

Credit Scores and Inequality Across the Life Cycle

Satyajit Chatterjee

Federal Reserve Bank of Philadelphia

Dean Corbae

University of Wisconsin–Madison

Kyle Dempsey

Ohio State University

José-Víctor Ríos-Rull

University of Pennsylvania

WP 26-07

PUBLISHED

February 2026



ISSN: 1962-5361

Disclaimer: This Philadelphia Fed working paper represents preliminary research that is being circulated for discussion purposes. The views expressed in these papers are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors. Philadelphia Fed working papers are free to download at: <https://www.philadelphiafed.org/search-results/all-work?searchtype=working-papers>.

DOI: <https://doi.org/10.21799/frbp.wp.2026.07>

Credit Scores and Inequality across the Life Cycle ^{*}

Satyajit Chatterjee [†] Dean Corbae [‡] Kyle Dempsey [§] José-Víctor Ríos-Rull [¶]

August 15, 2025

Abstract

Credit scores are a primary screening device for the allocation of credit, housing, and sometimes even employment. In the data, credit scores grow and fan out with age; at the same time, income and consumption inequality also increase with a cohort's age. We postulate a simple model with hidden information to explore the joint determination of credit scores, income, and consumption over an individual's lifetime which can replicate these empirical facts. We use the model to understand the role of technologies like big data or legal restrictions limiting information on certain adverse events like medical expenses intended to increase credit market access.

Keywords: Adverse Selection, Moral Hazard, Bayesian Learning, Reputation.

JEL Classification Numbers: D82, E21, G51.

^{*}We wish to thank the editors John Leahy and Valerie Ramey for valuable comments on an earlier version of this paper. We also wish to thank our discussants Stefania Albanesi and Kyle Herkenhoff for very helpful comments which have made this a better paper. And finally, thanks to John Ryan for excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect views of the Federal Reserve Bank of Philadelphia or of the Federal Reserve System.

[†]Federal Reserve Bank of Philadelphia

[‡]University of Wisconsin-Madison and NBER

[§]The Ohio State University

[¶]University of Pennsylvania, CAERP, CEPR, NBER, and UCL

1 Introduction

Credit scores are a fundamental ingredient of a borrower’s access to credit and housing. They are also widely used elsewhere: in pre-employment credit screens, the determination of insurance rates, and even in the choice of partners. Broadly, they are often seen as proof of “character,” even though adverse events outside a person’s direct control (like hospitalizations that result in financial distress) may enter an individual’s credit report and weaken one’s score. Despite their widespread use, credit scores as a signal of reputation are conspicuously absent from standard quantitative models used to evaluate consumer credit policy recommendations.

In this paper, we build on our recent work ([Chatterjee et al., 2023](#)) to explore the link between reputational inequality, as evidenced in the distribution of credit scores, and inequality in income and consumption, both in the cross-section and over the life cycle. To do this, we develop a model with defaultable consumption loans, savings, and endogenous effort in which an individual privately knows her preferences (hidden information) and effort choice (hidden action). This framework has the minimum ingredients necessary to generate a notion of reputation akin to credit scores whose distribution evolves alongside the distributions of income and consumption and in which we can analyze how reputational concerns impact a person’s effort choice and her choice to smooth consumption across time and states.

To motivate our analysis, we document two key patterns regarding reputational inequality over the life cycle. Dispersion in credit scores within a given age bracket rises across brackets, mirroring well-known patterns in income and consumption. The increase is particularly strong between the ages of 20 and 40, which suggests that learning or reputation formation may play an important role. Second, while there is a strong correlation between income and credit score in general, this correlation is weakest early in life (between the ages of 20 and 30). While credit scores reflect the likelihood of repayment and are therefore correlated with income, the low correlation early in life suggests there are other payoff-relevant factors that affect one’s repayment behavior, about which less is known early in life.

Since credit scores are complex, equilibrium objects, it is generally quite difficult to understand what drives the empirical relationships we document above with reduced-form data analysis. Therefore, to shed light on the mechanisms underpinning these relationships, we develop a series of simple models in which reputational incentives directly affect income and consumption inequality. We begin with a setting in which the only signals of one’s type—and therefore the only things that affect one’s reputation—are savings and default decisions. In this simplified version of [Chatterjee et al. \(2023\)](#), the presence of adverse selection, combined with dynamically updated reputations, changes individuals’ consumption and savings incentives: the greater the value of reputation, the greater is consumption and wealth inequality. This version of the model, however, is silent on the relationship between reputational incentives and income inequality. We address the link between reputational incentives and income inequality by adding a hidden costly effort choice; in a modified setting with both moral hazard and adverse selection, income realizations may become powerful signals of one’s type in addition to the saving and default decisions from the baseline version of the model. We

show that as the cost of mimicking high-type effort choices rises, low types take less effort, leading to higher income and consumption inequality. Further, as the value of reputation increases (for example, keeping track of reputation through time), income inequality increases alongside consumption inequality, illustrating how technological innovations in scoring or legal restrictions on what information can be accessed in credit reports can have important implications for inequality.

As in our previous work, asset market transactions, such as taking out a loan or making a deposit, carry a signal about a person’s hidden information. As is well known, formal analysis of signaling often entails specifying “off-equilibrium-path” (oep) beliefs (see, for instance, [Kreps and Sobel \(1994\)](#)). However, as we show with an example later in the paper, exogenously specified oep beliefs can cause the nonexistence of equilibria, a serious problem for quantitative work. There is also the possibility that different oep beliefs can lead to different equilibria ([Cho and Kreps, 1987](#)). To avoid these issues, we add Type I extreme value preference shocks to payoffs from actions as in discrete choice models ([McFadden \(1974\)](#), ([Rust, 1987](#))). These shocks ensure that every budget-feasible action is chosen with positive probability similar to the structural quantal response literature as discussed in [Goeree et al. \(2016\)](#) (and reminiscent of [Selten \(1975\)](#) and [Myerson \(1978\)](#)). A necessary consequence is that beliefs following any feasible action are well-defined and determined in equilibrium; thus, there is no need to supply oep beliefs. The assumption that the shocks are drawn from a Type I extreme value distribution implies the familiar logit choice probability function and delivers computational traceability. It is worth noting that while we have chosen to adopt a structural approach, the logit choice probability function that figures prominently in our paper can be given a bounded rationality interpretation wherein the deviations from optimal choice are a result of errors as in nonstructural quantum response models, or follow from models of control costs ([Mattson and Weibull, 2002](#)) or rational inattention ([Matejka and McKay, 2015](#)).

We take seriously the idea that credit scores embed some notion of “character”: For us, this means scores measure some aspect of a person’s preferences. But what could “character” mean in terms of preferences? This is a wide-open question the burgeoning field of personality economics ([Heckman et al., 2023](#)) may eventually have much to say. For instance, it could be that one of the so-called “Big Five” personality traits, conscientiousness, is determinative of prudent behavior in asset markets. Conscientiousness within OCEAN¹ is defined to be the desire to be careful, diligent and to regulate immediate gratification with self-discipline. [Ameriks et al. \(2007\)](#) use survey methods to measure self-control problems and find it to be correlated with measures of conscientiousness. It is thought to generate behavior patterns like ambition, discipline, consistency, and reliability. At this point, however, we don’t know how personality traits map to utility functions in economic models. In this paper and our previous work, we take clues from what credit scorers say credit scores depend on and the (very) limited empirical evidence on the mapping between credit scores and preferences.

According to credit scoring companies such as Fair Isaac Corporation (FICO) and VantageScore, the most critical factor that affects credit scores negatively is missed or late payments, followed by high revolving balances on existing credit card accounts and requests, in quick succession, for new

¹The other four traits are (o)penness to experience, (e)xtraversion, (a)greeableness, and (n)euroticism.

credit accounts.² It appears, therefore, that taking on too much consumer debt and being late, or delinquent, on payments will lower a person’s credit score.

Since borrowing and default privilege current consumption over future consumption, *discount-factor heterogeneity* is a natural candidate, besides income, for explaining differences in credit scores. While this is a reasonable hypothesis, which we adopt in this and our previous work, we know of only one study that directly supports this view. Meier and Sprenger (2012) measured discount factors in an experimental setting and found that subjects with higher estimated discount factors also tended to have higher credit scores. Taking a broader view, the hypothesis implies that credit scores should have predictive power in explaining behavior in other choice domains where time discounting could be important.³ We note that the macroeconomics literature has invoked discount-factor heterogeneity as a driver of wealth inequality and differences in the marginal propensity to consume (Krusell and Smith (1998), Carroll et al. (2017), respectively, among others) and recent work has linked wealth inequality to genetic endowments (Barth et al., 2020). So, provisionally, we take the view that the evolution of credit scores reflects, in part, the evolution of lenders’ and scorers’ beliefs about a person’s discount factor.

This perspective brings to the fore the critical importance of record-keeping technologies and the law for the evolution of reputational inequality. In terms of record-keeping technology, computerization has been and continues to be a key enabler for credit scoring and reputational inequality (see Poon (2007) for a history of the role of computers in the development of FICO scores and Pasquale (2015) for a wide-ranging discussion of the role of algorithms in the advent of “digital reputations”). On the legal side, U.S. law prohibits the use of race, color, religion, national origin, and sex for credit scoring purposes.⁴ Thus the law has an important impact on how record-keeping technologies can be used to construct credit assessments.

As record-keeping technologies advance, we might expect more restrictions on their use due to privacy concerns. Indeed, privacy concerns triggered by widespread credit scoring led to the enactments of the Fair Credit Reporting Act (1970) and its predecessor, Title VI of the Consumer Credit Protection Act (1968), and to the restrictions noted above on what information scorers can use (Lauer (2017, Ch. 8)). If the use of such technologies is restricted, the increase in reputational inequality over the life cycle that might otherwise happen will be constrained, but the scope for welfare-improving trades might either narrow or widen. To expand on this point, it is well understood that efficiency is compromised if significant pay-off relevant information remains hidden.

²<https://www.myfico.com/credit-education/whats-in-your-credit-score>, <https://www.experian.com/blogs/ask-experian/what-is-a-vantagescore-credit-score/>

³As a factual matter, there is strong evidence that credit scores help predict accident insurance claims (Golden et al., 2016) and the onset of diseases (Israel et al., 2014), even after controlling for the conventional predictors of accident claims and morbidity. It has also been observed (Dokko et al., 2015) that individuals in committed relationships show substantial positive assortative matching with respect to credit scores, even when controlling for other socioeconomic and demographic characteristics. It is, of course, not known if discount factor differences are the causal factor in these correlations.

⁴While not explicitly prohibited by law, scorers do not include age, salary, occupation, employment history, place of residence, and marital status when constructing their scores. However, lenders are permitted to use some of this information along with a credit score when making credit decisions.

On the other hand if explicit insurance against adverse events is absent, there can be benefits to those affected if they are pooled with the unaffected (Hirschleifer, 1971).⁵ In addition to exploring the links between reputation and inequality, a second focus of our work is to understand how record-keeping technologies and the law could affect insurance and efficiency via the reputation channel.

Along these lines, we use our model to evaluate a significant recent development. Based on privacy and other concerns, the Consumer Finance Protection Bureau attempted to ban unpaid medical bills from being included in credit reports and prohibited lenders from requesting information on medical debts from prospective borrowers. Our framework provides a straightforward way to model legal constraints on information use as restrictions on the information set upon which inference regarding preferences is based. In this case, we show that the ban would have raised costs for the medical sector by increasing the incidence of medical delinquencies. Moreover, this regulation induces more pooling and less price dispersion in the unsecured credit market, since now there is the potential that a given borrower may have a “hidden debt” in the form of a medical liability, which increases their effective leverage on a loan of a given size.

In summary, this paper creates a laboratory to explore the links between consumption, income, and reputational inequality.⁶ We operationalize “reputation as character” via hidden differences in discount factors about which lenders learn over time from consumers’ choices. We explore the implications of restrictions on information sets of lenders for the evolution of reputational and consumption and income inequality. While we have focused on credit markets, the framework can be easily adapted to study issues in other areas where hidden information (and consequent adverse selection) plays an important role as in health and insurance markets.

1.1 Related literature

There is a rich literature on quantitative dynamic models of unsecured consumer credit and bankruptcy (see Exler and Tertilt (2020) for a survey). Early examples include Athreya (2002), Chatterjee et al. (2007) and Livshits et al. (2007). In these papers, while a credit score (viewed as the probability of not defaulting over some horizon) can be constructed, the score is not a signal of reputation. In this paper and our earlier work Chatterjee et al. (2023), we extend this class of models to have hidden information that asset market actions can only *partially* reveal. The inference based on these actions is summarized dynamically as probability assessments that resemble credit scores.

The partial revelation of type is an important aspect of our work because it makes it consistent with the behavior of lenders. If real-world credit contracts were successful in completely revealing a person’s type (as in the classic work of Rothschild and Stiglitz (1976) or more recently in Guerreri

⁵For instance, in Chatterjee et al. (2023) we showed that the average welfare of newborns (i.e., welfare behind the “veil of ignorance”) is higher in a “no-tracking” economy, despite loans being more expensive due to the elimination of incentives to repay that rely on maintaining a reputation. The reason for this was that individuals unlucky to be born as the impatient type could pool with individuals lucky to be born as the patient type and the welfare gain from this pooling overcame the efficiency loss from diminution in the incentive to repay.

⁶In fact, we operationalize the laboratory through a “User’s Manual” in the Appendix that includes code to run the laboratory.

et al. (2010)), it would be hard to understand why the credit industry devotes resources to creating and maintaining credit scores since the revelation of type would make that effort unnecessary. Furthermore, partial revelation of types is also consistent with studies that have documented the importance and relevance of adverse selection in credit markets (Ausubel (1999), Agarwal et al. (2010), Einav et al. (2012), and Hertzberg et al. (2018), Xin (2023), among others).

Of relevance to our current and predecessor papers are studies that have examined the role of improvements in information technology on credit access. In these papers, the “technology” is a noisy signal of a borrower’s actual characteristics, and the “improvement” is an increase in signal precision. These include Narajabad (2012), Livshits et al. (2016), and Sanchez (2018). For instance, Narajabad (2012) compares polar cases: one in which the credit market lacks information, resulting in a pooling equilibrium, and one in which sufficient information separates borrowers. Livshits et al. (2016) consider a simple asymmetric information model with costly contracting where borrowers know their types but uninformed lenders receive a noisy signal of a borrower’s type. As signal precision improves, the extent of pooling in a given contract falls.⁷

Technological change can improve signal precision. Albanesi and Vamossy (2019) develop a model to predict consumer default based on deep learning. They show that their deep learning approach outperforms conventional logistic regression approaches in predicting default because it can unearth non-linear patterns of interaction among factors affecting default in high-dimensional credit history data. One way to think of their approach is that it provides assessments that approach that of a full information model. We conduct one such exercise.

While previous quantitative theory models imposed exogenous punishment, we incorporate dynamic reputation as a means of disciplining borrowers along the lines of Diamond (1989). Our reputational environment, where everyone optimizes but people have hidden knowledge about their preferences, is closely linked to repeated games with incomplete information. Reputation in debt markets in which one player is a commitment type have been recently studied by Amador and Phelan (2021). The fact that reputation in one market may discipline behavior in another market has been considered in Cole and Kehoe (1998), Chatterjee et al. (2008), and Corbae and Glover (2025). A related paper that studies the reputational spillover from credit markets with the option to default to labor market earnings via its effect on search effort is Braxton et al. (2024).

1.2 Outline

The paper is structured as follows. Section 2 presents motivating facts for our paper, in particular new facts about inequality (as measured by cross-sectional variation) in credit rankings. There we also propose a simple reduced form “model” to help illustrate what forces might explain these facts. We then lay out our structural framework in a series of sections. Section 3 studies credit access and default with adverse selection in the simplest possible two-period version of Chatterjee et al. (2023),

⁷Kovrijnykh et al. (2024) build a simple but powerful model to rationalize the empirical fact that when there is little past information about the borrower’s creditworthiness, non-exclusive lenders with heterogeneous signals about a borrower’s type can learn from others’ approval decisions. By contrast, the credit history of an established borrower may provide sufficient information to supersede the need for others’ signals.

illustrating our methodological contribution to solving dynamic quantitative Bayesian equilibrium models with hidden information. Section 4 adds moral hazard to this model to endogenize how unobservable type interacts with unobservable effort. This lets us study the interaction between reputation in the credit market and income. Section 5 adds medical expense shocks to the basic adverse selection model to illustrate the costs and benefits of regulating credit record information as in a recent CFPB proposal. Section 6 puts the pieces together in a quantitative life cycle model to revisit the data of Section 2 and make quantitative statements rather than qualitative ones. To facilitate future work, all codes to replicate the model analysis in Sections 3 through 6 may be found at <https://zenodo.org/records/15731380>.

2 Reputational Inequality

It is well known that income and consumption inequality have a life-cycle pattern. Figure 1b displays these patterns as constructed from the dataset in Heathcote et al. (2023). While Figure 1b makes clear that both income and consumption inequality rise with age, it is also apparent that there is a divergence later in life between the two. One question is how much credit market access accounts for the divergence.


Less well-known, perhaps, is that people’s reputation in the credit market displays a similar pattern. Figures 1c and 1d show this pattern for credit rankings derived from credit scores, the latter being a commonly used measure of creditworthiness.⁸ For a person with score X , we define her credit ranking to be the fraction of people with scores equal to or less than X , i.e., the CDF of the overall score distribution at score X . The left panel shows that average credit rankings increase with age, and the right panel shows that credit rankings become more dispersed with age. What this means is that average scores in an age bracket increase across age brackets, and the dispersion of scores in a bracket around the bracket average also increases across age brackets.

The striking similarity between the patterns shown in the top and bottom panels of Figure 1 raises the question of whether they are related. A credit score is an inverse measure of the likelihood that an individual will become delinquent on some debt over the next two years.⁹ All else the same, we would expect higher-income individuals to be less leveraged (i.e., have less debt as a proportion of income) and, therefore, less likely to encounter debt problems. The increasing dispersion of credit scores with age could also result from the increasing dispersion of income with age. Thus, one possibility is that the life-cycle patterns in credit scores reflect the life-cycle patterns in income.


However, as already mentioned, the hypothesis that informs this paper and our previous work is that there is more to the life-cycle pattern of credit scores than just the life-cycle evolution of income. Specifically, the life-cycle of reputational inequality is driven, in part, by the market’s gradual learning of an individual’s preferences and, possibly, the life-cycle evolution of those preferences.

⁸We focus on credit rankings (an ordinal measure) rather than credit scores (a cardinal measure) since what does the number 780 really mean?


⁹More precisely, if δ is the probability of delinquency over the next two years, a credit score is the transform $a + b \cdot \ln(1 - \delta)$, where a and b are positive numbers.




(a) Mean, income and consumption



(b) Variance, income and consumption



(c) Mean, credit ranking




(d) Variance, credit ranking

Figure 1: Life cycle patterns in income, consumption, and credit rankings

Source for panels (a) and (b): Retrieved from <https://ideas.repec.org/c/red/ccodes/23-158.html>. The source of income data is the U.S. Census Bureau's Current Population Survey (CPS), and the source of expenditure data is the U.S. Bureau of Labor Statistics Consumer Expenditure Survey (CES), both for the sample period 1990-2021. The CPS sample consists of 1,486,687 observations after dropping 15,867 with 0 or negative income. The CES sample consists of 330,038 observations after dropping 124 with 0 or negative consumption values. The regression specification to obtain the figure included fixed time dummies. The intercepts, which were normalized for comparison purposes, are: (i) 10.5 and 8.6 for mean log income and consumption, respectively, and (ii) 0.66 and 0.37 for variance of log income and consumption, respectively.

Source for panels (c) and (d): A 2 percent subsample of the FRBNY Consumer Credit Panel (CCP)/Equifax panel for the 3rd quarter of 2024. The credit rankings are based on the VantageScore 4.0 credit score, which is a proprietary credit score of VantageScore Solutions LLC, similar to other credit scores used in the industry. Credit ranking is the percentile ranking of an individual's VantageScore in the overall distribution of VantageScores in our subsample.

Figure 2: Credit rankings and income



Source for panels (a) and (b): A 2 percent subsample of the FRBNY CCP/Equifax panel for the 3rd quarter of 2024. The credit rankings are based on the VantageScore 4.0 credit score, which is a proprietary credit score of VantageScore Solutions LLC, similar to other credit scores used in the industry. Credit ranking is the percentile ranking of an individual's VantageScore in the overall distribution of VantageScores. The income variable is an estimate of an individual's income provided by Equifax. An income ranking is the percentile of an individual's income in the overall distribution of estimated incomes in our subsample.


One reason for considering this hypothesis is shown in Figure 2a. It plots the percentage of people in the bottom, middle, and top thirds of the credit score distribution for each tercile of the income distribution. These segments of the credit score distribution correspond to the nonprime, prime, and superprime categories of borrowers, respectively.¹⁰ The plot shows that nonprime borrowers dominate the bottom tercile of income and that superprime borrowers dominate in the top. However, for the middle tercile, the three categories of borrowers are about even. For this middle group, there is little doubt that some factors other than income separate people by creditworthiness. More generally, the correlation between income and credit scores is 0.60, which, while positive and significant, leaves ample scope for factors other than income to play a role in determining creditworthiness.

