Microfounding Household Debt Cycles with Extrapolative Expectations

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

Jan 2024

Abstract

Combining transaction-level data with survey-based information from a large consumer panel, we show that on average consumers form excessively high expectations about future income relative to ex-post realizations after unexpected positive income shocks. This systematic bias in expectations leads to higher current consumption and debt accumulation as well as a higher likelihood of subsequent default on consumer debt. A consumption-saving model with defaultable unsecured debt and diagnostic Kalman filtering with consumers who over-extrapolate income shocks rationalizes these findings. The model predicts excessive leverage and higher subsequent default rates compared to a rational expectations benchmark. Over-extrapolation of income expectations can contribute to explaining state-dependent household debt cycles.

Keywords: Income Expectations, Consumption, Household Debt, Behavioral Economics, Behavioral Finance, Surveys, Income Volatility.

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

*D'Acunto: Georgetown University, francesco.dacunto@georgetown.edu; Weber: University of Chicago and NBER, michael.weber@chicagobooth.edu; Yin: University College London, xiao.yin@ucl.ac.uk. We deeply appreciate the valuable comments from Yuriy Gorodnichenko, Emi Nakamura, and Peter Maxted. We are thankful for the valuable comments from UC Berkeley, UCL, 2022 Society of Economic Dynamics Meeting, and 2023 American Economic Association Annual Meeting.

I Introduction

Household debt cycles have been a focus of research in economics and finance since the Global Financial Crisis. Spikes in household debt predict the eruption of economic crises across space and over time in the US (Mian and Sufi, 2014) and globally (Reuven and Lansing, 2010; Mian et al., 2017). Researchers have proposed arguments and produced evidence consistent with both supply-side (e.g., Agarwal et al. (2017), Mian et al. (2020)) and demand-side drivers of household debt cycles (e.g., Bordalo et al. (2018); Bianchi et al. (2021); Chodorow-Reich et al. (2021))

Due to their intrinsic co-determination, disentangling the roles of supply- and demand-side drivers of household debt cycles using observational data is empirically challenging. Separately identifying the roles of changes in credit supply and changes in the demand for credit is hard when using aggregate macroeconomic data on household debt and GDP growth, given that both channels might be affected by the same unobserved shocks and can influence each other. At the micro level, testing directly for a potential effect of demand-side channels, such as consumers' expectations, on households' consumption and debt accumulation requires observing individual-level data that include, at the same time, information about households' income flows, spending, and debt accumulation as well as direct elicitation of beliefs about future individual outcomes and the subsequent realizations of these outcomes.

In this paper, we propose a unique setting to tackle these empirical challenges. Our setting builds on transaction-level bank-account panel data for a large population, which includes households' debt and spending decisions as well as their consumer-credit limits, income inflows, and demographic characteristics. The unique feature of this setting is the direct elicitation of individuals' expectations for their own income through customized surveys that were administered repeatedly to the same consumer population over time. This panel structure enables us to study the dynamics of beliefs, consumption, debt accumulation, and debt outcomes within individuals, but also the comparison to past and future realized income. [1]

Our setting provides us with direct measures of consumers' beliefs, ex-post accuracy, and choices. In each period, we observe not only consumers' income inflows but also the difference between realized flows and previous-period numerical income expectations (*contemporaneous unexpected income shocks*). Moreover, in each period we can assess how expectations about future income differ from subsequent income realizations—a direct measure of the ex-post accuracy of consumers' income expectations. In this way, we can study directly how income expectations react when consumers face unexpected income shocks and which consumption and debt choices agents make based on such expectations dynamics.

We start by showing that the income expectations of the average consumer overreact to unexpected income shocks after both positive and negative shocks. That is, after a shock in either direction, expectations become systematically inaccurate in the direction of the current shock relative to future income realizations. Moreover, the size of this overreaction increases proportionally with the size of the unexpected income shock. The emergence of extrapolative income beliefs is not explained by a rich set of individual-level demographics, locations, or occupations, which we proxy with the industries in which consumers work. At the same time, demographics are important to understand the cross-sectional variation in the size of the belief mistake, which is larger for lower-income consumers, for consumers who face more volatile income flows, and for younger consumers. Facing larger income shocks increases also the second moment of consumers' income expectations—an effect that looms larger for negative shocks than for same-size positive shocks.

To better understand the economic channels that drive these results and obtain additional predictions to test empirically, we propose a simple theoretical framework of the intertemporal consumption choice problem featuring extrapolative expectations. In the framework, we use the diagnostic-expectations formulation of Bordalo et al. (2018) as a microfoundation of consumer expectation-formation process. This choice is motivated

[1] For a recent review of survey-based research in economics, see Haaland et al. (2021). Examples in macroeconomics include Bachmann et al. (2013), D'Acunto et al. (2021b), and Coibion et al. (2023) among others.

by our results that consumers overreact to the unexpected component of income shocks, which suggests that they put excessive weights on the states of the world that have objectively become more representative of future states based on observed realizations.

In this framework, a larger deviation between expected future income and realized future income predicts higher current consumption and, when consumers are allowed to borrow to finance current consumption, higher current borrowing and a higher likelihood of subsequent default. We find empirical evidence consistent with these three predictions: the larger is the difference between consumers' expected income and their realized future income, the larger is the present-day increase in spending and next period's stock of debt as well as the likelihood of default.

Once augmented with time-varying income volatility, the theoretical framework predicts that the consumption responses to unexpected income shocks should be asymmetric around zero: negative shocks should loom more than same-size positive shocks—a prediction for which we also find evidence in our data. This finding is important to microfound the aggregate dynamics of household debt cycles with inaccurate beliefs because existing literature shows that the buildup of spending and borrowing in times of positive GDP growth shocks is slower than the sudden drop in spending and borrowing after negative shocks (Mian, Sufi, and Verner, 2017).

Overall, our results provide direct micro-level evidence of a demand-side channel that might help explain the dynamics of household debt cycles across space and over time: the average consumer forms systematically excessive beliefs about future income when facing unexpected positive income shocks. She consequently raises more debt to finance higher current spending and, once subsequent income does not reach the expected levels, she is more likely to default.

Given the panel structure of the data, we can isolate the effects of biased belief by exploiting the emergence of idiosyncratic income shocks across consumers who make decisions at the same point in time by using within-individual variation in expectations and outcomes over time. Because we observe elicited expectations and outcomes for the same individual over time, our empirical analysis can also abstract from aggregate shocks that affect all consumers at the same point in time, such as the COVID-19 pandemic, which falls within our sample period.

At the same time, the individual-level channel we document can be important to understand aggregate dynamics of household debt cycles around economy-wide shocks: in the aggregate, in times of positive GDP growth more and more consumers might face unexpected income shocks and hence, through extrapolative income expectations, a larger fraction of consumers increases consumption and debt accumulation. By contrast, when GDP growth turns negative, more and more consumers face unexpectedly negative income shocks, drop their consumption and borrowing, and are more likely to default.

Inspired by this consideration, in the last part of the paper we investigate to what extent the micro-level demand-side channel we document might help us understand the dynamics of household debt cycles in the aggregate. To this aim, we estimate a structural model that incorporates diagnostic expectations with defaultable debt into a standard consumption model with incomplete markets and heterogeneous agents. We use this model to perform two exercises. First, we show that the model is successful in replicating the relationships between expectation errors and consumption, debt, and default choices we observe in the data with standard parametrizations.

As a second exercise, we simulate a heterogeneous-agent economy to isolate the effects of diagnostic expectations on aggregate consumption, debt, and default dynamics. We find that diagnostic expectations amplify the effects of positive income shock on income expectations and consumption. For this reason, agents who form expectations diagnostically start to accumulate more debt after positive income shocks relative to agents who learn rationally. Once we remove positive shocks, default risk increases substantially, but only with diagnostic expectations.

Overall, our simulation based on the structural model implies that diagnostic expectations can generate dynamics that are consistent with facts about household debt cycles documented in the literature. For instance, the fact that elevated consumer sentiment implies higher subsequent debt growth (López-Salido et al., 2017), and that the end of a booming period usually coincides with times of high financial fragility (Greenwood et al., 2019; Maxted, 2023).

Note that our paper does not argue that supply-side channels are irrelevant to explain the dynamics of household debt cycles. To the contrary, credit limits and, more broadly, consumers' debt capacity are determined not only by consumers' income but also by banks' credit decisions. For instance, Aydin (2022) and Yin (2022) document the impact of supply-side forces on credit uptake and consumer spending by exploiting randomized increases in credit-card limits. Also, we microfound extrapolative expectations with the diagnostic framework because such framework is portable to various theoretical and empirical applications in economics, but we do not argue that diagnostic expectations are the only theoretical underpinning of the expectations-formation process that can be consistent with our results. Other forms of extrapolative expectations might also be consistent with some or all our results (for instance, see Barberis et al. (2018) and Barberis (2018)).

Our paper contributes to at least three strands of literature. First, it contributes to the study of the drivers of credit cycles. Theoretically, two categories of explanations emerge. On the one hand, financial frictions in the corporate and household sectors can be an amplification mechanism that induces cycles in credit supply (for instance, see Kiyotaki and Moore (1997), Gertler and Kiyotaki (2010), Brunnermeier and Sannikov (2014), He and Krishnamurthy (2019), Li (2019), and Mian et al. (2020)). On the other hand, beliefs and hence a demand channel can be relevant. For instance, beliefs might change when the incentives of producing information change over time, thereby inducing swings in asset prices and macroeconomic fluctuations (Gorton and Ordoñez (2014), Dang et al. (2020)). Alternatively, consumers' extrapolative expectations can generate credit cycles (Bordalo et al. (2018), Bianchi et al. (2021), L'Huillier et al. (2023) and Bordalo et al. (2021)). Our paper provides empirical evidence at the micro level consistent with a demand channel based on consumers' beliefs even though, as Krishnamurthy and Li (2020) suggest, both supply- and demand-side channels are likely to be important to explain credit fluctuations.

On the empirical side, the analysis of credit-cycle fluctuations so far has mostly focused on aggregate economy-wide or regional-level data.[2] In our paper, we focus on

[2]See, for instance, see Bordo et al. (2001), Borio and Lowe (2002), Claessens et al. (2010), Reinhart and Rogoff (2009), Borio and Lowe (2013), Jordà et al. (2013), Baron and Xiong (2017), Greenwood et al. (2020), Krishnamurthy and Muir (2017), Mian and Sufi (2018), Mian et al. (2020), and Baron et al. (2020), among others.

the household sector and contribute by providing micro-level evidence in a panel dataset that allows us to uncover how unexpected income shocks can induce over-reaction about future income, which leads to overreaction in total borrowing and to higher default rates.

The second area to which we contribute is the rich literature on the marginal propensity to consume out of income shocks (Parker et al., 2013; Fuster et al., 2020; Kueng, 2018; Olafsson and Pagel, 2018; Baugh et al., 2021; Fagereng et al., 2021)[3]. While most previous literature looks directly at the relationship between shocks to income and consumption, we document evidence that income shocks affect consumption also through a biased belief channel. Moreover, thanks to the richness of our transaction-level data, we can also provide evidence on the effects of income shocks not only on consumption but also on borrowing outcomes and subsequent default.

Lastly, this paper contributes to the growing literature on the role of beliefs in explaining consumers' spending-saving decisions (see DellaVigna (2009) and Benjamin (2019) for a review). In earlier work, Ameriks et al. (2016), Ameriks et al. (2020), and Ameriks et al. (2020) document the role of survey-elicited beliefs on retirement choices. Manski (2004), Ameriks et al. (2020), Giglio et al. (2021), and Beutel and Weber (2022) study the relationship between beliefs and stock-market investments. Bucks and Pence (2008), Bailey et al. (2019), and Kuchler et al. (2022) analyze how beliefs affect mortgage leverage choices. Our work builds on this literature by exploring a series of quantitative surveys matched to transaction-level data on consumer spending-saving decisions. D'Acunto et al. (2021b, 2022) and Coibion et al. (2022) assess the effect of macroeconomic expectations, namely inflation expectations, on households' consumption, saving, and borrowing choices. In this paper, we provide direct evidence on the role of income expectations and their departure from rational expectations in driving consumption, borrowing, and default outcomes.

[3]See Attanasio and Weber (2010) and Jappelli and Pistaferri (2010) for a review before 2010.)

II Theoretical Framework

Before describing our empirical setting and results, we propose a simple theoretical framework to clarify the assumptions that, in a partial equilibrium intertemporal consumption optimization problem with extrapolative expectations, produces the three predictions we test in the data.

A. Intertemporal Consumption Optimization with Extrapolative Expectations

We consider the standard lifetime consumption optimization problem of a consumer whose preferences display constant absolute risk aversion (CARA) with coefficient of risk aversion γ subject to a standard budget constraint. In each period, agents' asset holding a depends on the previous period's savings, that is, the difference between past asset holding and consumption c, and income y:

$$\max_{c_t} E_0 \left[-\frac{1}{\gamma} \sum_{t=0}^{\infty} \beta^t e^{-\gamma c_t} \right],$$

where β reflects the rate of time preference.

$$s.t. \;\; a_{t+1} = R(a_t - c_t + y_t).$$

Moreover, we assume that income evolves as an AR(1) process that follows:

$$y_{t+1} = y_t + \epsilon_{t+1},$$

where ϵ_{t+1} is an innovation that follows a standard normal distribution, i.e. $N(0, \sigma^2)$. A neoclassical agent would form full-information rational expectations over this income process.

In this model instead, we consider a diagnostic-expectations agent, who overweights the most recent news about income because such news provide information about states of the world that have objectively become more likely after the innovation. This intuition captures the *kernel of truth* feature of the representativeness principle, which microfounds diagnostic expectations in settings such as Bordalo et al. (2018). We consider this specific form of extrapolative expectations in our framework because it has been shown portable across several economic settings and is consistent with evidence from the laboratory.

