenarios. Panel A analyzes the weight of the income-inference channel in explaining MPC^L. Panel B analyzes the counterfactual scenario when the precision of the signals the bank receives is larger. Δl and Δc are respectively the changes in credit limit and spending at age 38. l and c are respectively the equilibrium credit limit and spending at age 38. MPC^L is the ratio of the changes in spending to changes in the credit limit. The units of Δl, Δc, l, and c are in thousand dollars.

	Panel A: MPC^L to Bank-Belief Shock		
	Δl	Δc	MPC^L
	(1)	(2)	(3)
Baseline	1.833	0.647	0.353
No Income-Inference	1.833	0.425	0.232
	Panel B: Counter-Factual σ_ξ^2		
	σ_ξ^2	l	c
	(1)	(2)	(3)
Baseline	0.077	12.431	9.876
Hypothetical	0.096	10.677	9.470
Diff. %	25.00%	-14.11%	-4.11%

Online Appendix

for "Learning in the Limit: Income Inference from Credit Extension" by Xiao Yin

A. Households Consumption and Preferences Survey[31]

Credit cards are an important method for daily consumption. **To test their business strategies, banks often randomly select some people to have a change in their credit card limits and see how they change their spending**. To better understand the impact of credit cards on people's lives, we randomly selected a certain number of active credit card users from our bank to participate in a survey. We hope to use this survey to study the consumption behaviors and preferences of the residents generally. Therefore, we will focus only on highly summarized information for scientific research purposes, such as average values. We will not disclose the personal information of the participants in any respect. We will not, in any way, change the types of financial products we provide, including those regarding credit scores, credit limits, deposit rates, etc., based on the participants' personal answers.

1. **Income Information**
   (a) Your total income over the past 6 months is _____.
   (b) Your expected total income over the next 6 months is _____.
   (c) With a probability of 80%, your total income over the next 6 months will be between _____ and _____.
   (d) Your expected total income over the next 12 months is _____.
   (e) With a probability of 80%, your total income over the next 12 months will be between_____ and _____.
2. **Work Information**
   (a) How many hours do you usually work every day? _____ hours
   (b) How many days do you usually work every week? _____ days
   (c) How many hours on average will you work every week over the next 6 months? _____ hours
   (d) How many days on average will you work every day over the next 6 months? _____ days
   (e) How many hours on average will you work every day over the next 12 months? _____ hours
   (f) How many days on average will you work every week over the next 12 months? _____ days
3. **Other Information**
   (a) * How likely is it for you to get a credit card limit increase from the bank over the next 6 months?[32]
      i. Very unlikely
      ii. Unlikely
      iii. Somewhat likely
      iv. Likely
      v. Very likely
   (b) * How likely is it for you to have a binding borrowing limit in the next 6 months.
      i. Very unlikely
      ii. Unlikely
      iii. Somewhat likely

[31] The order of the choices in questions labeled with * was randomized. The sentence in bold was randomly sent to 15% of the subjects.

[32] Sent to a random 30% of the participants.

iv. Likely
v. Very likely

(c) * Why didn't you accept the credit card limit increase offered to you last week? (Ignore if you accepted) [33]

i. Too busy to accept
ii. Afraid of overspending
iii. Other

(d) Your total credit limit across all banks is _____.

(e) Your total saving across all banks is _____.

4. **Hypothetical Questions**[34]

(a) Suppose your bank increases your credit card limit by 5,000 Yuan this month. This would mean that the bank expects your total income to be changed by _____ in the next 6 months.

(b) Suppose your bank increases your credit card limit by 10,000 Yuan this month. This would mean that the bank expects your total income to be changed by _____ in the next 6 months.

(c) Suppose your bank increases your credit card limit by 5,000 Yuan this month. This would mean that the bank expects your total income to be changed by _____ in the next 12 months.

(d) Suppose your bank increases your credit card limit by 10,000 Yuan this month. This would mean that the bank expects your total income to be changed by _____ in the next 12 months.

(e) The bank assigns each customer a credit score to label the relative safeness for granting a loan. What would be the credit score you believe you have at the bank? (Please give a number between 0 and 10, 10 being the safest).

5. **Preferences**

(a) Rather than receiving 100 Yuan today, which of the following options would you choose. (select all that apply)

i. 100 Yuan in six months.
ii. 102.5 Yuan in six months.
iii. 105 Yuan in six months.
iv. 107.5 Yuan in six months.
v. 110 Yuan in six months.
vi. 112.5 Yuan in six months.
vii. 115 Yuan and more in six months.

(b) Rather than receiving 150 Yuan in six months, which of the following options would you choose. (select all that apply)

i. 150 Yuan in twelve months.
ii. 153 Yuan in twelve months.
iii. 156 Yuan in twelve months.
iv. 159 Yuan in twelve months.
v. 162 Yuan in twelve months.
vi. 165 Yuan in twelve months.
vii. 168 Yuan and more in twelve months.

(c) Suppose there is a game. With a 60% probability, you will win 150 Yuan, with a 40% probability you will receive 50 Yuan. Which of the following option would you choose over playing this game. (select all that apply) [35]

i. 70 Yuan for sure.
ii. 75 Yuan for sure.
iii. 80 Yuan for sure.
iv. 85 Yuan for sure.

[33] Only included in the post-experiment surveys.
[34] Sent to a random 50% of the participants.
[35] Question 3 (b) and questions 4 (a) to 4 (h) were only included in the pre-experiment surveys.

v. 90 Yuan for sure.
vi. 95 Yuan for sure.
vii. 100 Yuan for sure.
viii. 105 Yuan for sure.
ix. 110 Yuan for sure.
x. 115 Yuan for sure.
xi. 120 Yuan for sure.

Figure A1. Sanity Check for Survey Answers

This figure displays the basic sanity checks for the survey answers. Panel A is a binned scatter plot of the average income before the experiment calculated from the bank and the average income before the experiment from the administrative department. Panel B is a binned scatter plot of consumer average monthly income over the six months before the experiment and the answers from question 1 (a) of the survey before the experiment.

Panel A

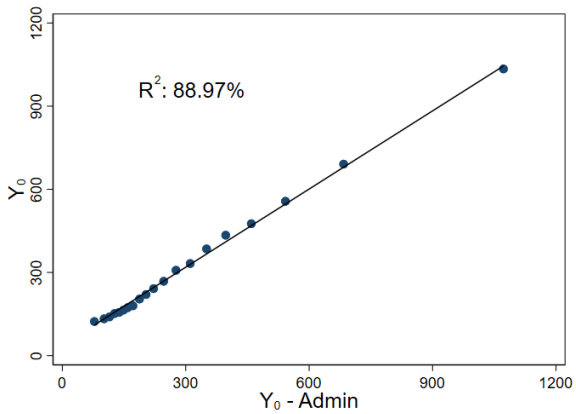

Panel B

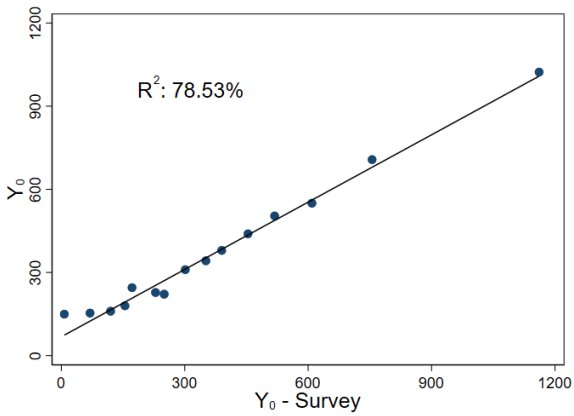

Figure A.2: Distribution of Acceptance Probabilities

This figure plots the distribution of the probabilities from predicting the acceptance decisions in the control group. Panel A focuses on the acceptance group, and Panel B focuses on the non-acceptance group.

