Learning in the Limit: Income Inference from Credit Extensions

Xiao Yin*

Jan, 2024

Abstract

Increases in credit limits raise consumption significantly, even for consumers who are not borrowing-constrained. Combining a randomized controlled trial with administrative and survey data, I show that credit limit extensions significantly increase consumers' expectations about future income. The increase in income expectations is due to beliefs about higher future labor productivity instead of planned increases in labor supply. Controlling for changes in consumer expectations regarding future income, the response of consumption to credit limit extensions is weakened by 37%.

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

*Yin: University College London, xiao.yin@ucl.ac.uk. I deeply appreciate the valuable comments from Ulrike Malmendier, Yuriy Gorodnichenko, David Sraer, and Michael Weber. I also appreciate the valuable comments from Deniz Aydin, Matteo Benetton, Markus Brunnermeier, Stefano DellaVigna, Rawley Heimer (discussant), Amir Kermani, Chen Lian, Peter Maxted, Maarten Meeuwis (discussant), Emi Nakamura, and Annette Vissing-Jørgensen. I am thankful for the valuable comments from various seminars and conferences.

I Introduction

Credit limit plays a crucial role in household consumption savings decisions, underpinning the extent to which consumers can borrow to smooth consumption. Empirical evidence shows a substantial average spending response to credit limit changes. Meanwhile, even for consumers who are far from being borrowing-constrained, extensions of the credit limit still induce non-trivial amounts of increases in total consumption[1]. This is contrary to predictions by the standard buffer-stock model, which posits that credit limit variations should not significantly impact overall spending for consumers who are not liquidity-constrained. This raises questions about the mechanisms at the micro-level through which credit limit extensions influence consumer spending.

Standard estimation of the spending responses to borrowing limit extensions relies on arguably random variations in credit limits. An implicit assumption in these settings is that consumers in the field also treat credit supply events as random. However, banks' credit extension decisions are rarely random and are usually a function of the economic conditions and consumer characteristics. An intriguing question is how consumers think about banks' credit supply decisions. That is, do consumers always treat credit supply, in the form of extended credit limits, as random shocks only to their borrowing constraints, or do they believe credit supply is an endogenous outcome that contains information the consumers are not fully informed about? Motivated by this question, this paper studies the effects of credit extensions on consumption by affecting expectations.

Studying how credit supply affects consumer beliefs is challenging, as one needs to identify belief changes around field credit supply events. To cope with this difficulty, this study involves a collaboration with a large commercial bank in China, focusing on how consumers modify their expectations in response to credit expansions. The methodology combines a randomized controlled trial (RCT) with administrative and survey data. In this setup, the bank had initially planned to increase the credit card limits of 17,000 customers, following its usual internal underwriting process. However, for experimental purposes, the limit increases were delayed by 12 months for a randomly

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

selected control group. The remaining customers (the treated group) received the planned credit limit increase. Given that the increases in credit supply are based on the bank's usual underwriting process, the setting gives a nice opportunity to identify the effects of limit extensions in the field.

To isolate the possible belief channel in credit supply, I use random information treatment that varies the degree of inferencing from limit extensions. The basic idea is that, to the extreme, if the consumers know that the credit supply decision is purely random, then they should not infer anything from it. To accomplish this, I separate those in the treatment group into two subgroups: T1 and T2. While for both T1 and T2, participants received a notice about the increase in their credit limit (Figure 1), as the bank customers would normally receive for such events, for T2, participants were also shown the following fact about the design of the study at receiving the limit-increase notice:

> *The increase in credit limit is based on a limit-increase event. In this event, among a portion of customers with a good credit score, we randomly selected a group of users, including yourself, and increased their credit limits.*

The information treatment informs that the limit increase is sent to a randomly selected group of customers, conditional on having a good credit score. It seeks to weaken, if any, how much consumers infer information from credit supply decisions. If there is no belief channel in credit supply, consumption responses should be statistically indifferent between T1 and T2.

I begin the analysis by studying responses of unsecured debt and spending to limit extensions. I find a large consumption response to limit extension. Specifically, for each CNY higher credit limit, consumers in T1 increase spending by 0.37 CNY and accumulate 0.15 CNY more unsecured debt over six months. These numbers are respectively close to the estimated marginal propensity to consume out of limit change (MPCL) and marginal propensity to borrow out of limit change (MPB) from the previous literature[2]. Comparing the consumption responses between T1 and T2 sheds light on the existence of a belief channel in credit supply. In particular, I find that the consumption responses are about

[2]For example, estimated MPCL is between 0.2 and 0.6 in Agarwal et al. (2017) over 12 months; estimated 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 0.16 over nine months in Aydin (2022).

37% smaller for T2 as compared with T1. Therefore, informing about the randomness in the credit expansion decision has large negative effects on total consumption, indicating a non-trivial weight of a belief channel in the consumption responses to credit supply.

To study the effects of limit increase on beliefs, I sent a survey to about 70% of the participants in all groups within ten days after the experiments. The survey aimed to elicit the beliefs of the participants' future perspectives. It mainly asked about expectations about different components of consumer budget constraints (e.g., consumption, saving, income, delinquency probability, etc.). The results on expectations can be summarized as

1. For T1:
   (a) limit extensions significantly increase expectations of future consumption and income conditional on having the same job.
   (b) there is no significant effect on expectations of future default rate and a marginally positive effect on expected future savings.
   (c) there is no significant effect on the number of hours planned to work, but significant negative effects on forecasted unemployment possibility.
   (d) there is no significant effect on the highest attainable credit limit in the short-run or long-run.
2. For T2, there is no significant effect on any elements.

The findings of this study are notable in several respects. Firstly, results 1.a and 1.b suggest that after receiving an increased credit limit, consumers anticipate higher future consumption, which they believe will be financed by increased income rather than by drawing down savings. This challenges the buffer stock model, where an extension in credit limit is expected to boost total consumption by reducing savings, based on the mechanism that a higher credit limit enables smoother consumption and lessens the need for precautionary savings.

However, results 1.a and 1.b do not explain why consumers expect their income to rise following a credit limit increase. Possible explanations might include increased labor supply due to factors like enhanced entrepreneurship (Herkenhoff et al., 2021), better labor mobility (Doornik et al., 2021), or reduced financial distress (Sergeyev et al., 2023). Result 1.c, however, dismisses these supply-side explanations by indicating that consumers do not expect to work more hours after a credit limit increase. This rules out the idea that a relaxed financial constraint or reduced financial distress leads to an increased labor supply. Moreover, the expectation of higher future income, conditional on retaining the same job, negates the hypothesis of improved labor mobility. Instead, the anticipation of lower unemployment risk and heightened future earnings suggests that consumers may interpret credit expansions as signals of higher labor demand or increased marginal productivity. The results, therefore, posit an *income-inference* channel through which credit limit extensions affect consumption.

I rationalize the results with a simple model, focusing on how consumers infer information from credit supply. In this model, consumers with variable income and potential borrowing constraints make consumption decisions. Their income is influenced by both idiosyncratic and systematic factors, with the latter varying among consumer types. A key assumption is that consumers are imperfectly informed about the impact of the systematic components on their income, possibly due to inattention to economic conditions (Mankiw and Reis, 2002; Reis, 2006; Coibion and Gorodnichenko, 2012) or lack of understanding of economic dynamics (Hansen, 2007; Collin-Dufresne et al., 2016). Banks, while not informed of the idiosyncratic income shocks, have partial insights into systematic shocks. In this framework, credit supply decisions convey the bankâ€™s beliefs about the economy and, thus, future income trends, which consumers then factor into their expectations.

The model suggests that credit limit supply influences consumption through two channels: one is the precautionary channel, which eases the inability to smooth consumption, and the other is the income-inference channel, which alters beliefs about future income. It predicts that credit limit increases impact consumption even for individuals unlikely to reach their borrowing limits, a phenomenon that is inconsistent with the buffer stock model but documented in prior empirical studies (Agarwal et al., 2017; Aydin, 2022; D'Acunto et al., 2020). For people with no income volatility, the model implies that the weight of the income inference channel is zero. Consequently, for such individuals, changes in income expectations should be negligible, and consumption responses should be close to zero if they also have substantial liquidity. This prediction aligns with the empirical finding that belief changes are insignificant for those with no income variation before the experiment, and the marginal propensity to consume (MPCL) is significantly positive for those with high liquid savings relative to annual income, yet becomes insignificant further conditional on low pre-experiment income volatility.

I further investigate the type of information consumers might be inferring from credit limit increases. One possibility is that banks can assess macroeconomic shocks, such as productivity variations across business cycles. In this scenario, consumers believe that banks are capable of predicting macroeconomic outcomes, either due to being more attentive to current economic states (Mankiw and Reis, 2002; Reis, 2006; Coibion and Gorodnichenko, 2012) or possessing a different degree of understanding of the structural parameters that govern macroeconomic trends (Hansen, 2007; Collin-Dufresne et al., 2016). Alternatively, banks could identify specific households likely to experience an income increase. If consumers are inferring information about macroeconomic conditions, those with a more significant change in income expectations should be less confident in assessing current economic performance, more likely to update their beliefs about the economy following credit expansions, and inclined to believe that macroeconomic shocks significantly impact their income. I find evidence supporting this channel. Notably, income expectations are minimal among individuals who are 1. highly confident in their evaluation of the macroeconomy, 2. do not adjust their beliefs about macroeconomic performance after the limit shock, and 3. do not perceive macroeconomic fluctuations as having a significant impact on their income. In sum, the findings suggest a mechanism where lending standards vary across the business cycle (Bassett et al., 2014; Fishman et al., 2020; Weitzner and Howes, 2023). Consumers, imperfectly informed about macroeconomic states, deduce information about productivity from credit supply decisions, leading to altered income expectations and subsequent changes in consumption.

The extent to which consumers infer income from credit supply depends on the banks' decision-making processes. More sophisticated banks may have a deeper understanding of overall economic conditions, while less sophisticated ones might rely primarily on payment history. An interesting question is whether Chinese banks differ in their lending practices compared to others. To explore the applicability of the income-inference channel in different settings, the study concludes with survey results from the United States, collected through online platforms. These surveys assessed responses to randomized hypothetical scenarios of credit limit increases. The findings indicate that, in hypothetical situations, larger credit limit increases are associated with expectations of higher future income and consumption without changes in savings, default rates, or labor supply. However, when these increases are perceived as purely random, they do not significantly impact expectations. These outcomes align with the main study's findings in the Chinese context, suggesting that US participants also adjust their income expectations following higher credit limit increases.

Related Literature This paper mainly contributes to two 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). Recent major progress is Aydin (2022), which provides a clean empirical estimation of the marginal propensity to borrow using an RCT in Turkey. Although previous literature mainly relies on the buffer-stock model to explain how credit limits affect consumption, the effect of credit expansions on consumer spending through changing beliefs is still an open question. The lack of evidence lies in the difficulties of combining an RCT with both observational and expectation data. The paper combines field credit supply events with survey data to provide a complete picture of how consumers change their beliefs about credit limit extension. The findings facilitate direct testing of the effects of credit supply on consumers' beliefs. It also helps provide new insights into macroeconomic models incorporating credit supply shocks.