A second reason is displayed in Figure 2b, which charts the correlation between credit and income rankings by age group. The correlation is lowest in the early to mid-twenties, rises quickly, and stabilizes by the early to mid-thirties. This pattern suggests that creditors learn about people's attitudes toward credit. Early in life, information on credit market attitude is scarce: people haven't made many credit decisions. This shows up as a weak correlation between income and credit rankings. People build up credit history as time passes, and the correlation between income and credit rankings strengthens. Still, the correlation rises to only about 0.66, which reiterates that there are non-income factors that influence credit rankings.

We conclude this section by noting that reputational inequality has consequences for credit

¹⁰For the credit scores used in this figure, the bottom third would correspond to individuals with scores below 660, the middle third to scores between 660 and 780, and the top third to scores between 780 and 850.

Figure 3: **Credit Card APRs by FICO Score and Income Quartile**



Source: [Dempsey and Ionescu \(2024\)](#). This figure is constructed using the Federal Reserve Board’s Y-14M data as described in Appendix A of the source article. The figure splits the data into 80 bins: 20 FICO bins $\times$ 4 income quartiles, and each data point corresponds to the within-bin average of the APR on the associated credit card account.

market access. Figure 3, taken from [Dempsey and Ionescu \(2024\)](#), uses administrative data from the Federal Reserve Board’s Y-14M data set to plot the relationship between FICO scores, income, and interest rates on credit card accounts. The plot shows that, conditional on an income quartile, interest rates decline with FICO scores. There is significant residual dispersion in APR based on incomes: for example, even with a high FICO score of 790, the gap in APR between the highest and lowest income quartiles is 1.5 percentage points.

2.1 A simple reduced form approach to the data

How can we explain some of the facts in Figures 1 and 2b? Let’s assume that person i ’s score $S_{i,n}$ in age bracket n is a linear function of their income $Y_{i,n}$ and a noisy signal about some unobservable trait $U_{i,n}$ inferred by individual’s observable credit market actions. Specifically, suppose

$$S_{i,n} = \alpha_n Y_{i,n} + U_{i,n}. \quad (1)$$

We think of equation (1) as a linear approximation in which the likelihood of repayment reflected in a credit score is increasing in both income and a noisy signal about some unobservable trait, say conscientiousness (reflecting, for example, patience). In that case, we can compute within age bracket means, variances, and correlations between score and income as in Figures 1 and 2b. For example, the correlation between score and income for age bracket n is given by:

$$\text{corr}(S_{i,n}, Y_{i,n}) = \frac{\text{cov}(S_{i,n}, Y_{i,n})}{\text{sd}(S_{i,n})\text{sd}(Y_{i,n})} = \frac{\alpha_n \sigma_{Y_n}^2 + \sigma_{Y_n, U_n}}{\sigma_{S_n} \sigma_{Y_n}}, \quad (2)$$

where $\sigma_{X_n}^2$ denotes the variance of X within bracket n and σ_{X_n, Z_n} denotes the covariance of X and Z within bracket n . Further, we know

$$\sigma_{S_n} = [\alpha_n^2 \sigma_{Y_n}^2 + 2\alpha_n \sigma_{Y_n, U_n} + \sigma_{U_n}^2]^{1/2}. \quad (3)$$

Finally, we can define $\bar{X}_n$ to be the mean of $X_{i,n}$ within bracket n .

What can we learn about the data through this simple reduced form approach? Figure 1a shows the well-documented result that mean income is hump-shaped. Figure 1c shows that mean credit rankings are rising throughout one's working life. The fact that credit rankings rise early in life along with income is not surprising since many models attribute higher income with higher ability to repay. The fact that credit scores keep rising while mean income falls later in life provides some evidence for the idea that learning about unobserved characteristics that are correlated with repayment for a rising segment of the population. That is, while $\bar{S}_n$ grows through time as a consequence of rising $\bar{Y}_n$ in (1) early in life, it may continue to grow through learning about an individual's unobservable type $\bar{U}_n$. We provide a structural model in later sections with hidden information and noisy public signals of unobservable type to consider this decomposition.

Figures 1b and 1d imply that $\sigma_{Y_n}^2$ and $\sigma_{S_n}^2$, respectively, are both weakly increasing across higher age brackets. Of course, equation (3) involves noisy signals about unobservables. If one believes credit scorers can learn about unobservable traits like conscientiousness (e.g., who is patient and does not want to risk their reputation versus who is impatient and doesn't care much about their reputation), then $\sigma_{U_n}^2$ will be increasing with age brackets since scorers are able to separate individuals by their credit market actions. Further, if unobservable type also affects effort choices (say through moral hazard), we might expect $\sigma_{Y_n, U_n} > 0$. Signing how σ_{Y_n, U_n} changes over age brackets is harder, but may be understood through the lens of our structural model in later sections.

Using data estimates from Figures 1a – 1d and 2b, we can infer the reduced form data generation process (DGP) in (1). We observe $\bar{S}_n$, $\bar{Y}_n$, $\sigma_{S_n}^2$, $\sigma_{Y_n}^2$, and $\text{corr}(S_{i,n}, Y_{i,n})$, but not $\sigma_{U_n}^2$, α_n , σ_{Y_n, U_n} , and $\bar{U}_n$. In this sense, we have 3 equations (1) – (3) in 4 unknowns. Therefore, with an assumption about one of these variables, the others are identified. To that end, if much of the learning about unobservable type from income realizations occurs early in an individual's life, we may assume that $\sigma_{Y_n, U_n} > 0$ is a decreasing function of age, after which we can solve for the others as:

$$\begin{aligned} \alpha_n &= \frac{\sigma_{S_n, Y_n} - \sigma_{Y_n, U_n}}{\sigma_{Y_n}^2} \\ \bar{U}_n &= \bar{S}_n - \alpha_n \bar{Y}_n \\ \sigma_{U_n}^2 &= \sigma_{S_n}^2 - \alpha_n^2 \sigma_{Y_n}^2 - \alpha_n \sigma_{Y_n, U_n}. \end{aligned}$$

To take the simple reduced form model to the data in Figures 1 and 2b, we estimate a linear fit on the bin-scatter of age and $\bar{S}_n$, a quadratic fit for $\bar{Y}_n$, $\sigma_{Y_n}^2$, and $\sigma_{S_n}^2$, and an exponential approach function for $\text{corr}(S_{i,n}, Y_{i,n}) = a - b \exp(-c \cdot n)$. Specifically, the first five panels in Figure 4 ((a) – (e)) provide the data (blue dots) and approximations (red lines). Panel 4(f) graphs our identifying




Figure 4: **Reduced form model**

assumption that learning about unobservable type from observing income realizations occurs early in life and the remaining 3 panels (4(g)-4(i)) graph the implied income sensitivity α_n and noisy signal moments $\bar{U}_n$ and $\sigma_{U_n}^2$ all in green.

The resulting unobservables in Figure 4 in green are consistent with an increase in the proportion of good risks (i.e., conscientious individuals) over later age bins (4(g)), increased separation (i.e., learning) over almost all age bins (4(h)), and less sensitivity of signal to income over later age bins due, for instance, to the accumulation of assets or lowering of borrowing costs for a growing proportion of conscientious individuals over time.

3 Reputation in a Simple Adverse Selection Model

We start studying reputation acquisition in the simplest possible two-period version of [Chatterjee et al. \(2023\)](#). That paper introduces unobservable shocks to budget feasible choices to get around the problem of off-the-equilibrium path beliefs in a hidden information model.

3.1 Environment

An individual lives two periods $n \in \{1, 2\}$.¹¹ There is a unit measure of two types $\tau \in \{H, L\}$ of individuals who differ in their patience, $\beta_H > \beta_L$. The fraction of type $\tau = H$ in the population is ρ . Importantly, type is unobservable. Individuals are risk averse with increasing and concave preferences $U(c_n)$. Individuals are endowed with income y_n . At $n = 1$, $y_1 = y_\ell$ for all individuals. At $n = 2$, $y_2 \in \{y_\ell, y_h\}$ is drawn with probability $Q_y(y_2) \equiv \mathbb{P}(y_2)$ with $y_\ell < y_h$.¹²

An asset market in unsecured discount bonds opens at $n = 1$ where individuals can choose $a_2 \in A \equiv \{\underline{a}, \dots, 0, \dots, \bar{a}\}$ at price q_1 . Let $A_- = \{a_2 \in A : a_2 < 0\}$ and $A_+ = A \setminus A_-$. Competitive lenders discount the future at rate r and recover nothing if a borrower defaults ($d_2 = 1$) on their debt $a_2 < 0$ at $n = 2$. Individuals begin $n = 1$ with no assets, $a_1 = 0$. We assume that $y_\ell > |\underline{a}|$ so that it is always budget feasible to pay back an individual's debt at $t = 1$; that is, all default is strategic and not due to an empty budget set.

Since type τ is unobservable and impatient individuals are less likely to care about the possible negative future consequences of default, lenders form assessments about an individual's type conditional on any observable correlated with their type when pricing debt. In this simple environment, the only information potentially correlated with type at the loan issuance phase is the individual's asset market choice a_2 . Free entry implies the competitive price is thus given by

$$q_1^{a_2} = \begin{cases} \frac{1}{1+r} & \text{if } a_2 \geq 0 \\ \frac{\mathbb{P}(d_2=0 \mid a_2)}{1+r} & \text{if } a_2 < 0 \end{cases}. \quad (4)$$

Notice that prices in our environment satisfy “block recursivity,” as in [Menzio and Shi \(2010\)](#). Specifically, under certain conditions on the information structure, our menu of nonlinear competitive prices plays the same role as free entry into submarkets in the competitive search paradigm; in neither case do we need to know the endogenous cross-sectional distribution of individuals to price debt, nor solve for equilibrium allocations. This is a general result in default pricing models like [Chatterjee et al. \(2007\)](#) and [Chatterjee et al. \(2023\)](#), except when information restrictions make it necessary for the lender to use the economy-wide cross-sectional distribution to form a prior.¹³

Since the likelihood of repayment $\mathbb{P}(d_2 = 0 \mid a_2)$ depends on an individual's unobservable type, lenders use updated assessments of an individual's type based on observable actions:

$$\psi_1^{a_2} \equiv \mathbb{P}(\tau = H \mid a_2). \quad (5)$$

¹¹Section 6 extends the number of periods agents live beyond 2.

¹²Throughout this paper, we denote probability with $\mathbb{P}(\cdot)$ and expectations with $\mathbb{E}(\cdot)$.

¹³See Appendix B.6.1 of ([Chatterjee et al., 2023](#)) for an example.

We call these updated assessments an individual's beginning-of-next-period *type score* $s_2 = \psi_1^{a_2}$. We assume these assessments satisfy Bayes' law. At $n = 1$, the prior probability an individual is of type H is given by the population measure of H types: $s_1 = \rho$.

At the beginning of $n = 2$, individuals receive an income realization y_2 drawn from Q_y and start with assets a_2 with type score s_2 . At that point, if $a_2 < 0$, the individual makes an observable choice d_2 . This provides another possible signal about an individual's type that can be used to form an updated assessment. Those with the opportunity to default (i.e. $a_2 < 0$) have their scores updated conditional on beginning-of-period observables $\omega_2 = (a_2, y_2, s_2)$ according to

$$s_3 = \psi_2^{d_2}(\omega_2) \equiv \mathbb{P}(\tau = H \mid d_2, \omega_2), \quad (6)$$

while individuals who take no default action (i.e., with $a_2 \geq 0$) have no new information upon which to condition an updated assessment, so their type score is simply retained and $s_3 = s_2$.

Given that agents live only two periods in this simple environment, why would anyone ever pay back debt? Unlike the infinite horizon environment in [Chatterjee et al. \(2023\)](#), here we simply assume individuals have preferences over their reputation $\phi(s_3)$ at the end of period $n = 2$ with $\phi'(s_3) > 0$. One can think of this reduced form preference over reputation as capturing how a parent's reputation may affect their access to other markets like insurance as in [Chatterjee et al. \(2008\)](#), labor as in [Corbae and Glover \(2025\)](#), or even children's access to future credit much like an intergenerational gift. The perpetual youth model in [Chatterjee et al. \(2023\)](#) provides a micro-foundation for $\phi(\cdot)$ via an individual's value function that is increasing in their posterior assessment due to enhanced credit access.

To summarize our information assumptions, in this simple case everything is observable – income (y_n), actions (a_2, d_2), and type scores (s_n) – except an individual's type (τ). The timing is:

1. At $n = 1$,
 - (a) an individual chooses a_2 taking the price schedule $q_1^{a_2}$ as given.
 - (b) individual type assessments $s_2 = \psi_1^{a_2}$ are updated according to Bayes' Law.
2. At $n = 2$,
 - (a) y_2 is realized from $Q_y(y_2)$, pinning down the observable state $\omega_2 = (a_2, y_2, s_2)$.
 - (b) if $a_2 < 0$, d_2 choice is made.
 - (c) observing d_2 , assessments of an individual's type are updated according to Bayes' Law: $s_3 = \psi_2^{d_2}(\omega_2)$ if $a(\omega_2) < 0$, and $s_3 = s_2$ otherwise.

3.2 Perfect Bayesian equilibrium

At age $n = 1$, an individual of type τ solves

$$V_1(\tau) = \max_{a_2 \in A} U(c_1^{a_2}) + \beta_\tau \sum_{y_2} Q_y(y_2) V_2(a_2, y_2, \psi_1^{a_2}, \tau) \quad \text{subj. to:} \quad c_1^{a_2} = y_\ell - q_1^{a_2} a_2, \quad (7)$$

with the optimal policy denoted $a_2^*(\tau)$. At timing stage 2(b), the individual solves:

$$V_2(\omega_2, \tau) = \max_{d_2 \in \{0,1\}} U\left(c_2^{d_2}(\omega_2)\right) + \beta_\tau \phi\left(\psi_2^{d_2}(\omega_2)\right) \quad \text{subj. to:} \quad c_2^{d_2}(\omega_2) = y(\omega_2) + (1 - d_2)a(\omega_2), \quad (8)$$

with the optimal policy $d_2^*(\omega_2; \tau)$ (if $a_2 \geq 0$ then $d_2^* = 0$ trivially).¹⁴ If $a_2 < 0$, the solution to (8) is simple: $d_2^*(\omega_2; \tau) = 0$ if

$$\phi\left(\psi_2^{d_2=0}(\omega_2)\right) - \phi\left(\psi_2^{d_2=1}(\omega_2)\right) \geq \frac{U(y(\omega_2)) - U(y(\omega_2) + a(\omega_2))}{\beta_\tau}, \quad (9)$$

where the left hand side of (9) is the reputational benefit of paying back while the right hand side is the cost of paying back. Notice that with $\beta_H > \beta_L$, the signaling cost is higher for the L type. Hence, L types are more likely to default. Notice further that given a concave utility function, $U(y_h) - U(y_h + a_2) < U(y_\ell) - U(y_\ell + a_2)$ so that the cost of paying back is higher in low income states, making it more likely that default happens in low income states.

3.3 Bayesian issues and a fix

We now explain a methodological contribution of our framework that makes quantitative analysis in this environment possible. Concretely, we illustrate an issue with the existence of a Bayesian equilibrium by considering the individual's decision to default in (8) on a particular level of debt $a_2 = -a$. We then show that incorporating extreme value shocks into the environment circumvents this issue.

The problem Equation (9) implies that default is more likely in low earnings states and for the riskier people (here the L type) as in real world data. Hence, consider individuals of both types who start with the same score s_2 and income y_2 , and conjecture that type H never defaults, while the riskier type L defaults when $y_2 = y_\ell$ but not when $y_2 = y_h$. Under this realistic conjecture, there is pooling when $y_2 = y_h$, and our Bayesian posteriors are:¹⁵

$$\begin{aligned} \psi_2^{d_2=1}(-a, y_h, s_2) &= \frac{0 \cdot s_2}{0 \cdot s_2 + 0 \cdot (1 - s_2)} = \frac{0}{0} \\ \psi_2^{d_2=0}(-a, y_h, s_2) &= \frac{1 \cdot s_2}{1 \cdot s_2 + 1 \cdot (1 - s_2)} = s_2. \end{aligned} \quad (10)$$

It is evident from (10) that we have an off-equilibrium-path (oep) beliefs issue when $d_2 = 1$. In that case, many researchers simply assume since there is no new information in actions to separate people, the oep posterior is unchanged:

$$s_3^{\text{oep}}(d_2 = 1, -a, y_h, s_2) = s_2. \quad (11)$$

¹⁴Note that we have adopted the notation $y_2 = y(\omega_2)$, for example.

¹⁵Note that since there is separation when $y_2 = y_\ell$, type is perfectly revealed in the low income state. Formally, the posteriors are $\psi_2^{d_2=1}(-a, y_\ell, s_2) = \frac{0 \cdot s_2}{0 \cdot s_2 + 1 \cdot (1 - s_2)} = 0$ and $\psi_2^{d_2=0}(-a, y_\ell, s_2) = \frac{1 \cdot s_2}{1 \cdot s_2 + 0 \cdot (1 - s_2)} = 1$.

A fundamental question arises: is this conjectured behavior – no default in high income states and infrequent default in low income states, similar to what we see in the data – consistent with a Bayesian equilibrium given what are often considered “reasonable” oep beliefs? In the $y_2 = y_\ell$ case where separation yields well-defined beliefs, the following conditions must hold for the high type to repay and low type to default, respectively:

$$\frac{U(y_\ell) - U(y_\ell - a)}{\beta_L} > \underbrace{\phi\left(\psi_2^{d_2=0}(-a, y_\ell, s_2)\right) - \phi\left(\psi_2^{d_2=1}(-a, y_\ell, s_2)\right)}_{=\phi(1)-\phi(0)} \geq \frac{U(y_\ell) - U(y_\ell - a)}{\beta_H}.$$

Both conditions can hold with sufficiently low β_L . For the troublesome $y_2 = y_h$ case where agents are pooled since no one defaults, the following must hold to support the conjectured behavior:

$$\underbrace{\phi\left(\psi_2^{d_2=0}(-a, y_h, s_2)\right) - \phi\left(\psi_2^{d_2=1}(-a, y_h, s_2)\right)}_{=\phi(s_2)-\phi(s_3^{\text{oep}})} \geq \frac{U(y_h) - U(y_h - a)}{\beta_\tau} \text{ for all } \tau. \quad (12)$$

What conditions on oep beliefs make this an equilibrium? If we take the oep belief $s_3^{\text{oep}} = s_2$ from (11), then the left hand side of (12) becomes $\phi(s_2) - \phi(s_2) = 0$, and the required inequality cannot hold. Therefore, this reasonable assumption for oep beliefs results in *non-existence* of our conjectured Bayesian equilibrium where risky types default in low earnings states.¹⁶

The solution To continue our simple example, there are two discrete choice problems: (i) the asset choice ($a_2 \in A_- \cup A_+$) and (ii) whether to default or not ($d_2 \in \{0, 1\}$) on debt $a_2 \in A_-$. Following the discrete choice literature, we assume that an individual’s discrete choices over assets and default are subject to shocks $(\varepsilon^{a_2}, \varepsilon^{d_2})$ drawn from a Type I extreme value distribution that are unobservable to anyone (e.g., econometricians performing risk assessments) except the individual.¹⁷ Economically, these shocks may proxy other events besides income realizations, such as unexpected expenditures (medical bills, auto breakdowns, etc.), which are plausibly unobserved (to the econometrician working at the credit scoring agency).

How do these shocks resolve the non-existence problem described above? Recall that the core issue was the undefined posterior in (10) in state $y_2 = y_h$. Introducing these shocks eliminates perfect pooling (i.e., $d_2^*(-a, y_h, s_2; \tau) = 0$ for all τ): some of each type will draw publicly unobservable preference shocks that drive them to default even if the “fundamental” value of doing so is lower than repaying. The decision rule will be defined not by binary outcomes in $\{0, 1\}$, but rather by

¹⁶An alternative is to assume the worst possible oep belief (akin to a harsher punishment) in the $y_2 = y_h$ pooling case where not all feasible actions are taken in equilibrium. In order for both types to repay when $s_3^{\text{oep}} = 0$, a necessary condition to make this an equilibrium is given by $\phi(s_2) - \phi(0) \geq \frac{U(y_h) - U(y_h - a)}{\beta_L}$. While with this specification of oep beliefs it may be possible to find a set of parameters to support our conjectured Bayesian equilibrium, it makes existence of an equilibrium dependent on the value of reputation encoded in our reduced form function ϕ . Such dependence could make existence a delicate matter in more general models.

¹⁷The Type I extreme value cumulative distribution function for discrete choice $x \in X$ is given by $F_\epsilon(\epsilon^x; \bar{x}, \alpha_x) = \exp\left\{-\exp\left(-\frac{\epsilon^x - \bar{x}}{\alpha_x}\right)\right\}$, where $\bar{x}$ is the location parameter and α_x is the scale parameter that governs its variance.

choice probabilities in the range (0,1), which depend on type. That is, the presence of these shocks rules out perfect separation of types, slowing down learning about agents' types. Without such shocks early in life, perfect separation would eliminate the need for updated type scores.

