Credit Scores and Inequality Across the Life Cycle *

Satyajit Chatterjee † Dean Corbae ‡ Kyle Dempsey §

José-Víctor Ríos-Rull ¶

This draft: March 31, 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 John Leahy for valuable comments on an earlier version of this paper. 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 with respect to 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 in general a strong correlation between income and credit score, 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 in general 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 which 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. These shocks ensure that every budget feasible action is chosen with positive probability, much like in quantum response models (Goeree et al., 2016). An important consequence is that beliefs following any feasible action are well-defined and determined in equilibrium (reminiscent of Selten (1975) and Myerson (1978)): 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 softmax choice probability function familiar from discrete choice models and delivers computational tractability (Rust, 1987).

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 that 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 OCEAN[1] 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

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

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 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

[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/

[3] As 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.

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 its 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. 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 has banned unpaid medical bills from being included in credit reports. It also prohibits lenders from requesting information on medical debts from prospective borrowers.[6] 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 raises 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 aims to explore the links between consumption, income, and

[4]While 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.

[5]For 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.

[6]The ruling went into effect on January 7, 2025 (https://www.consumerfinance.gov/about-us/newsroom/cfpb-finalizes-rule-to-remove-medical-bills-from-credit-reports/).

reputational inequality. 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.

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 encountering a default 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 partially reveal. The inference based on these actions is summarized dynamically as probability assessments that resemble credit scores.

Of more 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) and 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 which approach that of a full information model. We conduct one such exercise.

While previous quantitative theory models imposed exogenous punishment, we incorpo-

[7] Kovrijnykh 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.

rate 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 (forthcoming). 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 some new facts about inequality (as measured by cross-sectional variation) in credit rankings. There we also propose a very simple reduced form "model" that can help us see what forces might be needed to explain these facts. We then lay out our structural framework in a series of sections. Section 3 starts by studying credit access and default with adverse selection in the simplest possible two-period version of Chatterjee et al. (2023). This section illustrates our methodological contribution to solving dynamic quantitative Bayesian equilibrium models with hidden information. Section 4 adds moral hazard to the basic framework. We do so to endogenize how unobservable type may interact with unobservable effort in order to study the interaction between reputation in the credit market and income. We add medical expense shocks to the basic adverse selection model in Section 5 in order to illustrate the costs and benefits of regulation of credit record information along the lines of the recent CFPB ruling. Section 6 puts all the pieces together in an N-period lived agent model to get closer to the data of Section 2 in order to make quantitative statements rather than qualitative ones.

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 Heathecote 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

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

(a) Mean, income and consumption

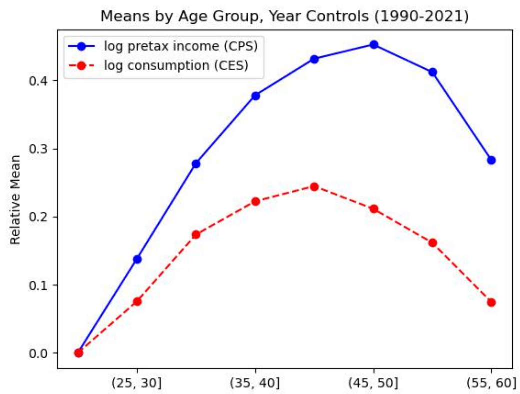

(b) Variance, income and consumption

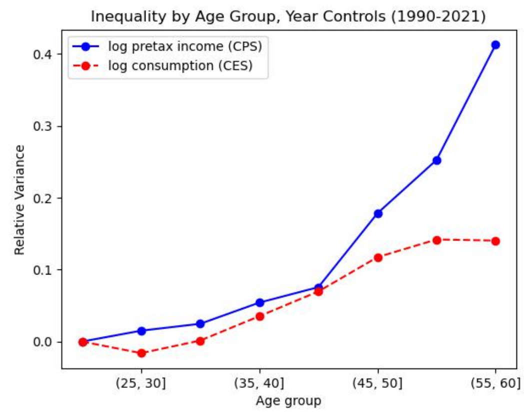

(c) Mean, credit ranking

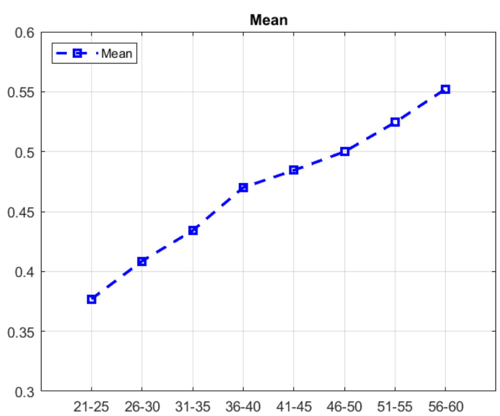

(d) Variance, credit ranking

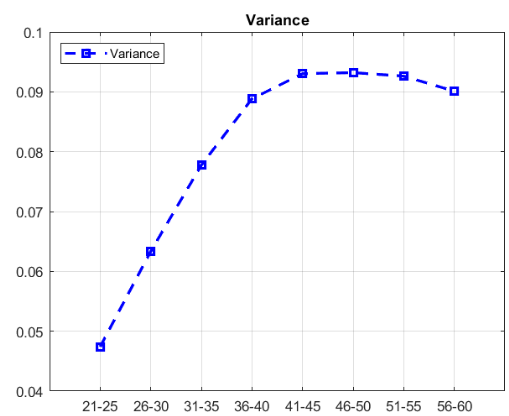

Source for panels (a) and (b): Retrieved from https://ideas.repec.org/c/red/ccodes/23-158.html. The source income data is the U.S. Census Bureau's Current Population Survey (CPS) and source expenditure data is the U.S. Bureau of Labor Statistics Consumer Expenditure Survey (CES), both for the sample 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 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. The income variable is an estimate of an individual's income provided by Equifax.

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 increases across age brackets and the dispersion of scores in a bracket around the bracket average also increases increases across age brackets.

The striking similarity between the patterns shown in the top and bottom panels of Figure 1 raises the question if 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.

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 dominates 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.

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

[9]More 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.

[10]For 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 2: **Credit rankings and income**

(a) Credit rankings by income groups

(b) Correlation between income and credit rankings

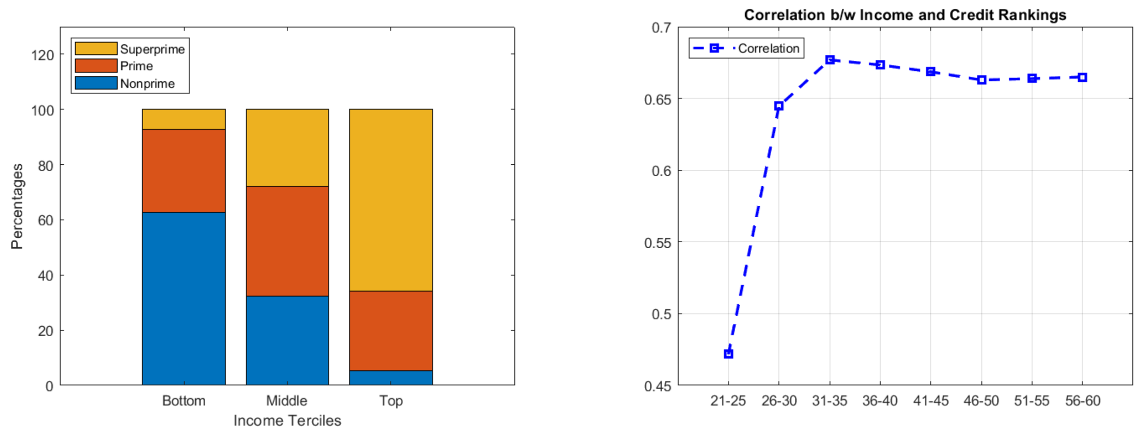

Source for panels (a) and (b): A 2 percent subsample of the FRBNY CCP/Equifax panel for the 3rd quarter of 2024. The credit rankings are based on VantageScore 4.0 credit score, which is a proprietary credit score of VantageScore Solutions LLC, similar to other credit scores used in the industry. Credit 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.

A second reason is displayed in Figure 2b. This figure 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 lacking as people haven't made many credit decisions. This shows up as a weak correlation between their income and credit rankings. People build up a credit history as time passes, and the correlation between income and credit rankings strengthens. Significantly, though, the correlation rises to only about 0.66, which reiterates the point made above that there are non-income factors that influence credit rankings.

We conclude this section by noting that repuational inequality has consequences for credit 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.[11] The plot shows that, conditional on an income quartile, interest rates decline with FICO scores. And, there is significant residual dispersion in APR based on incomes: For example, even conditional on a high FICO score of 790, the gap in APR between the highest and lowest income quartiles is 1.5

[11] The figure splits the data into 80 bins: 20 FICO bins × 4 income quartiles, and each data point corresponds to the within-bin average of the APR on the associated credit card account.

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

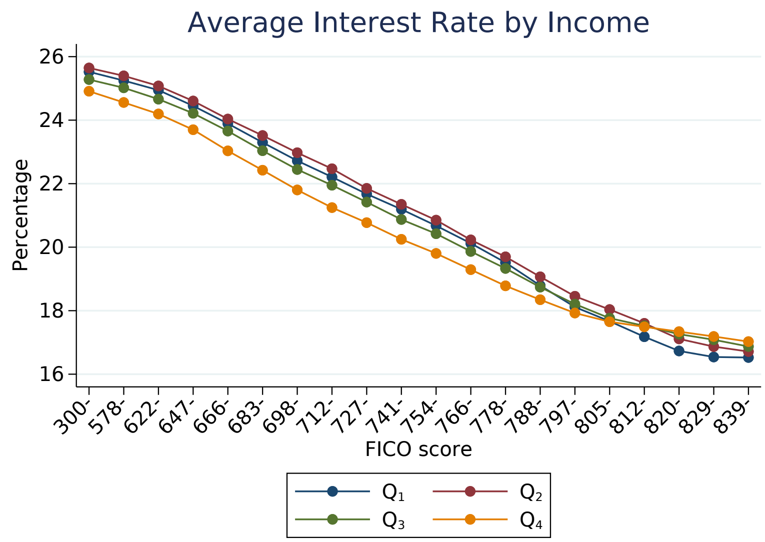

Source: Dempsey and Ionescu (2024). This figure is constructed using the Federal Reserve Board's Y-14M data as described in Appendix A of the source article.

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 a person'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 $U_{i,n}$. Specifically, suppose

$$S_{i,n} = \alpha Y_{i,n} + U_{i,n} \tag{1}$$

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:

$$\begin{aligned} \text{corr}(S_{i,n}, Y_{i,n}) &= \frac{\text{cov}(S_{i,n}, Y_{i,n})}{\text{sd}(S_{i,n})\text{sd}(Y_{i,n})} \\ &= \frac{\alpha\sigma_{Y_n}^2 + \sigma_{Y_n,U_n}}{\sigma_{Sn}\sigma_{Yn}} \end{aligned} \tag{2}$$

where $\sigma^2_{X_n}$ 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

$$\sigma_{S_n} = [\alpha^2\sigma^2_{Y_n} + 2\alpha\sigma_{Y_n,U_n} + \sigma^2_{U_n}]^{1/2} \tag{3}$$

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

What can we learn about the data through the lens of this simple 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 age life. The fact that observable credit rankings are rising early in life along with observable 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 which are correlated with repayment for a rising segment of the population.[12] That is, while $\overline{S}_n$ grows through time as a consequence of rising $\overline{Y}_n$ in equation (1) for early age brackets, it may continue to grow through learning about an individual's noisy unobservable type $\overline{U}_n$. In later sections of this paper we provide a structural model with adverse selection and noisy signals of unobservable type to consider this type of decomposition.

Notice that Figure 1b implies that σ_{Y_n} and Figure 1d implies that σ_{S_n} are both weakly increasing across higher age brackets in the data. Of course, equation (3) involves noisy signals about unobservables. If one believes in the ability of credit scorers to learn about unobservable traits like conscientiousness (e.g. who in the age bracket is patient and does not want to risk their reputation versus who in the age bracket is impatient and doesn't care much about their reputation), then $\sigma^2_{U_n}$ will be increasing with age brackets since scorers are able to separate individuals on the basis of risk by their credit market actions. Further, if unobservable conscientiousness also affects effort choices (say through moral hazard), we might expect σ_{Y_n,U_n} to be positively correlated. Signing how σ_{Y_n,U_n} changes over time is harder but something that may be understood through the lens of a structural model that we take up in Section 4. This also brings up the challenge of endogeneity and causal identification, which we address as well in Section 4.

