

Subjective Income Expectations and Household Debt Choices

Francesco D'Acunto, Michael Weber, and Xiao Yin *

May 7, 2026

Abstract

Matched transaction-level, credit-registry, and survey data show that consumers over-extrapolate from unexpected income shocks to their income expectations. Excessively optimistic income forecasts raise current spending, and when income realizations fall short, lead to greater debt accumulation and higher default. A consumption-savings model with defaultable unsecured debt and extrapolative Kalman filtering, in which consumers overreact to income surprises, reconciles our empirical findings and shows extrapolative income expectations can generate state-dependent household debt cycles.

JEL codes: D14, D84, E21, E71, G51.

Keywords: Income Expectations, Behavioral Macroeconomics, Consumption, Household Debt, Behavioral Economics, Extrapolative Beliefs, Surveys.

*D'Acunto: Georgetown University, francesco.dacunto@georgetown.edu; Weber: ESMT, Purdue University, CEPR, and NBER, weber366@purdue.edu; Yin: UCL, xiao.yin@ucl.ac.uk. For very helpful comments and suggestions, we thank Elena Afanasayeva, Deniz Aydin, Richard Blundell, Pedro Bordalo, Anthony DeFusco, Laura Gáti, Michael Gelman, Yuriy Gorodnichenko, Sasha Indarte, Eungik Lee, Peter Maxted, Ben Moll, Emi Nakamura, Sean Myers, Søren Leth-Petersen, Stijn Van Nieuwerburgh, Arna Olafsson, Michaela Pagel, Alessandro Previtto, Stephen Terry, Mirko Wiederholt as well as seminar and conference participants at UC Berkeley, UCL, 2022 SED Meeting, 2023 AEA Annual Meeting, 2024 Duke-UNC Asset Pricing Conference, 2024 SFS Finance Cavalcade, FRB Conference on the Macroeconomic Seeds of Financial Imbalances, LSE-IFS-UCL-CEPR-Imperial Business School workshop on Household Finance, 2024 CES-ifo Venice Summer Institute on Expectations Formation, Stockholm Debt Relief Workshop, 2024 NBER Conference on Innovative Data in Household Finance, 2025 CEPR Belief and Macroeconomy conference, WFA 2025, Purdue University, Boston College, ITAM, Auckland University of Technology, University of Auckland, Monash University, UC Irvine, 2025 NBER SI Behavioral Macro, 2026 LSU Finance Conference, the 2026 Spring Finance Workshop, and NTU Behavioral Macroeconomics Workshop 2026. All errors are our own.

I Introduction

Understanding the causes and consequences of household debt decisions is important because spikes in household leverage reliably predict economic contractions across space and over time. Existing evidence is consistent with both supply- and demand-side channels as key factors shaping debt dynamics.¹ Disentangling these channels empirically is challenging: unobserved shocks may affect both credit supply and household borrowing demand, and these two forces often reinforce each other.

This paper develops a micro-to-macro approach to test if and how subjective income expectations affect household debt and to quantify the aggregate implications. We combine, to the best of our knowledge for the first time, transaction-level records on income and spending from bank accounts, debt positions and defaults from administrative credit-registry data, and repeated surveys eliciting account-holders' income forecasts.² From 2020 to 2025, we survey the same individuals about their six-month-ahead income expectations multiple times, which allows tracing the evolution of individuals' income forecast errors. We link these expectations data to realized income, consumption, debt, and default outcomes that we observe in administrative data, which enables us to study how income forecast errors relate to individuals' consumption, borrowing, and default decisions.

We start by documenting that the average consumer's income expectations overreact to unexpected income shocks after both positive and negative shocks: after an unexpected shock in either direction today, expectations about future income are systematically inaccurate in the direction of the shock relative to subsequent realizations. The size of the forecast error, defined as the difference between subjective income expectations over the subsequent six months and the realized income over the same period, increases proportionally with the size of the income shock. These results hold after keeping constant a rich set of individual-level characteristics. At the same time, the extent of overreaction is heterogeneous across demographic groups: it is larger for consumers who face more

¹See, e.g., Glick et al. (2010), Mian and Sufi (2014), Mian et al. (2017), Agarwal et al. (2017), Mian et al. (2020), Bordalo et al. (2018), Bianchi et al. (2021), and Chodorow-Reich et al. (2023).

²For recent reviews of survey-based research in economics, see Haaland et al. (2021) and D'Acunto and Weber (2024).

volatile incomes and for younger consumers. These patterns hold when we only exploit variation within individuals over time.

We then examine whether inaccurate income expectations are associated with consumption and borrowing behavior and find they are: the larger the forecast error, the greater the contemporaneous increase in spending. Moreover, consumers with larger forecast errors subsequently borrow more and face a higher likelihood of default.

These patterns are a robust feature of the data. For example, our results hold when restricting our sample to 2023 and subsequent years, which excludes the COVID-19-pandemic period and its aftermath. The results also stay nearly unchanged if we focus on individuals with small survey reporting errors on past income, indicating that inattention to current income is unlikely to be a relevant driver of our results.

To further strengthen a causal interpretation of our findings, we instrument individual income shocks and forecast errors with the average shock among similar consumers in other locations, who are not exposed to the same local demand and supply shocks, and its interaction with a high-income-volatility indicator. This instrument does not capture the component of income shocks induced by local business cycles. The IV estimates are very similar to the OLS counterparts.

Overall, the first part of the paper provides direct evidence at the micro level of an expectations-based channel of household debt accumulation. In particular, not only do consumers form systematically more optimistic expectations about future income following unexpected positive income shocks, but such over-optimistic beliefs lead to higher contemporaneous consumption. When, later, realized income falls short of expectations, consumers borrow more and are more likely to default.

The second part of the paper assesses the aggregate implications of our micro-level evidence. Our reduced-form empirical analysis mainly exploits variation in income shocks within individuals and over time. In the aggregate, such shocks are likely correlated across individuals: during periods of strong growth in the aggregate, many consumers experience positive income changes, part of which might be unanticipated. Given our reduced form results, a larger share of households should increase consumption and accumulate more debt. Conversely, when GDP growth slows or turns negative, a higher share of consumers

experience unexpected income declines, and should respond by cutting spending and reducing borrowing.

Based on this intuition, we augment a standard consumption model in partial equilibrium with incomplete markets and heterogeneous agents with defaultable debt to assess the role and dynamics of subjective income expectations for aggregate household debt cycles. In the model, consumers learn about the permanent innovation of their income process from their realized income. Meanwhile, we allow consumers to over-extrapolate income shocks by overreacting to signal surprise. To provide a benchmark, we compare the implications of this model with a version in which consumers are Bayesian learners. Our model matches several empirical moments of the liquidity distribution, such as the fraction of debtors in the economy as well as the average and median wealth to income ratios. The model also matches successfully the relationship between forecast errors and economic choices we observe in the micro data, but only in the version in which consumers overreact to signal surprise.

We use the model to simulate a heterogeneous-agent economy and study the effects of overreaction in income expectations on the dynamics of aggregate consumption, debt, and defaults. This form of expectations amplifies the effect of income shocks on income expectations and consumption. For this reason, agents start accumulating more debt after positive income shocks relative to the Bayesian benchmark. Once the economy is hit with a small negative shock, debt and default increase substantially relative to the Bayesian benchmark. These results are consistent with many patterns that have motivated the recent strand of research on household debt cycles in macroeconomics. For instance, elevated consumer sentiment implies higher subsequent growth in debt (López-Salido et al., 2017) and the end of a boom usually coincides with times of high financial fragility (Greenwood et al., 2019; Maxted, 2023).

Our evidence of a demand-side channel of household debt cycles does not imply the irrelevance of supply-side channels. In contrast, Aydin (2022) and Yin (2025) document the impact of banks' credit choices on consumers' debt and spending using randomized increases in credit card limits. Bornstein and Indarte (2023) find that expanding social insurance via Medicaid eligibility increases households' financial resilience and hence credit

supply, which leads to higher household debt accumulation. Kluender et al. (2024) show both beliefs (demand) and constraints (supply) drive low-income workers' consumption decisions.

We contribute to three areas of economics. First, to the rich literature on the consumption response to income changes, and especially on the marginal propensity to consume (MPC) out of income shocks (Parker et al., 2013; Kueng, 2018; Olafsson and Pagel, 2018; Fuster et al., 2020; Jappelli and Pistaferri, 2020; Baugh et al., 2021; Fagereng et al., 2021; Colarieti et al., 2024).³ Our study is closely related to work on the consumption response to total income shocks rather than one-time wealth transfers (Ganong et al., 2025; Krueger et al., 2023). We are the first to combine repeated surveys eliciting income expectations within individuals over time with transaction-level bank data and comprehensive credit-registry information, allowing us to directly link belief formation to realized income and financial behavior. This unique design offers micro-level evidence on how income forecast errors transmit to spending, borrowing, and default decisions through an expectation channel, and informs models of household consumption and saving as well as the design of social-insurance and redistribution policies.

We also contribute to the literature on credit cycles in the household sector. This literature has proposed two broad mechanisms: financial frictions that amplify shocks and generate cycles in credit supply (e.g., Kiyotaki and Moore, 1997; Gertler and Kiyotaki, 2010; Mian et al., 2020), and shifts in subjective beliefs that drive fluctuations in credit demand (e.g., Bordalo et al., 2018, 2021; L'Huillier et al., 2023). Our paper provides micro-level evidence consistent with the latter mechanism. We document that unexpected income shocks lead consumers' income expectations to overreact, ultimately resulting in excessive borrowing and higher subsequent default. This individual-level evidence complements studies linking aggregate belief dynamics to credit booms and busts (e.g., Kaplan et al., 2020; Mian and Sufi, 2018).

Lastly, we contribute to the literature linking beliefs to intertemporal consumption choices (see DellaVigna (2009); Benjamin (2019)). Existing work relates subjective beliefs to discrete or long-horizon decisions such as retirement, portfolio allocation, housing,

³See Attanasio and Weber (2010) and Jappelli and Pistaferri (2010) for earlier reviews and Kaplan and Violante (2022) for a recent review of MPCs in heterogeneous agent models.

and saving (e.g., Ameriks et al., 2020; Giglio et al., 2021; Kuchler et al., 2022; Beutel and Weber, 2025; Rozsypal and Schlafmann, 2023; Bellifemine et al., 2024), or studies how macroeconomic expectations shape consumption and saving decisions (e.g., Roth and Wohlfart, 2020; Coibion et al., 2022; D’Acunto et al., 2022; Kluender et al., 2024). Relatedly, Exler et al. (2024) examine the welfare implications of over-optimism about income risk in a model with lender screening. Finally, Balleer et al. (2025) study biased beliefs of labor-market transitions and quantify the implications for wealth inequality in a life-cycle model.

While recent studies exploit rich survey data on income and labor-market expectations, to the best of our knowledge no prior work exists that jointly observes individual expectations, realized income, and subsequent economic choices within the same administrative panel. We therefore focus on biased income expectations at the within-individual margin and study their joint effects on consumption, borrowing, and default. By combining repeated individual-level survey measures with transaction data and comprehensive credit-registry records, we link belief distortions to dynamic financial behavior not jointly observed in prior work.

II Expectation Formation

Before describing our empirical setting and results, we lay out a conventional income process to highlight how income forecast errors behave when agents’ expectation formation deviates from Bayesian learning. This section introduces notation and disciplines the estimation equations we will lay out in the empirical analysis.

Let Y_t denote income and $y_t = \log(Y_t)$ log income. Suppose log income follows:

$$\begin{aligned} y_t &= \alpha + z_t + \epsilon_t \\ z_t &= z_{t-1} + \eta_t, \end{aligned} \tag{1}$$

as in a standard buffer-stock model (Carroll, 1997; Blundell and Preston, 1998). Here, z_t is the permanent component in the income process. ϵ_t and η_t are i.i.d. normal shocks, with $E[e^{\epsilon_t}] = E[e^{\eta_t}] = 1$ and variances σ_ϵ^2 and σ_η^2 , respectively. α represents a time-invariant deterministic component of log income, which we assume is common knowledge.

Consumers do not observe the true value of z_t . Instead, each period they observe y_t

and update beliefs about z_t according to

$$\hat{z}_t = \hat{z}_{t-1} + (1 + \theta)\kappa [y_t - \alpha - \hat{z}_{t-1}], \quad (2)$$

where κ is the steady-state Kalman gain in a standard Kalman filter (Güvenen, 2007). The parameter θ captures deviations from Bayesian learning: when $\theta > 0$, consumers overreact to signal surprises.⁴ Throughout, we use $E_t[\cdot]$ to denote expectations under Bayesian learning and $E_t^\theta[\cdot]$ for expectations that deviate from Bayesian learning.

Define the income shock as $s_t \equiv y_t - E_{t-1}[y_t]$ and the one-period-ahead forecast error under overreaction as $FE_{t+1}^{\theta,y} \equiv E_t^\theta[y_{t+1}] - y_{t+1}$. We then have:

Proposition 1: The regression coefficient of $FE_{t+1}^{\theta,y}$ on s_t is $b_{log} = \kappa\theta$.

The proof is in Online Appendix Section I.A. Proposition 1 states income shocks positively predict next-period forecast errors only under overreaction ($\theta > 0$), because when $\theta = 0$, the Kalman orthogonality condition holds, so that income shocks are uncorrelated with future forecast errors, that is, agents fully condition their forecasts on all available information.

Note s_t is not directly observable because beliefs about the latent state are unobserved. However, we can express the log income process as:

$$y_t = y_{t-1} + s_t - (1 - \kappa)s_{t-1}. \quad (3)$$

Thus, we can recover s_t by estimating an ARMA(1,1) model for log income.

Later, we study the effects of income shocks on consumption and borrowing behavior. Since borrowing often equals zero at the individual level, estimating borrowing elasticities using log values is problematic. We therefore derive level analogues of income shocks and forecast errors to obtain estimates for marginal propensities. If we define level shocks as $s_t^Y \equiv Y_t - E_{t-1}[Y_t]$ and level forecast errors as $FE_{t+1}^{\theta,Y} \equiv E_t^\theta[Y_{t+1}] - Y_{t+1}$, we have:

Proposition 2: The regression coefficient of $FE_{t+1}^{\theta,Y}$ on s_t^Y is $b_{level} = \kappa\theta\Gamma$, where $\Gamma > 0$. For small shock variances, $\Gamma \approx 1$.

The difference between b_{log} and b_{level} arises from Jensen's inequality when moving from

⁴This specification could be microfounded, for instance, through overconfidence about signal precision (Daniel et al., 1998; Daniel and Hirshleifer, 2015; Angeletos et al., 2021).

logs to levels. As the variance of income shocks approaches zero, b_{level} converges to b_{log} .

Propositions 1 and 2 characterize the relationship between next-period forecast errors and current income shocks. These shocks are measured relative to objective forecasts rather than consumers' own beliefs. Alternatively, we can characterize the relationship between future forecast errors and current *income surprises*, defined as $s_t^\theta \equiv y_t - E_{t-1}^\theta[y_t]$, which is the difference between realized income and consumers' own ex-ante forecasts. Note income surprises equal the negative of current forecast errors. Regressing future forecast errors on current income surprises is therefore equivalent to estimating the negative of the first-order autocorrelation of forecast errors.

Proposition 3: The regression coefficient of $FE_{t+1}^{\theta,y}$ on s_t^θ is $\kappa\theta\tilde{\Gamma}$, where $\tilde{\Gamma} \in (0, 1)$.

Using income surprises yields qualitatively similar but attenuated results compared to using income shocks, due to additional noise in subjective expectations induced by overreaction.

III Institutional Setting and Data

We obtain the data we use in the paper from several sources. We collaborate with a large national commercial bank in China that ranks among the country's top ten banks by total assets. As of 2023, the bank held more than one trillion U.S. dollars in assets and served over 70 million active account holders.

We combine transaction-level data from this bank with multiple waves of a customized survey that elicits consumers' subjective income expectations. The survey covers a population representative of the Chinese banked population, and our main sample period spans from January 2020 to June 2025.