Modified equilibrium with extreme value shocks This change to the environment requires us to change the optimization problems in (7) and (8). In the case of (8), we add unobservable preference shocks over the default choice (ε^{d_2}) drawn from an extreme value distribution with scale parameter set to α_d to now solve:

$$V_2(\omega_2, \tau) = \mathbb{E}_{\varepsilon^{d_2}} \left[\max_{d_2 \in \{0,1\}} \left\{ v_2^{d_2}(\omega_2, \tau) + \varepsilon^{d_2} \right\} \right], \quad (13)$$

$$\text{where } v_2^{d_2}(\omega_2, \tau) = U(y(\omega_2) + (1 - d_2)a(\omega_2)) + \beta_\tau \phi(\psi_2^{d_2}(\omega_2))$$

is the fundamental value of choosing action d_2 . Note that here and throughout the paper, we use the notational convention that the subscript denotes the age, the superscript denotes an action of the individual (or a set of actions), and the functional arguments in parentheses are individual state variables. In a simple 2-period environment in which all individuals have the same ω_1 , we suppress this argument for notational convenience. We also use the notation $y(\omega)$ or $a(\omega)$ to denote the y or a argument in the vector $\omega = (a, y, s)$. The solution to (13) yields choice probabilities:

$$\sigma_2^{d_2}(\omega_2, \tau) = \exp \left\{ \frac{v_2^{d_2}(\omega_2, \tau)}{\alpha_d} \right\} / \sum_{d_2} \exp \left\{ \frac{v_2^{d_2}(\omega_2, \tau)}{\alpha_d} \right\}. \quad (14)$$

In the case of no default, then, the choice probability is

$$\sigma_2^{d_2=0}(\omega_2, \tau) = \left[1 + \exp \left\{ \frac{U(y(\omega_2)) - U(y(\omega_2) + a(\omega_2)) - \beta_\tau \left[\phi(\psi_2^{d_2=0}(\omega_2)) - \phi(\psi_2^{d_2=1}(\omega_2)) \right]}{\alpha_d} \right\} \right]^{-1}.$$

This implies that as long as $\phi'(s_3) > 0$, type H individuals are less likely to default than type L when default lowers the Bayesian posterior of a person's type.¹⁸

Finally, notice that the choice probabilities in (14) are strictly bounded in (0,1) so there is no longer an issue with supplying off-equilibrium-path beliefs associated with Bayesian posteriors for the perfectly pooling case in (10). Specifically, posteriors are now given by

$$\psi_2^{d_2}(\omega_2) = \frac{s_2 \sigma_2^{d_2}(\omega_2; H)}{s_2 \sigma_2^{d_2}(\omega_2, H) + (1 - s_2) \sigma_2^{d_2}(\omega_2, L)}. \quad (15)$$

Since $V_2(\omega_2, \tau)$ is increasing in β_τ , $\sigma_2^{d_2=0}(\omega_2, H) > \sigma_2^{d_2=0}(\omega_2, L)$, and so $\psi_2^{d_2=0}(\omega_2) > \psi_2^{d_2=1}(\omega_2)$.

In order to keep the Bayesian posterior $\psi_1^{a_2}$ well-defined, we associate extreme value shocks to

¹⁸That is, $U(y(\omega_2)) - U(y(\omega_2) + a(\omega_2)) - \beta_H \left[\phi(\psi_2^{d_2=0}(\omega_2)) - \phi(\psi_2^{d_2=1}(\omega_2)) \right] < U(y(\omega_2)) - U(y(\omega_2) + a(\omega_2)) - \beta_L \left[\phi(\psi_2^{d_2=0}(\omega_2)) - \phi(\psi_2^{d_2=1}(\omega_2)) \right]$ if and only if $\phi(\psi_2^{d_2=1}(\omega_2)) < \phi(\psi_2^{d_2=0}(\omega_2))$.

the asset decisions at age 1 as well. This yields the adapted version of problem (7):

$$V_1(\tau) = \mathbb{E}_{\varepsilon^{a_2}} \left[\max_{a_2 \in A} \{v_1^{a_2}(\tau) + \varepsilon^{a_2}\} \right], \quad (16)$$

$$\text{where } v_1^{a_2}(\tau) = U(y_\ell - q_1^{a_2} a_2) + \beta_\tau \sum_{y_2} Q_y(y_2) V_2(a_2, y_2, \psi_1^{a_2}; \tau).$$

We refer to $v_1^{a_2}(\tau)$ as the “fundamental” value of choosing a_2 given type τ . This problem has as its solution the decision density over asset choices $\sigma_1^{a_2}(\tau)$. Note that the shocks to individuals’ choices result in choice probabilities that imply that equilibria are semi-separating (or partially pooling); if equilibria were fully separating, there would be no need to assess individuals’ creditworthiness via something like a credit score. Given this structure, we can write the equilibrium reputation updating function (5) and loan pricing function (4) as

$$\psi_1^{a_2} = \frac{\rho \sigma_1^{a_2}(H)}{\rho \sigma_1^{a_2}(H) + (1 - \rho) \sigma_1^{a_2}(L)} \quad (17)$$

$$q_1^{a_2} = \frac{1}{1 + r} [\psi_1^{a_2} \bar{p}^{a_2}(H) + (1 - \psi_1^{a_2}) \bar{p}^{a_2}(L)], \quad (18)$$

$$\text{where } \bar{p}^{a_2}(\tau) = \sum_{y_2} Q_y(y_2) \sigma_2^{d_2=0}(a_2, y_2, \psi_1^{a_2}, \tau) \quad (19)$$

is the type-specific expected probability of repayment on a loan of size a_2 .

3.4 Reputational incentives and price effects at age 1

We use this simple model to illustrate how reputational incentives affect individuals’ savings decisions. Ignoring extreme value shocks and any kinks in the continuation value associated with the default option, and suppressing the explicit state arguments in the value functions to ease notation, the net marginal value of an increase in savings a_2 is

$$\frac{\partial v_1^{a_2}}{\partial a_2} = \beta_\tau \sum_{y_2} Q_y(y_2) \left[\frac{\partial V_2}{\partial a_2} + \frac{\partial V_2}{\partial s_2} \frac{\partial \psi_1^{a_2}}{\partial a_2} \right] - U'(c_1) \left(q_1^{a_2} + \frac{\partial q_1^{a_2}}{\partial a_2} a_2 \right). \quad (20)$$

Given the competitive structure of the model, the individual takes q_1 , ψ_1 and their derivatives as given. This marginal value would be set equal to 0 to find the optimal a_2 in the “standard” model without extreme value shocks. Absent limited commitment (and therefore default), the price elasticity term $\frac{\partial q_1^{a_2}}{\partial a_2}$ would be equal to zero. Likewise, absent the information friction, the reputation elasticity term $\frac{\partial \psi_1^{a_2}}{\partial a_2}$ would be equal to zero. In this frictionless case, then, the optimal policy would equate the relative price of age 2 consumption in terms of age 1 consumption to the analogous marginal rate of substitution: $q_1^{a_2} = \beta_\tau \frac{\partial V_2}{\partial a_2} / U'(c_1)$.

How do the key frictions in our environment affect this tradeoff? As has been well-documented in the literatures on consumer and sovereign default, limited commitment induces the price elasticity term $\frac{\partial q_1^{a_2}}{\partial a_2}$ to be non-zero. In a world without information frictions, this elasticity is governed entirely

by a leverage effect: since the marginal change in expected repayment probability *conditional on type* $\frac{\partial \bar{p}^{a_2}(\tau)}{\partial a_2}$ is generically positive (smaller loans are less risky), limited commitment deters borrowing because larger loans are more expensive than smaller loans.

Our primary contribution is to study how the presence of hidden information further affects the key tradeoff described in (20). There are two separate effects. First, hidden information compounds the standard leverage effect on pricing associated with limited commitment described in the previous paragraph. Concretely, in addition to the type-specific term $\frac{\partial \bar{p}^{a_2}(\tau)}{\partial a_2}$, lenders now account for the fact that a change in the loan size requested may be revelatory about the borrower's type. Differentiating the pricing equation (18), we obtain

$$\frac{\partial q_1^{a_2}}{\partial a_2} = \frac{1}{1+r} \left[\underbrace{\psi_1^{a_2} \frac{\partial \bar{p}^{a_2}(H)}{\partial a_2} + (1 - \psi_1^{a_2}) \frac{\partial \bar{p}^{a_2}(L)}{\partial a_2}}_{\text{limited commitment}} + \underbrace{\frac{\partial \psi_1^{a_2}}{\partial a_2} (\bar{p}^{a_2}(H) - \bar{p}^{a_2}(L))}_{\text{reputation}} \right] \quad (21)$$

The first term is the standard limited commitment effect: conditional on type, smaller debts lower default probabilities and raise prices (or lower interest rates). The only difference in this first term relative to a full information model is that we must weight by the likelihood of the borrower being of each type, which depends on the updated type assessment according to (17). The second term adjusts for the revised type assessment: to the extent that H types are less likely to default than L types on a loan of a given size, and that H types are more likely to save more (or borrow less), this further increases the price elasticity term.

The second effect comes from the continuation value term $\frac{\partial V_2}{\partial s}$ in (20). To the extent that saving more improves one's reputation (i.e. $\frac{\partial \psi_1^{a_2}}{\partial a_2} > 0$), saving more has an additional reputational benefit on top of the standard increase in wealth accounted for in the $\frac{\partial V_2}{\partial a}$ term. Our framework provides a way to quantify both the marginal value of reputation and the elasticity of reputation with respect to a given action, both of which are critical for determining the strength of reputational incentives.

3.5 Model properties

This subsection illustrates the equilibrium implications of individuals' incentives to acquire reputation. Since we simply mean to illustrate model properties in this simple framework (leaving until Section 6 a more serious parameterization), we assume the fraction of H types $\rho = 2/3$ broadly in line with the fraction of prime borrowers in the U.S. economy. The full parameterization for these numerical illustrations can be found in <https://zenodo.org/records/15731380>.

3.5.1 On adverse selection

The extent of adverse selection depends on how fundamentally different the two types are. Figure 5 illustrates the role of type differences by holding β_H fixed and varying β_L in the range $[\underline{\beta}, \beta_H]$. The fact that H -types take on less debt and are less likely to default than L -types provides justification for listing them as "less risky." When $\beta_L = \beta_H$, there is no information problem: the two types




Figure 5: **Adverse selection affects borrowing and default behavior**

Notes: Both panels present moments of the model for a range of values of β_L between β and β_H . Panel (a) plots the fraction of agents in debt (choosing $a_2 < 0$) in aggregate and by type, while panel (b) plots the default probability conditional on being in debt (aggregated across debtors) in aggregate and by type. The vertical line in each graph corresponds to the benchmark parameter value for β_L used throughout the numerical illustrations.

behave the same, so there is nothing to infer. As *only* L -types become more impatient (starting from the right of the figure), though, adverse selection implies that the behavior of *both* types changes.

When the types are close, borrowing behavior is similar, but the gap in borrowing rates grows as the types become more different. While the L -types' change in behavior stems primarily from impatience, the H -types' change combines two effects. First, price schedules shift down to reflect the increased riskiness of the overall pool of borrowers, deterring borrowing. Second, H -types further ration borrowing in order to separate themselves from L -types to bolster their reputations.

The default behavior in Figure 5(b) is more subtle. When the types are close, increasing the distance between them induces little change in default probability for either type, although H -types do cut their default slightly to preserve their reputations. When the types become sufficiently far apart, though, reputational incentives take over and both types sharply cut their default rates. Once H -type default becomes sufficiently unlikely, though, the fundamental impatience of the L -types takes over, and their default rates rise once more.

Fixing the difference between types, adverse selection creates incentives to acquire reputation. In this 2-period model, the strength of those incentives is governed by ϕ : households do not value reputation at all when $\phi = 0$, while they value reputation a lot when $\phi = \bar{\phi}$. Figure 6 documents how strengthening the incentives to acquire a good reputation changes economic outcomes.

When the value of reputation is zero, there are no reputational incentives. Since H -types save more, saving improves one's reputation: as ϕ increases, this increases savings rates of both types (panel (a)). The only differences in default at age 2 (panel (b)) come from reputational effects. As ϕ increases, H -types default less. When reputation becomes sufficiently important, they cut default so sharply that L -types – who also value reputation, albeit less – also cut their default rates. This expands credit access: since both types default less, loan prices improve. This causes a substitution effect that makes the net effect on savings ambiguous: panel (a) shows that the first




Figure 6: **Adverse selection and the value of reputation**

Notes: Each panel presents the moment indicated in the title over a range of values of ϕ for both H and L types. All panels are averages across the equilibrium distribution. The vertical line in each graph corresponds to the benchmark parameter value for ϕ used throughout the numerical illustrations.

effect dominates for H -types, while the two effects approximately cancel for L -types. As ϕ increases even more, though, individuals' reputations are more determined early in life by savings decisions, and the marginal effects of default are second order: therefore, default rates increase again. Panel (c) shows that as ϕ increases, H -types who value reputation more increasingly separate themselves from L -types with their savings. L -type behavior changes similarly, but less than one-for-one. While L -types value reputations and take steps to preserve them, given the shift in H -type behavior, the lower weight they place on future reputation leads to increased separation. As the value of reputation increases, the equilibrium features more (but still far from complete) separation.

How do these dynamics shape inequality over the life cycle? Panels (d) and (e) of Figure 6 show that as ϕ increases, consumption inequality increases alongside wealth inequality, *despite the fact that income inequality is unchanged* by construction. The consumption effect mirrors the savings effect, as H -types reveal their patience by postponing consumption to the future. Moreover, “reputational inequality” – as measured by the dispersion in type scores between the two fundamental types – increases as well. This illustrates that reputational inequality directly relates to other (more commonly measured) forms of inequality.

3.5.2 Life cycle patterns

Figure 7 plots how wealth, reputation, consumption, and income evolve over the life cycle in the aggregate and by type. Even though the only heterogeneity at age 1 is type, asset choices induce fanning out of wealth at age 2 and consumption at both ages. This figure shows that even this simple model can make sharp predictions about the evolution of inequality over the life cycle.




Figure 7: **Life cycle in the baseline model**

Notes: Each panel in this figure plots the indicated moment in the title of the panel for ages 1 and 2 in aggregate and by type. The error bands in each figure correspond to plus or minus one standard deviation, aggregating across both high and low types. Type-specific error bands are not shown.

All agents begin with the same income $y_1 = y_\ell$, and income increases on average over the life cycle exogenously. Ex post some individuals are “lucky” ($y_2 = y_h$), while others are “unlucky” ($y_2 = y_\ell$), which generates income inequality (panel (d)). Notably, there is no difference in income across types. Savings, however, do differ across types. On average, individuals save modestly despite their expected income growth due to reputational incentives. Underlying this aggregated result, though, is divergence between types: H -types save, while L -types borrow. Panel (c) presents a corollary of this divergence: there is consumption inequality *across types* at both ages. In particular, impatient L types consume more than patient H types at age 1, and vice versa at age 2. Panel (b) shows how these behaviors impact the evolution of reputation. At age 1, everyone has a common type score, but asset choices lead to divergence in the updated assessments s_2 . Since H types tend to save more than L types, they tend to see an improvement in their reputation at age 2. Reversing this logic leads to a decline in reputation on average among L types.

3.5.3 Credit record information restrictions

An important policy debate centers around what can enter a credit record. As discussed in the introduction, payment history is the primary factor (41%) used in Vantage credit scoring models. Further, the Fair Credit Reporting Act (FCRA) mandates that a bankruptcy must be wiped from

	baseline			no tracking		
	mean	diff, $H - L$	std. dev.	mean	diff, $H - L$	std. dev.
A. Inequality moments						
age 2 wealth	0.01	0.11	0.11	-0.06	0.02	0.09
age 1 consumption	0.73	-0.09	0.09	0.75	-0.01	0.05
age 2 consumption	0.82	0.07	0.18	0.80	0.01	0.15
B. Driving behavior						
fraction in debt (%)	46.9	-44.9		78.5	-7.2	
default rate (%)	7.5	-21.7		52.0	-6.6	

Table 1: **The role of information in the baseline model**

Notes: This table reports the statistic in the column header for the moment in the row header for the model version indicated in the top row of the table. The mean is the population average, and the standard deviation is the cross-sectional standard deviation. The difference reported is the difference for the average of the moment for H -types less the average of the moment for L -types.

a consumer’s credit report after 10 years. From the perspective of our model, the factors that are included or excluded in the information set upon which lenders base their risk assessments matter crucially for credit access (here, interest rates).

To consider the impact of restricting the information upon which lenders can condition their assessments, we compare our baseline model to a version in which type scores are not tracked over time. In this “no tracking” case, there is no incentive to acquire a better reputation despite the presence of hidden information.¹⁹ The details of this model are specified fully in Appendix A.3.1.

Table 1 examines how the reputation acquisition incentive shapes consumption and savings over the life cycle by documenting patterns in both the baseline and no tracking economies. All the reported differences may be attributed to reputational incentives since the economies are otherwise identical. There are three key findings. First, agents *of both types* save more (or borrow less) in the baseline than in the no tracking economy. Since saving suggests patience, and reputation is valued, the marginal value of saving at age 1 is higher in the baseline than in the no tracking economy. Second, since reputation accrues in the future, the increase in wealth for patient H -types exceeds the increase for L -types, and so there is an increase in wealth inequality *across the two types*. Notably, there is also an increase in overall wealth inequality, as measured by the cross-sectional standard deviation. Third, consistent with the analysis above, these savings dynamics imply that

¹⁹In Chatterjee et al. (2023), we considered two different informational assumptions: “no-tracking,” as described above, and “full information,” in which types are directly observable. In this two-period model, these two economies are identical in the sense that they deliver the same allocation. In both cases, there is no type score state variable for the individual, and so there are no dynamic considerations relevant for the age 2 default decision. The only difference is that in the no tracking economy, loan prices cannot be conditioned directly on type, while in the full information economy they can. However, in this 2-period setting, the only difference in default behavior among the two types comes from reputational incentives at age 2. Since both the full information and no tracking economies have no reputation updates, then, there is no difference in how the two types behave conditional on being in debt at age 2 across the economies. Therefore, the loan price schedules faced at age 1 must be the same and hence the two economies deliver the same allocation.

consumption inequality – both between types and in the population as a whole – increases at both ages in the baseline relative to the no tracking economy.

Panel B of Table 1 documents the drivers of these outcomes. First, while the economy-wide share of borrowers drops in the baseline compared to the no tracking economy, the drop is steeper for H -types: in the no tracking world, only 7.2% more L -types borrow than L -types, while this difference is 44.9% in the baseline model. Second, age 2 default rates evolve similarly. This induces a price effect that amplifies the effect of the changes in borrowing on consumption: cheaper credit allows agents to transfer more resources across time.

4 Reputation and Income: Adding Moral Hazard

The earnings process in the previous section featured no dependence on an individual’s type, as in Chatterjee et al. (2023). Hence, there was no income inequality *between the two types* in that economy, as evident in Figure 7(b). In that case, consumption inequality (Figure 7(c)) was induced primarily by reputational concerns. Here, we extend the model to allow the income process to depend on type endogenously. Specifically, we assume an individual makes a hidden costly effort choice at age 1 that raises the probability of a high income realization at age 2. H -types who care more about the expected present discounted value of future income are more likely to bear that cost today, inducing dependence of earnings on type as in Corbae and Glover (2025). Thus, income inequality may be induced by different unobservable effort choices across unobservable type. This adds a mechanism by which type differences may amplify consumption inequality by inducing income inequality. At the same time, it complicates the relationship between reputation and income, as one’s income is informative about one’s type unlike the previous section where it is uninformative.

4.1 Environment

The environment is identical to Section 3.1 with one key difference. Here, individuals make an unobservable effort choice e_1 at the beginning of age 1 that raises the likelihood of receiving y_h at age 2 via the endogenous earnings process $Q_y^{e_1}(y_2) \equiv \mathbb{P}(y_2 \mid e_1)$. This replaces the exogenous earnings process from Section 3.1. Specifically, at $n = 1$ each individual chooses whether or not to exert effort $e_1 \in \{0, 1\}$ subject to an extreme value shock ε^{e_1} . This unobservable, costly effort choice influences the income process: exerting effort ($e_1 = 1$) raises the probability of receiving high income at $n = 2$; that is, $Q_y^{e_1=1}(y_h) > Q_y^{e_1=0}(y_h)$. The utility cost for exerting effort is κ , and the gain associated with higher future expected income is discounted by β_τ .

Because effort is unobservable, the assessment of an individual’s type after observing their asset choice is not the final revision before the age 2 default decision. Instead, an individual’s type score at age 2 depends also on the realization of their income at age 2, according to the function

$$s_2 = \Upsilon_1^{a_2}(y_2) \equiv \mathbb{P}(\tau = H \mid a_2, y_2). \quad (22)$$

This formula extends equation (5) from Section 3.1 to account for the fact that different types may have different effort probabilities and therefore different likelihoods of each income realization. Since the probability is computed differently, prices and assessments of type differ relative to the model in the previous section.

To summarize our information assumptions for this environment, an individual's type (τ) and effort choice (e_1) are unobservable: the former yields adverse selection and the latter moral hazard. Income y_n is observable, as are credit market actions (a_2, d_2) and type score (s_n). All actions (e_1, a_2, d_2) are subject to unobservable idiosyncratic shocks ($\varepsilon^{e_1}, \varepsilon^{a_2}, \varepsilon^{d_2}$). To summarize timing:

1. At $n = 1$,
 - (a) an individual chooses unobservable effort e_1 .
 - (b) an individual chooses a_2 taking the price schedule $q_1^{a_2}$ as given.
2. At $n = 2$,
 - (a) y_2 is realized from $Q_y^{e_1}(y_2)$.
 - (b) scorers update their assessment according to Bayes' Law via $s_2 = \Upsilon_1^{a_2}(y_2)$, yielding observable state $\omega_2 = (a_2, y_2, s_2)$.
 - (c) if $a_2 < 0$, d_2 choice is made.
 - (d) observing d_2 , individual type score is updated according to Bayes' Law: $s_3 = \psi_2^{d_2}(\omega_2)$ if $a(\omega_2) < 0$ or $s_3 = s_2$ otherwise.