In addition, this paper contributes to a growing literature that focuses on the role of beliefs in explaining consumers' spending-saving decisions[3]. For example, Ameriks et al. (2016), Ameriks et al. (2020), and Ameriks et al. (2020) provide recent advances by linking survey evidence to retirement choices. Manski (2004), Ameriks et al. (2020), and Giglio et al. (2021) study the relationship between investor beliefs and stock investment. Bucks and Pence (2008), Bailey et al. (2019), and Kuchler et al. (2022) analyze how beliefs affect mortgage-leverage choices. A related study is Soman and Cheema (2002), who show that participants' reported MPCL is larger when credit-limit assignments accurately reflect future earning potential. This paper builds on this literature by exploring a quantitative

[3]See DellaVigna (2009) and Benjamin (2019) for a review.

survey matched to administrative and transaction-level data to explore consumer spending and borrowing decisions.

II Conceptual Framework

A. Setup

In this section, I lay out a simple model to illustrate the main channels through which consumers change their spending after a credit constraint shock. The model spans three periods and consists of k types of consumers indexed by i. There are N_k consumers in each type k. Consumer i's utility in t has the form

$$u(C_{i,t}) = C_{i,t} - \frac{b}{2}C_{i,t}^2$$

where $C_{i,t}$ is consumer i's consumption in period t. i is endowed with an initial asset $A_{i,0}$ and receive income $Y_{i,t}$ at the beginning of each period. The budget constraints in the three periods are, respectively

$$A_{i,t} = A_{i,t-1} + Y_{i,t} - C_{i,t}$$

where $A_{i,t}$ is the total savings at the end of t. In this simple model, the discount factor and interest rate are both set to zero. At the beginning of period 3, $Y_{i,3}$ is realized. The game ends afterwards, and consumer i consumes everything and ends the game with zero saving, i.e. $A_{i,3} = 0$. Besides, consumer i faces a borrowing limit L_i such that

$$A_{i,t} \geq -L_i.$$

Income is stochastic and follows

$$\Delta Y_{i,t+1} = \alpha_{i,t} + \rho_k X_{t+1},$$
$$X_{t+1} = \rho_X \widetilde{X}_t.$$

$\alpha_{i,t} \sim N(\alpha_i, \sigma_\alpha^2)$ is the component of income for which the mean is only known to i. X_{t+1} captures the systematic shocks to income (e.g. macroeconomic shocks to growth, productivity, or inflation, etc.). $\widetilde{X}_t = X_t + \epsilon_{X,t}$, and $\epsilon_{X,t}$ is an innovation to the systematic factor. Each type k consumers have a different level of exposure ρ_k to the shock. The income process can be written as

$$\Delta Y_{i,t+1} = \alpha_{i,t} + \rho_{kX}\widetilde{X}_t.$$

The key information friction is that consumers have noisy perceptions about the contribution of current systematic shocks on future income growth. That is, the consumers are imperfectly informed about $\rho_{kX}\widetilde{X}_t$. There are two possibilities. First, consumers are imperfectly informed about $\widetilde{X}_t$ because of inattention to current information (Mankiw and Reis, 2002; Reis, 2006; Coibion and Gorodnichenko, 2012). Alternatively, consumers could be uncertain about ρ_{kX} due to lacking the knowledge about the structural parameters governing the dynamics of the model (Hansen, 2007; Collin-Dufresne et al., 2016). For simplicity, I do not differentiate between these two cases. Instead, let

$$\Delta Y_{i,t+1} = \alpha_{i,t} + \widetilde{\rho}_{i,t}.$$

where $\widetilde{\rho}_{i,t} = \rho_{kX}\widetilde{X}_t$ is the contribution of current economic states on consumers' future income. Consumers have a prior of $\widetilde{\rho}_{i,t}$ that follows $N\left(\widetilde{\rho}_{i,t}^0, \sigma_0^2\right)$.

B. Supply of Credit Limits

There is a monopolistic bank that decides on a level of borrowing limit L_i at the beginning of t_1, before the consumers' optimal decisions. For simplicity, let the bank be fully attentive to $\widetilde{X}_t$. The bank observes $\widetilde{\rho}_{i,t}$ with a different precision. This could be that the bank is less inattentive to the current economic states or that the bank can estimate ρ_{kX} with the regression

$$\Delta Y_{i,t} = \overline{\alpha} + \rho_{kX}\widetilde{X}_{t-1} + \epsilon_{i,t}.$$

In this case, assuming $\epsilon_{i,t} \sim N(0, \sigma_\epsilon^2)$, then the bank's estimate of ρ_{kX}, ρ_{kX}^B, follows $N(\rho_{kX}, \sigma_\rho^2)$, where $\sigma_\rho^2 = \sigma_\epsilon^2/(N_k s_X^2)$ and s_X^2 is the sample variance of X_t. Let $\widetilde{\rho}_{i,t}^B = \rho_{kX}^B\widetilde{X}_t$

be the bank's estimates of the systematic component in i's income. Then, from the bank's perspective, $\widetilde{\rho}^B_{i,t} \sim N(\widetilde{\rho}_{i,t}, \widetilde{\sigma}^2_\rho)$, where $\widetilde{\sigma}^2_\rho = \sigma^2_\rho X^2_t$.

Note that $\rho_{kX} = \rho_k \rho_X$. It is also possible that the consumers and the bank only differ in estimating ρ_X. A reason is that households and professionals have heterogeneous reasonings about how macroeconomic shocks would affect the future economy, which is consistent with recent findings in Andre et al. (2022).

After estimating $\widetilde{\rho}^B_{i,t}$, the bank sets L_i following the rule

$$L_i = f(\widetilde{\rho}^B_{i,t}, \theta_i), \tag{1}$$

where θ_i captures other characteristics that affect the credit supply decisions. I assume that L_i is monotonic in $\widetilde{\rho}^B_{i,t}$.

C. Learning from Credit Limit Changes

After receiving the credit limit L_i, consumer i infers the future income change as perceived by the bank via Bayesian learning. First, the consumers form the subjective beliefs of $\widetilde{\rho}^B_{i,t}$ as

$$E_c[\widetilde{\rho}^B_{i,t}] = f^{-1}(L_i) \equiv g(L_i).$$

With rational learning, consumers correctly infer the functional form of f, and $E_c[\widetilde{\rho}^B_{i,t}] = \widetilde{\rho}^B_{i,t}$. That is to say, rational learning indicates that the bank cannot change L_i to over-signal its beliefs.

With the supplied credit limit L_i, the posterior of consumer i's belief about future income growth has the expected value of

$$\widehat{\rho}_{i,t} = \widetilde{\rho}^0_{i,t} - \kappa_i[g(L_i) - \widetilde{\rho}^0_{i,t}], \tag{2}$$

where $\kappa_i = \sigma^2_0/(\sigma^2_0 + \widetilde{\sigma}^2_\rho)$ is the Kalman gain of the learning process.

Note that Bayesian learning does not require that the bank has better predictability of $\widetilde{\rho}_{i,t}$. As long as the bank's signal precision is not 0, and individuals are not perfectly informed about $\widetilde{\rho}_{i,t}$, credit supply that incorporates the bank's beliefs also changes consumers' beliefs.

In addition, whether g' is positive or negative is ambiguous and depends on the bank's objectives. When $g' > 0$, credit limits increase with the bank's belief about consumer future income. If the effects of negative shocks to income on default rate is low, then g' could be negative, as a lower income increases interest income from taking more debt. For simplicity, the model is loose on how credit limit is affected by bank beliefs about income growth. Instead, under the assumption of rational learning, consumers are correct in guessing the function form of $g(L_i)$. (3) implies the following proposition:

> *Proposition 1: Suppose g' >0, a higher credit limit increases posterior income expectations, and the change is larger when consumers have a higher signal-to-noise ratio, κ_i.*

D. Optimality Condition

Given a three-period setup and quadratic utility, the solution of the model is standard. Consumer i's optimal decision can be determined using backward induction. In period 3, consumer i consumes everything available. Optimal consumption in periods 2 and 1 can be written as

$$C^*_{i,2} = \min\left\{\frac{A_{i,1} + Y_{i,2} + E_{i,2}[Y_{i,3}]}{2}, A_{i,1} + Y_{i,2} + L_i\right\},$$
$$C^*_{i,1} = \min\left\{E_{i,1}[C^*_{i,2}], A_{i,0} + Y_{i,1} + L_i\right\}.$$

Suppose consumer i's t_1 consumption isn't binding; then the consumption rule in period 1 is the classic Hall (1978)'s martingale. When the borrowing limit is binding, consumer i consumes all resources available.

E. Marginal Propensity to Consume out of Liquidity

Borrowing the language from Gross and Souleles (2002), I analyze consumer i's MPCL as the effects of a one-unit increase in L_i on $C^*_{i,1}$. When borrowing is currently binding before and after the credit shock, MPCL equals one. Extensive literature documents that MPCL is large even if the borrowing limit is slack. To analyze the MPCL for financially unconstrained consumers, consider the case when

$$C_{i,1}^{*} = E_{i,1}[C_{i,2}^{*}]. \tag{3}$$

For brevity, I assume $\tilde{\rho}_{i,3}$ is zero. Given that future income is normal. The probability that consumption in the second period does not bind is

$$P_{i,2}(not\ binding) = P\left(\frac{A_{i,1} + Y_{i,2} + Y_{i,3}}{2} < A_{i,1} + Y_{i,2} + L_i\right) = \Phi\left(\frac{2L_i + A_{i,1} - \alpha_i}{\sigma_\alpha}\right) \tag{4}$$

where $\Phi(\cdot)$ is the standard normal CDF. From (5), the probability of a slack borrowing limit is larger if the saving is higher, the credit limit is larger, the income growth from period two to three is smaller, or the income volatility is smaller.

Combining (4) and (5) yields

$$C_{i,1}^{*} = \widetilde{C}_{i,2} - (1-\Phi)(\widetilde{C}_{i,2} - \widetilde{C}_{i,2}) \tag{5}$$

where $\widetilde{C}_{i,2} = (A_{i,1} + E_1[Y_{i,2}] + E_1[Y_{i,3}])/2$ is the optimal level of t_2 consumption when there is no borrowing limit, and $\widetilde{C}_{i,2} = A_{i,1} + E_1[Y_{i,2}] + L_i$ is the highest level of t_2 consumption when borrowing limit binds in t_2.

The MPCL for currently unconstrained consumers is then derived by total differentiating $C_{i,1}^{*}$ with respect to L_i, which yields

$$\frac{\partial C_{i,1}^{*}}{\partial L_i} = \underbrace{\frac{1}{\omega}\left[\left(\tilde{C}_{i,2} - \bar{C}_{i,2}\right)\frac{2\phi}{\sigma_\alpha} + (1-\Phi)\right]}_{precautionary} + \underbrace{\frac{1}{\omega}Kg'(L_i)}_{income-inference}, \tag{6}$$

where $\omega = 1 + 1/2 + (1-\Phi)/2 + (\tilde{C}_{i,2} - \bar{C}_{i,2})\phi/\sigma_\alpha$ is a number that is larger than one. As shown from (6), there are two channels through which credit expansion affects current consumption for unconstrained consumers. The first bracket captures the conventional precautionary channel. Through this channel, an increase in credit limit increases current consumption by both reducing the probability of a binding constraint and increasing the debt capacity in the future. In addition, an income-inference channel is captured by the second term on the right-hand side of (6). Suppose $g' > 0$; the bank would offer more credit limits if the bank perceives a higher income in the future. Then, a one-unit increase in credit limit signals to consumer i that the bank believes that consumer i's income will grow by g' units. The marginal effect of the signal on consumer i's income expectation is g'. (6) suggests the following proposition:

Proposition 2: Suppose credit expansion from the bank signals higher future income growth to the consumers. The unconditional level of MPCL is smaller than that when controlling for the effects of credit expansion on income expectations.