Turning to Figure 2b, consider the simple case where there is no effort choice and the earnings process is unaffected by unobservable type as in (Chatterjee et al., 2023) so that $\sigma_{Y_n,U_n} = 0$, an assumption we make in our simplest adverse selection model of Section 3. In

[12] That is, the correlation could arise if the proportion of less risky types grows with age brackets due possibly to more risky types having lower survival probabilities over time.

that case, equation (2) becomes

$$\operatorname{corr}(S_{i,n}, Y_{i,n}) = \frac{1}{\left[1 + \frac{\sigma^2_{U_n}}{\alpha^2 \sigma^2_{Y_n}}\right]^{1/2}}. \tag{4}$$

In this case, (4) implies the correlation between observed credit rankings and income in a given age bracket is increasing in the variance of earnings observed in Figure 1b and decreasing in the variance of the noisy signal. In that case, we can explain the correlation in Figure 2b by the variance of income ($\sigma^2_{Y_n}$) in Figure 1b rising faster than the variance of the signal of the unobserved type ($\sigma^2_{U_n}$) in early age brackets with $\sigma^2_{U_n}$ catching up to $\sigma^2_{Y_n}$ in later age brackets. This might suggest credit scorers are slow to learn early in borrower's lives and strong separation later in life.

While the preceding logic (i.e. both $\sigma^2_{U_n}$ and $\sigma^2_{Y_n}$ rising across age brackets) is consistent with the correlation in Figure 2b, it is not necessarily consistent with the behavior of the variance of credit rankings in equation (3) under the assumption that $\sigma_{Y_n,U_n} = 0$ and Figure 1d. Specifically while it works to explain the rising variance of credit rankings for early age brackets, the fact that the variance of credit rankings levels off in late age brackets would require a negative covariance of income and unobservable characteristic. If, for instance, effort happens today and the earnings payoff happens in the future, but one will be dead in the future, then there is no reason to exert effort, which may induce a negative correlation. We take up the more complicated case where unobservable characteristics affect earnings across age buckets in our structural model.

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 $n \in \{1, 2\}$.[13] There is a unit measure of two types $\tau \in \{H, L\}$ of individuals who differ in their patience where $\beta_H > \beta_L$. The fraction of type $\tau = H$ in the population is ρ. Importantly, in this adverse selection model type is unobservable. Individuals are risk averse with increasing and concave preferences $U(c_n)$.

[13]Section 6 extends the number of periods agents live beyond 2.

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 $\mathbb{P}(y_2)$ with $y_\ell < y_h$.[14]

An asset market in unsecured discount bonds opens at $n = 1$ where individuals can choose $a_2 \in A \equiv \{\underline{a}, ..., 0, ..., \overline{a}\}$ at price q_1. Let $A_- = \{a_2 \in A : a_2 < 0\}$ and $A_+ = A \backslash A_-$. Competitive lenders receive no recovery if a borrower defaults $d_2 \in \{0, 1\}$ on their debt $a_2 < 0$ at $n = 2$. Competitive lenders discount the future at rate r. Individuals begin $n = 1$ with no assets (i.e. $a_1 = 0$). 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, a decision to 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 a default choice $d_2 = 1$, 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 is the individual's asset market choice a_2. Free entry implies the competitive price is thus given by

$$q_1(a_2) = \begin{cases} \frac{1}{1+r} & \text{if } a_2 \geq 0 \\ \frac{\mathbb{P}(d_2=0 \mid a_2)}{1+r} & \text{if } a_2 < 0 \end{cases}. \tag{5}$$

Notice that prices in our environment satisfy what is called "block recursivity" in Menzio and Shi (2010). Specifically, under certain conditions on the information structure, our menu of nonlinear competitive prices conditional on asset choice 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 individual states in order to price debt in our case nor solve for equilibrium allocations in the search framework. This is a general result in the nonlinear pricing models of default like that in the baseline model of Chatterjee et al. (2007) and (Chatterjee et al., 2023), except when information restrictions exist where given the lack of individual information, it is necessary for the lender to use the economy-wide cross-sectional distribution to form her prior.[15]

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

$$\psi_1(a_2) \equiv \mathbb{P}(\tau = H \mid a_2) \tag{6}$$

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, $s_1 = \rho$, is given by the population measure of H types.

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

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

At the beginning of $n = 2$, individuals receive an income realization y_2 drawn from $\pi(y_2 \mid y_1)$ 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 which can be used to form an updated assessment at $n = 2$. Those with the opportunity to default (i.e. $a_2 < 0$) have their scores updated conditional on beginning-of-period observables (a_2, y_2, s_2) according to

$$s_3 = \psi_2(d_2; a_2 < 0, y_2, s_2) \equiv \mathbb{P}(\tau = H \mid d_2, a_2, y_2, s_2) \tag{7}$$

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, their type score is simply retained $s_3 = s_2$.

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 (forthcoming), 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 which is increasing in their posterior assessment due to enhanced credit access.

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

To summarize timing, we have:

1. At $n = 1$,
   (a) an individual chooses a_2 taking the price schedule $q(a_2)$ 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 $\mathbb{P}(y_2 \mid y_1)$.
   (b) if $a_2 < 0$, d_2 choice is made.
   (c) observing d_2, assessments of an individual's type $s_3 = \psi_2(d_2; a_2 < 0, y_2, s_2)$ or $s_3 = s_2$ if $a_2 \geq 0$ are updated according to Bayes' Law.

3.2 Perfect Bayesian equilibrium

We now describe the optimization problems faced by individuals and lenders. The individual problem in period $n = 1$ is

$$\begin{aligned} V_1(\tau) \quad &= \quad \max_{a_2 \in A} U(c_1) + \beta_\tau \mathbb{E}_{y_2|y_1}\left[V_2(a_2, y_2, s_2; \tau)\right] \\ \text{subject to:} \quad & \quad c_1 = y_\ell - q(a_2)a_2 \end{aligned} \tag{8}$$

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

$$\begin{aligned} V_2(a_2, y_2, s_2; \tau) \quad &= \quad \max_{d_2 \in \{0,1\}} U(c_2) + \beta_\tau \phi(s_3) \\ \text{subject to:} \quad & \quad c_2 = y_2 + (1 - d_2)a_2 \end{aligned} \tag{9}$$

with the optimal policy $d_2^*(a_2, y_2, s_2; \tau)$ (if $a_2 \geq 0$ then $d_2^* = 0$ trivially). If $a_2 < 0$, the solution to (9) is simple: $d_2^*(a_2 < 0, y_2, s_2; \tau) = 0$ if

$$U(y_2 + a_2) + \beta_\tau \phi\left(\psi_2(d_2 = 0; a_2 < 0, y_2, s_2)\right) \geq U(y_2) + \beta_\tau \phi\left(\psi_2(d_2 = 1; a_2 < 0, y_2, s_2)\right). \tag{10}$$

Consider some particular level of debt $a_2 = -a$. Then, in terms of signaling theory, we can write (10) as:

$$\phi\left(\psi_2(d_2 = 0; -a, y_2, s_2)\right) - \phi\left(\psi_2(d_2 = 1; -a, y_2, s_2)\right)) \geq \frac{U(y_2) - U(y_2 - a)}{\beta_\tau} \tag{11}$$

where the left hand side of (11) 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 then $U(y_h) - U(y_h - a) < U(y_\ell) - U(y_\ell - a)$ 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

The issue: We now explain a methodological contribution of our framework which has relevance for the empirical contribution of our paper. We illustrate existence issues in our Bayesian equilibrium working backwards on one portion of the individual's problem: the decision to default in (9) on a particular level of debt $a_2 = -a$.

As discussed above, (11) implies that default is more likely to happen 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 receive the same income y_2 and conjecture the following behavior:

$d_2^*(-a, y_2, s_2; \tau)$	$\tau = H$	$\tau = L$
y_h	0	0
y_ℓ	0	1

(12)

That is, suppose 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:

$$\psi_2(d_2 = 1, -a, y_h, s_2) = \frac{0 \cdot s_2}{0 \cdot s_2 + 0 \cdot (1 - s_2)} = \frac{0}{0}, \tag{13}$$

$$\psi_2(d_2 = 0, -a, y_h, s_2) = \frac{1 \cdot s_2}{1 \cdot s_2 + 1 \cdot (1 - s_2)} = s_2. \tag{14}$$

It is evident from (13) 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:

$$s_3^{\text{oep}}(d_2 = 1, -a, y_h, s_2) = s_2. \tag{15}$$

Note that since there is separation when $y_2 = y_\ell$, type is perfectly revealed in the low income state.[16]

A fundamental question arises. Is this conjectured behavior which is similar to what we see in the data – no default in high income states with infrequent (rate $1 - \rho$) default in low income states – consistent with a Bayesian equilibrium given what are often considered "reasonable" oep beliefs?

In the $y_2 = y_\ell$ case where beliefs are well defined because types are separated, for the

[16] Formally, the posteriors are formed according to:

$$\psi_2(d_2 = 1, -a, y_\ell, s_2) = \frac{0 \cdot s_2}{0 \cdot s_2 + 1 \cdot (1 - s_2)} = 0$$

$$\psi_2(d_2 = 0, -a, y_\ell, s_2) = \frac{1 \cdot s_2}{1 \cdot s_2 + 0 \cdot (1 - s_2)} = 1$$

high type to repay and low type to default, we need, respectively:

$$\phi\left(\psi_2(d_2=0,-a,y_\ell,s_2)\right)-\phi\left(\psi_2(d_2=1,-a,y_\ell,s_2)\right)=\phi(1)-\phi(0)\geq\frac{U(y_\ell)-U(y_\ell-a)}{\beta_H}$$
$$\phi\left(s_2(d_1=0,-a,y_\ell,s_1)\right)-\phi\left(s_2(d_1=1,-a,y_\ell,s_1)\right)=\phi(1)-\phi(0)<\frac{U(y_\ell)-U(y_\ell-a)}{\beta_L}.$$

Both conditions are possible under certain parameterizations with sufficiently low $\nu_L\beta_L$.

For the troublesome $y_2=y_h$ case where agents are pooled and not all feasible actions are taken in equilibrium, in order for both types to repay as conjectured we need:

$$\phi\left(\psi_2(d_2=0,-a,y_h,s_2)\right)-\phi\left(\psi_2(d_2=1,-a,y_h,s_2)\right)=\phi(s_2)-\phi(s_3^{\text{oep}})\geq\frac{U(y_h)-U(y_h-a)}{\beta_L}\tag{16}$$

So what conditions on oep beliefs in $\phi(s_3^{\text{oep}})$ make this an equilibrium? If we take the oep belief in (15) we have $s_3^{\text{oep}}=s_2$ in which case (16) is given by:

$$\phi(s_2)-\phi(s_2)=0<\frac{U(y_h)-U(y_h-a)}{\beta_\tau}\text{ for all }\tau\tag{17}$$

In that case, the standard assumption for oep beliefs results in *non-existence* of our conjectured Bayesian equilibrium where risky types default in low earnings states.

What happens if we use 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 we need:

$$\begin{aligned}\phi\left(\psi_2(d_2=0,-a,y_h,s_2)\right)-\phi\left(\psi_2(d_2=1,-a,y_h,s_2)\right)&=\phi(s_2)-\phi(s_3^{\text{oep}})\\&\geq\frac{U(y_h)-U(y_h-a)}{\beta_L}\end{aligned}\tag{18}$$

What conditions on oep beliefs in $\phi(s_3^{\text{oep}})$ make this an equilibrium? In the case where $s_3^{\text{oep}}=0$, a necessary condition to make this an equilibrium is given in (18) by $\phi(s_2)-\phi(0)\geq[U(y_h)-U(y_h-a)]/\beta_L$. Thus, with this specification of oep beliefs it may be possible to find a set of parameters to support our conjectured Bayesian equilibrium. On the other hand, 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.

The fix: 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 choice over assets and default are subject to shocks $(\epsilon^{a_2}, \epsilon^{d_2})$ drawn from a Type I extreme value distribution which are unobservable to anyone (e.g. econometricians performing risk assessments) except the individual.[17]

How should we interpret the shocks? There are numerous shocks 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). And how do they help fix the non-existence problem described above emanating from an undefined posterior in (13) in state $y_2 = y_h$ where an individual has debt $-a$? Introducing these shocks eliminates perfect pooling (i.e. $d_2^*(-a, y_h, s_2; \tau) = 0, \forall \tau$); some of each type will draw unobservable (to the econometrician assessing their risk) shocks for which they choose to default. The choice probabilities (and hence fractions of people choosing those actions) will depend on their type τ. Hence, the matrix of conjectured actions in (12) will be populated by probabilities in $(0, 1)$ not 0 or 1. Alternatively, the presence of such unobservable shocks implies that there is not perfect separation of types thus slowing down learning about agent's unobservable type. Without such shocks, perfect separation would imply that after the initial asset choice, there would be no need for updated type scores.