Thus, the only differences from the timing in Section 3.1 is the addition of 1(a), 2(a), and the substitution of Υ_2 in 2(b), which defines the law of motion for an individual's type score, for ψ_1 in 1(c) of the timing from the baseline model in Section 3. This accounts for the fact that the realization of y_2 is informative about type in this setting.

4.2 Equilibrium

At age 1, the individual first makes an effort choice, weighing the cost of effort against the continuation value associated with exerting that effort:

$$V_{E,1}(\tau) = \mathbb{E}_{\varepsilon^{e_1}} \left[\max_{e_1 \in \{0,1\}} \left\{ v_1^{e_1}(\tau) + \varepsilon^{e_1} \right\} \right] \text{ where } v_{E,1}^{e_1}(\tau) = -\kappa e_1 + V_{A,1}^{e_1}(\tau), \quad (23)$$

where the continuation value reflects the second decision over assets:

$$\begin{aligned} V_{A,1}^{e_1}(\tau) &= \mathbb{E}_{\varepsilon^{a_2}} \left[\max_{a_2} \left\{ v_{A,1}^{(a_2, e_1)}(\tau) + \varepsilon^{a_2} \right\} \right], \\ \text{where } v_{A,1}^{(a_2, e_1)}(\tau) &= U(y_\ell - q_1^{a_2} a_2) + \beta_\tau \sum_{y_2} Q_y^{e_1}(y_2) V_2(a_2, y_2, \Upsilon_1^{a_2}(y_2), \tau). \end{aligned} \quad (24)$$

The first decision problem generates the type-specific decision rule over effort $\sigma_{E,1}^{e_1}(\tau)$, while the second generates the effort-choice and type-specific decision rule over assets $\sigma_{A,1}^{(a_2, e_1)}(\tau)$. Note that

discounting occurs intertemporally, not between effort and asset choices at age 1. At age 2, if $a_2 < 0$ was chosen at age 1, the individual makes the exact same default decision (13) as in the baseline model, yielding decision rule $\sigma_2^{d_2}(\omega_2, \tau)$.

Note that type scores and prices are computed similarly to the baseline model from Section 3. In order to streamline notation in defining these objects, it is useful to define four objects. First, the age-1 joint decision rule over effort and assets is $\sigma_1^{(a_2, e_1)}(\tau) = \sigma_{E,1}^{e_1}(\tau) \sigma_{A,1}^{(a_2, e_1)}(\tau)$. Second, the effort-choice-weighted age-1 observable decision rule over assets is $\bar{\sigma}_1^{a_2}(\tau) = \sum_{e_1} \sigma_1^{(a_2, e_1)}(\tau)$. Third, the probability of a given y_2 realization based on type and asset choice is

$$\tilde{Q}_y^{a_2}(y_2; \tau) \equiv \mathbb{P}(y_2 \mid a_2, \tau) = \sum_{e_1} Q_y^{e_1}(y_2) \frac{\sigma_1^{(a_2, e_1)}(\tau)}{\bar{\sigma}_1^{a_2}(\tau)}.$$

Fourth, although the belief update based solely on the choice of a_2 no longer defines the evolution of one's type score in this model, it is still useful as an input to loan prices and the full type score update. We compute the type assessment made *after* the choice of a_2 but *before* the realization of y_2 exactly as in (17), with the only difference being that we need to sum over unobservable effort decisions: that is, we replace $\sigma_1^{a_2}(\tau)$ with $\bar{\sigma}_1^{a_2}(\tau)$.

With repeated application of Bayes' Law, we can write the type score updating equation (22) using the "interim" type score $\psi_1^{a_2}$ and the updated y_2 density $\tilde{Q}_y^{a_2}$:

$$\Upsilon_1^{a_2}(y_2) = \frac{\psi_1^{a_2} \tilde{Q}_y^{a_2}(y_2; H)}{\psi_1^{a_2} \tilde{Q}_y^{a_2}(y_2; H) + (1 - \psi_1^{a_2}) \tilde{Q}_y^{a_2}(y_2; L)}. \quad (25)$$

Conditional on the revised type assessment associated with the asset choice, $\psi_1^{a_2}$, equation (25) further revises the type assessment by evaluating how likely the individual is to be of a certain type given the income realization y_2 . Finally, prices in this model take the same form as (18), with the one change being that now the type-specific repayment probability takes into account the updated expectation over y_2 given the information from the a_2 choice: that is, equation (19) is now

$$\bar{p}_1^{a_2}(\tau) = \sum_{y_2} \tilde{Q}_y^{a_2}(y_2; \tau) \sigma_2^{d_2=0}(a_2, y_2, \Upsilon_1^{a_2}(y_2), \tau). \quad (26)$$

The type score updates after the age 2 default decision follow (15) exactly as in the baseline model.

4.3 Reputational incentives and income with moral hazard

As discussed in Section 3.4, one of the key aspects of our environment is how the incentive to acquire reputation affects choices. In the equilibrium described above, the same basic forces hold for an additional choice: effort. How do reputational incentives shape this choice?

Define the "fundamental" value of given effort and asset choices across stages at age 1 as

$$v_1^{(a_2, e_1)}(\tau) = -\kappa e_1 + U(y_\ell - q_1^{a_2} a_2) + \beta_\tau \sum_{y_2} Q_y^{e_1}(y_2) V_2(a_2, y_2, \Upsilon_1^{a_2}(y_2), \tau),$$

where we have written the expectation term explicitly. What are the net benefits of choosing to exert effort? As a binary choice, we can simply fix the asset choice and compute

$$v_1^{(a_2,1)}(\tau) - v_1^{(a_2,0)}(\tau) = \beta_\tau \sum_{y_2} [Q_y^1(y_2) - Q_y^0(y_2)] V_2(a_2, y_2, \Upsilon_1^{a_2}(y_2), \tau) - \kappa.$$

This shows that the benefit of increasing the likelihood of high y_2 realizations is weighed against the fundamental cost of exerting effort, conditional on the asset choice.

Since effort makes a high income tomorrow more likely, we must understand the value of having a higher income tomorrow. Taking a total derivative of the continuation value term yields

$$\frac{dV_2}{dy_2} = \frac{\partial V_2}{\partial y_2} + \underbrace{\frac{\partial V_2}{\partial s_2} \frac{\partial \Upsilon_1^{a_2}(y_2)}{\partial y_2}}_{\text{reputation}}.$$

The first term is standard: increasing income tomorrow expands the budget set. The second term, though, comes from reputational incentives: to the extent that reputation is valued ($\frac{\partial V_2}{\partial s_2} > 0$) and that high income improves one's reputation ($\frac{\partial \Upsilon_1}{\partial y_2} > 0$), the benefit of exerting effort is higher, making effort more likely to be exerted. The former effect is true so long as reputation is valued, for example because $\phi > 0$. The second is true so long as H types are more likely to have high y_2 realizations, which is generically true in our environment since the benefits of exerting effort are discounted, while the costs are not. Since a segment of the population responds to these incentives more strongly than another, such reputational incentives create scope for expanding income inequality.

4.4 Model properties

4.4.1 Interaction between adverse selection and moral hazard

One of the main features of the model from Section 3 is that adverse selection creates an incentive for individuals to acquire reputation. How does this incentive change in the presence of moral hazard? To answer this, Figure 8 repeats the exercise of Figure 5 – varying the fundamental difference between the two types – but in the richer model with moral hazard.

Figure 8(a) plots the share of each type exerting effort as a function of β_L , with β_H fixed. Steeper discounting means that L -types value the future gain in income less, and so they exert less effort. This further implies that the marginal benefit of exerting effort increases for H -types, since now a realization of $y_2 = y_h$ improves one's reputation. This drives up H -types' effort rate. This effect strengthens as the types separate further until the L -types become so impatient that the H -types separate more based on the borrowing behavior in panel (b), which is exactly consistent with the baseline model behavior in Figure 5(a).

Another way of understanding the role of moral hazard is to investigate how key aspects of the equilibrium change as we vary the strength of the friction. To this end, Figure 9 considers equilibria across a range of the effort costs κ . Each panel plots the average of the indicated variable (green, right axis), as well as the difference between the H - and L -types (black, left axis). Panel (a) shows




Figure 8: **The interaction between moral hazard and adverse selection**

Notes: This figure plots the rate at which effort is exerted (panel (a)) and the share of agents choosing to borrow (panel (b)) by type in the moral hazard model over a range of β_L values. The vertical line in each graph corresponds to the benchmark parameter value for β_L used throughout the numerical illustrations.

that while effort declines on average when its cost increases, the difference in effort between the types exhibits a hump-shaped pattern: H -types first maintain their effort levels more than L -types, but then cut effort when the cost increases too much. Income at age 2 (panel (b)) follows this exact pattern given our assumption about the distribution of income conditional on effort.

Panels (c) through (e) show the impacts of the increase in effort cost on savings and consumption patterns. The non-monotone patterns here suggest an interplay between the signaling value of exerting effort and that of saving. When effort is cheap, H -types improve their reputations by working harder rather than borrowing less. Low effort cost acts like a wealth effect: agents are effectively “richer” due to an increase in the value of their labor endowments, and so they want to borrow, reducing the signaling value coming from the a_2 choice. Therefore, borrowing actually increases for the initial increases in κ . As effort becomes more expensive and effort levels drop, though, this effect reverses: households no longer borrow against high future income, but instead save for the future to bolster their reputations. This effect is especially pronounced for the H -types, as in the baseline model.

These savings patterns from panel (c) are mirrored for age 1 consumption in panel (d): initially, age 1 consumption increases and the gap between types closes as the effort channel dominates, but thereafter the savings channel dominates, recovering the trends from the baseline model. Panel (e) reflects these same forces for age 2 consumption, but with the change that high effort levels for low κ increase income and therefore drive up consumption. Combining all these insights delivers the non-monotone pattern in type score dispersion presented in panel (f).




Figure 9: Varying the strength of the moral hazard friction

Notes: This figure shows a range of outcomes across a range of levels of the effort cost, κ , in the version of the model with moral hazard. Each panel plots the average (combining high and low types) of the variable indicated in the title on the right axis (green), as well as the difference between H - and L -types on the left axis (black). The vertical line in each graph corresponds to the benchmark parameter value for κ_e used throughout the numerical illustrations.

4.4.2 Life cycle and reputational incentives with moral hazard

The left side of Table 2 presents key life cycle metrics for the moral hazard model from this section. It reports the same moments as Table 1 for the baseline model from Section 3, with two exceptions.²⁰ First, Panel A adds in metrics on age 2 income, since this is now endogenous given the effort choice at age 1. Second, Panel B reports statistics on the exertion of effort in aggregate and by type. The moments reported in Table 1 follow the same basic patterns in this extended model: H -types save more than L -types on average, and as a result consume more than L -types at age 2 but less at age 1. This is driven by both signaling and price effects: a lower share of H -types borrow and default, facing the same price schedules and reputational incentives.

The key novelty in this version of the model documented in Table 2 is *type-specific income inequality*: on average, the income of H -types is about 15% higher than that of L -types. This comes entirely through the choice of effort: H -types exert effort at a rate 51.5% higher than L -types, bolstered simultaneously by patience and the incentive to acquire reputation. In the baseline model, where income was exogenous, there was no scope for dispersion in income across the types, and therefore no ability for income itself to be informative about reputation. In this sense, this version of the model delivers a mechanism that could be broadly consistent with the patterns documented in Section 2.

How important is the incentive to acquire reputation for this channel? To understand this, Table 2 compares the basic moral hazard model to the “no tracking” version, like we considered for the baseline in Section 3.5.3. The results for wealth and consumption broadly mirror those from Table

²⁰Note that the analog of Figure 7 can be found in <https://zenodo.org/records/15731380>.

	moral hazard benchmark			moral hazard, no tracking		
	mean	diff, $H - L$	std. dev.	mean	diff, $H - L$	std. dev.
A. Inequality moments						
age 2 income	0.91	0.10	0.23	0.87	0.07	0.22
age 2 wealth	-0.02	0.05	0.10	-0.06	0.01	0.09
age 1 consumption	0.74	-0.03	0.07	0.75	-0.01	0.05
age 2 consumption	0.92	0.13	0.24	0.87	0.07	0.21
B. Driving behavior						
effort rate (%)	54.9	51.5		36.9	33.8	
fraction in debt (%)	65.7	-20.6		80.9	-5.3	
default rate (%)	34.4	-22.6		50.5	-8.9	

Table 2: **Life cycle and the role of information in the moral hazard model**

Notes: This table reports the statistic in the column header for the moment in the row header for the model version indicated in the top row of the table. The mean is the population average, and the standard deviation is the cross-sectional standard deviation. The difference reported is the difference for the average of the moment for H -types less the average of the moment for L -types.

1: when there is incentive to acquire reputation, wealth and consumption inequality increase in the aggregate and across types. What is novel here is that the same is also true for *income* inequality. This is driven by changes in effort: in addition to increasing on average, effort tilts sharply towards H -types when we track reputations. This is because they value it not only for the increase in future income, but for the boost in reputation that higher future income may offer. Notably, tracking reputations also promotes credit access in this setting, lowering interest rates on average.

5 Regulating Information Sets: Adding Medical Expenses

There is survey evidence in Fulford and Low (2024) that unexpected expense shocks (primarily medical and auto) are cited as major reasons for delinquency. Research by the CFPB also finds that: (i) medical debts have little predictive value about borrowers' ability to repay other debts; and (ii) consumers frequently report receiving inaccurate bills or being asked to pay bills that should have been covered by insurance or financial assistance programs.²¹ These findings provided a rationale for a CFPB ruling that went into effect on January 7, 2025, which banned the inclusion of medical bills on credit reports used by lenders and prohibited lenders from using medical information in their decisions. In this section, we use our model as a laboratory to examine the implications of regulating such adverse information out of an individual's credit record. To do so, we add medical expense shocks and the possibility of going delinquent to the benchmark model of Section 3.

²¹For details see https://files.consumerfinance.gov/f/documents/cfpb_med-debt-final-rule_2025-01.pdf.

5.1 Environment

In order to isolate the economic mechanisms, we return to the Section 3.1 environment with exogenous income but add exogenous medical expenses at age 1. Agents may go delinquent on those medical expenses, transferring the cost to age 2. Since individuals discount the future at different rates, the delinquency decision may provide a signal of an individual's unobservable type.

Specifically, at the beginning of age 1, individuals realize a medical expense shock $m_1 \in \{0, \bar{m}\}$ with probability $Q_m(m_1) \equiv \mathbb{P}(m_1)$. We assume these shocks are independent of type: therefore although the shocks themselves are bad luck, the delinquency decision *conditional on a shock* may signal one's type. The delinquency decision $\delta_1 \in \{0, 1\}$ is subject to an extreme value shock ε^{δ_1} so that there is partial pooling of types within the set of medical delinquents. If $m_1 = 0$, there is nothing on which to go delinquent, so $\delta_1 = 0$. Since delinquency delays the possibility of repayment until age 2, we denote the individual's medical debt at age 2 by $b_2 = m_1 \delta_1$. We assume that $y_\ell > |\underline{a}| + \bar{m}$ so that it is always budget feasible to pay back debts (medical and otherwise) at age 2.

Let (m_1, δ_1) denote an individual's "medical record," which comprises both the expense itself and whether or not the individual went delinquent. Towards analyzing the role of the proposed regulation, we consider 2 versions of the model: one in which medical records are observable, and one in which they are not. We denote by μ_1 the set of publicly observable medical outcomes at age 1: when medical records are observable, $\mu_1 = (m_1, \delta_1)$, and when they are not, $\mu_1 = \emptyset$ and there is no conditioning on medical events.

Following the medical expense shock and delinquency decision, individuals can choose to borrow $a_2 < 0$ or save $a_2 \geq 0$. The discount price q_1 can be made contingent on an individual's medical record if it is observable: we denote the price schedule by $q_1^{a_2}(\mu_1)$ (which includes the case in which medical records are not observable and $\mu = \emptyset$). Similarly, after the asset choice, lenders update their assessment of the individual's type based on the asset choice and the observable component of the medical record: that is, and agent's type score evolves according to

$$s_2 = \Gamma_1^{a_2}(\mu_1) \equiv \mathbb{P}(\tau = H \mid a_2, \mu_1). \quad (27)$$

Equation (27) is the analog of (5) from the baseline model and (22) from the moral hazard model.

To summarize our information assumptions for this environment, we retain the assumption that an individual's type (τ) is unobservable (i.e., adverse selection) but return to the case of exogenous observable income (y_n). Medical shocks (m_1) are observable as are credit market actions (δ_1, a_2, d_2). However, all actions (δ_1, a_2, d_2) are subject to unobservable idiosyncratic shocks ($\varepsilon^{\delta_1}, \varepsilon^{a_2}, \varepsilon^{d_2}$). As before, type scores (s_n) are observable. To summarize the new timing:

1. At $n = 1$,
 - (a) each individual's medical expense shock m_1 is realized.
 - (b) if $m_1 = \bar{m}$, the individual makes a delinquency decision δ_1 .
 - (c) each individual chooses a_2 taking the price schedule $q_1^{a_2}(\mu_1)$ as given.

- (d) observing a_2 and μ_1 , assessments of an individual's type $s_2 = \Gamma_1^{a_2}(\mu_1)$ are updated according to Bayes' Law.

2. At $n = 2$,

- (a) y_2 is realized from $Q_y(y_2)$, determining observable state $\omega_2 = (a_2, b_2, y_2, s_2)$.
(b) if $a_2 - b_2 < 0$, the individual makes a default choice d_2 .
(c) observing d_2 , individual type score is updated according to Bayes Law: $s_3 = \psi_2^{d_2}(\omega_2)$ if $a(\omega_2) - b(\omega_2) < 0$ or $s_3 = s_2$ otherwise.

Thus, the only differences from the timing in Section 3.1 are the addition of 1(a) and 1(b). The only other notable difference is conditioning prices on medical records.

5.2 Equilibrium

As in the baseline model, in this version of the model we associate extreme value preference shocks to all the relevant decisions: medical delinquency, asset choice, and default. At age 1, the ex ante value of the individual sums over the likelihood of medical expense shocks:

$$V_1(\tau) = \sum_{m_1} Q_m(m_1) V_{\Delta,1}(m, \tau), \quad (28)$$

where $V_{\Delta,1}(m, \tau)$ is the value function at the medical delinquency choice stage, which reflects that the individual must make a delinquency decision if hit by a medical expense shock:

$$V_{\Delta,1}(m_1, \tau) = \begin{cases} V_{A,1}^0(m_1, \tau) & \text{if } m_1 = 0 \\ \mathbb{E}_{\varepsilon^{\delta_1}} \left[\max_{\delta_1 \in \{0,1\}} \left\{ V_{A,1}^{\delta_1}(\bar{m}, \tau) + \varepsilon^{\delta_1} \right\} \right] & \text{otherwise} \end{cases} \quad (29)$$

and $V_{A,1}^{\delta_1}(m_1, \tau)$ is the value function at the asset choice stage

$$V_{A,1}^{\delta_1}(m_1, \tau) = \mathbb{E}_{\varepsilon^{a_2}} \left[\max_{a_2} \left\{ v_{A,1}^{(a_2, \delta_1)}(m_1, \tau) + \varepsilon^{a_2} \right\} \right], \quad (30)$$

$$\begin{aligned} \text{where } v_{A,1}^{(a_2, \delta_1)}(m_1, \tau) &= U(y_\ell - (1 - \delta_1)m_1 - q_1^{a_2}(\mu(m_1, \delta_1))a_2) \\ &\quad + \beta_\tau \sum_{y_2} Q_y(y_2) V_2(a_2, m_1 \delta_1, y_2, \Gamma_1^{a_2}(\mu(m_1, \delta_1)), \tau). \end{aligned} \quad (31)$$

Problem (29) generates the type-specific decision rule over medical delinquency $\sigma_{\Delta,1}^{\delta_1}(m, \tau)$, while (30) generates the medical-record- and type-specific decision rule over assets $\sigma_{A,1}^{(a_2, \delta_1)}(m_1, \tau)$. If $m_1 = 0$, we adapt the convention that $\sigma_{\Delta,1}^{\delta_1=1}(0, \tau) = 1$ for all τ since there is no notion of delinquency. We use the notation $\mu_1 = \mu(\delta_1, m_1)$ to denote the mapping from medical events to observable medical records: this evaluates to (m_1, δ_1) when medical records are observable and $\emptyset$ for all (m_1, δ_1) when they are not. Second, we incorporate the law of motion for medical debt $b_2 = m_1 \delta_1$.

At age 2, the individual makes a default decision if $a_2 - b_2 < 0$:

$$V_2(\omega_2, \tau) = \mathbb{E}_{\varepsilon^{d_2}} \left[\max_{d_2 \in \{0,1\}} \left\{ v_2^{d_2}(\omega_2, \tau) + \varepsilon^{d_2} \right\} \right], \quad (32)$$

where $v_2^{d_2}(\omega_2, \tau) = U(y(\omega_2) + (1 - d_2)(a(\omega_2) - b(\omega_2))) + \beta_\tau \phi(\psi_2^{d_2}(\omega_2))$.