Panel A: Acceptance

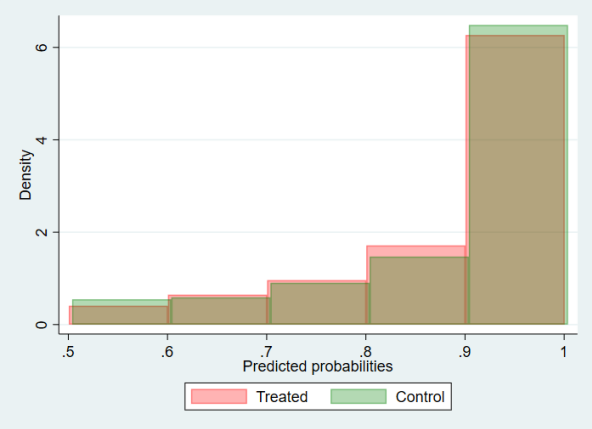

Panel B: Non-Acceptance

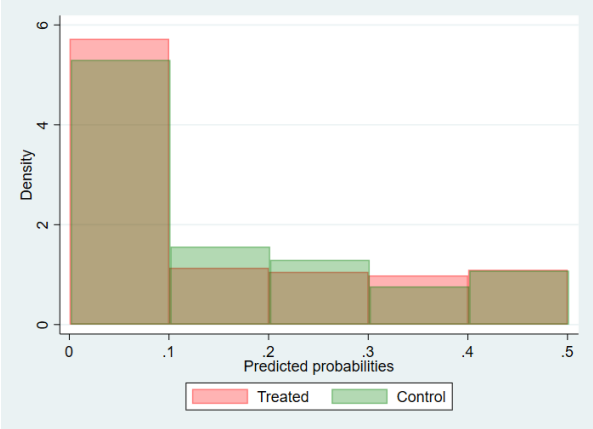

Figure A.3: Evolution of Spending and Borrowing – No Survey

This figure plots the evolution of total interest-incurring debt and cumulative non-debt-financed spending on both sides of the experimental period for those who did not fill out the survey. In each panel, the x-axis plots the number of months away from the experiment. The solid red line shows the evolution of the treatment group, and the blue dotted line shows the evolution of the control group. All lines are vertically shifted so that the value for the control group at time zero is 0.

Panel A: Outstanding Debt – All

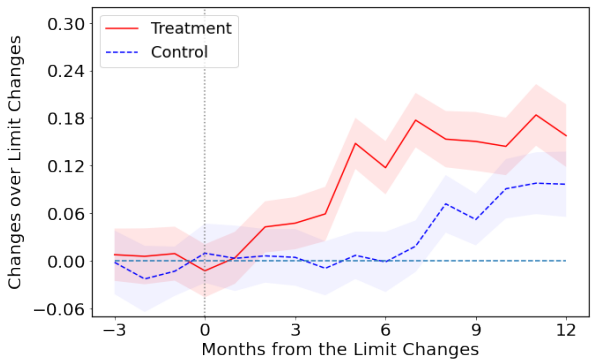

Panel B: Cumulative Spending – All

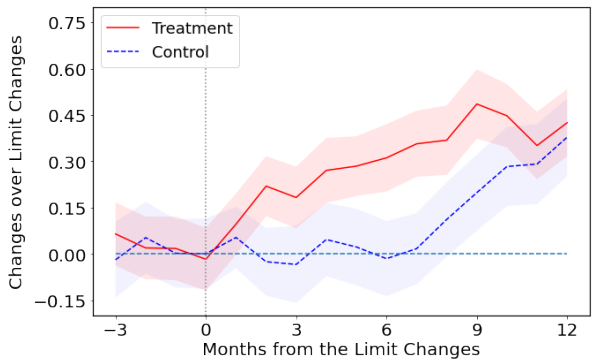

Panel C: Outstanding Debt – Acceptance

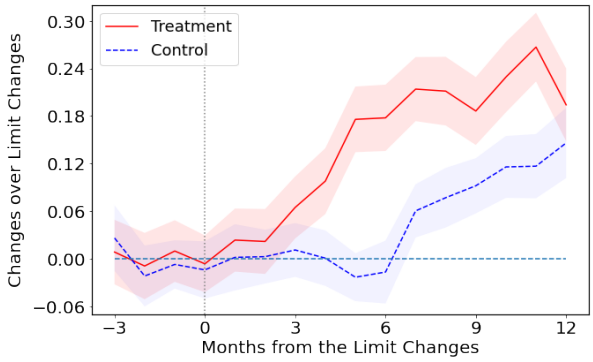

Panel D: Cumulative Spending – Acceptance

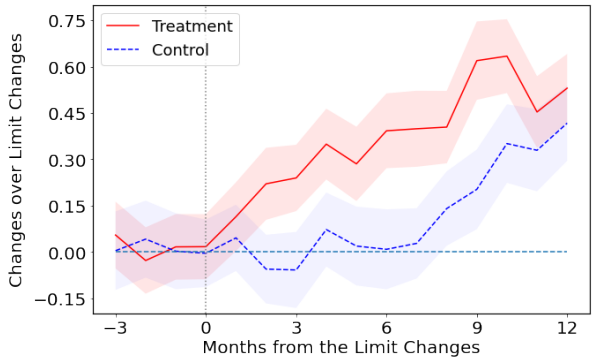

Panel E: Outstanding Debt – Non-Acceptance

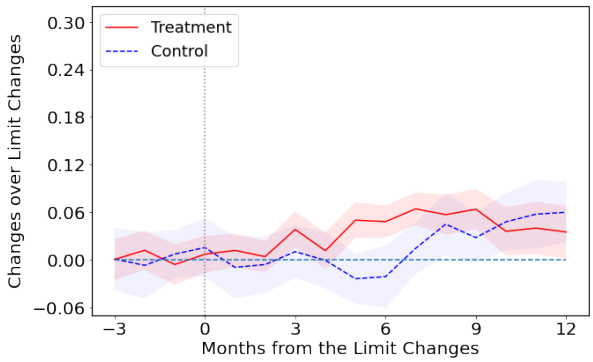

Panel F: Cumulative Spending – Non-Acceptance

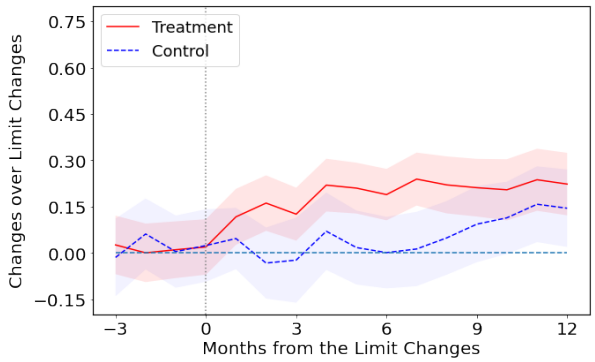

Figure A.4: Evolution of Spending and Borrowing – No Survey 1

This figure plots the evolution of total interest-incurring debt and cumulative non-debt-financed spending on both sides of the experimental period for those who did not fill out the first survey. In each panel, the x-axis plots the number of months away from the experiment. The solid red line shows the evolution of the treatment group, and the blue dotted line shows the evolution of the control group. All lines are vertically shifted so that the value for the control group at time zero is 0.