This change to the environment requires us to change the optimization problems in (8) and (9). In the case of (9), 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:

$$V_2(a_2, y_2, s_2; \tau) = \mathbb{E}_{\epsilon^{d_2}} \left[\max_{d_2 \in \{0,1\}} \left\{ U(y_2 - (1 - d_2)a_2) + \varepsilon^{d_2} + \beta_\tau \phi\Big(\psi_2(d_2, a_2, y_2, s_2)\Big) \right\} \right] \quad (19)$$

Letting $V_2^{d_2}(a_2, y_2, s_2; \tau)$ denote the value of choosing a given action d_2, the solution to (19) yields choice probabilities given by, for the no-default case:

$$\sigma^{d_2=0}(-a, y_2, s_2; \tau) = \frac{\exp\left(\frac{V_2^{d_2=0}(-a, y_2, s_2; \tau)}{\alpha_d}\right)}{\sum_{d_2} \exp\left(\frac{V_2^{d_2}(-a, y_2, s_2; \tau)}{\alpha_d}\right)} \quad (20)$$
$$= \left[1 + \exp\left[\frac{U(y_2) + \beta_\tau \phi\left(\psi_2(d_2 = 1, -a, y_2, s_2)\right) - U(y_2 - a) - \beta_\tau \phi\left(\psi_2(d_2 = 0, -a, y_2, s_2)\right)}{\alpha_d}\right]\right]^{-1}.$$

Then, it is simple to see from the last line of (20) that as long as $\phi'(s_3) > 0$, type H

[17]Recall the Type 1 extreme value cumulative distribution function for discrete choice $x \in X$ is given by

$$F_\epsilon(\epsilon^x; \overline{x}, \alpha_x) = \exp\left\{ -\exp\left(-\frac{\epsilon^x - \overline{x}}{\alpha_x} \right) \right\}$$

where $\overline{x}$ is the location parameter and α_x is the scale parameter which governs its variance.

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 (20) 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 (13). Specifically, posteriors are now given by

$$\psi_2(d_2, -a, y_2, s_2) = \frac{s_2 \sigma^{d_2}(-a, y_2, s_2; H)}{s_2 \sigma^{d_2}(-a, y_2, s_2; H) + (1 - s_2)\sigma^{d_2}(-a, y_2, s_2; L)}.$$

The fact that $V_2(a_2, y_2, s_2; \tau)$ is increasing in β_τ implies $\sigma^{d_2=0}(-a, y_2, s_2; H) > \sigma^{d_2=0}(-a, y_2, s_2; L)$ which yields $\psi_2(d_2 = 0, -a, y_2, s_2) > \psi_2(d_2 = 1, -a, y_2, s_2)$.

3.4 Model properties

This subsection illustrates the equilibrium implications of an individual's incentives to acquire reputation given their hidden type. In what follows, all timing and information assumptions remain as in the environment except for the introduction of unobservable shocks $(\epsilon^{a_2}, \epsilon^{d_2})$. Those shocks to individuals' discrete choices result in choice probabilities $\sigma_1^{a_2}(\tau)$ and $\sigma_2^{d_2}(a_2, y_2, s_2; \tau)$ which 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.

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. Since these exercises are numerical illustrations, we relegate the parameterization of the model to Appendix B.1.

3.4.1 On adverse selection

Since the extent of adverse selection depends on how fundamentally different the two types are, Figure 4 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 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.

[18]That is, $U(y_2) + \beta_H \phi(\psi_2(d_2 = 1, -a, y_2, s_2)) - U(y_2 - a) - \beta_H \phi(\psi_2(d_2 = 0, -a, y_2, s_2)) < U(y_2) - \beta_L \phi(\psi_2(d_2 = 1, -a, y_2, s_2)) - U(y_2 - a) - \beta_L \phi(\psi_2(d_2 = 0, -a, y_2, s_2))$ if and only if $\phi(\psi_2(d_2 = 1, -a, y_2, s_2)) < \phi(\psi_2(d_2 = 0, -a, y_2, s_2))$.

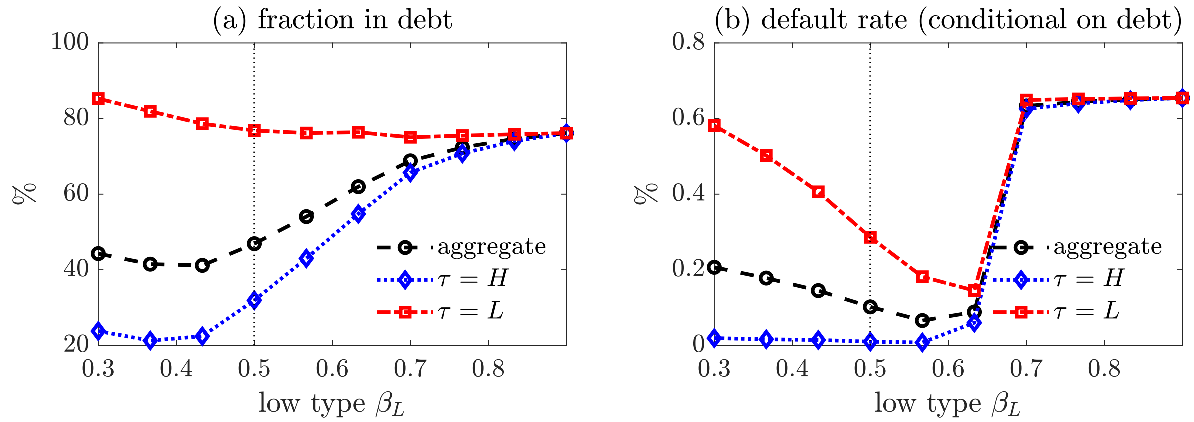

Figure 4: **Adverse selection affects borrowing and default behavior**

Notes: Both panels present moments of the model for a range of values of β_L between $\underline{\beta}$ 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.

When the types are fairly close, borrowing behavior is similar, but the gap in the share of each type borrowing grows as the distance between the types grows. 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 4(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 the parameter ϕ: households do not value reputation at all when $\phi = 0$, while they value reputation a lot when $\phi = \overline{\phi}$. Figure 5 documents how strengthening the incentives to acquire a good reputation changes economic outcomes.

When the value of reputation ϕ is zero, savings and default choices are not affected by reputational considerations, but this changes as ϕ increases. Since H-types save more, saving improves one's reputation: to the extent that reputation is valued, this increases savings rates of both types (panel (a)). The only differences in default behavior at age 2 (panel (b)) come

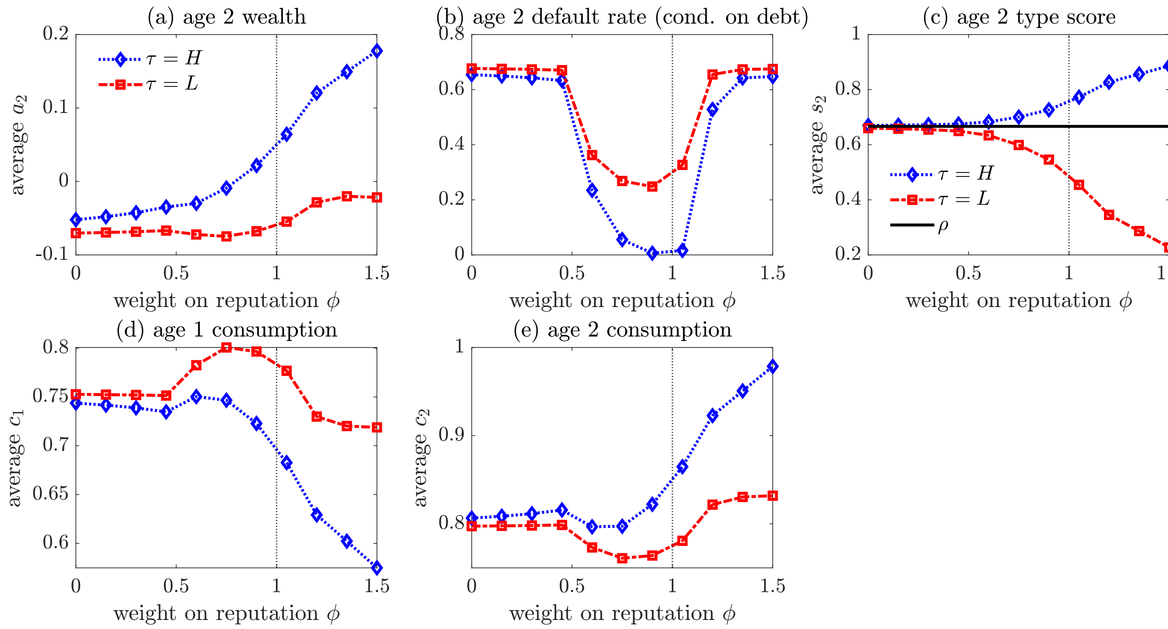

Figure 5: **Adverse selection affects borrowing and default behavior**

Notes: Each panel presents the moment indicated in the title over a range of values of ϕ for both the high and low 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.

from reputational effects. As the weight on reputation increases, H-types also cut their default rates. 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 induces an expansion of credit access: since both types default less, price schedules shift up and borrowing terms improve. This causes a substitution effect which makes the net effect on savings ambiguous: panel (a) shows that the first effect always dominates for H-types, while the two effects approximately cancel for L-types. As reputation becomes even more valued, 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 the value of reputation increases, the H-types who value reputation more increasingly separate themselves from L-types with their savings decisions. L-type behavior changes similarly, but less than one-for-one. While L-types value reputations and take steps to preserve it 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, then, the equilibrium features more (though far from complete) separation.

What are the implications of these dynamics for inequality over the life cycle? Panels (d) and (e) of Figure 5 show that as the value of reputation increases, consumption inequality increases alongside wealth inequality, *despite the fact that income inequality is unchanged*

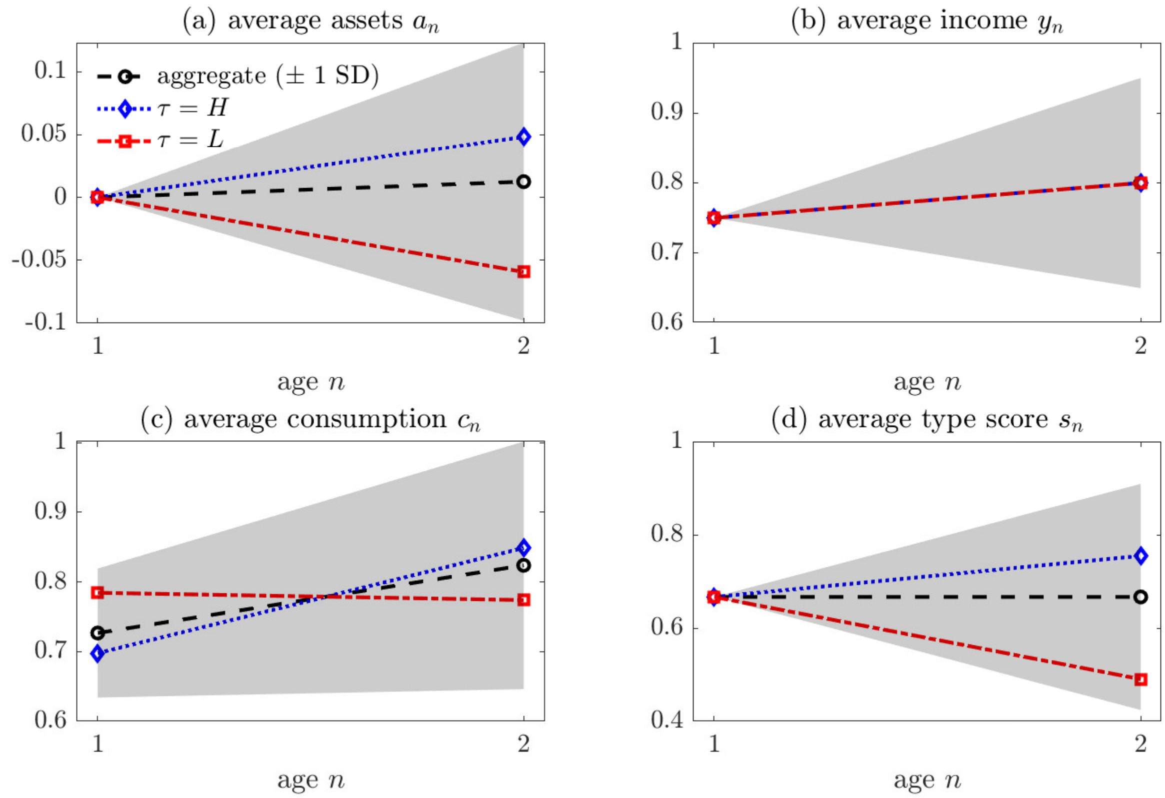

Figure 6: **Life cycle in the baseline model**

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.

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 a key mechanism in our analysis: reputational inequality is directly related to other forms of inequality which are more commonly measured.

3.4.2 Life cycle patterns

Figure 6 plots the evolution of key variables in the aggregate and across types. 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. At a high level, the figure shows that even this simple model can make sharp predictions about the evolution of inequality over the life cycle.

Panels (a) and (b) plot the resources for consumption: wealth and income, respectively. All agents begin with the same income $y_1 = y_\ell$, and income increases on average over the life cycle exogenously. While ex post some individuals are "lucky," receiving $y_2 = y_h$, while others are "unlucky," and this generates an increase in income inequality. Notably, though, since income does not depend on type, there is no difference in income across types.

Saving decisions, however, do differ across types. On average in this economy, individuals save modestly despite their expected income growth, due to reputational incentives. Underlying this aggregated result, though, is divergence between the types: H-types tend to save more, while L-types borrow a lot. Panel (c) presents a corollary of this divergence: there is meaningful consumption inequality *across types* at both ages. In particular, the impatient L-types consume more than patient H-types at age 1, while this pattern flips at age 2.

Panel (d) shows how these behaviors impact the evolution of reputation. At age 1, all individuals have a common type score, but the divergence in asset choices leads to divergence in the updated assessments s_2. Since H-types tend to save on average, and since saving is more likely among H-types 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.

