

A Quantitative Theory of Credit Scoring*

Kartik Athreya[†]

Research Department

Federal Reserve Bank of Richmond

Xuan S. Tam[‡]

Eric R. Young[§]

Centre for Financial Analysis and Policy

Department of Economics

University of Cambridge

University of Virginia

February 12, 2013

Preliminary, Do Not Cite

* The opinions expressed here are not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

[†]E-mail: kartik.athreya@rich.frb.org.

[‡]E-mail: xst20@cam.ac.uk.

[§]E-mail: ey2d@virginia.edu.

Abstract

Starting in the early 1990s credit scoring became widespread and central in credit granting decisions. Credit scores are scalar representations of default risk. They are used, in turn, to price credit, and as a result alter household borrowing and default decisions. We build on recent work on defaultable consumer credit under asymmetric information to develop a quantitative theory of credit scores. We construct and solve a rich and quantitatively-disciplined lifecycle model of consumption in which households have access to defaultable debt, and lenders are asymmetrically informed about household characteristics relevant to predicting default. We then allow lenders to keep record of inferences on the hidden type of a borrower, as well as a binary “flag” indicating a past default. These inferences arise endogenously from a signalling game induced by borrowers’ need to obtain loans. We show how lenders’ inferences evolve over the lifecycle as a function of household behavior in a way that can be naturally interpreted as “credit scores.” In particular, we first show that lenders’ assessments that a household has relatively low default risk matter significantly for the interest rates households pay. We then show that such assessments rise most sharply—and interest rates paid by borrowers fall most sharply (on the order of 5 – 6 percentage points)—when the bankruptcy flag is removed, consistent with work of Musto (2005). Lastly, we compare allocations across information regimes to provide a measure of the social value of credit scores, and the dependence of these measures on lenders’ ability to observe borrower characteristics.

Keywords: *Bankruptcy, Unsecured Credit, Credit Scoring*

JEL Codes: *D82, D91, E21.*

1 Introduction

An individual's FICO score is the most common method of assessing creditworthiness (default risk). Credit scores matter for terms, as noted in Chatterjee, Corbae, and Ríos-Rull (2010):

1. Interest rates are lower for households with higher credit scores
2. Bankruptcy lowers the credit score, removal increases it
3. Increasing debt lowers credit score, reducing debt raises it
4. Credit scores are mean reverting

Anecdotally, using a simple on-line calculator one can obtain a distribution of interest rate offers by credit score:

Table 1

FICO Score	Auto Loan	Mortgage
720 – 850	4.94%	5.55%
700 – 719	5.67%	5.68%
675 – 699	7.91%	6.21%
620 – 674	10.84%	7.36%
560 – 619	15.14%	8.53%
500 – 559	18.60%	9.29%

It is clear that rates rise as FICO scores fall, even for debts that are (at least partially) secured by collateral. Figure (1) shows that in fact FICO scores are negatively correlated with default rates (this figure appears in Chatterjee, Corbae, and Ríos-Rull 2011).

Despite this (and more) evidence, the existing quantitative literature on consumer default has not introduced credit scoring, even in environments with asymmetric information where it would be useful (such as Athreya, Tam, and Young 2012, Sánchez 2011, Narajabad 2012, and Livshits, MacGee, and Tertilt 2011). Only one paper attempts to introduce dynamic record keeping (Chatterjee, Corbae, and Ríos-Rull 2011), but it does so in a very simple quantitative environment.¹ In contrast, this paper integrates dynamic credit scoring into a quantitative model of default with asymmetric information (namely, the model of Athreya, Tam, and Young 2012); specifically, we assume the existence of a credit bureau that assembles a one-dimensional statistic summarizing the assessed probability that a given individual is a "bad type" given their actions (including the choice to go bankrupt). This statistic evolves over time as agents borrow, save, and file for bankruptcy; similar to Chatterjee, Corbae, and Ríos-Rull (2011), our goal is to develop a model-specific object whose behavior matches the four observations about credit scores.

