ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF ST. LOUIS

WORKING PAPER SERIES

What Drives Household Financial Distress? The Role of Earnings Misperceptions

Authors	Juan M. Sánchez, and Stefan Song
Working Paper Number	2025-030D
Revision Date	July 2026
Citable Link	https://doi.org/10.20955/wp.2025.030
Suggested Citation	Sanchez, J.M., Song, S., 2026; What Drives Household Financial Distress? The Role of Earnings Misperceptions, Federal Reserve Bank of St. Louis Working Paper 2025-030. URL https://doi.org/10.20955/wp.2025.030

Federal Reserve Bank of St. Louis, Research Division, P.O. Box 442, St. Louis, MO 63166

What Drives Household Financial Distress?
The Role of Earnings Misperceptions

Juan M. Sánchez
Federal Reserve Bank of St. Louis

Yingrui (Stefan) Song
Washington University in St. Louis

July 2026*

Abstract

Why do households borrow heavily and experience financial distress? We develop and estimate a heterogeneous-agent model of unsecured credit and default in which households learn about income persistence and overreact to recent income realizations. We estimate overreaction using survey measures of income expectations and forecast errors, and the remaining parameters to match household debt and financial distress. Information frictions and estimated overreaction account for roughly half of delinquencies and one-third of bankruptcies, and improve the model's fit of the observed negative relationship between income and interest rates. Because the models with earnings misperceptions do not need extremely impatient households to match financial distress, policies that substantially reduce financial distress by restricting risky borrowing entail much smaller welfare costs.

Keywords: Income, Information Frictions, Overreaction, Delinquency, Bankruptcy.
JEL Codes: D14, D84, E21, G51.

*The views expressed herein are those of the authors and do not reflect the views of the FRB of St. Louis or the Federal Reserve System. The authors acknowledge helpful comments by Mark Bils, Philipp Grübener, Julian Kozlowski, Eungik Lee, Paulo Lins, Yueran Ma, Matias Moretti, Jose Mustre-del-Río, Luigi Pistaferri, Joaquín Saldain, Yongseok Shin, Ping Wang, and participants in the Midwest Macro at Kansas City Fed, the EGSC and Macro Study Group at Washington University in St. Louis, the brown bag seminar series at the St. Louis Fed, the Stockman Conference at the University of Rochester, the Philadelphia Fed Conference on Consumer Finance and Macroeconomics, the North America Summer Meeting of the Econometric Society in Altanta, and the Asia Meeting of the Econometric Society in Hong Kong.

1 Introduction

Understanding whether an earnings change is temporary or persistent is central to household decisions. When information frictions blur this distinction, households may misperceive the persistence of shocks, and belief biases exacerbate the error. This paper studies how such misperceptions amplify household financial distress.

The analysis proceeds in several steps. First, we use a simple two-period model with linear utility to illustrate the main mechanisms at work. Households that must learn about income persistence and overweight recent outcomes when forming their beliefs tend to misjudge their future income. These misperceptions prompt households to borrow too aggressively, increasing the likelihood of default when income falls short in the second period. In this way, both information frictions and overreaction contribute to increased household financial distress.

Second, we introduce a quantitative model that incorporates both formal and informal defaults, transitory and persistent income shocks, and a life-cycle profile. This standard model, labeled **FI**, assumes *full information* and *rational expectations*. Two extensions relax these assumptions. The first removes full information: Households and lenders learn the persistent component of income from earnings histories. This version, the *rational learning* (**RL**) model, maintains the standard structure while adding a rational updating rule for beliefs. Building on that, the second extension removes rational expectations: Households and lenders overweight recent income changes when forming beliefs about the persistent component of income. This version, the *diagnostic expectations* (**DE**) model, introduces a distorted beliefs updating rule that affects current-period decisions alongside the rational updating rule, which continues to initialize the next period with unbiased beliefs.

Third, we estimate the diagnostic expectation parameter θ, which measures the degree of overreaction, by aligning the model's (negative) forecast error autocorrelation with its empirical counterpart in the NYFED *Survey of Consumer Expectations* (SCE).[1] The other three key parameters are estimated by indirect inference, using moments related to debt and financial distress at different ages to capture their life-cycle profiles, based on the FRBNY *Consumer Credit Panel/Equifax* (CCP) and the *Survey of Consumer Finances* (SCF).

Fourth, we quantify the contribution of information frictions and diagnostic expectations to household financial distress by comparing the DE model with the RL and FI

[1] We interpret this empirical pattern through the lens of *diagnostic expectations*. However, the empirical moments we target identify the extent of overreaction rather than its underlying source. For example, the same forecast-error autocorrelation could also arise if households overestimate the persistence of income shocks. Our quantitative results, therefore, do not depend on diagnostic expectations per se, but on the broader implication that households overreact to recent income realizations when forming expectations.

models while holding all other parameters fixed. In partial equilibrium, both mechanisms are quantitatively important. Among households aged 35-44, moving from the DE benchmark to the RL model lowers the delinquency rate from 12.1 to 9.1 percent and the bankruptcy rate from 0.56 to 0.41 percent. Eliminating information frictions altogether further reduces delinquency to 5.5 percent and bankruptcy to 0.30 percent. Overall, diagnostic expectations and information frictions account for roughly half of delinquencies and one-third of bankruptcies. Because these comparisons are conducted in partial equilibrium, one concern is that the reduction in financial distress simply reflects lower borrowing. To address this concern, we solve the models in general equilibrium, allowing the interest rate to adjust to clear the asset market. The main conclusions are unchanged: removing diagnostic expectations and information frictions continues to substantially reduce financial distress, while the equilibrium interest rate rises only modestly.

Fifth, we use illustrative examples to highlight the model's mechanism. In learning-based models, young households that experience a favorable transitory income shock mistakenly interpret part of the increase as permanent. As a result, they borrow more, both because they expect higher future income (credit demand) and because lenders perceive them as less risky and offer cheaper credit (credit supply). When income reverts to its trend, households are left overleveraged and become more likely to enter delinquency or bankruptcy. Diagnostic expectations amplify this mechanism by causing both borrowers and lenders to overreact to recent income realizations, leading to larger debt buildups, sharper reversals, and ultimately greater financial distress.

Sixth, we estimate the RL and FI models separately, matching the same moments on household borrowing and financial distress as in the DE model. Although all three models fit these moments well, the estimated discount factor rises monotonically from the FI model ($\beta \approx 0.78$), to the RL model ($\beta \approx 0.81$), to the DE model ($\beta \approx 0.86$). Equivalently, accounting for information frictions and diagnostic expectations lowers the implied rate of impatience by more than 40%, substantially alleviating the unrealistically high impatience often required in quantitative models of household borrowing.[2] As we show below, these differences in estimated impatience are central to policy evaluation, as they substantially alter the welfare costs of policies that reduce financial distress by restricting borrowing.

To further distinguish the models, we compare their predictions with the empirical correlation between interest rates and household income from the SCF. The FI model yields only a weak negative relationship: Households with lower income do not face

[2] The implied rate of impatience is defined as $\beta^{-1} - 1$. Thus, increasing the discount factor from $\beta = 0.78$ to $\beta = 0.86$ reduces the implied rate of impatience from 28.2% to 16.3%, a decline of more than 40%.

significantly higher rates, as part of that low income may be transitory. Learning-based models (RL and DE) generate a stronger negative link, as lenders read low income as a signal of persistent risk. Diagnostic expectations amplify this effect: By overweighting recent shocks, the benchmark DE model predicts an even steeper penalty for low-income borrowers. The data aligns with this pattern, favoring learning-based models as more accurate descriptions of household behavior and interest rate determination.

Finally, we turn to policy analysis. The substantial differences in the estimated degree of impatience across models imply markedly different policy evaluations. We study interest rate caps, a commonly proposed intervention to reduce household financial distress by restricting borrowing. Although these caps reduce debt, delinquency, and bankruptcy in both models, they are considerably more effective in the diagnostic-expectations (DE) economy while imposing substantially smaller welfare costs. The reason is that the FI model requires a much higher rate of impatience to match household borrowing and financial distress. Overall, accounting for diagnostic expectations substantially improves the policy trade-off between reducing financial distress and preserving household welfare. We also quantify the welfare cost of diagnostic expectations themselves: households in the DE economy would require a 1.5 percent permanent increase in consumption to attain the welfare achieved under full information.

There is a large literature on heterogeneous-agent models with debt non-repayment. Athreya (2002) developed the first dynamic, stochastic general-equilibrium model of personal bankruptcy. Livshits, MacGee and Tertilt (2007) and Chatterjee, Corbae, Nakajima and Ríos-Rull (2007) later introduced the framework most widely used today, building on the debt repudiation setup of Eaton and Gersovitz (1981). In these models, default represents the U.S. bankruptcy system. Subsequent work by Athreya, Sánchez, Tam and Young (2015, 2017) extended the analysis by adding credit card delinquency, which raised the modeled incidence of default from less than 1 percent (bankruptcies) to more than 10 percent (delinquencies). The models in this paper build on Athreya, Mustre-del Río and Sánchez (2019), which incorporates both bankruptcy and delinquency into a life-cycle model similar to Kaplan and Violante (2010) as distinct forms of credit card non-repayment.

To study the role of earnings-shock misperceptions, this paper extends the existing framework in two novel directions. The first is rational learning. Guvenen (2007) and Guvenen and Smith (2014) provide the primary references for adding this feature into heterogeneous-agent models to capture the dynamics of earnings and consumption. However, the focus was on learning about a fixed characteristic. Learning in our setup is more related to a recent working paper, Lee and Sæverud (2023), which documents a declining Kalman gain across the age profile for households' income expectations using Danish data, offering empirical support for a learning mechanism similar to the one introduced here. Moreover, in our model, financial intermediaries also face *information frictions* and must infer households' persistent income component. This connects to the literature on *credit scores* as signals of borrower quality. For instance, Chatterjee, Corbae and Rios-Rull (2008) and Chatterjee, Corbae, Dempsey and Ríos-Rull (2023) show that credit histories help lenders infer a borrower's type. Similarly, in our framework, the inferred mean of the persistent income component reflects households' repayment ability, serving a role analogous to a credit score. However, unlike their models, we abstract from ex-ante heterogeneity in household types and focus instead on the effects of income dynamics. The second extension introduces *diagnostic expectations*. These distortions arise from Kahneman and Tversky (1974)'s representativeness heuristic and base-rate neglect. Unlike rational expectations, households form beliefs through diagnostic expectations, as in Bordalo, Gennaioli, Ma and Shleifer (2020). Our work complements this literature by following Ma, Ropele, Sraer and Thesmar (2020) to discipline the belief-distortion parameter and by applying it within a heterogeneous-agent framework with both formal and informal defaults.

Our paper also relates to a growing literature that incorporates behavioral biases into heterogeneous-agent models of household borrowing and default. The closest paper is D'Acunto, Weber and Yin (2024), which uses Chinese household data to estimate diagnostic expectations and shows that overreaction in income beliefs contributes to household default. Like them, we discipline belief distortions using survey data and embed them in a structural model. We differ by studying U.S. households, modeling both delinquency and bankruptcy, allowing for endogenous interest rates, and emphasizing how learning and diagnostic expectations affect the evaluation of policies aimed at reducing financial distress. Other behavioral approaches include Exler, Livshits, MacGee and Tertilt (2024), who study permanently optimistic borrowers, and Saldain (2025), who model self-control and temptation. Unlike these papers, we estimate the magnitude of the behavioral distortion directly from survey data and quantify its implications for household financial distress. Moreover, whereas they assume that only a subset of households exhibits the behavioral friction, we model a common belief-formation process for all households. This specification is supported by the empirical finding that overreaction exhibits little systematic heterogeneity across households, which allows us to isolate the effects of distorted expectations without introducing an additional screening problem in which lenders must infer borrowers' behavioral types.

Our paper also connects to a broader literature applying behavioral economics to household decision-making. For example, Ameriks, Caplin, Lee, Shapiro and Tonetti (2023) explores behavioral frictions, such as limited awareness, in financial choices; Berger, Milbradt, Tourre and Vavra (2024) incorporates behavioral inattention in refinancing decisions; Drozd and Kowalik (2023) shows that time-inconsistent consumption behavior can rationalize the design of promotional debt products; Nakajima (2012) shows how present-bias alone can generate excessive borrowing and welfare losses; and He and Kircher (2025) shows that belief over-reaction is key to reconciling survey data and model in unemployed workers' job finding rate expectation.

2 The Mechanism in a Simple Two-Period Model

This section presents a two-period economy to show how information frictions and diagnostic expectations amplify financial distress. Proofs of the results appear in Section A in the Appendix.

There is a continuum of ex-ante identical households, each living for two periods. Households begin life with zero assets and derive utility from consumption in both periods. Preferences are given by:

$$V = c_1 + \beta \mathbb{E}[c_2],$$

where β is the subjective discount factor.

Income in the first period is denoted by y and consists of two components: A persistent part z and a transitory shock e, such that $y = z + e$. In the second period, households receive only the persistent component z. For simplicity, we assume that z remains constant across periods. The persistent component z is drawn once at the beginning of life from a Bernoulli distribution with support $\{z_L, z_H\}$, where $z_H > z_L > 0$, $\mathbb{P}(z = z_H) = \pi$, and $\mathbb{P}(z = z_L) = 1 - \pi$. The transitory component e is an idiosyncratic shock drawn from a standard normal distribution: $e \sim \mathcal{N}(0, 1)$.

Risk-neutral lenders discount the future at the risk-free rate r. Households start life with zero assets, so default is not possible in the first period. In the second period, however, households may default on their obligations. Debt prices, therefore, reflect the expected probability of default.[3]

Suppose a household chooses to default in the second period by not repaying its debt obligation, denoted by d'. In that case, it incurs a penalty equal to a fraction $(1 - \nu)$ of its income, where income in the second period is given by the persistent component z. If,

[3]Households can file for bankruptcy but cannot become delinquent, since they live only two periods. Delinquency is incorporated in the quantitative model.

instead, the household repays, it must fully honor the debt d' and receive the full second-period income z. Therefore, the household chooses to default if the repayment burden exceeds the value of income under default—that is, $d' > (1-\nu)z$—and chooses to repay otherwise.

2.1 Full-Information Rational-Expectation (FI) Economy

In this economy, households and lenders observe the realization of z after income y is realized. Conditional on z, the price of debt, denoted by $q(z, d')$, is given by:

$$q(z, d') = \begin{cases} \frac{1}{1+r}, & \text{if } d' \leq (1-\nu)z, \\ 0, & \text{if } d' > (1-\nu)z. \end{cases}$$

Since z is fully observable to the lender, any debt level d' that exceeds the default threshold $(1-\nu)z$ is expected to be defaulted on, thus commanding a zero price. Conversely, if d' is below the threshold, repayment is guaranteed, and the debt is priced at the risk-free rate.

Conditional on the realization of z and e, the household's lifetime utility from consumption over the two periods is given by:

$$V(e, z, d') = \begin{cases} z + e + \beta z + \left(\frac{1}{1+r} - \beta\right) d', & \text{if } d' \leq (1-\nu)z, \text{ (repayment)} \\ z + e + \beta \nu z, & \text{if } d' > (1-\nu)z. \text{ (default)} \end{cases}$$

In the repayment case, the household receives d' in period 1 and repays it in period 2, with utility discounted at β, whereas the lender prices the debt using the market interest rate r. To ensure that households choose to borrow in equilibrium, we assume that households are sufficiently impatient; i.e., $\frac{1}{1+r} > \beta$. Figure 1 depicts a stylized pattern of $q(\cdot)$ and $V(\cdot)$ along d', conditional on the observed persistent income component z.

Since households are impatient, they want to borrow a positive amount. However, since the lender fully observes their income z in the second period, they borrow only until the debt price remains positive. This gives us the central result regarding the full information economy. Households borrow exactly up to the threshold at which lenders are still willing to provide credit at a positive price, and the aggregate default rate is zero.[4]

Lemma 1. *Conditional on z, the household's optimal debt choice is given by $d' = (1-\nu)z$, and, therefore, the aggregate default rate is zero.*

[4] We assume that households at the margin, with $d' = (1-\nu)z$, always choose to repay.