III Methodology

A. Data and Institutional Environment

The data for this study is sourced from a large national commercial bank in China and is one of the country's top ten banks in terms of total assets. In 2023, the bank reported assets exceeding \$1 trillion, serving over 50 million active customers and managing 80 million active credit cards. This extensive customer base ensures that the sample is representative of the diverse demographic distribution of consumers in China.

For daily transactions, most people in China use Alipay or Weixin Pay as payment methods. Using such payment tools requires the users to link their accounts with bank cards or credit cards, similar to using PayPal and Apple Pay in the US. 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 cashback for purchasing certain types of goods or services. At the end of each billing cycle, a minimum repayment is required (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 cashback 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.

Credit card use in China has grown significantly since 2016. From 2016 to 2022, the total outstanding balance from credit cards in China has grown from 3.6 trillion CNY to 8.7 trillion CNY. During the same time, the total credit limits has increased from 9.1 trillion CNY to 22.3 trillion CNY. Credit cards and other personal credit from commercial banks in China remain the most used method for consumption-based unsecured debt. Similar products from FinTech platforms and consumption debt companies, including Alibaba's Huabei, have been gaining market share recently. However, the total market share from these companies is still relatively small, taking around 20% of all consumption-based credit debt in 2023.[4]

B. Measuring Income, Debt, and Spending

I follow the steps the bank use to classify income. Specifically, individual income is classified based on regular inflows. The bank classifies income into two main categories: salary and business cash flows. Salary is defined as the periodical monthly inflows of income, bonuses, and commissions if the consumer declares that they work as an employee. The bank calculates this number in one of two ways. First, if income is paid through direct deposit to this bank, then the number is directly labeled as salary. 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[5]. As for income from business operations, the measure 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, given the requirements of income information for the analysis. I restrict the sample to those whose income information is observed at the bank. This restriction drops the sample by around 35%.

When all the incomes in our sample are aggregated, the split of the two components comes out to be 70.16% from salary and 29.84% from business operations. To verify that 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

[4]See here and here for the sources.

[5]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 like the retirement savings plan in other countries. The monthly contribution is 8% of income. However, the income base is usually capped at the two tails of the income distribution. The numbers differ for different geographic areas. 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. This only causes an around 7% drop of the sample.

administrative government agency[6]. The results of this comparison are shown in Panel B of Figure B.1. in the Online Appendix. The results show a very high relationship between the income from the bank and that from the administrative agency. Fitting a regression between the two yields an R^2 of 0.80.

Debt data is from the Credit Reference Center of the People's Bank of China (the official credit registry), based on the credit reports retrieved by the bank. The Credit Reference Center aggregates personal credit information from all financial institutions. Therefore, studies of debt behavior are expected to capture the overall borrowing outlook of the consumers.

The main analysis is based on interest-incurring debt due to complete coverage. In addition, I calculate total spending as the sum of all purchasing transactions in a given period. There are advantages and drawbacks of this measure. First, since the data is from a single provider, it has issues covering all spending histories of the participants. However, testing the buffer-stock model necessitates spending data. In particular, the buffer-stock model suggests that the consumption response to limit increases for high-liquidity consumers is also positive, but debt needs are, in general, close to zero for these individuals. In this case, testing the mechanisms of how limit extensions affect consumption requires a focus on the spending patterns of high-liquidity consumers.

To eliminate the sample-coverage problems associated with using data from one financial institution, I leverage the findings that many consumers only use one bank for daily transactions[7], and focus on consumers who use this bank as their main source of transactions. I select this sample based on two criteria. First, participants are asked the following survey question

How many banks do you usually use for transaction purposes?

The first filter is that the participants answered *one* to this question. Following Ganong and Noel (2019), the second filter is that there are at least 15 spending transactions on average each month over the 12 months before the surveys. The two filters ensure that

[6]Data from the administrative government agency is only available for 37% of the participants who agreed to share the data with the bank.

[7]Nelson (2022) shows that, depending on their FICO scores, at least 80% to over 90% of the consumers in the US hold only one primary credit card account.

the consumers use only one bank for transaction purposes, and this bank is the one the participants are referring to.

I verify the effectiveness of the selected samples with two tests. First, I also elicit total past consumption with the following question

> *What was the total amount of your spending during the past 12 months (excluding investment and purchases of durable goods including housing, cars, etc.)?*

I then compare this number with the total spending amount based on summing up all purchasing transactions for the selected sample. Panel C of Figure B.1 in the online appendix shows the binned scatter plots of the two. A regression between the two yields an intercept of zero, a slope of 0.98, and an R^2 of 0.58. This indicates a very high correlation between the two measures, especially when the survey measure of consumption is, in general, very noisy.

The second test is to study the changes in cross-bank transfers before and after the experiment. If the participants start to use this bank more after the experiment, I should see a positive change in the net inflow transfer for the treatment groups. Table B.4 in the online appendix shows that the changes in transfer for the treatment groups, focusing on the selected samples, are insignificantly from zero, indicating that changes in consumption are unlikely results of changes in the banks the participants choose to use for transactions.

C. Experimental Design

Figure 1 describes the procedure. There are four steps in the procedure. Specifically,

1. **Sample construction**: From Jun 19 to Jun 23, the bank selected a group of consumers (around 50000 from 57 cities) and decided to increase their credit limits. The amounts of increase were based on the bank's credit-scoring rules. Then 17000 individuals were selected as the subjects in this study. The selected individuals were then grouped into two subsamples. In each subsample, subjects were assigned to either one of a control group, treatment group 1 (T1), and treatment group 2 (T2). The numbers of subjects in each group are shown in the table in Figure 1.

2. **Treatment**: On Jul 03, credit limits were changed to the pre-determined level for participants in the two treatment groups. Treated participants were sent a

Figure 1. Timeline of the Experiments

This figure gives the timeline of the two experiments. Sub-figure A is for the pilot study, which spans from May 2019 to June 2019. Sub-figure B is for the main study, which spans from March 2020 to October 2020.

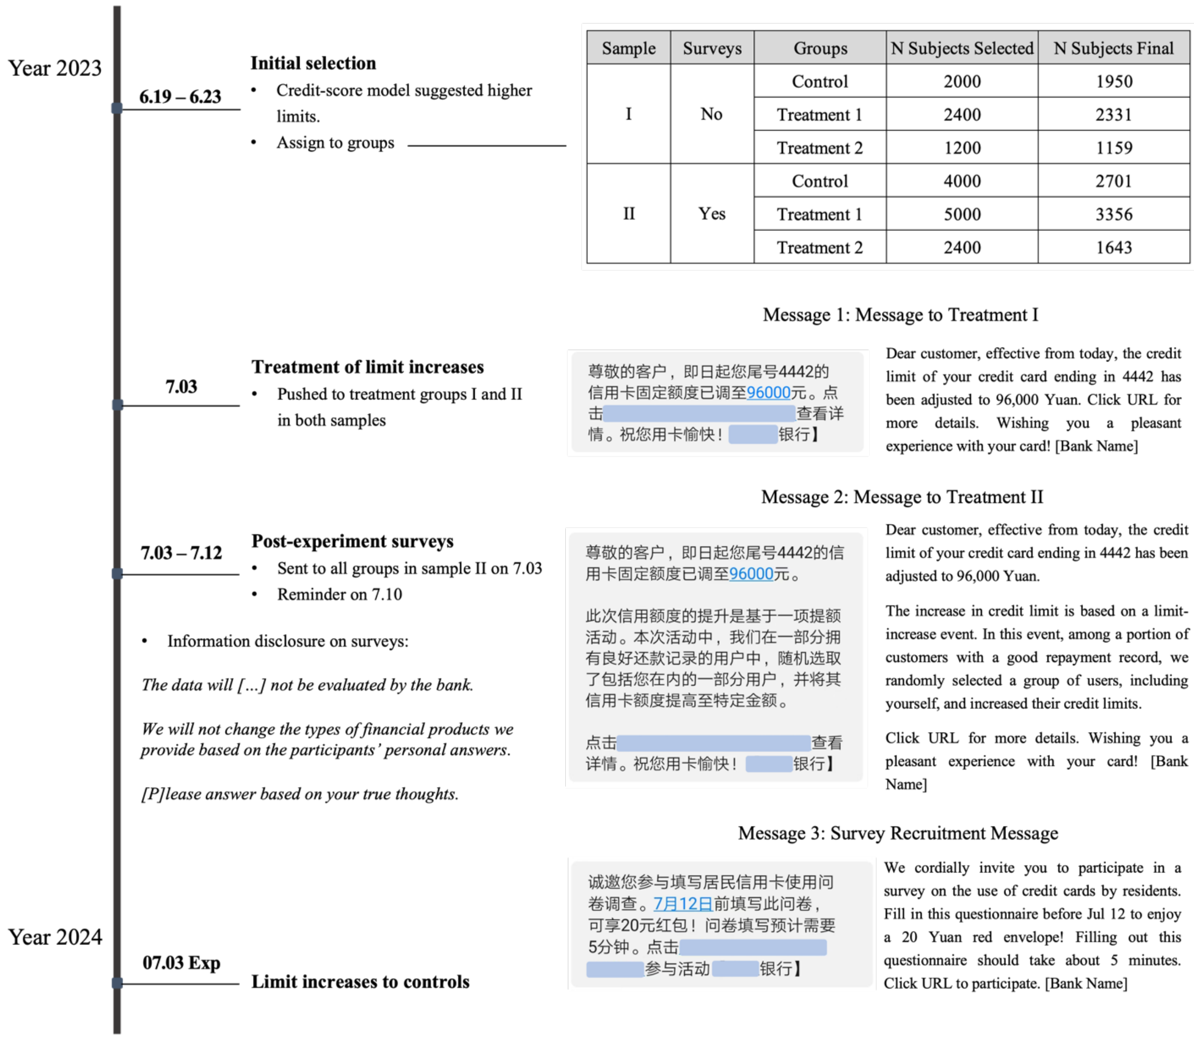

text message about such changes. At the same time, participants in T2 were also informed that the changes were based on a research project. The additional disclosed information was

The increase in credit limit is based on a limit-increase event. In this event, among a portion of customers with a good credit record, we randomly selected some users, including yourself, and increased their credit limits.

3. **Post-experiment survey**: On Jul 03, after the treatment notice, participants in sample II were invited to fill out a survey[8] through text messages. The survey had

[8] Section A in the online appendix shows the survey in English.

to be completed before Jul 12. A reminder text about filling out the surveys was sent on Jul 10.