3.4.3 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 a consumer's credit report after 10 years. From the perspective of our model, the factors which 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.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

[19] In 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.

	baseline			no tracking		
	mean	diff, $H-L$	std. dev.	mean	diff, $H-L$	std. dev.
Panel 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
Panel 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.

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 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 which 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 6(b). In that case, consumption inequality (Figure 6(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 which 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 (forthcoming). Thus, income inequality may be induced by different effort choices across the types. 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 may be directly informative about one's type.

4.1 Environment

We consider an environment that is identical to that in Section 3.1 with one key difference. All individuals start age 1 with the same income y_ℓ in both cases. Here, though, individuals make an unobservable effort choice e_1 at the beginning of age 1 which raises the likelihood of a receiving y_h at the beginning of age 2 via the endogenous earnings process $\mathbb{P}(y_2 \mid e_1)$. This replaces the exogenous earnings process $\mathbb{P}(y_2)$ 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, $\mathbb{P}(y_h \mid e_1 = 1) > \mathbb{P}(y_h \mid e_1 = 0)$. 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 of one's assessed type before the age 2 default decision. Instead, an individual's type score at age 2 depends not only on their asset choice, but on the realization of their income at age 2, according to the function

$$s_2 = \Upsilon_2(a_2, y_2) = \mathbb{P}(\tau = H \mid a_2, y_2) \tag{21}$$

This formula extends equation (6) from Section 3.1 to account for the fact that different types may have different choice probabilities $\sigma^{e_1}(\tau)$ 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, as before an individual's type (τ) is unobservable (i.e. adverse selection). But now we add an unobservable effort (e_1) choice (i.e. moral hazard) which affects next period's observable income (y_2). While credit market actions (a_2, d_2) are observable, all actions (e_1, a_2, d_2) are subject to unobservable idiosyncratic shocks ($\epsilon^{e_1}, \epsilon^{a_2}, \epsilon^{d_2}$). As before, income ($y_n$) and type scores ($s_n$) are observable providing noisy signals of the unobservable type.

To summarize the new timing:

1. At $n = 1$,
   (a) an individual chooses unobservable effort e_1.
   (b) an individual chooses a_2 taking the price schedule $q(a_2)$ as given.
2. At $n = 2$,
   (a) y_2 is realized from $\mathbb{P}(y_2 \mid e_1)$.
   (b) scorers update their assessment according to Bayes' Law via $s_2 = \Upsilon_2(a_2, y_2)$.
   (c) if $a_2 < 0$, d_2 choice is made.
   (d) observing d_2, individual type score $s_3 = \psi_2(d_2; a_2 < 0, y_2, s_2)$ is updated according to Bayes' Law (or $s_3 = s_2$ if $a_2 \geq 0$).

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

As in the baseline model, in this version of the model we associate extreme value preference shocks to all the relevant decisions: effort, asset choice, and default. At age 1, the individual first makes an effort choice, weighing the cost of effort against the continuation value associated with exerting that effort:

$$V_1(\tau) = \mathbb{E}_{\epsilon^{e_1}} \left[\max_{e_1 \in \{0,1\}} \Big\{ -\kappa e_1 + \epsilon^{e_1} + W_1(\tau, e_1) \Big\} \right], \tag{22}$$

where the continuation value reflects the second decision over assets:

$$\begin{aligned} W_1(\tau, e_1) \quad &= \quad \mathbb{E}_{\epsilon^{a_2}}\left[\max_{a_2}\left\{U(c_1)+\epsilon^{a_2}+\beta_\tau \mathbb{E}_{y_2|e_1} V_2(a_2, y_2, \Upsilon_2(a_2, y_2); \tau)\right\}\right] \\ \text{subject to:} \quad &\quad c_1 = y_\ell - q(a_2)a_2 \end{aligned} \tag{23}$$

The first decision problem generates the type-specific decision rule over effort $\sigma^{e_1}(\tau)$, while the second generates the effort-choice and type-specific decision rule over assets $\sigma^{a_2}(e_1, \tau)$. 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 a default decision:

$$\begin{aligned} V_2(a_2, y_2, s_2; \tau) \quad &= \quad \mathbb{E}_{\epsilon^{d_2}}\left[\max_{d_2\in\{0,1\}}\left\{U(c_2)+\epsilon^{d_2}+\beta_\tau \phi\left(s_3\right)\right\}\right] \\ \text{subject to:} \quad &\quad c_2 = y_2 + (1-d_2)a_2 \end{aligned} \tag{24}$$

The decision rule associated with this choice problem is $\sigma^{d_2}(a_2, y_2, s_2; \tau)$.

Note that type scores and prices are computed similarly to the baseline model from Section 3. In particular, the loan price is still defined as in (5) as a function of the probability of repayment conditional on the observation of a_2. It can be useful to define an "interim" type score as in (6) which serves as an input to this pricing function, but does not govern the transition of an individual's type score state variable, which is now governed by (21). This interim type score also serves as an input to the individual's age 2 type score. See Appendix A.2 for details and explicit formulas.

4.3 Model properties

4.3.1 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 7 repeats the exercise of Figure 4 – varying the fundamental difference between the two types – but in the richer model with moral hazard.

Figure 7(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 4(a).

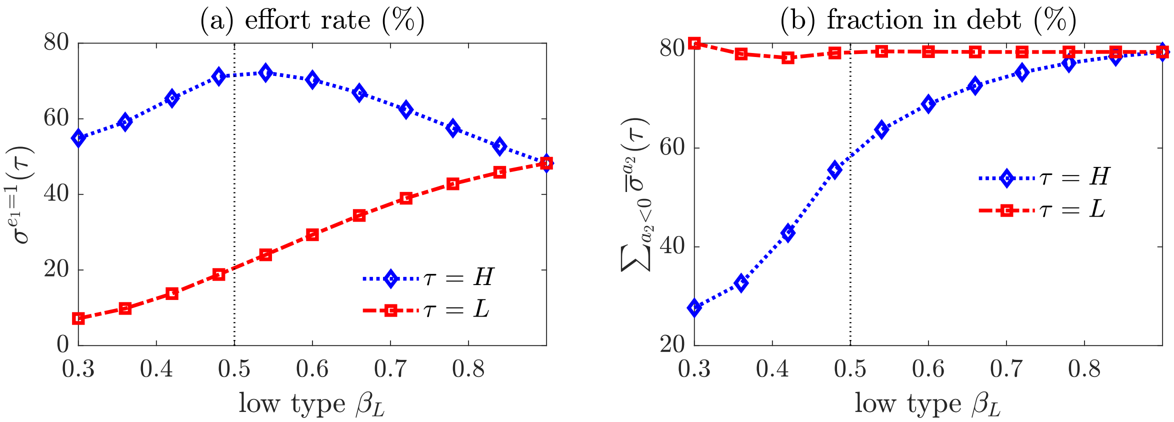

Figure 7: **The interaction between moral hazard and adverse selection**

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.

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 8 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 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 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

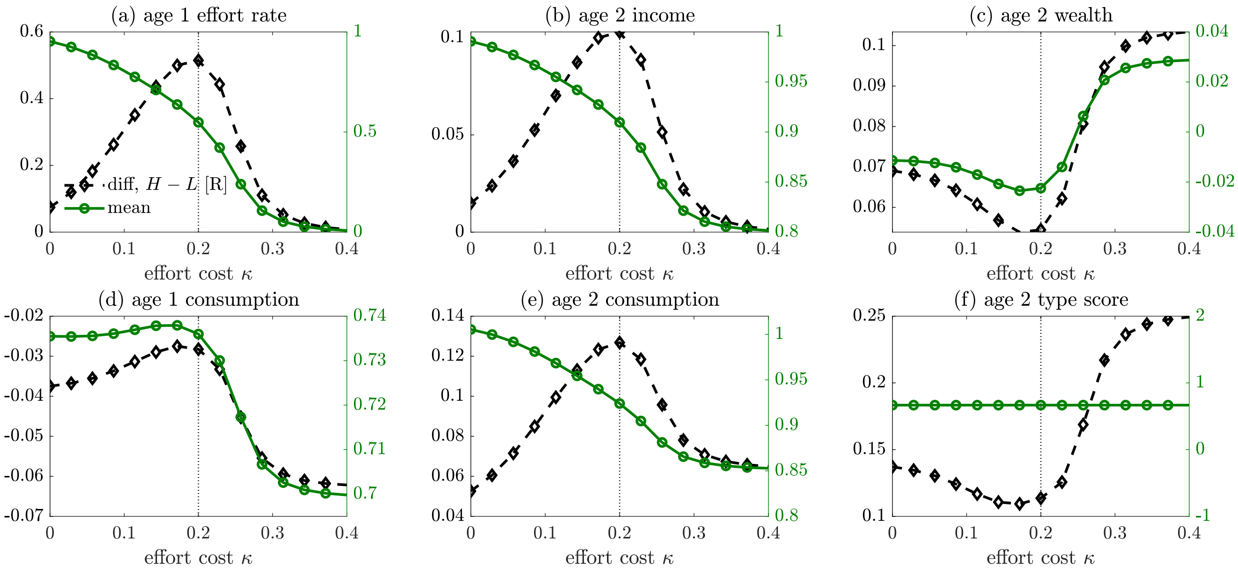

Figure 8: **Varying the strength of the moral hazard friction**

Notes: This figure shows a range of outcomes across a range of levels of the effort cost, κ_e, 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 the indicated metric between H- and L-types on the right axis (black). The vertical line in each graph corresponds to the benchmark parameter value for κ_e used throughout the numerical illustrations.

baseline model. Panel (e) reflects these same force 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).

4.3.2 Causal identification

An important question to consider when assessing any credit policy is: how will it affect individuals' terms of credit by changing how reputations are assessed via credit scores? In this subsection, we argue that our framework is well-suited to address this question by allowing for a crucial decomposition: how much of the adjustment in an individual's credit score is attributable to his own actions, conditional on his state, versus the outcomes of random events. Put differently, to what extent do individuals "make their own luck" in terms of their perceived creditworthiness? In this section, we use the models from Sections 3 and 4 to formalize how our model can be used to assess these two channels in the context of savings and default actions versus income shocks, but the basic idea can be extended to other contexts (for example, medical shocks in Section 5).

Section 3 makes an identifying assumption that earnings are exogenous. In this context, one's income is uninformative about one's type. In Section 4, however, there is a *causal* link from type to income: since effort (i) affects the likelihood of each income realization and (ii) is (generically) taken at different rates by type, income is informative about one's type.

Our model, then, suggests a simple approach to measuring how much of credit score adjustment is within individual's span of control. Suppose we observed credit rankings χ, income y, asset market choices (d, a'), and a vector of observable states ω for individuals i at time t. Suppose further that two individuals, i and j, begin period t with the same state ($\omega_{it} = \omega_{jt}$) and credit ranking ($\chi_{it} = \chi_{jt}$) and make the same choice ($(d_{it}, a_{it+1}) = (d_{jt}, a_{jt+1})$), but receive different income realizations at date $t+1$ ($y_{it+1} \neq y_{jt+1}$).[20] Then, the difference $\chi_{it+1} - \chi_{jt+1}$ measures the change in credit score attributable solely to the income realization. These differences could then be averaged over the distribution of individuals to compute an aggregate metric, and compared to an analogous number computed by holding future income realizations fixed across individuals but allowing their asset market choices to vary. Moreover, such metrics could be used to inform key parameters in our model, such as the noisiness of effort and asset market choices, the utility cost of exerting effort, or the shift in the income distribution between effort choices.

4.3.3 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.[21] 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 which were 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 which could be

[20] Of course, a practical implementation would have to "match" individuals along these dimensions, perhaps by binning.

[21] Note that the analog of Figure 6 is presented in the appendix in Figure B.3.

	moral hazard benchmark			moral hazard, no tracking		
	mean	diff, $H-L$	std. dev.	mean	diff, $H-L$	std. dev.
Panel 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
Panel 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.

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 what we considered for the baseline model in Section 3.4.3. The results for wealth and consumption broadly mirror those from Table 1: when there is an incentive to acquire reputation, wealth and consumption inequality in the aggregate and across types both increase. 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.[22] 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.

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 $m_1 \in \{0, \overline{m}\}$ with probability $\zeta = \mathbb{P}(m_1 = \overline{m})$. We assume these shocks are independent of type τ, consistent with the idea that while 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, and 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 $y_\ell > |\underline{a}| + \overline{m}$ 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

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

asset choice and the observable component of the medical record: that is, and agent's type score evolves according to

$$s_2 = \Gamma_1(a_2, \mu_1) \equiv \mathbb{P}(\tau = H \mid a_2, \mu_1) \tag{25}$$

Equation (25) is the analog of equation (6) from the baseline model and equation (21) 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 ($\epsilon^{\delta_1}, \epsilon^{a_2}, \epsilon^{d_2}$). As before, type scores (s_n) are observable.