To do so, we simplify our previous model by excluding expenditure shocks; we also assume that income is observable for computational reasons (quantitative results from Athreya, Tam, and Young 2012 suggest income is not crucial to the distortions created by asymmetric information). Here we assume that the current nonpecuniary cost of default is unobservable (along with current debt, which is assumed unobservable in order to ensure a nontrivial inference problem). The credit score is the probability that a given household has a low cost of bankruptcy, and this probability is updated using Bayes' rule conditioned on observable states (age, education, income) and observable actions (new borrowing and the default decision). In addition to the credit score, we assume that past bankruptcies are recorded with a flag that is removed stochastically.

¹Ad hoc methods of introducing scoring exist, such as Ionescu and Simpson (2011) for student loans, where the evolution of the credit scores (and their influence on interest rates) is exogenously specified. D'Erasmus (2011) studies a one-agent sovereign debt model.

Our preliminary results show that the presence of a credit score, in addition to the bankruptcy flag that we used in our previous work, does help mitigate the asymmetric information problem, but it does not eliminate it. Asymmetric information still reduces the amount of borrowing in the unsecured credit market, decreases the bankruptcy rate, and decreases the size of debts discharged through bankruptcy.

We then evaluate three alternative information settings. First, we assume the credit score is inactive, leading to a model similar to Athreya, Tam, and Young (2012). Second, we assume that past bankruptcy are not recorded, leaving only the credit score. Third, we assume that past bankruptcies are *permanently* recorded, so that information about past defaults never leaves the borrower. Chatterjee, Corbae, and Ríos-Rull (2011) run a related experiment, where they examine the effect of forcing past default information to be dropped – in their model restricting information benefits bad type (myopic) agents at the expense of good type (patient) agents, and on average information restriction is welfare-improving.

We can compare our model’s predictions to the empirical work in Musto (2004), who studies the effect of a bankruptcy marker on credit scores and borrowing. In his data, the removal of a flag associated with a past bankruptcy generates an immediate increase in the credit score of a borrower, but is also associated with a future *decline*; in turn, these movements are associated with short-term decreases and long-run increases in credit card delinquency rates. Quantitatively, the removal of the flag of a borrower between the 40th and 50th percentiles of the credit score distribution jumps over 19 percent of the population of nonfilers in the data; in turn their debt limit increases by \$1800 on average.

2 Model Environment

There is a continuum of *ex ante* identical households who each live for a maximum of $J < \infty$ periods. Households supply labor inelastically until they retire at age $j^* < J$, and differ

in their human capital type, y . A household's human capital type governs the mean of income at each age over the life-cycle. A household of age j and human capital type y has a probability $\psi_{j,y} < 1$ of surviving to age $j + 1$ and has a pure time discount factor $\beta < 1$. Households also vary over their lifecycle in size. Let n_j denote the number of adult-equivalent members present when the head of the household is age- j . Consumption per person at age- j , c_j , is a purely private good and therefore produces less utility as the number of household members grows.

The economy is one in which all agents take a risk-free rate r as given, and we take this rate to be exogenous. There exists a competitive market of intermediaries who offer one-period savings contracts that promise a deterministic rate of return r , and also offer contracts for one-period loans that agents may default on by invoking personal bankruptcy protection. Loans are modeled as arising from the purchase of debt from households that is discounted at rate q , capturing both default risk and the maturity of the loan. Specifically, if $I \in \mathcal{I}$ denotes the information available to lenders, then let $q(b, I) \in [0, q^f]$ be the discount applied to the bond issuance of face value b from a household. In other words, if a household issues b and the market discounts this issuance at $q(b, I)$, the household then receives $q(b, I)b$ today and owes, outside of bankruptcy, b units next period.