4. **Limit changes to control**: the new credit limits for the control group as determined in step 1 are expected to be pushed on Jul 03 2024.

Mapped to (6), the treatment effect on T1 estimates the total effects of the credit limit on consumption. The information treatment to T2 seeks to exogenously vary $g'(L_i)$. T1 and T2, therefore, enable the decomposition of the two channels in (6).

1. demand effects and selective responding

The use of surveys helps study consumer beliefs around credit supply. However, survey collections could come with potential problems. For example, receiving the survey might induce the participants to respond or behave differently based on the anticipation of the survey senders' intention (survey demand effects). In addition, since taking the survey is time-consuming, the response rate is always less than perfect. If the decision to respond to the survey systematically varies with participants' characteristics, treatment effects will suffer from selection biases.

Several features of the survey design aim to eliminate potential confounding effects from filling out the surveys. For example, since the survey is sent through the bank, participants might want to use the survey answers to signal better creditworthiness. To prevent the participants from developing such strategic motives, the survey starts by showing the participants

> *This study is in collaboration with [the author], an Assistant Professor in Economics [...]. The data will only be analyzed by [the author] for scientific research purposes and will not be evaluated by this bank. We will not disclose participants' personal information in any respect. We will not, to any extent, change the types of financial products we provide, including credit scores, credit limits, deposit and borrowing interesting rates, etc., based on the participants' personal answers. Please answer the survey based on your true thoughts.*

This explicit framing, especially the disclosure that the data would not be analyzed by the bank, was designed to minimize the possibility that consumers provide answers that depart from their true beliefs in the hope of obtaining better services from the bank. I further check this concern by focusing on consumers whose borrowing relationship are with other banks both before and after the experiments in Section 4.C. Since these consumers do not borrow from this bank, they are expected to have less incentive to cater to the bank in the sample.

Another concern is the loss of sample representativeness due to selectively responding to the survey. To exclude this problem, the survey design tried to minimize the time to complete with a very large compensation. In total, participants have to take 15 questions. In addition, there are six questions that are sent to a random 30% of the participants. The average time to complete the survey was around 6 minutes, and the compensation was 15 CNY. This is equivalent to an hourly rate of around 150 CNY, which is higher than the 95th income percentile of all urban residents in China.

In the end, the survey response rate is around 67%. This is a very high number compared with previous literature. The high response rate is due to the high reward associated with completing the survey. However, due to imperfect response rates, samples with and without surveys differ in some dimensions. Section 3.D. compares the distribution of the surveyed sample and the unsurveyed sample. In general, there are more males completing the surveys. Those who complete the survey are more likely to be younger, less educated, earn less, and less wealthy. However, the differences are not especially large between the two samples.

2. sanity checks of survey answers

To check the quality of the surveys, I compare the survey answers about past income with the information from the bank's database. Panel A of Figure B.1 from the Appendix presents the binned scatter plot of consumers' average monthly incomes over the 12 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 0.74. This finding indicates the high quality of the survey.

3. external validity of the experiment

Since the experiment is based on the bank's usual internal underwriting process, the selected consumers might be systematically different from the average Chinese consumer.

This is because credit supply usually targets those whom the banks perceive as having more needs for credit. A potentially selective sample casts doubt on the external validity of the experiment.

To assess the representativeness of the sample, I compare the demographics between the sample and a 3% random sample for the bank database covering all customers. Since the bank is one of the largest banks in China, its customer base should be representative of the overall Chinese urban residents. Table B.1 presents the results, as expected, in general, the participants in the sample have fewer spending, income, saving, and credit limits, and more debt. That is, the participants in the experiments seem to have a larger need for more credit. However, the differences are not excessively large. For all characteristics, the differences are below 15%. Therefore, the sample is broadly representative of the whole Chinese urban population.

D. Summary Statistics

Table 1 gives the summary statistics based on pre-experiment characteristics. Panel A summarizes those in Sample I, and Panel B summarizes those in Sample II. Sub-panels 1, 2, and 3 respectively describe the control group, T1, and T2. For the sample without surveys, the average age of the participants is around 39 years old, and around 50% are female. Over half of the participants have college degrees. The average outstanding interest-incurring debt is about 7.5 thousand CNY, and around 17.5 CNY if conditional on holding a positive amount of debt. A simple calculation indicates around 43% of the participants hold positive unsecured 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). The average increase in credit limit is around 11.8 thousand CNY. This magnitude is economically significant. It is around 14% of the pre-experiment average total credit limit and around 10% of the average pre-experiment annual income. Columns (7) and (12) give the t-statistics testing the differences between the control and treatment groups. All samples are quite balanced, with no statistical differences among any of the dimensions. This indicates the effectiveness of the randomization.

TABLE 1. Summary Statistics

This table gives the summary statistics of the sample. Panel A is based on the sample without surveys (sample I), and panel B is based on the sample with the surveys (sample II). The units of the variables excluding Age, Female, and College are in thousands of CNY. The column Diff gives the differences in the average values between the given group and the control group. t-stats are the associated t-statistics, testing the significance of the differences in the means. All variables are winsorized at 1% level.

	Mean	SD	N	Mean	SD	Diff	t-stats	N	Mean	SD	Diff	t-stats	N
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
	Panel A: Sample I – Without Survey												
	Panel A1: Control			Panel A2: T1					Panel A3: T2				
Age	38.82	9.85	2050	38.71	9.49	-0.11	-0.38	2331	38.81	9.91	-0.01	-0.01	1119
Female	0.50	0.50	2050	0.50	0.50	0.00	0.18	2331	0.50	0.50	-0.01	-0.32	1119
College	0.56	0.50	2050	0.55	0.50	-0.02	-1.20	2331	0.54	0.50	-0.02	-1.01	1119
Income	10.46	8.41	2050	10.18	7.19	-0.04	-1.17	2331	10.66	8.20	0.21	0.70	1119
Saving	168.13	214.53	2050	171.39	184.64	0.53	0.31	2331	172.42	208.51	4.29	0.57	1119
Debt	7.34	13.48	2050	7.25	9.83	-0.01	-0.24	2331	7.91	14.45	0.57	1.24	1119
Debt|Debt>0	17.69	15.95	851	16.86	7.90	-0.13	-1.33	1003	18.03	17.14	0.34	0.44	491
Limit	87.29	99.45	2050	85.99	103.65	-0.20	-0.41	2331	89.39	117.43	2.10	0.54	1119
Δ Limit	11.93	8.75	2050	11.62	7.73	-0.05	-1.25	2331	12.02	9.02	0.09	0.27	1119
	Panel B: Sample II – With Survey												
	Panel B1: Control			Panel A2: T1					Panel A3: T2				
Age	37.91	10.25	1588	37.62	9.50	-0.29	-0.74	1875	37.83	9.75	-0.08	-0.16	1097
Female	0.43	0.50	1588	0.41	0.49	-0.02	-1.01	1875	0.42	0.49	-0.01	-0.52	1097
College	0.46	0.50	1588	0.47	0.50	0.02	0.92	1875	0.47	0.50	0.01	0.52	1097
Income	9.69	8.74	1588	9.48	6.99	-0.21	-0.65	1875	9.83	8.65	0.15	0.35	1097
Saving	139.05	149.84	1588	139.91	138.71	0.86	0.15	1875	146.14	139.33	7.09	0.97	1097
Debt	7.40	13.37	1588	7.01	10.10	-0.39	-0.79	1875	7.07	13.40	-0.33	-0.54	1097
Debt|Debt>0	16.54	15.76	711	16.99	8.81	0.45	0.54	774	16.39	16.25	-0.15	-0.15	473
Limit	23.25	27.19	1588	22.33	26.53	-0.92	-0.81	1875	24.18	31.89	0.94	0.65	1097
Δ Limit	12.01	9.09	1588	11.73	8.20	-0.29	-0.82	1875	12.40	8.94	0.39	0.87	1097
	Survey												
Spending	7.77	13.57	1588	7.94	12.92	0.17	0.32	1875	8.09	11.09	0.32	0.48	1097
Income	9.61	9.60	1588	9.54	7.75	-0.07	-0.19	1875	9.95	10.88	0.34	0.73	1097
Liquid Wealth	155.59	266.48	1588	150.71	188.84	-4.88	-0.53	1875	154.00	233.84	-1.59	-0.14	1097
Total Wealth	431.39	900.12	1588	426.85	678.23	-4.54	-0.14	1875	443.92	725.35	12.53	0.31	1097

IV Results

A. Consumption Responses to Limit Extensions

I first present results about the consumption dynamics around the experiment. The analysis only focuses on Sample I, to which surveys are not sent. As guided by Proposition 2, suppose credit limit affects consumption only through the precautionary motive, as usually suggested in the buffer stock model, then one should expect similar spending dynamics for both treatment groups because the realized changes in credit limits are statistically indifferent between the two groups. However, if the supply of credit limits

Figure 2. Evolution of Debt and Spending - Unsurveyed Sample

This figure plots the evolution of total non-durable debt and spending on both sides of the experimental period for Sample I. In each panel, the x-axis gives the months in 2023. The solid red line shows the evolution of T1, the blue dashed line shows the evolution of T2, and the gray dotted line shows the evolution of the control group. The gray vertical line gives the time of the treatment. All lines are vertically shifted so that the value for the control group at the treatment time is 0.

A: Non-durable Debt B: Spending

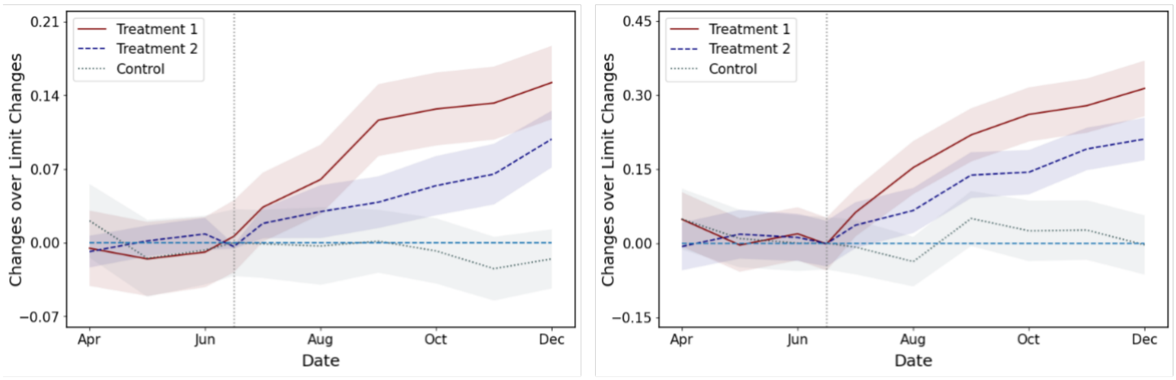

affects consumer beliefs, then the consumption response for those in treatment group 2, after informing about the randomness in the supply decisions, should be different.