The decision rule associated with (32) is $\sigma_2^{d_2}(\omega_2, \tau)$. The only difference between this problem and the default choices in the models in Sections 3 and 4 is the presence of the medical debt term, b_2 .

Observable medical records Analogously to the moral hazard case, define the ex-ante decision rule over (a_2, δ_1) given the medical shock m_1 by $\sigma_1^{(a_2, \delta_1)}(m_1, \tau) = \sigma_{A,1}^{(a_2, \delta_1)}(m_1, \tau) \sigma_{\Delta,1}^{\delta_1}(m_1, \tau)$. When medical records are observable, the type score exactly mirrors the baseline in (17), with the modification that the inference is medical-record-specific:²²

$$\Gamma_1^{a_2}(\mu_1) = \frac{\rho \sigma_1^{(a_2, \delta(\mu_1))}(m(\mu_1), H)}{\rho \sigma_1^{(a_2, \delta(\mu_1))}(m(\mu_1), H) + (1 - \rho) \sigma_1^{(a_2, \delta(\mu_1))}(m(\mu_1), L)}. \quad (33)$$

Similarly, the price schedule directly mirrors the one from the baseline model:

$$q_1^{a_2}(\mu_1) = \frac{1}{1+r} [\Gamma_1^{a_2}(\mu_1) \bar{p}^{a_2}(\mu_1; H) + (1 - \Gamma_1^{a_2}(\mu_1)) \bar{p}^{a_2}(\mu_1; L)], \quad (34)$$

$$\text{where } \bar{p}^{a_2}(\mu_1; \tau) = \sum_{y_2} Q_y(y_2) \sigma^{d_2=0}(a_2, b(\mu_1), y_2, \Gamma_1^{a_2}(\mu_1), \tau). \quad (35)$$

The only difference between (34) and the baseline price (18) is that the type update that is relevant now depends on the observed medical record. Similarly, the only differences between (35) and the baseline repayment probability (19) are the type score update and the presence of the medical debts $b(\mu_1) = m(\mu_1) \delta(\mu_1)$ in tomorrow's default decision.

Unobservable medical events When medical records are unobservable ($\mu = \emptyset$), it is useful to define the τ -specific probability of choosing a_2 analogous to the moral hazard case:

$$\bar{\sigma}^{a_2}(\tau) = \sum_{m_1} Q_m(m_1) \sum_{\delta_1} \sigma_1^{(a_2, \delta_1)}(m_1, \tau).$$

This includes summation over both the realization of the medical expense shock and the likelihood of delinquency given the medical expense shock. Since neither is observed, lenders may only condition on a_2 . Therefore, the updated type score at the beginning of age 2 is simply $\Gamma_1^{a_2} = \psi_1^{a_2}$ from equation (17) for the baseline model, with the modification that $\sigma_1^{a_2}$ is replaced by $\bar{\sigma}_1^{a_2}$ as in the moral hazard model.

Towards determining equilibrium loan prices, another useful object to define is the type-specific

²²Note that this derivation relies on the assumption that medical expense shocks are independent of type; a more general formulation would have to correct for this.

likelihood of a specific medical record μ given an observation of a_2 :

$$\tilde{Q}_\mu^{a_2}(\mu_1; \tau) \equiv \mathbb{P}(\mu_1 \mid a_2, \tau) = Q_m(m(\mu_1)) \frac{\sigma_1^{(a_2, \delta(\mu_1))}(m(\mu_1), \tau)}{\bar{\sigma}_1^{a_2}(\tau)}.$$

This object serves a similar purpose to the $\tilde{Q}_y^{a_2}$ object in the moral hazard model by updating the likelihood of an unobservable individual-level outcome based on an observable one. We can write the loan price function for this case as

$$q_1^{a_2} = \frac{1}{1+r} [\Gamma_1^{a_2} \bar{p}^{a_2}(H) + (1 - \Gamma_1^{a_2}) \bar{p}^{a_2}(H)], \quad (36)$$

$$\text{where } \bar{p}^{a_2}(\tau) = \sum_{y_2} Q(y_2) \sum_{\mu_1} \tilde{Q}_\mu^{a_2}(\mu_1; \tau) \sigma_2^{d_2=0}(a_2, b(\mu_1), y_2, \Gamma_1^{a_2}, \tau). \quad (37)$$

The pricing function (36) is identical to the baseline (18) once we account for the fact that the type score update is the same. The main difference between (36) and its analog from the observable medical record case (34) is that the repayment probability in (37) sums over the (adjusted) likelihood of each possible medical record, while the analog (35) is medical-record-specific.

5.3 Reputational incentives, observability, and medical delinquency

How do reputational incentives shape the medical delinquency decision? To see this, define the fundamental value of going delinquent conditional on realizing the medical expense shock in the economy in which medical records are observable as

$$v_1^{(a_2, \delta_1)}(\tau) = U(y_\ell - (1 - \delta_1)m_1 - q_1^{a_2}(\mu(m_1, \delta_1))a_2) + \beta_\tau \sum_{y_2} Q_y(y_2) V_2(a_2, m_1 \delta_1, y_2, \Gamma_1^{a_2}(\mu_1(m_1, \delta_1)), \tau).$$

For intuition, suppose δ_1 is a continuous choice. Then, the net marginal benefit of delinquency is

$$U'(c_1) \left(m_1 - \frac{\partial q_1^{a_2}(\mu_1)}{\partial \delta_1} \right) + \beta_\tau \sum_{y_2} Q_y(y_2) \left[\frac{\partial V_2}{\partial b_2} m_1 + \frac{\partial V_2}{\partial s_2} \frac{\partial \Gamma_1^{a_2}(\mu_1)}{\partial \delta_1} \right].$$

Reputational incentives come in two dimensions. First, going delinquent may affect the prices an individual faces on loans today, reflected in the $\frac{\partial q_1}{\partial \delta_1}$ term. This arises because going delinquent decreases the borrower's "effective" net assets tomorrow, $a_2 - b_2$ and may also revise down the assessment of the likelihood that the individual is of type H . Second, to the extent that reputation is valued, this latter effect directly affects the continuation value to the extent that going delinquent is more likely among L types (i.e. $\frac{\partial \Gamma_1}{\partial \delta_1} < 0$). Both of these effects are absent when medical records cannot be observed.




Figure 10: **Credit prices in the medical expense shock model**

Notes: This figure plots the loan price schedules across debt levels in the medical expense shock model. The first three lines correspond to the case in which medical records are observable, which has 3 distinct cases, one for each possible realization of μ_1 . The last line is for the unobservable medical record model, in which there is only one price schedule in equilibrium.

5.4 What happens when medical records are unobservable?

Our goal is to understand the equilibrium effects of making medical records unobservable. Therefore, in this section, we compare the observable and unobservable medical record cases described above. The main effect of precluding lenders from considering medical records is pooling in the credit market: lenders know that some share of the borrowers seeking a given a_2 actually have a “hidden debt” coming from medical delinquency. Figure 10 shows how this pooling affects credit prices by plotting price schedules from both cases.

The first three lines in the graph correspond to the case when medical records are observable. Prices here follow a predictable pattern: when a borrower gets no medical expense shock, or does get the medical expense shock but chooses not to go delinquent, he gets more favorable terms. By contrast, credit terms worsen materially when a medical delinquency is observed. The explanation is simple: borrowers have lower effective leverage for a given a_2 choice in the former two cases as opposed to the latter case. When medical records are unobservable, of course, lenders are forced to pool borrowers of each type, and in equilibrium prices are less favorable for borrowers with no medical debt and more favorable for borrowers with medical debt.

Table 3 examines how these information partitions affect aggregate and distributional outcomes. The outcomes are intuitive: since lenders are unable to ration credit specifically for higher-levered borrowers with medical debt, the terms of credit for these borrowers improve. This eliminates a powerful incentive to avoid medical delinquency, and as a result the rate of medical delinquency increases sharply, particularly for the relatively impatient L -types. Since the only way to accrue medical debt in this model (conditional on being hit by the expense shock) is to go delinquent, Panel B shows that this behavior is driven entirely by changes in medical delinquency behavior.

The rest of Table 3 shows that there is very little change to other aggregate moments and

	observable medical			unobservable medical		
	mean	diff, $H - L$	std. dev.	mean	diff, $H - L$	std. dev.
A. Inequality moments						
share with medical debt (%)	0.5	-1.4		12.0	-2.5	
age 2 mean wealth	0.02	0.10	0.10	0.02	0.10	0.11
age 1 mean consumption	0.69	-0.06	0.08	0.69	-0.06	0.08
age 2 mean consumption	0.84	0.06	0.17	0.84	0.05	0.17
B. Driving behavior						
medical delinquency rate (%)	1.3	-4.0		35.2	-7.3	
fraction in debt (%)	44.5	-41.9		44.7	-41.6	
default rate (%)	22.8	-32.2		24.5	-32.9	

Table 3: The role of observability of medical events

Notes: This table reports the statistic in the column header for the moment in the row header for the model version indicated in the top row of the table. The mean is the population average, and the standard deviation is the cross-sectional standard deviation. The difference reported is the difference for the average of the moment for H -types less the average of the moment for L -types.

measures of inequality examined elsewhere in the paper: the primary effect of this regulation is borne out in medical payments. This makes sense given that (uninsured) medical expenses, while an important motivator for self-insurance and credit usage at the individual level, do not comprise a large share of overall expenditures in the economy.

6 N -Period Model with Credit Scores

Can the forces described in the preceding sections help us understand the empirical patterns documented in Section 2? To address this question, here we expand on the moral hazard model from Section 4 in four ways. First, we extend the model to have $N > 2$ periods of life. This allows us to generate realistic life cycle profiles as in Figure 1: we set $N \geq 8$ and have each “age” correspond to 5 years, consistent with the age bins in that figure. In practice, we set $N = 12$ to be consistent with a life span of 60 years (age 20 – 80) and so that the last period we compare with the data ($N = 8$, or age bin 55 – 60) is not the terminal period in the model.

Second, we introduce “churn” in individuals’ types as in (Chatterjee et al., 2023) to help match rising mean credit rankings over age bins in the data. Specifically, we introduce a Markov process over type at the individual level.²³ This allows for the composition of types to differ across age bins, reflecting the possibility of compositional changes in unobservable type alluded to in Section 2.1.

Third, although the model is still expressed in terms of *type scores*, for comparability with the data we also measure reputation using model-implied *credit rankings*. We define these credit rankings formally in Section 6.3.1 after laying out the environment and the equilibrium, but the

²³Obviously, we nest the case in Sections 3 – 5 if the Markov process is an identity matrix.

idea is that a credit score is a measure of the likelihood an individual experiences an adverse credit event over a specified time horizon, and a credit ranking is an ordinal ranking of this measure over the population. Note that comparing the dynamics of these measures with the dynamics we observe in the data requires at least two periods in which there is a probability of default in the next period, another benefit of considering $N > 2$. The credit score is a murkier measure than the type score, since the likelihood of default depends not only on the fundamentally unobservable type, but also income and wealth. Chatterjee et al. (2023) show that under certain conditions, equilibria like those described in the previous sections can be implemented with lenders using the model equivalent of a credit ranking in place of the type score.

Fourth, since individuals value future credit access, which depends on reputation, we drop the simplifying assumption that individuals directly value terminal reputation via $\phi(s_{N+1})$: here we set $\phi(s_{N+1}) = 0$. This eases welfare analysis across information structures, since this reputational utility bonus can be present in a model with reputation (like our baseline) but not in one without it (like our full information counterfactual, where the absence of hidden information obviates the need for type scores). Instead, the value of reputation works through equilibrium credit prices, which are determined endogenously in each model we consider, regardless of the information structure.²⁴

6.1 Environment

We consider a model with both adverse selection and moral hazard as in Section 4, but no medical expenses as in Section 5.²⁵ As discussed above, we let $\tau_i \in \{H, L\}$ evolve as a Markov process $Q_\tau(\tau_{i+1}; \tau_i)$, drawn independently across individuals. We assume that the fraction of high types at age 1, ρ_1 , is lower than the ergodic share of high types implied by $Q_\tau, \bar{\rho}$, so that the share of H types increases with age. Individuals begin age 1 with state (ω_1, τ_1) drawn from a distribution $\lambda_1(\omega, \tau)$.²⁶ At the terminal age N , individuals cannot borrow. Since there are multiple periods of asset choice and default unlike the previous subsections, we assume as in Chatterjee et al. (2023) that an individual who defaults is excluded from the credit market in the current period.

To summarize the information structure, an individual's type (τ) and effort choice (e) are unobservable. Income (y), wealth (a), and credit market actions (a' and d) are observable. All actions are subject to unobservable idiosyncratic shocks ($\varepsilon^e, \varepsilon^d, \varepsilon^{a'}$). Finally, type scores (s) are observable, summarizing the assessment of unobservable type through the noisy signals observable choices and state realizations. To summarize the timing, at each age $n = 1, \dots, N - 1$:

1. Individuals begin with unobservable type τ_n and observable state $\omega_n = (a_n, y_n, s_n)$.
2. Individuals choose unobservable effort e_n .

²⁴Indeed, in the fully dynamic model we lay out below, it's possible to show that the dynamic marginal value of reputation – $\frac{\partial V}{\partial s}$ from our simpler models, e.g. Section 3.4 – is a (fairly complex) discounted sum of the elasticities of equilibrium prices with respect to reputation, adjusted by the elasticity of reputation with respect to actions and the likelihood of borrowing.

²⁵It is straightforward to extend our environment to include these expenses; we omit them for ease of exposition.

²⁶Note that the initial distribution assumed in the previous sections is a special case that is degenerate on $\omega_1 = (a_1, y_1, s_1) = (0, y_\ell, \rho_1)$ with the fraction of H -types equal to ρ_1 , which was equal to $\bar{\rho}$.

3. Individuals make default choice d_n on debt a_n .
4. If $d_n = 0$, then individuals make an asset choice a_{n+1} at price $q_n^{a_{n+1}}(\omega_n)$; otherwise $a_{n+1} = 0$ by assumption.
5. Individuals learn their next-period income y_{n+1} , drawn from Q_y^e and their unobservable type τ_{n+1} drawn from Q_τ .
6. Individuals' type assessments are updated according to Bayes' Law:

$$s_{n+1} = \Upsilon_n^{(a_{n+1}, d_n)}(y_{n+1}, \omega_n) \equiv \mathbb{P}(\tau_{n+1} = H \mid a_{n+1}, d_n, y_{n+1}, \omega_n). \quad (38)$$

7. Individuals leave period n in observable state $\omega_{n+1} = (a_{n+1}, y_{n+1}, s_{n+1})$.²⁷

6.2 Equilibrium

Ages $n = 1, \dots, N - 1$ in this general model are similar to the initial age in the two-period models in the earlier sections, and the terminal period $n = N$ is similar to the terminal age $n = 2$ from these models. Therefore, here we only lay out the problem for the interim periods $n \in \{1, \dots, N - 1\}$ and relegate much to Appendix A.3.2. In what follows, we use recursive notation where $x_n = x$ and $x_{n+1} = x'$. An agent of (observable) age n in observable state $\omega = (a, y, s)$ makes effort, default, and asset choices consistent with the timeline laid out above. While we model these decisions as sequential, we assume that the timing of each decision is such that the agent's assessed type (score) is only adjusted after all the observable actions of the period.²⁸

Value functions and decision rules For any age $n < N$, the individual first makes an effort choice, weighing the cost of effort against the continuation value associated with exerting that effort:

$$V_{E,n}(\omega, \tau) = \mathbb{E}_{\varepsilon^e} \left[\max_{e \in \{0,1\}} \left\{ v_{E,n}^e(\omega, \tau) + \varepsilon^e \right\} \right] \text{ where } v_{E,n}^e(\omega, \tau) = -\kappa e + V_{D,n}^e(\omega, \tau), \quad (39)$$

where the continuation value at the default stage is:

$$V_{D,n}^e(\omega, \tau) = \begin{cases} \mathbb{E}_{\varepsilon^d} \left[\max_{d \in \{0,1\}} \left\{ v_{D,n}^{(d,e)}(\omega, \tau) + \varepsilon^d \right\} \right] & \text{if } a(\omega) < 0 \\ V_{A,n}^e(\omega, \tau) & \text{if } a(\omega) \geq 0. \end{cases} \quad (40)$$

The action-specific value of default is

$$v_{D,n}^{(d=1,e)}(\omega, \tau) = U(y(\omega)) + \beta_\tau \sum_{\tau'} Q_\tau(\tau'; \tau) \sum_{y'} Q_{y,n}^e(y') V_{E,n+1} \left(0, y', \Upsilon_n^{(0,1)}(y'; \omega), \tau' \right) \quad (41)$$

²⁷Note that the terminal condition must be $\omega_{N+1} = (0, 0, s_{N+1})$.

²⁸We model these decisions as sequential, but this can of course be done with nesting. Our formulation easily allows for different degrees of noise in each decision.

and the action-specific value of no default is $v_{D,n}^{(d=0,e)}(\omega, \tau) = V_{A,n}^e(\omega, \tau)$, where

$$V_{A,n}^e(\omega, \tau) = \mathbb{E}_{\varepsilon^{a'}} \left[\max_{a' \in \mathcal{F}_n(\omega)} \left\{ v_{A,n}^{(a',e)}(\omega, \tau) + \varepsilon^{a'} \right\} \right] \quad (42)$$

$$v_{A,n}^{(a',e)}(\omega, \tau) = U \left(c_n^{a'}(\omega) \right) + \beta_\tau \sum_{\tau'} Q_\tau(\tau'; \tau) \sum_{y'} Q_{y,n}^e(y') V_{E,n+1} \left(a', y', \Upsilon_n^{(a',0)}(y'; \omega), \tau' \right), \quad (43)$$

where consumption satisfies the budget constraint

$$c_n^{a'}(\omega) = a(\omega) + y(\omega) - q_n^{a'}(\omega) a' \quad (44)$$

and the feasible set is $\mathcal{F}_n(\omega) \equiv \{a' \mid c_n^{a'}(\omega) > 0\}$.

The first decision problem generates the type-specific decision rule over effort $\sigma_{E,n}^e(\omega, \tau)$. The second generates the effort- and type-specific decision rule over default, $\sigma_{D,n}^{(d,e)}(\omega, \tau)$. Of course, there is no weight on default if $a(\omega) \geq 0$. The third generates the effort- and type-specific decision rule over assets $\sigma_{A,n}^{(a',e)}(\omega, \tau)$. Note that an asset choice decision is only made if $d = 0$, and so this is implicitly conditional on $d = 0$. Discounting occurs intertemporally.

At age N , the only decision is whether to default if $a < 0$:

$$V_N(\omega, \tau) = \mathbb{E}_{\varepsilon^d} \left[\max_{d \in \{0,1\}} \left\{ U(y(\omega)) + \varepsilon^{d=1}, U(c_N^0(\omega)) + \varepsilon^{d=0} \right\} \right]. \quad (45)$$

The decision rule associated with this choice problem is $\sigma_N^d(\omega, \tau)$.

Loan price and type score updating functions Type scores and loan prices are computed similarly to our 2-period economy from Section 4:

$$\Upsilon_n^{(a',d)}(y'; \omega) = \mathbb{P}(\tau' = H \mid a', d, y', \omega, n) \quad (46)$$

$$q_n^{a'}(\omega) = \frac{1}{1+r} \mathbb{P}(d' = 0 \mid a', \omega, n). \quad (47)$$

To solve for these equilibrium functions, it is useful to define several preliminary objects. These objects are generalizations of the ones described in Section 4.2. The age- j joint decision rule over all decisions (effort, default, and assets) is

$$\sigma_n^{(a',d,e)}(\omega, \tau) = \sigma_{E,n}^e(\omega, \tau) \sigma_{D,n}^{(d,e)}(\omega, \tau) \sigma_{A,n}^{(a',e)}(\omega, \tau), \quad (48)$$

where it is understood that $\sigma_{D,n}^{(1,e)}(\omega, \tau)$ evaluates to zero if $a(\omega) \geq 0$ and that $\sigma_{A,n}^{(0,e)}(\omega, \tau) = 1$ if $d = 1$ at the default stage. The effort-weighted decision rule over the observable decision (a', d) is

$$\bar{\sigma}_n^{(a',d)}(\omega, \tau) = \sum_e \sigma_n^{(a',d,e)}(\omega, \tau). \quad (49)$$

The probability of a given y' realization based on type and observable choice is

$$\tilde{Q}_{y,n}^{(a',d)}(\omega, \tau) \equiv \mathbb{P}(y' \mid a', d, \omega, \tau, n) = \sum_e Q_{y,n}^e(y') \frac{\sigma_n^{(a',d,e)}(\omega, \tau)}{\bar{\sigma}_n^{(a',d)}(\omega, \tau)}. \quad (50)$$

Lastly, the “interim” reputation update of *today’s* type based only on this period’s observable actions is useful as an input to loan prices and the full type score update. The assessment made *after* the choice of (a', d) but *before* the realization of y' is:

$$\psi_n^{(a',d)}(\tau; \omega) \equiv \mathbb{P}(\tau \mid a', d, \omega, n) = \frac{s(\tau; \omega) \bar{\sigma}_n^{(a',d)}(\omega, \tau)}{\sum_{\tilde{\tau}} s(\tilde{\tau}; \omega) \bar{\sigma}_n^{(a',d)}(\omega, \tilde{\tau})},$$

where we have slightly extended the existing notation to be defined for either type τ (e.g. $s(H; \omega) = s(\omega)$ and $s(L; \omega) = 1 - s(\omega)$ in our earlier notation), that is useful in the calculations below which account for churn in types.