Panel A: Outstanding Debt – All

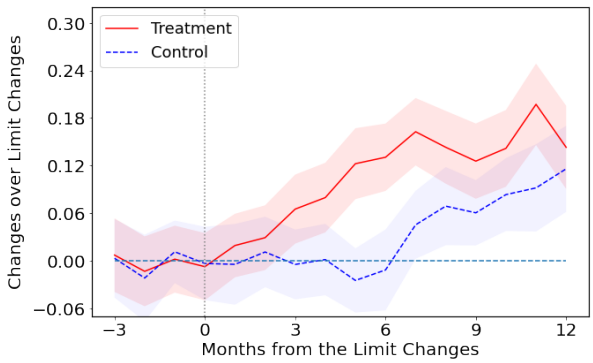

Panel B: Cumulative Spending – All

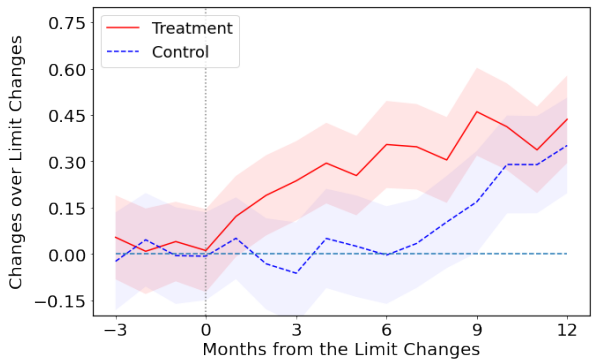

Panel C: Outstanding Debt – Acceptance

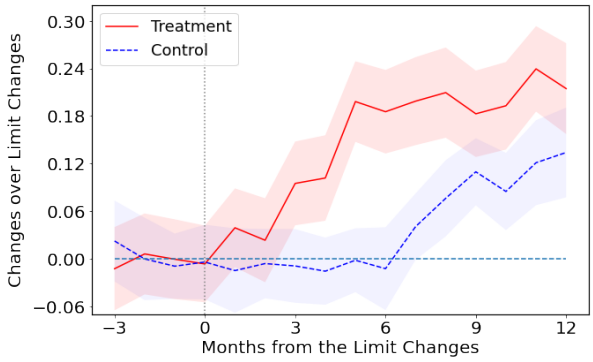

Panel D: Cumulative Spending – Acceptance

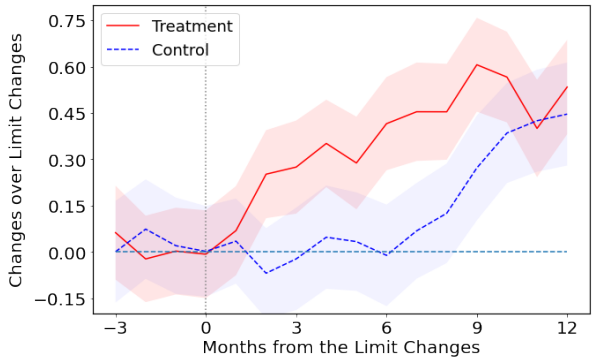

Panel E: Outstanding Debt – Non-Acceptance

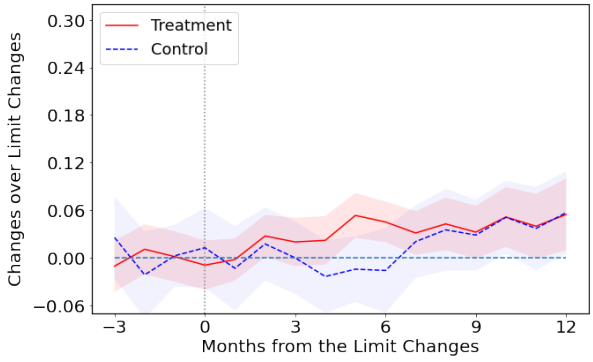

Panel F: Cumulative Spending – Non-Acceptance

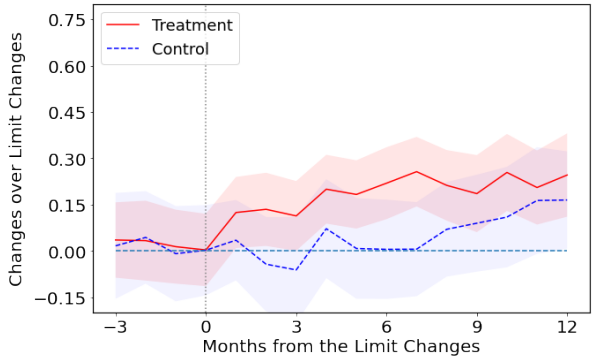

Figure A.5: Evolution of Spending and Borrowing – No Survey 2

This figure plots the evolution of total interest-incurring debt and cumulative non-debt-financed spending on both sides of the experimental period for those who did not fill out the second survey. In each panel, the x-axis plots the number of months away from the experiment. The solid red line shows the evolution of the treatment group, and the blue dotted line shows the evolution of the control group. All lines are vertically shifted so that the value for the control group at time zero is 0.

Panel A: Outstanding Debt – All

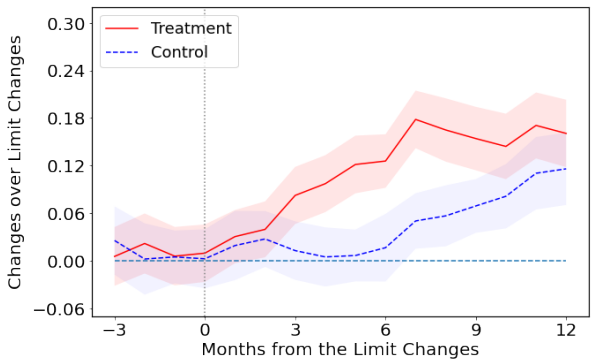

Panel B: Cumulative Spending – All

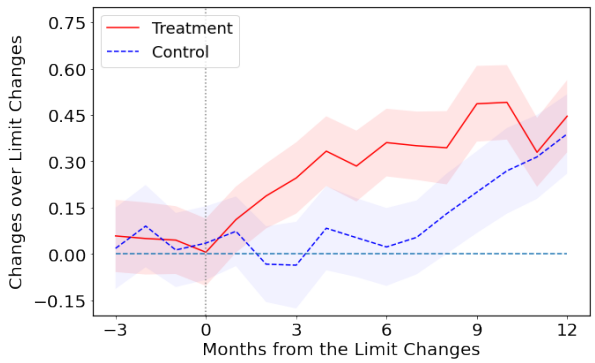

Panel C: Outstanding Debt – Acceptance

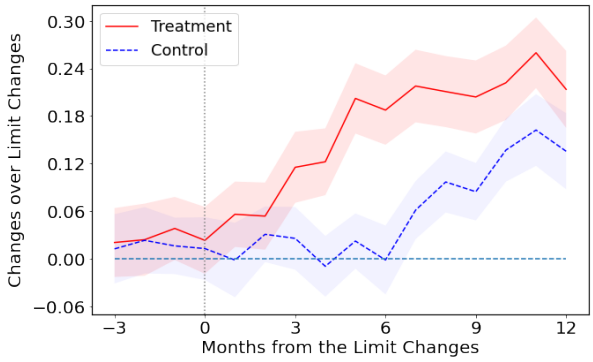

Panel D: Cumulative Spending – Acceptance

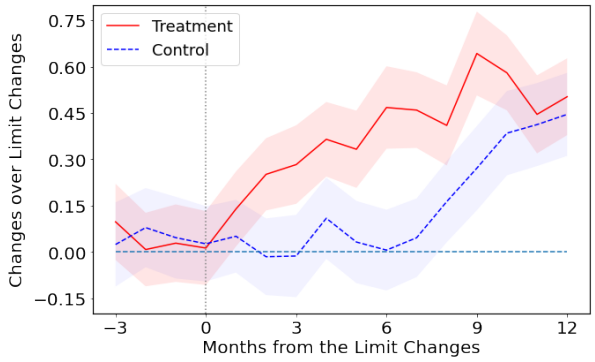

Panel E: Outstanding Debt – Non-Acceptance

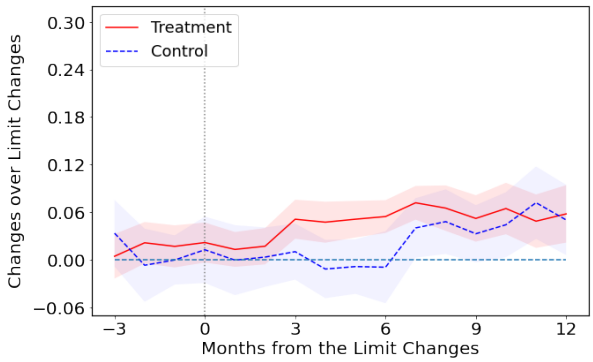

Panel F: Cumulative Spending – Non-Acceptance

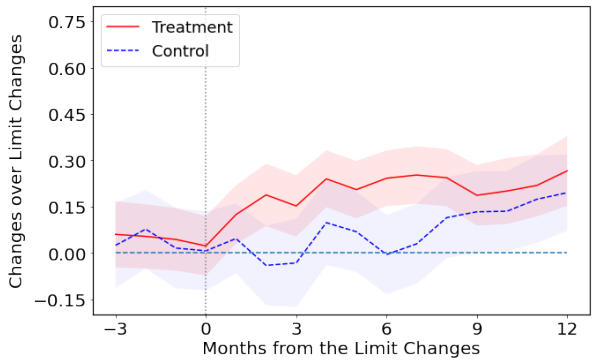

Figure A.6: Changes in Net Transfer

This figure plots the changes in the net transfers from other banks after the experiment. The dots are the averages and the segments are two times the standard errors. The red segments plot the treatment group, and blue segments plot the control group. The solid segments plot the acceptance group, and the dashed segments lot the non-acceptance group.