Figure 1: Asset price $q(z, d')$ and value $V(e, z, d')$ as functions of d'

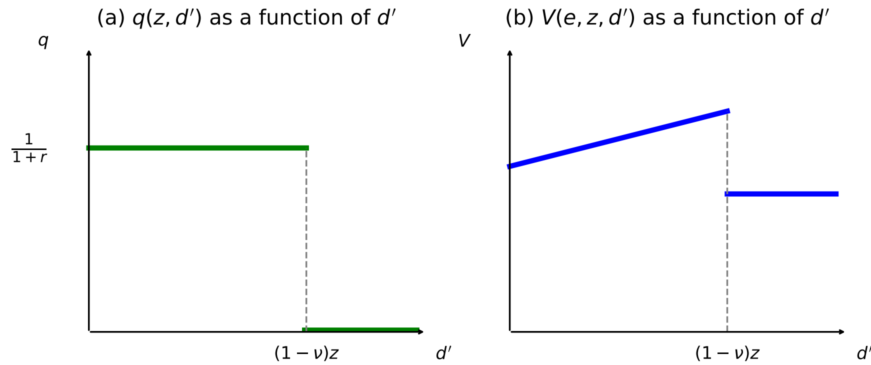

2.2 Rational Learning (RL) Economy

We now consider an economy with information frictions. In the first period, neither households nor lenders observe the persistent income component z. Instead, both form rational beliefs about z based on the observed income realization $y = z + e$. In the second period, $y' = z$, so the persistent component is fully revealed to both households and lenders.

Conditional on y, the posterior belief that the household got the shock z_H, given the observed income y, is given by:

$$Prob(z = z_H|y) = \frac{\phi(y - z_H)\pi}{\phi(y - z_L)(1 - \pi) + \phi(y - z_H)\pi} \equiv x(y),$$

where $\phi(\cdot)$ denotes the probability density function (pdf) of the standard normal distribution. Note also that we define $x(y) \equiv \mathbb{P}(z = z_H \mid y)$ to simplify notation. Importantly, it is simple to show that $x(y)$ is a strictly increasing function of the income signal y.

Now, the debt price $q(y, d')$ depends on the amount borrowed d', as before, and on the income realization y, which will determine the posterior belief about the persistent income component z. The debt pricing function is given by:

$$q(y, d') = \begin{cases} \frac{1}{1+r}, & \text{if } d' \leq (1 - \nu)z_L, \\ \frac{x(y)}{1+r}, & \text{if } (1 - \nu)z_L < d' \leq (1 - \nu)z_H, \\ 0, & \text{if } d' > (1 - \nu)z_H. \end{cases}$$

If the debt level is below or equal to $(1 - \nu)z_L$, repayment is guaranteed regardless of the household's type, and the debt is risk-free. When the debt level falls between $(1 - \nu)z_L$ and $(1 - \nu)z_H$, only households with $z = z_H$ will repay, while those with $z = z_L$ will default. The debt price in this region is thus adjusted downward to reflect the probability of repayment, which equals the posterior probability that $z = z_H$ given the realization of income $x(y)$. Finally, if the debt exceeds $(1-\nu)z_H$, all households are expected to default, and the debt price is zero.

Conditional on the income realization y, the household's expected lifetime utility from consumption over two periods is given by:

$$V(y,d') = \begin{cases} y + \beta\left[x(y)z_H + (1-x(y))z_L\right] + \left(\frac{1}{1+r} - \beta\right)d', & \text{if } d' \le (1-\nu)z_L, \\ y + \beta\left[x(y)z_H + (1-x(y))\nu z_L\right] + x(y)\left(\frac{1}{1+r} - \beta\right)d', & \text{if } (1-\nu)z_L < d' \le (1-\nu)z_H, \\ y + \beta\nu\left[x(y)z_H + (1-x(y))z_L\right], & \text{if } d' > (1-\nu)z_H. \end{cases}$$

Note that the first case corresponds to full repayment by all types, the second to repayment by only high-type households, and the third to default by all types. Figure 2 illustrates the debt pricing function $q(\cdot)$ and the lifetime utility function $V(\cdot)$ as functions of the debt level d', conditional on the observed income y.

Figure 2: Asset price $q(y,d')$ and value $V(y,d')$ as functions of d'

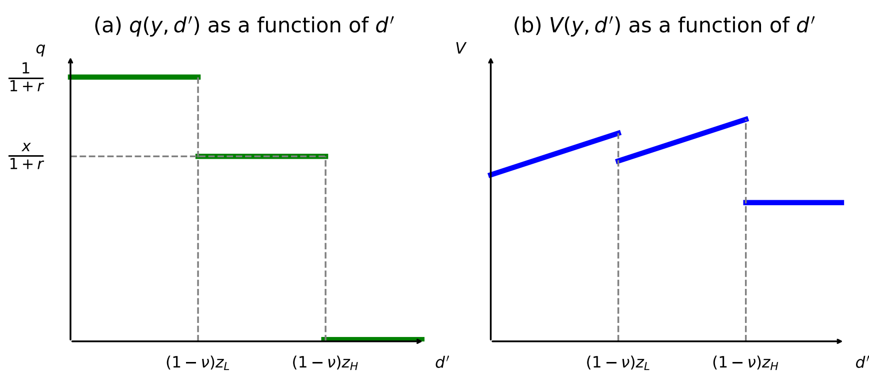

Because $\frac{1}{1+r} > \beta$, households always borrow the maximum allowed within each repayment region. The choice between borrowing $(1-\nu)z_L$ and $(1-\nu)z_H$ depends on the posterior belief $x(y)$. If the posterior is high enough, the household borrows more. Since $x(y)$ is strictly increasing in the income signal y, there exists a threshold level of first-period income above which the household chooses to borrow more.

Lemma 2. *There exists a unique income threshold y^* such that:*

1. *If $y \ge y^*$, the household chooses $d' = (1-\nu)z_H$.*

2. *If $y < y^*$, the household chooses $d' = (1-\nu)z_L$.*

And the aggregate default rate in the RL economy is given by:

$$\int_{y^*}^{\infty} (1 - x(y)) f(y)\, \mathrm{d}y > 0,$$

where $f(y)$ is the probability density function of the income realization y, defined as:

$$f(y) = \pi\,\phi(y - z_H) + (1 - \pi)\,\phi(y - z_L),$$

and $\phi(\cdot)$ denotes the standard normal density function.
Proof: See Appendix A.1.

Default arises only when households receive a high enough income ($y \geq y^*$), borrow the larger amount $(1 - \nu)z_H$, and later discover that the transitory component was high and the persistent component was low ($z = z_L$). The key takeaway is simple: Misperceptions of income persistence trigger default.

2.3 Diagnostic Expectation (DE) Economy

In this economy, both households and lenders overreact to income realizations.[5] We denote the distorted posterior belief that $z = z_H$ by $x^{\theta}(y)$, where $x^0(y) = x(y)$ corresponds to the rational expectations benchmark.

The logic for the borrowing decision is the same as in the case of rational learning. Specifically, in the DE economy, there also exists a income realization threshold, $y^{*,\theta}$, such that borrowing depends on the realization of income in the following way:

1. if $y \geq y^{*,\theta}$, the household chooses $d' = (1 - \nu)z_H$,
2. if $y < y^{*,\theta}$, the household chooses $d' = (1 - \nu)z_L$,

where the threshold $y^{*,\theta}$ corresponds to the signal level at which the posterior belief $x^{\theta}(y^{*,\theta})$ equals the critical threshold at which agents are indifferent between borrowing $(1 - \nu)z_L$ or $(1 - \nu)z_H$.

A key feature of diagnostic expectation is that agents overreact asymmetrically to income realizations. When the first-period income y is high enough, they become overly optimistic about the persistent component z (their second-period income). In contrast, it is low enough when the first-period income, y, leads to excessive pessimism. As a result, there exists a threshold first-period income level at which the posterior probability that $z = z_H$ is identical under both rational learning and diagnostic expectation. This threshold is denoted by $\tilde{y}$ and illustrated in Figure 3. As shown in the figure, diagnostic posterior (solid red line) will be larger than its rational expectation counterpart (yellow

[5] In our quantitative analysis, we also consider a setting where households form diagnostic expectations, while lenders have rational learning. This case is excluded from the current section due to its greater analytical complexity.

dashed line) if income realization is such that $y > \tilde{y}$, vice versa. In the proof of Lemma 3, we will derive a closed form solution for the $\tilde{y}$.

Figure 3: Posterior Beliefs $x(y)$ under RL and DE

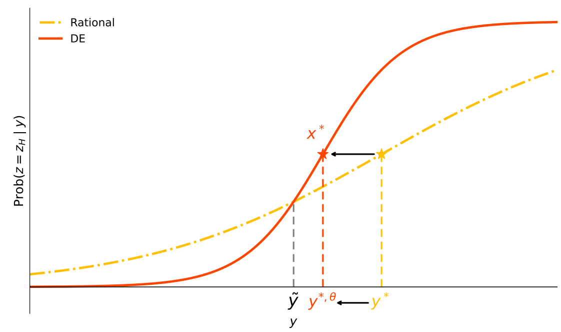

To derive the main result of this section, we impose the following assumption, which implies that $y^{*,\theta} > \tilde{y}$ as in Figure 3.

Assumption 1. $(1+r)\beta > \frac{z_H - z_L/\pi}{z_H - z_L}$.

This condition is not particularly restrictive. The left-hand side is less than one under our assumption that households are more impatient than lenders; i.e., $\beta < \frac{1}{1+r}$. The inequality fails when the prior places full weight on the high state, $\pi = 1$, so one interpretation is that the prior π should not be too close to one. Alternatively, the condition is easily satisfied when the difference between z_H and z_L is modest—that is, when z_H is not substantially larger than z_L. In particular, if $\pi z_H < z_L$, the right-hand side becomes negative, reducing the condition to $\beta > 0$, which is trivially satisfied. For reasonable values of z_H, z_L, and π, the assumption requires that households are not dramatically more impatient than lenders.

Lemma 3. *If Assumption 1 holds, then for any $\theta > 0$, it follows that $y^{*,\theta} < y^*$, and:*

$$\int_{y^{*,\theta}}^{\infty} (1 - x(y)) f(y)\,\mathrm{d}y \; > \; \int_{y^*}^{\infty} (1 - x(y)) f(y)\,\mathrm{d}y \; > \; 0.$$

This implies that the default incidence is strictly higher in the DE economy than in the RL economy.

Proof: See Appendix A.2.

Figure 3 illustrates the intuition behind the result. Under diagnostic expectation, agents overstate the likelihood that $z = z_H$ for any income realization such that $y > \tilde{y}$, so achieving the same posterior threshold (x^*) requires a lower income realization threshold

$(y^{*,\theta})$ than that in the rational expectation benchmark (y^*). Since the ex-ante distribution of y is identical across economies, this shift implies that more households will choose higher levels of debt, even though they have a low, persistent component, thereby increasing the incidence of default.

2.4 Distinguishing Implication

Later in the paper, we empirically evaluate the relationship between interest rates and income (conditional on the debt amount) as a distinguishing implication across models.

The difference across models regarding this relationship is evident even in this simple example. In the FI economy, all households face a risk-free interest rate, regardless of income. As a result, the conditional correlation between income and the interest rate is zero. In the RL economy, the correlation remains zero for households that borrow $d' = (1-\nu)z_L$, as this level of debt is always repaid and priced at the risk-free rate. However, the situation differs for households that receive a high signal and choose to borrow $d' = (1-\nu)z_H$. In this region, the debt price $q(y, d')$ increases (the interest rate decreases) with income y, which implies a higher posterior probability of having z_H. That is, $x(y) = \mathbb{P}(z = z_H \mid y)$ is increasing in y. Therefore, after controlling for debt levels, there is a stronger *negative* correlation between income and interest rates in the learning economy. The same logic applies to the economy with diagnostic expectation, but since beliefs vary more with the signal y, the relationship between interest rates and income is stronger (more negative).

In Section 7.1, we confirm this prediction by running regressions in data simulated with the alternative models and empirically using data from the SCF.

3 Quantitative Model

In this section, we set up the model for an economy with information friction and diagnostic expectation. The formulation of the FI economy, along with the value functions at and after retirement, is detailed in Section A in the Online Appendix.

The economy consists of a continuum of finitely lived, risk-averse households. Each individual enters the labor market with an initial balance a_0 at age $n = n_0$, retires at age $n = T$, and lives up to age $n = N$.

3.1 Income Process

Household income follows a restricted income profile (RIP) process proposed by Blundell and Preston (1998) and more recently used by Kaplan and Violante (2010):[6]

$$\tilde{y}_n = \exp\left(\ell_n + z_n + e_n\right),$$

where ℓ_n represents the deterministic component that depends solely on age, z_n is the persistent stochastic component, and e_n is an i.i.d. transitory shock, distributed as $\mathcal{N}(0, \sigma_e^2)$.

The persistent component evolves as a random walk $z_n = z_{n-1} + \epsilon_n$, where $\epsilon_n \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is an i.i.d. innovation.

We assume that ℓ_n is common knowledge. Each period, agents observe their realized income $\tilde{y}_n$. Additionally, we define the residual income realization y_n as

$$y_n \equiv z_n + e_n = \ln(\tilde{y}_n) - \ell_n,$$

which is also observable since both $\tilde{y}_n$ and ℓ_n are known. However, agents cannot distinguish between the stochastic components of income—persistent and transitory. That is, they do not directly observe either z_n or e_n. This is what we call "information friction."

3.2 Rational Belief Updating

At any age n, a rational household enters the period with prior belief mean about z_n, which we refer to as $z_{n|n-1}$, and has the probability density $f(z_n|Y_{n-1})$, where Y_{n-1} denotes the full history of residual income observed up to the previous period. Then the household and the lender observe a new signal—its residual income realization y_n—and update its estimate of the persistent income component using the Bayes' rule:

$$f(z_n|Y_n) = \frac{f(y_n|z_n) f(z_n|Y_{n-1})}{\int f(y_n|z) f(z|Y_{n-1})\,\mathrm{d}z}, \tag{1}$$

which iteratively defines the household's rational beliefs.

With normal priors and shocks, the rational household's and lender's learning problem reduces to a textbook Kalman filter process. Specifically, at age n, the prior about z_n

[6]Recent studies, such as Braxton, Herkenhoff, Rothbaum and Schmidt (2021), Huang, Braxton, Herkenhoff, Nattinger, Rothbaum and Schmidt (2025), and Ganong, Noel, Patterson, Vavra and Weinberg (2025), have made significant progress in estimating U.S. households' income processes and assessing their macroeconomic implications. We view our chosen income process as a benchmark exercise and expect the qualitative insights of our results to extend to richer specifications of income dynamics.

is summarized by

$$\mathcal{N}(z_{n|n-1}, \Sigma_{n|n-1}).$$

Crucially, in a Kalman filter framework, the variance depends only on the number of received signals, rather than on the specific history of realized signals. In our setting, the number of observed realizations can be mapped directly to age n (since shocks occur yearly), allowing us to express the variance solely as a function of age. This property is fundamental: To fully characterize the distribution of z_n, it suffices to know the mean and the age.

Given the prior $\mathcal{N}(z_{n|n-1}, \Sigma_{n|n-1})$, their belief updating follows:

$$\begin{aligned}
z_{n|n} &= z_{n|n-1} + K_n \left(y_n - z_{n|n-1} \right), \\
\Sigma_{n|n} &= \frac{\sigma_e^2}{\Sigma_{n|n-1} + \sigma_e^2} \Sigma_{n|n-1}, \\
\Sigma_{n|n-1} &= \Sigma_{n-1|n-1} + \sigma_\epsilon^2, \\
K_n &= \frac{\Sigma_{n|n-1}}{\Sigma_{n|n-1} + \sigma_e^2}.
\end{aligned}$$

Given the posterior, a rational household's belief about next period's residual income realization follows a normal distribution with mean equal to the updated belief about income, $z_{n|n}$, and with the variance incorporating the updated uncertainty $\Sigma_{n|n}$ along with the variances of the income shocks: $y_{n+1} \sim \mathcal{N}\left(z_{n|n}, \Sigma_{n|n} + \sigma_\epsilon^2 + \sigma_e^2\right)$.