In addition, we obtained data from the Credit Reference Center (CRC) of the People's Bank of China, the country's official credit registry, which aggregates borrowing information from all financial institutions. These data allow us to observe the full liabilities of borrowers, that is, their total outstanding debt and its composition across lenders above and beyond the bank with which we collaborate, which eliminates concerns about unobserved debt accounts when using data from a single financial institution.

III.A Identifying Primary Bank Users

The credit registry does not collect spending information. To study consumption choices, we thus need to rely on the transaction-level data from the bank. Consumers might have multiple banking relationships and multiple spending accounts, though, which would be a concern for our analysis if consumers had systematically different spending habits across accounts and these habits were a function of income expectations and past income shocks.

To alleviate this concern, we follow recent research using single-provider transaction-level data (Ganong and Noel, 2019; Boehm et al., 2025) and impose two restrictions on the accounts that enter our empirical analysis to capture consumers who are likely to use the bank as their primary banking institution. First, we only consider consumers who have their regular income directly deposited at our partner bank. Second, we only consider active users that use the accounts at the bank for at least 20 transactions per month. In Section V.E, we verify that using different thresholds has no material impact on our results.

III.B Measuring Income, Spending, and Debt

We retrieve income information from account-level data. For employees, we calculate income based on social insurance contributions, because these payments typically represent a fixed fraction of total income, which includes salary, commissions, and bonuses. For self-employed individuals or employees with business-related income, this part of income is calculated from tax payments and the associated tax rates. This approach provides a reliable estimate of non-financial income.⁵ In Online Appendix Figure A.1, Panel A, we compare the income data used in our analysis with administrative income records from a provincial tax authority, which cover about 13% of our sample, to validate our measure. The plot shows a strong linear relationship with an R^2 of 0.92, indicating that our income measure closely matches the administrative benchmark for the overlapping consumers.

⁵In China, social security consists of five types of insurance and a housing fund. Each insurance is usually paid as a fixed share of monthly income; the retirement insurance, similar to U.S. retirement plans, mostly requires an 8% contribution. The income base is typically capped at both tails, with regional caps set at about 30%-300% or 40%-400% of the previous year's local average income. Since these caps affect only around 10% of workers, we exclude them from the final sample.

We compute monthly total spending as the sum of all purchase transactions excluding installment payments on mortgages and vehicles plus the total repayments of linked credit cards’ end-of-month balances between the end of the last billing cycle and the current billing cycle.⁶

Debt data come from the CRC, based on credit reports retrieved by the bank. CRC aggregates personal credit information from all financial institutions, including detailed monthly information on credit accounts (term loans, credit cards, and other personal credit lines), bank names, outstanding balances and limits, recent credit utilization, repayment history, end-of-billing cycle total amount payable and unpaid balances for each account, housing fund and social security contributions, etc. Given its comprehensive coverage, this dataset provides a complete view of consumers’ borrowing behavior. We define debt as the unpaid balance at the end of the interest-free billing cycle.⁷

The credit reports have information about the banks from which consumers obtain debt. For the subset of individuals who only have debt from our partner bank, we can check that the debt amounts reported on the credit report are accurate relative to the figure we see from the bank. In Online Appendix Figure A.2, for this subsample, we compare unsecured debt computed based on our partner’s bank data with the unsecured debt on the credit reports. We observe an almost perfect linear relationship, which corroborates the reliability of the credit registry data.

III.C Survey Design and Answers’ Reliability

To elicit consumers’ expectations, we designed a customized survey administered by the partner bank. Figure 1 shows the survey timeline. The bank conducted two waves per year, each at the beginning of January and July, from 2020 to 2025, with the exception of 2022. Each wave elicited expectations for the subsequent six months.⁸

To limit survey fatigue, the bank required a maximum of 12 questions to target a

⁶During our sample period, most Chinese residents made payments through Alipay or WeChat Pay. These platforms can be used either by linking debit or credit cards or by storing small balances in their built-in “change” wallets. Our data capture the card-linked transactions, which account for the vast majority of payments. A supplementary bank survey indicates that only 5.46% of transactions are made using wallet balances.

⁷Section IV in the Online Appendix provides an example of a credit report translated to English.

⁸The bank requested six-month gaps in between survey waves, which align with their business auditing frequency.

response time below 5 minutes for most respondents.⁹ We report the English version of the survey (which was administered in Mandarin) in Appendix III. On average, survey participants spent about 4 minutes answering the survey. The average compensation was 10 CNY, which was above the 95th percentile of the hourly wage rate in China and led to a high response rate (68%).

The survey starts by describing its purpose. Besides ensuring informed participation, this step aims to eliminate potential strategic motives by respondents: without explanations, respondents might incorrectly infer that their answers would be used by the bank to design tailored offers and rebates. To avoid this possibility, the survey starts as follows:

The data will be analyzed by third-party research scholars for scientific research purposes and will not be evaluated by this bank. We will not disclose participants' identifiable information in any respect. We will not, to any extent, change the types of financial products we provide, including credit scores, credit limits, deposit rates, etc., based on participants' answers. Therefore, please answer based on your actual opinion.

Respondents then report their total income over the previous 12 months. We use this question to assess directly the accuracy of survey responses, because we can directly compare survey answers with transactions-based income data from the bank. If respondents answer carelessly or intentionally misreport their income, reported incomes would differ from observed inflows. Panel B of Figure A.1 in the Online Appendix compares self-declared income with inflows from account-level data, showing a strong linear relationship. The corresponding regression yields an R^2 of 0.74, supporting the reliability of our survey data.¹⁰

For subjective income expectations, we follow standard elicitation methods to measure both the first and second moments (Manski, 2004). First moments are elicited based on the following question:

⁹The overall sample receives nine questions. 29% of participants with transactions in investment accounts over the six months before fielding the surveys receive three additional questions about financial market returns.

¹⁰Table A.1 in the Online Appendix shows that forecast errors are not correlated with proxies for the need to obtain credit, suggesting misreporting to obtain more credit is unlikely.

What will your total income most likely be over the next 6 months?

To ensure participants have the same definition of income in mind when answering, which corresponds to the one we use to calculate total income in the account-level data, we add:

Note: income includes wages, salaries, bonuses, commissions, etc., excluding earnings from financial investments.

In addition, unemployed consumers and job switchers might be unsure about whether to report a zero income value or the expected income from a possible future job. To avoid these issues, our analysis focuses on employed consumers and we add the following clarification:

For the following three questions, we would like to ask you about your income expectations over the next six months. Please assume that you will not change your current job.

A potential concern is that respondents anticipating a job change within six months may report their current income but spend according to the income expected from their new job. In our data, only about 4% of respondents switched jobs during each forecasting period, and results are virtually unchanged when we include or exclude them. To err on the side of caution, we exclude job switchers from the main analysis.

To elicit the second moment of income expectations, we use a minimum-maximum elicitation procedure, in which respondents report the lowest and highest expected possible realizations of their future income (e.g., Guiso et al., 2002; Christelis et al., 2020). We then measure income uncertainty as the standard deviation assuming that subjective income growth follows a triangular distribution.¹¹ This simple format keeps the response burden low and facilitates consistency across survey waves and across respondents, despite variation in cognitive constraints (D’Acunto et al., 2019, 2023).¹²

Specifically, we ask the following two questions:

¹¹Denote $a = \log(E_t^\theta[\min Y_{t+1}]/Y_t)$, $b = \log(E_t^\theta[Y_{t+1}]/Y_t)$, and $c = \log(E_t^\theta[\max Y_{t+1}]/Y_t)$, then the squared income uncertainty is $(a^2 + b^2 + c^2 - ab - ac - bc)/18$.

¹²This coarse measure may induce measurement error. However, we use subjective income uncertainty only as a sorting variable to study heterogeneity in overreaction. As long as the reported uncertainty is positively correlated with true subjective uncertainty, such noise should be largely immaterial for our analysis.

What would be the lowest possible level of total income you believe you could make over the next 6 months?

What would be the highest possible level of total income you believe you could make over the next 6 months?

These two estimates allow us to compute the second moment of subjective expectations.¹³

We restrict the sample to individuals who completed at least two survey waves. This panel structure allows us to exploit within-person variation in forecast errors for income, income shocks, and economic choices over time. By adding individual fixed effects, we absorb time-invariant unobserved traits, such as cognitive abilities and financial literacy, that could confound the relationship between expectations, consumption, and debt. In addition, observing multiple cross-sections over time enables us to estimate results within each period, netting out common aggregate shocks. This feature is particularly important because our sample begins at the onset of the COVID-19 pandemic.

III.D Summary Statistics

Table I reports summary statistics for our working sample. All currency values are in thousand CNY, and all continuous variables are winsorized at the 1st and 99th percentiles to limit the influence of outliers.

The final sample includes 6,285 unique individuals and 26,449 total observations. On average, each person completed 4.2 of the nine possible survey waves.¹⁴ Demographically, the sample is balanced and representative of the working-age Chinese population. The age distribution is centered around 40 years (interquartile range: 34-46). Women account for 52% of the sample. These statistics eliminate concerns that transaction data from banks with online operations sometimes overrepresent men and younger users. Representativeness matters for external validity because prior work shows that individuals of different genders and ages vary systematically in their expectations and in how they respond to economic news (e.g., D’Acunto et al., 2021; Coibion et al., 2022).

Turning to the transaction data, the average monthly income is 13.0 thousand CNY and the average monthly spending is 7.1 thousand CNY. Both distributions are

¹³We do not additionally elicit the subjective probability that income exceeds the mode or midpoint, as both versions yield nearly identical first and second moments (Weber et al., 2022; Coibion et al., 2024).

¹⁴See Figure A.3 for the histogram of the number of completed waves.

right-skewed: the medians are 8.3 and 5.1 thousand, respectively. Respondents spend on average 2.8 thousand CNY per month on installment payments for mortgages or vehicle loans. The average stock of interest-bearing unsecured debt is 5.0 thousand CNY, though the median consumer has zero debt. Conditional on having any debt, the average balance is 12.5 thousand CNY. About 40% of respondents carry positive credit card debt, broadly in line with estimates from the U.S. (40-60%) (Gross and Souleles, 2002; Zinman, 2009; Fulford, 2015).

Average total credit limits across all banks are about 81.3 thousand CNY, suggesting that most consumers have substantial unused borrowing capacity.¹⁵ Liquidity is defined as checking plus savings account balances minus unsecured debt, divided by annual income. On average, participants hold liquid wealth equal to 1.13 times their annual income, though roughly 30% are net debtors with negative liquidity. 17% of respondents simultaneously hold positive liquid assets and credit card debt (co-holders), consistent with the well-known co-holding puzzle (Gross and Souleles, 2002; Telyukova, 2013; Gorbachev and Luengo-Prado, 2019; Gathergood and Olafsson, 2024). About a third of the sample are hand-to-mouth consumers, which we define as those with liquid assets below two months of income.

Finally, income expectations are on average close to realizations: the mean forecast error is 270 CNY. However, variation is sizable, with an interquartile range from negative 1,500 CNY to 2,500 CNY. We also verify that the surveyed consumers are broadly comparable to the bank’s overall customer base (see Table A.2), with the exception of slightly lower income and spending.

IV Income Shocks and Subjective Income Expectations

Matched transaction-level data and survey-based beliefs allow us to relate income shocks to income forecast errors at the individual level. To identify income shocks, we estimate equation (3) using the following ARMA specification for log income:

$$y_{i,t} = \rho_{j,k,a} y_{i,t-1} + \Gamma X_{i,t-1} + s_{i,t} + \omega s_{i,t-1}. \quad (4)$$

¹⁵The ratio of average credit limit to average personal income, which is around 52% in our sample, is slightly larger but close to the roughly 43% in the U.S. in 2023.

This approach is similar to the estimation of expected and unexpected tax refunds in Baugh et al. (2021), but it is specifically tied to our income process. The vector $X_{i,t}$ includes age, age squared, education, gender, the log of savings and credit limits in the previous period, city fixed effects, and industry-by-period fixed effects. The period is defined as half a year to align with the survey frequency. We estimate the persistence parameter $\rho_{j,k,a}$ at the industry-city-age-quintile level on a random 10% sample of all customers in the bank’s database between January 2020 and June 2025.

We construct log income shocks as $\hat{s}_{i,t}^y = y_{i,t} - E_{i,t-1}[y_{i,t}]$, where $E_{i,t-1}$ is the predicted income from equation (4). We define level income shocks analogously as $\hat{s}_{i,t}^Y = Y_{i,t} - E_{i,t-1}[Y_{i,t}]$, where $E_{i,t-1}[Y_{i,t}] = \exp\{y_{i,t} - \hat{s}_{i,t} + \hat{\sigma}_s^2/2\}$ is the expected value of $Y_{i,t}$ implied by equation (4), and $\hat{\sigma}_s^2$ is the estimated variance of $s_{i,t}$ at the industry-city-age-quintile level.¹⁶

Finally, we use the negative of forecast errors in period t as income surprises. Panels A and C of Figure 2 show binned scatter plots of income shocks and income surprises in logs and levels, respectively; Panels B and D absorb individual fixed effects. The shaded regions correspond to 95% confidence intervals for the fitted lines. Without residualization, the R^2 is around 0.1, indicating a positive but imperfect correlation between the two measures. Including individual fixed effects reduces measurement noise and modestly increases the R^2 , but the two measures remain distinct. Overall, consumers’ income surprises are positively related to realized income shocks, yet the correlation is far from one, leaving room for systematic errors in subjective expectations.

Serially correlated income shocks are a potential concern, which would bias the estimated effects of these shocks on subsequent outcomes. To assess this concern, we examine the autocorrelation of $\hat{s}_{i,t}$ in Online Appendix Figure A.4 and find no evidence of serial dependence. This validation supports our focus on income shocks rather than income surprises in the main analysis. When individuals overreact to past information, income surprises, defined as the negative of forecast errors, would exhibit negative autocorrelation, which would mechanically attenuate estimates of behavioral

¹⁶Results are unaffected if we assume a common $\hat{\sigma}_s^2$ across individuals. This normalization only rescales the predicted level income by a constant factor and therefore does not affect the time-series variation in income shocks that identifies our estimates.

responses—an issue that we avoid by using income shocks.

IV.A Extrapolative Subjective Income Expectations

We move on to study how income shocks relate to income forecast errors: the measures of realized income, subjective income expectations, and income shocks allow us to assess how unexpected income changes relate to the ex-post accuracy of consumers’ expectations.

IV.A.1 Income shocks, income surprises, and forecast errors

We begin by documenting systematic deviations of income forecasts from subsequent realizations. Panel A (B) of Figure 3 shows a binned scatter plot of ex-post realized income changes (growth) against ex-ante expected income changes (growth). The shaded regions give the 95% confidence intervals for the fitted lines. In both cases, the correlation is positive: consumers’ forecasts tend to move in the same direction as realized changes (see also Caplin et al. (2023)). At the same time, the cross-sectional dispersion in forecast errors is substantial. Panel C of Figure 3 displays the distribution of forecast errors, which is wide, centered at 0, and symmetric. Panel D shows this dispersion remains even after absorbing individual fixed effects. The large heterogeneity in forecast errors motivates our subsequent analysis of how forecast errors shape economic behavior.

We next examine the relationship between current-period income shocks and subsequent forecast errors by estimating the following linear specification:

$$FE_{i,t+1}^{\theta} = bs_{i,t} + X'_{i,t}\delta + \nu_{i,t+1}^{FE}, \quad (5)$$

where $FE_{i,t+1}^{\theta} \in \{FE_{i,t+1}^{\theta,y}, FE_{i,t+1}^{\theta,Y}\}$ are forecast errors in logs and levels, respectively. $FE_{i,t+1}^{\theta}$ is defined as the difference between forecasted and realized income in period $t + 1$, with forecasts made at the beginning of that period. For instance, if $t + 1$ corresponds to the first half of 2024, $FE_{i,t+1}^{\theta,y}$ equals the log of the income forecast from January to June 2024, made in early January, minus the log realized income over the same period. $s_{i,t} \in \{\hat{s}_{i,t}^y, \hat{s}_{i,t}^Y, \hat{s}_{i,t}^{\theta,y}, \hat{s}_{i,t}^{\theta,Y}\}$ denotes income shocks and income surprises in logs and levels, respectively.