Figure 2 plots the evolution of the changes in unsecured debt and total spending around the experiment. I scale the changes around the experiment by the pre-determined limit changes. So, the magnitudes give an interpretation in terms of marginal propensity. The x-axis is the months in 2023. In both plots, the solid red line and the dashed blue line represent T1 and T2, and the dotted gray line represents the control group. The shaded regions are two times the standard errors. Both debt and spending are residualized by month-fixed effects. As shown, the sharp increase in spending right after the experiment for the two treatment groups indicates the effectiveness of the experiment. Besides, the spending response of T2 is significantly smaller than that in T1. A divergence in the evolution of debt and spending between T1 and T2 indicates changes in credit limit affect factors other than instant borrowing capacity.

I continue to study the average treatment effects (ATE) of credit limits on spending. Table 2 gives the intent-to-treat (ITT) estimates of the experiment. I scale the estimates by the average changes in credit limit in the sample. In this case, the numbers in Table 2 have the interpretation of MPB and MPCL. That is, average changes in borrowing and spending for each CNY higher credit limit. Panels A and B, respectively, give the

TABLE 2. The Effects of Limit Increase on Debt and Spending

This table assesses the effects of credit extension on non-durable debt and spending. Panel A is the three-month changes and panel B is the six-month changes. T1 and T2 are respectively the two treatment group identifiers. Coefficients are divided by the pre-determined average increase in credit limit to give an interpretation of marginal propensity. In each column, Difference is the difference in the estimates between T1 and T2. Controls include gender, province fixed effects, industry fixed effects, a dummy variable labeling if the participants are younger than 38, and a dummy variable for having at least a college degree. All variables are winsorized at the 1% level.

	Panel A: 3 Months				Panel B: 6 Months			
	ΔB	ΔB	ΔC	ΔC	ΔB	ΔB	ΔC	ΔC
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T1	0.088***	0.086***	0.201***	0.197***	0.157***	0.152***	0.374***	0.367***
	(0.007)	(0.009)	(0.028)	(0.033)	(0.012)	(0.015)	(0.042)	(0.050)
T2	0.059***	0.055***	0.126***	0.124***	0.104***	0.100***	0.240***	0.232***
	(0.007)	(0.008)	(0.034)	(0.034)	(0.012)	(0.013)	(0.044)	(0.044)
Difference	0.029***	0.031***	0.074**	0.073**	0.054***	0.052***	0.135***	0.134**
	(0.007)	(0.011)	(0.031)	(0.035)	(0.014)	(0.018)	(0.051)	(0.058)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
N	5500	5500	2221	2221	5500	5500	2221	2221

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

3-month and 6-month responses. As shown, each CNY higher credit limit increases the borrowing of T1 by 8.6 cents over three months and 15.2 cents over six months. At the same time, each CNY higher credit limit increases the spending of T1 by 19.7 cents over three months and 36.7 cents over six months. These estimates are close to the documented MPB out of credit limit and MPCL in the previous literature (Gross and Souleles, 2002; Agarwal et al., 2017; Aydin, 2022), which is usually in the range of \$0.09 to \$0.20 for MPB and 0.2 to 0.6 for MPCL(Agarwal et al., 2017). The spending responses are much larger than the debt responses. This is consistent with both the buffer-stock model and credit limit changing beliefs. For example, in the buffer-stock model, even for consumers with high liquidity, a larger credit limit reduces the precautionary motive and increases consumption by reducing total savings.

For comparison, each CNY higher credit limit increases the borrowing of T2 by 5.5 cents over three months and 10 cents over six months. At the same time, each CNY higher credit limit increases the spending of T2 by 12.4 cents over three months and 23.2 cents over six months. These differences are significantly different from zero, indicating that there is a belief channel that affects spending responses to limit changes.

Figure 3. Evolution of Debt and Spending - Surveyed Sample

This figure plots the evolution of total non-durable debt and spending on both sides of the experimental period for Sample II. In each panel, the x-axis gives the months in 2023. The solid red line shows the evolution of T1, the blue dashed line shows the evolution of T2, and the gray dotted line shows the evolution of the control group. The gray vertical line gives the time of the treatment. All lines are vertically shifted so that the value for the control group at the treatment time is 0. The analysis is based on the sample of participants to whom questions 16 to 21 are not sent.

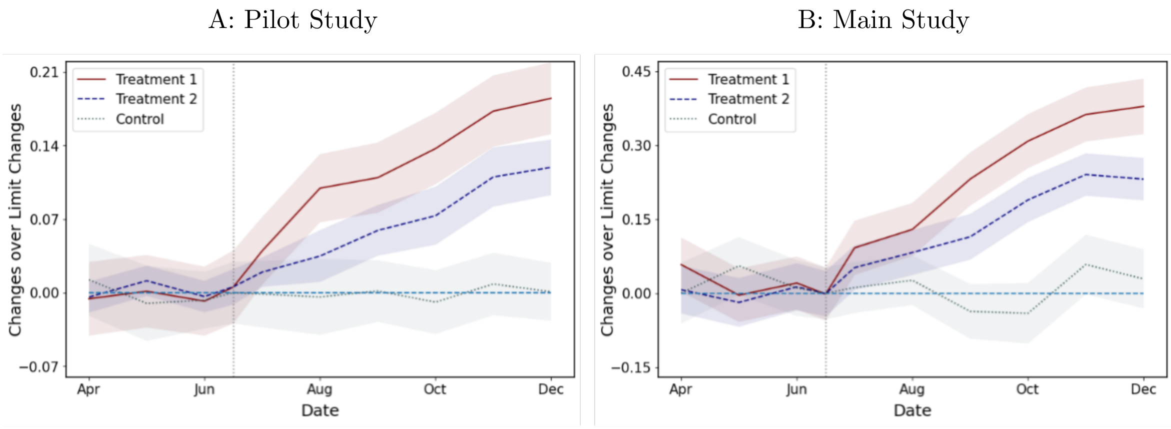

B. Expectation Responses to Limit Changes

Informing that the credit supply decision involves randomization reduces consumption responses to limit extensions by more than one-third. This indicates that credit supply affects consumption decisions in addition to relaxing the instantaneous borrowing constraints. In this section, I use the survey data from Sample II to dissect the effects of credit supply on consumer subjective beliefs about various components of their budget constraints.

Since the survey response rate is not perfect, I first study the consumption responses of the surveyed sample to compare the two samples. The evolution of debt and spending for the surveyed sample is in Figure 3. Consumption responses are generally slightly larger for the surveyed sample, as these consumers have relatively less liquidity. However, the differences are not significantly large, with the 6-month responses only around 13% larger. In addition, the patterns are very similar across the two samples.

The ITT estimates of limit changes on expectations are in Table 3. Similar to the estimation of MPB and MPCL, I scale all estimates by the average changes in credit limits. So, the coefficients are in terms of each CNY increase in credit limit. For changes

TABLE 3. The Effects of Limit Increase on Expectation Changes

This table assesses the effects of credit extension on expectations. E[ΔC], E[ΔY], and E[ΔHrs] are respectively the difference between expected total spending, total income, and hours to work every week over the 12 months after and before the experiment. E[ΔL.W] and E[ΔT.W] are respectively the difference between expected liquid wealth and total wealth 12 months after the experiment and right before the experiment. E[u] and E[p(d)] are the expected unemployment probability and delinquent probability over the 12 months after the experiment. E[ΔL]-1Y and E[ΔL]-5Y are the expected growth rate of one-year and five-year credit limits. T1 and T2 are respectively the two treatment group identifiers. Coefficients are divided by the pre-determined average increase in credit limit to give an interpretation of marginal propensity. All variables are winsorized at the 1% level. The analysis is based on the sample of participants to whom questions 16 to 21 are not sent.

	E[ΔC]	E[ΔY]	E[ΔL.W]	E[ΔT.W]	E[ΔHrs]	E[u]	E[p(d)]	E[ΔL]-1Y	E[ΔL]-5Y
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
T1	0.277**	0.381***	0.001*	0.001	0.000	-0.283*	-0.064	0.943	0.474
	(0.136)	(0.080)	(0.000)	(0.002)	(0.000)	(0.156)	(0.170)	(9.874)	(2.042)
T2	-0.030	0.079	0.000	-0.002	0.000	-0.020	0.133	0.958	0.695
	(0.102)	(0.100)	(0.001)	(0.002)	(0.000)	(0.214)	(0.209)	(0.707)	(3.044)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	5125	5125	5125	5125	5125	5125	5125	5125	5125

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

in wealth and credit limits, the units are in terms of each thousand CNY increase in credit limit. The results from T1 show that a higher credit limit significantly increases subjective expectations about future consumption and income, and marginally higher expected liquid saving and lower expected unemployment rate. However, there are no significant changes in subjective labor supply as captured by the number of hours likely to work. At the same time, future borrowing capacity, as captured by the one-year and five-year change in total credit limit or default probability all stays unchanged. As for T2, when informed about the randomization in credit supply, ex-ante expectations about consumption, income, saving, and unemployment all become insignificant.

The results in Table 3 suggest that, in response to a higher credit limit, consumers believe that they will consume more in the future, consistent with the empirical findings in the literature. In addition, the higher consumption is believed to be financed by more income in the future due to either higher marginal productivity of labor or lower unemployment risk, but not through drawing down savings, increasing default frequency, or increasing labor supply. Indifferent responses to subjective limit growth suggest that the information treatment attenuates consumption responses by erasing consumers' updates about future earnings ability rather than indirectly informing a less persistent increase in credit supply.

The findings show that, from the consumer perspective, the reason for more spending after a higher credit limit is inconsistent with the buffer-stock model. In the buffer-stock model, a higher credit limit increases consumption by reducing savings. This is because a higher credit limit alleviates precautionary motives by increasing the ability to smooth consumption. However, from Table 3, subjective beliefs about total wealth do not decrease; rather, they increase marginally. The results, therefore, suggest that the precautionary motive is unlikely the sole reason for which the supply of credit lines affects consumption.

In addition, tables 2 and 3 show that increases in credit limits raise consumption and expectations only about future total income, either through a lower unemployment rate or higher income conditional on having the same job. Informing the randomness of the supply decision eliminates the belief changes and reduces consumption responses by around 37%. Mapped to (5). The results indicate a weight of 37% for the income-inference channel through which limit extensions affect total consumption.

A concern of sending out surveys through the bank is that consumers might want to misreport creditworthiness to signal a lower risk type. This is unlikely in this study for two reasons. First, the disclosure information on the survey explicitly informs the participants that the bank will not analyze the data. Second, from Table 3, even though income is reported to be higher, subjective beliefs about default rates aren't smaller[9]. Therefore, it is improbable that consumers are signaling to have lower risk. An additional test is to focus on the expectation and debt responses for those whose borrowing relationship is with other banks. If consumers use the survey to strategically misreport income, then those who do not use this bank for daily transactions would have less incentive to misreport. In the online appendix Table B.3, I restrict the sample to consumers who only use credit cards at other banks both before and after the experiment. Panel A gives the results for the unsurveyed sample (Sample I), and Panel B focuses on the surveyed sample (Sample

[9] A possibility is that many consumers who have self-control problems but are sophisticated to be aware of this problem would conjecture an over-spending behavior, which is followed by a higher default probability.