To summarize the new timing:

1. At $n = 1$,
   (a) each individual's the medical expense shock m_1 is realized.
   (b) if $m_1 = \overline{m}$, the individual makes a delinquency decision δ_1.
   (c) each individual chooses a_2 taking the price schedule $q(a_2, \mu_1)$ 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 $\mathbb{P}(y_2 \mid y_1)$.
   (b) if $a_2 - b_2 < 0$, the individual makes a default choice d_2
   (c) observing d_2, individual type score $s_3 = \psi_2(d_2; a_2, b_2, y_2, s_2)$ is updated according to Bayes Law or $s_3 = s_2$ if $a_2 - b_2 \geq 0$.

Thus, the only differences from the timing in Section 3.1 is 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:

$$V_1(\tau) = \zeta J_1(\tau) + (1 - \zeta) W_1(0, 0; \tau) \tag{26}$$

where $W_1(m_1, \delta_1; \tau)$ is the value function at the asset choice stage and $J_1(\tau)$ is the value function at the medical delinquency choice stage conditional on $m_1 = \overline{m}$:

$$J_1(\tau) = \mathbb{E}_{\epsilon^{\delta_1}} \left[\max_{\delta_1 \in \{0,1\}} \left\{ \epsilon^{\delta_1} + W_1(\overline{m}, \delta_1; \tau) \right\} \right] \tag{27}$$

$$\begin{aligned} W_1(m_1, \delta_1; \tau) &= \mathbb{E}_{\epsilon^{a_2}} \left[\max_{a_2} \left\{ U(c_1) + \epsilon^{a_2} + \beta_\tau \mathbb{E}_{y_2|y_1} V_2(a_2, y_2, b_2(m_1, \delta_1), \Gamma_1(a_2, \mu_1(m_1, \delta_1)); \tau) \right\} \right] \\ &\quad \text{subject to: } c_1 = y_\ell - (1 - \delta_1) m_1 - q_1(a_2, \mu_1(m_1, \delta_1)) a_2 \end{aligned} \tag{28}$$

Problem (27) generates the type-specific decision rule over medical delinquency $\sigma^{\delta_1}(\tau)$, while (28) generates the medical-record- and type-specific decision rule over assets $\sigma^{a_2}(m_1, \delta_1; \tau)$. We use the notation $\mu_1(\delta_1, m_1)$ 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 use the notation $b_2(m_1, \delta_1) = m_1 \delta_1$ to describe the evolution of medical *debt* based on expense shock realizations and delinquency decisions.

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

$$\begin{aligned} V_2(a_2, b_2, y_2, s_2; \tau) &= \mathbb{E}_{\epsilon^{d_2}} \left[\max_{d_2 \in \{0,1\}} \left\{ U(c_2) + \epsilon^{d_2} + \beta_\tau \phi \left(\psi_2(d_2; a_2, b_2, y_2, s_2) \right) \right\} \right] \\ &\quad \text{subject to: } c_2 = y_2 + (1 - d_2)(a_2 - b_2) \end{aligned} \tag{29}$$

The decision rule associated with this choice problem is $\sigma^{d_2}(a_2, b_2, y_2, s_2; \tau)$.

5.3 What happens when medical records are unobservable?

Out 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 9 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

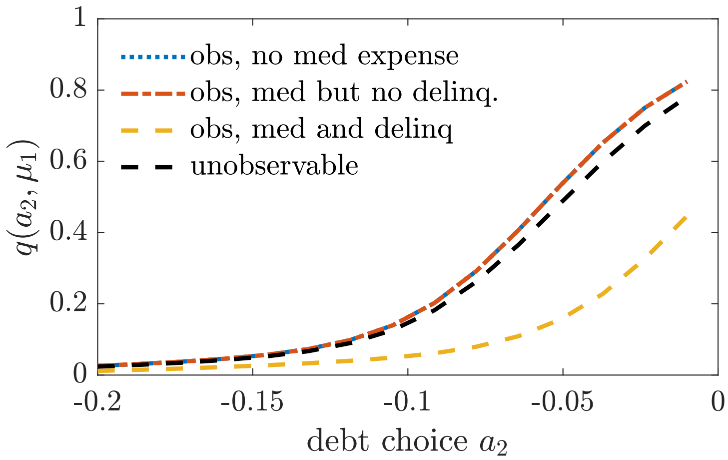

Figure 9: **Credit prices in the medical expense shock model**

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.

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 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 Putting It Together: N-Period Model with Credit Scores

Can the forces described in the models in Sections 3 through 5 help us understand the empirical patterns documented in Section 2? Can we quantify the impacts of regulating lenders' information sets on the credit market and medical payments? To address these questions, this section expands our analysis in four ways. First, we combine the individual

	observable medical events			unobservable medical events		
	mean	diff, $H-L$	std. dev.	mean	diff, $H-L$	std. dev.
Panel 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
Panel 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.

components we've studied along the way. Therefore, the model presented in this section includes adverse selection (as in Sections 3 – 5), costly effort (Section 4), and medical expense shocks and a delinquency decision (Section 5).

Second, we extend the model to have $N > 2$ periods of life. We do this for two reasons. First, this allows us to generate more realistic life cycle profiles as in Figure 1: for example, $N = 8$ where each "age" corresponds to 5 years which allows us to replicate 5 year age bins from $21-25$ until $55-60$. To do so, we simply add periods *between* $n = 1$ and $n = 2$ from the previous sections. Second, and perhaps more importantly, while the two-period models of Sections 3 through 5 can illustrate reputational or type score dynamics, they cannot illustrate credit score dynamics given the constraints of the definition of a credit score outlined above. Put simply, to illustrate credit score dynamics, we need at least two periods of credit scores, which means we need at least two periods in which there is a likelihood of default in the next period. Clearly our 2-period environments do not meet this criterion as constructed.

Third, we make a minor modification by introducing some some churn in individuals types as in (Chatterjee et al., 2023) to match the data better. To that end, we introduce a Markov process over type and assume it is drawn independently across individuals. Obviously, we nest the case in Sections 3-5 if the Markov process is an identity matrix. This allows for the possibility that the composition of H and L types changes over time to model the possibility of compositional changes in unobservable type alluded to in Section 2.1.

Finally, for comparability with the data, we measure reputation using model-implied *credit rankings*, rather than the simpler *type scores* we have studied thus far. Specifically, the probability of repayment (and hence debt prices) depends in part on the lender's assessment of an individuals type encapsulated in their type score s_n that we have focused on in earlier sections. While we observe "cardinal" credit scores in the data ranging from 300 to 850, as described in Section 2 we represent them "ordinally" as credit rankings which depend on observable credit market actions and priors given by the model prior s_n. In Chatterjee et al. (2023), we showed that under certain conditions, equilibria like those described in the previous sections can be implemented via an arrangement in which lenders use the model's equivalent of a credit ranking in place of the type score. We formally define credit rankings in the model in Section 6.3 below after laying out the general environment and the equilibrium. At a high level, though, we define credit rankings by computing the likelihood that an individual will not default or go delinquent next period, then rank individuals in the population according to this likelihood.

6.1 Environment

As discussed above, to introduce the possibility of compositional change in unobservable type, let $\tau_n \in \{H, L\}$ evolve as a Markov Process $Q^\tau(\tau_{n+1}|\tau_n)$. The τ_n are drawn independently across individuals. We assume that the transition probabilities in $Q^\tau(\tau_{n+1}|\tau_n)$ are not symmetric and the initial fraction of high types $\rho_1 < \rho$ so that the composition of type H individuals is growing over time. This process induces changes in β_{τ_n}.

Individuals begin age 1 with state $(a_1, b_1, y_1, s_1, \tau_1)$ drawn from a distribution G_1. Note that the initial distribution assumed in the previous sections is a special case which is degenerate on $(a_1, b_1, y_1, s_1) = (0, 0, y_\ell, \rho_1)$ with the fraction of H-types equal to ρ_1. At the terminal age N we assume that besides valuing consumption c_N, individuals cannot borrow or save, have Medicare pay for any possible medical expense (i.e. they do not realize any medical expense shocks), and value leaving a good reputation $\phi(s_{N+1})$. Since there are multiple periods of asset choice and default unlike the previous subsections, we assume as in Chatterjee et al. (2023) that if an individual defaults they are excluded from the credit market for one period. The only complicated part becomes passing state variables over periods consistent with timing.

To summarize our information assumptions for the general environment, an individual's type (τ_n) and effort (e_n) choice are unobservable. Medical shocks (m_n) and income (y_n) are observable as are credit market actions (δ_n, a_n, d_n). All actions are subject to unobservable idiosyncratic shocks $(\epsilon^{e_n}, \epsilon^{\delta_n}, \epsilon^{a_n}, \epsilon^{d_2})$. Finally, type scores (s_n) are observable providing noisy signals of the unobservable type.

To summarize the timing for this section, at any time $n \in \{1, ..., N-1\}$, then:

1. individuals begin with unobservable type τ_n and observable state $\omega_n = (a_n, b_n, y_n, s_n)$.

2. individuals choose unobservable effort $e_n(\omega_n; \tau_n)$.

3. individuals make default choice $d_n(e_n, \omega_n; \tau_n)$ on private and/or medical debt $|a_n| + b_n$.

4. individuals realize medical expense shock m_n and make a delinquency choice $\delta_n(m_n, d_n, e_n, \omega_n; \tau_n)$, which determines tomorrow's stock of medical debt $b_{n+1} = m_n \delta_n$.[23]

5. if $d_n(e_n, \omega_n; \tau_n) = 0$, then individuals make an asset choice $a_{n+1}(\delta_n, m_n, e_n, \omega_n; \tau_n)$ at price $q_n(a_{n+1}, \omega_n, \mu_n)$; otherwise $a_{n+1} = 0$ by assumption.

6. individuals learn their next-period income y_{n+1}, drawn from $\mathbb{P}(y_{n+1} \mid e_n)$ and their unobservable state τ_{n+1} drawn from $Q^{\tau}(\tau_{n+1}|\tau_n)$.

7. individuals' type assessments are updated according to Bayes' Law:

$$s_{n+1} = \psi_n^{d_n}(y_{n+1}, a_{n+1}, \mu_n; \omega_n) \equiv \mathbb{P}(\tau_{n+1} = H \mid y_{n+1}, a_{n+1}, \mu_n, d_n, \omega_n) \tag{30}$$

8. individuals leave period n in observable state $\omega_{n+1} = (a_{n+1}, b_{n+1}, y_{n+1}, s_{n+1})$.[24]

6.2 Equilibrium

Ages $n = 1, ..., N-1$ 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 $n \in \{1, ..., N-1\}$ and relegate much to Appendix A.4. 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 $\omega = (y, a, b, s)$ makes effort, default, delinquency, 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.[25]

[23] Note that we assume that medical debt can't accumulate: after going delinquent, it is either repaid or defaulted on. This simplifies the state space in our analysis.

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

[25] We 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.

Beginning with the effort choice, an agent first solves

$$V_n(\omega,\tau) = \mathbb{E}_{\epsilon^e}\left[\max_{e\in\{0,1\}}\Big\{-\kappa_e e + W_n(e;\omega,\tau) + \epsilon^e\Big\}\right] \tag{31}$$

where $W_n(e,\omega;\tau_n)$ is the ex-ante value at the default stage, described below. The effort choice decision density is denoted by $\sigma_n^e(\omega;\tau)$. At the default stage, the individual's ex ante value conditional on effort is

$$W_n(e;\omega,\tau) = \mathbb{E}_{\epsilon^d}\left[\max_{d\in\{0,1\}}\Big\{-\kappa_d d + J_n^d(e;\omega,\tau) + \epsilon^d\Big\}\right] \tag{32}$$

where $J_n^d(e;\omega,\tau)$ is the ex ante value at the medical event stage, described below. The default decision density is denoted by $\sigma_n^d(e;\omega,\tau)$. Note that equation (32) only applies if $a(\omega) - b(\omega) < 0$; otherwise, $W_n(e;\omega,\tau) = J_n^{d=0}(e;\omega,\tau)$.

The ex-ante value at the medical event stage accounts for both the probability of a medical expense shock and the potential delinquency choices conditional on the shock:

$$J_n^d(e;\omega,\tau) = \zeta X_n^d(e;\omega,\tau) + (1-\zeta) Z_n^d(0,0,e;\omega,\tau) \tag{33}$$

where $Z_n^d(\delta,m,e;\omega,\tau)$ is the ex ante value at the asset choice stage (here evaluated conditional on $m=0$, which implies $\delta=0$) and $X_n^d(e;\omega,\tau)$ is the ex ante value at the delinquency choice stage conditional on being hit by the medical expense shock:

$$X_n^d(e;\omega,\tau) = \mathbb{E}_{\epsilon^\delta}\left[\max_{\delta\in\{0,1\}}\Big\{-\kappa_\delta \delta + Z_n^d(\delta,\overline{m},e;\omega,\tau) + \epsilon^\delta\Big\}\right], \tag{34}$$

which induces the delinquency decision density $\sigma_n^\delta(d,e;\omega,\tau)$.