Specifically, our framework assumes that the representative consumer puts a positive weight ($\theta > 0$) on contemporaneous income shocks when forming expectations about future income, which implies the following form of income expectations[4]:

$$y_{\theta,t+1} = y_t + \theta\epsilon_t. \tag{1}$$

The first order conditions with diagnostic expectations imply a closed-form optimal value of current consumption that differs from its neoclassical counterpart: the consumption rule is augmented so that it also depends on current unexpected income shocks (news relative to past expectations, ϵ_t):

$$c_t = \frac{r}{R}a_t + y_t - \frac{\log(\beta R)}{r\gamma} - \frac{\gamma\sigma^2}{2r} + \frac{\theta}{r}\epsilon_t. \tag{2}$$

Due to the additional last term relative to the standard neoclassical consumption rule, the current consumption of an agent with diagnostic expectations will be higher than a neoclassical agent's if current unexpected income shocks are positive ($\epsilon_t > 0$), and lower otherwise, which leads to the following proposition:

Prediction 1 (Consumption). *For a diagnostic-expectation agent, current unexpected positive income shocks increase current consumption relative to a neoclassical agent, whereas current unexpected negative income shocks reduce it. These differences increase with the size of the shock.*

In this setting, agents can borrow to finance current consumption at the expense of future wealth. Optimal consumption rule implies

[4]The detailed derivation of the functional form based on diagnostic expectation follows directly from Bordalo et al. (2018).

$$a_{t+1} - a_t = R\left(\frac{\log(\beta R)}{r\gamma} + \frac{\gamma\sigma^2}{2r} - \frac{\theta}{r}\epsilon_t\right).$$

Therefore, for a diagnostic-expectation agent, higher current consumption due to over-optimistic beliefs implies lower wealth, thus weakly higher debt, in the subsequent period.

We also assess the probability that a diagnostic-expectations agent default on her debt relative to the neoclassical counterpart. Given the standard budget constraint, we can define the probability of default of the representative consumer as the probability that the growth of wealth in the next period falls below an exogenous threshold (ψ)[5], which is distributed according to a standard normal distribution function:

$$\begin{aligned} p_{d,t+1} &= Pr\left(a_{t+1} + y_{t+1} < \psi\right) \\ &= \Phi\left(\frac{\psi - R(a_t - c_t + y_t) - y_t}{\sigma}\right). \end{aligned}$$

Because current consumption, c_t, is higher after unexpected positive income shocks than for a neoclassical agent, future wealth is lower and hence the likelihood that the sum of future wealth and future income falls below the default threshold is higher than for a neoclassical agent, ceteris paribus.[6] Note that the larger is the current unexpected positive income shock, the higher is the probability of default.

Prediction 2 (Debt and Default). *A diagnostic-expectation agent incurs more debt in the subsequent period after unexpected positive income shocks and faces a higher probability of default. These differences increase with the size of the shock.*

[5]Our baseline framework does not allow for strategic default. We assume a consumer's default choice follows a threshold rule over the amount of savings, ceteris paribus. This assumption is motivated by the micro-foundation of the previous literature that studies household debt and default behaviors (e.g., Livshits et al. (2007)). When the value of choosing to default does not change with the amount of debt and the value of repaying is decreasing in the debt level, the default decision follows a simple threshold rule over total wealth.

[6]The opposite is true if the diagnostic-expectations agent faces an unexpected negative income shock.

B. Introducing Time-Varying Income Volatility

So far, we have compared the consumption and borrowing choices of a diagnostic-expectations consumer relative to a neoclassical consumer in a setting in which income followed a simple AR(1) process. Because all the differences between these two agents boil down to the assumptions we make about the dynamics of income, and hence income expectations, assessing the role of diagnostic expectations in a setting in which income follows a law of motion that is closer to what earlier research has documented in observational data is important.

Time-varying volatility of individual-level life-cycle income is an important feature that earlier research based on micro-level data has documented (for instance, see Guvenen and Smith (2014); Fagereng et al. (2018); Chang et al. (2021)). Incorporating this feature in our simple setting provides richer predictions on the differences between diagnostic-expectations and neoclassical agents' consumption choices. To see this, assume that income follows a process in which volatility increases with the size of the current income shock, which is distributed according to a standard normal distribution:

$$
\begin{aligned}
y_{t+1} &= y_t + \sigma_{t+1}\epsilon_{t+1} \\
\epsilon_{t+1} &\sim N(0,1) \\
\sigma^2_{t+1} &= \alpha_0 + \alpha_1(\sigma_t\epsilon_t)^2,
\end{aligned}
$$

with the restriction that[7]

$$\omega_t = 1 + \frac{\alpha_1}{\alpha_0}(\sigma_t\epsilon_t)^2 < \left(1 + \frac{1}{\theta}\right)^{\frac{1}{2}}.$$

Then the expectations of future income of a diagnostic-expectations agent follows

[7]This condition ensures that the variance does not increase excessively, and thus diagnostic expectations are normalizable. This condition always holds in the limit of rational expectations. See Bordalo et al. (2018) for more details.

$$
\begin{aligned}
E_t[y_{\theta,t+1}] &= E_t[y_{t+1}] + \tilde{\theta}\left(y_t - E_{t-1}[y_t]\right) \\
\sigma^2_{\theta,t+1} &= \sigma^2_{t+1}\frac{1}{1-\theta(\omega_t^2-1)} \\
\tilde{\theta} &= \theta\frac{\omega_t^2}{1-\theta(\omega_t^2-1)}.
\end{aligned}
$$

We can thus write the dynamics of income for a diagnostic-expectations agent as follows:

$$
y_{\theta,t+1} = y_t + \tilde{\theta}\epsilon_t + \epsilon_{\theta,t+1},
$$

and the first order conditions from the Euler equation imply the following closed-form expression for optimal current consumption:

$$
c_t = \frac{r}{R}a_t + y_t - \frac{\log(\beta R)}{r\gamma} - \frac{\gamma\sigma^2_{\theta,t+1}}{2r} + \frac{\tilde{\theta}}{r}\epsilon_t. \tag{3}
$$

We can see that in this case higher current-income volatility, σ_t^2, increases both $\sigma^2_{\theta,t+1}$ and $\tilde{\theta}$ and hence negative realizations of ϵ_t reduce optimal current consumption by more than same-sized positive realizations increase it.

Prediction 3 (Asymmetry). *When facing time-varying income volatility, a diagnostic-expectations agent cuts current consumption after negative unexpected income shocks by more than she increases current consumption after same-sized positive unexpected income shocks.*

III Institutional Setting and Data

We collaborate with a large Chinese commercial bank that operate across the whole country to obtain transaction-level information on a large representative sample of consumers. We also field a customized survey on these customers to elicit a set of individual-level economic expectations through multiple survey waves. The bank operates nationally and is among the top 10 commercial banks in China by total assets. In 2023, the bank's total assets amounted to more than one trillion dollars, and the number of active account holders were over 70 million. Because of the broad customer base of this bank, the random sample for which we obtain data and elicit expectations is representative of the Chinese banked population.

We obtain debt data at the consumer level from the Credit Reference Center of the People's Bank of China (China's official credit registry), based on the reference reports retrieved by the bank. The Credit Reference Center aggregates personal credit information from all financial institutions. This feature is crucial for our analysis because we can observe the amount of debt raised in each period by the consumer across all the financial institutions she can access and not just the specific bank with which we cooperate.

A. Primary Banking Institution: Sample Restrictions

The credit registry does not collect spending information. To study consumption choices, we thus need to rely on transaction-level data within the bank. This feature of our design raises the concern that consumers might have multiple bank relationships and multiple spending accounts. In this case, we would only observe a fraction of consumers' overall spending. To tackle this concern, we follow recent research that also uses single-provider transaction-level data (e.g., see Ganong and Noel (2019)) and impose two restrictions on the accounts that enter our empirical analysis to capture consumers who are most likely to use the bank with which we collaborate as their primary banking institution.

First, we only consider consumers whose bank accounts include at least 15 outflow transactions during the sample period. An outflow is any debit from a checking, saving, or credit card account, including a cash withdrawal, an electronic payment, or a card transaction. Imposing this criterion reduces the original sample by approximately 35%. The second restriction we impose is that consumers' income can be identified and calculated directly by observing regular inflows to the bank's checking accounts. This restriction leads to a drop of about 10% of the observations in the overall sample.

B. Measuring Income and Spending

For our working sample, the transaction-level data allows direct measurement of income inflows and spending outflows. In terms of income, we follow the steps the bank with which we cooperate uses, which identify individual income following a classification rule of regular inflows. The bank classifies income into two main categories: salary and business cash flows.

Salary is defined as the regular monthly income flow and bonuses if the customer declares working as an employee. The bank calculates this number in one of two alternatives ways. First, if income is paid as a direct deposit from the consumers' employers into the bank account, the number is directly labeled as salary in the bank's system. Otherwise, the bank can identify monthly income if the consumer's social security insurance, which is a fixed portion of the consumer's income, is paid through the bank.[8] Removing customers whose income cannot be identified and computed with certainty reduces the sample by 10.3%. As for income from business operations, it is calculated as the difference between total inflows and total outflows of transactions that are categorized as business operations.

When aggregating all incomes in our sample, the split of the three components is 70.48% from salary and 29.52% from business operations. We can directly verify that these figures are not only representative of the Chinese banked population but also accurately computed at the individual level by matching the income computations at the consumer-year level from the bank to individual-level data from the Chinese tax administrative agency. We report the results of this comparison in Panel A of Figure A.1. The figure is a bin scatter plot that compares the income computed by the bank based on transaction-level data based on the procedures described above and the income the same individuals

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

report to the Chinese tax authority. We see a strong linear relationship between these two values, with a R^2 of 78.26%, which corroborates the quality of our income data and the fact that for the primary bank user sample on which we focus, the bank does not appear to miss systematic sources of income that are instead declared by users in their tax filings.

Moving on to the measurement of spending, we compute monthly total spending as the sum of all nondurable purchasing transactions from consumers' checking account plus the total amount of repayment of linked credit cards' end-of-month balances between the end of the last billing cycle and the current billing cycle. Debt is the sum of outstanding interest-incurring balances on all credit cards and other unsecured personal loans that, as we discussed above, we observe across all financial institutions with which consumers have a relationship through the credit registry data.

C. Eliciting Income Expectations: Incentives and Data Quality

To elicit consumers' expectations, we designed a short survey that the bank administered to the consumers in our sample. To avoid consumers' cognitive overload and disaffection due to a time-intensive request, the bank limited the number of questions we could pose to a total of 12 questions[9]. We report our own English translation of the full survey (which was administered in mandarin) in Appendix *B.*.

The survey starts with indicating its purpose. On top of guaranteeing that survey participants are fully informed about the aims of the study, this step also aims to avoid that participants develop a strategic motive when answering. Indeed, respondents might incorrectly infer that their answers would affect the types and quantities of financial services the bank would offer them going forward. To reach both aims, the survey starts by stating that the results will be used for academic research purposes. Specifically, the participants are shown

[9]The whole sample receives nine questions. For the 29% participants who have transactions with financial accounts over the six months before collecting the surveys, an additional three questions about financial market returns.

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' personal information in any respect. We will not, to any extent, change the types of financial products we provide, including credit scores, credit limits, deposit rates, etc., based on the participants' personal answers. Therefore, please answer based on your true thoughts.

This explicit framing was designed to minimize the possibility that consumers provide answers that depart from their true beliefs in the hope of obtaining better services from the bank if they provided distorted expectations.

Moving on to the questions, respondents are first asked to report their average income over the previous six months. Because we observe income inflows in the data and we restrict the sample to primary bank users, we can compare the answer to this question to the actual income flows in the respondent's account, which serves as a data quality check: if we were concerned that many respondents answered our questions at random to finish quickly and/or to provide false information on purpose, we would be able to detect this behavior from the income question. Panel B of Figure A.1 is a binned scatter plot that compares reported income values in the survey with the same respondents' income flows computed from the bank's account-level administrative data. The plot documents a strong linear relationship. A regression between the two variables yields an R^2 of around 72.56%. This evidence corroborates the validity and reliability of the survey's answers.

As far as expectations about individual economic outcomes are concerned, we followed state-of-the-art elicitation methodologies to elicit not only the first moment of income expectations but also a full-blown subjective beliefs distribution (Manski, 2004). To obtain a first-moment point estimate, we ask:

What would your average monthly earnings most likely be in the next 6 months?

To elicit the full-blown beliefs distribution, asking consumers to report a probability distribution directly is highly cognitive demanding and faces the concern that most consumers might not have enough numerical literacy to understand what is a probability distribution, which would invalidate the exercise by confounding actual beliefs with a measure of respondents' cognitive abilities, which in turn shape beliefs (D'Acunto et al., 2019, 2021b).

To tackle this concern, we rely on the triangular-distribution question design that was recently proposed in economics research (for instance, see Guiso et al. (2002); Christelis et al. (2020)). This design consists of asking respondents for a point estimate of the numerical value of the expected minimum of a variable and the expected maximum. This design allows us to compute the second moment of individual-level subjective expectations after imposing the assumption that the distribution of beliefs is symmetric around the midpoint between the minimum and maximum possible expect values.[10]

We elicit these value with the following two questions:

What would be the lowest possible level of average monthly earnings you believe you would get over the next 6 months?

What would be the highest possible level of average monthly income you believe you would get over the next 6 months?

Moreover, the survey includes a similar set of three questions that elicit respondents' expectations about the future credit limits the bank is likely to set on their credit cards to assess directly respondents' beliefs about potential changes in the supply-side of debt, which would affect respondents' spending and borrowing choices above and beyond changes in their own income.

The survey is filed in three rounds. Each round contains two waves sent to the same participants. The first set of two waves were sent in January 2020 and July 2020, the second set were sent in January 2021 and July 2021, and the last set were sent in January 2023 and July 2023. Overall, the surveys cover consumers' expectation over at least four six-month period. We supplement each of the two rounds of surveys with administrative and consumer bank-account data over the periods before and after the surveys.

Because our administrative and survey-based data cover the same individuals across several time periods, we can exploit variation in income expectations and actual realizations and economic choices within individuals over time. Hence, we can

[10]Eliciting also a subjective probability that the outcome falls above the midpoint typically results in identical first and second moments, see, Coibion et al. (2023).

absorb systematic time-invariant unobserved characteristics across individuals that might confound the relationship between income expectations and consumption/debt choices, such as cognitive abilities and financial literacy. To ensure that we can control for individual fixed effects, we include individuals that have completed at least three of the four waves of surveys.

Moreover, by observing cross-sections of respondents across multiple time periods, we can assess our baseline predictions within time periods. This allows us to absorb the effects of aggregate economic shocks that all consumers faced at the same point in time. This feature is important in our setting given that the early sample spans the times before and after the start of the COVID-19 pandemic as well as periods in which the pandemic had its utmost negative effect on economic outcomes and, possibly, on beliefs (January to June 2020).

D. Summary Statistics

Table I provides summary statistics for our sample. Panel A summarizes consumers' demographic characteristics and panel B the same consumers' expectations as elicited from the survey. All variables are converted to US dollars for ease of interpretation and are winsorized at 1-99% levels to reduce the influence of potentially extreme outliers.

In terms of demographics, the age distribution of the consumers in our sample is symmetric around its mean (about 38 years old) and most consumers are in their active working age—the interquantile range is between 29 and 48 years. Moreover, the gender distribution of the sample is quite balanced and includes 52% women and 48% men. These demographic characteristics of the sample dismiss a common concern with transaction-level banking samples that tend to oversample male and younger consumers. This balancing is important for the external validity of our analysis given that men and women might differ systematically in the extent to which their expectations depart from subsequent realizations, for instance due to the higher prevalence of overconfident beliefs in men documented in other settings or relying on alternative information sources to form expectations (D'Acunto et al., 2021).