$X_{i,t}$ includes individual characteristics such as credit limit, income growth, and liquidity to control for borrowing constraints and the demand for debt. The specification includes individual fixed effects, which absorb time-invariant traits, such as persistent

optimism, cognitive abilities, and financial sophistication, as well as industry \times wave fixed effects, which capture aggregate shocks common to all individuals employed within the same industry, particularly relevant during the COVID-19 period. Standard errors are clustered at the industry \times wave level to allow for correlation of unknown form across consumers whose incomes face the same industry-level economic shocks at the same point in time.

We first present plots that describe relationships in the raw data. Panels A and B of Figure 4 plot next-period forecast errors against current income shocks in levels and logs, respectively, whereas Panel C and Panel D plot next-period forecast errors against current-period income surprises. All plots show the residuals after absorbing individual fixed effects. Across specifications, both income shocks and income surprises positively predict next-period forecast errors, both in levels and in logs.¹⁷

Table II reports the results from estimating equation (5). All columns include individual fixed effects. In the even-numbered columns, we additionally control for time-varying demographics and industry \times wave fixed effects. Across all columns, the estimated coefficients \hat{b} are positive and statistically significant. A positive \hat{b} indicates that consumers' income expectations tend to move excessively in the direction of recent shocks, relative to subsequent realizations. Focusing on column (2), the inclusion of industry \times wave fixed effects absorbs economic shocks that may shift the income processes of all consumers subject to the same shocks, whereas individual fixed effects account for time-invariant personal traits that could jointly affect forecast errors and the severity of income shocks. This specification therefore identifies how the size of income shocks relates to the accuracy of the same individual's expectations over time. The estimate $\hat{b} = 0.431$ implies that, on average, a one-CNY unexpected higher income in the current period leads the average consumer to over-predict their next-period income by about 0.43 CNY.

Because the estimates are similar across specifications, we use income shocks in levels as our preferred measure of income innovations for two reasons. First, many individuals do not borrow, that is, borrowing frequently equals zero at the individual level, making log transformations for estimating elasticities problematic. Level measures

¹⁷Bellifemine et al. (2024) document a similar pattern using Italian survey data.

of income shocks and forecast errors allow us to obtain interpretable marginal-propensity estimates. Second, since income surprises are the opposite of forecast errors, the positive predictability of future forecast errors implies negative serial correlation. Using surprises would thus bias our estimates of behavioral responses, whereas level income shocks avoid this issue.¹⁸

IV.A.2 Heterogeneous responses by individual characteristics

We next ask if the relationship between income shocks and forecast errors is heterogeneous across consumers. Indeed, research shows that agents tend to overreact more when forming beliefs using noisier and more volatile signals (Ba et al., 2022; Augenblick et al., 2024). We therefore test whether consumers facing more volatile incomes overreact more strongly to unexpected income shocks. We use three proxies for income volatility: (i) age, since younger individuals typically experience more volatile earnings; (ii) the standard deviation of expected income growth; and (iii) a government-sector dummy, as government employees face stable incomes in our setting. We interact these proxies with income shocks.

Columns (1)-(3) of Table III report the results. We define *Young* as a dummy that is one for those with age below the sample median. For *More Volatile*, we first calculate the participants' subjective income uncertainty in each wave and then compute the average subjective uncertainty for each participant. *More Volatile* is one for individuals with average subjective income uncertainty above the sample median. Overreaction is systematically stronger among consumers with more volatile incomes: income shocks predict forecast errors more strongly for younger individuals and those with greater volatility in expected income growth. In contrast, no economically significant relationship between income shocks and forecast errors exists for government employees, that is, the interaction term and the baseline estimate are similar in magnitude and have opposite signs, summing up to zero. Overall, these findings show the degree of overreaction is

¹⁸In addition, Liao (2024) highlights that regressing forecast errors on forecast revisions can introduce a mechanical bias. Our approach avoids this concern by using realized income shocks from the previous period as the regressor, which has no mechanical overlap with the dependent variable.

larger when income uncertainty is higher.¹⁹

We then move on to consider other proxies inspired by recent research. For instance, Andresen et al. (2025) and Ganong et al. (2025) document substantial within-job earnings volatility, largely driven by commissions, bonuses, and other irregular income components.²⁰ Because salaried income is more predictable and varies less from one period to the next, individuals who receive a larger share of their income from salaries might extrapolate less from income shocks. We test this intuition by defining salaried workers as those in the top quartile of the regular-income share and examining whether survey participants in this group form expectations differently than others.²¹ Column (4) of Table III indeed shows salaried workers exhibit significantly weaker overreaction to income shocks.

Our sample period includes the COVID-19 pandemic, which could affect the interpretation of our baseline results. For instance, if this period happened to coincide with widespread pessimism, our estimates might capture unobserved belief heterogeneity rather than actual overreaction to unexpected income shocks. The inclusion of industry×wave fixed effects partially alleviates this concern by absorbing time-specific shocks common to all workers employed in the same industry at the same time. To further alleviate this concern, we interact income shocks with a post-2023 dummy to test whether the baseline pattern we document persists after the pandemic, where the post-2023 dummy is one if individuals participate in the survey in waves from 2023 to 2025. Column (5) shows the degree of extrapolation is only slightly smaller after 2023 and the difference is not statistically significant.

Finally, column (6) of Table III tests whether extrapolation differs by the sign of the shock. We cannot reject equality: the degree of overreaction is similar following positive

¹⁹In Online Appendix Section I.B, we show that stronger overreaction among individuals with more volatile income is consistent with greater cognitive uncertainty about the informativeness of current income.

²⁰Using administrative data from the U.S., Ganong et al. (2025) document substantial monthly variation in total earnings. In Online Appendix Figure A.5, we plot the histogram of the standard deviation of log income growth at the individual level for the participants in our sample and find a similar pattern.

²¹We classify income as regular if an individual receives at least four payments labeled as salary payments within the previous 12 months with identical amounts, capturing stable salary income, and exclude irregular components such as bonuses, commissions, and other variable transfers. The regular income share is the ratio of total regular income to total income.

and negative income shocks, both economically and statistically.

V Income Expectations, Spending, Debt, and Defaults

To assess whether income forecast errors shape economic choices, we ask if individuals adjust current consumption in line with their expectations. If a relationship exists, we would expect overly optimistic (pessimistic) expectations to raise (lower) contemporaneous spending. Using transactions-level data on actual expenditure allows us to test this prediction directly.

V.A Spending Responses

We estimate variations of the following specification:

$$\Delta C_{i,t} = b_1 \hat{s}_{i,t}^Y + b_2 FE_{i,t+1}^{\theta,Y} + X'_{i,t} \delta_c + \nu_{i,t}^C, \quad (6)$$

where $\Delta C_{i,t}$ is the change in consumption from period $t - 1$ to period t . The coefficient b_2 captures the effect of forecast errors about next period's income on current consumption, controlling for current-period income shocks $\hat{s}_{i,t}^Y$. All other variables are defined as in equation (5).²²

Table IV reports the results. All columns include industry×wave and individual fixed effects, which absorb time-varying local shocks common to all consumers employed in the same industries as well as time-invariant individual characteristics that could induce spurious trends in consumption. Column (1) relates consumption to income shocks only, yielding an estimate of the MPC: a one-CNY income increase relates to higher contemporaneous consumption by about 0.254 CNY.²³ Column (2) shows that income forecast errors are also positively correlated with current spending: a one-CNY higher forecast error relates to higher contemporaneous consumption by about 0.207 CNY.

²²Note that this specification estimates the total effects of forecast errors on current consumption. That is, we are not controlling for expectation changes further in the future. This specification is well defined under the (conventional) assumption that the state of income follows an AR(1) process and consumption is chosen according to the standard first-order Markov property with frequency of half a year, so that expectations of next-period income suffice to summarize the continuation value. This specification also assumes that the expectations elicited at the beginning of period $t + 1$ drive spending decision during period t . It is possible for consumers to have expectations for income in period $t + 1$ that change during period t and hence the expectations that drive decisions differ from those we elicit at the beginning of period $t + 1$. This mechanism would introduce measurement error in our elicited expectations and we would estimate lower bounds on the effect of forecast errors on current spending.

²³Estimates are consistent with recent evidence; Online Appendix Table A.3 reports elasticities similar to those in U.S. data (e.g., Ganong et al. (2025)).

Column (3) includes both current income shocks and forecast errors about next period's income. The MPC out of income shocks declines by 35% (from 0.254 to 0.166). At the same time, a one-CNY higher forecasted-than-realized income is associated with about 0.196 CNY higher current consumption.

The estimates in Tables II and IV decompose the MPC out of income shocks into a *belief channel* and a *residual (liquidity) channel*. The overall MPC is defined as

$$\frac{dC_{i,t}}{d\hat{s}_{i,t}} = \frac{\partial C_{i,t}}{\partial \hat{s}_{i,t}} + \frac{\partial C_{i,t}}{\partial FE_{i,t+1}^\theta} \times \frac{\partial FE_{i,t+1}^\theta}{\partial \hat{s}_{i,t}},$$

where the first term captures the direct liquidity-driven response and the second term captures the indirect response through expectation changes. From Table II column (2), we estimate $\partial FE_{i,t+1}^\theta / \partial \hat{s}_{i,t} = 0.431$; from Table IV column (3), $\partial C_{i,t} / \partial \hat{s}_{i,t} = 0.166$ and $\partial C_{i,t} / \partial FE_{i,t+1}^\theta = 0.196$. Combining these estimates yields a back-of-the-envelope MPC of $dC_{i,t} / d\hat{s}_{i,t} = 0.250$, which closely matches the total effect of 0.254 reported in Table IV column (1). Meanwhile, controlling for forecast errors, \hat{b}_1 reflects the conventional MPC to income shocks, a key measure of consumption insurance (Blundell et al., 2008).

Columns (4)-(7) examine heterogeneity by liquidity constraints. We proxy constraints by using (i) the average credit line utilization rate and (ii) the ratio of liquid savings to income, both measured over the six months prior to period t . *High Util_t* equals one for utilization above the sample median in period t , and *High Liq_t* equals one for liquid savings-to-income above the median. The effects of income shocks and forecast errors on consumption are larger for consumers with tighter constraints (high utilization or low liquidity), which is consistent with higher MPCs among more constrained households.

V.B Debt and Default

Because forecast errors raise current consumption while future realized income increases by less than expected, positive forecast errors lead consumption to deviate from its optimal path absent forecast errors. The combination of higher current spending and lower-than-expected subsequent income implies that debt should rise, on average, when expectations are overly optimistic. We therefore test whether extrapolative income expectations predict higher subsequent borrowing. To do so, we estimate a version of equation (6), whereby the dependent variable is the change in total unsecured debt, $\Delta B_{i,t+1}$. For example, if

$t + 1$ corresponds to the first half of 2024, $\Delta B_{i,t+1}$ is the difference between consumer i 's total outstanding interest-bearing unsecured debt at the end of June 2024 and that at the end of December 2023.

Column (1) in Table V shows that time- t income shocks weakly but insignificantly predict next-period debt. Columns (2) and (3) instead show that the larger is the difference between expected and realized income, the greater is the subsequent increase in unsecured borrowing. In column (3), which includes the full set of controls and fixed effects, a one-CNY higher income forecast error is associated with a 6.2-cent increase in unsecured debt. This magnitude is economically meaningful, given that unconstrained consumers do not need to borrow to finance additional spending, and therefore exhibit virtually no response.

A remaining concern is that forecast errors could be negatively correlated with current income shocks, which would mechanically reduce net assets. To isolate the effect of forecast errors, we include income shocks in the same period as forecast errors as an additional control in column (4). In this specification, we estimate the impact of forecast errors on debt accumulation conditional on contemporaneous income shocks. We find that a one-CNY positive income shock lowers debt by about 3 cents, while a one-CNY higher forecast error for income raises debt by roughly 6.5 cents, which is indistinguishable from the baseline estimate.

The results on consumption and debt indicate that positive income shocks lead to an MPC smaller than one. And yet, debt on average does not decline after these shocks. While this pattern may seem puzzling at first, it arises because debt is bounded below at zero; consumers with high saving rates often have little or no debt in the first place. In contrast, increased spending driven by forecast errors typically results in higher debt among consumers with limited savings, which suggests that, on average, debt can increase but barely decline.

The third outcome we examine is default. When realized income is lower than expected, consumers who cannot draw on savings may fail to repay their debt in full. We estimate equation (5) using a dummy variable that equals one if a consumer defaults

over the next period as the dependent variable.²⁴ By construction, we observe outcomes only up to period $t + 1$, while some defaults may occur later. Our estimates therefore likely represent a lower bound of the true relationship between income forecast errors and default. For ease of interpretation, we multiply the default indicator by 100 and standardize income shocks and forecast errors.

Columns (5)-(8) of Table V show income shocks do not significantly predict future default. Meanwhile, larger gaps between income forecasts and realized incomes are associated with a higher likelihood of default. A one-standard-deviation increase in income forecast errors raises the default rate by about 0.85 percentage points (column (7)). This effect is economically significant: given an average default rate of around 3.6%, a one-standard-deviation higher income forecast error implies roughly a 24% increase in the likelihood of default. These results align with aggregate evidence linking optimistic credit-market sentiment to financial fragility (López-Salido et al., 2017). Similar to the debt results, column (8) shows controlling for contemporaneous income shocks barely affects the estimated relationship between forecast errors and default.²⁵

In Table A.5 in the Online Appendix, we show that income shocks in period t reduce contemporaneous debt and default, consistent with positive wealth effects. Controlling for income shocks, forecast errors modestly increase current debt but do not affect default. This pattern is also consistent with optimistic expectations raising spending, and therefore weakly borrowing on impact. However, the rise in debt and default mainly occurs in $t + 1$, when optimistic forecasts do not meet realizations.

V.C Addressing Endogeneity: 2SLS Estimation

Spending and borrowing responses to forecast errors may be endogenous if unobserved traits or common shocks jointly affect income expectations and behavior. We mitigate this concern by including individual fixed effects, which absorb time-invariant heterogeneity, and industry \times wave fixed effects, which capture aggregate shocks common to the same

²⁴For example, the default indicator in the first half of 2020 reflects 90-day delinquencies recorded between October 2020 and March 2021. We focus on delinquency given that formal default does not exist in China.

²⁵Similar to spending, Table A.4 in the Online Appendix shows debt and default also react more strongly for those with higher utilization rates.

industry.

To further improve identification, we employ an instrumental-variables (IV) approach that allows us to focus on quasi-exogenous variation in income shocks and forecast errors. Our IV design uses two instruments and is closely related to the location-by-treatment interaction strategies commonly used to isolate indirect treatment channels when shocks affect outcomes through multiple channels (e.g., Kling et al. (2007); Abdulkadiroğlu et al. (2014); Kline and Walters (2016)).

First, to proxy for exogenous income shocks, we divide the 10% bank panel for which we estimate equation (4) into 37 consumer types and compute the mean income shock within each type-period cell, excluding (i) all individuals in our sample and (ii) all individuals in the same city as the survey participants.²⁶ This type-level average shock, similar in spirit to co-worker income shocks (Ganong et al., 2025), captures common innovations faced by similar consumers while minimizing local spillovers. In addition, we expect this instrument to only capture within-type systematic innovations to income, and hence it should be uncorrelated with individuals' private information. As a result, the 2SLS design also alleviates the concern that part of the income shocks we extract might capture participants' private information.

Second, following Yin (2025), we interact the type-level income shock with an indicator for provinces with above-median average income volatility. The validity of the location-by-treatment strategy requires that the degree of expectation updating varies systematically across locations with different levels of income uncertainty, and this variation is different from that of the type-specific income shocks. This assumption is plausible because the interaction exploits our finding that consumers facing higher income uncertainty update their beliefs more strongly in response to a given income shock. As a result, the interaction term provides additional, plausibly exogenous variation in belief updating.