Note that the mean of residual income in the next period equals the mean of the belief about z in the current period, as z follows a random walk and both e and ϵ have zero means. Similarly, since the shocks are independent, the variance is the sum of the individual variances.

3.3 Distorted Belief Updating

We consider the possibility that both households and lenders form diagnostic beliefs with the extent parameterized by $\theta \geq 0$, following Bordalo, Gennaioli, Ma and Shleifer (2020). In line with previous literature using a diagnostic Kalman filter, we define the representativeness of a state at age n as the likelihood ratio:

$$R_n(z) = \frac{f(z|Y_n)}{f\left(z|Y_{n-1} \cup \left\{z_{n|n-1}\right\}\right)}.$$

Here, $Y_n = Y_{n-1} \cup \{y_n\}$ represents the history up to today and $Y_{n-1} \cup \{z_{n|n-1}\}$ represents the history up to today under the assumption that the new residual income equals the ex-ante forecast, $y_n = z_{n|n-1}$. This implies that state z is more representative at age n if the new signal y_n increases its probability relative to the scenario in which the residual income matches the forecast. Intuitively, the most representative states z are those whose likelihood has increased the most in response to the latest residual income realization.

Agents then overweight representative states by using a distorted posterior, replacing the rational posterior (1) with $f^\theta(z_n|Y_n) = f(z_n|Y_n)\,[R_n(z_n)]^\theta \frac{1}{Z_n}$, where Z_n is a normalization factor ensuring that $f^\theta(z_n|Y_n)$ integrates to one.

The parameter $\theta \geq 0$ measures the extent to which beliefs deviate from rational updating due to representativeness. When $\theta = 0$, beliefs remain rational and are given by the Bayesian conditional distribution $f(z_n|Y_n)$. For $\theta > 0$, the diagnostic density $f^\theta(z_n|Y_n)$ inflates the probability of representative states while deflating that of unrepresentative ones. Crucially, this distortion leads to systematic mistakes: States that have become relatively more likely are overweighted, even if they remain unlikely in absolute terms.

Note that after observing the current residual income realization y_n, the diagnostic posterior follows a normal distribution $\mathcal{N}(z^\theta_{n|n}\,,\,\Sigma_{n|n})$, where the posterior mean $z^\theta_{n|n}$ is adjusted according to diagnostic expectation[7]:

$$
\begin{aligned}
z^\theta_{n|n} &= z_{n|n-1} + (1+\theta)K_n\left(y_n - z_{n|n-1}\right),\\
\Sigma_{n|n} &= \frac{\sigma_e^2}{\Sigma_{n|n-1} + \sigma_e^2}\Sigma_{n|n-1},\\
K_n &= \frac{\Sigma_{n|n-1}}{\Sigma_{n|n-1} + \sigma_e^2}.
\end{aligned}
$$

Thus, an agent's distorted belief about next period's residual income realization follows:

$$
y^\theta_{n+1} = \; z^\theta_{n|n} + \epsilon_{n+1} + e_{n+1} \sim \mathcal{N}(z^\theta_{n|n}\,,\,\Sigma_{n|n} + \sigma_\epsilon^2 + \sigma_e^2).
$$

Since z follows a random walk, the mean of residual income in the next period equals the mean of the belief about z in the current period. However, this belief is distorted due to agents' diagnostic expectations.

A key feature of diagnostic expectation is what the literature refers to as *rational reversion*. The diagnostic distortion applies only to the posterior update within each period. At the end of the period, households reassess their beliefs rationally, as if re-evaluating the entire history of realizations. This reversion ensures that the prior for the next period

[7]For proof, see Proposition 1 of Bordalo et al. (2020).

remains rational. As a direct implication, diagnostic expectations always start with a rational prior, and there is no need to carry the value of the distorted belief from one period to another, therefore reducing the state space.

For our benchmark results with diagnostic expectations, we assume that lenders also hold diagnostic expectations to the same extent (i.e., the same θ) as households.[8] But we also analyze the case in which the lenders have rational learning while households have diagnostic expectations, where our key results remain valid.

3.4 Timeline

Figure 4 illustrates the sequence of events within a period (indexed by age n). Each period begins with households and lenders arriving with a prior belief about the persistent component of income, denoted $z_{-1} \sim \mathcal{N}(z_{-1}, \Sigma_{n|n-1})$, and the household with asset holdings a. At the start of the period, income y_n is realized and observed by both households and lenders. Based on this income signal, both agents update their beliefs and form a distorted posterior with mean z^θ. Next, households choose between bankruptcy (BK), delinquency (DQ), or repayment. Lenders then offer debt prices to those opting to repay, reflecting the updated beliefs and endogenous risk. Households who choose repayment proceed to select consumption and savings (or borrowing) levels, conditional on the prices offered. Finally, a rational reversion step occurs at the end of the period: Beliefs revert to the rational posterior. This structure, as shown in the timeline below, repeats at each period and governs the dynamic evolution of decisions and beliefs within the model.

Figure 4: Timeline Within a Period

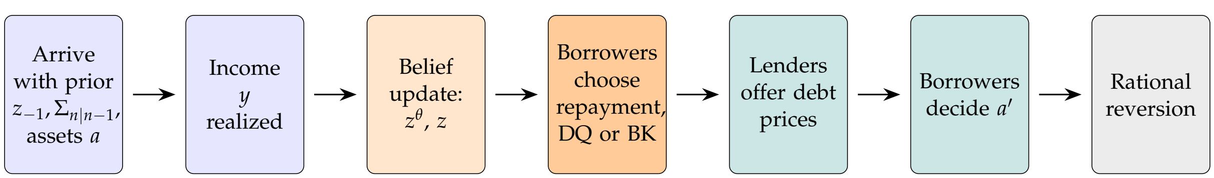

3.5 Value Functions

We write the value functions so that households' distorted beliefs can arise from diagnostic expectations when $\theta > 0$. This formulation accommodates the presence of cognitive biases, where agents may hold biased beliefs. The case of rational learning is nested

[8] This assumption is in line with Bordalo, Gennaioli, Shleifer and Terry (2024) and D'Acunto, Weber and Yin (2024).

within this framework and can be obtained by setting $\theta = 0$. In this case, agents' expectations are updated entirely rationally. Therefore, the model captures both scenarios: One with households' distorted beliefs ($\theta > 0$) and the other with rational expectations ($\theta = 0$).

At each age n, the household decides to either declare bankruptcy, choose delinquency, or repay its debt. The value functions associated with these choices incorporate beliefs about future income and the persistence of income shocks. In particular, given the current residual income realization y, the prior about the persistent component carried over from the previous period $(z_{-1}, \Sigma_{n|n-1})$, and the asset or debt level a, the household's lifetime utility at age n is given by:

$$G_n(y, z_{-1}, a) = \max \Big\{ \underbrace{V_n(y, z_{-1}, a)}_{\text{Repayment}}, \underbrace{B_n(y, z_{-1})}_{\text{Bankruptcy}}, \underbrace{D_n(y, z_{-1}, a)}_{\text{Delinquency}} \Big\}.$$

These value functions evolve with age n, reflecting changes in the income profile, the learning process, and the finite lifespan of households.

The lifetime utility from declaring bankruptcy at age n is:

$$B_n(y, z_{-1}) = u\left(\frac{c}{h_n}\right) + \varrho_n \beta \int G_{n+1}(y', z, 0)\, \mathrm{d}\Phi(y'; z^\theta, \Sigma_{n|n} + \sigma_\epsilon^2 + \sigma_e^2),$$

subject to:

$$\begin{aligned} z &= z_{-1} + K_n(y - z_{-1}), \\ z^\theta &= z_{-1} + (1+\theta) K_n (y - z_{-1}), \\ c &= \underbrace{\exp\left(\ell_n + y\right)}_{\text{Income}} + L_s - \mathbf{c}_{BK}. \end{aligned}$$

Here, $\mathbf{c}_{BK}$ denotes the cost associated with filing for bankruptcy, under which debt is reset to zero. The function $\Phi(y; \mu, \sigma^2)$ represents the cumulative distribution function (CDF) of a normal distribution with mean μ and variance σ^2. The term h_n captures household size in equivalent adults.[9] The variable L_s represents a transfer to households that reflects lenders' profits.[10]

Compared with rational learning ($\theta = 0$), households with diagnostic expectations set the mean of the residual income distribution at z^θ instead of z. Thus, when $\theta > 0$,

[9]The inclusion of a changing family size over the life-cycle is standard in this literature; see, for example, Attanasio and Weber (1995), Livshits, MacGee and Tertilt (2007), and Drozd and Serrano-Padial (2017).

[10]When lenders are assumed also to form diagnostic expectations, their profits are not necessarily zero. Section B in the not for publication online appendix details the algorithm used to compute L_s. Since expectations are not biased in one direction (as in pessimism or optimism), this object is quantitatively very small and has a negligible impact on the equilibrium.

households *overreact* to the current income realization when assessing the persistent income component. Note also that *rational reversion* implies that the next period's prior is z instead of today's posterior z^θ.

Similarly, the lifetime utility from delinquency is given by:

$$\begin{aligned} D_n(y, z_{-1}, a) = u\left(\frac{c}{h_n}\right) + \varrho_n \beta \int \Big[& \gamma G_{n+1}(y', z, 0) \\ & + (1-\gamma) G_{n+1}(y', z, a') \Big] \, \mathrm{d}\Phi(y'; z^\theta, \Sigma_{n|n} + \sigma_\epsilon^2 + \sigma_e^2), \end{aligned}$$

subject to:

$$\begin{aligned} a' &= (1+\eta) a, \\ z &= z_{-1} + K_n (y - z_{-1}), \\ z^\theta &= z_{-1} + (1+\theta) K_n (y - z_{-1}), \\ c &= \underbrace{\min\{\exp(\ell_n + y) + L_s, \tau\}}_{\text{Income net of garnishment}}. \end{aligned}$$

With probability γ, the debt is forgiven, while with probability $1-\gamma$, the debt rolls over at a penalty interest rate η.

If the household chooses to repay its debt, it must settle the outstanding balance and may borrow from the lender's menu of borrowing costs, q_n^θ, or save, in which case $q_n^\theta = 1/(1+r)$. The lifetime utility from repayment is given by:

$$V_n(y, z_{-1}, a) = \max_{a', c} u\left(\frac{c}{h_n}\right) + \varrho_n \beta \int G_{n+1}(y', z, a') \, \mathrm{d}\Phi(y'; z^\theta, \Sigma_{n|n} + \sigma_\epsilon^2 + \sigma_e^2),$$

subject to:

$$\begin{aligned} z &= z_{-1} + K_n (y - z_{-1}), \\ z^\theta &= z_{-1} + (1+\theta) K_n (y - z_{-1}), \\ c &= a - a' q_n^\theta(y, z_{-1}, a') + exp(\ell_n + y) + L_s. \end{aligned}$$

The price of debt is determined by the lender's posterior on the household's persistent income component, the amount of debt taken, and age, as policy functions vary with age.

The repayment decision function is:

$$R_n(y, z_{-1}, a) = \begin{cases} 1, & \text{if } V_n = \max\{V_n, D_n, B_n\}, \\ 2, & \text{if } D_n = \max\{V_n, D_n, B_n\}, \\ 0, & \text{otherwise.} \end{cases}$$

Since default and delinquency are household choices, a risk-neutral lender charges a premium to ensure zero expected profit. Debt prices depend on the household's probability of default or delinquency. Consequently, debt prices are set as

$$q_n^\theta(y, z_{-1}, a') = \frac{1}{1+r} \int \Big[1(R_{n+1}(y', z, a') = 1) \\ + 1(R_{n+1}(y', z, a') = 2)(1-\gamma)(1+\eta) q_{n+1}^\theta(y', z, a'') \Big] \mathrm{d}\Phi(y'; z^\theta, \Sigma_{n|n} + \sigma_\epsilon^2 + \sigma_e^2),$$

subject to:

$$\begin{aligned} z &= z_{-1} + K_n(y - z_{-1}), \\ z^\theta &= z_{-1} + (1+\theta) K_n(y - z_{-1}), \\ a'' &= (1+\eta) a' \end{aligned}$$

Note that when lenders also have diagnostic expectations, the equilibrium implies that lenders must expect to earn zero profits. As these expectations are subjective and lenders do not necessarily make zero profits, we transfer the profits as a lump sum to households. However, since expectations are not biased in one direction (as in pessimism or optimism), this object is quantitatively very small and has a negligible impact on the equilibrium.[11]

4 Calibration/Estimation

This section outlines the calibration and estimation strategy. We begin by estimating the diagnostic expectation parameter θ using income expectations data from the New York Fed's Survey of Consumer Expectations. Second, we use the same survey to calibrate the level of uncertainty households have about the initial value of z before entering the labor market. Then, we calibrate a set of common parameters across all four models using values drawn directly from the data or established in the literature. Finally, we estimate three model-specific parameters for each of the four frameworks by targeting moments related to household debt and financial distress, minimizing the distance between model-implied and empirical moments.

[11] See Section B in the not for publication online appendix for details on the algorithm used to compute it.

4.1 Estimation of the Diagnostic Expectation Parameter

To estimate the parameter θ, we use data from the labor market module of the Survey of Consumer Expectations (SCE), administered by the Federal Reserve Bank of New York since June 2013. This module is fielded every four months, with each respondent participating in three consecutive survey waves. Each wave consists of a rotating cross-section of U.S. adults, recruited through address-based sampling to construct a nationally representative panel. Post-stratification weights are applied to align the sample with the Current Population Survey (CPS) margins along key demographics: Age, gender, education, race, and region.

We construct our measure of forecast error using responses to two questions in the Labor Market Survey. First, we use question *L3*, which asks respondents about their current total annual income. Second, we use question *oo2e2*, which elicits expectations about their annual income four months from now.[12] We restrict our sample to waves conducted after the second question was introduced in March 2015. The original wording of both questions is reproduced in Section B.1 in the Appendix.

We denote realized income in period t as $\tilde{y}_t$, and the corresponding expectation formed in period $t-1$ as $\tilde{E}_{t-1}[\tilde{y}_t]$, where the tilde on the expectation reflects the possibility of distortion. The forecast error is then defined as the difference between the realized and expected income:

$$\text{Forecast Error}_{it} = \tilde{y}_{it} - \tilde{E}_{t-1}[\tilde{y}_{it}].$$

Our main empirical analysis follows Ma, Ropele, Sraer and Thesmar (2020) by examining the autocorrelation of forecast errors using the following specification:

$$\text{Forecast Error}_{it+1} = \beta_0 + \beta_1 \,\text{Forecast Error}_{it} + {\beta_2}'\mathbf{X}_{it} + u_t + \xi_{it+1}, \tag{2}$$

where $\mathbf{X}_{it}$ denotes a vector of observable household characteristics, including gender, education, employer type, and industry, sourced from the core SCE survey. The term u_t captures time fixed effects.[13]

If respondents form rational expectations about their income, forecast errors at time $t+1$ should be uncorrelated with any information available at time t. However, the results in the first column of Table 1 indicate a statistically significant negative autocorre-

[12] Ganong, Noel, Patterson, Vavra and Weinberg (2025) use high-frequency administrative data to show that most U.S. workers experience substantial month-to-month fluctuations in pay, even within ongoing employment relationships. This suggests that, even over a four-month interval, households in our sample are very likely to experience income shocks and update their beliefs accordingly.

[13] Because of data limitations, each household contributes only one observation to the regression. As a result, we cannot include household-level fixed effects.

lation in forecast errors, suggesting that household expectations deviate from rationality. In particular, the negative coefficient implies that households tend to overreact to income news—an empirical pattern consistent with the presence of diagnostic expectations. In Section B.3 in the Appendix, we include additional empirical specifications to validate the robustness of our results. Additionally, in Section B.4 in the Appendix, we consider regression results by age and education group and find little evidence of differences in overreaction across these groups.

Table 1: Forecast Error Regression Results

Dependent Variable	Forecast Error at $t+1$		
	SCE	$\theta = 1.145$	$\theta = 0$
Forecast Error at t	-0.353^{***}	-0.353^{***}	-0.003^{***}
	(0.117)	(0.001)	(0.001)
Controls	Yes	Yes	Yes
Observations	622	5,399,616	5,399,660
R^2	0.230	0.125	0.000

Notes: Standard errors in parentheses; *** $p < 0.01$, * $p < 0.10$.