Moving on to the actual administrative data on spending, income, debt, and saving based on bank accounts, we see that the average monthly income is around $2067.28, whereas the average monthly spending is around $1,245. Both distributions are right skewed, even though the spending distribution is more skewed. This can be seen by the fact that the median income is about $1,227, whereas the median spending amount is only about $762. Consumers have accumulated on average $17818 in savings at the time of the first round of surveys, but even in this case a fat right tail emerges—the median consumer has only accumulated about $4,261 in savings.

On the debt side, we see that the average outstanding interest-incurring unsecured debt is around $1,000. This figure masks substantial heterogeneity given that the median consumer, in fact, has no interest-bearing unsecured debt outstanding. And, indeed, conditional on holding a positive amount of debt, the average consumer had accumulated about $2,300 in debt before the survey was administered. A simple calculation indicates that around 43% of the consumers in the sample held positive credit card debt. This figure is similar to the range of 40% to 60% found in the previous literature using data on US consumers (Gross and Souleles, 2002; Zinman, 2009; Fulford, 2015). Note that the amounts of debt consumers accumulate are, on average, substantially lower than the credit limits they are assigned by all banks (on average about $12,000) and hence the maximum amounts of debt they could raise. For this reason, most consumers in our setting have substantial untapped debt capacity.

The bottom part of Table I reports statistics for the elicited point estimates and ranges of consumers' expectations about their income and updated credit limit over the following six months. We can see that both the average and median expected income are higher than the average and median incomes measured based on administrative bank data for the six month period before the survey was administered. We can see this point more directly from the distribution of the difference between consumers' income expectation over the subsequent six months and the realized value: the average and median values are respectively $231 and $381. Reported expected future credit limits align with consumers' income expectations in that they are also higher than the existing credit limits at the time consumers are surveyed.

E. Sample Representativeness

Since the sample collection relies on surveys, the analysis might cast doubt on how representative the final sample is to the whole Chinese economy due to selectively responding to the survey. Table A.5 in the online appendix compares the demographics between our sample and the a 5% random sample from the bank database. Overall, our sample is very close to the whole sample at the bank. Meanwhile, the surveyed participants are a little younger with slightly lower income than the average consumer. Specifically, the age at our sample is 39, while it is 41 for the average consumer. The income of surveyed individuals is around 10% fewer.

The reason for a very close distribution between the surveyed participants and the average consumer is the high reward rate. On average, the survey takes around 4 minutes, with a reward of 10 CNY. This is higher than the 95-th percentile of the hourly wage rate in China, which is around 300 thousand CNY. In the end, the response rate is 68%.

IV Unexpected Shocks and Expectations

The main empirical challenge the econometrician faces when trying to bring the three predictions we propose to the data is the combination of consumer-level income shocks in the current period as well as subjective expectation errors in the next period. That is, we need to study how *objective* income shocks affect *subjective* forecast errors. Such information typically cannot be measured in standard observational data that do not include information on income expectations.

In our setting, we address this issue directly by first measuring consumers' quantitative expectations of future income through our survey waves and observing their ex-post income realizations based on transaction-level data. In this way, at each point in time we have continuous measures of two crucial dimensions at the individual consumer level: (i) the realized (and observed by consumers) income changes between the previous period and the current period (ΔY_t), and (ii) the deviation between current expectations of future income and actual subsequent realizations of future income ($E_C[Y_{t+1}] - Y_{t+1}$).

Our estimates of income shock in period t is based on the following specification:

$$Y_{i,t+1} = \rho_{j,k,a} Y_{i,t} + \Gamma X_{i,t} + \epsilon_{i,t+1}. \tag{4}$$

The estimation strategy is similar to the estimation of expected and unexpected tax refunds in Baugh et al. (2021). In equation (4), $Y_{i,t+1}$ is consumer i's income in period $t+1$, $X_{i,t}$ is a set of consumer demographics that includes age, age-squared, educational attainment, gender, the log of savings in the previous period, the log of the credit limit in the previous period, city fixed effects, and industry × period fixed effects. The period is defined as a half year to be consistent with the design of the survey. $\rho_{j,k,a}$ is the persistence of income at industry-city-age quintile level. We use the residual $\epsilon_{i,t+1}$ as our measure of objective income shocks in period $t+1$.

We first compare the objective income shocks derived from equation (4) with the subjective income shocks we compute from the survey answers, which are the difference between a consumer's expected average income over the subsequent six months and her actual average income over the subsequent six months. Figure 2 is a binned scatter plot that plots the objective income shocks against the subjective ones. Panel A is the raw measure, while Panel B uses the measures residualized by consumer demographics. The plots show a linear relationship between the two measures. For the unresidualized measure, the R^2 is 0.25, indicating a positive but far from perfect correlation between the two measures.

Overall, although consumers appear to form expectations about income that are on average positively correlated with actual future income shocks, the correlation is far from 1 and there is scope for systematic deviations of expectations from ex-post realized outcomes, which we investigate further in the next subsection.

A. Extrapolative Income Expectations

The measures of realized income, subjective income, and objective income dynamics allow us to test directly whether unexpected current income shocks determine a larger deviation of future income expectations from actual future income realizations. Hence, we can test whether, on average, the income expectations of the consumers in our data are extrapolative like those of diagnostic-expectations agents. Running this test before moving to the predictions about economic decisions (consumption and debt) is important to assess whether our diagnostic-expectations interpretation has any scope in the data and hence whether the predictions about consumption and debt we derived in our simple theoretical framework can be brought to the data.

We start by presenting motivating evidence about systematic mistakes in consumers' income expectations (*income misbeliefs*). Panel A of Figure 1 depicts a binned scatter plot of ex-post realized income against ex-ante income expectations as elicited through the survey. Both dimensions are positively associated, which indicates that consumers' forecasts go on average in the same direction as realized values. However, Panel B plots the forecast errors. The label $E_C[Y_{t+1}]$ indicates that the expectation is subjective and is measured from the perspective of the consumers. Figure 1 shows that, despite a linear relationship between forecasted income and realized income, the distribution of forecast errors is wide: the standard deviation of forecast errors is about one third of the size of the standard deviation of the distribution of income.

We continue by assessing the relationship between current-period unexpected income and misbeliefs about next-period income. Figure 3 reports a binned-scatter plots of forecasts errors and unexpected income. In each of the four panels, the y-axes report subjective forecast errors measured using next-period income realizations. In Panel A and Panel B, the x-axes report the objective income shocks in the current period as retrieved from equation (4). In Panel C and Panel D, the x-axes report the subjective unexpected income shock in the current period. Consistent with Proposition 1, the plots show a positive relationship between income shocks and income forecast errors irrespective of whether income shocks are measured objectively based on the specification in equation (4) or subjectively based on survey answers.

To move one step forward relative to the raw data, we estimate the following specification:

$$E_C[Y_{i,t+1}] - Y_{i,t+1} = \beta(Y_{i,t} - E[Y_{i,t}]) + X'_{i,t}\delta + \boldsymbol{\eta} + \nu^Y_{i,t}, \tag{5}$$

where X is a vector of individual-level characteristics that includes age and its square, educational attainment, female indicator, the log number of weekly hours worked in period t, the logarithms of monthly income and credit-card limits in period $t-1$, which proxies for consumer's debt capacity, consumer expected income changes from period $t-1$ to t, and different sets of fixed effects ($\boldsymbol{\eta}$).

The left-hand-side variable in equation (5), $E_C[Y_{i,t+1}]-Y_{i,t+1}$, measures the subjective expectation error at time $t+1$. On the right-hand-side, $Y_{i,t} - E[Y_{i,t}] = \epsilon_{i,t}$ measures the income shocks at time t. Our coefficient of interest is β, which estimates the marginal relationship between current-period income shock and next-period expectation error. As predicted by Proposition 1, β should be zero for an agent with rational expectations but positive for an agent with extrapolative expectations.

Table II reports the results for estimating equation (5). Column (1) only includes the unexpected income changes in period t as the right-hand-side variable. Column (2) adds in individual characteristics X as controls. Column (3) controls for the city $\times$ year fixed effects and industry fixed effects. In the end, column (4) further adds individual fixed effects. Across all columns, $\hat{\beta}$ is significantly larger than zero and quite stable regardless of the specific set of characteristics that are kept constant.

Focusing on column (4), the inclusion city $\times$ year fixed effects absorbs shocks in a year that might induce a structural change in the income processes of all consumers in the same city. At the same time, the inclusion of individual fixed effects absorbs unobserved individual-level characteristics that might induce negative auto-correlation of subjective expectations across periods. As a result, our specification compares how larger income shocks would induce the same individual to form larger expectation errors for future outcomes. The estimate of $\hat{\beta}$ in column (4) is 0.4, which means that for a one-dollar unexpected income change in the current period the average consumer over-estimates her income in the next period by about 40 cents. Mapped into equation (1), this estimate implies a θ of 0.4 in the framework presented in Section II.

Note that in equation (5), we regress future forecast errors on current income innovations. Because we observe subjective income errors for the same consumers over more than two periods, we can also study the relationship between future forecast errors and current subjective income shocks. We report the results in Table A.1 in the Online Appendix. The results of using subjective income shocks at time t are very close to those we obtained when using the objective income innovation $\epsilon_{i,t}$.[11]

The results so far suggest that the expectations about future income of the average consumer in our sample overreact to shocks to current income, which is consistent with the notion of diagnostic expectations. If these patterns were attributable to diagnostic expectations, though, we would also observe substantial heterogeneity in the association between the size of observed income news and the inaccuracy of expectations about future income across consumers. In particular, consumers who face more volatile incomes should overreact more to unexpected income shocks relative to others, because they are more likely to observe larger unexpected shocks.

To assess this potential source of heterogeneity in the data, we consider four proxies for consumers' income volatility – the actual implied standard deviation of the logarithm of expected income growth; whether the consumer belongs to the bottom half of the income distribution, which based on earlier research faces more volatile income realizations (Fermand et al., 2023); consumers' age, because incomes tend to be more volatile among younger individuals; and consumers' educational attainment, because incomes tend to be more volatile among non-college-educated individuals.

Columns (1) to (4) of Table III reports the results for estimating equation (5) in a form that includes interactions with our four proxies for income volatility. The results show that the extent of overreaction to unexpected news about income appears systematically lower for consumers whose incomes are less volatile: consumers above the median of the income distribution overreact less than others by about half; those whose implied expected income growth is higher overreact by more; and, older consumers and college-educated consumers overreact by less.

We continue to study the effects of income shocks on misbeliefs separately for the two rounds of the survey, which allows us to assess if the extent of extrapolation varies across the business cycle. The first wave, which happened at the onset of the COVID

[11] If the data generating process and expectation process indeed follow the assumptions in Section II, *beta* estimated using subjective income shocks would be biased due to serial correlation between expectation errors. However, the bias when θ is around 0.4, as estimated in Table II, is very small.

crisis, is characterized by heightened uncertainty about future economic activities and incomes given the shutdown of multiple types of economic activities.[12] At the time of the second and the third waves, instead, China had mostly moved back to relatively normal business-cycle conditions. We can thus use results from the two rounds of surveys to assess extrapolative attitudes for income expectations during both recessions and expansions. The results are shown in column (5) of Table III. In general, we observe over-extrapolation across both rounds of surveys. At the same time, there is a significant heterogeneity across COVID and non-COVID periods. The degree of extrapolation is larger during the first round, indicating a larger degree of extrapolation during periods with heightened aggregate uncertainty.

Overall, the expectations-formation process of the average consumer in our setting, which we can observe directly rather than through revealed choices, is extrapolative and appears consistent with the diagnostic-expectations framework based on which we have obtained predictions for consumption and debt choices as well as subsequent likelihoods of default.

V Extrapolative Expectations and Consumption, Debt, Defaults

In this section, we move forward to bring to the data the predictions of our theoretical framework regarding the role of unexpected income shocks on consumption and debt choices.

Our first prediction relates to current consumption choices: because diagnostic-expectations agents overestimate future income when facing positive unexpected income shocks, we should expect higher current consumption induced by excess expected income, and vice versa. We estimate variations of the following linear specification:

$$\Delta C_{i,t} = \gamma(E[Y_{i,t+1}] - Y_{i,t+1}) + X'_{i,t}\delta + \boldsymbol{\eta} + \nu^C_{i,t}, \tag{6}$$

[12] In China, COVID induced a nation-wide lockdown starting at the end of January 2020. However, most areas turned to relatively normal conditions around late February 2020 to early March 2020. Wuhan was the latest city for which the lockdown policy was removed, which happened on April 8, 2020.

where $\Delta C_{i,t}$ is the change in total non-durable consumption at time t relative to the previous period $(t-1)$, and all other variables are defined as in equation (5). We report the estimates in Table IV. Column (1) considers a baseline univariate specification. In this case, we can see that consumers who have income expectations $1,000 higher than future actual realizations increase their current consumption by about $208 dollars more than other consumers. The statistical significance and size of this association is similar once, in column (2), we keep constant the characteristics we observe and a saturated set of fixed effects. Adding these controls increases the explained variation of the individual current change in spending from an R^2 of 3.27% to about 54.76%, and yet the estimated association between excess income expectations and change in spending barely changes.

Because expectation errors increase current consumption but not future realized income, positive forecast errors induce consumption to deviate from the optimal path. The positive response in current consumption paired with negative surprises in future income suggests that debt should increase when expectations errors are positive. We thus also assess if excessive income expectations predict an increase in borrowing in the next period relative to the current period. Note that, ex ante, this association could be zero if agents were liquidity and financially unconstrained and could finance the full amount of current spending increase with available cash. For this test, we estimate a version of equation (6) in which the left-hand-side variable is $\Delta B_{i,t+1}$, i.e. the difference between the average outstanding interest-bearing unsecured debt over the six months between period $t+1$ and period t.

Table IV reveals that the larger is the difference between income expectations and actual ex-post income realization, the higher is the increase in the unsecured debt consumers raise, and the estimated magnitude is similar irrespective of whether we consider the univariate association (column (3)) or we include the full set of observables in our analysis (column (4)). The estimate indicates that, for each dollar higher misbelief in the average monthly income over the following period, average monthly unsecured debt in the same period increases by about 7.5 cents.