Finally, at the asset choice stage the individual solves the following problem if he has not defaulted in this period:

$$\begin{aligned} Z_n^{d=0}(\delta,m,e,;\omega,\tau) = \mathbb{E}_{\epsilon^{a'}}\Big[\max_{a'}\Big\{&U(c) + \epsilon^{a'} \\ &\beta_\tau \mathbb{E}_{y',\tau'|e}\Big(V_{n+1}(a',b'(m,\delta),y',\psi_n^{d=0}(y',a',\mu(m,\delta);\omega);\tau')\Big)\Big\}\Big] \\ \text{subject to: } c = y + a - b - (1-\delta)m - q_n(a';\mu(m,\delta),\omega)a' \end{aligned} \tag{35}$$

The budget constraint adjusts the household's current liabilities for both medical delinquency and default decisions, as well as the equilibrium loan price (discussed below). The interior expectation is over realizations of y', whose likelihood depends on the effort exerted earlier in the period as well as any type switch τ'. The V_{n+1} term spells out all the endogenously evolving parts of the individual state, ω_{n+1}: assets evolve based on the a' choice, medical debts evolve based on the medical record, income evolves based on effort, and the type score evolves based on all observable decisions from this period and the observable beginning-of-period state as per equation (30). The asset density associated with problem (35) is $\sigma_n^{a'}(\delta, m, e; \omega, \tau)$.

If the individual has defaulted earlier in the period, there is no asset choice to be made since $a' = 0$ by assumption, and we have

$$\begin{aligned} Z_n^{d=1}(\delta, m, e, ; \omega, \tau) \quad &= \quad U(c) + \beta_\tau \mathbb{E}_{y', \tau' | e}\Big(V_{n+1}(a', b'(m, \delta), y', \psi_n^{d=1}(y', 0, \mu(m, \delta); \omega); \tau') \Big) \Big\} \Big] \\ \text{subject to:} \quad &\quad c = y - (1 - \delta) m \end{aligned} \tag{36}$$

As discussed in Section 6.1, type scores evolve in a similar fashion to our 2-period model economies. Loan prices in this case are determined in the same way. In particular, the analog of equation (5) in this environment for debt levels $a' < 0$ is simply

$$q_n(a'; \mu, \omega) = \frac{1}{1+r} \mathbb{P}(d' = 0 \mid a', \mu, d = 0, \omega, n) \tag{37}$$

The conditioning on no default this period reflects the fact that this decision is an observable signal of one's type, and that borrowing is only allowed conditional on not defaulting.

6.3 Defining credit scores and credit rankings in the model

Credit scores measure 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 either a default or a medical delinquency, $d = 1$ or $\delta = 1$, and we use the time 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 analogue of a *credit score* which lies in $[0, 1]$ not $[300, 850]$:

$$\tilde{\chi}_n(\omega) \equiv \mathbb{P}(d' = 0, \delta' = 0 \mid \omega, n). \tag{38}$$

Then, in order to take our model to the data, we construct a *credit ranking* by taking the probability $\tilde{\chi}$ from equation (38) and determines the fraction of people who have probability equal to or less than $\tilde{\chi}$ across the equilibrium distribution of people over observable states, $\lambda_n(\omega)$:

$$\chi_n(\omega) \equiv \sum_{\{\tilde{n}, \tilde{\omega} \ \mid \ \tilde{\chi}_{\tilde{n}}(\tilde{\omega}) \leq \tilde{\chi}_n(\omega)\}} \lambda_{\tilde{n}}(\tilde{\omega}) \tag{39}$$

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.

6.4 Model properties

To focus on the information frictions behind credit scoring, here we study a version of the general model with only adverse selection and moral hazard and leave the more general case with medical debt to future research. We compute the model for $N = 10$ but report the model properties for 8 age bins corresponding to the 8 bins in Figure 1. Notably, since reputation affects the terms individuals face in the credit market over their lifetimes in this model, we do not need to introduce the reduced form reputational parameter ϕ in this version of the model: the reputational benefits of good behavior are endogenous here.[26]

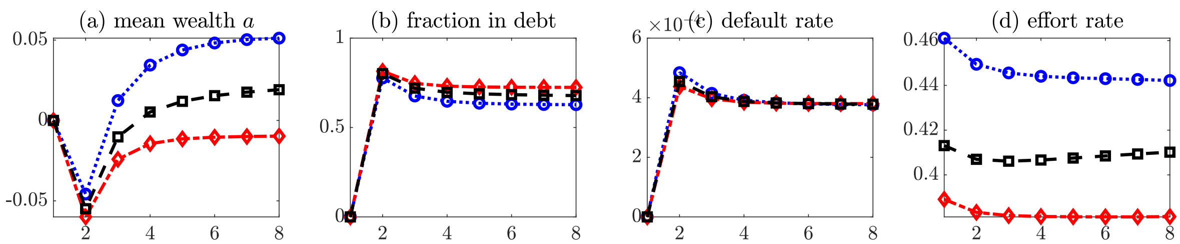

Figure 10: **Life cycle: driving behavior**

Notes: This figure plots the life cycle for the given variables in the aggregate and by type.

To understand the mechanics behind the life cycle pattens of income, consumption, and rankings, we begin with the basic type-dependent choices over assets, default, and effort in the model. Given that people start their working life with low earnings, they borrow early and save late in panel (a) and (b) of Figure 10. Panel (d) shows that there can be big differences in unobservable effort which can lead to big differences in observable income in later age brackets. Since income realizations are a noisy signal of type, there can be early learning and separation of types evident in Figure 11 for both type score in panels (g) – (h) and credit rankings in panels (e) – (f). Panel (c) of Figure 10 suggests that for this parameterization, the default decision is not very informative relative to the asset and effort choices.

Figure 11 uses our model to understand how type differences influence the mean and variance of income, consumption, and credit ranking documented in Section 2. Specifically, as one would expect, type H have higher type score, credit ranking, income, and consumption than type L. We summarize these implications of the model for consumption, income, and

[26]Concretely, we set $\phi = 0$ when computing the continuation value in period N.

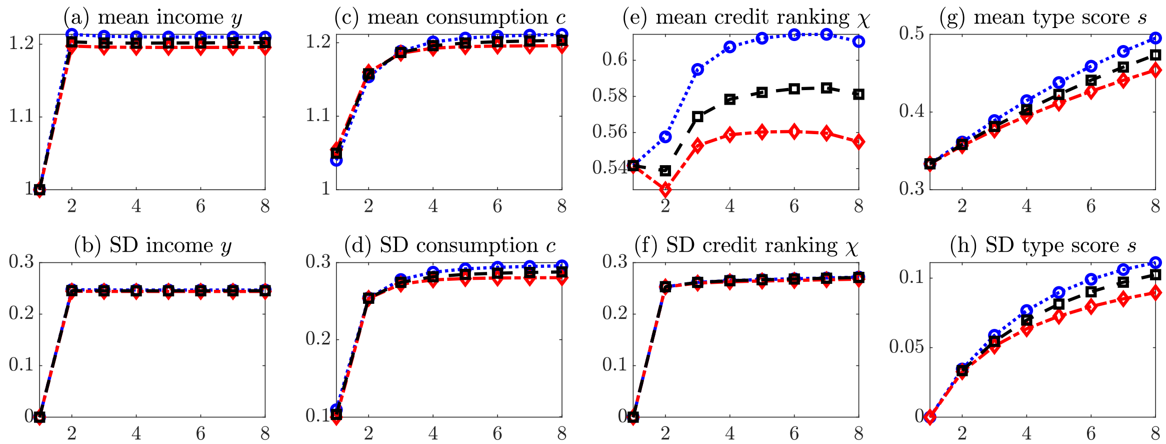

Figure 11: **Life cycle moments in the full model**

Notes: This figure plots the mean and standard deviation over the life cycle for the given variables in the aggregate and by type.

reputational inequality in Figure 12 – the model analog of Figure 1. One interesting fact evident in panel (a) is that there is an upward trend in credit rankings across age despite the leveling off in income, in broad agreement with the data.

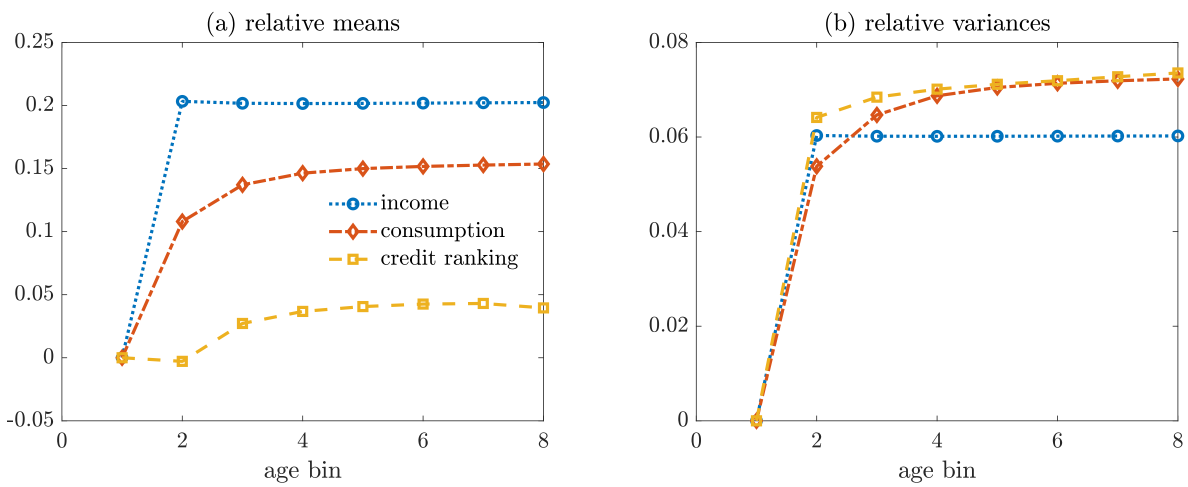

Figure 12: **Relative mean and variance of income, consumption, and credit ranking**

Notes: This figure plots path of the indicated (aggregate) moment for each age, with the first age bin normalized to 0 as in Figure 1.

Finally, Figure 13 shows the joint distribution of income and credit rankings meant to correspond to Figure 2a. Its stark nature arises from the model's two point distribution as opposed to the worlds much richer earnings process. Consistent with the data, though, individuals in the model tend to have higher credit scores (or, the share of high income individuals rises as we move up the credit score distribution.

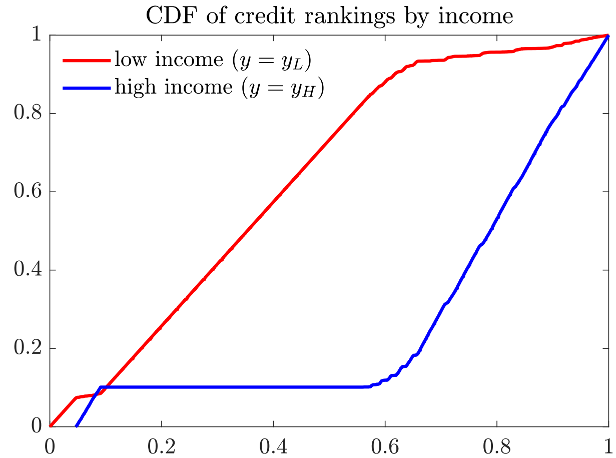

Figure 13: **Distribution of credit rankings by income**

Notes: This figure plots the CDF of credit rankings in the equilibrium distribution of the model conditional on a given level of income. The distribution of credit rankings for high income individuals first order stochastically dominates the one for low income individuals, as in Figure 2a.

References

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.

Athreya, K. (2002): "Welfare Implications of Bankruptcy Reform Act of 1999," *Journal of Monetary Economics*, 49, 1567–1595.

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.

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. Ríos-Rull (2023): "A Quantitative Theory of the Credit Score," *Econometrica*, 91, 1803–1804.

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 (forthcoming): "Employer Credit Checks: Poverty Traps versus Matching Efficiency," *Review of Economic Studies*.

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.

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.

Heathecote, 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*.

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 stuies 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.

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," *The Sociological Review*, 55, 284–306.

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.