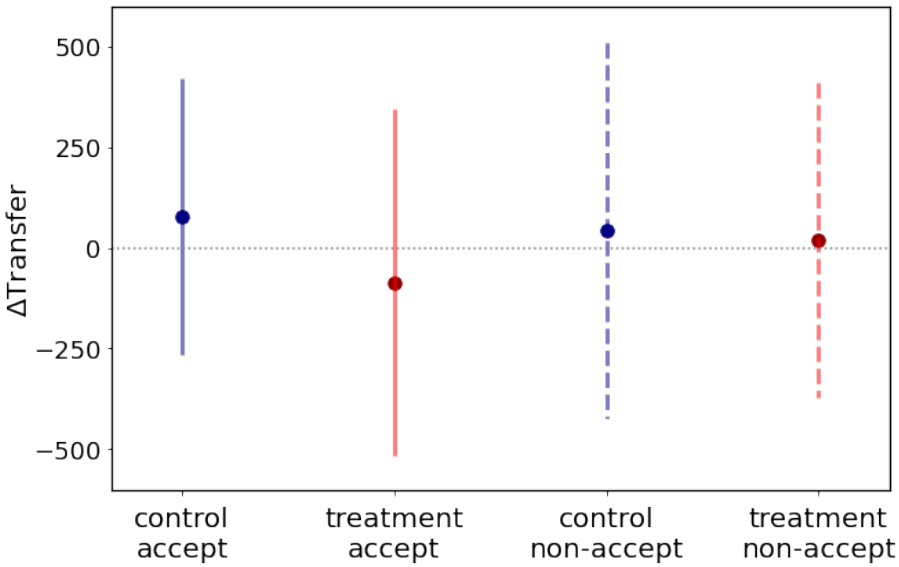

Figure A.7: Reasons to Not Accepting the Offers

This figure plots the answers to Q3c by those who didn't accept the offers in the treatment group.

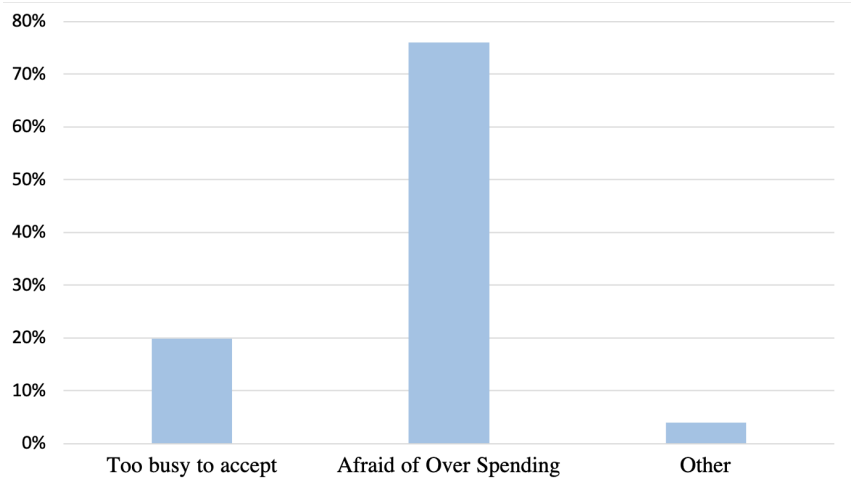

Figure A.8: Policy Functions

This figure plots policy functions of the simulated consumers as functions of the wealth-to-income ratio at age 38. The top-left panel plots consumption; the top-right panel plots debt; the bottom-left panel plots the credit limit; and the bottom-right panel plots the default probability.

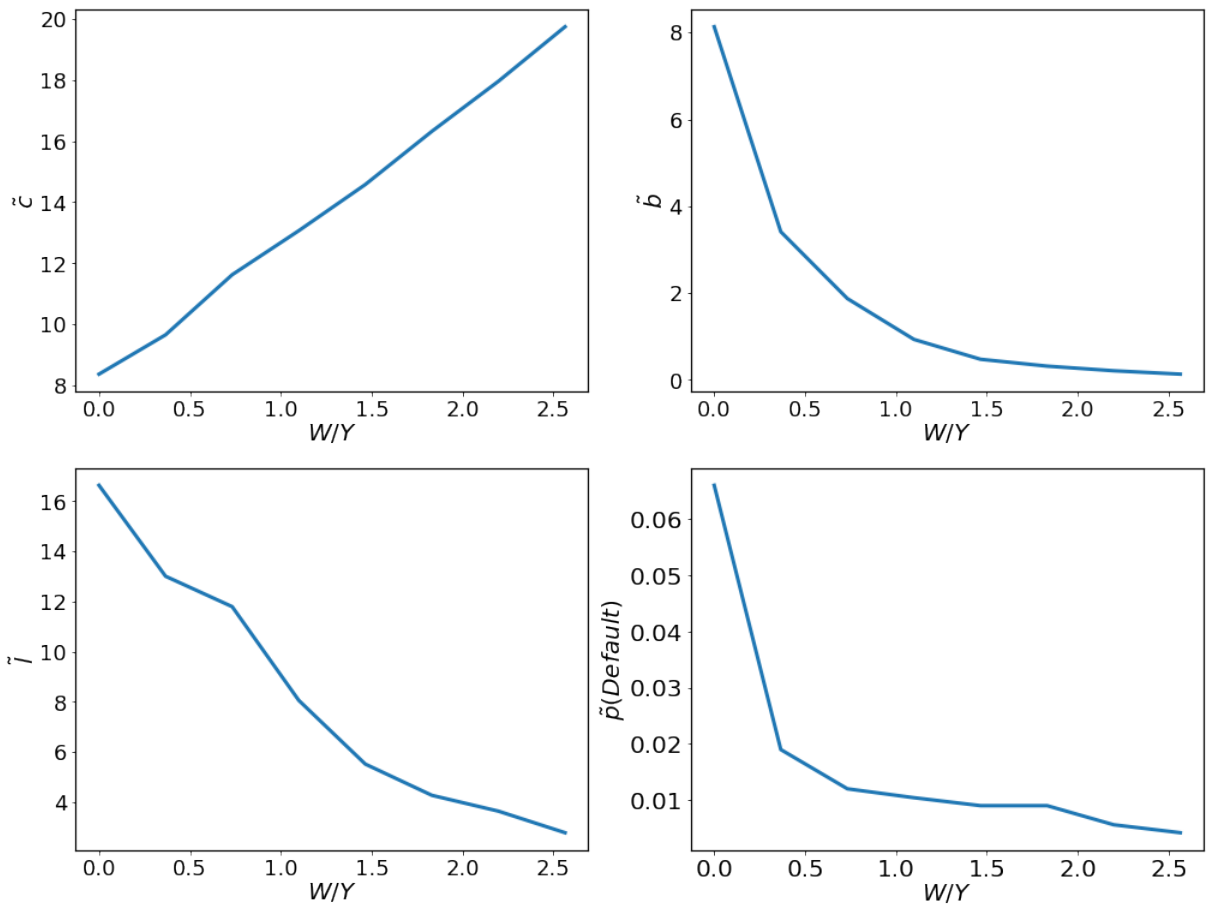

Table A.1: Credit Extension and Expectations – Pilot Study

This table assesses the effects of credit expansion on consumer expectations. The sample is based on a pilot study in March 2019. Column (1) is the expected changes in six-month total spending. Column (2) is the expected changes in six-month total income. Column (3) is the expected changes in total savings six months after the experiment. Column (4) is the expected changes in default probabilities. The experimental setting is the same to that in the main study, while with a smaller sample of 2,500 potential participants. The survey is sent only once after the experiment. The differences are based on the survey question and the true value before the experiments. All variables are winsorized at 1% level.