Figure 4. Distributions of Belief Changes

This figures plots the expected changes in consumption (left column) and income (right column) using the sample receiving the post-experiment surveys (sample II). Panels A and B give the control group; panels C and D give the treatment group 1; panels E and F give the treatment group 2. The illustration is based on samples winsorized at 5% level. The analysis is based on the sample of participants to whom questions 16 to 21 are not sent.

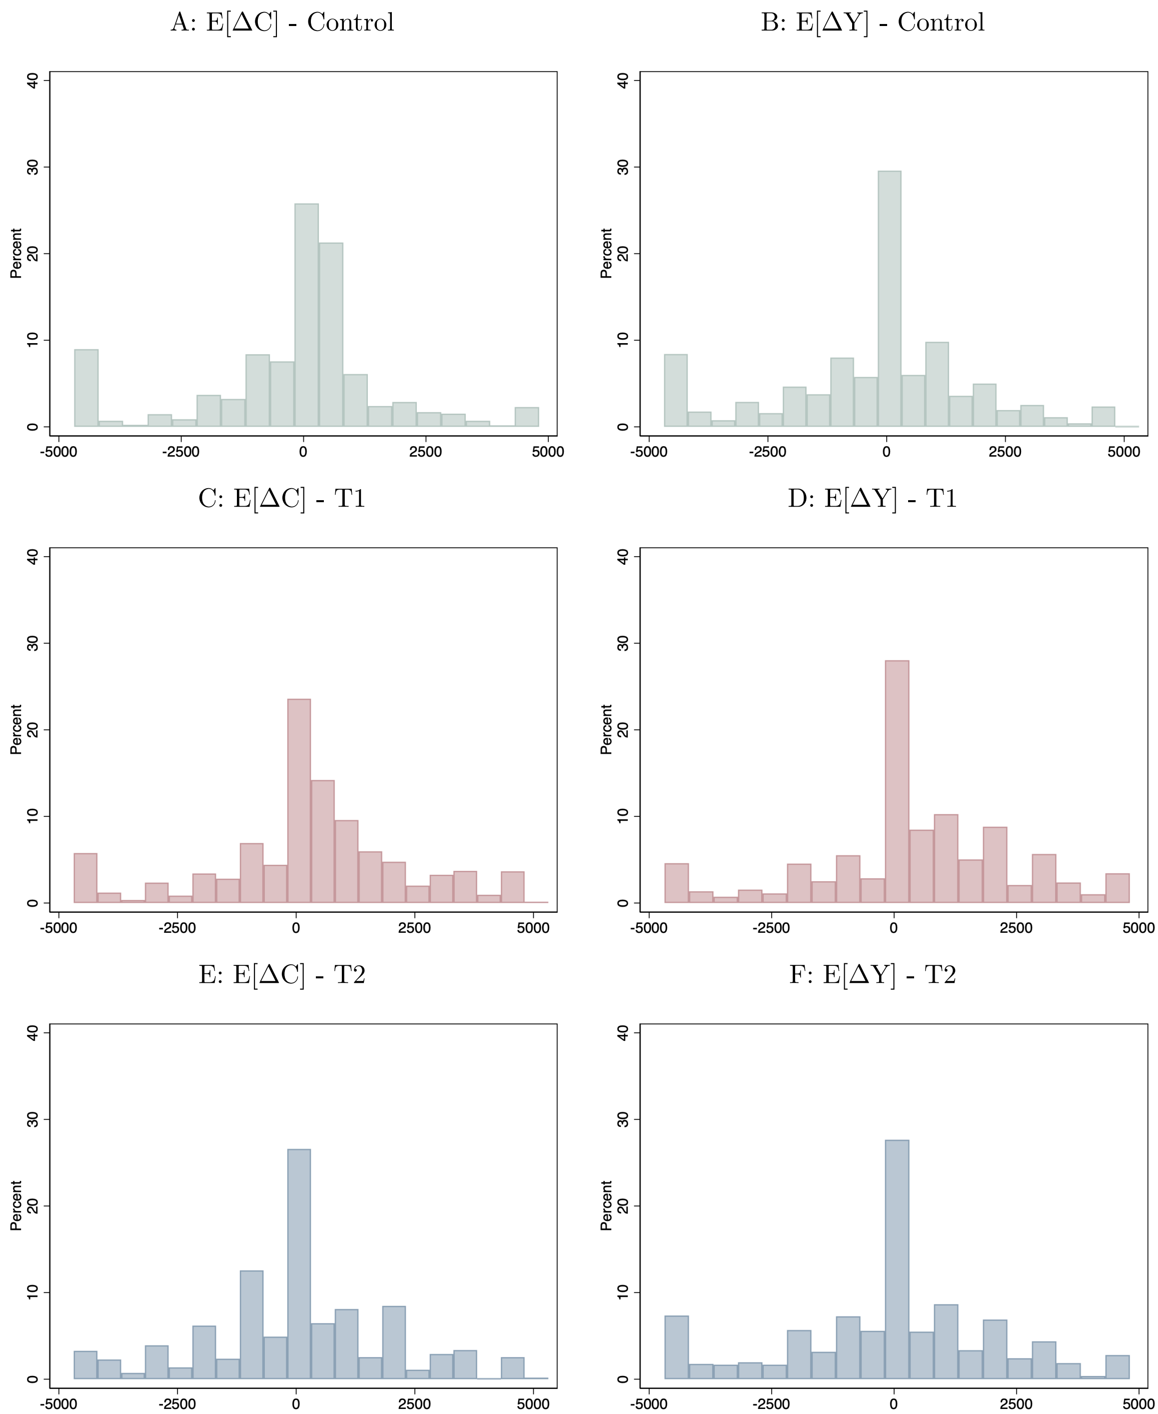

II). The results are close to that using the whole sample. This finding suggests that the reported higher income is unlikely a result of strategic misreporting to cater to the bank. In sum, the results in Table 3 suggest that credit expansions positively change consumers' beliefs about labor productivity, increasing consumption through a higher expected income.

Limit shocks have a large average impact on income expectations. However, behaviors of inferencing from credit supply should be heterogeneous and largely depend on factors like ex-ante degrees of income uncertainty. Heterogeneity in belief changes is also evident from directly exploring the distributions. Figure 4 plots the histograms of expectation changes separately for the three groups. As shown, for both the control group and treatment group 2, the changes in beliefs are more symmetric around zero. For treatment group 1, expectation changes with respect to consumption and income are more distributed to the positive region. However, the belief changes are not entirely positive, with around 45% of the participants having a negative or zero change in the belief of future income after the experiment. At the same time, the impacts are very heterogeneous, with many people having no belief change. Despite measurement errors from surveys, Figure 4 shows that the large average belief change is due to a large proportion of consumers having large belief changes, instead of everyone having similar belief changes.

C. Heterogeneity by Income Uncertainty

As suggested by Proposition 1, if consumers infer information about income growth from credit supply, then the degree of updating is larger when the signal-to-noise ratio is larger. Directly testing this hypothesis requires observing the signal-to-noise ratio. Since I do not observe the precisions of bank signals, I instead focus on income uncertainty. In particular, when consumer income is less uncertain, the difference between the ATEs of treatments 1 and 2 is supposed to be smaller. In the extreme case, when income has zero variation, there should be no difference between the ATEs of treatments 1 and 2.

To test this hypothesis, I first study how the ATEs associated with the two treatments differ by consumer income volatility. Since, for many participants, income is only observed for a short number of years, measuring individual income volatility is difficult. Therefore, I use two proxies for income uncertainty. First, to measure income variability

TABLE 4. The Effects of Limit Increase on Debt and Spending by Income Uncertainties

This table assesses the effects of credit extension on non-durable debt and spending by income uncertainties. Panels A1 and A2 split the sample by whether pre-experiment income has zero or positive volatility. Panels B1 and B2 split the sample by whether consumers have the type that has a low or high cross-sectional income variability. Splits are conditional on deciles of pre-determined limit changes. Coefficients are divided by the pre-determined average increase in credit limit to give an interpretation of marginal propensity. All variables are winsorized at the 1% level. The analysis is based on the sample of participants to whom questions 16 to 21 are not sent.

	ΔB	ΔC	E[ΔY]	ΔB	ΔC	E[ΔY]
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A1: SD(Y) = 0			Panel A2: SD(Y) > 0		
T1	0.081***	0.249***	-0.031	0.172***	0.433***	0.610***
	(0.025)	(0.083)	(0.030)	(0.018)	(0.060)	(0.118)
T2	0.078**	0.223**	-0.007	0.107***	0.249***	-0.073
	(0.035)	(0.090)	(0.035)	(0.014)	(0.057)	(0.149)
Diff	0.003	0.026	-0.024	0.064***	0.184**	0.683**
	(0.039)	(0.122)	(0.042)	(0.021)	(0.079)	(0.192)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	1165	1152	1141	4326	1068	3984
	Panel B1: Low XS Var			Panel B2: High XS Var		
T1	0.121***	0.255***	0.203**	0.197***	0.517***	0.594***
	(0.018)	(0.048)	(0.103)	(0.024)	(0.092)	(0.144)
T2	0.091***	0.203***	0.087	0.106***	0.270***	0.058
	(0.017)	(0.054)	(0.112)	(0.017)	(0.072)	(0.211)
Diff	0.029	0.051	0.116	0.091***	0.247**	0.536**
	(0.021)	(0.058)	(0.157)	(0.026)	(0.110)	(0.247)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	3084	1269	2562	2410	952	2563

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

within individuals, I use a coarser measure by splitting individuals based on whether the variability of monthly income is zero or positive within the 12 months before the experiment. A caveat is that since at the intensive margin, the amounts of increased credit limit are not random but based on the bank's assessment of consumers' creditworthiness, unconditionally separating the customers by income volatility may be inducing selection bias due to different amounts of increases in credit limit. In this case, ATEs could be different only because MPCL is not linear. To cope with this concern, I first separate consumers into ten groups by the proposed amounts of limit increase and then split the consumers by income volatility within the ten limit-increase groups. Doing so controls the bank's assessment of the consumers' change in creditworthiness.

The results are in Table 4. Columns (1) to (6) give the ATEs separately for those whose log monthly income has zero and positive variation. As shown, for those whose income has zero variation, there is no change in income expectations for either treatment group. In addition, differences between the ATEs of the two treatments in spending responses become insignificant. Meanwhile, for those with positive income variations, changes in income expectations become strongly and significantly different from zero for T1. At the same time, differences between the ATEs of the two treatments in spending responses are also larger than the baseline estimates. The results in columns (1) to (6) are consistent with Proposition 1, such that when prior uncertainty about future income growth is likely to be low, information content in credit supply has less effect on consumption.

As a second measure, I construct a statistic of cross-sectional income variability. Specifically, I use a deep neural network to cluster individuals into 20 types using a random 3% sample of the bank's database over the five years before the experiment (around 5 million observations with more than 200 variables). Then, I calculate a cross-sectional variance of income (*XS Var*) as the within-type variance of log income growth in 2023. Given that consumers are usually less advantaged to observe income information about other consumers of similar type, this measure of cross-sectional income variability also suggests to what extent the bank is able to observe income patterns for similar individuals. Columns (7) to (12) of Table 4 give the ATEs separately for those whose *XS Var* is below and above the sample median. Consistent with the conjecture, the difference between the ATEs is larger when the *XS Var* is larger.