Lenders utilize available information to assess default risk and offer individualized credit pricing that is competitive given this risk. Bankruptcy is costly, and has both a pecuniary cost Δ , and an individual-specific and stochastic non-pecuniary component (that will also be persistent in the quantitative analysis) denoted by the term $\lambda \in \Lambda \subseteq [0, 1]$. The explicit resource costs of bankruptcy represent legal fees, court costs, and other direct expenses associated with filing. The existence of nonpecuniary costs of bankruptcy, represented by λ , is strongly suggested by a range of recent work. First, Fay, Hurst, and White (1998) find that a large measure of households would have “financially benefited” from filing for bankruptcy but did not. Second, Gross and Souleles (2002) and Fay, Hurst, and White (1998)

document significant unexplained variability in the probability of default across households, even after controlling for a large number of observables. These results imply the presence of implicit collateral, which may or may not be observable and heterogeneous across households; λ reflects any such collateral, including (but not limited to) any stigma associated with bankruptcy. We model λ as a multiplicative factor that alters the value of consumption in the period in which a household files for bankruptcy.² We assume that λ follows a two-state Markov chain with education-invariant transition probabilities and realizations that differ across education levels.

Household preferences are represented by the expected utility function

$$\sum_{j=1}^J \sum_{s^j} \left(\prod_{i=0}^j \beta \psi_{j,y} \right) \Pi(s^j) \left[\frac{n_j}{1-\sigma} \left(\frac{I_{\mathcal{D}}(\lambda_{j,y}) c_j}{n_j} \right)^{1-\sigma} \right] \quad (1)$$

where $\Pi(s^j)$ is the probability of a given history of events s^j , and $\sigma \geq 0$ is the Arrow-Pratt coefficient of relative risk aversion. Letting the current period bankruptcy decision be denoted by $\mathcal{D} \in \{0, 1\}$, we set $I_{\mathcal{D}}(\lambda_{j,y}) = 1$ if the household does *not* choose bankruptcy ($\mathcal{D}=0$), and $I_{\mathcal{D}}(\lambda_{j,y}) = \lambda_{j,y} < 1$ otherwise. Our specification of non-pecuniary bankruptcy costs is such that the higher the value of λ , the *higher* the risk of bankruptcy, because the effective “tax” on consumption is smaller when λ takes a relatively high value than a low one.

Our model allows us to study improvements in information arising from improvements in the ability of lenders to assess factors *orthogonal* to income but germane for the prediction of default risk. As to the latter, the now-common use of detailed expenditure patterns and general “data-mining” by lenders suggests that they are indeed interested in gleaning

²An alternative method of capturing the facts noted by Fay, Hurst, and White (1998) is used in Athreya *et al.* (2012), which permits households to go delinquent and stop servicing their debt without actually discharging it. We have experimented with introducing delinquency into the model we are using here, but the performance of the algorithm used to compute prices was not stable.

differences in default risk amongst groups whose observables, have already been forecasted as accurately by lenders as by the borrower.³

Creditors will generally want to track the history of default, since it will be informative regarding the current value of λ (in incomplete market models, current wealth is a function of initial wealth, here 0, and the entire history of shocks unless there is a "renewal" event where agents exit the event with their history having been wiped out).⁴ Our first imperfect tracking mechanism is via a binary marker $m \in \{0, 1\}$, where $m = 1$ indicates the presence of bankruptcy in a borrower's past and $m = 0$ implies no record of past default. This marker will reset in some future period, capturing the effect of the Fair Credit Reporting Act requiring that bankruptcy filings disappear from one's credit score after 10 years. For tractability, we model the removal of the bankruptcy "flag" probabilistically. We denote by $\xi \in [0, 1]$ the likelihood of the bankruptcy flag disappearing tomorrow; having $m = 0$ does not prohibit the household from borrowing or defaulting. Even though current US law prohibits using Chapter 7 bankruptcy more than once every 8 years, the presence of Chapter 13 (which can be used every nine months if the borrower meets an asset test) and the option of informal bankruptcy (see Athreya *et al.* 2012) means that this abstraction, which is convenient computationally, is reasonable.