	$\Delta E_C[C]$ (1)	$\Delta E_C[Y]$ (2)	$\Delta E_C[W]$ (3)	$\Delta E_C[D]$ (3)
ΔLimit	0.372***	0.319***	0.012	-0.018
	(0.114)	(0.099)	(0.085)	(0.125)
Province FE	Y	Y	Y	Y
First-Stage F	332.43	332.43	332.43	332.43
N	1,486	1,486	1,486	1,486

Standard Errors Clustered at City level in Parentheses
* p <0.10 ** p <0.05 *** p <0.01

Table A.2: Confusion Matrix

This table gives the confusion matrix of using the LASSO logistic model to predict the acceptance decisions of the test sample of the treatment group. The LASSO logistic model is first fitted on a random 50% of sample of the treatment group to predict who would accept the offer given pre-experiment characteristics. Model fitting is based on 3-fold cross-validation. The confusion matrix is then based on the out-of-sample prediction using the same model over the test sample. Pre-experiment characteristics include gender, education, age, average income, average saving, average spending, average debt, average hours worked every week, credit score, changes in the credit score, number of credit credits owned, number of credit limit offers received, subjective income volatility, city, short-term and long-term discount rates, expectation about future income, and bank-proposed changes in the credit limits. All variables are before the experiment.

	Realized	
	(1)	(2)
Predicted	Accept	Not Accept
Accept	421	42
Not Accept	179	1,036

Table A.3: Credit Extension and Consumer Behavior – Information Treatment

This table assesses the effects of credit extension on consumer income expectation, borrowing, and spending. The specification is based on (1). ΔB is the difference between the total outstanding interest-incurring debt six months after the experiment and that at the beginning of the experiment. ΔC is the difference between the total spending minus newly accumulated interest-incurring debt six months after the experiment and that at the beginning of the experiment. $\Delta E_C[Y]$ is the differences between the answer of Q1b of the post-experiment survey and that of the pre-experiment survey. $\Delta Limit$ is the proposed changes in credit limit consumers see on the offers. It is positive for those in the treatment group, and zero for those in the control group. In columns (3) to (6), $\Delta Limit$ and $\Delta E_C[Y]$ are instrumented with the random assignment indicator and the information treatment indicators. All variables are winsorized at 1% level.

	$\Delta E_C[Y]$ Accept (1)	$\Delta E_C[Y]$ Not Accept (2)	ΔB Accept (3)	ΔB Not Accept (4)	ΔC Accept (5)	ΔC Not Accept (6)
ΔLimit	0.276***	0.097	0.115***	0.017	0.156***	0.076
	(0.055)	(0.085)	(0.008)	(0.010)	(0.050)	(0.076)
$\Delta E_C[Y]$			0.106*	0.153*	0.237*	0.324*
			(0.059)	(0.081)	(0.131)	(0.183)
Controls	Y	Y	Y	Y	Y	Y
First-Stage F	597.12	278.98	18.97	9.64	12.69	7.32
N	733	312	4,519	1,888	4,519	1,888

Standard Errors Clustered at City level in Parentheses
* p <0.10 ** p <0.05 *** p <0.01

Table A.4: Robustness Check – No Hypothetical Questions

This table assesses the effects of credit extension on income expectation, spending, and borrowing. The specification is the same as in tables 2, 5, and 7. The sample focuses on those who are not asked the hypothetical question (Section 4 of the Survey).

	$\Delta E_C[Y]$ Acceptance (1)	$\Delta E_C[Y]$ Non-Acceptance (2)	ΔB Acceptance (3)	ΔB Non-Acceptance (4)	ΔB Acceptance (5)	ΔB Non-Acceptance (6)	ΔC Acceptance (7)	ΔC Non-Acceptance (8)	ΔC Acceptance (9)	ΔC Non-Acceptance (10)
ΔLimit	0.451***	0.216**	0.183***	0.047*	0.134***	0.023	0.318***	0.114	0.228**	0.031
	(0.079)	(0.096)	(0.016)	(0.028)	(0.019)	(0.024)	(0.072)	(0.108)	(0.090)	(0.110)
$\Delta E_C[Y]$					0.105***	0.113**			0.186	0.254
					(0.028)	(0.050)			(0.127)	(0.166)
Province FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
First-Stage F	708.67	292.55	708.67	292.55	18.27	9.66	890.28	708.67	18.27	9.66
Overid. p-value					0.57	0.74			0.12	0.42
N	981	402	981	402	981	402	981	402	981	402

Standard Errors Clustered at City level in Parentheses
* p <0.10 ** p <0.05 *** p <0.01

Table A.5: Robustness Check – No Income Information

This table assesses the effects of credit extension on income expectation, spending, and borrowing. The specification is the same as in tables 2, 5, and 7. The sample focuses on those whose income information is Non-observed by the bank.

	$\Delta E_C[Y]$ Acceptance (1)	$\Delta E_C[Y]$ Non-Acceptance (2)	ΔB Acceptance (3)	ΔB Non-Acceptance (4)	ΔB Acceptance (5)	ΔB Non-Acceptance (6)	ΔC Acceptance (7)	ΔC Non-Acceptance (8)	ΔC Acceptance (9)	ΔC Non-Acceptance (10)
ΔLimit	0.482***	0.256***	0.193***	0.033*	0.132***	-0.002	0.302***	0.083	0.262***	0.000
	(0.068)	(0.057)	(0.014)	(0.019)	(0.014)	(0.016)	(0.041)	(0.077)	(0.054)	(0.093)
$\Delta E_C[Y]$					0.129***	0.136***			0.181**	0.333**
					(0.022)	(0.34)			(0.083)	(0.155)
Province FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
First-Stage F	890.28	421.25	890.28	421.25	32.69	24.17	890.28	421.25	32.69	24.17
Overid. p-value					0.30	0.24			0.30	0.24
N	1,941	792	1,941	792	1,941	792	1,941	792	1,941	792

Standard Errors Clustered at City level in Parentheses
* p <0.10 ** p <0.05 *** p <0.01

Table A.6: Credit Extension and Income Expectations – Fuzzy RD

This table assesses the effects of credit expansion on consumer expectations. All variables are winsorized at 1% level.

	$\Delta E_C[Y]$ All (1)	$\Delta E_C[Y]$ All (2)	$\Delta E_C[Y]$ Acceptance (3)	$\Delta E_C[Y]$ Acceptance (4)	$\Delta E_C[Y]$ Non-Acceptance (5)	$\Delta E_C[Y]$ Non-Acceptance (6)
ΔLimit	0.375***	0.346***	0.443***	0.427***	0.210**	0.203*
	(0.093)	(0.096)	(0.105)	(0.106)	(0.112)	(0.109)
Polynomials	2	3	2	3	2	3
Province FE	Y	Y	Y	Y	Y	Y
First-Stage F	23.39	19.87	25.63	17.79	19.87	14.36
N	3,355	3,355	2,270	2,270	1,085	1,085

Standard Errors Clustered at City level in Parentheses
* p <0.10 ** p <0.05 *** p <0.01

Table A.7: Income-Inference of Credit Extension – 12-Month Expectation

This table assesses the effects of credit expansion on the expectation of income growth over 12 months. The specification is based on (1). $\Delta E_C[Y]$ is the difference between the answer of Q1 (d) of the post-experiment survey and that of the pre-experiment survey. A. All variables are winsorized at 1% level.