As we discussed above, these results capture the debt raised across all possible sources covered in the credit registry and not just the bank with which we cooperate. By contrast, consumption data is based on the accounts at that bank. To dismiss this concern, on top of the screening filters we use to select consumers for which the bank we observe is the primary banking institution, we study the relationship between expectation errors and the changes in cash withdrawal/ net transfers between our bank and external bank accounts. The results are in Online Appendix Table A.5. We find that in our complete specifications no significant relationship between expectation errors and cash withdrawals or net transfers exists, which dismisses the concern that our results on consumption are driven by balance shifting between observed and unobserved bank accounts.

We then consider the likelihood of default on unsecured personal loans. This outcome is also important to study because if consumers were increasing their debt due to extrapolative income expectations but repay such debt fully, extrapolative expectations would still represent a microfoundation for household debt cycles but would not necessarily be worrisome from a policy perspective.

To assess the relationship between misbelief and the likelihood of default, we use a 90-day delinquency indicator in period $t+1$ as the default event, and re-estimate equation (5) using the default indicator as the left-hand-side variable. Default is multiplied by 100 and forecast errors are multiplied by 1,000 for easier interpretation. Note that, because by construction we are limited to observing outcomes up to time $t+1$ but defaults might happen at any subsequent time until the loan is required to be repaid, our estimates might represent a lower bound of the actual size of the relationship between excessive income expectations and the likelihood of future default.

Columns (5) - (6) of Table IV document that a higher distance between consumers' income expectations and ex-post realizations is associated with a higher probability of default. From column (6), for each $1,000 higher mistake in income expectations, default is about 0.92-percentage-point higher. This magnitude is large if we consider that the average default rate in our sample during our sample period is only 2.4%, and hence a $1,000 higher mistake in income expectations leads to a 38% higher likelihood of default relative to the sample mean. The findings are consistent with the aggregate dynamics such that credit-market sentiment leads financial fragility (López-Salido et al., 2017)

Overreaction in income expectations increase spending in the short run. However, in the long run, once consumers observe subsequent income realizations and note that they are lower than expected, debt and consumption amounts should revert back. We assess this possibility by considering the relationship between cumulative spending and income forecast errors at different horizons. Specifically, we consider the following specification:

$$C_{i,\tau+k} - C_{i,\tau-1} = \alpha + \beta_k(E_C[Y_{i,t+1}] - Y_{i,t+1}) + X'_{i,t}\delta + \boldsymbol{\eta} + \nu^C_{i,\tau+k}. \tag{7}$$

In equation (7), τ is the first quarter of period t. Recall that t is a time period that covers six months. Therefore, in equation (6), period t includes the quarters τ and $\tau+1$, whereas period $t+1$ includes the quarters $\tau+2$ and $\tau+3$. Thus, $C_{i,\tau-1}$ is the average monthly nondurable spending in the quarter before the first survey. $E_C[Y_{i,t+1}]$ is the expected average monthly income during quarters $\tau+2$ and $\tau+3$. Analogously, equation (6) can be written in the form of equation (7) if we replace the left-hand-side variable with $C_{i,\tau+1} + C_{i,\tau+2} - C_{i,\tau-1} - C_{i,\tau-2}$. We fit 12 regressions based on the specification of equation (6) for k ranging from -4 to 8 excluding $k=-1$, which we use as the benchmark consumption. Therefore, equation (7) measures the relationship between expectation errors at time $t+1$ and cumulative spending from quarter $\tau-1$ to quarter $\tau+k$.

The results are reported graphically in Figure 4. Panels A and B respectively plots the results for nondurable spending and borrowing. The red solid lines depict positive expectation errors, and the blue dashed lines negative expectation errors. Consistent with overreaction, consumers first increase their consumption during the two quarters before the time when expectation errors are measured (quarters 2 and 3) as well as during the two quarters when expectation errors are measured. However, afterwards, both consumption and debt start to revert and in fact revert almost fully to the level of three quarters before the time in which expectation errors are measured.

VI Asymmetric Effects: Positive and Negative Unexpected Income Shocks

An important feature of household debt cycles as described in aggregate-level data by earlier research is that these cycles and their correlations with households' spending decisions build up slowly in times of positive income growth (positive domain) but drop quickly in times of negative income growth (negative domain), that is, as soon as a recession hits.

In our theoretical framework, we have shown that diagnostic expectations can rationalize this asymmetric effect of income expectations once we allow for consumers' income volatility to vary over time. Not only should unexpected income shocks have a stronger effect on economic choices in the negative domain, but this effect should largely be driven by consumers who expect more volatile income growth going forward.

Before assessing these predictions, we test in the data whether unexpected income shocks affect consumers' subjective income volatility. We do so by estimating versions of equation (5) in which the outcome variable is the standard deviation of consumers' expected income growth, which we compute based on our survey question under the assumption that income beliefs follow a simple triangular distribution. Because we are considering the second moment of the distribution, we use the absolute value of the unexpected income shocks as the outcome variable.

In Table V, we find that as the size of the unexpected income shock increases, irrespective of its sign, the standard deviation of expected future income growth increases as well. Moreover, to assess whether negative and positive unexpected income shocks of the same size relate to expectations differently, in columns (4), we interact the absolute value of the unexpected income shock with a dummy variable for whether the shock is negative. We find that, indeed, not only the second moment of future income growth is higher after negative income shocks, but the association between the unexpected income shock and income-expectations volatility is higher for negative shocks than for positive shocks of the same size. Specifically, for each $1,000 higher income shock, subjective income volatility is 18% larger for positive shocks but about 35.4% higher for negative shocks.

We now move on to test whether unexpected income shocks in the negative domain have an effect on consumption and debt choices of larger size relative to income shocks in the positive domain. Columns (1) and (3) of Table VI reveal that, indeed, the consumption and debt accumulation response to unexpected income shocks is stronger for negative shocks relative to positive shocks. Note that, for negative shocks, the control variable is negative and hence the shock is larger the lower, rather than higher the variable is. For this reason, a positive coefficient means that the larger is the shock, the more negative is the change in current consumption.

VII Extending the Time Series: Income Innovations and Consumption

Our analysis is based on differences in expectations and consumption within individuals when we add individual fixed effects to our specifications, which rules out that systematic unobserved heterogeneity across consumers might affect at the same time the levels and changes of consumption and expectations. However, our sample includes a relatively short panel. In such cases, the assumption that income innovations average out to zero might be challenged (Chamberlain, 1982; Keane and Runkle, 1998; Souleles, 2004).

To tackle this concern, in robustness analysis we study the relationship between income innovations and spending decisions, for which we observe a much longer time series from the bank-level data, including periods in which no surveys were administered. Specifically, we focus on the same set of consumers that answered the survey waves but extend the bank-level data to the longest time horizon we can observe in the bank-level data, which is on average about 5.2 years (10 or more observations) for each consumer.

The results are shown in Table A.3. The odd columns focus on the same sampling periods as in the main study, while the even columns extend the data to include all the periods we have available in the bank-level data. We can see that the conditional association between income innovations and consumption choices is, though slightly smaller with a longer sample, still quite similar across columns. Therefore, even if not directly observed over a longer period of time, Table A.3 suggests that the relationship between expectations mistakes and consumption choices might also be stable over longer sample period than we have available based on the survey waves we could run on the banks' consumers.

VIII From Reduced Form to Quantitative Analysis: A Structural Estimation

Our analysis so far has produced reduced-form evidence that is consistent with the predictions of our simple theoretical framework in which we microfound extrapolative expectations with diagnostic expectations. To assess whether this microfoundation can provide plausible quantitative predictions on top of the directional results of our reduced-form analysis, in this section we introduce diagnostic expectations with defaultable debt into a standard consumption model with incomplete market and heterogeneous agents and we use the model to study the relationship between unexpected income shocks, expectation errors, and consumption-decision making quantitatively.

A. Income Process and Expectation Formation

In the model, time is discrete and denoted by $t = 1, 2, ...$. A unit mass of consumers exists that are subject to idiosyncratic income risk. For each individual i, the process of income $y_{i,t}$ follows:

$$\begin{aligned} \log y_{i,t} &= \alpha + z_{i,t} + \epsilon_{i,t} \\ z_{i,t} &= \rho z_{i,t-1} + \eta_{i,t}, \end{aligned} \tag{8}$$

where $\epsilon_{i,t}$ and $\eta_{i,t}$ are i.i.d. normal shocks with $\mathbb{E}[e^{\epsilon_{i,t}}] = 1$ and $\mathbb{E}[e^{\eta_{i,t}}] = 1$. The variances of $\epsilon_{i,t}$ and $\eta_{i,t}$ are σ_ϵ^2 and σ_η^2, respectively. α is the life-cycle component, which we assume is constant and common knowledge.

Consumers do not know the true value of $z_{i,t}$ and need to make inferences based on Bayesian learning. The Kalman-filtering problem with respect to the persistent component of $\log y_{i,t}$ here adapts Guvenen (2007) and Bordalo et al. (2019). In each period, consumers observe $y_{i,t}$ and update their beliefs about $z_{i,t}$. Without diagnostic expectation, consumers' current forecasted value of $z_{i,t}$ is normally distributed with variance σ_z^2 and mean

$$\hat{z}_{i,t} = \rho \hat{z}_{i,t-1} + \kappa \left[y_{i,t} - \alpha - \rho \hat{z}_{i,t-1} \right], \tag{9}$$

where κ is the Kalman gain of the learning process. 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 this case,

$$\begin{aligned} \kappa &= \frac{\rho^2 \sigma_z^2 + \sigma_\eta^2}{\rho^2 \sigma_z^2 + \sigma_\eta^2 + \sigma_\epsilon^2}, \\ \sigma_z^2 &= \frac{(1-\kappa)\sigma_\eta^2}{1-(1-\kappa)\rho^2}. \end{aligned} \tag{10}$$

Under diagnostic expectation, consumers overreact to surprises in income realizations. The posterior average of $z_{i,t}$ becomes

$$\hat{z}_{i,t}^{\theta} = \hat{z}_{i,t} + \theta \kappa \left[y_{i,t} - \alpha - \rho \hat{z}_{i,t-1} \right], \tag{11}$$

where $\hat{z}_{i,t}^{\theta}$ is the expectation of $z_{i,t}$ under diagnostic expectations, and θ is the degree of representativeness. When $\theta > 0$, consumers overweight representative states, their beliefs exaggerate the signal-to-noise ratio relative to the standard Kalman filter, inflating the persistent component of income upon receiving good news and deflating those while receiving bad news. Exaggeration of the signal-to-noise ratio is reminiscent of overconfidence, with overreaction to news increases in θ. At $\theta = 0$, the model reduces to rational learning.

B. Consumer Preferences

B.1 Preferences

Household preferences follow the literature on consumer credit and default (e.g. Chatterjee et al. (2007) and Livshits et al. (2007)). Consumers at time t maximize their expected lifetime utility with period-s utility flow of

$$E_{i,t}^{\theta}\left[\beta^{s-t}\frac{c_{i,s}^{1-\gamma}-1}{1-\gamma}-\chi d_{i,s}\right],$$

where the superscript θ indicates that the expectation is taken according to diagnostic expectation with degree of representativeness θ. β is the per-period discount rate, γ is the coefficient of relative risk aversion, and $d_{i,t}=1$ if consumer i chooses to default at the end of period t. When default occurs, consumer i incurs a fixed non-pecuniary utility cost ("stigma") $\chi>0$. In addition, consumers receive a pair of additively separable i.i.d. shocks, $\xi=\{\xi^D,\xi^N\}$, which are attached to the options to default or repay and are drawn from a type one extreme value distribution with scale parameter of one. These shocks aim to capture the fact that many defaults are associated not with income shocks, but with events such as marital disruptions and medical expenses which we do not model explicitly. With these shocks, the model generates a positive probability of default across the whole range of borrowing. In addition, as suggested in Dempsey and Ionescu (2023), the introduction of utility shocks associated with defaulting smoothes out individuals' repayment probability functions, which eases computation of the model.

The budget constraint in each period t is

$$\begin{aligned} a_{i,t+1} &= \begin{cases} (1+r_{i,t})(a_{i,t}-c_{i,t})+y_{i,t+1} & \text{if } d_{i,t}=0 \\ (1-\nu)y_{i,t+1} & \text{if } d_{i,t}=1 \end{cases} \\ a_{i,t} &\geq -l_{i,t}, \end{aligned} \tag{12}$$

where $a_{i,t}$ is the total amount of available resources. $l_{i,t}$ is the credit limit, and $\nu\in[0,1]$ is the marginal rate of garnishment. Equation (12) states that when consumers do not default, their wealth in the next period is the sum of their income and gross saving. When consumers default, their saving becomes zero; at the same time, they need to pay for a garnishment cost equal to ν times their income in the next period.[13]. The interest rate is different for saving and borrowing and takes the value

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

B.2 Optimality Conditions

Consumers' problem is characterized by a set of four state variables $\Theta_{i,t} = (a_{i,t}, \hat{z}_{i,t}, z_{i,t}, \epsilon_{i,t})$. Given the overall state $\theta_{i,t}$, consumer i's value function at time t is

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

The continuation value from defaulting is

$$V_D(\Theta_{i,t}) = \max_{c_{i,t}} \frac{c_{i,t}^{1-\gamma} - 1}{1-\gamma} - \chi + \beta\, \mathbb{E}_{i,t}^{\theta}[V((1-\nu) \times y_{i,t+1}, \hat{z}_{i,t+1}, z_{i,t+1}, \epsilon_{i,t+1})] + \xi_{i,t}^{D}. \tag{14}$$

The continuation value from not defaulting is

$$V_N(\Theta_{i,t}) = \max_{c_{i,t}} \frac{c_{i,t}^{1-\gamma} - 1}{1-\gamma} + \beta\, \mathbb{E}_{i,t}^{\theta}[V(a_{i,t+1}, \hat{z}_{i,t+1}, z_{i,t+1}, \epsilon_{i,t+1}) | I_{i,t}] + \xi_{i,t}^{N}. \tag{15}$$

Given that ξ follows a type one extreme value distribution, the probability of default is

$$pd(\Theta_{i,t}) = \left[1 + \exp\{V_N(\Theta_{i,t}) - V_D(\Theta_{i,t})\}\right]^{-1}. \tag{16}$$

[13]Some previous studies also assume that defaults go hand in hand with a temporary inability of borrowing, i.e. $l_{i,t} = 0$ (Chatterjee et al., 2007; Livshits et al., 2007; Dempsey and Ionescu, 2023) However, for simplicity, we assume that consumers' borrowing capacity does not change when defaulting. This assumption allows us to discard one additional state variable. In addition, Livshits et al. (2007) show that the costs of default from changing borrowing capacities is usually quantitatively not important.

C. Calibration

The estimation consists of two stages. In the first stage, we rely on the income information as retrieved by the bank to pin down parameters associated with the income process, including the degree of extrapolation θ, and the marginal rate of garnishment. In the second stage, we use the simulated method of moments (SMM) to get the estimates of consumers' coefficient of risk aversion γ and non-pecuniary costs of default χ.