A second tracking device is introduced in this paper, namely a "credit score" μ (although perhaps not obvious, the quantitative part of the paper will show that the credit score and the bankruptcy flag represent distinct pieces of information). μ represents the probability, as assessed by the market, that any particular individual is the bad type, and evolves according to some function to be detailed later. For future reference, a "bad type" in this model is a household with high λ , since that individual does not regard bankruptcy as particularly

³As in our previous work, we ignore regulations that require certain characteristics not be reflected in credit terms, such as those proscribed by the Equal Credit Opportunity Act. The economic effects of the ECOA are the subject of ongoing work.

⁴See Aiyagari and McGrattan (1998) for a discussion of renewal events in incomplete market models.

painful. Because types shift over time according to the Markov process and agents have only finite lifetimes, it is not the case that μ will converge to 0 or 1.

2.1 Income, Consumption, and Financial Market Arrangement

The timing of decisions within a period is as follows. Agents first draw shocks to their current period income. Log labor income is the sum of four terms: the aggregate wage index $\log(W)$, a deterministic age term $\log \omega_{j,y}$, a persistent shock $\log(e) \in \mathcal{E}$ that evolves as an AR(1)

$$\log(e_j) = \rho \log(e_{j-1}) + \epsilon_j, \quad (2)$$

and a purely transitory shock $\log(\nu) \in \mathcal{V}$; the parameters of the processes for e and ν , as well as the path for $\omega_{j,y}$, depend on the permanent human capital indicator $y \in \mathcal{Y}$ realized prior to entry into the labor market). Both ϵ and $\log(\nu)$ are independent mean zero normal random variables with variances (σ_ϵ^2 and σ_ν^2) that are y -dependent. In the quantitative exercises, we will interpret y as differentiating between non-high school, high school, and college education levels, as in Hubbard, Skinner, and Zeldes (1994), and the differences in these lifecycle parameters will generate different incentives to borrow across types. In particular, college workers will have higher survival rates and a steeper hump in earnings; the second is critically important as it generates a strong desire to borrow early in the lifecycle, exactly when default is highest. The deterministic age-income terms $\omega_{j,y}$ (as well as the survival probabilities $\psi_{j,y}$) also differ according to the realization of y .

After receiving income shocks and paying their taxes, the household makes a decision regarding bankruptcy: if there is debt maturing in the current period, it may be repudiated. Conditional on the default decision, the household makes a consumption-saving decision and then the period ends. Given the timing and our restriction to one-period debt, if lenders observe all household attributes relevant for predicting default any bad household-

level outcome that can be observed or inferred will immediately be reflected in the terms of credit (to the extent that information is available or inferred by lenders). As a result consumption smoothing in response to bad shocks is more difficult all else equal: credit tightens exactly at times which it is most needed. While this may not be an ideal abstraction, we make this assumption for tractability and because the level of commitment on the part of lenders to not readjusting credit terms to be *ex post* optimal is not easily observed. What *is* observed, however, is that until recently (since the CARD Act of 2009), credit contracts explicitly permitted repricing by the lender at will.⁵

Denote by $b \in \mathcal{B}$ the face value of debt ($b < 0$) or savings ($b > 0$) that matures today. Let primes denote one-period-ahead variables (that is, b' is debt that will mature tomorrow). If the household chooses bankruptcy, all their debts are removed. After the bankruptcy decision, a household's income and asset position for the current period are fully determined. Given this vector, a household of age- j chooses current consumption c , and savings or borrowing, b' . If the household chose bankruptcy at the beginning of the current period, they are prohibited *for this period only* from borrowing or saving.⁶ Given the asset structure and timing described above, and using $\mathcal{D} \in \{0, 1\}$ (defined earlier) to indicate whether an agent elected to file for bankruptcy in the current period or not, the household budget constraint during working age and *prior* to the bankruptcy decision is given by

$$c + q(b', I)b'(1 - \mathcal{D}) + \Delta\mathcal{D} \leq (1 - \mathcal{D})b + (1 - \tau)W\omega_{j,y}e\nu, \quad (3)$$

where $q(\cdot)$ is the locus of bond prices (a pricing function) that awaits an individual with

⁵Tam (2009) considers the consequences of permitting lenders to offer contracts that are fixed for up to two periods (that is, they get the same q function in the second period of the contract, so that the change in the state does not alter the cost of borrowing a given level of debt); he finds that few households would choose such contracts because they imply high interest rates, so we feel justified in ignoring lender commitment.