	$\Delta E_C[Y]$ All (1)	$\Delta E_C[Y]$ All (2)	$\Delta E_C[Y]$ Acceptance (3)	$\Delta E_C[Y]$ Acceptance (4)	$\Delta E_C[Y]$ Non-Acceptance (5)	$\Delta E_C[Y]$ Non-Acceptance (6)
ΔLimit	0.425***	0.414***	0.511***	0.499***	0.257***	0.245***
	(0.030)	(0.060)	(0.043)	(0.048)	(0.041)	(0.048)
Province FE	N	Y	N	Y	N	Y
First-Stage F	2403.73	1293.45	1040.27	2403.73	1293.45	1040.27
N	5,363	5,363	3,786	3,786	1,577	1,577

Standard Errors Clustered at City level in Parentheses
* p <0.10 ** p <0.05 *** p <0.01

Table A.8: Spending Responses after Credit Extension Shocks – Heterogeneity

This table assesses the effects of credit extension on non-debt financed spending by consumer characteristics before the experiment focusing on the participants in the acceptance group. The specification is based on (1) and (2). ΔC is the difference between the total spending minus newly accumulated interest-incurring debt six months after the experiment and that at the beginning of the experiment. $\Delta E_C[Y]$ is the differences between the answer of Q1b of the post-experiment survey and that of the pre-experiment survey. $\Delta Limit$ is the proposed changes in credit limit consumers see on the offers. It is positive for those in the treatment group, and zero for those in the control group. The four panels splits the consumers into high and low groups based on their pre-experiment subjective beliefs of income growth uncertainty, within-industry standard deviation of income growth, wealth-to-income, and utilization rate, respectively. In the treatment group, Acceptance group contains those who accept the offers. In the control group, the classification of acceptance is explained in section III. A. All variables are winsorized at 1% level.

	ΔC Low (1)	ΔC High (2)	ΔC Low (3)	ΔC High (4)
	Panel A: Industry-Level Volatility			
ΔLimit	0.233***	0.511***	0.174***	0.342***
	(0.043)	(0.030)	(0.048)	(0.055)
$\Delta E_C[Y]$			0.311***	0.297***
			(0.081)	(0.053)
First-Stage F	1618.68	504.67	35.01	47.41
Overid. p-value			0.25	0.11
N	2,117	1,669	2,117	1,669
	Panel B:Subjective Volatility			
ΔLimit	0.235***	0.456***	0.180***	0.306***
	(0.036)	(0.049)	(0.042)	(0.061)
$\Delta E_C[Y]$			0.321***	0.270***
			(0.049)	(0.089)
First-Stage F	2021.18	741.79	36.60	45.05
Overid. p-value			0.54	0.11
N	1,893	1,893	1,893	1,893
	Panel C: Wealth/Income			
ΔLimit	0.403***	0.269***	0.357***	0.088*
	(0.048)	(0.041)	(0.046)	(0.049)
$\Delta E_C[Y]$			0.328***	0.390***
			(0.052)	(0.087)
First-Stage F	1056.36	980.63	55.16	45.84
Overid. p-value			0.83	0.07
N	1,893	1,893	1,893	1,893
	Panel D: Utilization Rate			
ΔLimit	0.398***	0.271***	0.347***	0.079*
	(0.049)	(0.042)	(0.046)	(0.049)
$\Delta E_C[Y]$			0.332***	0.379***
			(0.067)	(0.095)
First-Stage F	1233.89	897.11	50.50	41.32
Overid. p-value			0.49	0.16
N	1,893	1,893	1,893	1,893

Standard Errors Clustered at City level in Parentheses
* p <0.10 ** p <0.05 *** p <0.01

B. US Survey from SurveyMonkey

B.1 Survey 1

1. What was you total income in 2020?
2. What was you total income in 2019?
3. What's your age?
4. When a bank actively offers to increase the credit card limit of someone, how much do you think the income of this person would change in the next year?

   - decrease by more than 30%.
   - decrease by 20% to 30%.
   - decrease by 15% to 20%.
   - decrease by 10% to 15%.
   - decrease by 5% to 10%.
   - decrease by less than 5%.
   - about the same.
   - increase by less than 5%.
   - increase by 5% to 10%.
   - increase by 10% to 15%.
   - increase by 15% to 20%.
   - increase by 20% to 30%.
   - increase by more than 30%.

5. When a bank actively offers to increase the credit card limit of someone, how much do you think the saving of this person would change in the next year?

   - decrease by more than 30%.
   - decrease by 20% to 30%.
   - decrease by 15% to 20%.
   - decrease by 10% to 15%.
   - decrease by 5% to 10%.
   - decrease by less than 5%.
   - about the same.
   - increase by less than 5%.
   - increase by 5% to 10%.
   - increase by 10% to 15%.
   - increase by 15% to 20%.
   - increase by 20% to 30%.
   - increase by more than 30%.

6. When a bank actively offers to increase the credit card limit of someone, how much do you think the spending of this person would change in the next year?

   - decrease by more than 30%.
   - decrease by 20% to 30%.
   - decrease by 15% to 20%.
   - decrease by 10% to 15%.
   - decrease by 5% to 10%.

- decrease by less than 5%.
- about the same.
- increase by less than 5%.
- increase by 5% to 10%.
- increase by 10% to 15%.
- increase by 15% to 20%.
- increase by 20% to 30%.
- increase by more than 30%.

7. When a bank actively offers to increase the credit card limit of someone, how much do you think the default rate of this person would change in the next year?[36]

- decrease by more than 30%.
- decrease by 20% to 30%.
- decrease by 15% to 20%.
- decrease by 10% to 15%.
- decrease by 5% to 10%.
- decrease by less than 5%.
- about the same.
- increase by less than 5%.
- increase by 5% to 10%.
- increase by 10% to 15%.
- increase by 15% to 20%.
- increase by 20% to 30%.
- increase by more than 30%.

[36]The order of the choices in questions 4 to 7 is flipped randomly, and the order of questions 4 to 7 is randomized.

B.2 Survey 2

1. What was you total income in 2020?
2. What was you total income in 2019?
3. What's your age?
4. Over the following categories, select the ones that the banks have a better predictive ability about than you (select all that apply)[37]

   - The near-future growth rate of your income.
   - The near-future growth rate of the macroeconomy.
   - The near-future growth rate of the local economy where you live.
   - The near-future growth rate of the sector/industry you are working in.
   - The near-future income growth of someone who is similar to you.
   - None of the above.

5. Suppose your bank increases your credit card limit by 5,000 dollars this month. This would mean that the bank expects your total income to be changed by _____ in the next 12 months.
6. Suppose your bank increases your credit card limit by 10,000 dollars this month. This would mean that the bank expects your total income to be changed by _____ in the next 12 months.
7. Suppose your bank increases your credit card limit by 5,000 dollars this month, how much are you going to increase your total spending in the next year?
8. Suppose your bank increases your credit card limit by 10,000 dollars this month, how much are you going to increase your total spending in the next year?

[37] The order of the choices except for the last option is randomized.

Figure B.1: Income Distribution by Age

This figure plots the income profile by age. The round dots plot the average income from the participants in each of the five-year bins from age 20 to age 65. The diamond dots plot the median income of consumers in the four age groups defined in the 2020 US Census.