D. Expectation Changes and Spending Responses

In response to a higher credit limit pushed by banks, consumers increase total consumption substantially. Survey results show that consumers also increase their expectations about future consumption and labor productivity. A natural question is whether the increases in consumption are, at least partly, driven by the changes in expectation. The comparison of the two treatment groups suggests that the differences in the spending responses are likely driven by belief differences after receiving credit extensions. In this section, I further test the relationship between spending responses and expectation changes. In

TABLE 5. The Effects of Limit Increase on Debt and Spending by Expectation Changes

This table assesses the effects of credit extension on non-durable debt and spending by expectation changes. Columns (1) and (2) split the sample by expectations of consumption changes. Columns (3) and (4) split the sample by expectations of income changes. Coefficients are divided by the pre-determined average increase in credit limit to give an interpretation of marginal propensity. All variables are winsorized at the 1% level. The analysis is based on the sample of participants to whom questions 16 to 21 are not sent.

	E[ΔC]		E[ΔY]	
	Low	High	Low	High
	(1)	(2)	(3)	(4)
	Panel A: ΔB - 6M			
T1	0.101**	0.252***	0.093**	0.259***
	(0.043)	(0.049)	(0.042)	(0.052)
T2	0.063	0.163***	0.053	0.173***
	(0.044)	(0.047)	(0.043)	(0.051)
Controls	Yes	Yes	Yes	Yes
N	2872	2253	2925	2200
	Panel B: ΔC			
T1	0.172*	0.562***	0.153*	0.586***
	(0.086)	(0.091)	(0.089)	(0.099)
T2	0.099	0.372***	0.073	0.398***
	(0.093)	(0.104)	(0.094)	(0.107)
Controls	Yes	Yes	Yes	Yes
N	1148	911	1162	897

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

Table 5, I estimate the ATEs of limit extensions on debt and spending by changes in expected consumption and income. If spending responses are partially driven by changes in expectations, then one should expect a larger spending response when changes in expectations are also large. Columns (1) and (2) split the sample by changes in consumption expectations by each group, and columns (3) and (4) split the sample by changes in income expectations by each group. The results show that for those who have larger changes in expectation, spending responses are also higher. This is consistent with expectation changes affecting consumption. Furthermore, in Table B.3 in the online appendix, I show that there is no difference in spending responses if ranked by changes in expectations with respect to other factors.

E. Subjective Beliefs about Credit Limit Extensions

As a direct test of the income-inference channel, I elicit consumer subjective beliefs about credit supply as a function of bank-perceived consumer future income growth. I rely on the following two questions from the survey:

> *Suppose banks increase your credit card limit by 5000 CNY this month. This would mean that the banks expect your total income to change by _____ over the next 12 months.*
>
> *Note: use a negative number for decreases.*

> *Suppose banks increase your credit card limit by 10000 CNY this month. This would mean that the banks expect your total income to change by _____ over the next 12 months.*
>
> *Note: use a negative number for decreases.*

These questions are sent to a 30% random sample of the participants. Suppose the answers from the two questions are respectively $x_{1,i}$ and $x_{2,i}$, I then calculate the consumers' subjective beliefs about the credit limit sensitivity to bank-perceived income growth, λ_i, as

$$\lambda_i = \frac{x_{2,i} - x_{1,i}}{5000}. \tag{7}$$

Mapped to (2), $\lambda_i = g'(L_i)$ is the marginal relationship between credit limit and bank beliefs about consumers' future income growth. When $\lambda_i = 1$, consumer i believes that the bank's supply of credit limit moves one-for-one with the bank's prediction about consumers' future income changes.

Figure 5 plots the distribution of λ_i. It shows that there is a large heterogeneity in consumers' subjective beliefs about the sensitivity of credit supply to bank-perceived income growth. Around 35% of the consumers believe $\lambda_i \leq 0$. However, most of the participants believe credit-limit extensions are associated with higher income growth in the future. The economic significance of λ_i is large. Its average value is 0.81, and the median is 0.60. Thus, for a 1-CNY increase in the credit limit, consumers, on average, believe the bank expects their income to increase by 0.81-CNY over the next 12 months. From a Bayesian-learning perspective, Figure 5 Panel A suggests that consumers, on average, learn about their future income from credit limit extension as a signal of income

Figure 5. Timeline of the Experiments

In this figure, panel A plots the distribution of consumer subjective beliefs about the sensitivity of income growth as perceived by the bank to credit extension, λ_i, which is calculated using (7). The plot is cut at 1% level. The right plot gives the changes in income expectations for each CNY higher pre-determined increase in credit limit. The estimates are conditional on four λ_i groups. Splits of λ_i groups are conditional on treatment groups and limit-increase deciles.

A: Pilot Study

B: Main Study

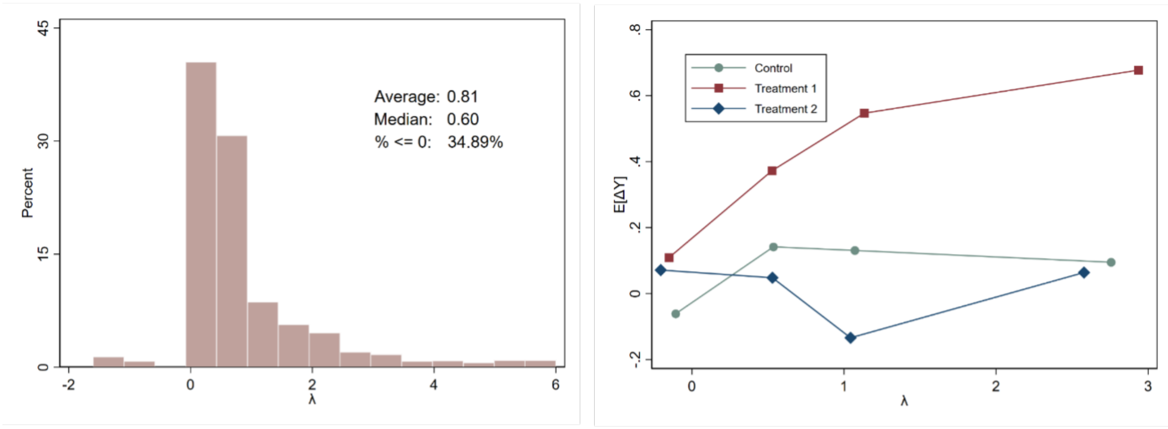

changes, with a signal sensitivity of 0.81. Given that the posterior income expectation is 0.381, the Kalman gain of the learning process is around 0.47.

Proposition 1 states that change in income expectations after receiving limit extensions should move positively with the signal sensitivity of income growth λ_i. In Panel B of Figure 5, I split the sample by λ_i into four groups and then plot the average change in income expectations by λ_i-groups within each treatment group. Consistent with proposition 1, participants in T1 have a larger change in income expectations after the experiment, and this change increases with λ_i. Income changes are also close to zero, especially when λ_i is close to zero. At the same time, there is no apparent association between λ_i and changes in income expectation for the other two groups. In sum, the results from Figure 5 indicate that consumers believe that limit extensions are positively associated with banks' beliefs about future income growth. Consistent with Bayesian learning, consumers with uncertain income streams adjust income expectations upwards in response to a positive credit supply shock.

F. What Information is Inferred?

Consumers, in spite of the private information about earning ability, could still have very noisy expectations about future income. Nonetheless, (1) does not take a stand on the source of the systematic components, especially from the consumers' subjective perspective. One explanation is that the systematic components represent macroeconomic shocks, including productivity shocks over the business cycles. If this is the case, consumers believe that banks can predict the macroeconomic outcomes, either due to being less inattentive to the current economic status or having better knowledge about the aggregate process's structural parameters (Mankiw et al., 2003; Giglio et al., 2021; Andre et al., 2022; Link et al., 2023). To test this possibility, I use the following question

> *Q16: How confident are you in evaluating whether the overall economy is functioning effectively at the moment?*

This question is sent to a random 30% of the participants. It measures the degree to which the consumers are informed about the current macroeconomic state. If income inference is due to consumers updating perceptions about the current macroeconomic state from credit supply, then expectation changes should be larger for those who are less confident in evaluating the macroeconomic performance. To test this conjecture, I define those who answer *very confident* as having low macroeconomic uncertainty and the rest as having high macroeconomic uncertainty. I then study expectation changes separately for these two groups. The results are in columns (1) and (2) of Table 6. Consistent with the conjecture, those in T1 who report to have high macroeconomic uncertainty also have a larger change in income expectations.

I rely on the following four questions to test the macroeconomic channel further. These questions capture how credit extensions are associated with subjective beliefs about the macroeconomy and how macroeconomic shocks affect individual incomes.

> *Q14/15: How much will the overall Chinese economy/unemployment rate change (in percentage relative to the current level) over the next year?*

> *Q17: Suppose the overall economy in China grows by 5% relative to the current level over the next year. How would this affect your total income over the next year?*

> *Q18: Suppose the unemployment rate in China decreases by 10% relative to the current level over the next year, how would this affect your total income over the next year?*

TABLE 6. Limit Increase and Income Expectations – Macroeconomic Channel

This table assesses the effects of credit extensions on income expectations by how limit changes affect personal income through providing information on macroeconomic states. Macro Uncertainty is low if the participants have high confidence about the current macroeconomic state. g'_Y is the product of changes in GDP growth after the experiment and the perceived effects of GDP growth on individual income. g'_U is the product of changes in the unemployment rate after the experiment and the perceived effects of decreases in the unemployment rate. g'_Y (g'_U) is low if it is non-positive.

	Macro Uncertainty		g'_Y		g'_U	
	Low	High	Low	High	Low	High
	(1)	(2)	(3)	(4)	(5)	(6)
T1	0.199	0.446**	0.123	0.523***	0.093	0.575***
	(0.231)	(0.228)	(0.263)	(0.212)	(0.278)	(0.203)
T2	0.077	0.089	0.033	0.113	0.019	0.132
	(0.240)	(0.219)	(0.241)	(0.197)	(0.252)	(0.183)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	943	1367	844	1466	693	1617

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

Q17 and Q18 are sent to a random 30% of the participants. Q14 (Q15) gives the effects of limit extensions on the macroeconomic movements; Q17 (Q18) in turn gives how macroeconomic movements affect individual beliefs about income. The product of Q14 (Q15) and Q17 (Q18), adjusted by the changes in credit limits and hypothetical changes in GDP (unemployment rate), denoted as g'_Y (g'_U), gives the effects of limit extension on future income through changing perceptions of the macroeconomic condition. If limit extensions affect income expectations by signaling the heterogeneous impacts of macroeconomic states on individual income, then one should see those with positive changes in income expectation are also those with large g'_Y (g'_U).

To test this conjecture, I split participants into four groups: those with non-positive g'_Y (g'_U) and those with positive g'_Y (g'_U), conditional on treatment groups and proposed limit-change deciles. I then report the changes in income expectations by the four groups. The results are in columns (3) to (6) in Table 6. The results are consistent with that consumers infer information about the macroeconomic states from credit extensions. That is, income expectations are larger when g'_Y (g'_U) are larger. More notably, for those with non-positive g'_Y or g'_U), the effects of limit extensions on income expectations become much smaller and insignificant from zero. This indicates that information about the macroeconomic state is the economically more important factor consumers are inferring from limit extensions.