With repeated application of Bayes’ Law, we can write the likelihood that an individual is of type τ today *and* earns income y' tomorrow based on her observable states and actions as

$$v_n^{(a',d)}(y', \tau; \omega) \equiv \mathbb{P}(y', \tau \mid a', d, \omega, n) = \psi_n^{(a',d)}(\tau; \omega) \tilde{Q}_{y,n}^{(a',d)}(\omega, \tau).$$

Then, to get to the probability that the agent will be of type H *tomorrow*, we need only weight across the likelihood of each current type today and account for the type transition probability between today and tomorrow:

$$\Upsilon_n^{(a',d)}(y'; \omega) = \sum_{\tau} Q_{\tau}(H; \tau) \frac{v_n^{(a',d)}(y', \tau; \omega)}{\sum_{\tilde{\tau}} v_n^{(a',d)}(y', \tilde{\tau}; \omega)}. \quad (51)$$

Using the same analysis, we can compute loan prices as

$$q_n^{a'}(\omega) = \frac{1}{1+r} \sum_{\tau', y', \tau} Q_{\tau}(\tau'; \tau) v_n^{(a',d)}(y', \tau; \omega) \left(1 - \bar{\sigma}_{n+1}^{(0,1)} \left(a', y', \Upsilon_n^{(0,a')}(y'; \omega), \tau' \right) \right). \quad (52)$$

Equilibrium distribution Given the block recursive structure of our model, the equilibrium distribution is not necessary to solve for the equilibrium pricing functions above. However, in order to compute model moments, we must solve forward for the equilibrium distribution of agents at each age, $\lambda_n(\omega, \tau)$, using the individual law of motion implied by the total decision rule $\sigma_{n-1}^{(a',d,e)}(\omega, \tau)$, the earnings probability function $Q_{y,n}^e(y')$, and the type scoring function $\Upsilon_{n-1}^{(a',d)}(y'; \omega)$, as well as the previous period’s distribution, $\lambda_{n-1}(\omega, \tau)$. For $n = 2, \dots, N$, this transition operator is

$$\lambda_n(a', y', s', \tau') = \sum_{\omega, \tau} Q_{\tau}(\tau'; \tau) Q_{y,n}^e(y') \mathbf{1} \left[s' = \Upsilon_{n-1}^{(a',d)}(y'; \omega) \right] \sigma_{n-1}^{(a',d,e)}(\omega, \tau) \lambda_{n-1}(\omega, \tau). \quad (53)$$

The first two terms account for exogenous transitions over y and τ , the third imposes that type scores evolve according to the equilibrium type score function (51), the fourth accounts for individual decision rules, and the fifth weights across the mass of individuals in the indicated state in the prior period. Note that the distribution at age 1, $\lambda_1(\omega, \tau)$ is exogenously specified.

6.3 Mapping the model to data

6.3.1 Defining credit scores and credit rankings in the model

In our model, assessments of an individual's type are encapsulated in a type score s that evolves based on observable credit market actions (d, a') and observable state variables $(\omega = (a, y, s))$. Our empirical measure of reputation (e.g., Figure 1), however, is the credit score, which measures creditworthiness as the likelihood of an adverse credit event over a given time horizon. To formalize this notion in the model, we define an “adverse credit event” as a default and use a horizon of one period. In our context, we take the probability of not encountering an adverse event at the beginning of age n as the model analog of a *credit score* that lies in $[0, 1]$ not $[300, 850]$:

$$\tilde{\chi}_n(\omega) \equiv \mathbb{P}(d' = 0 \mid \omega, n). \quad (54)$$

Note that this object is effectively the choice-weighted probability of repayment and is therefore closely related to the loan price.²⁹ How does this object relate to one's type score? In equilibrium, default is more common among individuals with more debt and lower income. Since L types are more likely to be in debt and have lower income, they tend to have lower credit scores. Furthermore, conditional on observable state ω , L types are more likely to borrow and default, and so a lower type score also lowers one's credit score. To summarize, then, the correlations between credit score $\tilde{\chi}_n(\omega)$ and type score s , income y , and wealth a are all generically positive.

To take our model to the data, we construct a *credit ranking* by taking the probability $\tilde{\chi}$ from equation (54) and determining the fraction of individuals with a credit score less than or equal to $\tilde{\chi}$ across the equilibrium distribution of individuals over observable states:

$$\chi_n(\omega) \equiv \sum_{\{\hat{n}, \hat{\omega} \mid \tilde{\chi}_{\hat{n}}(\hat{\omega}) \leq \tilde{\chi}_n(\omega)\}} \bar{\lambda}_{\hat{n}}(\hat{\omega}), \quad (55)$$

where $\bar{\lambda}_n(\omega) = \sum_{\tau} \lambda_n(\omega, \tau)$ using the cross-sectional distribution from (53). Note that this formulation assumes the “overlapping generations” structure of our model economy: in each period, the mass of individuals born (and dying) is $1/N$, and the distribution of individuals evolves according to the decision rules outlined in Section 6.2.

Alternative constructions While we view the credit score we define in (54) as the closest possible to the analog in the data, there are in principle alternative ways to construct a credit score in the

²⁹Formally, by applying conditional expectations we can write $\mathbb{P}(d' = 0 \mid \omega, n) = \sum_{a'} \mathbb{P}(d' = 0 \mid a', \omega, n) \mathbb{P}(a' \mid \omega, n)$. In model notation, this can be written as $\tilde{\chi}_n(\omega) = (1 + r) \sum_{a'} q_n^{a'}(\omega) \left[s \bar{\sigma}_n^{(a', 0)}(\omega, H) + (1 - s) \bar{\sigma}_n^{(a', 0)}(\omega, L) \right]$.

model that may have somewhat different properties. For example, in [Chatterjee et al. \(2023\)](#) we defined an individual’s credit score as the likelihood of repayment tomorrow on a *standard* loan contract of size $\bar{a} < 0$.³⁰ This has the advantage of controlling for the fact that individuals may face different price schedules that shape their borrowing choices today and therefore how likely they are to default tomorrow. Put differently, the measure we adopt in (54) is “clouded” relative to this measure by the likelihood that individuals will respond to the *actual* prices they face by adjusting their borrowing decisions. At the same time, many individuals in our model have high credit scores not because they are unlikely to default conditional on borrowing, but because they are unlikely to borrow in the first place.

It is worth noting also that both these definitions use more information than what credit scorers are able to use in the data.³¹ For example, income technically cannot be included on a borrower’s credit report. We assume that income is observable in our model for two reasons. First, from a modeling perspective, it allows the model to retain a block recursive structure that facilitates computation by eliminating the need to use the equilibrium distribution to form posteriors. Second, from a practical perspective, it stands to reason that given the wealth of other types of information available to lenders, there is reasonably good inference about individuals’ income. Certainly, though, exploring such information restrictions is a promising area for future research.

6.3.2 Parameterization

To be consistent with the age brackets in Figure 1, a model period is 5 years, and focus on the age ranges of [21-25] ($n = 1$) through [56-60] ($n = 8$). We implement a simple, stylized income process that replicates the paths of the mean and dispersion in log income in the data in Figure 1. Specifically, we assume that at each age n , log income is normally distributed with effort-dependent mean μ_n^e and effort-independent standard deviation σ_n . We assume a path $\{\Delta_n^e\}$ of the difference between the mean of log income conditional on effort relative to no effort, $\Delta_n^e \equiv \mu_n^1 - \mu_n^0 > 0$, which is linearly decreasing from $\Delta_1^e = \bar{\Delta}^e$ to $\Delta_8^e = \underline{\Delta}^e$. This captures the idea that the returns to effort may be largest early in life. Given Δ_n^e , the path of average log income from the data, and a conjecture of the path of age-specific average effort rates, we compute the implied effort-specific paths for the means.³² In calibrating, we ensure that our model effort rates are close to the conjecture.

The remaining 10 parameters are chosen to match 15 moments in the data, 8 of which concern the life cycle patterns in log income, log consumption, and credit rankings documented in Section 2 and 7 of which are standard credit market statistics. The calibrated parameters govern the process for unobservable types, the effort cost, and the scale parameters governing the noisiness of each decision. Given we have posed a simple, stylized income process, our calibration in this section is meant only to provide a rough basis for the model properties to follow. In [Chatterjee et al. \(2023\)](#)

³⁰Formally, given the pricing schedule in (52), this alternative credit score would be $\chi_n(\omega) = (1 + r)q_n^{\bar{a}}(\omega)$.

³¹See [Albanesi and Vamossy \(2019\)](#) for an excellent discussion of what’s included in credit reports and how good credit scores are at actually sorting people based on the default risk that can be assessed using these reports.

³²That is, for a guess of age-specific average effort rate $\bar{e}_n$, we set $\mu_n^0 = \mu_n - \bar{e}_n \Delta^e$ and $\mu_n^1 = \mu_n^0 + \Delta^e$, where μ_n is the age-specific mean log income from the data.

parameter		value	notes / target	data	model
A. Assigned externally					
γ	risk aversion	2	standard CRRA utility		
r_{ann}	risk-free rate	2%	2% annualized, $r = (1 + r_{\text{ann}})^5 - 1$		
$\bar{\rho}$	long run share of H -types	2/3	prime share of population		
$\{\mu_n^0\}_{n=1}^N$	average log income, age j , no effort		path of average log income by age, no effort		
$\{\sigma_n\}_{n=1}^N$	SD log income, age j		path of SD income by age, effort-independent		
B. Calibrated internally					
B.1 Parameters			B.2 Targets		
<i>Type process</i>			<i>Life cycle</i>		
β_H	high discount factor	0.906	log income, mean, slope	0.047	0.044
β_L	low discount factor	0.815		variance, slope	0.009
ρ_1	initial share H types	0.444	log consumption, mean, slope	0.013	0.057
q	$\mathbb{P}(\tau' = H \mid \tau = L)$	0.100		variance, slope	0.026
			credit ranking, mean, intercept	0.385	0.184
<i>Extreme value</i>				mean, slope	0.024
α_e	noise, effort choice	0.001		variance, intercept	0.060
α_d	noise, default choice	0.030		variance, slope	0.006
$\alpha_{a'}$	noise, asset choice	0.005			0.007
<i>Earnings process</i>			<i>Credit market</i>		
			debt to income (%)	0.4	0.1
κ	effort cost	0.060	fraction in debt (%)	7.9	13.7
$\bar{\Delta}^e$	max log income gain from effort	0.400	average interest rate (pp ann)	11.9	17.6
$\underline{\Delta}^e$	min log income gain from effort	0.100	SD interest rate (pp ann)	7.0	5.2
			bankruptcy rate (pp ann)	1.0	1.2
			corr(credit ranking, income)	0.64	0.77

Table 4: **Parameters for quantitative model (Section 6)**

Notes: This table reports the parameter values used in the quantitative model in Section 6. Values in Panel A are calibrated outside of the model, values in Panel B are calibrated within the model. The average debt-to-income and variance of interest rate moments come from the Survey of Consumer Finances (SCF). The default rate is simply the annual rate of Ch. 7 bankruptcy filings from the U.S. Courts bankruptcy statistics. All other targeted moments come from the analysis in Section 2: the mean income growth is the slope of Figure 1a, the average correlation between credit ranking and income is the average from Figure 2b, the variance of credit rankings among the young is the intercept from Figure 1d, and the slope of the credit ranking is the slope from Figure 1c.

we employed a simulated method of moments estimation strategy to match credit market facts as in Table 4. There we also explore the sensitivity of parameter estimates to the moments of the data using the local methods in Andrews et al. (2017).

The idea behind this approach is as follows. Conditional on lifetime trends in income, which the model matches, credit market behavior is closely related to the average level of patience in the economy. Therefore, intensive margin moments such as the debt to income ratio and average interest rate and extensive margin moments such as the fraction of the population in debt help pin down the discount factor process. As discussed in the context of our numerical illustrations in Section 3, the rate at which types are revealed depends on how noisy decisions are relative to the fundamental difference in patience between the types. Therefore, conditional on the average level of patience informed by the credit market moments, that describe the changes in mean credit

rankings and their dispersion over the life cycle are very informative about the extreme value scale parameters associated with each choice. As in [Chatterjee et al. \(2023\)](#), our calibration is consistent with more noise in the default choice than other choices. In this richer setting with both adverse selection and moral hazard, this logic extends to the effort cost and benefit parameters, κ , $\overline{\Delta}^e$, and $\underline{\Delta}^e$.

6.3.3 On quantitative implementation of the extreme value shocks

While extreme value shocks resolve issues associated with Bayesian posteriors, they do create the potential for distortions in individual decisions. The source of the distortion is straightforward: in discrete choice environments, indirect utility is increasing in the number of choices, which is an endogenous function of the individual’s state. This implies that the marginal value of certain actions reflects the potential to benefit from more preference shock draws. This affects the default choice (40) and the asset choice (42), discussed individually below, but not the effort choice, since the number of choices ($2, e \in \{0, 1\}$) at this stage of the decision problem is unaffected by the state.

Savings Since our framework is dynamic, the savings effect strengthens individuals’ precautionary savings motive: saving more today means that the number of choices in tomorrow’s budget set is larger, which increases the marginal value of saving more. At a basic level, this makes individuals appear more patient for a given subjective discount factor β .³³ In practice, this can be handled by calibrating to match a desired target moment closely related to patience, such as the fraction of agents in debt. Second, these shocks introduce a novel form of grid sensitivity into standard grid-based solution methods. For example, simply increasing the fineness of the asset grid may exacerbate the issue, rather than increasing precision.

How can we address this latter issue quantitatively? In both [Chatterjee et al. \(2023\)](#) and [Briglia et al. \(2021\)](#), we treat the extreme value shocks as structural to the environment, capturing all the fundamental factors unobservable to the econometrician that affect an individual’s choice of a given level of consumption, rather than as “trembles” whose variance we’d like to take to zero. The basic idea of this approach is that the utility bonus associated with the extreme value draws should scale with the range of feasible consumption, not the number of feasible choices: while the latter is a function of the computational method used to solve the model, the former is a primitive of the economic environment. In Appendix A.1 of this paper, we provide detailed instructions for how to implement this method by adjusting the means of the extreme value shocks.

Default In our environment, default is only feasible for individuals in debt, i.e. with $a < 0$ in (40). Consider then the value of having some small amount of debt ($a = -\varepsilon$) versus having no debt at all ($a = 0$). If the extreme value shock associated with the default decision is noisy enough, then,

³³In [Briglia et al. \(2021\)](#), for example, we show that in a simple example with log utility, we show that the marginal value of wealth scales with $\frac{1}{a}$ at a rate of $\frac{1+\alpha}{1-\beta}$, compared to $\frac{1}{1-\beta}$ in the case without extreme value shocks.

the additional option to default may make the value associated with the former state higher than the latter. This would lead to a non-monotonicity of the value function in wealth.

While there have been several approaches to address this (see [Herkenhoff and Raveendranathan \(2024\)](#) and [Chatterjee et al. \(2023\)](#) for two examples), in [Appendix A.1](#) we provide an adjustment to the mean of the extreme value shock that eliminates the discontinuity in the value function at $a = 0$. The only complication with this approach is that the adjustment depends on the endogenous value function and so must be solved iteratively, but we find this process is quite stable and efficient.

6.4 Model properties

In this section we use our quantitative model to highlight four key properties of our framework. First, we analyze how reputational incentives contribute to inequality in income and consumption over the life cycle. Second, we demonstrate how the relationship between credit rankings and the fundamental state variables in our model – particularly type score and income – changes with features of the economic environment. Third, we use our model to determine the welfare effects of switching to an environment without adverse selection. Fourth, we show how our model generates relationships between credit scores, prices, and income that are broadly consistent with what we observe in the data.

6.4.1 How do reputational incentives contribute to inequality over the life cycle?

The illustrative analysis in [Sections 3 – 5](#) demonstrated how reputational incentives can increase dispersion throughout the life cycle in income, consumption, and wealth. This section examines these forces using our baseline quantitative model alongside 2 counterfactual models: one with low moral hazard (“low MH”), in which we lower the value of κ such that both types always choose to exert effort, and one with low adverse selection (“low AS”) in which the value of β_L is set close to the value of β_H to minimize the incentives to separate and therefore acquire a good reputation.

Panels (a), (b), and (c) of [Figure 11](#) show how age-bin-specific cross-sectional variances of log income, log consumption, and credit rankings vary over the life cycle for our baseline and two counterfactuals. By construction, all three economies exhibit similar income inequality. Relative to the baseline, the low moral hazard model exhibits slightly more consumption inequality. This is due to reputational incentives: since all agents exert effort in this economy, agents now separate more through savings and default decisions, which directly translate into consumption inequality. By contrast, the model with low adverse selection exhibits reduced consumption inequality compared to the baseline. This arises for two related reasons: the types are more similar in their “fundamental” preferences, which in turn implies that the cost of a bad reputation is lower. In terms of consumption inequality, these effects reinforce each other.

Reputational inequality is lower early in life in the low moral hazard model than in either the baseline or low adverse selection model. The main insight we gain from this finding is that the moral hazard model may actually make inference more challenging early in life. This comes from two channels. First, the presence of moral hazard creates reputational incentives around earning




Figure 11: **Inequality over the life cycle**

Notes: Each point in each panel of this figure corresponds to the age-bin-specific cross-sectional variance of the variable indicated in the title. The low moral hazard (low MH) model lowers κ so that all types always choose effort. The low adverse selection (low AS) model sets β_L very close to β_H to mute the reputational incentives.

higher income as discussed in Section 4.3. Weakening this channel then limits separation, since the high types respond to these reputational incentives more strongly. Second, as all agents work harder and therefore become richer in expectation in the future, they borrow more against that income early in life. This further makes the asset choices less informative, reinforcing the reduction in separation.

6.4.2 What do credit rankings measure?

As discussed in Section 6.3.1, real-world credit rankings in principle include information about much more than the assessment of an unobservable characteristic, like the type score we use in our model. In this subsection, we consider three distinct questions related to credit rankings in our model. First, what is the relationship between individuals' credit rankings and their type assessments? In our baseline model, this correlation is low but positive, averaging about 0.25 across the entire population. Figure 12(a) shows that this correlation increases over the life cycle, consistent with the idea that there is something fundamental and unobservable about individuals that lenders learn about over time, as discussed in Section 2. The level and steepness of this profile of correlations, however, depend on the specifics of the economic environment, which should be considered in future work. To establish this, consider the low moral hazard and low adverse selection models from the previous section alongside the baseline. With low moral hazard, there is convergence of incomes across types and income is higher on average, so the assessment of type becomes more important and the correlation of credit ranking with type score increases. By contrast, with low adverse selection, this correlation is quite low over the entire life span since types are not very different.

Second, harkening back to the discussion in Section 2, how important is the assessment of unobservable type relative to income in driving credit rankings? Figure 12(b) shows that the correlation between income and credit ranking in the model is much higher on average than the correlation between type score and credit ranking. Notably, it is also much higher early in life than later in life, because in our model income is the only separating characteristic early in life, which is




Figure 12: Frictions, learning and credit rankings over the life cycle

Notes: The low moral hazard (low MH) model lowers κ so that all types always choose effort. The low adverse selection (low AS) model sets β_L very close to β_H to mute the reputational incentives. The correlations in both panels are age-bin specific. In both panels, the correlation is undefined for the first age bin since there is no dispersion in type score or income, and therefore no dispersion in credit rankings.

not necessarily true in the data. This creates a gap between our model and the data that may be fruitful to explore in future research.

Third, what is the correlation between income and the noisy signal of unobservable type (which we take to be the type score in our model)? This is a version of the object σ_{Y_n, U_n} from Section 2 that we took to be positive but decreasing over the age profile in Figure 4(f). Figure 12(c) shows that this correlation generically takes a hump-shape over the life cycle. This is because early in life (when Δ_n^e is high), the reputational incentives associated with effort are strongest, and therefore the high types are more likely to exert effort and income is a strong signal of type. This intuition is confirmed by the alternative models: the pattern is virtually absent with low moral hazard, and muted when adverse selection is sufficiently small to dampen reputational incentives.

6.4.3 Welfare analysis: is more information better?

How costly (or beneficial) is the incentive to acquire a good reputation? To address this question, we conduct a welfare analysis that asks: how much wealth would we have to give an individual born into the baseline economy to be indifferent to being born into the full information economy (with no reputational incentives)?³⁴ The sign convention for this welfare measure is that a positive number means that the full information economy is preferred, and we express the wealth transfers as a percent of average income in the baseline model. Table 5 presents the results of this analysis in the aggregate and by type for two variants of this welfare metric.