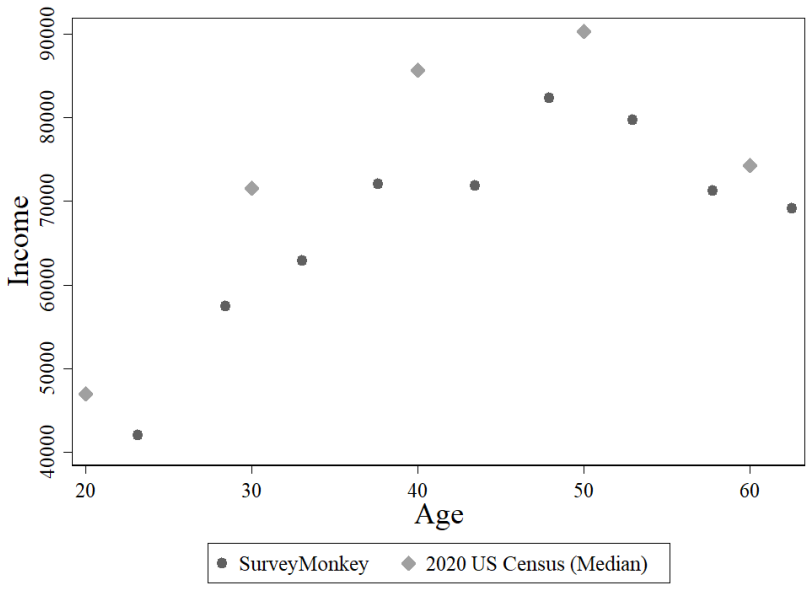

Figure B.2: Reported Effects of Credit Extension

This figure plots the histogram of the answers from US Survey 1 questions 4 to 7. Panel A gives the result of changes in total spending. Panel B shows the result of changes in default probability. Panel C gives the result of changes in total saving. Panel D gives the result of changes in total income. Standard errors are in parentheses.

Panel A: Changes in Spending

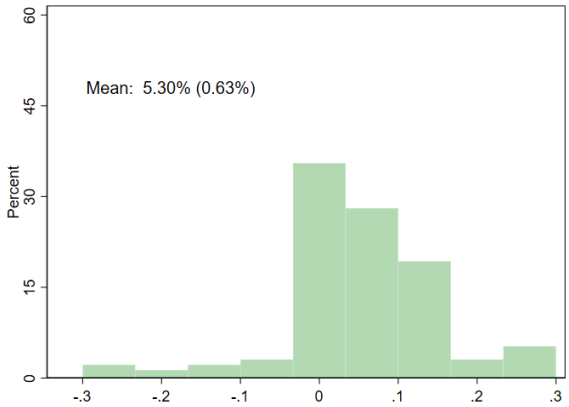

Panel B: Changes in Default Rate

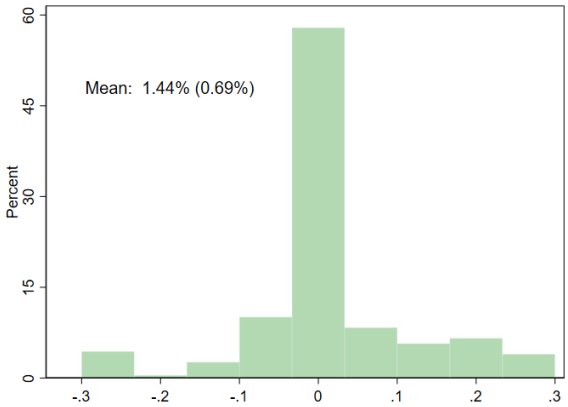

Panel C: Changes in Saving

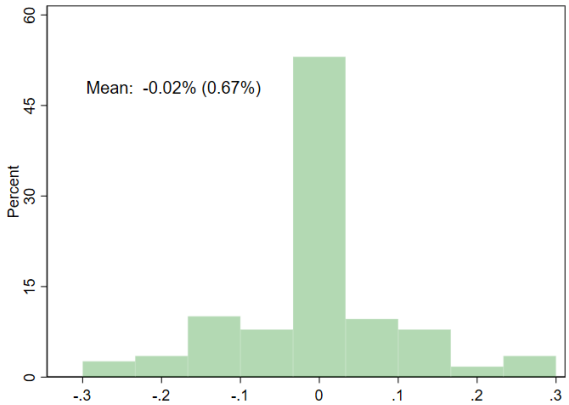

Panel D: Changes in Income

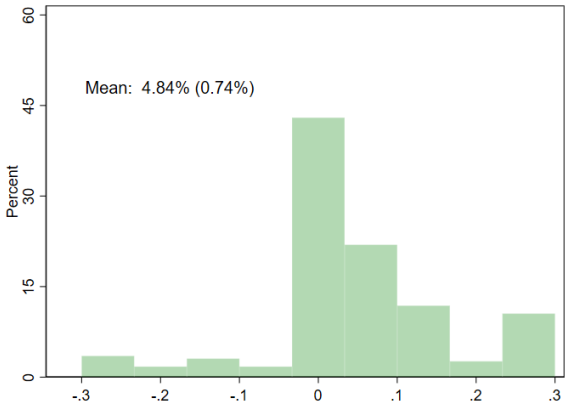

Figure B.3: Distribution of λ

This figure plots the distribution of the subjective sensitivity of bank-perceived income growth to credit supply from the US data. The calculation is based on (3) in the main text. The survey questions are Q5 and Q6 in section B.2.

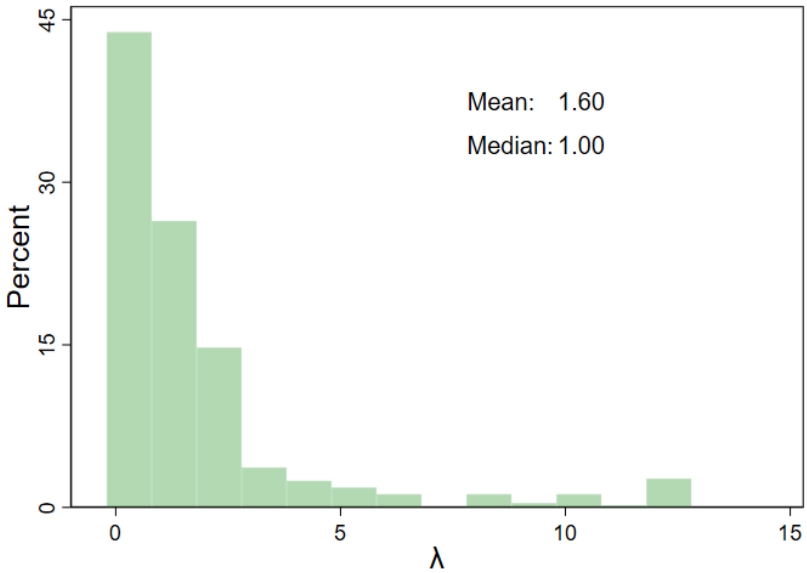

Figure B.4: Can Banks Predict the Future Better?

This figure plots the distribution of the choices from US Survey 2 Q4. The six choices from Q4 are labeled as *Income*, *Macro*, *Local*, *Industry*, *Peer*, and *None*.

C. Solving the Model and Estimation

C.1 Model Solution

The algorithm for solving the model is similar to that in Guvenen (2007). The procedure of setting up the grid is as follows:

1. Simulate $J = 1000$ income paths $\{y^j_{i,t}, t = 1, ..., T, j = 1, ..., J\}$ based on (6).
2. Given σ^2_ξ, get bank beliefs $y^j_{B,i,t}$.
3. For each j, get consumer belief paths based on (6) and (11).
4. Discretize the states into $n_w = 500$ grid points over $-\bar{b}$ and $3 \times \bar{y}_{i,t}$. The maximum value is set to match the 95th percentile in the data.
5. Discretize $\log \hat{y}_{i,t+1}$ and $\log \hat{y}_{Bi,t+1}$ into five values using the Tauchen methods.
6. Discretize $\tilde{w}_{i,t}$ into 20 values using the Tauchen methods.