Table VI reports the results. The first-stage estimates in columns (1) and (2) confirm the relevance of the instruments. The within-type average income shocks are strongly positively correlated with individual income shocks. At the same time, the

²⁶We do so based on K-means clustering over age, gender, education, industry, location, income, spending, debt, and credit scores, with the number of clusters chosen by ten-fold cross-validation.

instrument is negatively correlated with forecast errors in low-volatility environments. A potential explanation is that when income risk is limited, consumers devote less attention to systematic income innovations and thus react less, consistent with rational inattention. In contrast, the positive interaction term in high-volatility cases indicates stronger belief updating, as consumers react more to the same income news when uncertainty is higher.

The second-stage estimates produce magnitudes and signs similar to the baseline OLS results. Columns (3)-(5) show that each one-CNY higher income shock raises current spending by 0.126 CNY but has no significant effect on subsequent debt accumulation or default. Controlling for income shocks, a one-CNY higher income forecast error increases contemporaneous spending by 0.216 CNY. Moreover, subsequent debt accumulation increases by 0.088 CNY, and defaults rise by 1.276 percentage points for each one-standard-deviation higher forecast error at $t + 1$. Overall, these findings corroborate the validity and interpretation of our empirical results.

V.D Dynamic Effects

Our analysis so far has focused on six-month horizons. In the longer run, consumers observe income realizations and hence learn that their expectations were incorrect. Over time, we might expect spending and borrowing to revert as biased expectations are corrected. To examine whether and how quickly consumption adjusts to income forecast errors, we estimate the relationship between cumulative consumption and forecast errors for income over different horizons:

$$C_{i,\tau+k} - C_{i,\tau-1} = b^{(k)} FE_{i,t+1}^{\theta,Y} + X'_{i,t} \delta^{(k)} + \nu_{i,\tau+k}^C. \quad (7)$$

The specification combines two time frequencies: forecast errors for income are observed at a half-year frequency, indexed by t , while for the inspection of adjustment dynamics, we use quarterly consumption, indexed by τ . Let τ denote the first quarter of half-year t , so that period t corresponds to quarters τ and $\tau + 1$, and period $t + 1$ to quarters $\tau + 2$ and $\tau + 3$, etc. The forecast error $FE_{i,t+1}^{\theta,Y}$ therefore refers to income realizations in quarters $\tau + 2$ and $\tau + 3$.

The left-hand side measures cumulative consumption from the quarter preceding the survey time ($\tau - 1$) up to quarter $\tau + k$. We express both consumption and forecast

errors as average monthly level to keep units comparable. We estimate equation (7) for horizons $k = -4, \dots, 8$, excluding $k = -1$, which serves as the benchmark. The resulting coefficients $b^{(k)}$ measure the cumulative changes of consumption from quarter $\tau - 1$ to quarter $\tau + k$ in response to each CNY higher forecast error in period $t + 1$.

Panel A of Figure 5 plots the results. The red solid line represents positive forecast errors and the blue dashed line negative errors, with shaded areas denoting two standard-error bands. Consistent with overreaction, consumers increase spending in the two quarters preceding and during the period when forecast errors are measured (quarters 2-3), followed by a gradual reversion that nearly eliminates the consumption response within five quarters.

Panel B shows similar results for cumulative borrowing. The dynamic pattern mirrors that of consumption: borrowing initially rises with optimistic expectations and later reverts as observing realized income corrects beliefs. Overall, these findings are consistent with consumers' overreaction to income news being followed by subsequent belief adjustment.²⁷

V.E Robustness Checks

In this section, we conduct several robustness tests. We first verify that our results are qualitatively similar if we add no controls to the right-hand-side of our baseline specification or absorb different sets of fixed effects (see Table A.6 in the Online Appendix). We also show that the heterogeneity results are similar when using income surprises (see Table A.7 in the Online Appendix). In terms of statistical inference, our main analysis clusters at industry \times wave level because many consumers are only observed a couple of times based on their survey responses, so one might worry that clustering at the individual level might not be conservative enough due to limited within-cluster variation. Nevertheless, Panel A of Table A.8 in the Online Appendix reports the results when clustering standard errors at the individual level, and yields similar statistical inference as in our baseline specification.

In terms of potential confounding factors, as was discussed above, the fact that

²⁷While this pattern aligns with the overreaction mechanism in our model, it is descriptive and does not by itself rule out complementary forces, such as liquidity constraints or transitory adjustment frictions, that may also contribute to the observed dynamics.

our sample spans the COVID-19 pandemic period might raise concerns. In the baseline analysis, we included industry \times wave fixed effects to absorb aggregate shocks affecting consumers within the same industry and period. To further verify that our results are similar at all times, irrespective of the pandemic shock, we also report results when excluding observations between 2020 and 2022 from the sample, and re-estimating income shocks using a 10% bank sample from January 2023 to June 2025. Panel B of Online Appendix Table A.8 shows that the results are nearly identical to our baseline when implementing this restriction.

Moreover, one might be concerned that some forecast errors may reflect inattention or reporting mistakes when respondents report current or past income. To assess this concern, we re-estimate the main regressions keeping only individuals whose reporting errors for past income are below the median within each survey wave. Panel C of Online Appendix Table A.8 shows that the results remain unchanged.

An additional concern is that income shocks estimated from equation (4) might capture innovations to the econometrician's information set rather than true surprises to the individual. Part of the residual could thus reflect private information known to respondents. Our IV strategy, as we discussed above, helps to address this concern. As a more direct test, we re-estimate equation (4) including the respondent's forecast for lowest, highest, and most likely income for the next period. The results, shown in Panel D of Online Appendix Table A.8, are again very similar to our baseline results, suggesting that private information is unlikely to be a relevant confounding factor.

We then move on to assess robustness of our definition of main bank users, which drives the initial sample selection. Following earlier research, our baseline sample only includes individuals with at least 20 transactions per month during the sample period. This filter aims to limit the sample to active account users. At the same time, it might raise the concern that our sample is tilted towards consumers with larger spending habits. Panel E of Online Appendix Table A.8 shows that our results are nearly identical when we change the filter to at least 10 transactions per month. In addition, in Panel F, we focus on participants who indicate that they only use one bank for their daily transactions. Again, the results remain virtually unchanged.

Note that we estimate equation (4) on a separate sample that excludes survey participants to avoid mechanical endogeneity in the construction of shocks: using the same individuals to estimate (4) would yield in-sample residuals that absorb individual-specific income dynamics and correlate with subsequent forecast errors and spending. Because we cannot include individual fixed effects in this out-of-sample estimation, we incorporate a rich set of demographic controls to absorb cross-sectional heterogeneity. As a robustness check, we re-estimate equation (4) in-sample using the same set of controls, and Panel G of Online Appendix Table A.8 shows that the results remain virtually unchanged.

The level of decision-making also raises a potential concern. Spending decisions might be made at the household level, while we observe individual-level accounts and individuals' forecasts for their own income, so our baseline results might be contaminated by intra-household transfers and decision making at the household level. To tackle this concern directly, in Panel H of Online Appendix Table A.8, we only consider single consumers, who account for 31% of our sample. We find similar results.

Moreover, our spending data is based on the accounts at the bank with which we cooperate. To assess concerns of systematic mismeasurement not eliminated by our screening filters to select primary account users, we study the relationship between forecast errors and the changes in net transfers between our bank and external bank accounts (see Online Appendix Table A.9). We find no relationship between forecast errors and cash withdrawals or net transfers to external accounts, which dismisses the concern that our bank-level data miss relevant transfers to potentially unobserved external accounts. Meanwhile, as a placebo test, Online Appendix Table A.9 shows that installment payments on mortgages and vehicle loans, which are rigid in the short run and typically set at the beginning of the debt contract, do not respond significantly to forecast errors at the frequency we observe in our setting.

Finally, our survey-based panel has a relatively short time dimension. Although results are similar when exploiting either within- or across-consumer variation, income innovations may not fully average out within individuals in short panels (Chamberlain, 1982; Souleles, 2004). To address this concern, we leverage the longer transaction-level panel available from the bank, which spans on average 7.9 years per consumer and includes

at least ten periods per individual. Online Appendix Table A.10 compares estimates from the extended panel (even columns) to those from the baseline survey sample (odd columns). The conditional relationships between income shocks and spending, borrowing, and default remain virtually unchanged.

V.F Transitory vs Permanent Shocks

As a final exercise, we decompose the income shocks consumers face into permanent and transitory components and test whether consumers update their subjective expectations based only on persistent shocks or on both types of shocks. The decomposition procedure follows Pistaferri (2001), and we describe in detail how we adapt it to non-rational expectations in Online Appendix Section V. Online Appendix Table A.11 shows that transitory shocks account for roughly two-thirds of the variation in unexpected income growth in our data.

Building on this decomposition, we regress forecast errors for income on both current transitory and permanent shocks to evaluate their respective contributions. The results, reported in Online Appendix Table A.12, indicate that despite transitory shocks explaining the majority of income variation, both types of shocks significantly predict forecast errors. Meanwhile, we find a 0.22 consumption elasticity to transitory shocks and a 0.44 consumption elasticity to permanent shocks. Overall, the findings are consistent with the framework in Section II, suggesting that consumers have limited ability to distinguish between transitory and persistent income shocks and therefore react to both.

In sum, the empirical analysis shows that consumers excessively extrapolate income shocks, both transitory and persistent, into their expectations. Following positive income shocks, they form subjective expectations of future income that are excessively optimistic relative to realizations. This tendency is stronger among younger consumers, those facing higher income volatility, and for those not working for the government. Moreover, forecast errors for income, defined as the gap between ex-ante subjective expectations and ex-post realizations, significantly increase current spending, leading to higher subsequent debt and default.

VI Structural Estimation With Overreaction

In the last part of the paper, we turn to a quantitative framework to assess the extent to which the micro-level, belief-driven demand channel we document empirically might help explain the aggregate patterns of household debt observed in the data.

To do so, we embed extrapolative income expectations and defaultable debt into a standard partial-equilibrium consumption model with incomplete markets. The model allows us to quantify how unexpected income shocks, forecast errors, and economic choices behave in the presence of extrapolative expectations.

The model has discrete time and infinite horizon. A unit mass of consumers are subject to idiosyncratic income risk. The income process follows the one in Section II. Given an infinite horizon, we follow the common assumption that a sufficient number of periods have passed such that consumers are in a learning steady state, that is, consumers' Kalman gain is constant each period, in which case,

$$\kappa = \frac{\sigma_z^2 + \sigma_\eta^2}{\sigma_z^2 + \sigma_\eta^2 + \sigma_\epsilon^2}, \quad \sigma_z^2 = \frac{(1 - \kappa)\sigma_\eta^2}{\kappa}. \quad (8)$$

VI.A Preferences and Optimality Conditions

In this subsection, we introduce consumer preferences and discuss the optimality conditions.

VI.A.1 Preferences

Household preferences follow the standard framework used in models of consumer credit and default (e.g., Chatterjee et al., 2007; Livshits et al., 2007). Consumers maximize expected lifetime utility with per-period utility given by

$$\frac{C^{1-\gamma}}{1-\gamma} - \chi d,$$

and discount factor β . Here, γ denotes the coefficient of relative risk aversion, and $d = 1$ if the consumer defaults at the end of period t .

Default entails a fixed non-pecuniary utility cost (“stigma”) $\chi > 0$. In addition, consumers receive i.i.d. additive utility shocks $\{\xi^D, \xi^N\}$ attached to the default and repayment options, respectively. These shocks are drawn from a Type-I extreme value distribution with scale parameter λ and capture unobserved events, such as marital

disruptions or medical expenses, that can trigger default. The inclusion of these shocks ensures a positive probability of default across the entire borrowing range. As noted by Dempsey and Ionescu (2023), smoothing the default decision facilitates model computation.

The budget constraint in period t is

$$A' = \begin{cases} (1+r)(A-C) + Y' & \text{if } d = 0 \\ (1-\nu)Y' & \text{if } d = 1 \end{cases} \quad (9)$$

$$A \geq -L,$$

where A is the total amount of available resources. L is the credit limit, and $\nu \in [0, 1]$ is the marginal rate of garnishment. Equation (9) states that when consumers do not default, their wealth in the next period is the sum of their income and gross savings. When consumers default, their savings become zero; at the same time, they need to pay a garnishment cost equal to ν times their income in the following period. For simplicity, we assume consumers' borrowing capacity does not change upon default, which allows us to discard one additional state variable.²⁸ The interest rates differ for saving and borrowing and take the values

$$r = \begin{cases} r_b & \text{if } A < 0 \\ r_s & \text{if } A \geq 0. \end{cases}$$

VI.A.2 Optimality Conditions

Consumers' problem is characterized by a set of two state variables $\Theta = (a, \hat{z})$. Given the state Θ , the consumer's value function is

$$V(\Theta) = \max \{V_D(\Theta), V_N(\Theta)\}. \quad (10)$$

The continuation value from defaulting and not defaulting are, respectively,

$$V_D(\Theta) = \max_C \frac{C^{1-\gamma}}{1-\gamma} - \chi + \beta \mathbb{E}_\theta[V(\Theta')] + \xi^D. \quad (11)$$

$$V_N(\Theta) = \max_C \frac{C^{1-\gamma}}{1-\gamma} + \beta \mathbb{E}_\theta[V(\Theta')] + \xi^N. \quad (12)$$

²⁸Some studies assume defaults go hand in hand with a temporary inability to borrow, that is, $L = 0$ (Chatterjee et al., 2007; Livshits et al., 2007; Dempsey and Ionescu, 2023), but Livshits et al. (2007) show the costs of default from changing borrowing capacities are quantitatively small.

The subscript θ indicates that expectations are taken under the agents' perceived law of motion, which is parameterized by the extrapolation parameter θ . Given the utility shocks follow a type I extreme value distribution, the probability of default is

$$pd(\Theta) = [1 + \exp\{[V_N(\Theta) - V_D(\Theta)]/\lambda\}]^{-1}. \quad (13)$$

VI.B Reproducing the Empirical Results

We next evaluate the model fit on targeted and untargeted moments and how well the model replicates the empirical results.

VI.B.1 Goodness of Fit

We estimate the parameters in two steps.²⁹ The first step does not rely on solving the model. In this step, we pin down the income and belief formation process and set interest rates using transaction data and survey expectations. We also fix the preference parameters β and γ to conventional values. In the second step, we calibrate the two default-related parameters, χ and λ , using the simulated method-of-moments (SMM).

We discipline the degree of overreaction in beliefs, θ , by matching the regression coefficient of forecast errors for income at $t + 1$ on income shocks at t . Specifically, we simulate 100,000 periods from equation (1), form expectations using equation (2), and drop the first 100 periods as burn-in. This procedure yields $\theta = 2.54$ (Table VII, Panel A), implying that the subjective Kalman gain is 2.54 times larger than the Bayesian counterpart.³⁰

Table VII shows the model closely matches both targeted and non-targeted moments. Conditional on the belief and income process parameters, we match the average wealth-to-half-year consumption ratio and the default rate to characterize default behavior. Panels B and C indicate a good fit of these targeted moments: the average wealth-to-half-year consumption ratio is 0.957 and the default rate is 3.60%, implying $\chi = 2.30$ and $\lambda = 0.25$.

Panel D compares a set of non-targeted moments and shows a similarly good fit. The average (median) wealth-to-six-month income ratio is 1.211 (0.912) in the data and 1.131

²⁹Online appendix Section VI details the solution, estimation, and simulation procedures.

³⁰Our estimate of θ lies within the range reported in previous literature in related settings, which is 0.8–2.75 (Bordalo et al., 2019; Chodorow-Reich et al., 2023). However, magnitudes depend on sample frequency and the equilibrium signal-to-noise ratio and are therefore not directly comparable across studies.