⁶Exclusion in the filing period follows the literature (Livshits, MacGee, and Tertilt 2007), and reflects the legal practice of debtors facing judgements for fraudulent bankruptcies. Unlike this literature, however, we do not impose exclusion in any subsequent period.

characteristics I that are *directly* observable (i.e. objects that do not need to be inferred from behavior). When the household saves ($b > 0$) it receives the risk-free price $q = 1/(1+r)$. The budget constraint during retirement is

$$c + q(b', I)b'(1 - \mathcal{D}) + \Delta \mathcal{D} \leq (1 - \mathcal{D})b + \theta W \omega_{j^*-1, y} e_{j^*-1} \nu_{j^*-1} + \Theta W, \quad (4)$$

where for simplicity we assume that pension benefits are composed of a fraction $\theta \in (0, 1)$ of income in the last period of working life plus a fraction $\Theta \in (0, 1)$ of average income (which has been normalized to one). There are no markets for insurance against any of the stochastic shocks.⁷

The recursive formulation of the household problem is therefore

$$v(b, e, \nu, \lambda, j, m, y, \mu) = \max \{V^{\mathcal{D}=0}(b, e, \nu, \lambda, j, m, y, \mu), V^{\mathcal{D}=1}(e, \nu, \lambda, j, m, y, \mu)\}$$

where

$$V^{\mathcal{D}=0}(b, e, \nu, \lambda, j, m, y, \mu) = \max_{b', c} \left\{ \frac{n_j}{1 - \sigma} \left(\frac{c}{n_j} \right)^{1 - \sigma} + \beta \psi_{j, y} E[v(b', e', \nu', \lambda', j + 1, m', y, \mu')] \right\}$$

and

$$V^{\mathcal{D}=1}(e, \nu, \lambda, j, m, y, \mu) = \frac{n_j}{1 - \sigma} \left(\frac{\lambda c}{n_j} \right)^{1 - \sigma} + \beta \psi_{j, y} E \left[\begin{array}{c} \xi v(0, e', \nu', \lambda', j + 1, 0, y, \mu') + \\ (1 - \xi) v(0, e', \nu', \lambda', j + 1, 1, y, \mu') \end{array} \right].$$

⁷The savings of deceased households is assumed to be taxed at 100 percent and used to finance some wasteful government spending.

2.2 Government Budget Constraint

The only purpose of government in this model is to fund pension payments to retirees using a proportional tax on labor earnings, τ . The government budget constraint is

$$\begin{aligned} & \tau W \left(\sum_{y,e,\nu,\lambda,j,m} \int_{b,\mu} \omega_{j,y} e \nu d\Gamma(b, y, e, \nu, \lambda, j < j^*, m, \mu) \right) \\ &= W \sum_{y,e,\nu,\lambda,j,m} \int_{b,\mu} (\theta \omega_{j^*-1,y} e_{j^*-1} \nu_{j^*-1} + \Theta) d\Gamma(b, y, e, \nu, \lambda, j \geq j^*, m, \mu). \end{aligned} \quad (5)$$

2.3 Loan Pricing

We now detail the construction of the equilibrium prices that will be quoted to agents attempting to borrow b' . All households take loan prices as parametric, and given by the function $q(b', I)$. Recall that I denotes the information *directly* observable to a lender. In the full information case, I includes all components of the household state vector $I = (y, e, \nu, \mu, \lambda, j, m)$, while only a subset of these variables are directly observed under asymmetric information. Under both symmetric and asymmetric information, we focus on competitive lending arrangements in which lenders must have zero profit opportunities.