We estimate the diagnostic expectation parameter θ by simulating the model to match the observed autocorrelation of forecast errors. Specifically, we simulate a four-month income process for a large number of households, using the same RIP as in our annual model, adjusted to a four-month frequency, as detailed in Section B.2 of the Appendix. We then estimate Equation (2) using the simulated data, applying the same sample restrictions as in the empirical analysis. It also includes age as a control variable, which is the variable available in the model of those included in the data. Note that in the model, the counterpart to the data's forecast of income at $t+1$ is computed based on the distorted posterior formed at time t, before rational reversion. The inferred value of θ to reproduce the autocorrelation is 1.145, which aligns closely with values reported in the literature.[14]

The second and third columns of Table 1 show the implied autocorrelation of forecast errors in the model with DE ($\theta = 1.145$) and RL ($\theta = 0$). When $\theta = 1.145$, the model reproduces the empirical autocorrelation precisely; by contrast, setting $\theta = 0$ yields a forecast-error autocorrelation that is essentially zero in magnitude, suggesting no eco-

[14] D'Acunto, Weber and Yin (2024) estimates $\theta = 1.618$ for household income expectations using Chinese survey data. Bordalo, Gennaioli, Ma and Shleifer (2020) report values between 0.5 and 1 for professional analysts' forecasts of macroeconomic variables. Bordalo, Gennaioli, Shleifer and Terry (2024) estimates $\theta = 0.913$ for expectations about future revenue among U.S. public firm managers.

nomically meaningful persistence in forecast errors absent diagnostic expectations.

4.2 Calibration of Initial Beliefs

In all four economies, the initial distribution of income—composed of two normally distributed components, $y = z + e$—is determined by the calibration of $\sigma_{z,0}$ and σ_e, both taken from Kaplan and Violante (2010). In the RL and DE economies, after observing the first realization of income y, households update their beliefs about the persistent component z. The extent of this adjustment, captured by the Kalman gain, depends critically on how uncertain households are about the initial value of z. This subsection describes how we calibrate that uncertainty in the RL and DE economies.

We take the position that not all of the initial dispersion in z is necessarily uncertain. Agents may have information about how well they will perform in the labor market that the econometrician does not capture. Thus, the initial value of z has an individual-specific component μ_i that agents know in advance. This means that in the first period in the labor market,

$$z_{i,0} = \mu_i + \kappa_i,$$

with both components being normally distributed with mean zero, and variances σ_μ^2 and σ_κ^2. Note that this is a generalization that nests the case in which $\sigma_\mu^2 = 0$ so prior beliefs follow the same distribution for all agents, $\mathcal{N}(0, \sigma_{z,0}^2)$.

We calibrate σ_μ^2 and σ_κ^2 such that the prior beliefs upon entering the labor markets:

1. align with the objective distribution of z_{n_0}, which means $\sigma_{z,0}^2 = \sigma_\mu^2 + \sigma_\kappa^2$, and

2. replicate the empirical life-cycle pattern of the variance of forecast errors.

In particular, we calibrate σ_μ^2 and σ_κ^2 using data on the change in the variance of income forecast errors during the first years after labor market entry. If $\sigma_\mu^2 = 0$, forecast errors are very large at young ages, since households begin with little information about their z. As σ_μ^2 increases, σ_κ^2 must fall to satisfy $\sigma_{z,0}^2 = \sigma_\mu^2 + \sigma_\kappa^2$, reducing forecast errors. In the limiting case where $\sigma_\mu^2 = \sigma_{z,0}^2$, households know their exact z upon entering the labor market.

We compute forecast errors using the same approach as in the previous subsection. Then, we calculate the average variance of forecast errors for ages 23-26 relative to that for ages 27-30. We use multiple years to increase precision, as individual age groups contain few observations. The age of 23 marks the first cohort with sufficient observations after labor market entry. The resulting ratio is 1.045.

Figure 5 plots the evolution of prior uncertainty and subjective Kalman gains as functions of age. The green dashed line illustrates the case in which all the dispersion in the

Figure 5: Subjective Kalman Gains and Prior Uncertainty as Functions of Age

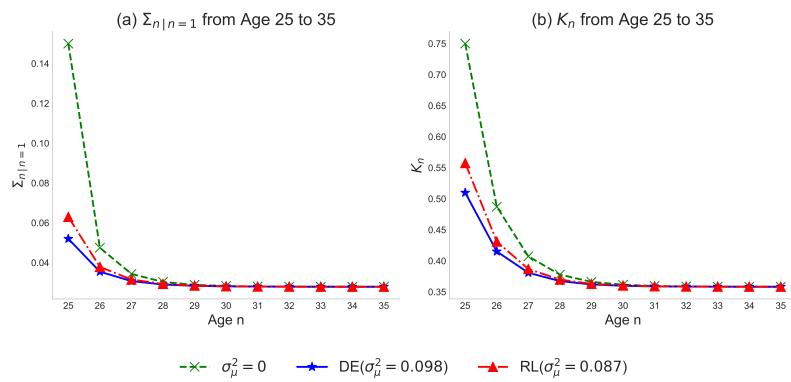

initial z is initial uncertainty ($\sigma_\mu^2 = 0$). In that case, uncertainty about the persistent component is very high, declines rapidly, and stabilizes around age 30. With these values, the variance of forecast errors for ages 23-26 relative to ages 27-30 equals 1.11 in the RL economy and 1.18 in the DE economy.

Matching the empirical ratio instead requires setting $\sigma_\mu^2 = 0.098$ and $\sigma_\kappa^2 = 0.052$ in the DE economy, and $\sigma_\mu^2 = 0.087$ and $\sigma_\kappa^2 = 0.063$ in the RL economy. The resulting evolution of uncertainty about z and the corresponding Kalman gains are also shown in Figure 5. Under these calibrations, the decline in uncertainty and the variation in the Kalman gain over the life-cycle are substantially smaller.

4.3 Calibration of Common Parameters

Table 2 includes parameters that are either well-established in the literature or based on empirical estimates. The risk-free rate ($r = 0.03$) reflects long-term averages. The risk aversion parameter is set to the standard value ($\sigma = 2$). Parameters governing income processes, such as the standard deviations of persistent and transitory shocks ($\sigma_\epsilon = 0.10$, $\sigma_e = 0.22$), follow from Kaplan and Violante (2010). The initial distribution of z is also taken from Kaplan and Violante (2010) ($\mu_{z,0} = 0, \sigma_{z,0} = 0.39$), as mentioned above. The debt roll-over rate ($\eta = 0.20$) takes the same value as in Livshits, MacGee and Tertilt (2007), and we also set the dismissal probability of debt γ as 0.23, which we take from Athreya, Mustre-del Río and Sánchez (2019). The bankruptcy cost ($\mathbf{c}_{BK} = 0.165$) and the normalized family size profile h_n are obtained directly from the data, and the life-cycle income profile ℓ_n and survival rates after retirement ϱ_n are taken from Kaplan and Violante (2010). In Section C in the Appendix, we plot h_n, ϱ_n, and $exp(\ell_n)$ as functions of age.

Table 2: Calibrated Parameters

Parameter	Value	Meaning	Basis
T	65	Retirement age	Standard
σ	2.00	Risk aversion	Standard
r	0.03	Risk-free rate	Standard
γ	0.23	DQ dismissal probability	Athreya, Mustre-del Río and Sánchez (2019)
η	0.20	Debt roll-over rate	Livshits, MacGee and Tertilt (2007)
σ_ϵ	0.10	Std of persistent shock	Kaplan and Violante (2010)
σ_e	0.22	Std of transitory shock	Kaplan and Violante (2010)
$\mu_{z,0}$	0	Initial mean of persistent shock	Kaplan and Violante (2010)
ϱ_n	Fig. 15	Survival rate	Kaplan and Violante (2010)
ℓ_n	Fig. 15	Life-cycle income profile	Kaplan and Violante (2010)
h_n	Fig. 15	Household size in equiv. adults	SCF & Buhmann et al. (1988)
$\mathbf{c}_{BK}$	0.165	Bankruptcy cost	BK cost to mean income

4.4 Estimation of Key Parameters by Indirect Inference

The remaining three parameters, $\{\beta, \tau, a_0\}$, are estimated by minimizing the distance between model-generated and empirical moments. The estimation targets three sets of moments: The share of households with debt, the share in delinquency, and the share in bankruptcy. Each moment is calculated across three age groups to capture the life-cycle patterns of borrowing and financial distress. Debt moments are sourced from the Survey of Consumer Finances (SCF), while delinquency and bankruptcy moments are constructed from the FRBNY Consumer Credit Panel/Equifax. Section C in the Online Appendix provides detailed information on the data sources, sample selection, and the construction of the targeted moments.

Since the three parameters $\{\beta, \tau, a_0\}$ interact within the model, we estimate them jointly using the Simulated Method of Moments. To gain insight into identification, consider how each parameter maps to the targeted moments. The discount factor β governs household impatience and thus influences borrowing behavior. Both the level and slope of the age profile for the share of households with debt are informative about the appropriate value of β. The consumption garnishment parameter in delinquency, τ, determines the relative attractiveness of delinquency compared with repayment and bankruptcy. It is therefore primarily identified by the overall incidence of delinquency and the relative frequency of delinquency versus bankruptcy. Finally, the initial asset position a_0 plays a key role in matching the financial distress rates of younger households and is especially influential in shaping the slopes of both delinquency and bankruptcy profiles throughout the life cycle.

Table 3 reports the estimated parameters. The estimated value of β is somewhat lower than standard values in quantitative macroeconomic models but higher than those typically used in default models (more on this later). The estimate of τ suggests that income garnishment during delinquency is only applicable to households with very high incomes, thereby discouraging such households from becoming delinquent. Finally, the positive value of the initial asset level a_0 implies that, even though households are impatient, it takes a few years for them to accumulate debt and potentially enter financial distress.[15]

Table 3: Estimated Parameters for the Benchmark DE Model with $\theta = 1.145$

	β	τ	a_0
Value	0.859	3.942	2.966
Std. Errors	(0.002)	(0.179)	(0.338)

Table 4 reports the fit of the model to the targeted moments. The model is disciplined by only three estimated parameters to jointly match nine moments describing the life-cycle profiles of delinquency, bankruptcy, and debt participation. Given this parsimonious parameterization, the fit is naturally not perfect. Nevertheless, the model reproduces the incidence of financial distress remarkably well. Delinquency rates are matched almost exactly for the youngest age group (11.87% in the data versus 11.85% in the model) and remain within 0.6 percentage points of the data for the remaining groups. Likewise, bankruptcy rates, despite being rare events, are matched closely across all ages, with deviations of at most 0.11 percentage points. These results indicate that the model successfully captures both the level and life-cycle profile of financial distress.

The main discrepancy concerns debt participation. While the model closely matches the fraction of young households in debt (40.8% versus 41.5% in the data), it increasingly underpredicts debt holding among middle-aged and older households, particularly for ages 45–54 (27.9% versus 44.7%). This pattern suggests that households in the model repay unsecured debt more rapidly over the life cycle than observed in the data. Importantly, this shortcoming does not translate into a poor fit for financial distress, indicating that the model successfully identifies the subset of borrowers most likely to become

[15] An alternative calibration strategy, followed in Athreya, Mustre-del Río and Sánchez (2019), sets the initial distribution of assets to match SCF data. The results are qualitatively similar but imply a higher incidence of financial distress among young households. Since young agents may receive direct transfers from their parents, we prefer to infer the initial asset level that is consistent with the observed initial incidence of financial distress.

Table 4: Fit of Targeted Moments

Moments	Data	DE Model ($\theta = 1.145$)
% in DQ, age 25–34	11.87	11.85
% in DQ, age 35–44	11.50	12.07
% in DQ, age 45–55	9.49	9.76
% in BK, age 25–34	0.51	0.44
% in BK, age 35–44	0.62	0.56
% in BK, age 45–54	0.51	0.40
% in Debt, age 25–34	41.49	40.79
% in Debt, age 35–44	44.55	39.46
% in Debt, age 45–54	44.70	27.89

delinquent or declare bankruptcy. In other words, although the model understates the extensive margin of borrowing later in life, it captures the intensive margin most relevant to understanding household financial distress.

The model effectively captures the incidence of financial distress (delinquency and bankruptcy) throughout the life cycle. The model also accurately captures the percentage of individuals in debt in the age groups 25-34 and 35-44. The most significant discrepancies appear among older households, mainly because the model may overlook reasons to borrow at older ages beyond insurance against income shocks. However, given that only three parameters were chosen to fit nine moments, the model's ability to match key financial outcomes with such precision underscores its robustness in capturing the core mechanisms driving default and borrowing behavior.

5 The Role of Shock Misperception in Financial Distress

In this section, we analyze the role of income shock misperception in financial distress, first in partial equilibrium and then in general equilibrium. To do so, we maintain the same parameters as estimated for the benchmark DE economy and compute two alternative economies: One with rational learning and the other with full information. We also examine specific examples to understand *how* income-shock misperception affects financial distress.

5.1 Partial-Equilibrium Results

This section compares alternative frameworks that hold all parameters fixed (estimated from the DE economy) and keep the risk-free interest rate constant. The results reveal that incorporating information friction and belief distortions significantly increases the incidence of both informal (delinquency) and formal (bankruptcy) default relative to the economy with complete information. We also include an economy where households have diagnostic expectations, but lenders have rational learning.

Table 5 compares the level of financial distress across alternative economies. The first column represents our benchmark economy, where households and lenders have diagnostic expectations with $\theta = 1.145$. We use the columns to the right of that one to estimate the share of financial distress attributable to each model feature. The second column presents results for the case where lenders have rational learning rather than diagnostic expectations. The differences relative to the benchmark case—where both households and lenders have DE—are generally slight: slightly negative for DQ and slightly positive for BK. Overall, DQ decreases by 0.63 percentage points, while BK increases by 0.21 percentage points. These modest changes suggest that lenders' beliefs play a limited role in determining the overall incidence of financial distress.

The third column reports results for the RL economy. Comparing this economy with the case in which only households have diagnostic expectations isolates the effect of households' belief distortions. Relative to the household-only DE economy, the RL economy exhibits a markedly lower incidence of both delinquency and bankruptcy across all age groups. For households aged 25–54, removing diagnostic expectations from households reduces the delinquency rate by 1.34 percentage points (from 10.60 to 9.26 percent) and the bankruptcy rate by 0.31 percentage points (from 0.68 to 0.37 percent). Thus, conditional on lenders having rational learning, household diagnostic expectations account for a substantial share of financial distress, with an especially pronounced effect on bankruptcies.

The fourth column of Table 5 reports the results for the FI economy, which can be compared with those under RL to isolate the effects of information frictions. Removing these frictions substantially reduces DQ rates across all age groups: For each of the three age groups, the DQ rate falls by more than 3 percentage points. Among households aged 25–54, the overall DQ rate declines from 9.26 percent under RL to 5.91 percent under FI—about 30 percent of the benchmark DE economy value. By contrast, the effect on BK rates is modest. For the same age range, the overall BK rate decreases only slightly, from 0.37 percent with RL to 0.32 percent with FI. The significant reduction in DQ and the minor change in BK highlight that information frictions primarily raise financial distress

Table 5: Misperceptions and Financial Distress, Partial Equilibrium

Statistic	Diagnostic Expectations Hhds & Lenders	Diagnostic Expectations Only Hhds	Rational Learning	Full Information
% in DQ, age 25–34	11.85	10.77	9.73	6.63
% in DQ, age 35–44	12.07	11.11	9.13	5.52
% in DQ, age 45–54	9.76	9.92	8.91	5.58
% in BK, age 25–34	0.44	0.62	0.34	0.31
% in BK, age 35–44	0.56	0.81	0.41	0.30
% in BK, age 45–54	0.40	0.62	0.36	0.37
Overall % in DQ, age 25–54	11.23	10.60	9.26	5.91
Overall % in BK, age 25–54	0.47	0.68	0.37	0.32

by causing more households to miss debt payments and fall into delinquency, rather than by directly increasing bankruptcies.