Appendix for "Credit Scores and Inequality over the Life Cycle"

A Model Appendix

A.1 Formulas for baseline environment from Section 3

Given the age-1 decision rule over $\sigma^{a_2}(\tau)$, type score updates are computed via (6):

$$\psi_1(a_2) \equiv \mathbb{P}(\tau = H \mid a_2) = \frac{\mathbb{P}(a_2 \mid \tau = H)\mathbb{P}(\tau = H)}{\sum_\tau \mathbb{P}(a_2 \mid \tau)\mathbb{P}(\tau)} = \frac{\rho\sigma_1(a_2; H)}{\rho\sigma_1(a_2; H) + (1-\rho)\sigma_1(a_2; L)} \tag{A.1}$$

These revised assessments are used to compute prices via (5), where the probability term is:

$$\mathbb{P}(d_2 = 0 \mid a_2) = \sum_{y_2,\tau} \mathbb{P}(y_2, \tau \mid a_2) \cdot \sigma^{d_2=0}(a_2, y_2, \psi_1(a_2); \tau)$$

where $\sigma^{d_2=0}(a, y, s; \tau)$ is the probability of repayment at age 2. Since y_2 is drawn randomly independent of age 1 choices, the probability term may be written $\mathbb{P}(y_2, \tau \mid a_2) = \mathbb{P}(y_2 \mid y_\ell) \cdot \mathbb{P}(\tau \mid a_2)$, where the last term is just a transformation of the type score; that is, $\mathbb{P}(\tau = H \mid a_2) = \psi_1(a_2)$ and $\mathbb{P}(\tau = L \mid a_2) = 1 - \psi_1(a_2)$. Combining these insights, we obtain

$$\begin{aligned} q_1(a_2) = \frac{1}{1+r} \sum_{y_2} \Bigg(\mathbb{P}(y_2 \mid y_\ell) \Big[&\psi_1(a_2)\sigma^{d_2=0}(a_2, y_2, \psi_1(a_2); H) \\ &+ (1 - \psi_1(a_2))\sigma^{d_2=0}(a_2, y_2, \psi_1(a_2); L) \Big] \Bigg) \end{aligned} \tag{A.2}$$

Lastly, the type score update based on the age 2 default decision is

$$\psi_2(d_2; a_2, y_2, s_2) = \frac{s_2\sigma^{d_2}(a_2, y_2, s_2; H)}{s_2\sigma^{d_2}(a_2, y_2, s_2; H) + (1 - s_2)\sigma^{d_2}(a_2, y_2, s_2; L)} \tag{A.3}$$

No tracking 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

(6), solved for in equation (A.2). In this case, the loan price is

$$q_1^{\mathrm{NT}}(a_2) = \frac{1}{1+r}\sum_{y_2}\left(\mathbb{P}(y_2 \mid y_\ell)\Big[\psi_1(a_2)\sigma^{d_2=0}(a_2, y_2; H) + (1-\psi_1(a_2))\sigma^{d_2=0}(a_2, y_2; L)\Big]\right) \tag{A.4}$$

Note that 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 the preclusion of tracking implies there is no scope for reputational considerations in the default decision at age 2, equation (A.3) is no longer relevant.

A.2 Formulas for moral hazard environment from Section 4

In order to streamline cumbersome notation, it is useful to define four objects. First, the age-1 joint decision rule over effort and assets is

$$\tilde{\sigma}^{(a_2,e_1}(\tau) \equiv \mathbb{P}(a_2, e_1 \mid \tau) = \sigma^{e_1}(\tau)\sigma^{a_2}(e, \tau)$$

Second, the effort-choice-weighted age-1 decision rule over assets is

$$\overline{\sigma}^{a_2}(\tau) \equiv \mathbb{P}(a_2 \mid \tau) = \sum_{e_1} \tilde{\sigma}^{(a_2,e_1)}(\tau)$$

Third, the probability of a given y_2 realization based on type and asset choice is

$$\pi^{y_2}(a_2; \tau) \equiv \mathbb{P}(y_2 \mid a_2, \tau) = \sum_{e_1} \mathbb{P}(y_2 \mid e_1) \frac{\tilde{\sigma}^{(a_2,e_1)}(\tau)}{\overline{\sigma}^{a_2}(\tau)}$$

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 version of the model, it is still useful as an input to loan prices and the full type score update Υ_2. We compute the type assessment made *after* the choice of a_2 but *before* the realization of y_2 similarly to (A.1), with the only difference that we need to sum over unobservable effort decisions:

$$\psi_1(a_2) \equiv \mathbb{P}(\tau \mid a_2) = \frac{\rho\overline{\sigma}^{a_2}(H)}{\rho\overline{\sigma}^{a_2}(H) + (1-\rho)\overline{\sigma}^{a_2}(L)} \tag{A.5}$$

With repeated application of Bayes' Rule, we can write equation (21) as

$$\mathbb{P}(\tau = H \mid a_2, y_2) = \frac{\mathbb{P}(y_2 \mid \tau = H, a_2)\mathbb{P}(\tau = H \mid a_2)}{\sum_\tau \mathbb{P}(y_2 \mid \tau, a_2)\mathbb{P}(\tau \mid a_2)}$$

which uses both the "interim" type score ψ_1 and the updated y_2 density π. Combining these expressions yields

$$\Upsilon_2(a_2, y_2) = \frac{\psi_1(a_2)\pi^{y_2}(a_2, H)}{\psi_1(a_2)\pi^{y_2}(a_2, H) + (1 - \psi_1(a_2))\pi^{y_2}(a_2; L)} \tag{A.6}$$

In this setting, to compute the loan prices we compute

$$\mathbb{P}(d_2 = 0 \mid a_2) = \sum_{y_2, \tau} \mathbb{P}(y_2, \tau \mid a_2) \cdot \sigma^{d_2=0}\left(a_2, y_2, \Upsilon_2(a_2, y_2); \tau\right)$$

where we can compute the probability term easily using objects derived above $\mathbb{P}(y_2, \tau \mid a_2) = \mathbb{P}(y_2 \mid a_2, \tau)\mathbb{P}(\tau \mid a_2) = \pi^{y_2}(a_2, \tau)\psi_1(a_2)$. Combining these yields

$$\begin{aligned} q_1(a_2) = \frac{1}{1+r} \sum_{y_2} \Big[&\psi_1(a_2)\pi^{y_2}(a_2, H)\sigma^{d_2=0}(a_2, y_2, \Upsilon_2(a_2, y_2); H) \\ &+ (1 - \psi_1(a_2))\pi^{y_2}(a_2, L)\sigma^{d_2=0}(a_2, y_2, \Upsilon_2(a_2, y_2); L) \Big] \end{aligned} \tag{A.7}$$

Lastly, the type score updates after the default / no default decision at age 2 are exactly as in the baseline model, following equation (A.3).

A.3 Formulas for medical event environment from Section 5

As in the moral hazard case, define the ex-ante decision rule over (a_2, δ_1) given the realization of the medical shock m_1 by

$$\tilde{\sigma}^{(a_2, \delta_1)}(m_1, \tau) \equiv \mathbb{P}(a_2, \delta_1 \mid m_1, \tau) = \sigma^{a_2}(\delta_1, m_1, \tau) \cdot \sigma^{\delta_1}(m_1, \tau)$$

One difference between this expression and the one from the moral hazard case is that it could in principle be extended to include the likelihood of the medical expense shock, but we exclude this for now since it is exogenous.

A.3.1 Observable medical events

Consider first the case in which medical debts are observable. In this case, we can write the type score as

$$\begin{aligned} \mathbb{P}(\tau = H \mid a_2, m_1, \delta_1) &= \frac{\mathbb{P}(\tau = H, a_2, m_1, \delta_1)}{\sum_\tau \mathbb{P}(\tau, a_2, m_1, \delta_1)} \\ \text{where } \mathbb{P}(\tau, a_2, m_1, \delta_1) &= \mathbb{P}(a_2 \mid m_1, \delta_1, \tau)\mathbb{P}(\delta_1 \mid m_1, \tau)\mathbb{P}(m_1 \mid \tau)\mathbb{P}(\tau) \end{aligned}$$

Then, using the definitions of the decision rules above and recognizing that medical expense shocks are independent of type, we obtain

$$\lambda_1(a_2, m_1, \delta_1) = \frac{\rho\tilde{\sigma}^{(a_2,\delta_1)}(m_1, H)}{\rho\tilde{\sigma}^{(a_2,\delta_1)}(m_1, H) + (1-\rho)\tilde{\sigma}^{(a_2,\delta_1)}(m_1, L)} \tag{A.8}$$

In this case, the price schedule directly mirrors the one from the baseline version of the model in (A.2), with the difference that there are more states to account for in the type scoring function and subsequent decision rule:

$$\begin{aligned} q_1(a_2, \mu_1) = \frac{1}{1+r}\sum_{y_2}\Bigg(&\mathbb{P}(y_2 \mid y_\ell)\Big[\lambda_1(a_2,\mu_1)\sigma^{d_2=0}\Big(a_2, b(\mu_1), \lambda_1(a_2,\mu_1), y_2; H\Big) \\ &+ (1-\lambda_1(a_2,\mu_1))\sigma^{d_2=0}\Big(a_2, b(\mu_1), \lambda_1(a_2,\mu_1), y_2; L\Big)\Big]\Bigg) \end{aligned} \tag{A.9}$$

A.3.2 Unobservable medical events

In this case, it is useful to define the τ-specific probability of choosing a_2 as

$$\begin{aligned} \overline{\sigma}^{a_2}(\tau) \equiv \mathbb{P}(a \mid \tau) &= \sum_{m,\delta}\mathbb{P}(a \mid m, \delta, \tau)\mathbb{P}(\delta \mid m, \tau)\mathbb{P}(m) \\ &= (1-\chi)\tilde{\sigma}^{(a_2,0)}(0,\tau) + \chi\left[\tilde{\sigma}^{(a_2,0)}(\overline{m},\tau) + \tilde{\sigma}^{(a_2,1)}(\overline{m},\tau)\right] \end{aligned}$$

Note that 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, this is the relevant conditioning for the lender / credit scorer.

Another useful object to define is the type-specific likelihood of a specific medical event (m_1, δ_1) given an observation of a_2:

$$\begin{aligned} \pi^{(m_1,\delta_1)}(a_2,\tau) \equiv \mathbb{P}(m_1, \delta_1 \mid a_2, \tau) &= \frac{\mathbb{P}(a_2 \mid m_1, \delta_1, \tau)\mathbb{P}(\delta_1 \mid m_1, \tau)\mathbb{P}(m_1)\mathbb{P}(\tau)}{\mathbb{P}(a_2 \mid \tau)\mathbb{P}(\tau)} \\ &= \mathbb{P}(m)\frac{\tilde{\sigma}^{(a_2,\delta_1)}(m_1,\tau)}{\overline{\sigma}^{a_2}(\tau)} \end{aligned}$$

This object serves a similar purpose to the π object in the moral hazard model.

When medical events are unobservable, we have to sum over their likelihood in making assessments. Individual's type scores in this version of the model evolve exactly as in the baseline model (equation (A.1)), with the caveat that the relevant decision rule is now $\overline{\sigma}^{a_2}$, the weighted sum over effort-specific choices:

$$\psi_1(a_2) = \frac{\rho\overline{\sigma}_1(a_2; H)}{\rho\overline{\sigma}_1(a_2; H) + (1-\rho)\overline{\sigma}_1(a_2; L)} \tag{A.10}$$

For loan prices, the formula resembles the moral hazard model with unobservable effort. In this case, we must compute

$$\begin{aligned} \mathbb{P}(d_2 = 0 \mid a_2) &= \sum_{y_2,\tau,m_1,\delta_1} \mathbb{P}(y_2, \tau, m_1, \delta_1 \mid a_2) \cdot \sigma^{d_2=0}(a_2, b(m_1, \delta_1), \psi_1(a_2), y_2; \tau) \\ \text{where } \mathbb{P}(y_2, \tau, m_1, \delta_1 \mid a_2) &= \mathbb{P}(y_2 \mid y_\ell)\mathbb{P}(\tau \mid a_2)\mathbb{P}(m_1, \delta_1 \mid a_2, \tau) \end{aligned}$$

where the second term is a function of the type score and the third term is the inferred probability of a medical event having occurred computed above. Combining these results, we can write the loan price function for this case as

$$\begin{aligned} q_1(a_2) = \frac{1}{1+r} \sum_{y_2,m_1} \Bigg(&\mathbb{P}(y_2 \mid y_\ell)\Big[\psi_1(a_2)\pi^{(m_1,\delta_1)}(a_2, H)\sigma^{d_2=0}\Big(a_2, b(m_1, \delta_1), \psi_1(a_2), y_2; H\Big) \\ &+ (1-\psi_1(a_2))\pi^{(m_1,\delta_1)}(a_2, L)\sigma^{d_2=0}\Big(a_2, b(m_1, \delta_1), \psi_1(a_2), y_2; L\Big)\Big)\Big]\Bigg) \end{aligned} \tag{A.11}$$

Lastly, in both the observable and unobservable cases, the type score updates after the default / no default decision at age 2 are exactly as in the baseline model, following equation (A.3), with the one modification that the age 2 decision rule depends also on the quantity of medical debt.

A.4 Formulas for General Case

To be added

B Additional Numerical Results

B.1 Parameterization for numerical illustrations

Table B.2 reports the parameter values used throughout the numerical illustrations in the paper. Many parameter values are standard (e.g. risk aversion, risk-free rate). The share of H-types in the population is $2/3$, to capture the idea that broadly H-types are the prime population while L-types are the subprime population. Both types are relatively impatient, a

parameter		value	notes
Panel A. Common across all models (Sections 3 – 5)			
risk aversion	γ	2	standard CRRA utility
risk-free rate	r	2%	avg real T-Bill rate
share of H-types	ρ	2/3	prime share of population
high discount factor	β_H	0.9	
low discount factor	β_L	0.5	
low income	y_ℓ	0.75	normalization, center income process around 1
high income	y_h	1.25	
probability of high income	$\mathbb{P}(y_h)$	0.1	equal to $\mathbb{P}(y_h \mid e = 0)$ in MH model
extreme value scale, default	α_d	0.05	
extreme value scale, saving	α_a	0.15	
marginal value of reputation	ϕ	1	
Panel B. Unique to moral hazard model (Section 4)			
probability of high income, effort	$\mathbb{P}(y_h \mid e = 1)$	0.5	
utility cost of effort	κ_e	0.2	
extreme value scale, effort choice	α_e	0.05	equal to α_d
Panel C. Unique to medical shock model (Section 5)			
probability of medical expense	χ	34.1%	Fulford and Low (2024)
size of medical expense	$\overline{m}$	0.05	8% of y_ℓ, Fulford and Low (2024)
utility cost of medical delinquency	κ_δ	0.1	equal to κ_d
extreme value scale, medical delinquency	α_δ	0.05	equal to α_d

Table B.1: **Parameters for numerical illustrations**

Notes: This table reports the parameter values used throughout the numerical illustrations in the paper.

standard assumption in the unsecured credit literature. The extreme value preference shocks are consistent with the idea that the default, delinquency, and effort choices are all less noisy than asset choices, which reflects scope for un-modeled expense or preference shocks driving consumption in a given period.