With full information, a variety of pricing arrangements will, under competitive conditions, lead to the same price function. Full information is also the case previously studied in the literature (see Chatterjee *et al.* 2007 or Livshits, MacGee, and Tertilt 2010), and is a special case of our model. In contrast, under asymmetric information it is well known that outcomes often depend on the particular “microstructure” being used to model the interaction of lenders and borrowers (Hellwig 1989). Specifically, since we have modeled households as issuing debt to the credit market, we must take into account the fact that the size of any debt issuance itself conveys information about the household’s current state. In other words, we study a signaling game in which loan size b' is the signal. As we will detail further below,

the lender's task is to form estimates of the current realizations of the two persistent shocks (e, λ) , given this signal. Given an estimate and knowledge of household decision-making, lenders can then compute the likelihood of default, and in turn, the conditional expectation of profits obtaining from any loan price q they may ask of the borrower.

Lenders and borrowers play a two-stage game. In the first stage, borrowers name a level of debt b' that they wish to issue in the current period. Second, a continuum of lenders compete in an auction where they simultaneously post a price for the desired debt issuance of the household and are committed to delivering the amount b' in the event their "bid" is accepted; that is, the lenders are engaging in Bertrand competition for borrowers. In equilibrium borrowers choose the lender who posts the highest q (lowest interest rate, $r = q^{-1} - 1$) for the desired amount of borrowing. Thus, households view the pricing functions as schedules and understand how changes in their desired borrowing will alter the terms of credit because they compute the locus of Nash equilibria under price competition.

The Bertrand competition spelled out above leads to equilibrium prices that must not permit lenders to do better than break even on loans, given their (common) estimate of default risk, once that estimate is updated to reflect the signal sent by households. Let $\hat{\pi}^{b'}: b'_I \rightarrow [0, 1]$ denote the function that provides the best estimate of the probability of default, conditional on surviving, to a loan of size b' under information regime I . Since default is irrelevant for savers, $\hat{\pi}^{b'}$ is identically zero for positive levels of net worth. In contrast, $\hat{\pi}^{b'}$ is equal to 1 for all debt levels exceeding some sufficiently large threshold. Given the exogenous risk-free saving rate r , let ϕ denote the proportional transaction cost associated with lending, so that $r + \phi$ is the risk-free borrowing rate; the pricing function takes into account the automatic default by those households that die at the end of the period (and we implicitly assume any positive accidental bequests are used to finance some wasteful government spending).

Given any $\widehat{\pi}^{b'}$, the break-even pricing function must satisfy

$$q(b', I) = \begin{cases} \frac{1}{1+r} & \text{if } b \geq 0 \\ \frac{(1-\widehat{\pi}^{b'})\psi_{j,y}}{1+r+\phi} & \text{if } b < 0 \end{cases} \quad (6)$$

2.3.1 Full Information

In the full information setting, given debt issuance b' and knowledge of the two *persistent* shocks e and λ , the lender does not actually need to know the current realizations of the transitory shock and expenditure shock as they will not help forecast next period's realization of the household's state. We include them here to maintain consistent notation with the partial information setting to be detailed further below.

Zero profit for the intermediary requires that the probability of default used to price debt must be consistent with that observed in the stationary equilibrium, implying that

$$\widehat{\pi}^{b'} = \sum_{e', \nu', \lambda'} \pi_e(e'|e) \pi_\nu(\nu') \pi_\lambda(\lambda'|\lambda) \mathcal{D}(b', e', \nu', \lambda', \mu'). \quad (7)$$

Since summing $\mathcal{D}(b', e', \nu', \lambda', \mu')$ yields the probability that the agent will default in state (e', ν', λ') tomorrow given current loan request b' , integrating over all such events *tomorrow* is what is necessary to estimate relevant default risk. This expression also makes clear that knowledge of the persistent components e and λ is critical for predicting default probabilities; the more persistent they are, the more useful knowledge of their current values becomes in assessing default risk. We now turn to loan pricing under partial information.