Overall, the sequential comparisons indicate that lenders' diagnostic expectations play only a minor role in determining financial distress, whereas households' diagnostic expectations and information frictions are quantitatively much more important. Comparing the benchmark economy with the full-information economy shows that informational frictions and belief distortions together account for a substantial fraction of financial distress. Relative to the benchmark, the full-information economy reproduces only 53 percent of delinquency and 68 percent of bankruptcy, implying that the remaining financial distress is generated by informational frictions and belief distortions rather than income risk alone.

5.2 General-Equilibrium Results

One possible reason for the increased financial distress under RL and DE is that these economies generate higher levels of household borrowing. To isolate the effect of beliefs from changes in aggregate borrowing, this section presents an alternative equilibrium in which the interest rate adjusts to keep aggregate household asset holdings fixed across economies, following the spirit of Huggett (1993). Specifically, we first compute the benchmark DE economy and calculate household asset holdings, $\bar{a} \equiv \int a \, d\Gamma^{DE}$, where the integral is taken over the stationary distribution of households and includes both borrowers ($a < 0$) and savers ($a > 0$). We then consider the j economy ($j =$ RL or FI), holding all preference and technology parameters fixed, and choose the interest rate in each economy so that $\int a \, d\Gamma^j = \bar{a}$.

Thus, aggregate net asset holdings are identical across economies, allowing us to isolate the effects of differences in beliefs and information from differences in the overall level of household borrowing and saving.

Comparing the general equilibrium results (Table 6) with the partial equilibrium results (Table 5) reveals that the results for the cases in which lenders are given rational learning are almost identical.

Removing diagnostic expectations from households reduces overall financial distress more. The most significant amplification is in terms of DQ. The DQ rate declines by 1.55 percentage points in general equilibrium, rather than 1.34 percentage points in partial equilibrium (a 16 percent amplification). The amplification for the BK rate is much smaller, only 5 percent. The impact of removing information frictions is also amplified for the DQ rate, which declines by 3.47 percentage points in GE and 3.35 percentage points in PE. The difference in the BK is minimal, as it declines 0.05 percentage points in PE and 0.04 in GE.

As shown in Table 6, removing DE for lenders does not change the equilibrium interest rate. Removing DE from households increases the equilibrium rate from 3 percent in the benchmark economy to 3.53 percent, and further eliminating information frictions raises it to 4.03 percent in the FI economy. The rise in interest rates in the general equilibrium case is a consequence that, without the interest rate adjustment (the partial equilibrium case), aggregate savings decline.

Table 6: Misperceptions and Financial Distress, General Equilibrium

Statistic	Diagnostic Expectations Hhds & Lenders	Only Hhds	Rational Learning	Full Information
% in DQ, age 25–34	11.85	10.76	9.51	6.33
% in DQ, age 35–44	12.07	11.10	8.91	5.26
% in DQ, age 45–54	9.76	9.89	8.66	5.11
% in BK, age 25–34	0.44	0.62	0.35	0.30
% in BK, age 35–44	0.56	0.81	0.37	0.29
% in BK, age 45–54	0.40	0.62	0.34	0.37
Overall % in DQ, age 25–54	11.23	10.58	9.03	5.56
Overall % in BK, age 25–54	0.47	0.68	0.36	0.32
Interest Rate (%)	3.00	3.02	3.53	4.03

6 Mechanisms at Work

This section uses the calibrated model to illustrate the key mechanisms at work through a series of figures. We first examine how alternative assumptions about information and expectations shape the consumption function. We then explore how income shocks, information frictions, and diagnostic expectations push households into financial distress.

6.1 Consumption: Income Shocks, Information, and Expectations

To illustrate how information, learning, and expectation formation shape consumption choices, Figure 6 plots consumption as a function of income across four cases.[16]

The red and green lines correspond to the FI model; in the red case, income variation arises from changes in the persistent component z; in the green case, it arises from changes in the transitory component e. As expected, consumption responds more strongly to persistent shocks, as reflected in the steeper slope of the red dotted line relative to the green dash-dotted line. The blue and yellow lines represent environments in which households must infer whether income changes are transitory or persistent. Under DE ($\theta > 0$, blue solid line), households assign greater weight to the change being persistent. As a result, consumption reacts more strongly than under RL (yellow dashed line). This amplification highlights how biased inference magnifies the response of consumption to income shocks.

Figure 6: Consumption as a Function of Income in the Alternative Models

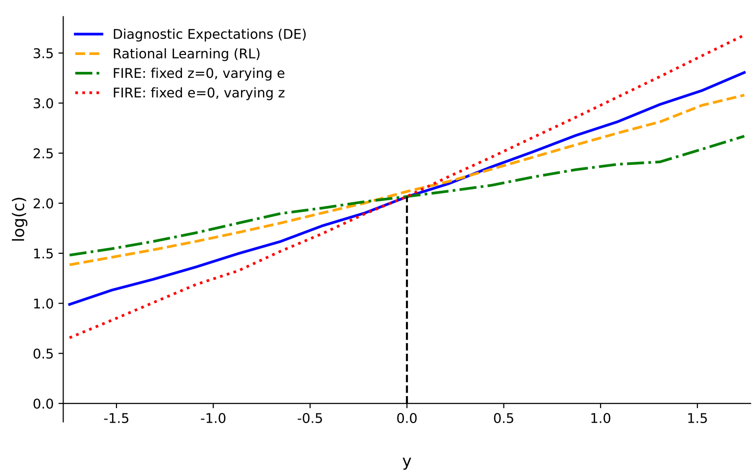

[16] Age is fixed at 40, and the values of z and e are each held at their median when the other varies. The value of $a > 0$.

6.2 Financial Distress: Income Shocks, Information, and Expectations

To illustrate how income shocks, information frictions, and diagnostic expectations can push households into financial distress, the following three subsections present illustrative examples of the model's dynamics. We use the parameters estimated from the benchmark DE economy and conduct the experiments in partial equilibrium. Each example isolates one mechanism: Income shocks, information frictions, or diagnostic expectations. Households in four scenarios—FI (□), RL (○), households only DE (△), and DE ($\star$)—start with zero debt and face the same sequence of income shocks.

6.2.1 How Earnings Shocks Trigger Financial Distress

Figure 7 displays an example in which income shocks alone trigger financial distress. Initially, households receive a large positive persistent income shock ($z_{25} \approx 0.17$) and a small positive transitory shock ($e_{25} \approx 0.04$). Expecting income to remain high, all households borrow ($a_{26} < 0$). Because lenders in the FI economy fully recognize that the rise in income is primarily due to the persistent component, interest rates are lowest in that economy. In contrast, lenders in the RL and DE economies assign some probability that the income surge reflects a transitory factor and therefore charge higher rates. Among those two, lenders in the DE economy offer slightly better terms, as they also tend to overreact to positive income news.

Figure 7: Labor Earnings Shocks and Financial Distress

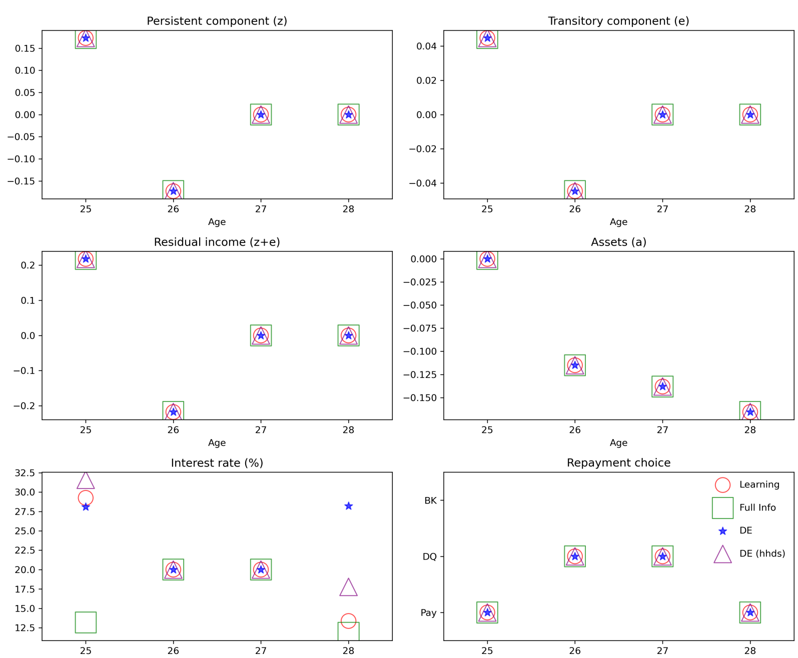

In the following year (age 26), the persistent shock reverses sharply ($z_{26} \approx -0.16$), accompanied by a small negative transitory shock ($e_{26} \approx -0.04$). Faced with an unexpectedly steep income decline, all households skip their debt payments and become delinquent on their accounts. Debt accumulates rapidly, as unpaid balances are rolled over at a punitive 20 percent penalty rate. At age 27, income partially recovers ($z_{27} = 0, e_{27} = 0$), but the high debt burden leaves households' disposable income insufficient to repay, prompting continued delinquency to smooth consumption. By age 28, with income continuing at the average levels, households are finally able to exit delinquency.

This example illustrates how a sharp reversal in income can lead households—regardless of the information structure—to follow similar borrowing paths and experience synchronized episodes of financial distress.

6.2.2 How Misperception of Income Shocks Leads to Financial Distress

The example in Figure 8 illustrates how misperceptions of income shocks—or, more broadly, information frictions—can lead households into financial distress. Initially, the individual experiences no persistent income shock ($z_{25} = 0$) but receives a positive transitory shock ($e_{25} \approx 0.2$). Consequently, residual income at age 25 is higher than anticipated ($z_{25} + e_{25} \approx 0.2$).

As shown in the middle-right panel, debt choices at age 26 differ markedly between the learning-based and FI economies. Under FI, the household recognizes the transitory nature of the income shock and saves part of it. In contrast, households in learning-based economies begin to borrow, as both households and lenders overestimate future income prospects. This borrowing raises interest rates because of higher perceived default risk, as depicted in the bottom-left panel. However, the increase is smaller under DE because lenders' overreaction to positive income news partially offsets the effect.

In the following year (age 26), households face a negative transitory shock ($e_{26} \approx -0.2$) with no change in the persistent component ($z_{26} = 0$). Both households and lenders in learning-based economies revise their income expectations downward. Lenders become reluctant to extend new credit that would allow households to roll over their debt. As a consequence, households resort to the contractual rollover at the punitive 20 percent interest rate that applies during delinquency. Meanwhile, the household with FI, having initially avoided debt, now borrows to maintain stable consumption.

At age 27, income returns to its average level ($z_{27} = 0, e_{27} = 0$), but households have accumulated debts that are too large to be serviced by their current income, resulting in continued delinquency as a form of consumption smoothing. By age 28, income continues at the average level and now enables households to repay their obligations and exit delinquency.

Figure 8: Learning the Persistence of Income and Financial Distress

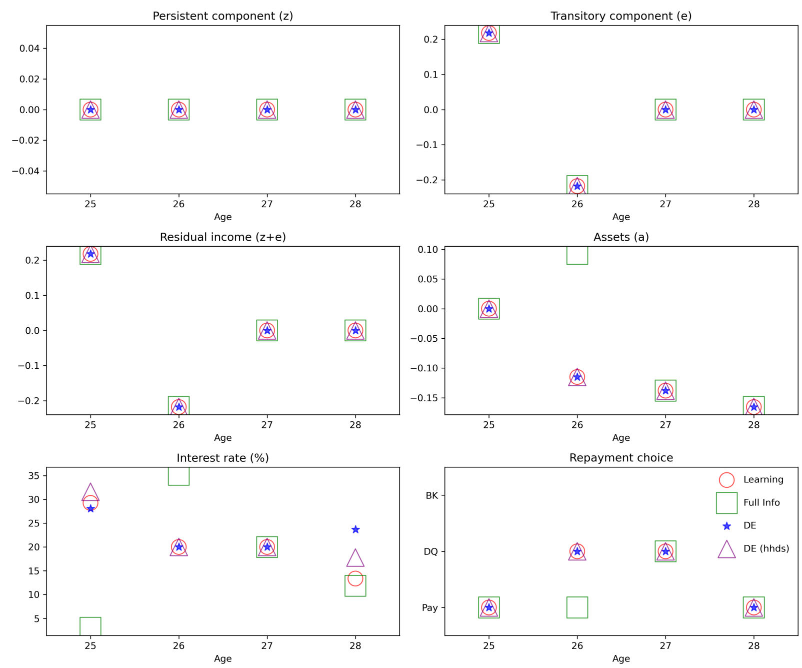

This example illustrates how information frictions lead both households and lenders in learning-based economies to misperceive the persistence of income shocks, resulting in excessive borrowing and a higher frequency of delinquency compared with the FI economy.

6.2.3 Exaggerated Optimism and the Path to Financial Distress

The last scenario, presented in Figure 9, illustrates how diagnostic expectations amplify financial distress relative to the RL and FI economies. The example begins at age 25 with no persistent income shock ($z_{25} = 0$) and a sizable positive transitory shock ($e_{25} \approx 0.4$). As a result, residual income that year is unusually high ($z_{25} + e_{25} \approx 0.4$), well above prior expectations.

The middle-right panel highlights stark differences in debt decisions across the four economies. Under FI, agents recognize that the income boost is temporary and therefore save aggressively. Under RL, moderate optimism about future income leads to consumption exceeding current income, resulting in modest borrowing. Under DE, however, exaggerated optimism fuels excessive consumption financed by borrowing, despite already elevated income. As expected, when both households and lenders expect high income to persist, borrowing is a bit larger than when only households have diagnostic expectations.

In the following year (age 26), agents do not experience a persistent income shock ($z_{26} = 0$) but a negative transitory shock ($e_{26} \approx -0.2$). Households with FI or RL who hold little or moderate debt borrow to smooth consumption. In contrast, households under DE, burdened by substantial debt and facing low bond prices, choose delinquency.

At age 27, with income slightly improving ($z_{27} = 0$, $e_{27} = 0$), households in the DE economy remain in delinquency because of the weight of accumulated debt. Those in the RL and FI economies, now also carrying some debt, roll it over through temporary delinquency. At age 28, when income remains unchanged ($z_{28} = 0$, $e_{28} = 0$), households in economies with RL and FI, and households with only DE, can repay and resume normal borrowing. In contrast, in the economy in which households and lenders hold DE, households resort to bankruptcy as a way to reset.

Figure 9: Exaggerated Optimism and Financial Distress

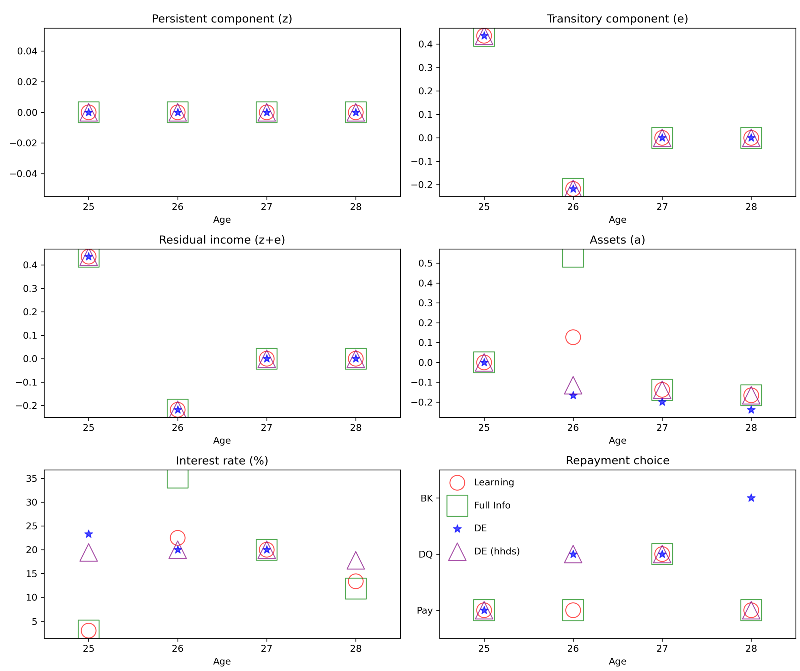

This example illustrates how diagnostic expectations, by amplifying misperceptions about income prospects, can lead to excessive borrowing, delinquency, and ultimately bankruptcy, thereby exacerbating financial distress over the life cycle. It also shows that differences in lenders' expectations can potentially shape households' financial outcomes.