The first metric, reported in the first line of the table, uses the exact value functions from the individual problem in (39) – (45). Metrics are reported as percentages of average income: that is, on average, an individual would need to receive a (one-time) wealth transfer equal to 0.10% of average income in the baseline economy to be as well off as if she were born into the full information

³⁴Formally, we solve for $\{\chi_1(\omega, \tau)\}$ such that $V_{E,1}(a(\omega) + \chi_1, y(\omega), s(\omega), \tau) = V_{E,1}^{FI}(a(\omega), y(\omega), \tau)$, where V^{FI} is the full information value function, then weight the χ_1 's by the distribution of newborns, $\lambda_1(\omega, \tau)$. In the full information economy, τ is directly observable (but effort remains unobservable), and so there is no need for type scores.

	aggregate	H types	L types
baseline	0.10	0.19	0.04
baseline, no EV shocks	0.13	0.23	0.05

Table 5: **Welfare analysis: baseline vs full information**

Notes: The units of this wealth-equivalent welfare measure are percentage points of economy-wide average income in the baseline model (e.g. the average newborn would have to be given wealth of 0.05% of average income in the baseline economy to be just as well off as being born into the full information economy).

economy. Thus, the full information economy is preferred to the baseline. Why is this the case? Recall that in the presence of adverse selection and moral hazard in this model, individuals have an extra incentive to ration consumption and work harder to build and preserve a good reputation to facilitate access to credit. Since both of these actions are directly costly in terms of utility, agents tend to prefer the full information world with no adverse selection.

In principle, this first effect is counteracted by a second: the cross-subsidization of types in credit pricing, which occurs only in the baseline. This effect is negative for H types and positive for L types: in the baseline, partial pooling implies that H types face worse terms of credit since they are lumped in with L types who are more likely to default. On the margin, this leads the H types to prefer the full information economy relative to the baseline *even more* than the L types.

In light of the discussion of the effects that extreme value shocks have on the indirect utility function from Section 6.3.3, one might find it desirable to evaluate welfare effects using the value function that shuts down these shocks. To address this, rather than solve for the full equilibrium with no extreme value shocks – which is not generically possible given the issues with Bayesian equilibrium described in Section 3.3 – we simply take the equilibrium loan price and reputation updating functions from the model with extreme value shocks and solve for the optimal policies and value functions assuming that the variance of the ε shocks are all zero.³⁵ The results of this analysis, presented in the second line of Table 5, are qualitatively similar to our baseline results, but with modestly larger magnitudes.

6.4.4 Interest rates, income, and credit rankings

At the core of our analysis is that reputation matters because it affects access to credit. Figure 3 documents that this is true in the data: conditional on income, higher FICO scores imply lower APRs on credit cards. Figure 13(a) is the model analog of Figure 3, plotting the balance-weighted average interest rate for each type score tercile-by-income quartile bin in our model economy. We use type score rankings in this analysis rather than credit rankings in order to isolate the role of reputation separately from income and wealth, which also affect credit rankings in our model.

Broadly, the key patterns in the data are replicated in our model. First, conditional on income quartile, average interest rates are generally decreasing in type score ranking: that is, the more likely

³⁵We thank our discussant, Kyle Herkenhoff, for this useful suggestion.




Figure 13: **Credit prices, borrowing, and reputation by income and credit ranking**

Notes: Panel (a) plots the balance-weighted average interest rate for the indicated type score ranking tercile for the indicated income quartile, analogous to the empirical profile in Figure 3. Type score ranking bins are computed using the equilibrium CDF of type scores across the population, analogously to the credit ranking procedure discussed in Section 6.3.1. Panel (b) plots the (smoothed) share of economy-wide debt balances accounted for by the indicated type score ranking by income quartile bin. In both panels (a) and (b), the top income quartile is excluded, because total debt balances among this group are less than 0.01% of economy-wide debt. Panel (c) is the model’s analog of Figure 2a, plotting the share of nonprime / prime / superprime borrowers in each income tercile, where nonprime / prime / superprime refers to the bottom / middle / top third of the credit ranking distribution.

an individual is perceived to be the H type, all else equal, the lower is the interest rate they pay when borrowing. The first and second income quartiles account for 68.7% and 25.4% of economy-wide debt, respectively, and the first and second terciles of the type score ranking distribution account for 48.2% and 43.1% of total debt, respectively. Second, conditional on reputation, as in the data individuals with higher income tend to pay lower interest rates. For example, in the bottom third of the type score distribution, the bottom income quartile’s average interest rate is 20.5%, compared to 18.3% and 17.1% for the second and third quartiles, respectively.

While not shown in this figure, the only exceptions to this decreasing pattern occur with high reputation and high income, where borrowing is infrequent. In these parts of the distribution, borrowing is driven by the noisiness in decisions, which makes these borrowers less sensitive to the interest rate, leading to higher balance-weighted averages.

Figure 2a also shows that the composition of borrowers in terms of their credit scores improves as we move up the income distribution. Figure 13(c) replicates this analysis in our model by plotting the share of individuals in each third of the credit ranking distribution in each third of the income distribution. Just as in the data, the bottom income tercile is dominated by nonprime borrowers, while the top is dominated by superprime borrowers.

7 Directions for Future Research

How much does reputation affect inequality? In this paper, we associate reputation with an individual’s conscientiousness, represented by their discount rate β ; higher β translates to less immediate gratification, which makes an individual less likely to borrow and default, and more likely to expend effort to generate higher future earnings. In our model, an individual’s “type score (s)” summarizes

their reputation, which is an input into a Bayesian assessment of their likelihood of repayment represented by their “credit ranking (χ)”. But it is clear that the likelihood of repayment depends not just on an individual’s type but also on their income as well as other factors. This makes it challenging to learn about the life-cycle evolution of reputation from the life-cycle evolution of credit scores. Nevertheless, we believe understanding the role of reputation in a structural life cycle model is an important avenue for future research since screening on the basis of credit scores spills over to many other areas of everyday life, including housing and employment.

Another challenge is accounting for what information can be conditioned upon when constructing a credit score. Broadly speaking, this relates to privacy concerns that have motivated much information-related legislation in the US. Our baseline model assumes that all factors relevant for predicting default, such as income, assets, type scores, and credit market actions, are observable to lenders and may be used to determine interest rates and updated type scores (and, thus, one’s credit ranking). This makes it very easy to maintain a Markov structure. In addition, free entry by lenders gives the model a block recursive structure, making it easy to solve. Restrictions on that information, which are likely in the real world, can be incorporated at the cost of the loss of block recursivity: An equilibrium would be a fixed point not only of value functions, prices, and posterior functions, but also of the cross-sectional distribution, which would be needed to formulate risk assessment posteriors (the “no tracking” environment is an example). Such restrictions would imply that type scores and credit rankings would incorporate lenders’ assessments of other unobservable but payoff-relevant individual states. We suspect this would make reputations harder to acquire and easier to maintain and potentially cause them to be a more important force in the persistence of inequality operating through credit access. Be that as it may, another fruitful avenue for future research is the exploration of the tradeoffs between privacy, efficiency, and equity.

Beyond *what* information can be used, another important factor is *how much* information is available to be used. Specifically, there is generally less information available for the young when constructing a measure of reputation as they have made fewer choices as adults. As a consequence, it is harder to *learn* about a young adult’s type. This might be a good thing if greater pooling provides important insurance benefits. However, how much information is available on the young depends also on societal factors. The expansions of student loan programs in the US and elsewhere have exposed many more young adults to the logic of reputation formation via credit markets than earlier cohorts. Another fruitful area of research would be to explore if debt-finance human capital accumulation has sped up learning about type and whether this has resulted in greater reputational inequality among younger cohorts. Ultimately, however, one needs a model to quantify the costs and benefits of information over the age profile.

Finally, one important goal of this and our predecessor paper is to understand how personality traits underpin reputation in credit markets. Economists typically shy away from explanations for behavior that rely on differences in preferences, preferring instead to base explanations on differences in opportunity sets. But personality traits identified by psychologists, such as OCEAN, hold out the hope that preference differences can be incorporated in a disciplined way into models of reputation

formation. In future work, we plan to expand the notion of type to include differences along several aspects of preferences, including present bias, intertemporal elasticity of substitution, risk aversion, and disutility from effort. We hope that this expanded notion of type can provide a better account of the life-cycle evolution of reputation.

References

- AGARWAL, S., S. CHOMSISENGPHET, AND C. LIU (2010): “The Importance of Adverse Selection in the Credit Card Market: Evidence from Randomized Trials of Credit Card Solicitations,” *Journal of Money, Credit and Banking*, 42.
- ALBANESI, S. AND D. F. VAMOSSY (2019): “Predicting Consumer Default: A Deep Learning Approach,” Working Paper No. 26165, National Bureau of Economic Research.
- AMADOR, M. AND C. PHELAN (2021): “Reputation and Sovereign Default,” *Econometrica*, 89, 1979–2010.
- AMERIKS, J., A. CAPLIN, J. LEAHY, AND T. TYLER (2007): “Measuring Self-Control Problems,” *American Economic Review*, 97, 966–972.
- ANDREWS, I., M. GENTZKOW, AND J. SHAPIRO (2017): “Measuring the Sensitivity of Parameter Estimates to Estimation Moments,” *Quarterly Journal of Economics*, 132, 55–70.
- ATHREYA, K. (2002): “Welfare Implications of Bankruptcy Reform Act of 1999,” *Journal of Monetary Economics*, 49, 1567–1595.
- AUSUBEL, L. (1999): “Adverse Selection in the Credit Card Market,” Unpublished, University of Maryland.
- BARTH, D., N. W. PAPAGEORGE, AND K. THOM (2020): “Genetic Endowments and Wealth Inequality,” *Journal of Political Economy*, 128, 1474–1522.
- BRAXTON, J. C., K. HERKENHOFF, AND G. PHILLIPS (2024): “Can the Unemployed Borrow? Implications for Public Insurance,” *Journal of Political Economy*, 132, 3025–3076.
- BRIGLIA, L., S. CHATTERJEE, D. CORBAE, K. DEMPSEY, AND V. RIOS-RULL (2021): “Saving for a Sunny Day: An Alternative Theory of Precautionary Savings,” Unpublished.
- CARROLL, C., J. SLACALEK, K. TOKUOKA, AND M. N. WHITE (2017): “The Distribution of Wealth and the Marginal Propensity to Consume,” *Quantitative Economics*, 8, 977–1020.
- CHATTERJEE, S., D. CORBAE, K. DEMPSEY, AND J.-V. RIOS-RULL (2023): “A Quantitative Theory of the Credit Score,” *Econometrica*, 91, 1803–1840.
- CHATTERJEE, S., D. CORBAE, M. NAKAJIMA, AND J.-V. RIOS-RULL (2007): “A Quantitative Theory of Unsecured Consumer Credit with Risk of Default,” *Econometrica*, 75, 1525–1589.
- CHATTERJEE, S., D. CORBAE, AND J.-V. RÍOS-RULL (2008): “A finite-life private-information theory of unsecured consumer debt,” *Journal of Economic Theory*, 142, 149–177.
- CHO, I.-K. AND D. M. KREPS (1987): “Signaling Games and Stable Equilibria,” *Quarterly Journal of Economics*, 102, 179–222.
- COLE, H. AND P. KEHOE (1998): “Models of Sovereign Debt: Partial Versus General Reputations,” *International Economic Review*, 39, 55–70.
- CORBAE, D. AND A. GLOVER (2025): “Employer Credit Checks: Poverty Traps Versus Matching Efficiency,” *Review of Economic Studies*, 92, 1661–1698.

- DEMPSEY, K. AND F. IONESCU (2024): “Default Risk and Pricing in the U.S. Credit Card Market,” Unpublished.
- DIAMOND, D. W. (1989): “Reputation Acquisition in Debt Markets,” *Journal of Political Economy*, 97, 828–862.
- DOKKO, J., G. LI, AND J. HAYES (2015): “Credit Scores and Committed Relationships,” Finance and Economics Discussion Series No. 2015-081, Federal Reserve Board, Washington D.C.
- EINAV, L., M. JENKINS, AND J. LEVIN (2012): “Contract Pricing in Consumer Credit Market,” *Econometrica*, 80, 1387–1432.
- EXLER, F. AND M. TERTILT (2020): “Consumer Debt and Default: A Macro Perspective,” *Oxford Research Encyclopedia of Economics and Finance*.
- FULFORD, S. AND D. LOW (2024): “Expense Shocks Matter,” Unpublished.
- GOEREE, J. K., C. A. HOLT, AND T. R. PALFREY (2016): *Quantal Response Equilibrium*, Princeton and Oxford: Princeton University Press.
- GOLDEN, L. L., P. L. BROCKETT, J. AI, AND B. KELLISON (2016): “Empirical Evidence on the Use of Credit Scoring for Predicting Insurance Losses with Psycho-social and Biochemical Explanations,” *North American Actuarial Journal*, 20, 233–251.
- GUERRERI, V., R. SHIMER, AND R. WRIGHT (2010): “Adverse Selection in Competitive Search,” *Econometrica*, 78, 1823–1862.
- HEATHCOTE, J., F. PERRI, AND G. L. VIOLANTE (2023): “More Unequal We Stand? Inequality Dynamics in the United States, 1967-2021,” *Review of Economic Dynamics*, 50, 235–266.
- HECKMAN, J. J., B. GALATY, AND H. TIAN (2023): “Economic Approach to Personality, Character and Virtue,” *NBER Working Paper No. 31258*.
- HERKENHOFF, K. F. AND G. RAVEENDRANATHAN (2024): “Who Bears the Welfare Cost of Monopoly? The Case of the Credit Card Industry,” *Review of Economic Studies*.
- HERTZBERG, A., A. LIBERMAN, AND D. PARAVISINI (2018): “Screening on Loan Terms: Evidence from Maturity Choice in Consumer Credit,” *The Review of Financial Studies*, 31, 3532–3567.
- HIRSCHLEIFER, J. (1971): “The Private and Social Value of Information and the Reward to Inventive Activity,” *American Economic Review*, 61, 561–574.
- ISRAEL, S., A. CASPI, D. W. BELSKY, H. HARRINGTON, S. HOGAN, R. HOUTS, S. RAMRAKHA, S. SANDERS, R. POULTON, AND T. E. MOFFITT (2014): “Credit scores, cardiovascular disease risk, and human capital,” *Proceedings of the National Academy of Sciences*, 111, 17087–17092.
- KOVRIJNYKH, N., I. LIVSHITS, AND A. ZETLIN-JONES (2024): “Building Credit Histories,” Unpublished.
- KREPS, D. M. AND J. SOBEL (1994): “Signalling,” in *Handbook of Game Theory*, ed. by R. Aumann and S. Hart, Elsevier, vol. 2, chap. 25, 849–867.
- KRUSELL, P. AND A. A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106, 867–896.

- LAUER, J. (2017): *Creditworthy: A History of Consumer Surveillance and Financial Identity in America*, Columbia Studies in the History of U.S. Capitalism, New York: Columbia University Press.
- LIVSHITS, I., J. C. MACGEE, AND M. TERTILT (2007): “Consumer Bankruptcy: A Fresh Start,” *American Economic Review*, 97, 402–418.
- (2016): “The Democratization of Credit and the Rise in Consumer Bankruptcies,” *Review of Economic Studies*, 83, 1673–1710.
- MATEJKA, F. AND A. MCKAY (2015): “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model,” *American Economic Review*, 105, 272 – 298.
- MATTSON, L.-G. AND J. W. WEIBULL (2002): “Probabilistic Choice and Procedurally Bounded Rationality,” *Games and Economic Behavior*, 41, 61–78.
- MCFADDEN, D. (1974): “Conditional Logit Analysis of Qualitative Choice Behavior,” in *Frontiers in Econometrics*, ed. by P. Zarembka, Academic Press, New York, 105–142.
- MEIER, S. AND C. D. SPRENGER (2012): “Time Discounting Predicts Creditworthiness,” *Psychological Science*.
- MENZIO, G. AND S. SHI (2010): “Block Recursive Equilibria for Stochastic Models of Search on the Job,” *Journal of Economic Theory*, 145, 1453–1494.
- MYERSON, R. B. (1978): “Refinements of the Nash Equilibrium Concept,” *International Journal of Game Theory*, 7, 73–80.
- NARAJABAD, B. N. (2012): “Information Technology and the Rise of Household Bankruptcy,” *Review of Economic Dynamics*, 15, 526–550.
- PASQUALE, F. (2015): *The Black Box Society: The Secret Algorithms That Control Money and Information*, Cambridge, MA: Harvard University Press.
- POON, M. (2007): “Scorecards as Devices for Consumer Credit: The Case of Fair Isaacs Company Incorporated,” *Sociological Review*, 55, 284–306.
- ROTHSCHILD, M. AND J. STIGLITZ (1976): “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” *Quarterly Journal of Economics*, 90, 629–649.
- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999–1033.
- SANCHEZ, J. M. (2018): “The Information Technology Revolution and the Unsecured Credit Market,” *Economic Inquiry*, 56, 914–930.
- SELTEN, R. (1975): “Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games,” *International Journal of Game Theory*, 4, 25–55.
- XIN, Y. (2023): “Asymmetric Information, Reputation, and Online Credit Markets,” working paper.

Appendix for “Credit Scores and Inequality across the Life Cycle”

A Model Appendix

This appendix is intended to serve as a “user guide” for solving a quantitative model with hidden information and Bayesian updating, with the Bayesian issues handled via extreme value shocks as described in Section 3.3. To this end, we organize the Appendix as follows. First, we describe in detail the issues with the extreme value shocks discussed in Section 6.3.3. Then, we lay out the detailed computational algorithm for solving the quantitative model from Section 6. Finally, we include the details of some of the alternative information structures we use throughout the paper. This Appendix is intended to be used alongside the replication package for all the results in the paper, located at <https://zenodo.org/records/15731380>.

A.1 On quantitative implementation of the extreme value shocks

A.1.1 Savings decision

The problem As discussed in Section 6.3.3, the fact that indirect utility in a discrete choice environment is increasing in the number of choices creates two issues in problem (42) that are absent in a standard consumption-savings problem. First, saving more today increases the number of choices in tomorrow’s budget set, which increases the marginal value of saving. As discussed in the main text, this can be handled via calibration, adjusting β to match certain targets. Second, these shocks make the value function and decision rules sensitive to the grid used to solve the model. Researchers typically approximate the value function over a discrete grid for the individual states. For a given grid density, this approximation is good when the function is close to linear, but worse when the function is more concave. A common solution is to increase grid density in regions of the state space with high curvature. One example is the use of log-spaced (rather than linearly-spaced) grids in wealth for the value function. Given that the number of choices matters, though, such a switch is not innocuous: holding the number of grid points fixed, switching from linear to log spacing effectively increases the value function at low wealth levels and lowers it at high wealth levels. Relatedly, increasing the fineness of the grid may exacerbate the problem.

The solution How can we resolve these tensions? In both Chatterjee et al. (2023) and Briglia et al. (2021), we treat the extreme value shocks as structural to the environment rather than as “trembles” whose variance we’d like to take to zero but for computation. We assume that the “utility bonus” associated with the extreme value shocks scales with the size of the budget set rather than the number of feasible choices. This resolves the grid sensitivity induced by the extreme value shocks: the size of the budget set at a given state is a fundamental feature of the model, not a vestige of the number or location of grid points used to solve it.

An attractive feature of this approach is that we can control for the issue of the number of

grid points by simply adjusting the means of the extreme value shocks.³⁶ The standard assumption is that the shocks associated with each choice i , ε^i , are distributed type one extreme value with scale parameter α and location parameter $-\alpha\gamma_E$, where γ_E is the Euler-Mascheroni constant. This implies that the expected maximum of N independent draws from this distribution satisfies $\mathbb{E}_{\varepsilon^i} [\max_{i=1,\dots,N} \varepsilon^i] = \alpha \ln N$. The issue is that if we double the number of feasible grid points at a given point in the state space, then the analogous expectation becomes $\alpha \ln 2N$, which introduces a potential distortion. The idea behind our approach is that in both the N and $2N$ cases, the utility bonus should be the same, since the fundamental feasible set is the same.

To illustrate further, assume that the model is solved on a discrete grid for wealth $A = \{a_1, \dots, 0, \dots, a_N\}$, and suppose we attach a fundamental, state-dependent, action-specific value $v_A^i(x)$ to each action.³⁷ Suppose further that $n_A(x) = |\mathcal{F}(x)|$ is the number of feasible choices on grid A for an individual in state x . Maintaining the assumptions on the distribution of ε_i outlined above, the ex ante expected value is:

$$V_A(x) = \mathbb{E}_{\varepsilon^i} \left[\max_{i=1,\dots,n_A(x)} \{v_A^i(x) + \varepsilon^i\} \right] = \alpha \ln \left(\sum_{i=1}^{n_A(x)} \exp \left(\frac{v_A^i(x)}{\alpha} \right) \right).$$

To make the point simply, suppose $v_A^i(x) = \bar{v}(x)$ for all i . Then, $V_A(x) = \bar{v}(x) + \alpha \ln n_A(x)$, but the analogous value for a grid $\tilde{A}$ with twice as many feasible points is $V_{\tilde{A}}(x) = \bar{v}(x) + \alpha \ln 2n_A(x)$. The difference in these results stems entirely from the difference between grids A and $\tilde{A}$, not a fundamental of the underlying model.