(1.000) in the model. Because liquidity shapes MPCs empirically, the stationary savings distribution must align with its empirical counterpart. The model also reproduces key features of the wealth distribution: in the data, 28.97% of total liquid assets are held by the top 5% of savers versus 26.71% in the model, and 19.33% (16.10% in the model) of consumers have negative net liquid wealth. Finally, the average debt-income ratio is 0.230 in the data and 0.235 in the model. Overall, the model matches several distributional features of the empirical data beyond the first moment.

VI.B.2 Extrapolation of Income Shocks

We use the model to examine how extrapolative income expectations influence economic choices. We first compare the relationship between income shocks and forecast errors under extrapolative expectations and Bayesian learning.

Column (1) of Table VIII reports the results. Panel B reports results for a model in which agents have extrapolative expectations (overreaction) and Panel C represents the Bayesian benchmark, with Panel A showing the empirical counterpart. By construction, the parameter θ is calibrated to match the empirical relationship between forecast errors and income shocks: when $\theta = 2.54$, each one-CNY income shock leads to a 0.431-CNY forecast error, perfectly replicating the data. In contrast, when $\theta = 0$ (Bayesian learning), the relationship is small and negative. Each one-CNY income shock produces a -0.038-CNY forecast error, confirming that a positive θ is necessary to capture the empirical degree of overreaction.³¹

VI.B.3 Consumption, Borrowing, and Defaults

We next assess whether the model can reproduce the empirical relationship between forecast errors and economic choices. We simulate 20,000 individuals for 1,000 periods after 100 periods of burn-in for the model to reach a stationary distribution. We then drop simulated observations with savings-to-income ratios above 10, which is the maximum we observe in the sample.

³¹In Section VII in the Online Appendix, we replicate the core forecast-error regression assuming expectation formation follows 1. diagnostic Kalman filtering, 2. full-information rational expectations, 3. over-persistence bias, and 4. learning about long-run growth. We find that overreaction best matches the empirical results. Specifically, while diagnostic Kalman filtering is also able to match the link between income shocks and forecast errors on impact, it fails to match the temporal patterns due to the kernel of truth assumption. See the appendix for more details.

Using the simulated data, we regress changes in total consumption at t , total debt at the beginning of $t + 1$, and default at the end of $t + 1$ on forecast errors for income for period $t + 1$. Debt at the beginning of $t + 1$ is defined as $-A_t$ when $A_t < 0$.

Columns (2)-(4) of Table VIII report the results, with 2SLS counterparts in Panel A. The model with extrapolative expectations (Panel B) replicates the empirical relationships closely. A one-CNY higher forecast error for income at $t+1$ increases current consumption by 0.174 CNY in the model (0.216 CNY in the data) and debt by 0.068 CNY (0.088 CNY in the data). A one-standard-deviation higher forecast error raises the default rate by 1.32 percentage points in the model versus 1.28 percentage points in the data.

By contrast, without extrapolative expectations ($\theta = 0$), the corresponding coefficients shrink substantially: forecast errors do not significantly affect current consumption and next-period debt, and increase next-period default by 53 basis points, which is around 60% smaller than the counterparts under extrapolation.

Overall, incorporating extrapolative expectations can replicate the empirical relationships between forecast errors for income and consumption, debt, and default, whereas the model with Bayesian learning cannot.

VI.C Welfare Analysis

Extrapolative income expectations can impose substantial costs on households. By amplifying perceived income fluctuations, overreaction increases borrowing following positive shocks and raises the likelihood of default when income is lower than expected. Higher debt burdens and more frequent financial distress reduce wealth accumulation, while distorted expectations impair intertemporal consumption smoothing. We therefore quantify the welfare implications of overreaction relative to the Bayesian benchmark.

For the welfare analysis, after discarding the first 100 periods as burn-in, we then simulate 20,000 agents over 100 periods, which roughly corresponds to a 50-year lifecycle. We first compare the stationary equilibrium distributions of consumption and wealth. The results are in Panel A of Table IX. Moments are computed from the stationary distribution. Default probability is the unconditional default rate across agents and time. C/Y denotes the consumption-to-income ratio and A/Y the wealth-to-income ratio. Relative to the Bayesian benchmark, overreaction increases the default rate

(3.60% vs. 2.91%) and the share of debtors (16.10% vs. 14.92%), while reducing both consumption-to-income and wealth-to-income ratios. Meanwhile, with overreaction, log consumption volatility is substantially larger due to consumers' excessive extrapolation. These patterns suggest that overreaction leads to higher default rates and lower wealth and consumption in equilibrium.

We then quantify welfare costs using consumption-equivalent variations (Panel B). For each individual i , we get the welfare costs as the λ_i that solves

$$\sum_t \beta^t u((1 + \lambda_i)C_{i,t}^O, d_{i,t}^O) = \sum_t \beta^t u(C_{i,t}^B, d_{i,t}^B),$$

where $u(c, d) = \frac{c^{1-\gamma}}{1-\gamma} - \chi d$, and superscripts O and B denote overreaction and Bayesian regimes, respectively.

We report welfare losses in Panel B of Table IX. Welfare losses are economically large and heterogeneous. The median consumption-equivalent loss is 8.37%, and the mean loss is 12.27%, with losses exceeding 24% at the 90th percentile.

Note that welfare losses operate through three channels: (i) more frequent defaults, (ii) lower average consumption in equilibrium, and (iii) inefficient consumption smoothing. To decompose these mechanisms, we first shut down default penalties by setting $\chi = 0$. The median welfare loss declines from 8.37% to 5.76%, implying that higher default risk accounts for roughly one-third of the total cost. Next, in the "Fixed C Level" exercise, we rescale each individual's overreaction consumption path so that average consumption matches the Bayesian benchmark. This step removes differences in mean consumption and isolates smoothing distortions. The median welfare loss further declines to 4.30%.

Note that our welfare cost in the inability-to-smooth case is larger than that in Ganong et al. (2025) because the consumption volatility implied by our model is substantially higher, reflecting households' failure to fully internalize and smooth income risk due to extrapolative expectations of income.

Because both higher default and lower average consumption originate from lower equilibrium wealth under overreaction, we interpret them jointly as a wealth channel. The remaining loss reflects inefficient intertemporal smoothing. Quantitatively, the wealth channel and the smoothing channel each account for roughly half of the total welfare cost.

VI.D Aggregate Household Debt Cycles

A recent literature argues that financial crises often arise after sequences of favorable news that fuel optimistic beliefs, credit expansion, and leverage accumulation, so that subsequent negative shocks are endogenously amplified rather than absorbed smoothly (Baron and Xiong, 2017; Bordalo et al., 2018; Gennaioli and Shleifer, 2018; Mian et al., 2020). In this view, adverse shocks matter not primarily because they are large but because they arrive when expectations and balance sheets have already been shaped by preceding good times.

Motivated by this literature, we use our model to study boom-bust dynamics with a controlled simulation experiment. Starting from the stochastic steady state, we simulate 20,000 consumers under extrapolative expectations ($\theta = 2.54$) for 1,000 half-year periods. We then compare this baseline to an otherwise identical simulation in which consumers experience three consecutive positive transitory income shocks totaling 1.5 standard deviations (periods 901-903), followed by a 0.5 standard-deviation negative transitory income shock. The reported dynamics are computed as differences relative to average income in the no-shock baseline except for income expectation and default, which are percentage difference from the no-shock baseline levels. We repeat the same experiment under Bayesian learning ($\theta = 0$) to isolate the role of extrapolative expectations in shaping the economy's response to the shock dynamics.

Figure 6 plots the results. The top panels display the transitory income shocks and the percentage deviations of average income expectations. The bottom panels display changes in consumption-to-income ratios and debt-to-income ratios, and percent changes of default. Red solid lines denote the model in which consumers have extrapolative expectations ($\theta = 2.54$) and blue dashed lines represent the Bayesian learning benchmark ($\theta = 0$).³²

The figure reveals a strong amplification effect of extrapolative expectations on income beliefs. Panel B shows that, following positive shocks, consumers with overreaction (*O*) raise their expected income more than twice as much as Bayesian learners (*B*). Higher

³²In Figure A.6 in the Online Appendix, we show that the results are nearly identical if we re-calibrate the model while setting $\theta = 0$ before simulating the Bayesian case.

expected income translates into larger consumption responses. In Panels B and C, O 's income expectation rises by about 20% and consumption by about 15% relative to income. Because the model only has one asset, an MPC below one implies an initial decline in debt, consistent with Agarwal et al. (2007) and Coibion et al. (2020). From Panel E, due to a lower debt level, both O and B see lower default rates.

As positive shocks persist, both income expectations and consumption of B rise smoothly, and deleveraging continues. In contrast, with extrapolative expectations, income expectations are overly optimistic ex ante and only partially realized ex post. As a result, asset accumulation is weaker for some households, and those with limited savings begin to accumulate debt. Because households with zero debt cannot deleverage further while optimistic households borrow more, average debt gradually increases. Consequently, contrary to B , default rates for O do not continue to decline.

When the sequence of positive shocks ends with a negative shock (period 0), income expectations drop for both types of agents. For O , income expectations drop much more because (1) the same income shock implies a larger disappointment due to ex-ante elevated prior expectations, and (2) they react more strongly to the same-sized negative shock due to a larger Kalman gain. As a result, consumption falls more, and debt spikes as O attempt to smooth spending given lower expected income. The combination of lower current assets and lower future income expectations leads to a pronounced surge in default among extrapolative consumers. As expectations converge back to steady state, both debt and default rates for O gradually return to those of B .

These dynamics align with recent aggregate evidence. López-Salido et al. (2017) show that elevated credit-market sentiment predicts higher subsequent credit growth, mirrored in Panel E, where debt growth turns positive for O after the initial good news, but remains negative for B . Likewise, the post-expansion rise in default in Panel E echoes findings by Greenwood et al. (2019) and Maxted (2023) that financial fragility peaks at the end of economic booms.

VI.E Unsecured Debt around the Global Financial Crisis

Economists' renewed interest in the drivers of household debt cycles has been shaped by the experience of the Global Financial Crisis (GFC). Our final exercise therefore assesses

whether the model can reproduce the patterns of unsecured borrowing and default risk observed in the United States during the GFC. Since the parameters are estimated using Chinese data, this exercise is illustrative.

Figure 7 presents the results.³³ In each panel, the red solid and blue dotted lines denote simulations with $\theta = 2.54$ and $\theta = 0$, respectively, while black dashed lines represent the data. The model frequency is semiannual. In the left panel, $\overline{B}/\overline{Y}$ denotes the ratio of consumer, credit card, and revolving debt to income; the right panel reports default probabilities measured as delinquency rates on consumer loans.³⁴

Two findings stand out. First, aggregate debt in equilibrium (pre-2005) is higher under extrapolative expectations, because overreaction widens the cross-sectional dispersion of income expectations: more consumers become optimistic and borrow more, while pessimists cannot deleverage below zero debt. As a result, forecast errors do not offset in the aggregate, producing higher debt in steady-state.

Second, with extrapolative expectations, the model generates a pronounced boom-bust cycle in the unsecured-debt-to-income ratio that closely tracks the U.S. credit card debt-to-GDP pattern. Without extrapolation, these cycles largely disappear. That is, the 2009 peak never rises significantly above its 2004 level.

The right panel shows the same income shocks can also reproduce the surge in default rates at the end of the expansion, followed by a gradual decline. In contrast, under Bayesian learning ($\theta = 0$), default rates remain nearly flat from 2004 to 2012.

Overall, the model produces an aggregate boom-bust debt cycle akin to the one observed during the Global Financial Crisis only when income expectations are extrapolative, but not under Bayesian learning.

³³The corresponding income shocks, reverse-engineered to match the observed debt dynamics, are shown in Figure A.7 of the Online Appendix. While it is difficult to find the exact series of shocks that can generate Figure 7, Figure A.7 shows that the picked shocks are highly correlated with the unexpected nonfarm payroll employment growth.

³⁴Because we are studying average debt to personal income in the model, we multiply GDP by the labor share to get total income. The resulting series trends down after 2000, so we detrend it over our sample period. For the delinquency rate, the model-implied default probability is based on all consumers, but the delinquency rate from the data is conditional on debt-holders. To make the series comparable, we adjust the delinquency rate by dividing it by the proportion of debtors in the U.S. Hence, we get a debt delinquency rate for all consumers rather than only for debtors.

VII Conclusions

We combine repeated survey measures of subjective income expectations within individuals with transaction-level and credit-registry data to study the joint dynamics of income forecast errors, spending, borrowing, and default. Consumers form extrapolative income expectations by overreacting to income shocks, and this overreaction drives their subsequent consumption and debt choices. Embedding such belief dynamics into a standard consumption-saving model reproduces the patterns in the data when consumers form extrapolative expectations. The model can generate Minsky-Kindleberger-style credit cycles such that a series of positive shocks can build up financial fragility and then a negative shock results in a bust. It can also generate the unsecured-debt and default dynamics observed in the United States during the Global Financial Crisis. A version of our model with Bayesian learning cannot replicate any such patterns.

Future research should further explore both the reduced-form and structural implications of these findings. On the reduced-form side, understanding the heterogeneity of income extrapolation across different countries and over different periods, and isolating the specific expectation-formation process that induces overreaction are important. Structurally, while agents in our model share a common degree of extrapolation, this parameter is likely heterogeneous across different types of agents.³⁵ Extending the analysis to macroeconomic models that allow for heterogeneous extrapolation and assessing its aggregate consequences represent promising directions for future research. In addition, our specification adopts a parsimonious learning framework that is consistent with our empirical findings and is standard in consumption-saving models. Future work, including laboratory or field experiments, could shed further light on the underlying mechanisms of learning and expectation formation.

References

- Abdulkadiroğlu, A., J. Angrist, and P. Pathak (2014). The elite illusion: Achievement effects at boston and new york exam schools. *Econometrica* 82(1), 137–196.
- Afrouzi, H., S. Y. Kwon, A. Landier, Y. Ma, and D. Thesmar (2023). Overreaction

³⁵For subjective macroeconomic expectations, heterogeneity has been linked to cognitive ability (D’Acunto et al., 2019, 2021), socioeconomic status (Kuhnen and Miu, 2017), and local experiences (Kuchler and Zafar, 2019), among others.

- in expectations: Evidence and theory. *The Quarterly Journal of Economics* 138(3), 1713–1764.
- Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2017, 07). Do Banks Pass through Credit Expansions to Consumers Who want to Borrow? *The Quarterly Journal of Economics* 133(1), 129–190.
- Agarwal, S., C. Liu, and N. Souleles (2007). The reaction of consumer spending and debt to tax rebates-evidence from consumer credit data. *Journal of Political Economy* 115(6), 986–1019.
- 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.
- Andresen, M. E., A. R. KostÅl, R. T. Milton, C. Mommaerts, and L. Wallossek (2025, December). Monthly earnings volatility and household pooling. Working Paper 34563, National Bureau of Economic Research.
- Angeletos, G.-M., Z. Huo, and K. A. Sastry (2021). Imperfect macroeconomic expectations: Evidence and theory. *NBER Macroeconomics Annual* 35, 1–86.
- Attanasio, O. P. and G. Weber (2010, September). Consumption and saving: Models of intertemporal allocation and their implications for public policy. *Journal of Economic Literature* 48(3), 693–751.
- Augenblick, N., E. Lazarus, and M. Thaler (2024, 10). Overinference from weak signals and underinference from strong signals*. *The Quarterly Journal of Economics* 140(1), 335–401.
- 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.
- Ba, C., J. A. Bohren, and A. Imas (2022). Over-and underreaction to information. *Available at SSRN 4274617*.
- Balleer, A., G. Duernecker, S. Forstner, and J. Goensch (2025). The effects of biased labor market expectations on consumption, wealth inequality, and welfare. *American Economic Journal: Macroeconomics*. Forthcoming.
- Baron, M. and W. Xiong (2017). Credit expansion and neglected crash risk. *The Quarterly Journal of Economics* 132(2), 713–764.
- Baugh, B., I. Ben-David, H. Park, and J. A. Parker (2021, January). Asymmetric consumption smoothing. *American Economic Review* 111(1), 192–230.
- Bellifemine, M., A. Couturier, and S. H. Maasoum (2024). A distributional theory of household sentiment. *Technical Report*.
- 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.
- Beutel, J. and M. Weber (2025). Beliefs and portfolios: Causal evidence. Technical report.
- Bianchi, F., C. Ilut, and H. Saijo (2021). Diagnostic business cycles. *NBER Working Paper* 28604.