First Stage: The first-stage estimation requires the use of income and the survey data. We first set the frequency of time t to six months to be consistent with the survey frequency. For the parameters governing the income process, α, ρ, σ_ϵ^2, and σ_η^2, we residualize all individual income by age quitile, year, education, industry, city, and gender fixed effects, and estimate (8) using maximum likelihood estimation. After residualization, α is set to be 0, and when solving the model, we measure consumption and saving with respective to the level of average income.

The estimation results are in Panel A of Table VII. log income is highly serially correlated, with an AR(1) coefficient of 0.97. The average annualized interest rate on deposit in the data is around 2.8%. We therefore set $r_s = 1.4\%$. We use interest rates on credit card borrowing to determine r_b. The average daily interest rate on credit card debt is 5 basis points. After accounting for 45 days of interest-free period, a 2.5% of reward rate, and 1% of all types of debt-related fees, r_b is set at 5.5% per six months. We set a uniform credit limit to all consumers. We set $l_{i,t}$ to be 1.4, so that the average credit limit to average income in the model matches the average total credit limit over credit cards and other lines of credit across all banks of 73% of the annual income in the data.[14]. For the marginal garnishment rate ν, we directly calculate this number from the bank's database, which is around 50%, roughly corresponding to three months of average income. In the end, we set the discount rate $\beta = 0.95^{0.5}$, so that the annual rate is 0.95.

Second Stage: We use SMM to estimate γ and χ. The targeted moments are the

[14] The level of credit limit to income is larger than the 20%-35% range in many previous papers using SCF data in the US (Kaplan and Violante, 2014; Lee and Maxted, 2023). However, it is very close to the number in a recent report by Experian using administrative data. The report documents that the average credit card limit in 2019 for an average American was around \$31,400, which is around 60% of the average individual income of \$54,129 in 2019, and should be a lower bound of the ratio between total limits over all lines of credit and income. See here and here for the report.

average wealth-consumption ratio and average default rate. The logic is straightforward. The risk-aversion parameter γ captures the curvature of the utility function. Higher risk aversion increases consumer willingness to save, thereby increasing the wealth-to-consumption ratio. The stigma cost χ directly affects consumers' willingness to default. A higher χ indicates a higher cost of default, and therefore a lower default rate.[15]

Panels B and C of Table VII show the estimated parameters and targeted moments. The estimation fits the empirical moments closely. The average wealth to half-year total consumption is 0.818. The default rate is around 2.36%, yielding a γ of around 2.51 and a χ of 24.45.

D. Results

We now discuss the results of the structural model.

D.1 Goodness of Fit

Panels B and C of Table VII indicate the model does a good job in matching the targeted moments. In addition, the model is effective in matching many non-targeted moments, especially the distribution of liquid wealth. From Panel D of Table VII, the average and median wealth to six-month income are 1.351 and 0.731 in the data, respectively. The equivalent numbers in the model are 1.243 and 0.777, respectively. Since liquidity is an important factor affecting the average MPC in the sample, it's necessary for the stationary distribution of saving to be close to that in the data. Panel D of Table VII shows the model can match several moments about the empirical distribution of liquid savings. For example, in the data, about 31.38% of total liquid assets are held by consumers in the the top 5% asset percentile. In the model, this number is 29.85%. In addition, 29.89% of individual hold negative net liquid wealth in the data; in the model, the corresponding number is 32.34%. Therefore, the model is capable of matching both first and higher moments.

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

D.2 Over-Extrapolation of Income Shocks

We now use the model to study the effects of diagnostic expectations on consumer economic decisions. We first present the results of the relationship between objective income shocks and subjective forecast errors under diagnostic expectations with the calibrated level of θ, and that under rational expectations. Specifically, we simulate 100,000 periods of income data following (8), and then construct subjective and objective income expectations based on (11) and (10). We then regress subjective forecast errors at $t+1$ on the objective income shock at t. When preforming the analysis, we drop the first 100 periods, which serve as the burning periods in the simulation.

The results are in columns (1) of panels B and C in Table VIII. For comparison, Panel A gives the empirical counterparts. Given that θ is calibrated using the same sample as in the empirical analysis, the relationship between objective income shocks and subjective forecast errors under diagnostic expectations matches the empirical analysis perfectly. That is, when $\theta = 1.65$, each dollar higher objective income shock leads to a forecast error of future income of 40 cents. The estimate in column (1) of Panel B directly sheds light on the effects of a positive θ on the extrapolative behavior. When θ is set to zero, the relationship between objective income shock and forecast errors become economically insignificant. Specifically, each dollar higher objective income shock leads to a negative forecast error of future income of 4 cents.[16]

D.3 Consumption, Borrowing, and Default Decisions

We continue to explore the relationship between subjective forecast errors and consumer economic decisions in the model. We use this exercise to explore whether the model can generate the patterns we see in the data. To perform the analysis, we simulate the model for 20,000 individuals with 1,000 periods, after a burn-in period of 100. We then drop simulated data with saving to average income ratio larger than 8.13, which is maximum of the empirical counterparts. After obtaining the simulated sample, we run regressions of changes in total consumption at t, total debt at the beginning of period $t+1$, and default

[16] The low statistical significance is a results of high precision of simulated data with a large number of data points.

indicator at the end of $t+1$ on the forecast errors at $t+1$. Debt at the beginning of $t+1$ is defined as the negative of $a_{i,t}$, conditional on $a_{i,t}$ being negative. To be consistent with the analysis in Table IV, we control for log income at $t-1$, expected changes in income at t, and individual fixed effects. We report the results in columns (2) to (4) of Table VIII. For comparison, Panel A reports the empirical counterparts.

The relationship between forecast errors and economic decisions as implied by the model are quite close to those in the data. Specifically, each dollar higher forecast error at $t+1$ raises total consumption by around 20.6 cents in the data and 24.3 cents in the model. Meanwhile, each dollar higher forecast error at $t+1$ raises total debt in the next period by around 7.5 cents in the data and 7.4 cents in the model. For default, since the units are different, we standardize the forecast errors so that the coefficients measure the association of default probability with each standard deviation higher forecast errors. In the data, each standard deviation higher forecast errors increases default by 91.8 basis points, whereas the model implied quantity is 95.2 basis points. In panel C, we see that, when setting θ to zero, the relationship between forecast errors and consumption, debt, and default becomes much smaller. Specifically, when consumers do not have diagnostic expectations, forecast errors do not have economically large impacts on current consumption or debt in the following period. While negative surprises always lead to more default, without diagnostic expectations, the effects reduce by around 60%. Hence, Table VIII shows that the model is able to reproduce the patters we observe in the data.

D.4 Debt Cycles

In a Minsky-Kindleberger style of credit cycle, boom-bust pattern in household leverage starts with household over-optimism after positive shocks to fundamentals. That is, after receiving good news, households exaggerate the informativeness of the good news for future growth, inducing them to take on too much leverage. As a result, excess leverage starts to build up financial fragility. When the effects of positive shocks diminish, households receive a large negative surprise in their earnings, and the ability to repay the debt already accumulated gets lower. Consequently, at the end of expansion, both demand for more debt and financial fragility get accelerated.

We continue to study the model's ability to generate cycles in household borrowing after positive shocks to income. Following Maxted (2023), we use impulse response functions (IRF) to study consumer responses to a series of positive transitory income shocks after initially being at the stochastic steady state. Specifically, we first simulate 20,000 individuals for 1,000 periods at half-year frequency. This serves as our benchmark case. To derive the IRFs, we use the same simulation, but introduce a 3-year sequence of positive income shocks that results in a three standard deviation cumulative shock over the three years from periods 895 to 900. The IRFs are then the (percentage) differences between the sample average between the two simulations. Then, we redo this exercise while setting θ to zero. Comparing the IRFs across the two calibrations yields the effects of diagnostic expectations on economic outcomes.

The results are in Figure 5. The top left panel plots the transitory income shocks. The top right panel plots the updates in expected log income, $o_{i,t} = \kappa(1+\theta)(y_{i,t} - \alpha - \hat{z}_{i,t-1})$. The bottom four panels are the percentage differences in average income expectations, consumption, borrowing, and default probability relative to the no shock simulation, respectively. The red solid lines present results when $\theta = 1.65$; the blue dashed lines present results when $\theta = 0$.

Figure 5 shows a strong amplification effects of diagnostic expectations on expectations of future income. Panel B and Panel C show, that initially after receiving the positive shock, the average diagnostic-expectation agent (D) increases income expectation by more than twice as much as a rational-learning agent (R). Given a much higher expected income in the next period, consumption also increases much more for D. As for debt, Panel A and Panel D show that, when D receives the first positive shock, realized income increases by around 20%, and consumption increases by around 14%. Because we only have one asset in the model, a MPC smaller than one indicates an initial reduction in debt, a finding that is consistent with Agarwal et al. (2007) and Coibion et al. (2020). Turning to Panel F, as compared with R, the initial positive income shock leads to a much lower default probability for D. This is because the higher expectations about income in the next period induces a higher perceived garnishment cost associated with default.

As D continues to receive positive income shocks, income expectations and consumption continue to reach to a level that is much higher for D as compared with R (Panel B, C, and D). At the same time, the trajectories for debt and default rate keep diverging between D and R. For R, as income and income expectation increase smoothly, R continues to de-lever. An increase in current-period assets and expectations about future income decrease the default probability. At the same time, for D, given a much higher future income expectations and a smaller income surprises, D starts to accumulate more debt. From Panel F, a lower level of current-period assets tends to increase the motive for defaulting. Ultimately, the channel outweighs the marginal garnishment channel, and default rates start to increase.

In period 1, the series of positive shocks are removed and as a result income expectations become smaller. Panel B and Panel C show that this reduction is much larger for D as compared with R due to over-extrapolation. Consequently, consumption decreases much more for D. The larger negative surprise induces D to have a much higher need to smooth the negative shock and a much lower income expectation in the future, thus creating a shoot-up in debt. Compared with R, a lower current-period asset holding and lower expected future earnings increase D's default rate substantially. As time elapses, income expectations converge. Both debt and default probabilities start to decrease to the rational-learning level.

The findings in Figure 5 are consistent with many recent empirical findings. For example, López-Salido et al. (2017) show that elevated credit-market sentiment is associated with higher credit growth in subsequent years. As Panel E shows, as initial income shocks create more optimistic belief, the debt growth rate becomes positive for D, whereas it remains negative for R. In addition, the elevated default rate after the expansion (Panel F) is consistent with the findings in Greenwood et al. (2019) and Maxted (2023), who show that financial fragility arises at the end of economic expansions.

D.5 Simulating the 2007-2008 Financial Crisis

As our last exercise, we study the ability of our model to generate the cycles of the unsecured borrowing and the default risk around 2007-2008 financial crisis. Since we are using parameter values estimated using Chinese data, we mainly focus on the dynamics instead of the levels.

The results are shown in Figure 6. The top panel plots the selected shocks. In the middle and bottom panels, the red solid lines and blue dotted lines respectively present the simulation when $\theta = 1.65$ and $\theta = 0$ at half-year frequency; the black dashed lines give the data at annual frequency. On the left panel, $\overline{B}/\overline{Y}$ is the ratio of total debt from credit cards and other credit lines and GDP, divided by labor share. $\overline{B}/\overline{Y}$ in the data is detrended from 2003 to 2012. On the right panel $\overline{p}(default)$ is consumer debt delinquency rate multiplied by the proportion of individuals with positive consumer debt.[17]

There are two notable findings evident from the middle panel. First, aggregate debt level at the equilibrium (before 2014) is much larger with diagnostic expectations than that without diagnostic expectations. This is because, at the micro level, diagnostic expectations widens the cross-sectional distribution of income expectations. That is, in repsonse to the same distribution of income shocks, there are more people having more optimistic expectations and more people having more pessimistic expectations. Since debt is bounded below at zero, more optimistic people end up with more debt, while more pessimistic people can not de-accumulate debt to the negative level. Therefore, micro-level expectation errors do not wash out at the macro level, causing more aggregate debt in equilibrium.

Second, incorporating diagnostic expectation over surprises in income, the model is successful in re-producing the boom-bust cycles in the average unsecured-debt to average income ratio. The pattern matches very closely with the average credit-card debt to GDP ratio in the US. After removing over-extrapolation, the cyclical behavior gets much weaker, with the peak in 2009 never surpassing the equilibrium-level in 2004.

From the bottom panel, while the shocks are reverse engineered to match the dynamics in debt, these shocks can also generate the cyclical patterns in default rate. That is, at the end of the expansion, default rate shoots up in 2009, and then slowly declines. However, removing the extrapolative behaviors, default rate stays nearly constant from

[17]All data excluding the proportion of individuals with positive consumer debt is from FRED. The proportion of individuals with positive consumer debt is calculated from Survey of Consumer Finance. We linear interpolate the numbers between each survey wave.

2004 to 2012. Overall, Figure 6 indicates the effectiveness of our model to reproduce a boom-bust debt cycle like the one around 2007 and 2008.

IX Conclusions

When agents' belief-formation process about their income follows diagnostic-expectation, they have upward biased income expectations after unexpected positive income shocks and hence spend more (and borrow more to finance higher current spending) than what a neoclassical life-cycle consumption optimizer would do. Moreover, if income volatility is time-varying, the effect is stronger for negative shocks than for same-size positive shocks. These predictions align well with aggregate features of household debt cycles documented in the literature.

Combining survey-based elicitation of income expectations with transaction-level bank-account data, we bring these predictions to the field in a unique setting in which we observe, for the same consumers and at the same time, income and debt-capacity expectations as well as actual past, current, and future spending and borrowing choices. We find evidence consistent with these predictions.

Our results suggest that extrapolative expectations could act as a microfoundation for aggregate household debt cycles, because the micro-level evidence pairs with aggregate evidence on household debt cycles: in good times, a larger fraction of consumers is likely to face unexpected positive income shocks, whereas in bad times a larger fraction of consumers is likely to face unexpected negative income shocks, relative to normal times. We propose a framework based on diagnostic expectations but do not argue that this specific formulation is the *only* form of extrapolative expectations that could rationalize our results. We use diagnostic expectations as a beliefs-formation structure that economic research has found portable theoretically and empirically to other features of micro- and macroeconomic outcomes, rather than assessing ad-hoc explanations for the facts we want to interpret. We encourage future research on non-standard belief-formation mechanisms that could at the same time explain household debt cycles as well as other economic choices by consumers in the laboratory and in the field.