Suppose instead that for a state x and a grid A , we define $d_A^i(x)$ as the measure of consumption associated with choice $a' = a_i$. By construction, $\sum_{i \in \mathcal{F}(x)} d_A^i(x) = \bar{c}(x)$, where $\bar{c}(x)$ is the maximal attainable consumption for an individual in state x , which does not depend on the grid. We adjust the location parameter of the extreme value shock for each action: rather than setting it to $-\alpha\gamma_E$ for all i and x , we now set it to $-\alpha\gamma_E + \alpha \ln d_A^i(x)$. A straightforward derivation shows that now:³⁸

$$V_A(x) = \alpha \ln \left(\sum_{i=1}^{n_A(x)} d_A^i(x) \exp \left(\frac{v_A^i(x)}{\alpha} \right) \right).$$

Now, repeat the same thought experiment as before in which $v_A^i(x) = \bar{v}(x)$ for all i . Then, $V_A(x) = \bar{v}(x) + \alpha \ln \left(\sum_{i=1}^{n_A(x)} d_A^i(x) \right)$. Since $\sum_{i=1}^{n_A(x)} d_A^i(x) = \bar{c}(x)$ for all grids A by definition, then, the value does not depend on the grid: that is, $V_A(x) = V_{\tilde{A}}(x) = \bar{v}(x) + \alpha \ln \bar{c}(x)$. Of course, when

³⁶See Briglia et al. (2021) for an even more detailed explanation of the mechanics behind this approach.

³⁷To be explicit, in the full quantitative model, we have $x = (\omega, \tau, n)$, and each action i corresponds to a specific $a' = a_i$ in problem (42).

³⁸An analogous derivation can be performed for decision rules: the probability of choosing action i is

$$\sigma_{i,A}(x) = \frac{d_A^i(x) \exp \left\{ \frac{v_A^i(x)}{\alpha} \right\}}{\sum_{j=1}^{n_A(x)} d_A^j(x) \exp \left\{ \frac{v_A^j(x)}{\alpha} \right\}}.$$

$v_A^i(x)$ depends on i , there will still be some grid sensitivity as in a standard approach: however, this sensitivity will not arise as a vestige of the discrete choice shocks. Conceptually, the $d_A^i(x)$ coefficients serve as weights, recognizing that since the values of “off-grid” actions a' near a given a_i should be close to $v_A^i(x)$ by continuity, and so actions that “span” a larger portion of the budget set should be up-weighted accordingly.

How do we compute these weights in practice? The idea is to compute how much of the budget set – in terms of consumption – is accounted for by each grid point i on the asset grid A over which the value function and decision rules are approximated. To do this, define an “auxiliary” asset grid $\hat{A} = \{\hat{a}_1, \dots, \hat{a}_N\}$ of the midpoints between choices on the fundamental asset grid: that is, set

$$\hat{a}_1 = a_1 \quad \text{and} \quad \hat{a}_i = a_{i-1} + \frac{a_i - a_{i-1}}{2} \quad \text{for } i = 2, \dots, N.$$

Then, for each grid point i on A , assign the weight

$$d_A^i(x) = \begin{cases} c^{\hat{a}_i}(x) - c^{\hat{a}_{i+1}}(x) & \text{if } i < n_A(x) \\ c^{\hat{a}_i}(x) & \text{if } i = n_A(x). \end{cases} \quad (\text{A.1})$$

The idea behind this is that the range $[\hat{a}_i, \hat{a}_{i+1})$ on the auxiliary grid $\hat{A}$ defines the “attraction band” for each point a_i on the fundamental grid A : these are the “off-grid” actions that are closest to grid point a_i and therefore whose value is most closely approximated by $v_A^i(x)$. Conditional on the action-specific value, actions with larger attraction bands should be more likely to be chosen. The correction for the maximal feasible grid point (second branch in (A.1)) reflects the fact that the budget set may not include the entire attraction band of this point.

There are two issues with this approach in the context of the pricing of debt under limited commitment. First, there may be a set of debt choices on the “wrong side of the Laffer curve” in the sense that default risk is so large that there exists a smaller debt level that yields more current consumption and a lower future debt burden.³⁹ These actions would never be chosen without extreme value shocks, and they would imply negative weights in (A.1). One might consider assigning a weight of zero to such actions, but this would imply that such an action is not chosen, undermining the original goal of keeping Bayesian posteriors well-defined.⁴⁰ In practice, the fundamental values of such actions are so low relative to feasible alternatives that they are rarely chosen anyway, and so we avoid this issue by assigning these actions an arbitrarily small but positive weight. Second, there may be non-monotonicities in the price schedule that may also lead to negative weights. This problem is solved similarly: essentially, in practice we augment (A.1) with the condition that if $d_A^i(x) < 0$, then we assign $d_A^i(x) = \zeta$, an arbitrary small positive parameter.

³⁹Formally, there may exist $i < j$ such that $-q^{a_i}(x)a_i < -q^{a_j}(x)a_j$.

⁴⁰Note that there is no issue with this implementation in an environment without Bayesian updating.

A.1.2 Default decision

The problem Another problem related to the increase in value associated with more feasible choices comes from the default decision. In our environment, default is only feasible for individuals in debt, i.e., with $a < 0$: see problem (40). Consider then, the value of being some small amount in debt ($a = -\epsilon$) vs having no debt at all ($a = 0$). If the variance of the extreme value shock associated with the default decision is high enough, then, the additional option to default may make the value associated with the former state higher than the latter. This would break the standard result of monotonicity of the value function in wealth.

The solution To date, the literature has taken two approaches to addressing this, each with their own benefits and drawbacks. One approach, taken for example in [Herkenhoff and Raveendranathan \(2024\)](#), is to assume that individuals can choose to default on positive amounts of wealth. While this eliminates the mathematical issue of the non-monotonicity near $a = 0$ by keeping the set of options available constant across all wealth levels, one might question what it means to “default” on savings. The suggested interpretation is that the choice to default is in response to unmodeled large expense shocks (e.g., liabilities stemming from medical or divorce bills and lawsuits). Implicitly, this assumption requires that such unmodeled shocks are larger than the endogenous upper bound on wealth (i.e., the wealthiest individual in the model economy).

The second approach, taken in [Chatterjee et al. \(2023\)](#), does not fully resolve the non-monotonicity but dampens it by taking an approach similar in spirit to the approach described above for the asset choices. Specifically, assume that the extreme value shocks ε^d associated with the default decision have scale α_d and mean $-\alpha\gamma_E - \alpha_d \ln 2$, so that $\mathbb{E}_{\varepsilon^d} [\max_{d \in \{0,1\}} \varepsilon^d] = 0$. In this case, defining $V_D(a, x)$ as the value of defaulting (which is defined only for $a < 0$) and defining $V_R(a, x)$ as the value of repaying (which is defined for all a), the ex-ante value is equal to

$$V(a, x) = \begin{cases} \underbrace{\alpha_d \ln \left(\exp \left\{ \frac{V_D(a, x)}{\alpha_d} \right\} + \exp \left\{ \frac{V_R(a, x)}{\alpha_d} \right\} \right)}_{\equiv \tilde{V}(a, x)} - \alpha_d \ln 2 & \text{if } a < 0 \\ V_R(a, x) & \text{if } a \geq 0. \end{cases} \quad (\text{A.2})$$

Note that we have isolated the wealth state variable a away from the other state variables in x for ease of exposition later. The idea is that if default and no default have the same value, then the expressions in the top and bottom branches are equal to $V_R(a, x)$, and so the fundamental discrepancy between the top and bottom branches of the expression above is eliminated.

Here, we advocate for a new approach that resolves the fundamental problem using a different adjustment of the mean of the extreme value shocks. We want the ex ante value function in (A.2) to have the property that $\lim_{a \nearrow 0} V(a, x) = V_R(a, x)$ for all x , since this implies no discontinuity in the value function at $a = 0$, as in the case with no extreme value shocks. To achieve this, we simply assign the mean of the extreme value shocks to be $-\alpha\gamma_E - \eta(x)$, where $\eta(x) = \tilde{V}(a^*, x) - V_R(a^*, x)$, where $a^* < 0$ is an arbitrarily small debt level (a natural choice is the first grid point to the left

of 0). The one complication with this approach is that the adjustment depends on the endogenous value function. In practice, though, if we iterate on the object $\eta(x)$ alongside the value function as we solve the model, this process is quite stable and efficient. We described this procedure in our computational algorithm in Section A.2 below.

A.2 Computational algorithm for full model in Section 6

1. **Assign all parameters and grids.** We find it useful to to construct the asset grid A as log-spaced in both directions from 0, and for the type score grid S to be log-spaced in both directions from ρ_1 . Note that $S \subseteq [Q_\tau(H' \mid L), Q_\tau(H' \mid H)]$ since the type score cannot go below the likelihood that someone identified as a low type today will be a high type tomorrow. An important restriction is that $y_1 + a_1 > 0$, so that repaying debt is always feasible (even the largest debt from the smallest level of income). Also specify the auxiliary asset grid $\hat{A}$ as described in Section A.1 above.
2. **Construct initial guesses** of the decision rules $\sigma_{E,n}^e(\omega, \tau)$, $\sigma_{D,n}^{(d,e)}(\omega, \tau)$, and $\sigma_{A,n}^{(a',e)}(\omega, \tau)$ for all $n \in \{1, \dots, N\}$, $\omega = (a, y, s) \times Y \times S$, $\tau \in \{H, L\}$, $e \in \{0, 1\}$, $d \in \{0, 1\}$, and $a' \in A$, as well as the extreme value shock adjustment for default $\eta_n^e(y, s, \tau)$ for all e, n, y, s, τ .
 - For $n = N$, we rule out $a' < 0$ since there is no possibility of repayment. For consistency with earlier periods, though, we allow $a' > 0$ and $e = 1$ (even though it is suboptimal), to keep the degree of noise in individuals' decisions constant throughout their lifetimes.
 - These initial guesses must be positive for every feasible action in an individual's state. For example, $\sigma_{E,n}^e(\omega, \tau) \in (0, 1)$ for all e, ω, τ, n . Likewise, $\sigma_{D,n}^{(d,e)}(\omega, \tau) \in (0, 1)$ for all $d \in \{0, 1\}$ when $a(\omega) < 0$, while $\sigma_{D,n}^{(d=1,e)}(\omega, \tau) = 0$ when $a(\omega) \geq 0$. Finally, $\sigma_{A,n}^{(a',e)}(\omega, \tau) \in (0, 1)$ for all $a' \in \mathcal{F}_n(\omega)$, while $\sigma_{A,n}^{(a',e)}(\omega, \tau) = 0$ for all $a' \notin \mathcal{F}_n(\omega)$.
 - A good initial guess is $\eta_n^e(y, s, \tau) = 0$.
3. **Solve backward for the equilibrium value, decision, pricing, and type score updating functions.** In practice, this means iterating to convergence on the decision rules at each age. Note that the terminal age is slightly different since there is no continuation value; in all other periods, the continuation value is simply $V_{n+1}(\omega, \tau)$, which has already been solved. To keep things compact here, however, we consider a generic age $n < N$.
 - (a) Given an initial guess of the fundamental decision rules, compute the total decision rule $\sigma_n^{(a',d,e)}(\omega, \tau)$ via (48), the observed decision rule $\bar{\sigma}_n^{(a',d)}(\omega, \tau)$ via (49), and the inferred income process $\tilde{Q}_{y,n}^{(a',d)}(\omega, \tau)$ via (50).
 - (b) Given the auxiliary decision rules and inferred income process, compute the implied type scoring function $\Upsilon_n^{(a',d)}(y'; \omega)$ via (51) and loan pricing function $q_n^{a'}(\omega)$ via (52).

- (c) Compute the consumption associated with each borrowing or saving action implied by the current pricing function, $c_n^{a'}(\omega)$, via (44).⁴¹
- (d) Compute the action-specific weights for the extreme value adjustment for the borrowing / savings choice, $d_n^{a'}(\omega)$ via (A.1).
- (e) Work backwards through the timeline of individual choices within the period to compute the value functions and decision rules. Specifically, for each (ω, τ) :
 - i. For each choice of effort e , compute the action-specific value $v_{A,n}^{(a',e)}(\omega, \tau)$ for all $a' \in \mathcal{F}_n(\omega)$ via (43). Then, compute the ex ante value of not defaulting V_A^e and the decision density over feasible asset choices $\sigma_A^{(a',e)}$ conditional on effort e :⁴²

$$\begin{aligned}
 V_{A,n}^e(\omega, \tau) &= \alpha_{a'} \ln \left(\sum_{a' \in \mathcal{F}_n(\omega)} d_n^{a'}(\omega) \exp \left\{ \frac{v_{A,n}^{(a',e)}(\omega, \tau)}{\alpha_{a'}} \right\} \right) \\
 \sigma_{A,n}^{(a',e)}(\omega, \tau) &= \frac{d_n^{a'}(\omega) \exp \left\{ \frac{v_{A,n}^{(a',e)}(\omega, \tau)}{\alpha_{a'}} \right\}}{\sum_{\hat{a}' \in \mathcal{F}_n(\omega)} \exp \left\{ \frac{v_{A,n}^{(\hat{a}',e)}(\omega, \tau)}{\alpha_{\hat{a}'}} \right\}}.
 \end{aligned} \tag{A.3}$$

This expression for the ex ante value function is the closed form of (42) with the inclusion of the weights for the extreme value mean correction from Section A.1.

- ii. If $a(\omega) < 0$, compute the value of defaulting conditional on e , $V_{D,n}^e(\omega, \tau)$. Note that for consistency with the extreme value adjustment for the a' choice, rather than define the action-specific value of default via (41), we define it as

$$V_{D,n}^e(\omega, \tau) = \mathbb{E}_{\varepsilon_{a'}} \left[\max_{\{a' \geq 0 \mid c_n^{a'}(0, y(\omega), s(\omega)) > 0\}} \left\{ v_{D,n}^{(a',e)}(\omega, \tau) + \varepsilon_{a'} \right\} \right] \tag{A.4}$$

$$\begin{aligned}
 v_{D,n}^{(a',e)}(\omega, \tau) &= U \left(c_n^{a'}(0, y(\omega), s(\omega)) \right) \\
 &+ \beta_\tau \sum_{\tau'} Q_\tau(\tau'; \tau) \sum_{y'} Q_{y,n}^e(y') V_{E,n+1} \left(0, y', \Upsilon_n^{(0,1)}(y'; \omega), \tau' \right).
 \end{aligned} \tag{A.5}$$

⁴¹Note that the consumption for savings actions can be computed outside this loop since the price does not change as decision rules change, and the consumption conditional on default can be computed outside the loop as well.

⁴²In practice, to avoid numerical overflow or underflow, it is best to make the ex ante value and decision probability calculations as follows. First, identify the maximal value action among all feasible actions given the state: $\bar{v}_{A,n}^e(\omega, \tau) = \max_{a' \in \mathcal{F}_n(\omega)} v_{A,n}^{(a',e)}(\omega, \tau)$. Then, compute the ex ante value and decision rules as

$$\begin{aligned}
 V_{A,n}^e(\omega, \tau) &= \bar{v}_{A,n}^e(\omega, \tau) + \alpha_{a'} \ln \left(\sum_{a' \in \mathcal{F}_n(\omega)} d_n^{a'}(\omega) \exp \left\{ \frac{v_{A,n}^{(a',e)}(\omega, \tau) - \bar{v}_{A,n}^e(\omega, \tau)}{\alpha_{a'}} \right\} \right) \\
 \sigma_{A,n}^{(a',e)}(\omega, \tau) &= \frac{d_n^{a'}(\omega) \exp \left\{ \frac{v_{A,n}^{(a',e)}(\omega, \tau) - \bar{v}_{A,n}^e(\omega, \tau)}{\alpha_{a'}} \right\}}{\sum_{\hat{a}' \in \mathcal{F}_n(\omega)} \exp \left\{ \frac{v_{A,n}^{(\hat{a}',e)}(\omega, \tau) - \bar{v}_{A,n}^e(\omega, \tau)}{\alpha_{\hat{a}'}} \right\}}.
 \end{aligned}$$

In practice, then, the action-specific default value is similar to the ex ante no-default value, but with a different feasible set of a' choices. While we allow consumption to be less than $y(\omega)$, we impose that this current consumption reflects the discharge of the debt $a(\omega)$, and the continuation value still reflects that $a' = 0$ (there is no opportunity to save in default) and s' is still updated based on the default action.⁴³ Then, compute the inclusive value at the default stage,

$$\tilde{V}_{D,n}^e(\omega, \tau) = \alpha_d \ln \left(\exp \left\{ \frac{V_{A,n}^e(\omega, \tau)}{\alpha_d} \right\} + \exp \left\{ \frac{V_{D,n}^e(\omega, \tau)}{\alpha_d} \right\} \right)$$

and set the ex ante value at the default stage equal to

$$V_{D,n}^e(\omega, \tau) = \tilde{V}_{D,n}^e(\omega, \tau) - \eta_n^e(y(\omega), s(\omega), \tau)$$

Compute the probability of default

$$\sigma_{D,n}^{(d=1,e)}(\omega, \tau) = \frac{\exp \left\{ \frac{V_{D,n}^e(\omega, \tau)}{\alpha_d} \right\}}{\exp \left\{ \frac{V_{A,n}^e(\omega, \tau)}{\alpha_d} \right\} + \exp \left\{ \frac{V_{D,n}^e(\omega, \tau)}{\alpha_d} \right\}} \quad (\text{A.6})$$

and set $\sigma_{D,n}^{(d=0,e)}(\omega, \tau) = 1 - \sigma_{D,n}^{(d=1,e)}(\omega, \tau)$. If $a(\omega) \geq 0$, then set $V_{D,n}^e(\omega, \tau) = V_{A,n}^e(\omega, \tau)$.

- iii. Compute effort-specific (and beginning-of-period) value $v_{E,n}^e$ and ex ante value $V_{E,n}$ via (39), but with the ex ante value modified by the extreme value adjustment for default. Also compute the likelihood of choosing effort level e :

$$\begin{aligned} V_{E,n}(\omega, \tau) &= \alpha_e \ln \left(\sum_{e \in \{0,1\}} \exp \left\{ \frac{v_{E,n}^e(\omega, \tau)}{\alpha_e} \right\} \right) \\ \sigma_{E,n}^e(\omega, \tau) &= \frac{\exp \left\{ \frac{v_{A,n}^{(a',e)}(\omega, \tau)}{\alpha_{a'}} \right\}}{\sum_{\hat{e} \in \{0,1\}} \exp \left\{ \frac{v_{E,n}^{\hat{e}}(\omega, \tau)}{\alpha_e} \right\}}. \end{aligned} \quad (\text{A.7})$$

- (f) Compute the updated guess of the extreme value adjustment for default via

$$\tilde{\eta}_n^e(y, s, \tau) = \tilde{V}_{D,n}^e(a^*, y, s, \tau) - V_{A,n}^e(a^*, y, s, \tau), \quad (\text{A.8})$$

where a^* is the smallest debt level on the grid and the value functions used in the calculations are the ones updated this iteration.

- (g) Assess convergence of the fundamental decision rules by comparing the results computed

⁴³In practice, this adjustment is small since consuming less than $y(\omega)$ has no real benefit, but, it does make the comparison of default and no-default values “apples-to-apples” in light of the adjustment.

in (A.3), (A.6), and (A.7), as well as the updated default adjustment in (A.8), to the original guesses. If within tolerance, proceed to the next age, $n - 1$, otherwise, update the guesses of all four objects via relaxation and return to step 3(a).

4. For each age $n = 2, \dots, N$, solve forward for the equilibrium distribution of agents at each age, $\lambda_n(\omega, \tau)$, via (53). Note that the distribution at age 1, $\lambda_1(\omega, \tau)$ is exogenously specified. Along the way, compute any desired moments.

A.3 Additional formulas and derivations

A.3.1 No tracking model for Section 3

In the “no tracking” (NT) version of the model, type is unobservable, and lenders try to infer it, but this information is not tracked through time. Since types are permanent, this implies that all individuals have a common beginning-of-period type score equal to ρ in each age, though their assessment might be temporarily updated within age 1 based on their borrowing choice. In this case, we can suppress the state variable s_2 in the default decision at age 2, but it is still useful to define the type scoring function in equation (5), solved for in equation (17). The loan price is

$$q_{\text{NT},1}^{a_2} = \frac{1}{1+r} \left[\psi_1^{a_2} \bar{p}_{\text{NT},1}^{a_2}(H) + (1 - \psi_1^{a_2}) \bar{p}_{\text{NT},1}^{a_2}(L) \right]$$

$$\text{where } \bar{p}_{\text{NT},1}^{a_2}(\tau) = \sum_{y+2} Q_y(y_2) \sigma_{\text{NT},2}^{d_2=0}(a_2, y_2; \tau).$$

The default decision rules do not include the type assessments that weight the probabilities, since the type score state variable remains fixed at ρ . Since precluding tracking implies there is no scope for reputational considerations in the default decision at age 2, equation (15) is no longer relevant.

A.3.2 Full information quantitative model from Section 6

The observable state in this model is $\omega = (a, y, \tau)$. The decision problems are the same as in the baseline model, with the exception that since there is no role for reputation, there is no type score updating function. However, we must still account for the fact that effort is unobservable in this model. We can define the likelihood over all actions, $\sigma_n^{(a', d, e)}(\omega)$, the likelihood of observable actions, $\bar{\sigma}_n^{(a', d)}(\omega)$, and the adjusted probability of a given y' realization based on the observable choices, $\tilde{Q}_{y,n}^{(a', d)}(\omega)$, as in the baseline model above. We compute loan prices as

$$q_n^{a'}(\omega) = \frac{1}{1+r} \sum_{\tau', y'} Q_\tau(\tau'; \tau(\omega)) \tilde{Q}_{y,n}^{(a', d)}(\omega) \left(1 - \bar{\sigma}_{n+1}^{(0,1)}(a', y', \tau') \right). \quad (\text{A.9})$$