- Blundell, R., L. Pistaferri, and I. Preston (2008, December). Consumption inequality and partial insurance. *American Economic Review* 98(5), 1887–1921.
- Blundell, R. and I. Preston (1998, 05). Consumption Inequality and Income Uncertainty*. *The Quarterly Journal of Economics* 113(2), 603–640.
- Boehm, J., E. Fize, and X. Jaravel (2025, January). Five facts about mpcs: Evidence from a randomized experiment. *American Economic Review* 115(1), 1–42.
- Bordalo, P., N. Gennaioli, R. L. Porta, and A. Shleifer (2019). Diagnostic expectations and stock returns. *The Journal of Finance* 74(6), 2839–2874.
- Bordalo, P., N. Gennaioli, and A. Shleifer (2018). Diagnostic expectations and credit cycles. *The Journal of Finance* 73(1), 199–227.
- Bordalo, P., N. Gennaioli, A. Shleifer, and S. J. Terry (2021, January). Real credit cycles. Working Paper 28416, National Bureau of Economic Research.
- Bornstein, G. and S. Indarte (2023). The impact of social insurance on household debt. Available at SSRN 4205719.
- Caplin, A., V. Gregory, E. Lee, S. Leth-Petersen, and J. Sæverud (2023, March). Subjective earnings risk. Working Paper 31019, National Bureau of Economic Research.
- Carroll, C. D. (1997, 02). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis*. *The Quarterly Journal of Economics* 112(1), 1–55.
- Carroll, C. D. and P. Toche (2009, August). A tractable model of buffer stock saving. Working Paper 15265, National Bureau of Economic Research.
- Chamberlain, G. (1982, June). Panel data. Working Paper 913, National Bureau of Economic Research.
- Chatterjee, S., D. Corbae, M. Nakajima, and J.-V. Ríos-Rull (2007). A quantitative theory of unsecured consumer credit with risk of default. *Econometrica* 75(6), 1525–1589.
- Chodorow-Reich, G., A. M. Guren, and T. J. McQuade (2023, 04). The 2000s Housing Cycle with 2020 Hindsight: A Neo-Kindlebergerian View. *The Review of Economic Studies* 91(2), 785–816.
- Christelis, D., D. Georgarakos, T. Jappelli, and M. van Rooij (2020). Consumption uncertainty and precautionary saving. *Review of Economics and Statistics* 102(1), 148–161.
- Coibion, O., D. Georgarakos, Y. Gorodnichenko, G. Kenny, and M. Weber (2024). The effect of macroeconomic uncertainty on household spending. *American Economic Review* 114(3), 645–677.
- Coibion, O., Y. Gorodnichenko, and M. Weber (2020, August). How did u.s. consumers use their stimulus payments? Working Paper 27693, National Bureau of Economic Research.
- Coibion, O., Y. Gorodnichenko, and M. Weber (2022). Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy* 130(6), 1537–1584.
- Colarieti, R., P. Mei, and S. Stantcheva (2024, March). The how and why of household reactions to income shocks. Working Paper 32191, National Bureau of Economic Research.
- D’Acunto, F., D. Hoang, M. Paloviita, and M. Weber (2019). Cognitive abilities and

- inflation expectations. In *AEA Papers and Proceedings*, Volume 109, pp. 562–66.
- D’Acunto, F., D. Hoang, M. Paloviita, and M. Weber (2021). Cognitive constraints and the transmission of economic policies. Technical report, National Bureau of Economic Research.
- D’Acunto, F., D. Hoang, M. Paloviita, and M. Weber (2023). Iq, expectations, and choice. *The Review of Economic Studies* 90(5), 2292–2325.
- D’Acunto, F., D. Hoang, and M. Weber (2022). Managing households’ expectations with unconventional policies. *The Review of Financial Studies* 35(4), 1597–1642.
- D’Acunto, F., U. Malmendier, and M. Weber (2021). Gender roles produce divergent economic expectations. *Proceedings of the National Academy of Sciences* 118(21), e2008534118.
- D’Acunto, F. and M. Weber (2024). Why survey-based subjective expectations are meaningful and important. *Annual Review of Economics*.
- Daniel, K. and D. Hirshleifer (2015). Overconfident investors, predictable returns, and excessive trading. *Journal of Economic Perspectives* 29(4), 61–88.
- Daniel, K., D. Hirshleifer, and A. Subrahmanyam (1998). Investor psychology and security market under-and overreactions. *the Journal of Finance* 53(6), 1839–1885.
- DeFusco, A. A., B. Enriquez, and M. Yellen (2024, March). Wage garnishment in the united states: New facts from administrative payroll records. *American Economic Review: Insights* 6(1), 38–54.
- DellaVigna, S. (2009, June). Psychology and economics: Evidence from the field. *Journal of Economic Literature* 47(2), 315–72.
- Dempsey, K. and F. Ionescu (2023). Borrowing premia in unsecured credit markets. Technical report.
- Exler, F., I. Livshits, J. MacGee, and M. Tertilt (2024, 12). Consumer credit with over-optimistic borrowers. *Journal of the European Economic Association* 23(4), 1431–1478.
- Fagereng, A., M. B. Holm, and G. J. Natvik (2021, October). Mpc heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics* 13(4), 1–54.
- Fulford, S. L. (2015). How important is variability in consumer credit limits? *Journal of Monetary Economics* 72, 42–63.
- Fuster, A., G. Kaplan, and B. Zafar (2020, 11). What Would You Do with \$500? Spending Responses to Gains, Losses, News, and Loans. *The Review of Economic Studies* 88(4), 1760–1795.
- Ganong, P., D. Jones, P. J. Noel, F. E. Greig, D. Farrell, and C. Wheat (2025, August). Wealth, race, and consumption smoothing of typical income shocks. Working paper.
- Ganong, P. and P. Noel (2019, July). Consumer spending during unemployment: Positive and normative implications. *American Economic Review* 109(7), 2383–2424.
- Ganong, P., P. J. Noel, C. Patterson, J. S. Vavra, and A. Weinberg (2025). Earnings instability. *NBER Working Paper* 34227.
- Gathergood, J. and A. Olafsson (2024). The cohoding puzzle: New evidence from transaction-level data. *The Review of Financial Studies*, hhae016.
- Gennaioli, N. and A. Shleifer (2018). A crisis of beliefs: Investor psychology and financial fragility.

- Gertler, M. and N. Kiyotaki (2010). Chapter 11 - financial intermediation and credit policy in business cycle analysis. Volume 3 of *Handbook of Monetary Economics*, pp. 547–599. Elsevier.
- Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2021, May). Five facts about beliefs and portfolios. *American Economic Review* 111(5), 1481–1522.
- Glick, R., K. J. Lansing, et al. (2010). Global household leverage, house prices, and consumption. *FRBSF Economic Letter* 1, 1–5.
- Gorbachev, O. and M. J. Luengo-Prado (2019). The credit card debt puzzle: The role of preferences, credit access risk, and financial literacy. *Review of Economics and Statistics* 101(2), 294–309.
- Greenwood, R., S. G. Hanson, and L. J. Jin (2019, April). Reflexivity in credit markets. Working Paper 25747, National Bureau of Economic Research.
- 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.
- Guiso, L., T. Jappelli, and L. Pistaferri (2002). An empirical analysis of earnings and employment risk. *Journal of Business & Economic Statistics* 20(2), 241–253.
- Guvenen, F. (2007, June). Learning your earning: Are labor income shocks really very persistent? *American Economic Review* 97(3), 687–712.
- Haaland, I., C. Roth, and J. Wohlfart (2021). Designing information provision experiments. *Journal of Economic Literature*.
- Jappelli, T. and L. Pistaferri (2010). The consumption response to income changes. *Annual Review of Economics* 2(1), 479–506.
- Jappelli, T. and L. Pistaferri (2020). Reported mpc and unobserved heterogeneity. *American Economic Journal: Economic Policy* 12(4), 275–297.
- Kaplan, G., K. Mitman, and G. L. Violante (2020). The housing boom and bust: Model meets evidence. *Journal of Political Economy* 128(9), 3285–3345.
- Kaplan, G. and G. L. Violante (2022). The marginal propensity to consume in heterogeneous agent models. *Annual Review of Economics* 14, 747–775.
- Kiyotaki, N. and J. Moore (1997). Credit cycles. *Journal of Political Economy* 105(2), 211–248.
- Kline, P. and C. R. Walters (2016). Evaluating public programs with close substitutes: The case of head start. *The Quarterly Journal of Economics* 131(4), 1795–1848.
- Kling, J. R., J. B. Liebman, and L. F. Katz (2007). Experimental analysis of neighborhood effects. *Econometrica* 75(1), 83–119.
- Kluender, R., S. Indarte, U. Malmendier, and M. Stepner (2024). The roles of beliefs versus constraints in consumption decisions of low-income workers. *Working Paper*.
- Krueger, D., E. Malkov, and F. Perri (2023, November). How do households respond to income shocks? Working Paper 31894, National Bureau of Economic Research.
- Kuchler, T., M. Piazzesi, and J. Stroebel (2022, April). Housing market expectations. Working Paper 29909, National Bureau of Economic Research.
- Kuchler, T. and B. Zafar (2019). Personal experiences and expectations about aggregate outcomes. *The Journal of Finance* 74(5), 2491–2542.

- Kueng, L. (2018, 06). Excess Sensitivity of High-Income Consumers*. *The Quarterly Journal of Economics* 133(4), 1693–1751.
- Kuhnen, C. M. and A. C. Miu (2017). Socioeconomic status and learning from financial information. *Journal of Financial Economics* 124(2), 349–372.
- L’Huillier, J.-P., S. R. Singh, and D. Yoo (2023). Incorporating diagnostic expectations into the new keynesian framework. *Review of Economic Studies* (forthcoming).
- Liao, J. (2024). Over/underreaction to new information and noise in expectations formation. Available at SSRN 4673339.
- Livshits, I., J. MacGee, and M. Tertilt (2007, March). Consumer bankruptcy: A fresh start. *American Economic Review* 97(1), 402–418.
- López-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.
- Manski, C. F. (2004). Measuring expectations. *Econometrica* 72(5), 1329–1376.
- Maxted, P. (2023, 03). A Macro-Finance Model with Sentiment. *The Review of Economic Studies*, rdad023.
- Mian, A. and A. Sufi (2014). What explains the 2007-2009 drop in employment? *Econometrica* 82(6), 2197–2223.
- Mian, A. and A. Sufi (2018, August). Finance and business cycles: The credit-driven household demand channel. *Journal of Economic Perspectives* 32(3), 31–58.
- Mian, A., A. Sufi, and E. Verner (2017). Household debt and business cycles worldwide. *The Quarterly Journal of Economics* 132(4), 1755–1817.
- Mian, A., A. Sufi, and E. Verner (2020). How does credit supply expansion affect the real economy? the productive capacity and household demand channels. *The Journal of Finance* 75(2), 949–994.
- Olafsson, A. and M. Pagel (2018, 04). The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software. *The Review of Financial Studies* 31(11), 4398–4446.
- Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland (2013, October). Consumer spending and the economic stimulus payments of 2008. *American Economic Review* 103(6), 2530–53.
- Pistaferri, L. (2001, 08). Superior Information, Income Shocks, and the Permanent Income Hypothesis. *The Review of Economics and Statistics* 83(3), 465–476.
- Roth, C. and J. Wohlfart (2020). How do expectations about the macroeconomy affect personal expectations and behavior? *Review of Economics and Statistics* 102(4), 731–748.
- Rozsypal, F. and K. Schlafmann (2023, October). Overpersistence bias in individual income expectations and its aggregate implications. *American Economic Journal: Macroeconomics* 15(4), 331–71.
- Souleles, N. (2004). Expectations, heterogeneous forecast errors, and consumption: Micro evidence from the michigan consumer sentiment surveys. *Journal of Money, Credit and Banking* 36(1), 39–72.
- Telyukova, I. A. (2013, 01). Household Need for Liquidity and the Credit Card Debt Puzzle. *The Review of Economic Studies* 80(3), 1148–1177.

- Weber, M., F. D'Acunto, Y. Gorodnichenko, and O. Coibion (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives* 36(3), 157–184.
- Yin, X. (2025). Learning in the limit: Income inference from credit extensions. *The Journal of Finance*. Forthcoming. Available at SSRN: <https://ssrn.com/abstract=4254400>.
- Zinman, J. (2009). Debit or credit? *Journal of Banking & Finance* 33(2), 358–366.

Figure 1. Survey Design

This figure describes the timeline of the survey waves run in each year.

Jan 2020 to Jan 2025 (excluding 2022)

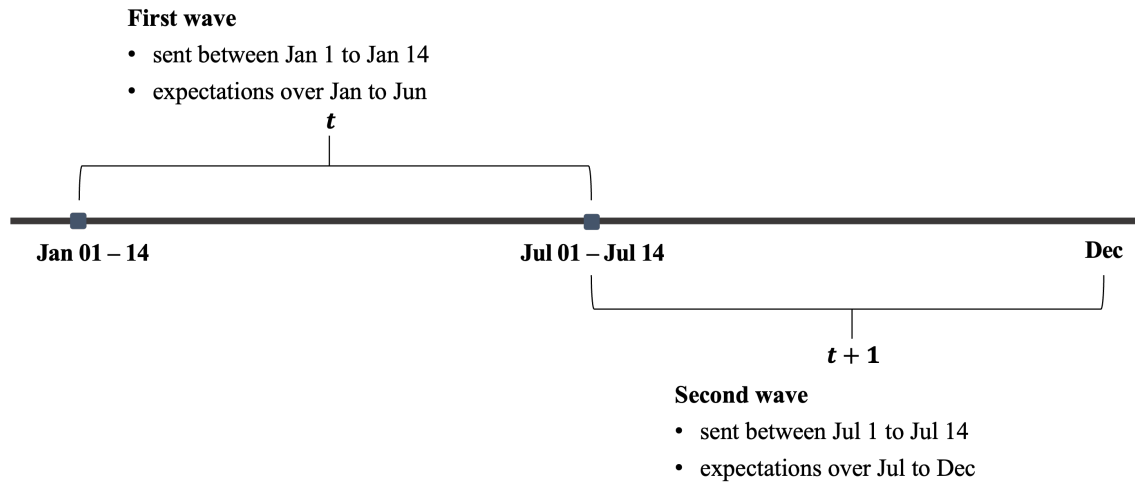


Figure 2. Income Shocks and Income Surprises

This figure plots income surprises vs income shocks. Income surprises - log are defined as the difference between log realized income and log subjective income expectations. Income shocks - log are defined as the unexpected income innovation from estimating (4). Income Surprises - Level and Income Shocks - Level are defined analogously in CNY thousands. Panels A and C plot the raw values, panels B and D control for individual fixed effects. The shaded regions give the 95% confidence intervals for the fitted lines. All variables are winsorized at the 1%-99% levels within each wave.