For the moral hazard model, we assume that the probability of high income conditional on not exerting effort is equal to the probability of high income from the baseline model, while the probability of high income conditional on exerting effort is much higher. This creates scope for the moral hazard friction to have bite. Lastly, for the model with medical expense shocks, we assume that the likelihood of being hit by a shock and the average size conditional on being hit by a shock are consistent with empirical measurements.

B.2 Baseline model mechanics

Figure B.1 presents key decision rules and equilibrium objects from the baseline model of Section 3. Panels (a) – (c) focus on decisions and equilibrium functions at age 1; panels (d) – (f) focus on age 2. Panel (a) presents the decision rules by type: it is clear that H-types

parameter		value	notes / target	data	model
Panel A. Assigned externally					
risk aversion	γ	2	standard CRRA utility		
risk-free rate	r_{ann}	2%	2% annualized, $r = (1 + r_{\text{ann}})^5 - 1$		
long run share of H-types	ρ	2/3	prime share of population		
high discount factor	β_H	0.906	$\beta_H = \frac{1}{1+r}$		
low income	y_ℓ	1	normalization		
probability of high income, no effort	$\mathbb{P}(y_h \mid e = 0)$	0.2	normalization		
Panel B. Calibrated internally					
low discount factor	β_L	0.7	average debt-to-income	0.4	0.1
share of high types among newborns	ρ_1	1/3	mean log income, slope	0.047	0.014
high income	y_h	1.5	mean credit ranking, intercept	0.385	0.547
			mean credit ranking, slope	0.024	0.007
transition probability, $\tau = L$ to $\tau' = H$	$\mathbb{P}(\tau' = H \mid \tau = L)$	0.05	average interest rate	11.9%	12.4%
probability of high income, effort	$\mathbb{P}(y_h \mid e = 1)$	0.7	variance of interest rates	7.0%	1.5%
utility cost of effort	κ_e	0.2	avg. corr between score & income	0.64	0.80
extreme value scale, default	α_d	0.05	variance of credit ranking, intercept	0.060	0.038
extreme value scale, saving	α_a	0.15	Ch 7 bankruptcy rate	1.0%	1.1%
extreme value scale, effort choice	α_e	0.05	variance credit ranking, slope	0.006	0.007

Table B.2: **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 calibrate 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.

tend to save more, while L-types are more likely to borrow. Panel (b) presents the type score update based on the asset choice, $\psi_1(a_2) - \rho$. Choices for which the number lies above (below) zero correspond to choices that lead to upward (downward) revision of one's type score. This figure inherits the shape of the "likelihood ratio" $\frac{\sigma^{a_2}(H)}{\sigma^{a_2}(H)+\sigma^{a_2}(L)}$, a transformation of panel (a). Panel (c) plots the loan pricing function, which takes a standard form for the unsecured credit literature.

Panels (d) and (e) plot the age 2 default probability across debt levels for H- and L-types for the realization $y_2 = y_\ell$. Panel (d) examines behavior for a relatively high type score, while panel (d) examines it for a relatively low type score. Both types become more likely default when more indebted, and are less likely to default when s_2 is at a more intermediate level (for which there is scope for a large downward revision of reputation upon default). Finally, panel (f) plots the type score update in the event of default for high and low values of s_2. Unsurprisingly, defaulting leads to a sharper downward revision of one's type score when one begins with a better reputation.

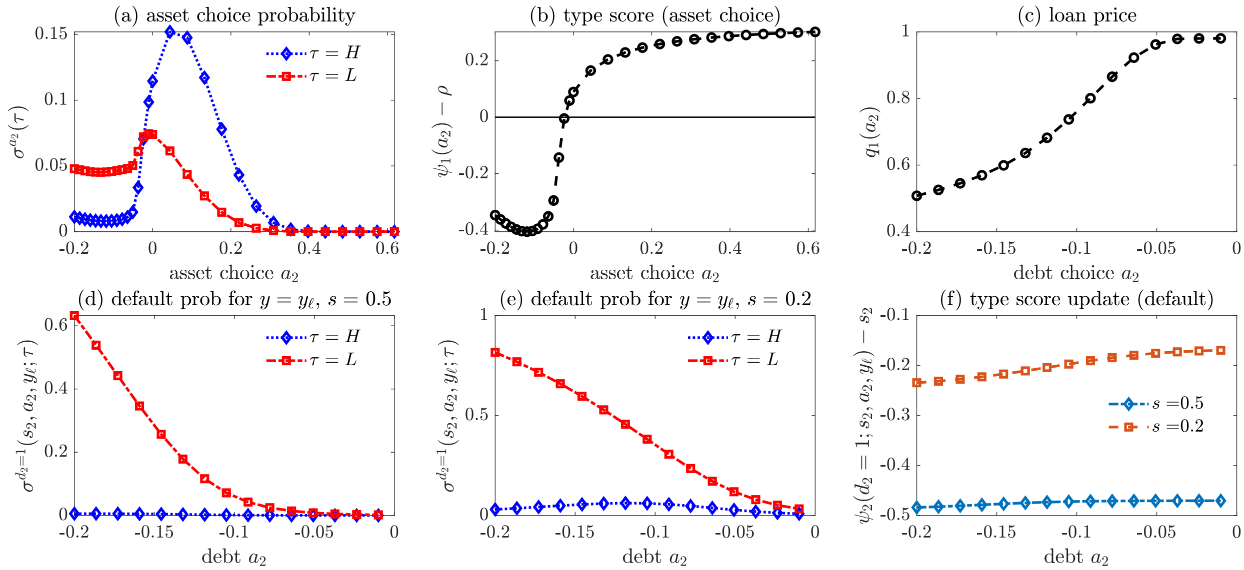

Figure B.1: **Equilibrium in the baseline model (Section 3)**

Notes: Panel (a) plots the decision density over a_2 for each type. Panel (b) plots the type score associated with each action a_2, expressed as a difference relative to the common initial type score ρ. Panel (c) plots the loan pricing function. Panels (d) and (e) plot the default probabilities conditional on having low income at age 2 over a range of debt levels for the high and low types for high and low type scores (10th and 50th percentiles), respectively. Panel (f) plots the type score update conditional on default for the same states as panels (d) and (e).

B.3 Moral hazard model mechanics

Equilibrium Figure B.2 describes the equilibrium of the model with moral hazard from Section 4. Since the main differences between this model and the baseline model from Section 3 affect behavior at age 1, we focus only on this age to ease exposition. Panel (a) is the analog of panel (a) from Figure B.1, with the caveat that the a_2 decision rule must be summed over effort levels. Still, the message is the same: H-types are more likely to save / borrow less than L-types. Panel (b) depicts the decision rule over effort, $\sigma^{e_1=1}(\tau)$. H-types, valuing the future increase in income, are more likely to incur the effort cost when young than L-types.

Panel (c) plots 2 type score updates: the "interim" one based only on the choice of a_2, $\psi_1(a_2)$ – which is relevant for pricing – and the "final" one based on both the choice of a_2 and the realization of y_2, $\Upsilon_2(a_2, y_2)$ – which determines the evolution of the type score state variable. All 3 lines show that the type score revision takes the shape implied by the asset choices (specificlly, their likelihood ratio representation) from panel (a). The blue line shows that a realization of high income at age 2 leads to a further upward revision of the assessment of one's type, while the red line shows that a realization of low income at age 2 leads to a further downward revision. This is driven by the differences in the effort decision highlighted in panel (b): H-types are more likely to exert effort, and high income is more likely given

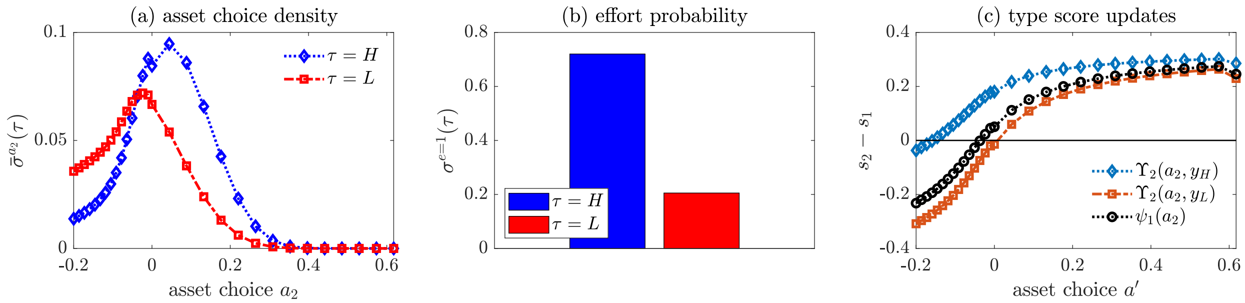

Figure B.2: **Equilibrium in the moral hazard model (Section 4)**

Notes: Panel (a) plots the decision rule over a_2 only, which is a weighted average of the a_2 choices made for each type across effort levels, with the weights as the probability of each effort choice. Panel (b) plots the probability of exerting effort by type, $\sigma^{e_1=1}(\tau)$. Panel (c) plots the type scoring functions, both the interim – after observing a_2 only – and the final – after observing both a_2 and the realization of y_2 at the beginning of age 2.

the exertion of effort, so a realization of high income is reputation-enhancing, all else equal.

Life cycle Figure B.3 presents the analog of Figure 6 from the main text. The main purpose of this figure is to provide a visual aid to accompany the moments from the baseline model in Table 2.

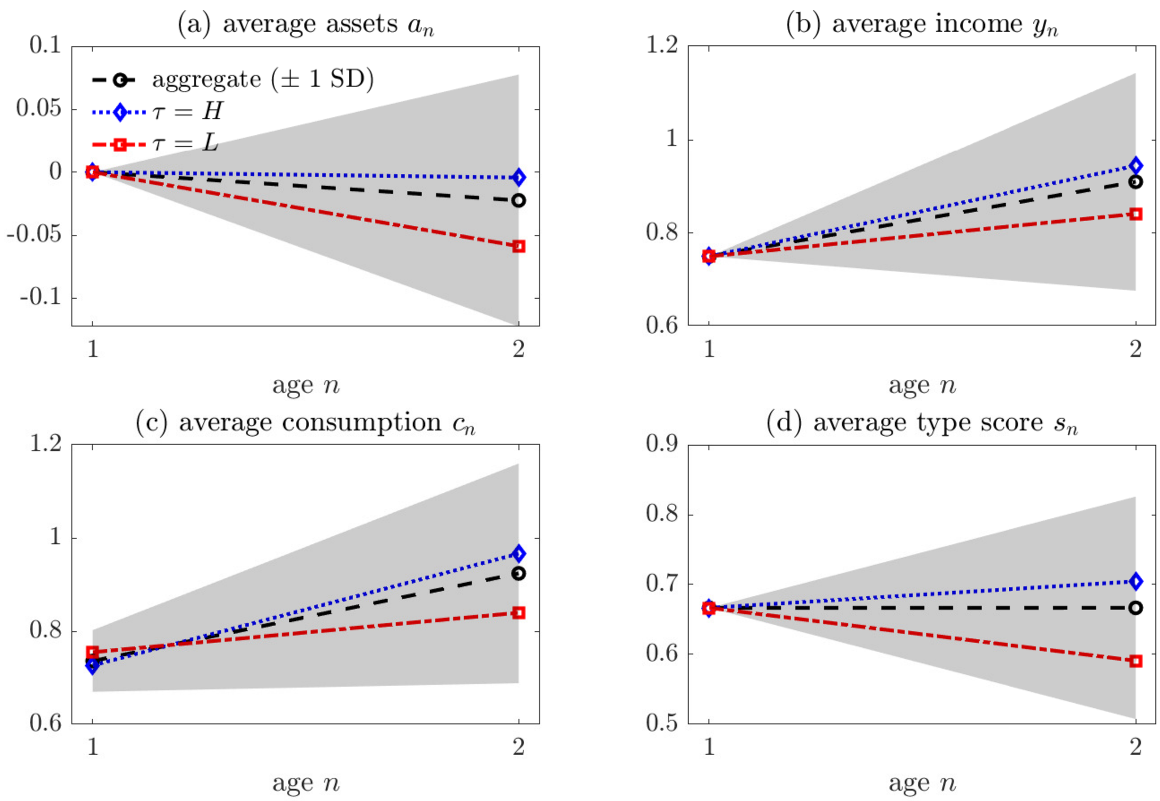

Figure B.3: **Life cycle in the moral hazard model (Section 4)**

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.