7 Distinguishing Model Implications

This section highlights two key distinguishing implications across the models. First, we estimate the three alternative economies to match the same data and show that they imply markedly different values for households' discount factors—a central differentiating feature of the models. Second, we identify the correlation between income and interest rates as another distinctive implication and compare it with its empirical counterpart.

7.1 Implications for Households' Discount Rate

In this section, we estimate the three alternative models. Table 7 reports the parameter estimates.[17] In the FI model, the discount factor is estimated at $\beta = 0.783$, implying an annual discount rate of $1/\beta - 1 \approx 27.7$ percent. This rate is unusually high relative to typical annual interest rates, indicating that, without learning, households would need to be implausibly impatient to replicate the observed levels of debt distress.

In the RL economy, the estimated discount factor rises to $\beta = 0.813$, corresponding to an annual discount rate of 23.1 percent. As shown earlier, in the DE economy, the estimated discount factor increases further to $\beta = 0.859$, implying a discount rate of 16.1 percent. We also estimate a mixed specification in which households exhibit diagnostic expectations while lenders learn rationally. The resulting estimate, $\beta = 0.851$, reinforces the conclusion that households' diagnostic expectations primarily account for the benchmark DE model's performance.

Overall, we find that learning-based models—and in particular the DE model—help rationalize more realistic degrees of household impatience. Moreover, the discount rate under DE is not much higher than observed borrowing costs. Based on Federal Reserve data, the average nominal APR on credit card plans between 1995 and 2025 is approximately 14.4 percent, while average inflation over the same period is 2.5 percent, implying a real APR of about 11.8 percent.[18]

7.2 Implications for the Correlation of Income and Interest Rates

One key prediction shared by all models is that, conditional on debt, households with higher income pay lower interest rates. The question, however, is how the strength of

[17] Table 13 in Section E of the Appendix compares the ability of the estimated models to replicate key empirical moments. Despite being overidentified, all models perform well at matching the targeted moments of financial distress.

[18] Nominal APR and inflation data are from FRED.

Table 7: Model Parameters Estimated

Parameter	Full Information	Rational Learning	Diagnostic Expectations Hhds & Lenders	Only Hhds
β	0.783	0.813	0.859	0.851
	(0.006)	(0.006)	(0.002)	(0.003)
τ	4.408	3.762	3.942	3.785
	(0.069)	(0.029)	(0.179)	(0.027)
a_0	2.964	2.096	2.966	2.501
	(0.371)	(0.269)	(0.338)	(0.260)

Notes: Standard errors in parentheses.

this relationship varies across models. A clear distinction emerges using binned scatterplots in Figure 10. In learning-based models, income is observable and directly shapes lenders' assessment of credit risk. Higher current income signals a higher expected future income through beliefs about the persistent income component, given a constant level of borrowing (as controlled for in the plots). This lowers the perceived probability of default and, in turn, the interest rate offered by lenders. In contrast, in the FI model, interest rates depend only on the persistent component z, while income also includes a transitory component e. Because e affects current income but not the interest rate, the relationship between income and borrowing costs is weaker in the FI model than in the learning-based economies.

Table 8 reports the results from regressions using simulated data generated by each of the model variants. The estimated standardized coefficients closely align with the patterns observed in the scatterplots. In the RL model, a one-standard-deviation increase in income is associated with a 0.21-standard-deviation decline in interest rates. Under DE only for households, the effect is stronger (0.39), and when both households and lenders have diagnostic expectations, it increases to 0.57. In contrast, the coefficient in the FI model is near zero (-0.02).

We evaluate this prediction using data from the SCF, which provides detailed information on household income, assets, liabilities, and borrowing costs. We keep the data details in Section C of the Online Appendix. Figure 11 highlights the main empirical finding of this section: In the SCF data, there is a strong negative relationship between interest rates and income. This robust negative association suggests that, conditional on debt, households with higher income tend to face lower borrowing costs.

To further assess this relationship, we complement the evidence with a formal regres-

Figure 10: Relationship Between Interest Rates and Income in the Models

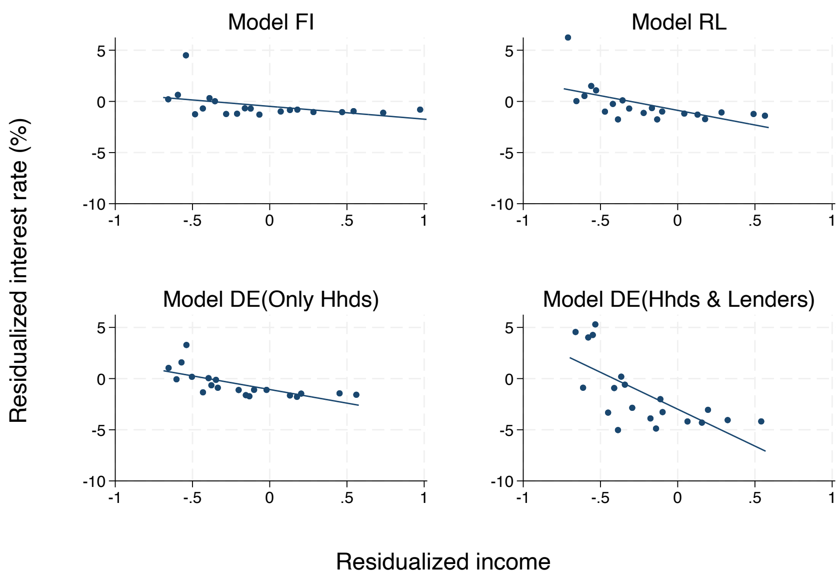

Note: The figure presents a binned scatterplot of residualized interest rates against residualized total income, controlling for age fixed effects, debt, and debt squared. The sample is winsorized at the 10th and 90th percentiles of the income distribution and includes households with individuals aged 25 to 65 who are indebted. Debt balances weight observations to reduce the influence of outliers with very low debt levels.

Table 8: Regressions of Interest Rates on Income, Across Models

	Full Information	Rational Learning	Diagnostic Expectations Hhds & Lenders	Only Hhds
Beta on Income	-0.021	-0.206	-0.567	-0.393

Note: Each column shows the standardized coefficient of interest rates on income. They represent the change in the interest rate (in terms of its standard deviation) for a one-standard-deviation increase in income. Controls include age, FE, and the quadratic amount borrowed. The sample consists of 20M simulated non-delinquent agents, winsorized at 1 and 99% of income.

Figure 11: Interest Rate and Income across Years

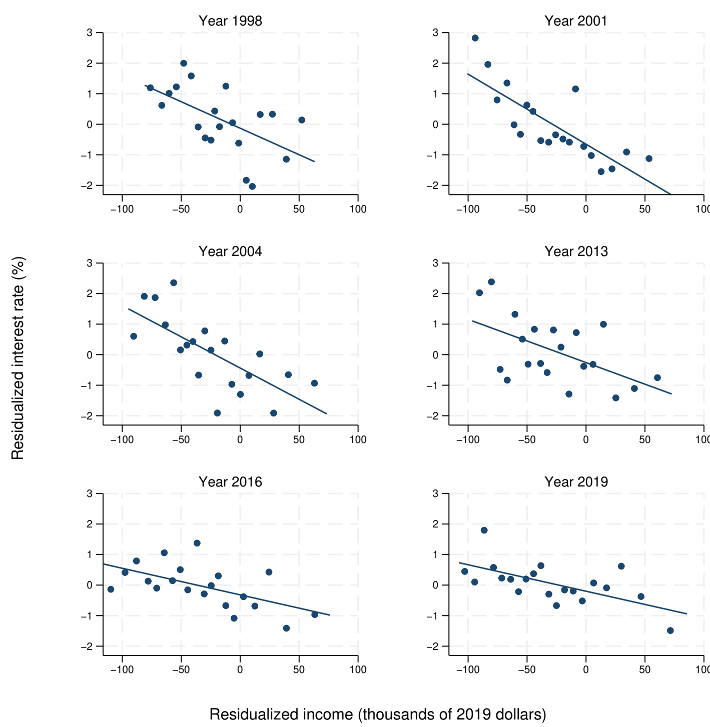

Note: This binned scatterplot displays residualized interest rates against income for households (ages 25-65) with credit card debt. The SCF-weighted sample is winsorized by income and controls for age and credit card debt (including its square).

sion analysis. The regression specification includes controls for age dummies, credit card debt, and its square to account for differences in both life-cycle borrowing patterns and debt size on interest rates. We estimate the relationship using two alternative measures of income—total income and labor earnings—to assess the robustness of the results across income definitions.[19] The results, reported in Table 12 in Section D of the Appendix,

[19] All regressions incorporate SCF sampling weights and include only households with positive credit

confirm the patterns shown in the scatterplots. Conditional on debt, higher household income is consistently associated with lower credit card interest rates. Quantitatively, a one-standard-deviation increase in income corresponds, on average across years, to a decline in interest rates of 0.30 standard deviations when using labor earnings and 0.43 when using total income. This empirical relationship aligns closely with the predictions of learning-based models, which hold that current income signals future repayment capacity and shapes perceived credit risk.

8 Policy Analysis

This section uses the calibrated model to evaluate alternative policy interventions to reduce household financial distress and mitigate the effects of biased beliefs and information frictions. The first subsection examines policies that directly target delinquency and bankruptcy by adjusting the credit terms. The second subsection examines policies aimed at correcting or mitigating belief distortions and informational barriers that, as illustrated above, exacerbate financial vulnerability.

8.1 A Cap on Revolving Interest Rates

We evaluate the effects of imposing a cap on revolving credit interest rates, focusing on the contrast between the FI and DE economies. Rather than directly changing the pricing schedule, we implement the policy by restricting the set of feasible borrowing choices. Specifically, a household can choose a debt level a' only if the equilibrium interest rate associated with that choice satisfies

$$\frac{1}{q_n^{\theta}(y, z_{-1}, a')} - 1 \leq r^{cap},$$

where r^{cap} is the interest rate cap.

Naturally, since an interest rate cap reduces risky borrowing, it substantially lowers financial distress in both economies. However, the policy is more effective and less costly in the economy with diagnostic expectations. Consequently, the trade-off between reducing financial distress and preserving welfare is much more favorable in the DE model than in the full-information model. [20]

card balances.

[20] Throughout this subsection, the policy exercise is implemented in general equilibrium, with equilibrium interest rate shown in Appendix F.

As shown in panel (a) of Figure 12, the interest rate cap reduces the fraction of households in debt in both models, with a substantially larger reduction in the DE economy. This difference emerges even for relatively loose caps. At a 30% cap, debt remains essentially unchanged in the FI economy but declines by about 2 percentage points in the DE economy. Tightening the cap to 15% amplifies the difference: the share of households in debt falls by 3.6 percentage points in the FI economy, compared with 9.4 percentage points in the DE economy. These findings suggest that borrowing is considerably more responsive to credit restrictions when households and lenders exhibit diagnostic expectations.

Panel (b) shows that interest rate caps generate substantially larger reductions in delinquency in the DE economy than in the FI economy. Under a 30% cap, delinquency falls by roughly 32% in the DE economy, compared with about 11% in the FI economy. The difference becomes more pronounced as the cap tightens: under a 15% cap, delinquency declines by about 68% under DE and 43% under FI, while under the tightest cap of 7%, the corresponding reductions are roughly 93% and 78%. These results suggest that borrowing restrictions are especially effective at reducing delinquency when excessive borrowing is driven in part by diagnostic expectations.

Panel (c) shows that bankruptcy responds in a similar, albeit somewhat stronger, way. A 30% cap reduces bankruptcies by roughly 48% in the DE economy, compared with about 26% in the FI economy. As with delinquency, the DE–FI gap widens under tighter caps: under a 15% cap, bankruptcies fall by about 78% under DE and 45% under FI, and under a 7% cap, the reductions reach roughly 95% and 86%, respectively. Thus, the bankruptcy results reinforce the main message from delinquency: rate caps have larger effects on financial distress when borrowing behavior is shaped by diagnostic expectations.

Finally, panel (d) reports the welfare costs of the interest rate cap. As expected, tighter borrowing restrictions reduce welfare in both economies by limiting households' ability to bring consumption forward over the life cycle. However, the welfare losses are substantially smaller in the DE economy. For example, at a 15% cap, the consumption-equivalent welfare loss is roughly 0.08% in the DE economy, compared with about 0.13% in the FI economy. At the tightest cap (7%), the corresponding losses are approximately 0.12% and 0.21%, respectively. The key reason is that the FI model requires a substantially lower estimated β ($\beta \approx 0.78$ for FI versus $\beta \approx 0.86$ for DE) to match household financial distress. As a result, households in the DE economy have a much weaker desire to borrow simply to advance consumption over the life cycle, making restrictions on credit substantially less costly in welfare terms.

Figure 12: The Impact of a Cap on Interest Rates

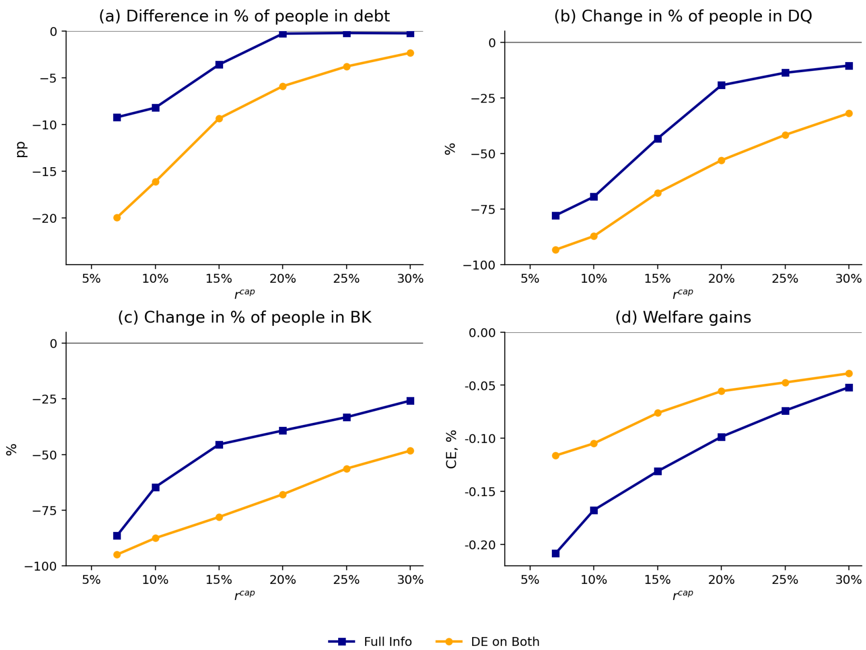

Overall, the DE economy delivers a strictly more favorable trade-off between reducing financial distress and preserving household welfare under interest rate caps.

8.2 Policies to Mitigate Belief Biases and Information Frictions

In this subsection, we quantify the welfare gains from eliminating belief overreaction and the information friction associated with income shocks.[21]

We begin by evaluating the welfare effects of removing diagnostic expectations from both households and lenders. Conceptually, this exercise asks what proportional increase in consumption a household born in the DE economy requires to be indifferent to being born in the RL economy. The first column of Table 9 shows that eliminating diagnostic expectations from both agents and lenders yields substantial welfare gains, equivalent to a 0.63 percent increase in consumption in every period during the life cycle.

Next, we decompose this welfare gain into the parts attributable to households and lenders. To do this, we compare an economy in which only households retain diagnostic expectations, while lenders have rational learning, with the benchmark economy in which both households and lenders retain diagnostic expectations. As shown in Column 2 of Table 9, removing diagnostic expectations from lenders alone yields a modest welfare gain of 0.08 percent. This finding suggests that households' behavioral biases are

[21] See Section G of the Appendix for computational details.

Table 9: Welfare Gains from Eliminating Diagnostic Expectations and Information Friction

	Drop Both Diagnostic Exp.	Drop Lenders' Diagnostic Exp.	Eliminate Info Fric.	Total Gain
Welfare Gains (CE, %)	0.63	0.08	0.90	1.54

Note: Welfare gains related to dropping diagnostic expectations are relative to the benchmark DE economy. The welfare gains of eliminating information friction are relative to the RL economy.

responsible for the majority of welfare losses.