Our research also beget follow-up work on the aggregate effects of household debt cycles both theoretically and empirically. Theoretically, it sheds light on the general-equilibrium macroeconomic models featuring diagnostic expectations to explain the mechanisms behind aggregate economic dynamics. On the empirical side, this study suggests the importance of structural and calibration analysis to characterize the functional forms and parameter sizes that would make extrapolative expectations best fit the wealth of micro and macro data. In addition, our structural analysis considers the prediction of agents with the same degree of extrapolative behavior, which pairs with the empirical analysis based on average expectations and choices in the field. The extent to which consumers' expectations-formation process deviates from the full-information rational-expectations paradigm might vary in the cross section of consumers, potentially based on dimensions that would shape the cross-sectional variation in the accuracy of macroeconomic expectations.[18] An interesting future avenue therefore would be studying more advanced heterogeneous macroeconomic models in the direction of heterogeneous belief-formation processes to explain why different consumers form expectations differently, and to explore the quantitative importance of the heterogeneity in the belief-formation processes.

[18]Examples include cognition (D'Acunto et al., 2019, 2021a)), socioeconomic status (Kuhnen and Miu, 2017; Das et al., 2020), and local experiences (Kuchler and Zafar, 2019).

References

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

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.

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

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

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

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

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

Bachmann, R., S. Elstner, and E. R. Sims (2013). Uncertainty and economic activity: Evidence from business survey data. *American Economic Journal: Macroeconomics 5*(2), 217–49.

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

Barberis, N. (2018). Psychology-based models of asset prices and trading volume. In *Handbook of behavioral economics: applications and foundations 1*, Volume 1, pp. 79–175. Elsevier.

Barberis, N., R. Greenwood, L. Jin, and A. Shleifer (2018). Extrapolation and bubbles. *Journal of Financial Economics 129*(2), 203–227.

Baron, M., E. Verner, and W. Xiong (2020, 10). Banking Crises Without Panics*. *The Quarterly Journal of Economics 136*(1), 51–113.

Baron, M. and W. Xiong (2017, 01). 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.

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 (2022). Beliefs and portfolios: Causal evidence. *Chicago Booth Research Paper* (22-08).

Bianchi, F., C. Ilut, and H. Saijo (2021). Diagnostic business cycles. *NBER Working Paper 28604*.

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.

Bordo, M., B. Eichengreen, D. Klingebiel, and M. S. Martinez-Peria (2001). Is the crisis problem growing more severe? *Economic Policy 16*(32), 52–82.

Borio, C. and P. Lowe (2002, 07). Asset Prices, Financial and Monetary Stability: Exploring the Nexus. *BIS Working Papers* (114).

Borio, C. and P. Lowe (2013). Systemic Banking Crises Database. *IMF Economic Review 61*(114), 225–270.

Brunnermeier, M. K. and Y. Sannikov (2014, February). A macroeconomic model with a financial sector. *American Economic Review 104*(2), 379–421.

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

Chamberlain, G. (1982, June). Panel data. Working Paper 913, National Bureau of Economic Research.

Chang, Y., J. H. Hong, M. Karabarbounis, Y. Wang, and T. Zhang (2021). Income volatility and portfolio choices. *Review of Economic Dynamics*.

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

Chodorow-Reich, G., A. M. Guren, and T. McQuade (2021, August). The 2000s housing cycle with 2020 hindsight: A neo-kindlebergerian view. Technical report.

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.

Claessens, S., M. Ayhan Kose, and M. E. Terrones (2010). The global financial crisis: How similar? how different? how costly? *Journal of Asian Economics 21*(3), 247–264. The Financial Crisis of 2008-09: Origins, Issues, and Prospects.

Coibion, O., D. Georgarakos, Y. Gorodnichenko, G. Kenny, and M. Weber (2023). The effect of macroeconomic uncertainty on household spending. *American Economic Review (forthcoming)*.

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.

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 (2021a). 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 (2021b). Iq, expectations, and choice. Technical report, Review of Economic Studies, forthcoming.

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.

Dang, T. V., G. Gorton, and B. Holmström (2020). The information view of financial crises. *Annual Review of Financial Economics 12*(1), 39–65.

Das, S., C. M. Kuhnen, and S. Nagel (2020). Socioeconomic status and macroeconomic expectations. *The Review of Financial Studies 33*(1), 395–432.

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.

Fagereng, A., L. Guiso, and L. Pistaferri (2018). Portfolio choices, firm shocks, and uninsurable wage risk. *The Review of Economic Studies 85*(1), 437–474.

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.

Fermand, E., C. M. Kuhnen, G. Li, and I. Ben-David (2023). Extrapolative uncertainty and household economic behavior. Available at SSRN: https://ssrn.com/abstract=3293399 or http://dx.doi.org/10.2139/ssrn.3293399.

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

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

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

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.

Gorton, G. and G. Ordoñez (2014, February). Collateral crises. *American Economic*

Review 104(2), 343–78.

Greenwood, R., S. G. Hanson, and L. J. Jin (2019, April). Reflexivity in credit markets. Working Paper 25747, National Bureau of Economic Research.

Greenwood, R., S. G. Hanson, A. Shleifer, and J. A. Sørensen (2020, June). Predictable financial crises. Working Paper 27396, 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.

Guvenen, F. and A. A. Smith (2014). Inferring labor income risk and partial insurance from economic choices. *Econometrica 82*(6), 2085–2129.

Haaland, I., C. Roth, and J. Wohlfart (2021). Designing information provision experiments. *Journal of Economic Literature*.

He, Z. and A. Krishnamurthy (2019, October). A macroeconomic framework for quantifying systemic risk. *American Economic Journal: Macroeconomics 11*(4), 1–37.

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

Jordà, O., M. Schularick, and A. M. Taylor (2013). When credit bites back. *Journal of Money, Credit and Banking 45*(s2), 3–28.

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Keane, M. P. and D. E. Runkle (1998). Are financial analysts' forecasts of corporate profits rational? *Journal of Political Economy 106*(4), 768–805.

Kiyotaki, N. and J. Moore (1997). Credit cycles. *Journal of Political Economy 105*(2), 211–248.

Krishnamurthy, A. and W. Li (2020, May). Dissecting mechanisms of financial crises: Intermediation and sentiment. Working Paper 27088, National Bureau of Economic Research.

Krishnamurthy, A. and T. Muir (2017, September). How credit cycles across a financial crisis. Working Paper 23850, 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.

Lee, S. C. and P. Maxted (2023, May). Household borrowing and mpcs in heterogeneous-agent models. Technical report.

L'Huillier, J.-P., S. R. Singh, and D. Yoo (2023). Incorporating diagnostic expectations into the new keynesian framework. *Review of Economic Studies (forthcoming)*.

Li, W. (2019). Public liquidity and financial crises. Available at SSRN: https://ssrn.com/abstract=3175101.

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.

Reinhart, C. M. and K. S. Rogoff (2009, May). The aftermath of financial crises. *American Economic Review 99*(2), 466–72.

Reuven, G. and K. J. Lansing (2010). Global household leverage, house prices, and consumption. *FRBSF Economic Letter*.

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.

Yin, X. (2022, October). Learning in the Limit: Income Inference from Credit Extension. Technical report.

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

A. Figures

Figure 1. Beliefs in Future Income

Panel A is a binned scatter plot of consumer ex ante income expectation and ex post income realization. Panel B plots the histograms of misbeliefs in future income. $E_C[Y_{t+1}]$ is the expected level of income (\$) in period $t+1$ based on survey question 8. Y_{t+1} is consumer realized income (\$) in period $t+1$. All variables are winsorized at 1% level by each wave.

Panel A: Expectations vs Realizations

Panel B: Misbeliefs in Future Income

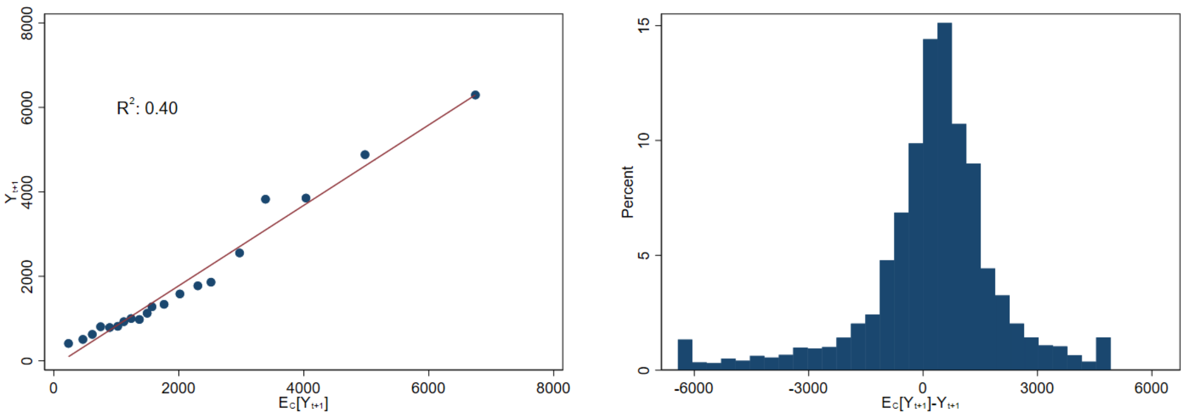

Figure 2. Subjective vs Objective Income Shocks

This figure presents the binned-scatter plots of subjective income shocks (those measured from the surverys) vs the objective income innovations (those measured based on (4)). Panel A gives the raw measures. In Panel B, variables are residualized by income changes in period t, age age-squared, degree, expected log standard deviation of expected income growth in period $t+1$, number of hours worked every week, log income in period t, industry fixed effects, and city $\times$ year fixed effects. All variables are winsorized at 1% level by each wave.

Panel A: Unresidualized　　　　Panel B: Residualized

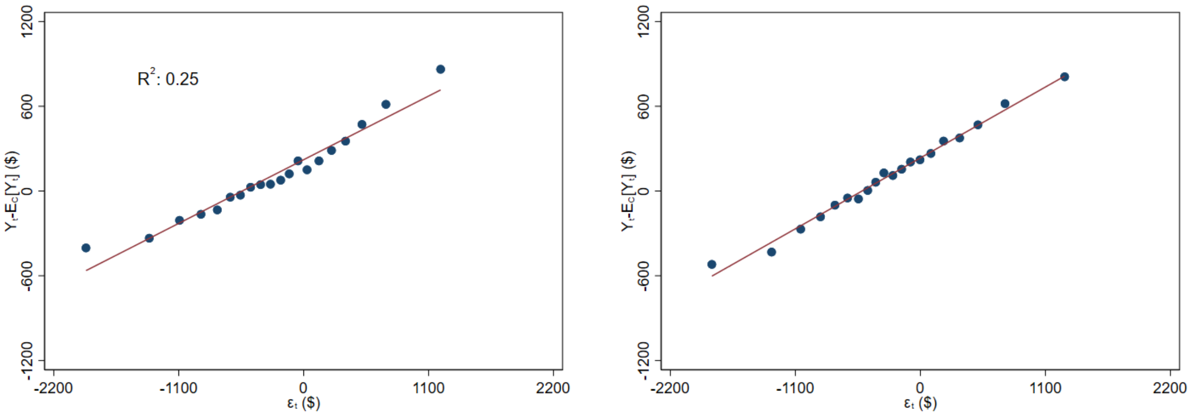

Figure 3. Income Shocks and Misbeliefs in Future Income

This figure presents the binned-scatter plots of misbeliefs in future income vs past income shocks. $E_C[Y_t]$ is the subjective income expectation ($) in period t based on survey question 8 sent in period $t-1$. Y_t is consumer realized income ($) in period t. $\epsilon_{i,t}$ is income innovations measured from (4). In panels B and D, variables are residualized by income changes in period t, age age-squared, degree, expected log standard deviation of expected income growth in period $t+1$, number of hours worked every week, log income in period t, industry fixed effects, and city $\times$ year fixed effects. All variables are winsorized at 1% level by each wave.

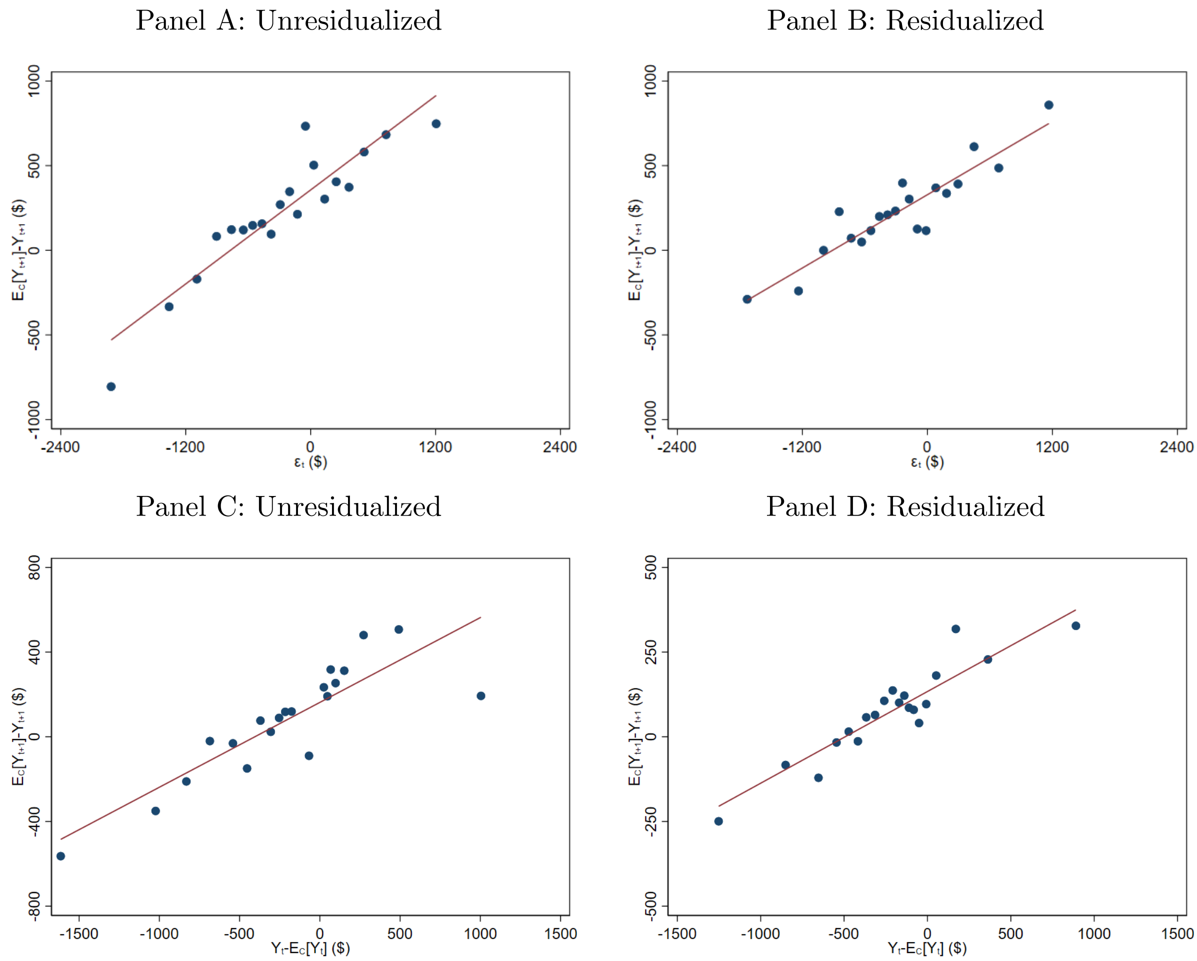

Figure 4. Beliefs in Future Income

This figure reports the effects of expectation errors on cumulative spending and debt. Panel A gives the results for total non-durable spending, and panel B gives those for unsecured debt. The red solid lines depict positive expectation errors, and the blue dashed lines gives the results for negative expectation errors. Estimation is based on (7).