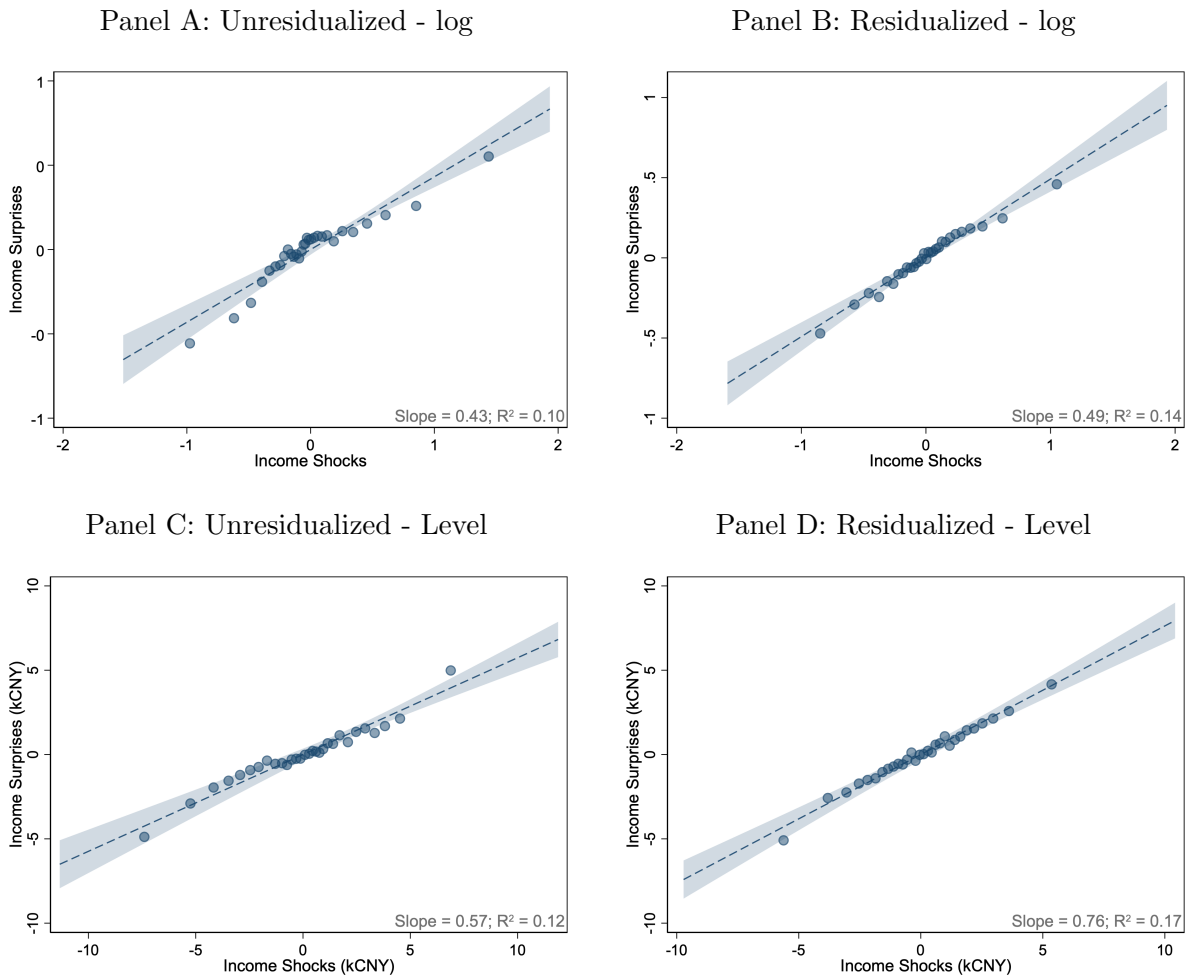
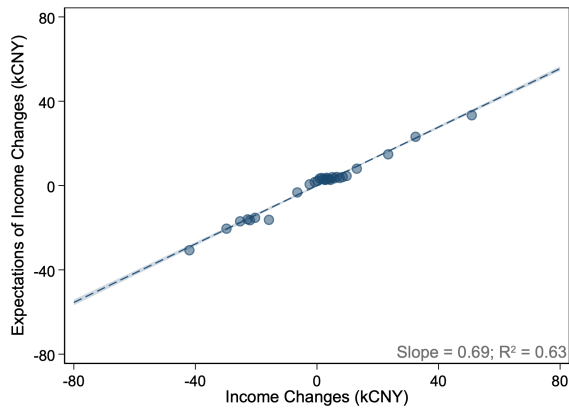



Figure 3. Errors in Subjective Income Expectations

Panels A and B are binned scatter plots of the ex-post realized income changes and income growth against ex-ante subjective expectations. Panels C and D are histograms of the individual-level income forecast errors. Panel D residualizes the forecast errors by individual fixed effects. The units in panels A, C, and D are CNY thousands. The shaded regions give the 95% confidence intervals for the fitted lines, where standard errors are clustered at the industry \times wave level. All variables are winsorized at the 1%-99% levels within each wave.

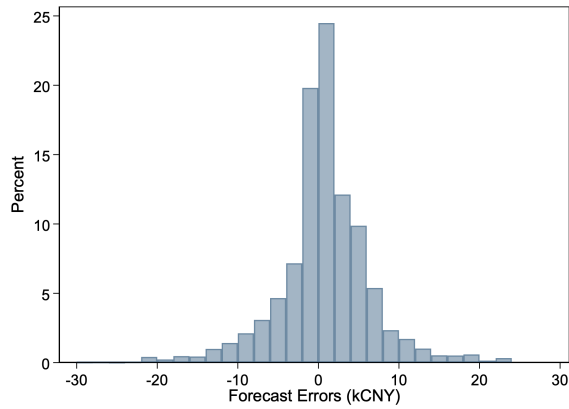
Panel A: Expectations vs Realizations – Changes



Panel B: Expectations vs Realizations – Growth



Panel C: Forecast Errors



Panel D: Forecast Errors – Residualized




Figure 4. Income Shocks, Income Surprises, and Forecast Errors

This figure plots binned scatter of future income forecast errors against current income shocks in panels A and B, and against current income surprises in panels C and D. The units in panels A and C are CNY thousands. The shaded regions give the 95% confidence intervals for the fitted lines, where standard errors are clustered at the industry \times wave level. All plots control for individual fixed effects. All variables are winsorized at the 1%-99% levels within each wave.

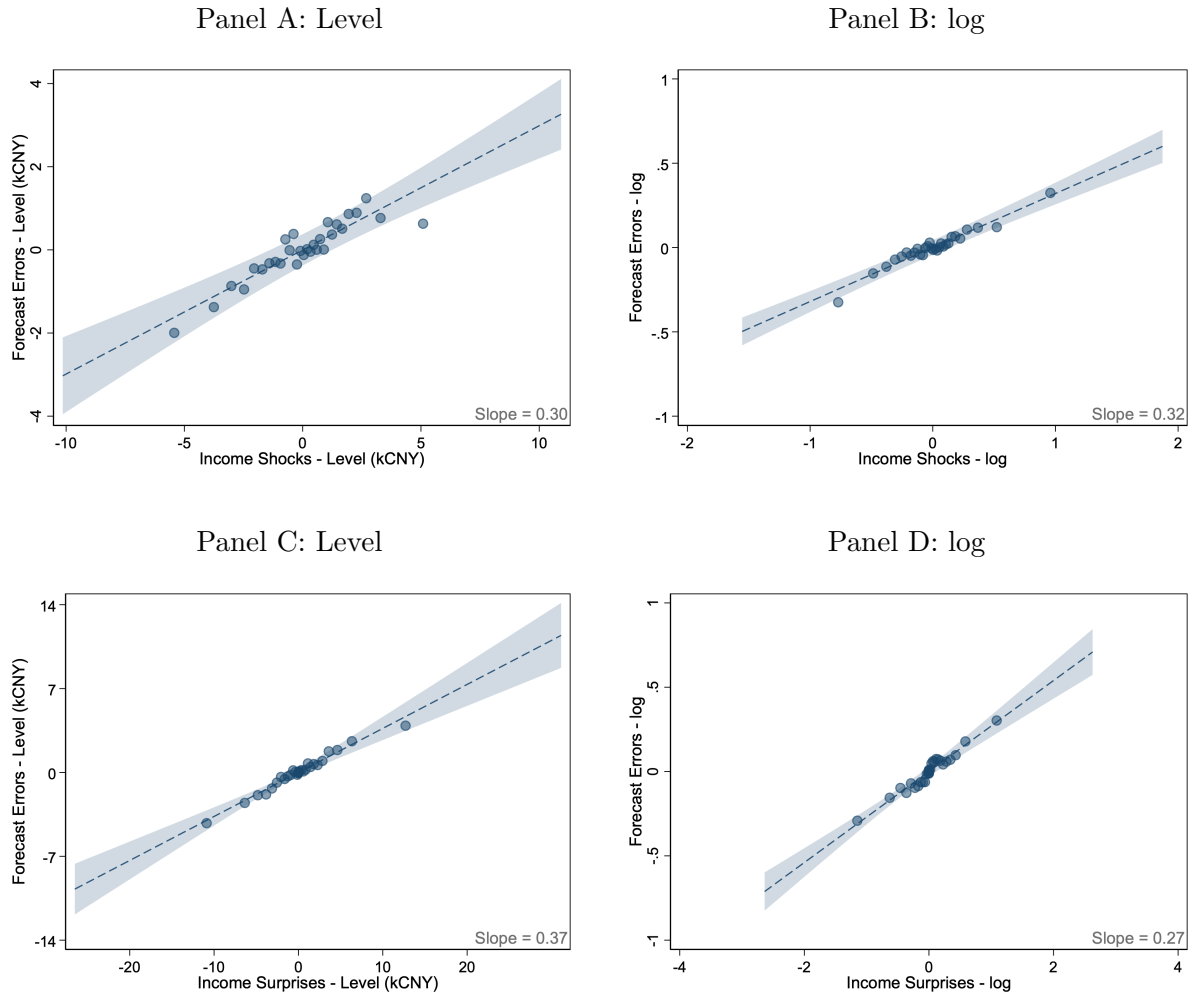


Figure 5. Dynamics of Consumption and Debt Accumulation around Income Forecast Errors

This figure reports the coefficient estimates of cumulative spending and debt on forecast errors. Estimation is based on the specification in equation (7). In Panel A, the outcome variable is total spending. In Panel B, it is the stock of unsecured debt. The red solid lines report the estimates for consumers with positive forecast errors and the blue dashed lines for consumers with negative forecast errors. The width of the shaded regions is twice the size of estimated standard errors.

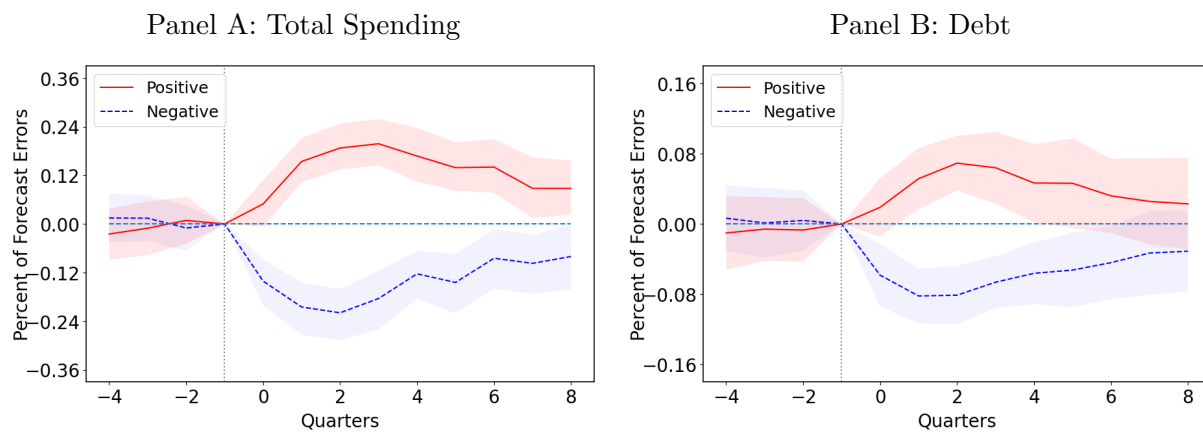


Figure 6. Structural Model - Boom-Bust Household Cycles

This figure shows simulated paths from the structural model following a sequence of transitory income shocks. Red solid lines correspond to overreaction ($\theta = 2.54$) and blue dashed lines to Bayesian learning ($\theta = 0$). The top panels display the income shocks and expected income, and the bottom panels report changes in consumption and debt relative to no-shock average income, and percentage deviations of default relative to the no-shock baseline.

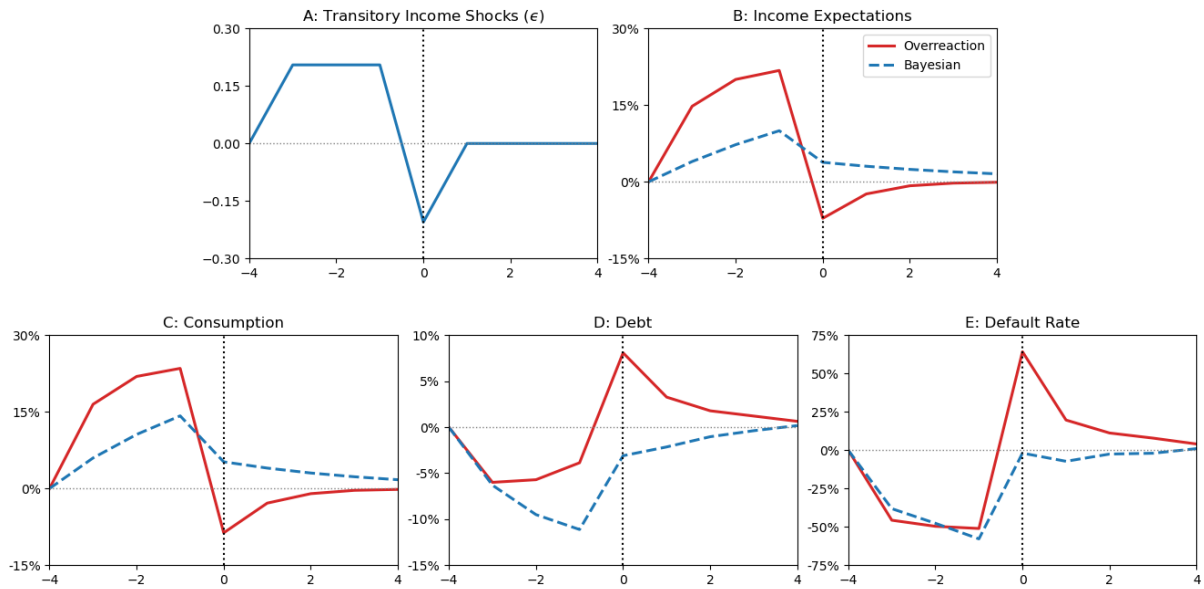


Figure 7. Simulating the Global Financial Crisis in the US

This figure simulates the 2008-2009 Global Financial Crisis using a sequence of transitory income shocks at the semiannual frequency. Red solid lines correspond to overreaction ($\theta = 2.54$), blue dotted lines to Bayesian learning ($\theta = 0$), and black dashed lines to the data. The left panel shows the detrended unsecured-debt-to-income ratio, and the right panel shows the aggregate default rate, defined as the delinquency rate multiplied by the share of debt holders.