We also examine the welfare effects of eliminating information frictions, thereby allowing households and lenders to perfectly distinguish between persistent and transitory income shocks. This exercise asks what proportional increase in consumption households born in the RL economy would require to be indifferent between being born in the RL and FI economies. Column 3 of Table 9 shows that removing information frictions generates substantial welfare gains, equivalent to a 0.90 percent permanent increase in consumption.

Finally, we evaluate the gains from moving from our benchmark DE economy to the FI economy. Eliminating both diagnostic expectations and information frictions yields welfare gains of 1.54 percent.

While it is difficult to directly change beliefs or information structures, these gains may be achievable through policies that improve expectation formation, such as financial literacy programs or labor market interventions. A growing literature shows that information provision improves labor market outcomes. Jäger et al. (2024) finds that providing workers with information about outside wages corrects beliefs and changes job search and wage bargaining behavior. Similarly, Belot et al. (2025) shows that informing unemployed workers about suitable occupations increases employment and earnings. These findings suggest that reducing information frictions can substantially improve economic outcomes.

Existing policies already move in this direction. Public Employment Services in Nordic countries provide detailed earnings information and career guidance that help workers form more accurate expectations about future income (OECD, 2024). Likewise, the UK's Behavioral Insights Team uses personalized debt feedback—such as informing borrowers how long repayment would take at current payment rates—to reduce optimism bias and improve financial decisions, offering a practical example of interventions that mitigate belief distortions in our model (Behavioural-Insights-Team, 2018).

9 Conclusion

Our findings suggest that household default is not solely driven by adverse income realizations, but also by misperceptions and biased beliefs about future income prospects. When positive income shocks are mistakenly perceived as permanent, households borrow excessively, driven by unwarranted optimism. As beliefs are later revised downward following negative income realizations, delinquency becomes a natural consequence. Households may remain in delinquency while adverse shocks persist and debt continues to accumulate. After a period of delinquency, a positive income shock can trigger bankruptcy as a form of "fresh start," allowing households to escape unsustainable debt burdens.

Two distinguishing implications favor the model with diagnostic expectations. First, its estimated discount factor implies a more realistic degree of household impatience. Second, it produces a correlation between income and interest rates that more closely matches the empirical evidence.

Finally, we show that policies designed to reduce financial distress are markedly more effective when diagnostic expectations are present. This highlights the importance of accounting for belief distortions when evaluating household behavior and designing credit market policies.

References

Ameriks, John, Andrew Caplin, Minjoon Lee, Matthew D Shapiro and Christopher Tonetti. 2023. "Cognitive Decline, Limited Awareness, Imperfect Agency, and Financial Well-being." *American Economic Review: Insights* 5(1):125–140.

Athreya, Kartik. 2002. "Welfare implications of the Bankrupcy Reform Act of 1999." *Journal of Monetary Economics* 49:1567–1595.

Athreya, Kartik, José Mustre-del Río and Juan M. Sánchez. 2019. "The Persistence of Financial Distress." *Review of Financial Studies* 32(10):3851–3883.

Athreya, Kartik, Juan M. Sánchez, Xuan S. Tam and Eric R. Young. 2015. "Labor Market Upheaval, Default Regulations, and Consumer Debt." *Review of Economic Dynamics* 18(1):32–52.

Athreya, Kartik, Juan M. Sánchez, Xuan S. Tam and Eric R. Young. 2017. "Bankruptcy and delinquency in a model of unsecured debt." *International Economic Review* 59(2):593–623.

Attanasio, Orazio P and Guglielmo Weber. 1995. "Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey." *Journal of political Economy* 103(6):1121–1157.

Behavioural-Insights-Team. 2018. "A Behavioural Approach to Managing Money: Ideas and Results from the Financial Capability Lab.".

Belot, Michèle, Bart K de Koning, Didier Fouarge, Philipp Kircher, Paul Muller and Sandra Phlippen. 2025. Advising job seekers in occupations with poor prospects: A field experiment. Technical report National Bureau of Economic Research.

Berger, David, Konstantin Milbradt, Fabrice Tourre and Joseph Vavra. 2024. "Optimal refinancing with inattention." *American Economic Review: Insights* . Forthcoming.

Blundell, Richard and Ian Preston. 1998. "Consumption Inequality and Income Uncertainty." *The Quarterly Journal of Economics* 113(2):603–640.

Bordalo, Pedro, Nicola Gennaioli, Andrei Shleifer and Stephen J. Terry. 2024. Real Credit Cycles. NBER Working Papers 28416 National Bureau of Economic Research, Inc.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma and Andrei Shleifer. 2020. "Overreaction in macroeconomic expectations." *American Economic Review* 110(9):2748–2782.

Braxton, J Carter, Kyle F Herkenhoff, Jonathan L Rothbaum and Lawrence Schmidt. 2021. Changing income risk across the US skill distribution: Evidence from a generalized Kalman filter. Technical report National Bureau of Economic Research.

Buhmann, Brigitte, Lee Rainwater, Guenther Schmaus and Timothy M. Smeeding. 1988. "Equivalence Scales, Well-Being, Inequality, and Poverty: Sensitivity Estimates Across Ten Countries Using the Luxembourg Income Study (LIS) Database." *Review of Income and Wealth* 34(2):115–142.

Chatterjee, Satyajit, Dean Corbae and Jose-Victor Rios-Rull. 2008. "A finite-life private-information theory of unsecured consumer debt." *Journal of Economic Theory* 142(1):149–177.

Chatterjee, Satyajit, Dean Corbae, Kyle Dempsey and José-Víctor Ríos-Rull. 2023. "A quantitative theory of the credit score." *Econometrica* 91(5):1803–1840.

Chatterjee, Satyajit, Dean Corbae, Makoto Nakajima and José-Víctor Ríos-Rull. 2007. "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default." *Econometrica* 75(6):1525–1591.

D'Acunto, Francesco, Michael Weber and Xiao Yin. 2024. Subjective Income Expectations and Household Debt Cycles. NBER Working Paper 32715 National Bureau of Economic Research, Inc.

Drozd, Lukasz A. and Michal K. Kowalik. 2023. "The Conundrum of Zero APR: An Analytical Framework." *FRB of Philadelphia Working Paper Series* (23-6).

Drozd, Lukasz A. and Ricardo Serrano-Padial. 2017. "Modeling the Revolving Revolution: The Debt Collection Channel." *American Economic Review* 107(3):897â930.

Eaton, Jonathan and Mark Gersovitz. 1981. "Debt with potential repudiation: theoretical and empirical analysis." *Review of Economic Studies* 48:289–309.

Exler, Florian, Igor Livshits, James MacGee and Michèle Tertilt. 2024. "Consumer credit with over-optimistic borrowers." *Journal of the European Economic Association* p. jvae057.

Ganong, Peter, Pascal J Noel, Christina Patterson, Joseph S Vavra and Alexander Weinberg. 2025. Earnings instability. Technical report National Bureau of Economic Research.

Guvenen, Fatih. 2007. "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" *American Economic Review* 97(3):687–712.

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

He, Qiwei and Philipp Kircher. 2025. "Understanding Expectations in Job Search: Subjective Duration Dependence, Aggregate Labor Market Shocks and Overreaction.".

Huang, Chengdai, J. Carter Braxton, Kyle Herkenhoff, Michael Nattinger, Jonathan Rothbaum and D. W. Schmidt. 2025. "Rising Risk Among the Rich: Implications for Wealth Inequality and Interest Rates." *Working Paper* .

Huggett, Mark. 1993. "The risk-free rate in heterogeneous-agent incomplete-insurance economies." *Journal of Economic Dynamics and Control* 17(5-6):953–969.

Jäger, Simon, Christopher Roth, Nina Roussille and Benjamin Schoefer. 2024. "Worker beliefs about outside options." *The Quarterly Journal of Economics* 139(3):1505–1556.

Kahneman, Daniel and Amos Tversky. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185(4157):1124–1131.

Kaplan, Greg and Giovanni L. Violante. 2010. "How Much Consumption Insurance beyond Self-Insurance?" *American Economic Journal: Macroeconomics* 2(4):53–87.

Lee, Eungik and Johan Sæverud. 2023. "Earnings Shocks, Expectations, and Spending.". Manuscript, Indiana University.

Livshits, Igor, James MacGee and Michèle Tertilt. 2007. "Consumer Bankruptcy: A Fresh Start." *American Economic Review* 97(1):402–18.

Ma, Yueran, Tiziano Ropele, David Sraer and David Thesmar. 2020. A quantitative analysis of distortions in managerial forecasts. Technical report National Bureau of Economic Research.

Nakajima, Makoto. 2012. "Rising Indebtedness and Hyperbolic Discounting: A Welfare Analysis." *Quantitative Economics* 3(2):257–288.

OECD. 2024. "The Role of Public Employment Services in Promoting an Inclusive Nordic Labour Market.".

Saldain, Joaquín. 2025. High-Cost Consumer Credit: Desperation, Temptation and Default. Staff Working Papers 25-6 Bank of Canada.

Appendix for

"What Drives Household Financial Distress? The Role of Earnings Misperceptions"

by Juan M. Sánchez and Stefan Song

A Proof of Propositions

A.1 Lemma 2

In the RL economy, the household's value function depends on current income y and next period's debt choice d', where d' serves as the control variable. Optimization implies that the optimal debt choice d' is a function of y. Since $x(\cdot)$ is itself a strictly increasing function of y, the inverse function theorem ensures that the optimal debt choice d' can equivalently be expressed as a function of x.

From the value function, notice first that households will never choose $d' > (1-\nu)z_H$, since any additional debt beyond this level has no present value. For the range $d' \leq (1-\nu)z_L$, because $\frac{1}{1+r} > \beta$, households optimally choose the maximum feasible debt level within this range, $d' = (1-\nu)z_L$. The associated optimized value is

$$V_1^* = y + \beta[xz_H + (1-x)z_L] + \left(\frac{1}{1+r} - \beta\right)(1-\nu)z_L.$$

For the intermediate range $(1-\nu)z_L < d' \leq (1-\nu)z_H$, again using $\frac{1}{1+r} > \beta$, households optimally choose the upper bound $d' = (1-\nu)z_H$, yielding

$$V_2^* = y + \beta[xz_H + (1-x)\nu z_L] + x\left(\frac{1}{1+r} - \beta\right)(1-\nu)z_H.$$

The optimal debt choice depends on which value is higher, V_1^* or V_2^*. Define x^* as the posterior threshold at which $V_1^* = V_2^*$, which solves

$$x^* = \frac{z_L}{(1+r)\left[\left(\frac{1}{1+r} - \beta\right)z_H + \beta z_L\right]}.$$

If $x > x^*$, then $V_2^* > V_1^*$ and households optimally choose $d' = (1-\nu)z_H$. Otherwise, they choose $d' = (1-\nu)z_L$.

A.2 Lemma 3

Since the distribution of z is Bernoulli, under diagnostic expectations, agents overreact to the change of posterior relative to the prior. That is, diagnostic agents exaggerate the probability of $z = z_H$ if $x(y) > \pi$, vice versa.

We are then interested in the threshold income $\tilde{y}$ such that $x(\tilde{y}) > \pi$, where superscript θ denotes diagnostic posterior. Using the expression for $x(y)$, we have

$$x(\tilde{y}) \equiv \frac{\phi(\tilde{y} - z_H)\pi}{\phi(\tilde{y} - z_L)(1-\pi) + \phi(\tilde{y} - z_H)\pi} = \pi \Rightarrow \tilde{y} = \frac{z_L + z_H}{2}$$

That is, if $y > \tilde{y}$, $Prob^{\theta}(z = z_H|y) > Prob(z = z_H|y)$, vice versa. Superscript θ here refers to diagnostic belief.

Given Assumption 1, we have

$$x(y^*) = x^* = \frac{z_L}{(1+r)\left[\left(\frac{1}{1+r} - \beta\right) z_H + \beta z_L\right]} > \pi = x(\tilde{y})$$

where x^* is the posterior threshold in the proof of Lemma 2. This immediately implies that $y^* > \tilde{y}$, where y^* is the income signal thresholds in the RL economy, since $x(\cdot)$ is a monotonically increasing function and since $y^* > \tilde{y}$, agents under diagnostic expectations will exaggerate the probability of $z = z_H$ if $y = y^*$ and therefore we have

$$Prob^{\theta}(z = z_H|y^*) > x^*.$$

Define the income signal threshold in the diagnostic expectations economy that triggers a higher level of borrowing as $y^{*,\theta}$ such that

$$Prob^{\theta}(z = z_H|y^{*,\theta}) = x^* < Prob^{\theta}(z = z_H|y^*),$$

which implies $y^{*,\theta} < y^*$ since $Prob^{\theta}(z = z_H|y)$ is also increasing in y.

As a result, we immediately have

$$\int_{y^{*,\theta}}^{\infty} (1 - x(y)) f(y) \, dy \;>\; \int_{y^*}^{\infty} (1 - x(y)) f(y) \, dy \;>\; 0.$$

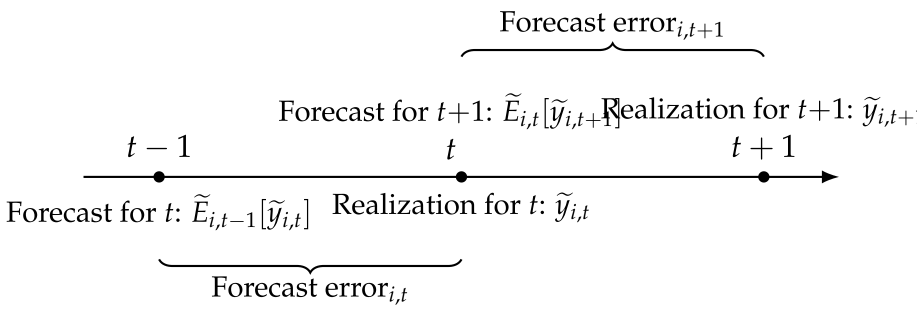

Figure 13: Timing of SCE income forecasts and realizations

B Use *SCE* to Estimate θ

B.1 Questionnaire Statement in Labor Market Module

L3 - L3

How much do you make **before** taxes and other deductions at your [main/current] job, on an annual basis? Please include any bonuses, overtime pay, tips or commissions.

________________________ dollars per year

OO2e2 (Added March 2015)

What do you believe your annual earnings will be in 4 months?

________________________ dollars per year

B.2 Aligning the Forecast Horizon

The SCE Labor Market Module employs a rotating panel in which each household is observed for one year, with three interview waves conducted at four-month intervals in March, July, and November. This panel structure allows us to construct income forecast errors using two consecutive waves and to estimate the autocorrelation of forecast errors at a four-month horizon. The timing is shown in Figure 13.

An important measurement issue is that the SCE earnings questions are stated in annual terms. Respondents report their current annual earnings and, in the previous wave, their expected annual earnings at the next interview. Since interviews are four months apart, the forecast error compares a forecast made four months earlier with the reported annual earnings at wave t. In this sense, the object being forecasted is annual earnings, with a four-month-ahead forecast horizon. We therefore keep the SCE variables in their reported annual units. Let $\widetilde{y}_{i,t}$ denote household i's reported annual earnings at wave t, and let $\widetilde{E}_{i,t-1}[\widetilde{y}_{i,t}]$ denote the expectation reported at wave $t-1$ for annual earnings at wave t. The SCE forecast error is constructed as

$$FE^{SCE}_{i,t} = \widetilde{y}_{i,t} - \widetilde{E}_{i,t-1}\left[\widetilde{y}_{i,t}\right].$$

We then estimate the empirical autocorrelation from

$$\mathrm{FE}^{SCE}_{i,t+1} = \beta^{SCE}_0 + \beta^{SCE}_1 \mathrm{FE}^{SCE}_{i,t} + \beta^{SCE\prime}_2 \mathbf{X}_{i,t} + u_t + \xi_{i,t+1},$$

where $\mathbf{X}_{i,t}$, as stated in the main text, denotes a vector of observable household characteristics.