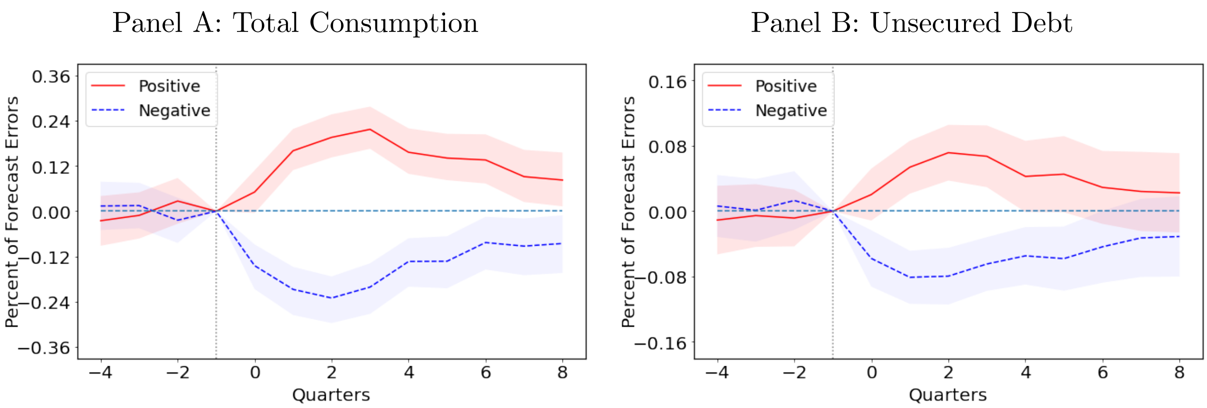

Figure 5. Impulse Responses after Transitory Income Shocks

This figure gives the IRFs after a series of transitory income shocks at the half-year frequency. The simulation is based on 20,000 individuals initially at the stochastic steady states. Starting at the stochastic steady states, each individual receives a 3-year sequence of positive shocks that result in a 3 standard deviation cumulative shock over three years. The top left panel gives the introduced transitory income shocks. The top right panel gives the updates in expected log income, where $o_{i,t} = \kappa(1+\theta)(y_{i,t}-\alpha-\hat{z}_{i,t-1})$. The bottom four panels are respectively the percentage difference in income expectations, consumption, borrowing, and default probability relative to when shocks are not introduced. The red solid lines present the results when $\theta = 1.65$; the blue dashed lines present the results when $\theta = 0$.

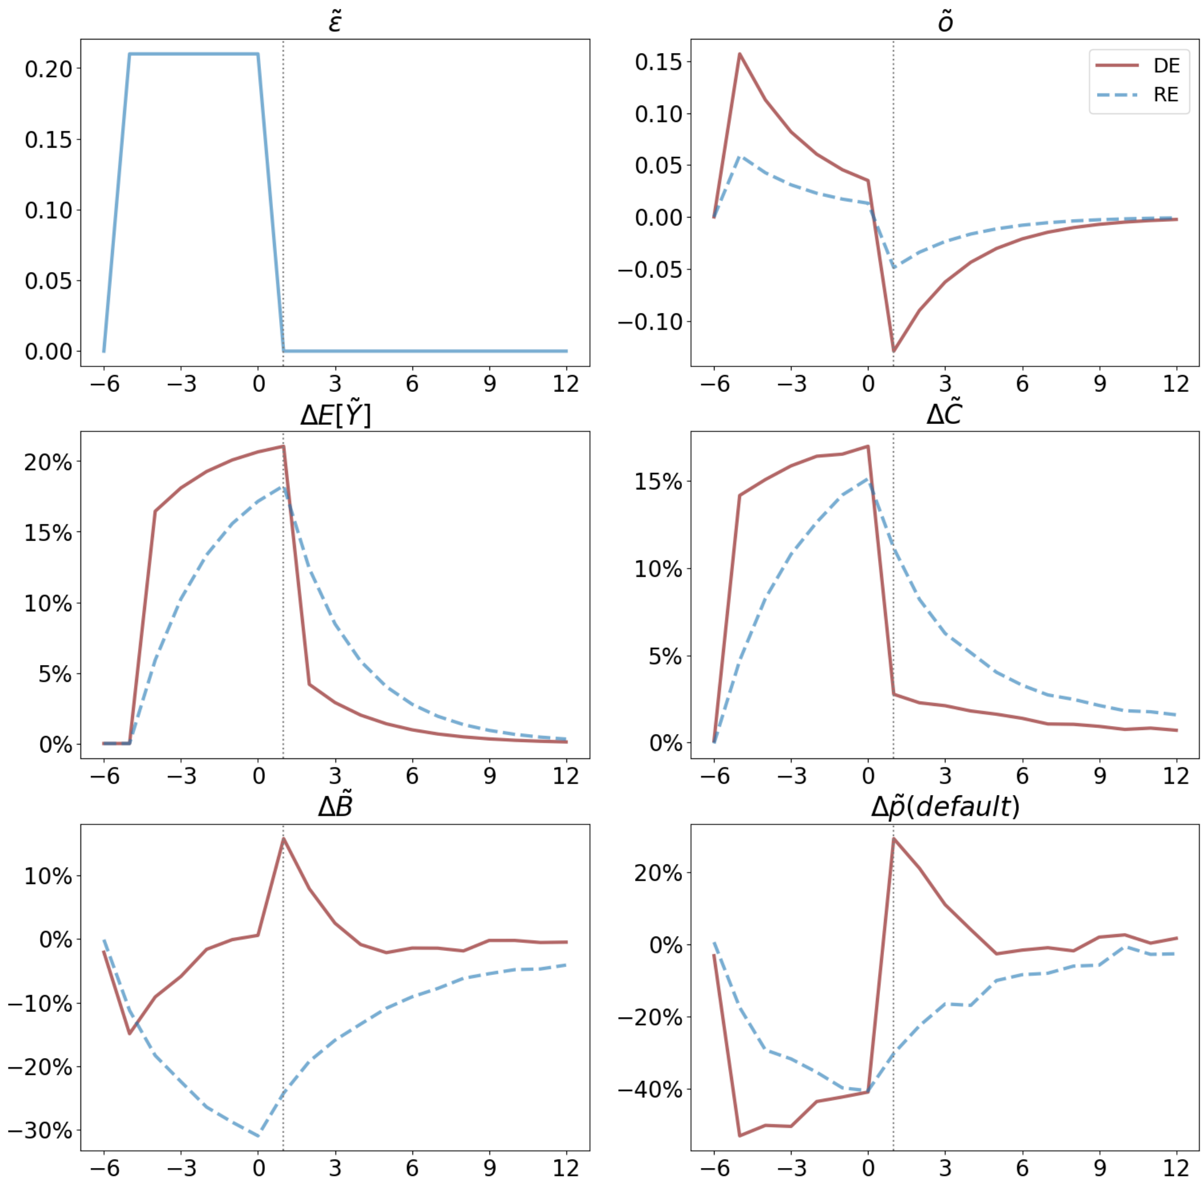

Figure 6. Simulating the 2007-2009 Financial Crisis in the US

This figure simulates the financial crisis given a series of transitory income shocks at the half-year frequency. The simulation is based on 20,000 individuals. Starting at the stochastic steady states, each individual receives a series of transitory shocks as indicated in the top panel. In the bottom two panels, the red solid lines present the simulation when $\theta = 1.65$; the blue dotted lines present the results when $\theta = 0$; the black dashed lines give the data. On the middle panel, $\overline{B}/\overline{Y}$ is the ratio of total debt from credit cards and other credit lines and GDP, divided by labor share. $\overline{B}/\overline{Y}$ is detrended over 2003 to 2012. On the bottom panel, $\overline{p}(default)$ is consumer debt delinquency rate multiplied by the proportion of individuals with positive consumer debt.

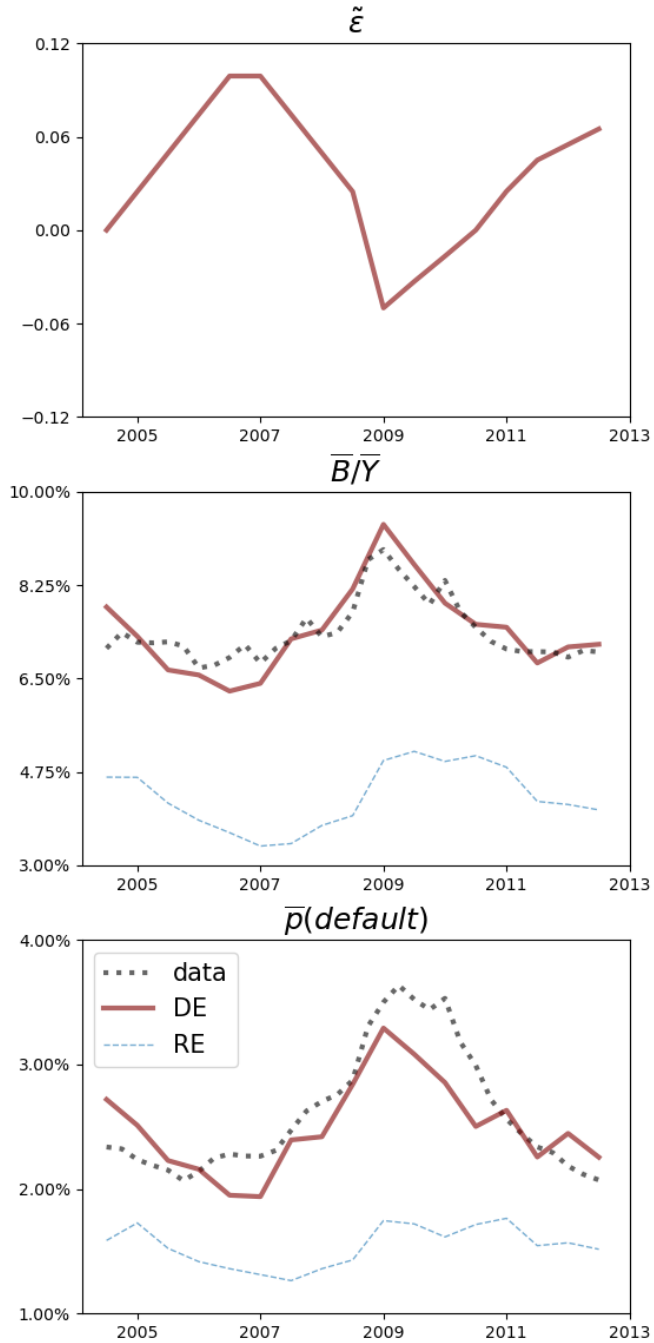

B. Tables

TABLE I. Summary Statistics

Spending is the average monthly spending from the consumers' bank account plus that from the credit card account. Income is the average monthly income. Saving is the average saving. Limit the credit card limit retrieved from the credit registry. Debt is the average interest-incurring credit card debt. Debt|Debt>0 is the average interest-incurring credit card debt for those that have positive debt. Income - E[Income] is the income shocks from fitting (4). E_C[Income] and E_C[Limit] are respectively based on the answers from survey Q8 and Q5. $SD_C[\Delta$ log Income] and $SD_C[\Delta$ log Limit] are the subjective standard deviation of expected income growth and limit growth. All level variables are in dollars and are winsorized at 1% level by each wave.

	(1) Mean	(2) SD	(3) p25	(4) Median	(5) p75	(6) N
Age	39.28	9.00	29	39	48	10497
Female	0.52	0.50	0	1	1	10497
Spending	1245.31	1588.01	292.73	761.68	1451.23	10497
Income	2067.28	2124.95	722.30	1227.24	2459.59	10497
Saving	17817.82	29855.22	1645.41	4261.38	19893.30	10497
Limit	12473.88	17380.76	2031.25	3906.25	15625.00	10497
Debt	1014.86	1655.40	0.00	0.00	1510.59	10497
Debt|Debt>0	2315.65	1830.48	651.09	1863.64	3717.34	4631
E_C[Income]	2039.69	1601.22	937.50	1546.88	2615.39	10497
$SD_C[\Delta$ log Income]	0.36	0.37	0.14	0.24	0.45	10497
E_C[Income]-Income	231.12	1789.11	-401.55	380.63	1111.82	10497
E_C[Limit]	15034.46	15429.41	4272.84	8461.54	21021.64	10497
$SD_C[\Delta$ log Limit]	0.55	0.77	0.12	0.27	0.61	10497
E_C[Limit]-Limit	2585.04	7905.94	-733.17	2203.13	6093.75	10497

TABLE II. Income Shocks and Excessive Income Expectations

$E_C[Y_{t+1}]$ is the expected level of income (\$ thousands) in period $t+1$ based on survey question 8 sent in period t. $E_C[\Delta Y_t]$ is the expected changes (\$ thousands) in income between period $t-1$ and period t. $E[Y_t]$ is the expected income at time t estimated from (4). $SD(E_C[\Delta \log Y_t])$ is the standard deviation of expected income growth in period t based on survey questions 6, 7, and 8 assuming income growth follows a Triangular distribution. $College$ is an indicator if the consumers' highest degree is college and above. $\log Hours_t$ is the number of hours the customers usually work every week in period t. $\log Y_t$ and $\log L_t$ are respectively log monthly income and log credit card limit in period t. All variables are winsorized at 1% level by each wave.