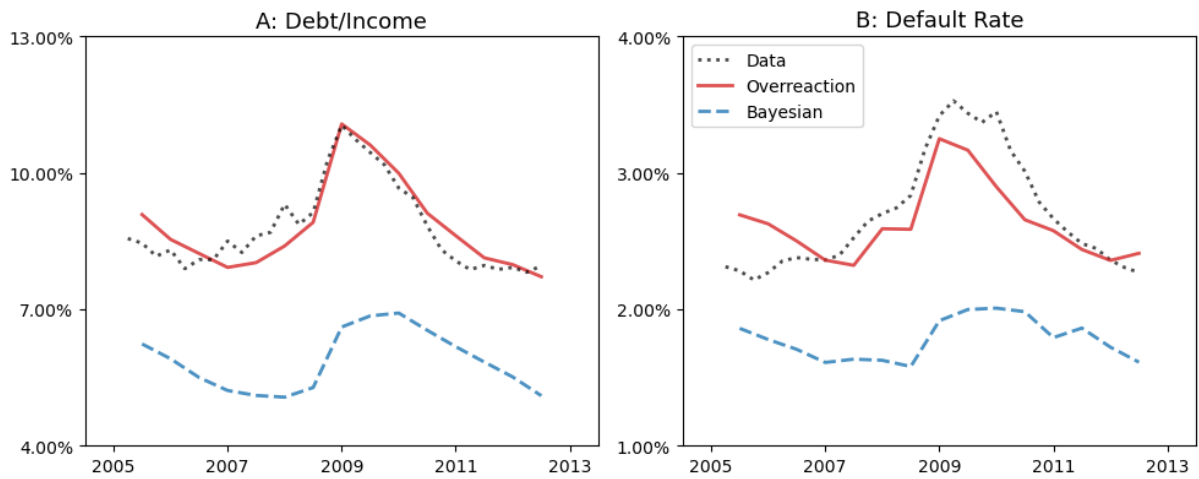


TABLE I. Summary Statistics

Income and Spending are respectively the average monthly income and spending. Credit Limit is the total credit limit across all financial institutions. Debt is the interest-incurring credit card debt for each consumer across all financial institutions. Debt|Debt > 0 is the interest-incurring credit card debt for those who hold positive debt. Liquidity is the ratio of liquid wealth to six-month income. Income Growth is the annual log income growth. Income Shock - log is the ϵ_t retrieved from fitting equation (4). Income Shock - Level is $Y_t - E_{t-1}[Y_t]$ implied by equation (4), where Y_t is the level of income. $E^\theta[\text{Income}]$ is the subjective expectations of average monthly income over the next six months. All continuous variables are measured in CNY thousands and winsorized at the 1%-99% levels within each wave.

	Mean	SD	25-Pctl	Median	75-Pctl	N
Age	39.58	9.24	33.00	39.00	46.00	26,449
Female	0.52	0.50	0.00	1.00	1.00	26,449
College	0.53	0.50	0.00	1.00	1.00	26,449
Income	13.03	13.26	4.20	8.27	16.66	26,449
Spending	7.08	6.53	2.66	5.08	9.27	26,449
Installment	2.82	13.32	0.00	1.55	18.21	26,449
Debt	5.01	11.06	0.00	0.00	4.65	26,449
Debt Debt > 0	12.46	14.53	2.63	7.37	16.83	10,650
Liquidity	1.19	1.54	0.01	0.91	2.54	26,449
Credit Limit	81.33	117.12	15.00	30.00	100.00	26,449
Income Growth	0.04	0.38	-0.00	0.00	0.06	26,449
Income Shocks - log	-0.05	0.39	-0.22	-0.08	0.05	26,449
Income Shocks - Level	-0.92	2.56	-2.13	-0.40	0.07	26,449
$E^\theta[\text{Income}] - \text{Income}$	0.27	5.45	-1.51	0.00	2.49	26,449

TABLE II. Income Shocks, Income Surprises, and Forecast Errors

Forecast Errors $_{t+1}$ - Level is the difference between the subjective expectation of income and realized income in $t + 1$. Forecast Errors $_{t+1}$ - log is the difference between the log of subjective income expectations and realized log income in $t + 1$. Shocks $_t$ - log is the ϵ_t retrieved from fitting (4). Shocks $_t$ - Level is $Y_t - E_{t-1}[Y_t]$ implied by (4), where Y_t is the level of income in t . Surprises $_t$ - Level and Surprises $_t$ - log are respectively the negative of Forecast Errors $_t$ - Level and Forecast Errors $_t$ - log. Credit Limit $_t$ is the total credit limit across all financial institutions. Liquidity $_t$ is the ratio of liquid wealth to six-month income. Income Growth $_t$ is the annual log income growth. All continuous variables are measured in CNY thousands and winsorized at the 1%-99% levels within each wave.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Forecast Errors $_{t+1}$ - Level				Forecast Errors $_{t+1}$ - log			
Shocks $_t$ - Level	0.351*** (0.042)	0.431*** (0.053)						
Surprises $_t$ - Level			0.393*** (0.057)	0.493*** (0.040)				
Shocks $_t$ - log					0.320*** (0.027)	0.389*** (0.030)		
Surprises $_t$ - log							0.264*** (0.029)	0.352*** (0.028)
Credit Limit $_t$		0.000 (0.000)		-0.001 (0.001)		0.000 (0.000)		-0.000 (0.000)
Liquidity $_t$		0.033 (0.077)		-0.331** (0.150)		0.014* (0.008)		0.043*** (0.016)
Income Growth $_t$		0.141 (0.401)		-1.248*** (0.232)		-0.060* (0.033)		-0.110*** (0.030)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry×Wave FE	No	Yes	No	Yes	No	Yes	No	Yes
N Obs	26,449	26,449	15,644	15,644	25,407	25,407	14,354	14,354

Standard errors clustered by industry \times wave in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE III. Income Shocks and Forecast Errors – Heterogeneity

Forecast Errors_{*t*+1} is the difference between the subjective expectation of income and realized income in *t* + 1. Shocks_{*t*} is $Y_t - E_{t-1}[Y_t]$ implied by (4), where Y_t is the level of income. Young is defined as age smaller than 39. More volatile is one if average subjective income volatility is above the median, and zero otherwise. Government is one if the participants work in the government industry. Salary Worker is one if the individual’s regular income payment as a fraction of total income is in the upper quartile. Controls include credit limit, liquidity, and income growth measured at time *t*. All continuous variables are measured in CNY thousands and winsorized at the 1%-99% levels within each wave.

	(1)	(2)	(3)	(4)	(5)	(6)
	Forecast Errors _{<i>t</i>+1}					
Shocks _{<i>t</i>}	0.370*** (0.060)	0.240*** (0.049)	0.465*** (0.053)	0.501*** (0.053)	0.397*** (0.091)	0.401*** (0.085)
Shocks _{<i>t</i>} × Young	0.137*** (0.041)					
Shocks _{<i>t</i>} × More Volatile		0.359*** (0.044)				
Shocks _{<i>t</i>} × Government			-0.469*** (0.060)			
Shocks _{<i>t</i>} × Salary Worker				-0.210*** (0.048)		
Shocks _{<i>t</i>} × (Year ≥ 2023)					-0.070 (0.079)	
Shocks _{<i>t</i>} × (Shocks _{<i>t</i>} < 0)						0.098 (0.088)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Wave FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
N Obs	26,449	25,551	26,449	26,449	26,449	26,449

Standard errors clustered by industry × wave in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE IV. Income Forecast Errors, Liquidity Constraints, and Spending

ΔC_t is the change in total spending from period $t - 1$ to t . Shocks_t is $Y_t - E_{t-1}[Y_t]$ implied by (4), where Y_t is the level of income. $\text{Forecast Errors}_{t+1}$ is the difference between the subjective income expectation and realized income in $t + 1$. High Util_t is one if utilization rate is above the median in t , and zero otherwise. High Liq_t is one if Liquidity is above the median in t , and zero otherwise. Controls include credit limit, liquidity, and income growth measured at time t . All continuous variables are measured in CNY thousands and winsorized at the 1%-99% levels within each wave.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	ΔC_t						
Shocks_t	0.254*** (0.032)		0.166*** (0.032)	0.195*** (0.049)	0.317*** (0.036)	0.115** (0.050)	0.221*** (0.038)
$\text{Forecast Errors}_{t+1}$		0.207*** (0.020)	0.196*** (0.022)			0.171*** (0.029)	0.219*** (0.022)
High Util_t				0.161 (0.187)		0.156 (0.185)	
$\text{Shocks}_t \times \text{High Util}_t$				0.118** (0.054)		0.104* (0.056)	
High Liq_t					-0.065 (0.251)		-0.067 (0.250)
$\text{Shocks}_t \times \text{High Liq}_t$					-0.124** (0.048)		-0.106** (0.049)
$\text{Forecast Errors}_{t+1} \times \text{High Util}_t$						0.049* (0.027)	
$\text{Forecast Errors}_{t+1} \times \text{High Liq}_t$							-0.046* (0.025)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry \times Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N Obs	26,449	26,449	26,449	26,449	26,449	26,449	26,449

Standard errors clustered by industry \times wave in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE V. Forecast Errors and Debt/Default Decisions

ΔDebt_{t+1} is the change in total unsecured borrowing between time t and $t + 1$. Default_{t+1} is one if there is a 90-day delinquency on the unsecured borrowing in $t + 1$, and zero otherwise. Shocks_t is $Y_t - E_{t-1}[Y_t]$ implied by (4), where Y_t is the level of income. $\text{Forecast Errors}_{t+1}$ is the difference between the subjective expectation of income and realized income in $t + 1$. Controls include credit limit, liquidity, and income growth measured at time t . In columns (5) - (8), Shocks_t and $\text{Forecast Errors}_{t+1}$ are standardized. All other continuous variables are measured in CNY thousands and winsorized at the 1%-99% levels within each wave.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	ΔDebt_{t+1}				Default_{t+1}			
Shocks_t	0.024 (0.020)	-0.003 (0.020)	-0.014 (0.019)	-0.015 (0.019)	0.155 (0.122)	0.004 (0.117)	0.002 (0.116)	0.003 (0.115)
$\text{Forecast Errors}_{t+1}$		0.061*** (0.009)	0.062*** (0.008)	0.065*** (0.009)		0.846*** (0.147)	0.846*** (0.147)	0.842*** (0.150)
Shocks_{t+1}				-0.030** (0.015)				0.022 (0.113)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry \times Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	No	Yes	Yes	No	No	Yes	Yes
N Obs	26,449	26,449	26,449	26,449	26,449	26,449	26,449	26,449

Standard errors clustered by industry \times wave in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE VI. Income Shocks, Forecast Errors, and Decisions - 2SLS

Shocks_{*t*} is $Y_t - E_{t-1}[Y_t]$ implied by (4), where Y_t is the level of income. Forecast Errors is the difference between the subjective expectation of income and realized income. ΔC_t is the change in total spending from period $t - 1$ to t . ΔDebt_{t+1} is the change in total unsecured borrowing from period t to $t + 1$. Default_{*t+1*} is one if there is a 90-day delinquency on the unsecured borrowing in $t + 1$, and zero otherwise. Shocks Same Type_{*t*} is the average of Shocks_{*t*} for individuals in the same type excluding the survey sample. High Vol_{*t*} is one if the participant is in a type with high income volatility, and zero otherwise. Controls include credit limit, liquidity, and income growth measured at time t . In column (5), Shocks_{*t*} and Forecast Errors_{*t+1*} are standardized. All other continuous variables are measured in CNY thousands and winsorized at the 1%-99% levels within each wave.

	(1)	(2)	(3)	(4)	(5)
	Shocks _{<i>t</i>}	Forecast Errors _{<i>t+1</i>}	ΔC_t	ΔDebt_{t+1}	Default _{<i>t+1</i>}
	First Stage		Second Stage		
Shocks Same Type	0.185*** (0.010)	-0.221*** (0.020)			
Shocks Same Type _{<i>t</i>} × High Vol _{<i>t</i>}	0.059*** (0.011)	0.453*** (0.032)			
Shocks _{<i>t</i>}			0.126** (0.064)	-0.022 (0.039)	-0.034 (0.240)
Forecast Errors _{<i>t+1</i>}			0.216*** (0.063)	0.088** (0.043)	1.276** (0.567)
Individual FE	Yes	Yes	Yes	Yes	Yes
Industry × Wave FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
First-Stage <i>F</i>				104.47	
N Obs	26,449	26,449	26,449	26,449	26,449

Standard errors clustered by industry × wave in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE VII. Structural Model: Estimation

This table reports the estimated parameters of the structural model. Panel A presents the parameters estimated in the first stage and Panel B those estimated in the second stage based on SMM. Panel C reports the matched moments and Panel D the moments not targeted directly. A/C is the average wealth-consumption ratio. P(Default) is the proportion of defaults. A/Y is the average wealth-income ratio. Median(A/Y) is the median of the wealth-income ratio. Top 5% liq share is the fraction of savings held by the top 5% individuals in the model. Debtor% is the fraction of consumers with positive debt. B/Y is the average debt-to-income ratio. Estimation of moments in the model is based on a simulation of 20,000 individuals with 1,000 periods, after a burn-in period of 100 periods for the distribution to reach the steady state. Estimation of moments in the data is based on a random 5% of active customers in the bank's database. Model moments are trimmed at savings to average income ratio larger than 20.

Panel A		Panel B		Panel D			
First-Stage Parameters		Second-Stage Parameters		Not Targeted Moments			
	Estimate		Estimate	S.E.			
	(1)		(2)	(3)		Data	Model
						(6)	(7)
γ	2.000	χ	2.300	(0.107)	A/Y	1.211	1.131
β	0.975	λ	0.250	(0.039)	Median(A/Y)	0.912	1.000
r_b	0.045	Panel C			Top 5% liq share	28.97%	26.71%
r_s	0.015	Targeted Moments			Debtor%	19.33%	16.10%
ν	0.050		Data	Model	B/Y	0.230	0.235
α	0.000		(4)	(5)			
σ_η	0.089	A/C	0.957	0.957			
σ_ϵ	0.410	P(Default)	3.60%	3.60%			
θ	2.540						

TABLE VIII. Linking Shocks, Forecast Errors, and Outcomes

ΔC_t is the change in total spending from period $t - 1$ to t . ΔDebt_{t+1} is the changes in total unsecured borrowing between time t and $t + 1$. Default_{t+1} is one if there is a 90-day delinquency on the unsecured borrowing in $t + 1$, and zero otherwise. Shocks_t is $Y_t - E_{t-1}[Y_t]$ implied by (4), where Y_t is the level of income. Forecast Errors is the difference between the subjective expectation of income and realized income. Controls include credit limit, liquidity, and income growth measured at time t . All continuous variables are measured in CNY thousands and winsorized at the 1%-99% levels within each wave.

	(1) Forecast Errors $_{t+1}$	(2) ΔC_t	(3) ΔDebt_{t+1}	(4) Default $_{t+1}$
Panel A: Data				
Shocks $_t$	0.431 (0.053)			
Forecast Errors $_{t+1}$		0.216 (0.063)	0.088 (0.045)	1.276 (0.567)
N Obs	26,449	26,449	26,449	26,449
Panel B: $\theta = 2.540$				
Shocks $_t$	0.431 (0.001)			
Forecast Errors $_{t+1}$		0.174 (0.003)	0.068 (0.000)	1.319 (0.005)
N Obs	100,000	9,920,072	9,920,072	9,920,072
Panel C: $\theta = 0$				
Shocks $_t$	-0.026 (0.003)			
Forecast Errors $_{t+1}$		0.036 (0.000)	0.009 (0.000)	0.532 (0.003)
N Obs	100,000	9,920,072	9,920,072	9,920,072

TABLE IX. Equilibrium Moments and Welfare Costs

Panel A reports moments from the stationary distribution. Default is the unconditional default rate. C/Y and A/Y denote consumption- and wealth-to-income ratios. SD log(C) is the average of within-individual standard deviation of log consumption. Panel B reports welfare costs of overreaction as consumption-equivalent variations (CEV). For each individual i , λ_i solves $\sum_t \beta^t u((1 + \lambda_i)C_{it}^O, d_{it}^O) = \sum_t \beta^t u(C_{it}^B, d_{it}^O)$, where $u(c, d) = \frac{c^{1-\gamma}}{1-\gamma} - \chi d$ and O, B denote overreaction and Bayesian regimes. The table reports percentiles and the mean of λ_i . Total includes default penalties. No Default sets $\chi = 0$. Fixed C Level rescales each individual's overreaction consumption to match the Bayesian mean, isolating smoothing effects. Simulations use 20,000 agents for 100 periods after 100 burn-in periods.

Panel A: Equilibrium Moments						
	P(Default)	Debtor%	C/Y	SD log(C)	Median(A/Y)	A/Y
Overreaction	3.60%	16.10%	1.092	0.593	1.090	1.307
Bayesian	2.91%	14.92%	1.119	0.515	1.232	1.442
Panel B: Welfare Costs						
	10%	25%	50%	75%	90%	mean
Total	1.63%	4.67%	8.37%	14.14%	24.83%	12.27%
No Default	0.99%	3.25%	5.76%	9.01%	14.91%	8.41%
Fixed C Level	-0.72%	1.63%	4.30%	7.63%	13.60%	6.85%