To map this moment to the model, we construct a four-month-frequency interpolation of the annual income process used in the quantitative model. The benchmark annual process is characterized by a transitory shock variance σ^2_e and a persistent innovation variance σ^2_ϵ. We assume that each annual shock can be represented as the sum of three independent four-month shocks. The auxiliary four-month process is therefore

$$\begin{aligned} y^{4m}_{j,t} &= z^{4m}_{j,t} + e^{4m}_{j,t}, \quad & e^{4m}_{j,t} &\sim \mathcal{N}(0, \sigma^2_e/3), \\ z^{4m}_{j,t} &= z^{4m}_{j,t-1} + \epsilon^{4m}_{j,t}, \quad & \epsilon^{4m}_{j,t} &\sim \mathcal{N}(0, \sigma^2_\epsilon/3), \end{aligned}$$

where t indexes four-month periods. This scaling preserves the annual shock variances in the benchmark model. In particular, the sum of three independent four-month transitory shocks has variance σ^2_e, and the sum of three independent four-month persistent innovations has variance σ^2_ϵ.

Using this auxiliary process, we construct the model-implied forecast error as

$$FE^{M}_{j,t+1} = y^{4m}_{j,t+1} - \widetilde{E}_{j,t}\left[y^{4m}_{j,t+1}\right].$$

where $\widetilde{E}_t[\cdot]$ denotes the expectation operator implied by the belief model. We then estimate the model counterpart of the SCE autocorrelation from [22]

$$\mathrm{FE}^{M}_{j,t+1} = \beta^{M}_0 + \beta^{M}_1 \mathrm{FE}^{M}_{j,t} + \beta^{M\prime}_2 \mathbf{X}^{M}_{j,t} + \xi^{M}_{j,t+1}.$$

The comparison between β^M_1 and β^{SCE}_1 relies on the interpretation that the SCE earnings response is an annualized earnings object observed at the survey date. Under this

[22]The model simulation does not include aggregate calendar-year variation, so the model regression omits year fixed effects. The empirical year fixed effects remove aggregate time variation from the SCE forecast errors, making the data moment closer to the idiosyncratic forecast-error persistence generated by the auxiliary model.

interpretation, the SCE and model forecast errors refer to the same four-month-ahead forecast horizon. The remaining difference is one of measurement units: the SCE forecast errors are constructed from reported annual earnings; in the auxiliary model, forecast errors are constructed from the four-month components of the interpolated income process. However, this unit difference does not affect the autocorrelation coefficient, which is invariant to common rescaling of the forecast error series. Thus, under the annualized-earnings interpretation, the model-implied coefficient β_1^M is directly comparable to the empirical coefficient β_1^{SCE}.

B.3 Robustness

For sample selection, we restrict the analysis to households aged 29 and above to mitigate the influence of potentially unformed belief priors. We exclude observations in which the absolute value of the forecast error exceeds five times the sample mean income, interpreting such values as likely survey noise.[23] Finally, we focus on households experiencing a positive income change, which we interpret as evidence of an income shock, consistent with the environment of our model.

To assess robustness, we estimate several alternative empirical specifications, reported in Table 10. In addition, Figure 14 presents a binscatter of forecast error autocorrelation based on our baseline regression.

Figure 14: Forecast Error Auto-Correlation Bin Scattered: SCE Labor Market Survey

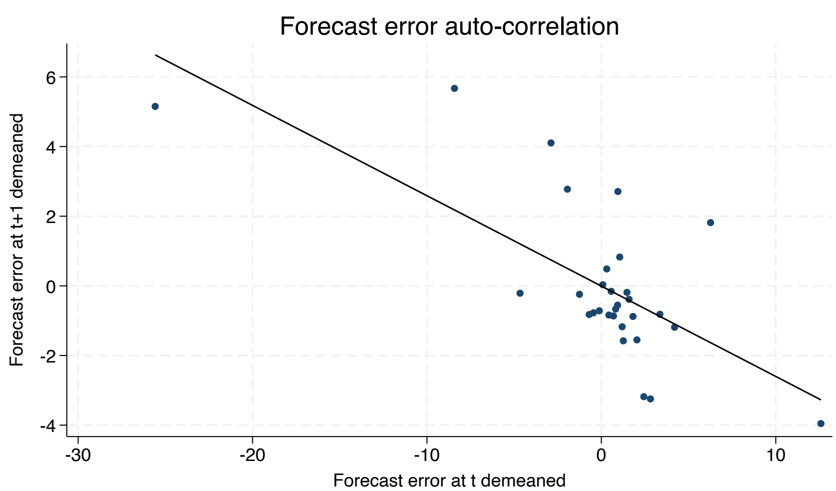

[23]This restriction removes only a small number of observations. The excluded cases appear visually implausible and are likely because of misreporting.

Table 10: Robustness: Forecast Error Regression Results

Dependent Variable: Forecast Error at $t+1$

	(1)	(2)	(3)	(4)	(5)	(6)
Forecast Error at t	-0.353^{***}	-0.355^{***}	-0.340^{***}	-0.332^{***}	-0.200	-0.294^{***}
	(0.117)	(0.117)	(0.124)	(0.054)	(0.125)	(0.051)
Age		-0.588				
		(0.326)				
Age^2		0.006				
		(0.003)				
$\lvert\Delta\tilde{y}\rvert > 0$	Yes	Yes	Yes	Yes	No	Yes
Age ≥ 29	Yes	Yes	Yes	Yes	Yes	No
Other Controls	Yes	Yes	No	Yes	Yes	Yes
Weight	Yes	Yes	Yes	No	Yes	Yes
Observations	622	622	622	622	989	681
R^2	0.230	0.235	0.060	0.228	0.136	0.204

Note: Standard errors in parentheses; *** $p < 0.01$. Other controls include gender, race, education level, tenure time, marital status, and state fixed effects. Column (1) is our benchmark regression. Column (2) adds age square controls. Column (3) drops other controls. Column (4) tests the importance of weight. Column (5) drops the main sample restriction of income change larger than 0. Column (6) removes the sample restriction of age 29 or older.

We also replace the previous sample selection criterion, which restricts the absolute value of forecast error to be within five times the mean income, with winsorization at the 1 percent level. Accordingly, we get an estimate of auto-correlation of -0.237 with standard error of 0.073 and $p-$value of 0.001.

B.4 Regression Results by Age and Education Groups

Table 11: Forecast Error Autocorrelation: Heterogeneity by Education and Age

	(1) No Controls	(2) With Controls
Panel A. Full Sample		
Lagged forecast error (FE_{t-1})	$-\mathbf{0.340}^{***}$	$-\mathbf{0.353}^{***}$
	(0.125)	(0.124)
Observations	622	622
R^2	0.059	0.230
Panel C. Education: College$\leq$ vs Postgraduate		
College $\leq$ (Q36 $\leq$ 5)	$-\mathbf{0.415}^{**}$	$-\mathbf{0.445}^{***}$
	(0.170)	(0.168)
Observations	479	479
R^2	0.077	0.216
Postgraduate (Q36 $>$ 5)	$-\mathbf{0.213}^{*}$	$-\mathbf{0.217}$
	(0.115)	(0.156)
Observations	143	143
R^2	0.032	0.298
Panel D. Age: $<$ 45 **vs** $\geq$ 45		
Age $<$ 45	$-\mathbf{0.476}^{**}$	$-\mathbf{0.517}^{**}$
	(0.197)	(0.209)
Observations	282	282
R^2	0.087	0.272
Age $\geq$ 45	$-\mathbf{0.264}^{*}$	$-\mathbf{0.281}^{**}$
	(0.144)	(0.140)
Observations	340	340
R^2	0.044	0.190

Notes: Each row reports a separate regression of the current forecast error on the lagged forecast error. Columns (1) and (2) report specifications without and with controls, respectively. Controls include demographic characteristics and other covariates described in the text. Education groups are defined using survey question Q36. Robust standard errors are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

C Calibrated Life-Cycle Profiles

Figure 15: Demographic and Economic Inputs: Household Size, Mortality, and Income

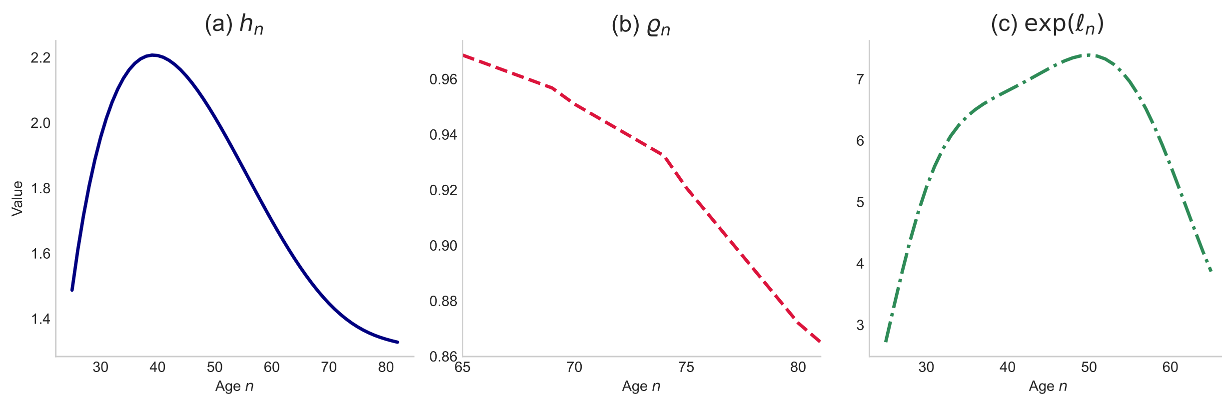

D Robustness of the Relationship Between Income and Interest Rate

As the measure of income, we consider the answer to the question: *"How much was the total income you (and your family living here) received from all sources before taxes and other deductions were made?"* In addition, we consider an alternative question about income to capture labor earnings: *"In total, what was your (family's) annual income from wages and salaries in year X, before deductions for taxes and anything else?"* As a measure of interest rate on credit card debt, the analysis uses the response to the SCF question: *"What interest rate do you pay on the card where you have the largest balance?"*

Table 12 reports the regression results by year. Overall, the estimated income betas for both labor earnings and total income align more closely with the predictions of the learning-based models.

E Fit of Alternative Models

This appendix documents the fit of alternative models.

Overall, the incidence of delinquency is closely matched—both in aggregate and across age groups —across all model variants. The models also capture the life-cycle pattern of bankruptcy reasonably well, though some over- or undershooting occurs depending on

Table 12: Regressions of Interest Rates on Income, Data

Year	Coeff. on Income	Standard Errors	Beta on income	Number of observations
		Income: Total Income		
1998	-0.007	0.001	-0.628	5,764
2001	-0.005	0.001	-0.349	6,182
2004	-0.009	0.001	-0.514	5,763
2013	-0.007	0.001	-0.492	6,484
2016	-0.002	0.000	-0.132	7,784
2019	-0.006	0.001	-0.462	7,247
Average	**-0.006**	–	**-0.430**	–
		Income: Labor Earnings		
1998	-0.009	0.001	-0.243	5,339
2001	-0.009	0.001	-0.310	5,714
2004	-0.016	0.001	-0.525	5,256
2013	-0.013	0.001	-0.417	5,693
2016	-0.004	0.001	-0.138	6,946
2019	-0.008	0.001	-0.188	6,380
Average	**-0.010**	–	**-0.303**	–

Note: Regression of interest rates on income, considering two measures of income: Labor earnings and total income. The regressions are controlled for age fixed effects, credit card debt, and its square. The sample includes households with individuals aged 25 to 65 and positive credit card balances, winsorized at the 1st and 99th percentiles of the income distribution. Observations are weighted using Survey of Consumer Finances (SCF) sampling weights. The reported beta on income is the coefficient from a regression of interest rates on income, normalized by the standard deviations of income and interest rates. This coefficient represents the change in interest rates, in standard deviation units, associated with a one-standard-deviation increase in income. Note that the years of the Global Financial Crisis and the COVID-19 crisis were omitted.

the specification. In contrast, while the FI and RL models effectively reproduce the share-in-debt profile, the diagnostic expectation model struggles to match the empirical flatness of that profile over the life cycle. This limitation stems from its higher estimated discount factor, which dampens borrowing, especially among older households. Additional life-cycle forces—such as health, housing, or intergenerational transfer expenses—may also contribute to the persistent use of credit among older households and help explain the empirical pattern.

Table 13: Fit of Targeted Moments in Alternative Economies

Moments	Data	Full Information	Rational Learning	Diagnostic Expectations Hhds & Lenders	Diagnostic Expectations Only Hhds
% in DQ, age 25–34	11.87	11.83	11.74	11.85	10.77
% in DQ, age 35–44	11.50	11.45	11.05	12.07	11.58
% in DQ, age 45–55	9.49	9.29	9.54	9.76	9.51
% in BK, age 25–34	0.51	0.46	0.65	0.44	0.62
% in BK, age 35–44	0.62	0.63	0.74	0.56	0.81
% in BK, age 45–54	0.51	0.52	0.52	0.40	0.50
% in Debt, age 25–34	41.49	42.09	41.64	40.79	45.79
% in Debt, age 35–44	44.55	49.28	44.71	39.46	41.03
% in Debt, age 45–54	44.70	47.13	41.94	27.89	32.06

F General Equilibrium Interest Rates in Section 8.1

Figure 16 plots the general equilibrium interest rate with policies on revolving interest rate cap.

Figure 16: Equilibrium interest rates with revolving interest cap policies

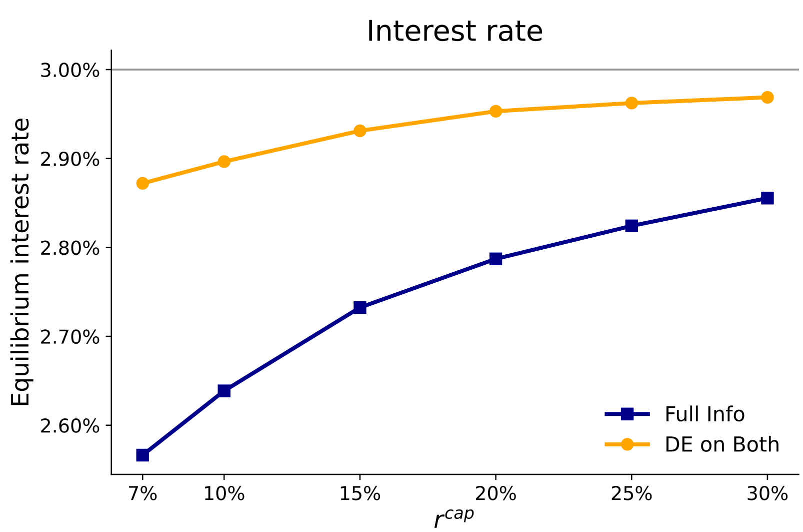

G Welfare Analysis

Welfare in each economy is defined as the expected discounted sum of utility from the consumption stream. Let $\{c_n^B\}_{n=n_0}^N$ denote the consumption path in the benchmark economy (B). Then, ex-ante welfare in economy B is given by

$$EV^B \equiv \frac{1}{M} \sum_i \sum_{n=n_0}^{N} \beta^n \varrho_n \frac{(c_n^B / h_n)^{1-\sigma}}{1-\sigma},$$

where β is the subjective discount factor, ϱ_n denotes the cohort weight or survival probability at age n, and σ is the coefficient of relative risk aversion. M is the total number of households in the economy.

Analogously, ex-ante welfare in the alternative economy (A) is defined as

$$EV^A \equiv \frac{1}{M} \sum_i \sum_{n=n_0}^{N} \beta^n \varrho_n \frac{(c_n^A / h_n)^{1-\sigma}}{1-\sigma}.$$

We measure welfare differences using the consumption-equivalent variation (CEV). Specifically, the welfare gain of economy A relative to economy B is defined as the constant percentage reduction in consumption in economy A that makes households indifferent between the two economies. Formally, Δ satisfies

$$EV^A (1+\Delta)^{1-\sigma} = EV^B,$$

where EV^i denotes expected lifetime utility in economy i. Solving for Δ yields

$$\Delta = \left(\frac{EV^B}{EV^A} \right)^{\frac{1}{1-\sigma}} - 1.$$

A positive value of Δ therefore indicates that households prefer economy A to economy B, since they would be willing to give up a constant fraction Δ of consumption in economy A and still be indifferent between the two economies. Under the benchmark preference parameter $\sigma = 2$, this expression simplifies to

$$\Delta = \frac{EV^A}{EV^B} - 1.$$