	(1) $E_C[Y_{t+1}] - Y_{t+1}$	(2) $E_C[Y_{t+1}] - Y_{t+1}$	(3) $E_C[Y_{t+1}] - Y_{t+1}$	(4) $E_C[Y_{t+1}] - Y_{t+1}$
$Y_t - E[Y_t]$	0.457***	0.424***	0.393***	0.400***
	(0.045)	(0.053)	(0.052)	(0.088)
$E_C[\Delta Y_t]$		-0.015	-0.026	0.043
		(0.052)	(0.054)	(0.086)
Age		-0.024**	-0.026***	
		(0.010)	(0.010)	
Age^2		0.000**	0.000**	
		(0.000)	(0.000)	
$Female$		-0.165***	-0.153***	
		(0.044)	(0.045)	
$College$		0.094**	0.067*	
		(0.040)	(0.039)	
$\log Hours_t$		-0.053	-0.011	-0.126
		(0.066)	(0.073)	(0.092)
$SD(E_C[\Delta \log Y_t])$		-0.021*	-0.018*	-0.002
		(0.011)	(0.009)	(0.015)
$\log Y_{t-1}$		-0.271***	-0.233***	-0.318**
		(0.031)	(0.029)	(0.125)
$\log L_{t-1}$		0.001	0.007	0.030
		(0.013)	(0.011)	(0.028)
N	10,497	10,497	10,497	10,497
Industry FE	No	No	Yes	No
City × Round FE	No	No	Yes	Yes
Individual FE	No	No	No	Yes
R^2	3.38%	5.94%	12.19%	60.75%

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

TABLE III. Income Shocks and Excessive Income Expectations—Heterogeneity Analysis

$E_C[Y_{t+1}]$ is the expected level of income ($ thousands) in period $t+1$ based on survey question 8 sent in period t. $E_C[\Delta Y_t]$ is the expected changes ($ thousands) in income between period $t-1$ and period t. $E[Y_t]$ is the expected income at time t estimated from (4). SD_H is a dummy variable indicating if $SD(E_C[\Delta \log Y_t])$ is above the median. Age_H is a dummy variable indicating if the age of the individual is above 38. $College$ is a dummy variable indicating if the highest degree earned is college or above. $1_{\{2nd\ Round\}}$ is a dummy variable indicating if the data is from the second round of survey collection. $1_{\{Y_t - E[Y_t]<0\}}$ is a dummy variable indicating if the income shock is negative. Controls include log number of hours worked, log income and credit limit in period $t-1$, and subjective uncertainty about income in period t, and the main effects. All variables are winsorized at 1% level by each wave.

	(1) $E_C[Y_{t+1}] - Y_{t+1}$	(2) $E_C[Y_{t+1}] - Y_{t+1}$	(3) $E_C[Y_{t+1}] - Y_{t+1}$	(4) $E_C[Y_{t+1}] - Y_{t+1}$	(5) $E_C[Y_{t+1}] - Y_{t+1}$	(6) $E_C[Y_{t+1}] - Y_{t+1}$
$Y_t - E[Y_t]$	0.505***	0.317***	0.542***	0.541***	0.376***	0.383***
	(0.073)	(0.092)	(0.064)	(0.139)	(0.088)	(0.107)
$Y_H \times (Y_t - E[Y_t])$	-0.228**					
	(0.112)					
$SD_H \times (Y_t - E[Y_t])$		0.163**				
		(0.082)				
$Age_H \times (Y_t - E[Y_t])$			-0.206*			
			(0.116)			
$Degree_H \times (Y_t - E[Y_t])$				-0.196		
				(0.131)		
$Year_{2020} \times (Y_t - E[Y_t])$					0.096*	
					(0.048)	
$1_{\{Y_t - E[Y_t]<0\}} \times (Y_t - E[Y_t])$						-0.105
						(0.177)
N	10,497	10,497	10,497	10,497	10,497	10,497
Controls	Yes	Yes	Yes	Yes	Yes	Yes
City × Round FE	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	60.74%	60.78%	60.77%	60.76%	60.78%	60.76%

Standard Errors Clustered at City Level in Parentheses

* p <0.10 ** p <0.05 *** p <0.01

TABLE IV. Excessive Income Expectations and Choices: Consumption, Debt, and Defaults

ΔC_t is the differences (\$ thousands) in the monthly average consumption between t and $t-1$. ΔB_{t+1} is the differences (\$ thousands) in the average interest-incurring unsecured debt between $t+1$ and t. $E_C[Y_{t+1}]$ is the expected level of income (\$ thousand) in period $t+1$ based on survey question 8 sent in period t. $E_C[\Delta Y_t]$ is the expected changes (\$ thousands) in income from period $t-1$ to period t. $SD(E_C[\Delta \log Y_t])$ is the standard deviation of expected income growth in period t based on survey questions 6, 7, and 8 assuming income growth follows a Triangular distribution. $\log Hours_t$ is the log number of hours the customers usually work every week in period t. $\log Y_{t-1}$ and $\log L_{t-1}$ are respectively log monthly income and log credit card limit in period $t-1$. All variables are winsorized at 1% level by each wave.

	(1) ΔC_t	(2) ΔC_t	(3) ΔB_{t+1}	(4) ΔB_{t+1}	(5) $Default_{t+1}$	(6) $Default_{t+1}$
$E_C[Y_{t+1}] - Y_{t+1}$	0.208***	0.206***	0.087***	0.075***	0.735***	0.918***
	(0.025)	(0.036)	(0.010)	(0.014)	(0.082)	(0.138)
$E_C[\Delta Y_t]$		0.015		-0.034		-0.645
		(0.081)		(0.027)		(0.522)
$\log Hours_t$		-0.169		-0.138***		-5.902**
		(0.137)		(0.050)		(2.748)
$SD(E_C[\Delta \log Y_t])$		0.011		-0.011		0.267
		(0.030)		(0.011)		(0.307)
$\log Y_{t-1}$		-0.141		-0.029		-16.484***
		(0.126)		(0.049)		(3.703)
$\log L_{t-1}$		-0.032		0.011		0.008
		(0.026)		(0.012)		(0.292)
N	10,497	10,497	10,497	10,497	10,497	10,497
Industry FE	No	Yes	No	Yes	No	Yes
City × Round FE	No	Yes	No	Yes	No	Yes
Individual FE	No	Yes	No	Yes	No	Yes
R^2	3.27%	54.76%	2.91%	61.00%	0.74%	53.65%

Standard Errors Clustered at City Level in Parentheses

* p <0.10 ** p <0.05 *** p <0.01

TABLE V. Income Shocks and Subjective Income Uncertainty

$E_C[Y_{t+1}]$ is the expected level of income ($ thousands) in period $t+1$ based on survey question 8 sent in period t. $E_C[\Delta Y_t]$ is the expected changes ($ thousands) in income between period $t-1$ and period t. $E[Y_t]$ is the expected income at time t estimated from (4). $abs(Y_t - E[Y_t])$ is the absolute value of income surprises at period t. $1_{\{Y_t - E[Y_t] < 0\}}$ is an indicator function that's equal to one if the income surprises at period t are negative. $SD(E[\Delta \log Y_{t+1}])$ is the standard deviation of the income growth in period $t+1$ based on survey questions 6, 7, and 8 assuming income growth follows a Triangular distribution. $\log Hours_t$ is the number of hours the customers usually work every week in period t. $\log Y_{t-1}$ and $\log L_{t-1}$ are respectively log monthly income and log credit card limit in period $t-1$. All variables are winsorized at 1% level by each wave.

	(1) $\Delta SD(E[\Delta \log Y_{t+1}])$	(2) $\Delta SD(E[\Delta \log Y_{t+1}])$	(3) $\Delta SD(E[\Delta \log Y_{t+1}])$	(4) $\Delta SD(E[\Delta \log Y_{t+1}])$
$abs(Y_t - E[Y_t])$	0.046	0.282***	0.309***	0.180**
	(0.029)	(0.073)	(0.071)	(0.084)
$1_{\{Y_t - E[Y_t] < 0\}}$				-0.050
				(0.052)
$abs(Y_t - E[Y_t]) \times 1_{\{Y_t - E[Y_t] < 0\}}$				0.174*
				(0.090)
$E_C[\Delta Y_t]$			-0.092***	-0.100***
			(0.016)	(0.019)
$\log Hours_t$			-0.010	-0.015
			(0.044)	(0.045)
$\log Y_{t-1}$			-0.449***	-0.447***
			(0.139)	(0.139)
$\log L_{t-1}$			0.002	0.001
			(0.012)	(0.011)
N	10,497	10,497	10,497	10,497
Industry FE	No	Yes	Yes	Yes
City × Round FE	No	Yes	Yes	Yes
Individual FE	No	Yes	Yes	Yes
R^2	0.05%	65.69%	66.07%	66.09%

Standard Errors Clustered at City Level in Parentheses

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

TABLE VI. Asymmetric Effects of Excessive Income Expectations

ΔC_t is the differences ($ thousands) in the monthly average consumption between t and $t-1$. ΔB_{t+1} is the differences ($ thousands) in the average interest-incurring unsecured debt between $t+1$ and t. $E[\Delta Y_t]$ is the expected changes ($ thousands) in income from period $t-1$ to t. $E[Y_t]$ is the expected income at time t estimated from (4). $1_{\{Y_t - E[Y_t]<0\}}$ is an indicator function that's equal to one if the income surprises at time 0 are negative. $SD(E[\Delta \log Y_t])$ is the standard deviation of expected income growth in period t based on survey questions 6, 7, and 8 assuming income growth follows a Triangular distribution. $\log Hours_t$ is the number of hours the customers usually work every week in period t. $\log Y_{t-1}$ and $\log L_{t-1}$ are respectively log monthly income and log credit card limit in period $t-1$. All variables are winsorized at 1% level by each wave.

	(1) ΔC_t	(2) ΔC_t	(3) ΔB_{t+1}	(4) ΔB_{t+1}
$Y_t - E[Y_t]$	0.352***	0.204*	0.167***	0.130*
	(0.078)	(0.113)	(0.039)	(0.066)
$1_{\{Y_t - E[Y_t]<0\}}$		-0.040		0.067
		(0.109)		(0.045)
$(Y_t - E[Y_t]) \times 1_{\{Y_t - E[Y_t]<0\}}$		0.198*		0.118*
		(0.118)		(0.069)
$E_C[\Delta Y_t]$	0.068	0.068	-0.008	-0.008
	(0.094)	(0.094)	(0.032)	(0.032)
$\log Hours_t$	-0.177	-0.182	-0.139***	-0.144***
	(0.143)	(0.144)	(0.050)	(0.049)
$SD(E_C[\Delta \log Y_t])$	0.008	0.009	-0.012	-0.011
	(0.030)	(0.030)	(0.012)	(0.011)
$\log Y_{t-1}$	-0.230*	-0.231*	-0.065	-0.059
	(0.119)	(0.119)	(0.052)	(0.051)
$\log L_{t-1}$	-0.024	-0.024	0.014	0.014
	(0.026)	(0.026)	(0.012)	(0.012)
N	10,497	10,497	10,497	10,497
Industry FE	Yes	Yes	Yes	Yes
City × Round FE	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes
R^2	53.59%	53.60%	60.26%	60.28%

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

TABLE VII. Estimation

This table gives the estimated parameters of the structural model. Panel A presents the parameters estimated in the first stage. Panel B gives the parameters estimated based on SMM. Panel C gives the matched moments. Panel D gives the moments not directly targeted. w/c is the average wealth-consumption ratio. p(default) is the proportion of defaults. w/y is the average wealth-income ratio. median(w/y) is the median of the wealth-income ratio. low liq share is the fraction of consumers with saving less than three months of income. top 5% liq share is the fraction of saving held by the top 5% individuals in the model. Estimates of moments in the model is based on a simulation of 5,000 individuals with 500 periods, after a burning period of 1,000 periods for the distribution to reach the steady state. Estimates of moments in the data is based on a random 5% of active customers in the bank's database. Model moments are trimmed at saving to average income ratio larger than 10.

Panel A First-Stage Parameters		Panel B Second-Stage Prameters			Panel C Targeted Moments			Panel D Not Targeted Moments		
	Estimates (1)		Estimates (2)	S.E. (3)		Data (4)	Model (5)		Data (6)	Model (7)
α	0	γ	2.511	(0.267)	w/c	0.818	0.818	w/y	1.351	1.243
ρ	0.970	χ	24.453	(1.201)	p(default)	2.36%	2.36%	median(w/y)	0.731	0.777
σ_ν	0.150							top 5% liq share	31.38%	29.85%
σ_ϵ	0.420							debtor %	29.89%	32.34%
β	0.975									
ν	0.500									
r_b	0.055									
r_s	0.014									
θ	1.650									

TABLE VIII. Linking Income Shocks, Forecast Errors, and Consumption-Debt Decisions in the Model

In Panel A, $E_C[Y_{t+1}]$ is the expected level of income ($ thousands) in period $t+1$ based on survey question 8 sent in period t. In panels B and C, $E_C[Y_{t+1}]$ is the expected level of income at time t under diagnostic expectation with $\theta = 1.65$ and that with $\theta = 0$. In Panel A, $E[Y_{t+1}]$ is the expected income at time t estimated by (4). In Panel B, $E[Y_{t+1}]$ is the expected income when $\theta = 0$. ΔC_t is the changed in total consumption, ΔB_{t+1} is the next-period changes in total liquid debt. $Default$ is a dummy variable that is equal to 100 if consumer i has a 90-day delinquency, and zero otherwise. All variables in Panel A are winsorized at 1% level by each wave. For columns (1) in panels B and C, data is based on 100,000 periods of simulation with a burning periods of 100. For columns (2) to (4) in panels B and C, analysis is based on a simulation of 20,000 individuals with 1,000 periods, after a burning period of 100 periods. Simulated data is right trimmed at saving to average income ratio larger than 8.13, which is the maximum number in the data.

	(1) $E_C[Y_{t+1}] - Y_{t+1}$	(2) ΔC_t	(3) ΔB_{t+1}	(4) $Default_{t+1}$
	Panel A: Data			
$Y_t - E[Y_t]$	0.400***			
	(0.096)			
$E_C[Y_{t+1}] - Y_{t+1}$		0.206***	0.075***	0.918***
		(0.036)	(0.014)	(0.138)
N	10,497	10,497	10,497	10,497
	Panel B: $\theta = 1.650$			
$Y_t - E[Y_t]$	0.400***			
	(0.004)			
$E_C[Y_{t+1}] - Y_{t+1}$		0.243***	0.074***	0.952***
		(0.000)	(0.000)	(0.003)
N	999,998	7,900,875	7,900,875	7,900,875
	Panel C: $\theta = 0$			
$Y_t - E[Y_t]$	-0.040***			
	(0.009)			
$E_C[Y_{t+1}] - Y_{t+1}$		-0.002***	0.024***	0.392***
		(0.000)	(0.000)	(0.002)
N	999,998	7,900,875	7,900,875	7,900,875

Standard Errors in Parentheses
* p <0.10 ** p <0.05 *** p <0.01