After setting up the grids, the procedure for solving the model is

1. Consumer problem:
   (a) Set up a grid of $n_b = 30$ points over $\bar{b}_{i,t}$ between $\bar{b}^{min}_{i,t}$ and $\bar{b}^{max}_{i,t}$, where $\bar{y}_{i,t}$ is the average income at age 38.
   (b) Given state variables $\theta_{i,t}$ and bank supply of credit $\bar{b}_{i,t}$, solve consumer Bellman equations (12) to (14) using backward induction.
   (c) Get policy functions $c^*(\theta_{i,t}, \bar{b}_{i,t})$, $b^*(\theta_{i,t}, \bar{b}_{i,t})$, and $d^*(\theta_{i,t}, \bar{b}_{i,t})$.
2. Bank problem:
   (a) Given each combination of state variables $\tilde{\theta}_{i,t}$ and $\bar{b}_{i,t}$, estimate policy functions.
   (b) Based on the estimated policy functions, calculate the optimal $\bar{b}^*_{i,t}$ based on (3) by linearly interpolating the policy function at each t and each $w_{i,t}$.
   (c) Linearly interpolate to get the optimal policy functions $c^*(\theta_{i,t}, \bar{b}^*_{i,t})$, $b^*(\theta_{i,t}, \bar{b}^*_{i,t})$, and $d^*(\theta_{i,t}, \bar{b}^*_{i,t})$.

C.2 Estimation

The estimation consists of two stages. In the first stage, I rely on the experiment to pin down the consumers' discount rates and the parameters associated with the learning process. In addition, I estimate parameters associated with bank profits from spending transactions based on the bank's realized income from the customers. In the second stage, I use the simulated method of moments (SMM) to get the estimates of the consumers' coefficient of risk aversion γ, rate of garnishment χ, ϕ_1, and ϕ_2 in the bank problem.

First Stage: The first-stage estimation requires the experiments and the survey data. A difference here is that the duration of the experimental period is six months, whereas the frequency in the model

is nine months to be consistent with the frequency of credit-limit changes. Given that the spending and borrowing responses nearly leveled off six months after the experiments (see Figure 4), when estimating the parameters, I assume the spending responses and the expectation changes are the same if the experimental period are over one year. In addition, Table A.5 in the Online Appendix shows the effects of the experiment on the 12-month expectations of consumers' future income (survey question Q1d) is similar to that of the six-month expectations. Therefore, the estimates based on the six-month responses should be close to that if the experimental period are over one year.

I first estimate discount rates δ as the average values based on the following survey question:

Q4 f: Rather than receiving 100 CNY today, which of the following options would you choose. (select all that apply).

- X CNY in six months.

$X \in \{95, 97.5, 100, 102.5, 105, 107.5, 110, 112.5, 115\}$. For each participant j, $\delta_j = (100/x)^2$, where x is the smallest choice. I then take the sample average of δ_j as the estimated δ.

For the parameters governing the income process, α, σ_ϵ^2, and σ_ν^2, I residualize all individual income by age, year, education, industry, city, and gender, and estimate (7) using maximum likelihood estimation.

To get the error variances of the bank belief σ_ξ^2, I use the following survey questions:

Q1 b: Your expected total income over the next 6 months is _____.

Q1 c: With a probability of 80%, your total income over the next 6 months will be between _____ and _____.

Questions Q1 b/c from the pre-experiment and post-experiment surveys give $\mathbb{E}[\log y_{0,i,t+1}]$, $\mathbb{E}[\log \hat{y}_{i,t+1}]$, $\hat{\sigma}^2_{0,y,t+1}$, and $\hat{\sigma}^2_{y,t+1}$. Then σ_ξ^2 can be retrieved based on (6b). For the misperception parameter η, the employed estimation strategy is similar to that in section IV.F. I first use survey questions Q4a and Q4b to get a partial derivative of the consumer's subjective beliefs with respect to bank credit supply, λ_i. I then use the bank's credit model and data before 2020 to predict $\log y_{B,i,t+1}$, which gives me the partial derivative of the bank's belief with respect to the bank's supply: $\frac{\partial \log y_{B,i,t+1}}{\partial \log l_{i,t}} = \tilde{g}'$. Taking the ratio of the two numbers gives η_i, and I use the sample average of η_i as η.

κ is directly estimated based on the bank contract. I estimate q to match the proportion of transactions using credit cards given credit limits. For errors in the bank's perception of consumers' total savings, σ_ω, I take the standard deviation of consumer total saving at the bank over the answers from question 3 (e).

Second Stage: I use SMM to estimate γ, χ, σ_ω^2, ϕ_0, and ϕ_1. The five matched moments are the average wealth-consumption ratio, average default rate, average credit limit-income ratio, the sensitivity of credit supply to shocks to the bank's beliefs, and the standard deviation of the income-limit ratio. The logic is explained ahead. The risk-aversion parameter γ captures the curvature of the utility function.

Higher risk aversion increases consumer willingness to save, thereby decreasing the wealth-to-consumption ratio. The marginal garnishment cost χ directly affects consumers' willingness to default. A higher χ indicates a higher cost of default, and therefore a lower default rate. ϕ_0 and ϕ_1 together determine the costs of supplying the credit limit. A larger ϕ_0 or ϕ_1 increases the marginal costs of supplying an additional unit of credit limit. In addition, ϕ_1 captures the convexity of the costs of credit supply, and a larger ϕ_1 increases the convexity of the costs. In response to a shock to bank belief, a larger ϕ_1 gives a smaller response to the change in credit supply.

The SMM procedure searches for the set of parameters that minimize the weighted deviation between the actual and simulated moments,

$$(m - \hat{m}(\Theta))'\widehat{W}(m - \hat{m}(\Theta)),$$

where $\hat{W}$ is the variance-covariance matrix of the data moments. The calculation of the empirical moments is straightforward and is based on the main sample of analysis. The weight matrix $\widehat{W}$ adjusts for the possibility that some moments are more precisely estimated than others. I calculate $\widehat{W}$ as the inverse of the variance-covariance matrix of the empirical moments based on 1,000 bootstrap draws with replacements. For simulated moments, I focus on the average consumer in the data and generate the moments based on the model at time $t = 0$, which is to match age 38 in the data. I fist split each variable into four groups based on the wealth-to-income ratio. The three thresholds are based on the quartiles in the data. I then take the median value in each quartile and take the average to get the simulated moments.

D. Bank Credit-Supply Rule

This section presents an overview of the bank's credit-supply rule. The bank's aggregate amount of credit to be extended to the household sector is guided by the People's Bank of China. With this soft aggregate credit limit, the bank decides the amount of credit to extend through credit cards. Credit supply at the extensive and intensive margins is guided by two different proprietary credit-supply algorithms. Because I study only the intensive margin of credit supply in this paper, I do not explain how the bank decides who receives credit cards at the extensive margin.

The bank's proprietary rule for extending a credit card limit is a collection of machine-learning techniques. Broadly speaking, the model is a supervised learning process with two objectives: maximizing utilization and minimizing the default rate, subject to the aggregate credit limits in the next year. The relative weights of the two objectives are determined based on the relative concern about profitability and financial stability during different periods. The model input is largely based on two aspects. The first aspect is consumption behavior. Consumption behavior is based on consumers' demographics, past spending, and repayment patterns. The key variables in this category include credit score from the Credit Reference Center of the People's Bank of China and the consumers' age, highest earned degree, occupation, total savings, default history, social network size, and behavioral preferences (types and amounts of goods and services consumed). In a consumption model, this aspect can be thought of as inferring the parameters of consumers' preferences, for example, their discount rates, risk aversions, and the distribution of taste shocks.

The second aspect relates to consumers' income perspectives and is estimated based on three categories. The first is an idiosyncratic category based on consumers' past inflows interacted with their consumption behaviors, as calculated from the first aspect. Consumers' past inflows consist of three variables: bank-calculated income, differences between bank-account inflows and outflows, and self-reported income. The second category is an occupation-specific component that is based on the aggregate wage growth and uncertainties of consumers' occupations. The third category is location specific and is based on the collection and analysis of the data about local economic activities. In a consumption model, the second aspect can be thought of as inferring the consumers' income distributions.

The ultimate amount of credit extension is based on a nonlinear combination of the two aspects, and within each aspect, the statistical model is highly nonlinear. Additionally, each subcategory within the two aspects can interact to generate a more precise estimation. Supposedly, in terms of consumption behaviors, credit extensions are larger for those whose consumption patterns imply lower default preferences or higher social awareness. In terms of income, credit extensions are larger for those with higher income perspectives in the near future.