Consumers believe that credit supply is associated with an expansionary aggregate economy. Those who think that a booming period affects their income to a greater extent have a larger change in income expectations. However, recent literature documents that expansionary credit conditions are usually associated with deteriorated instead of improved economic conditions in the future (Lóepez-Salido et al., 2017; Mian et al., 2017). Figure B.2 in the online appendix provides further evidence in the US such that periods with higher growth of credit limits are also times with higher subjective future income growth but lower realized future GDP growth. The recent literature on the extrapolative-expectation formation process is a possible explanation for this seeming inconsistency. Suppose lending standards are looser during booming periods (Bassett et al., 2014; Fishman et al., 2020; Weitzner and Howes, 2023), then credit supply reacts positively to current economic shocks. Consumers, equipped with incomplete information about the current state of the economy, update their beliefs accordingly in response to credit expansions. With extrapolative expectation formation process (Bordalo et al., 2018), this positive news would be over-extrapolated to the future, amplifying the positive relationship between current credit supply decisions and expected future income. When misbeliefs are later resolved, consumption growth decreases, inducing a boom-then-bust pattern.

G. Implications from US Data

Presumably, the existence of income inference from credit supply depends on the nature of banks' decision-making process. For more sophisticated banks, supply reacts more to lenders' beliefs about the economy and how economic shocks affect consumers, while for less sophisticated banks, credit supply could be purely based on individuals' payment history. A natural question is if China is special as the banks use more statistical analysis during its lending practices. To shed light on the likelihood of the existence of the income-inference channel in other settings, I present some survey results based on US data collected through Prolific, an online survey platform[10]. Recent studies note that data from online survey platforms can have high quality when compared with traditional, larger surveys and field experiments (Haaland et al., 2023; Douglas et al., 2023).

Without an experiment and administrative 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 reported decisions under randomized hypothetical scenarios to show consumers' expectation changes after credit supply. This strategy is similar to the reported preference approach in estimating MPC out of one-time wealth shock (Shapiro and Slemrod, 2003; Jappelli and Pistaferri, 2014; Graziani et al., 2016; Parker and Souleles, 2019; Fuster et al., 2020; Jappelli and Pistaferri, 2020).

The survey design is similar to the main study. Participants are required to be credit-card holders that are either part-time or full-time employed. For the recruited participants, I first randomly split them into two groups. Then, for group 1, participants are asked to consider the following scenarios:

> *For the following questions, please imagine a scenario where your bank has chosen you to raise your credit card limit by X%.*

For group 2, participants are shown the following text:

> *For the following questions, please imagine a scenario where your bank has chosen you at random to raise your credit card limit by X%. This decision by the bank is entirely random and not influenced by any assessment of pertinent factors.*

In both texts, X is the hypothetical limit change and is randomly picked from $\{10, 15, 20, 25, 30\}$.

The two scenarios mirror the two treatments in the main study. Specifically, in addition to the information shown to group 1, group 2 also receives an additional assumption that the credit supply is purely random. By informing about the randomness of the credit supply, the information shown to group 2 aims to isolate any belief channels in credit supply. To elicit belief changes, I then ask the following five questions.

> *Q1-3: How much do you think your Y would change over the next year?*
>
> *Q4: What's the probability that you would default on your debt over the next year?*

[10] The survey design, results, and the collection method are provided in the Online Appendix C.

Q5: How many hours would you work on average every week over the next year?

Y is an element of {income, spending, saving}. Comparing how answers to these questions vary with X between groups 1 and 2 sheds light on the informational effects of credit supply. The results are shown in Table C.1 in online appendix C. Consistent with the findings in the Chinese setting, for group 1, a higher hypothetical increase in credit limit increases expectations about future consumption and income. Meanwhile, there is no change in saving, default probability, or labor supply. For the results from group 2, when credit supply is assumed to be random, there is no significant association between reported future economic activities and limit extensions. The results are similar to those documented in the main study. That is, credit supply increases beliefs about future income and consumption but not saving default rate or labor supply. In addition, if credit supply is purely random, changes in expected future income and consumption become insignificant. Therefore, consumers in the US are also likely to change expected labor productivity after receiving limit extensions.

V Conclusion

This paper represents an initial practice into understanding how changes in credit supply causally impact subjective beliefs at the micro level and, consequently, how these altered beliefs influence consumer spending and borrowing behaviors. Traditional studies on the macroeconomic effects of credit supply often assume that economic agents possess full-information rational expectations, leaving the analysis of credit supply's impact on beliefs largely unexplored. This study documents a notable finding: extensions in credit limits boost consumers' beliefs regarding their future labor productivity. It estimates that approximately 37% of the increased consumption following higher credit limits can be attributed to this shift in income expectations.

However, this is just the beginning of a broader investigation. There's a need for further research to comprehensively understand the macroeconomic implications when lenders and borrowers have access to different sets of information. Additionally, the study touches on the nuances of banks' credit supply decisions, which may vary based on the statistical precision that can be achieved with different borrower characteristics. For instance, credit supply decisions grounded in statistical analysis might disproportionately favor individuals for whom the bank can derive more accurate predictions (Fuster et al., 2022). This aspect raises interesting questions about the potential asymmetric impacts of monetary policies across various industries, influenced by banks' abilities to make statistical inferences. Future research could beneficially explore the distributional effects of monetary policy in scenarios where banks depend on statistical analysis to make credit supply decisions, further illuminating the complex dynamics at play in credit markets.

References

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.

Ameriks, J., J. Briggs, A. Caplin, M. Lee, M. D. Shapiro, and C. Tonetti (2020, January). Older americans would work longer if jobs were flexible. *American Economic Journal: Macroeconomics 12*(1), 174–209.

Ameriks, J., J. Briggs, A. Caplin, M. D. Shapiro, and C. Tonetti (2016, October). The long-term-care insurance puzzle: Modeling and measurement. Working Paper 22726, National Bureau of Economic Research.

Ameriks, J., J. Briggs, A. Caplin, M. D. Shapiro, and C. Tonetti (2020). Long-term-care utility and late-in-life saving. *Journal of Political Economy 128*(6), 2375–2451.

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.

Andre, P., C. Pizzinelli, C. Roth, and J. Wohlfart (2022, 02). Subjective Models of the Macroeconomy: Evidence From Experts and Representative Samples. *The Review of Economic Studies 89*(6), 2958–2991.

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.

Bailey, M., E. Dávila, T. Kuchler, and J. Stroebel (2019). House price beliefs and mortgage leverage choice. *The Review of Economic Studies 86*(6), 2403–2452.

Bassett, W. F., M. B. Chosak, J. C. Driscoll, and E. ZakrajÅ¡ek (2014). Changes in bank lending standards and the macroeconomy. *Journal of Monetary Economics 62*, 23–40.

Benjamin, D. J. (2019). Chapter 2 - errors in probabilistic reasoning and judgment biases. In B. D. Bernheim, S. DellaVigna, and D. Laibson (Eds.), *Handbook of Behavioral Economics - Foundations and Applications 2*, Volume 2 of *Handbook of Behavioral Economics: Applications and Foundations 1*, pp. 69–186. North-Holland.

Bordalo, P., N. Gennaioli, and A. Shleifer (2018). Diagnostic expectations and credit cycles. *The Journal of Finance 73*(1), 199–227.

Bucks, B. and K. Pence (2008). Do borrowers know their mortgage terms? *Journal of Urban Economics 64*(2), 218–233.

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

Coibion, O. and Y. Gorodnichenko (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy 120*(1), 116–159.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer (2016, March). Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review 106*(3), 664–98.

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

DellaVigna, S. (2009, June). Psychology and economics: Evidence from the field. *Journal of Economic Literature 47*(2), 315–72.

Doornik, B. F. N. V., A. R. Gomes, A. R. Gomes, D. Schoenherr, and Skrastins (2021, March). Financial access and labor market outcomes: Evidence from credit lotteries. Available at SSRN: https://ssrn.com/abstract=3800020.

Douglas, B. D., P. J. Ewell, and M. Brauer (2023, 03). Data quality in online human-subjects research: Comparisons between mturk, prolific, cloudresearch, qualtrics, and sona. *PLOS ONE 18*(3), 1–17.

Fishman, M. J., J. A. Parker, and L. Straub (2020). A dynamic theory of lending standards. Available at SSRN: https://ssrn.com/abstract=3665881.

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

Fuster, A., P. Goldsmith-Pinkham, T. Ramadorai, and A. Walther (2022). Predictably unequal? the effects of machine learning on credit markets. *The Journal of Finance 77*(1), 5–47.

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.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2021, May). Five facts about beliefs and portfolios. *American Economic Review 111*(5), 1481–1522.

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.

Haaland, I., C. Roth, and J. Wohlfart (2023, March). Designing information provision experiments. *Journal of Economic Literature 61*(1), 3–40.

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy 86*(6), 971–987.

Hansen, L. P. (2007, May). Beliefs, doubts and learning: Valuing macroeconomic risk. *American Economic Review 97*(2), 1–30.

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.

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.

Kuchler, T., M. Piazzesi, and J. Stroebel (2022, April). Housing market expectations. Working Paper 29909, National Bureau of Economic Research.

Link, S., A. Peichl, C. Roth, and J. Wohlfart (2023). Information frictions among firms and households. *Journal of Monetary Economics 135*, 99–115.

Lóepez-Salido, D., J. C. Stein, and E. Zakrajšek (2017, 05). Credit-Market Sentiment and the Business Cycle*. *The Quarterly Journal of Economics 132*(3), 1373–1426.

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

Mankiw, N. G. and R. Reis (2002, 11). Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve*. *The Quarterly Journal of Economics 117*(4), 1295–1328.

Mankiw, N. G., R. Reis, and J. Wolfers (2003). Disagreement about inflation expectations. *NBER Macroeconomics Annual 18*, 209–248.

Manski, C. F. (2004). Measuring expectations. *Econometrica 72*(5), 1329–1376.

Mian, A., A. Sufi, and E. Verner (2017, 05). Household Debt and Business Cycles Worldwide*. *The Quarterly Journal of Economics 132*(4), 1755–1817.

Müller, K. and E. Verner (2023, 12). Credit Allocation and Macroeconomic Fluctuations. *The Review of Economic Studies*, rdad112.

Nelson, S. T. (2022). Private information and price regulation in the us credit card market.

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.

Reis, R. (2006). Inattentive consumers. *Journal of Monetary Economics 53*(8), 1761–1800.

Sergeyev, D., C. Lian, and Y. Gorodnichenko (2023, May). The economics of financial stress. Available at SSRN: https://ssrn.com/abstract=4461635.

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

Soman, D. and A. Cheema (2002). The effect of credit on spending decisions: The role of the credit limit and credibility. *Marketing Science 21*(1), 32–53.

Weitzner, G. and C. Howes (2023). Bank information production over the business cycle. Available at SSRN: https://ssrn.com/abstract=3934049 or http://dx.doi.org/10.2139/ssrn.3934049.

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

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