2.3.2 Partial Information

“Partial Information” in our model refers to cases in which lenders cannot directly observe the stigma term λ . Knowledge of λ would, given b' , collapse the model to the “Full-Information”

case as one could invert the decision rule to obtain the unique value of λ consistent with the optimality of b' given the other states (absent some zero probability flat spots in the decision rule). To capture the effects of limits on information held by creditors about households, we also rule out the ability of lenders to directly observe objects that through household decision rules, would be completely informative about a household's current λ , even when these objects are not *themselves* useful for forecasting default risk, in particular current debt b .

The signaling approach adopted here introduces problems with multiplicity of equilibria: there will in general be many outcomes satisfying the requirements of equilibrium. Multiplicity arises because the standard solution concept appropriate for games of incomplete information, Perfect Bayesian Equilibrium, does not fully discipline the beliefs players hold *off* the path of equilibrium play. As a result, given enough freedom to pick beliefs the modeler can deliver many outcomes. In our context, this freedom means that even in cases when the nature of private information is *known a priori* to be genuinely irrelevant – for example, when all currently privately-held information governing the costs and benefits of default costs one period ahead involves iid random variables) – one can still deliver equilibria in which such private information is “led” to matter simply by imposing off-equilibrium-path beliefs (in)appropriately.

Since our contribution is partly linked to the manner in which we select equilibria, something dependent on the iterative procedure we employ, it is important to check the nature of outcomes that are selected and the nature of the off-equilibrium-path beliefs that are *induced* by our algorithm. We check if our iterative procedure leads to an outcome in which private information is falsely led to matter in the iid case detailed above, and find that it does not. This test provides us with confidence that we are selecting off-equilibrium beliefs in a reasonable manner.

The lender's problem is to infer the λ of a household who has requested a loan of b' .

If lenders cannot directly observe the discount factor they must resort to constructing an estimate of λ for a household requesting b' . To make the best possible inference regarding λ , lenders will also use any knowledge they have of the decision rules for households and the distribution of households over the states.

To deal with the inference problem tractably, we will restrict attention throughout to stationary equilibria. Let this stationary joint distribution be denoted by $\Gamma(b, y, e, \nu, \lambda, j, m, \mu)$, and let the household decision rule for optimal borrowing be given by $b' = g(b, y, e, \nu, \lambda, j, m, \mu)$. For any *given* λ , not all values of the state vector are consistent with the observables, particularly the loan request b' . Therefore, let $S_{\{\lambda\}}$ denote the set of values for remaining household characteristics that, for a household with given λ , *could be* consistent with the observable vector $(b', y, e, \nu, j, m, \mu)$. This set is therefore the (possibly set-valued) pre-image of $g(\cdot)$ for a given λ .

Let $\Pr(\lambda|b', y, e, \nu, j, m, \mu)$ denote the probability of a borrower with observables $(b', y, e, \nu, j, m, \mu)$ having current state λ based on knowledge of the decision rules of agents. In the partial information environment the calculation of $\Pr(\lambda|b', y, e, \nu, j, m, \mu)$ is nontrivial because it involves the distribution of endogenous variables. In a stationary equilibrium, the conditional probability of a household having any particular λ , given both observables and lender inference from decision rules, is as follows.

$$\Pr(\lambda|b', y, e, \nu, j, m, \mu) = \int_{S_{\{\lambda\}}} d\Gamma(b, y, e, \nu, \lambda, j, m, \mu) \quad (8)$$

Given this assessment, the lender can compute the likelihood of default on a loan of size b' :

$$\hat{\pi}^{b'} = \sum_{\lambda} \left[\sum_{e', \nu', \lambda'} \pi_e(e'|e) \pi_{\nu}(\nu') \pi_{\lambda}(\lambda'|\lambda) d(b', e', \nu', \lambda', \mu') \right] \Pr(\lambda|b', y, e, \nu, j, m, \mu). \quad (9)$$

2.3.3 Equilibrium

Our specification of the game between borrowers and lenders makes it essentially identical to the standard education-signaling model as described in Mas-Colell, Whinston, and Green (1995), though with an important difference: in our model the signal is productive. Namely, debt not only may convey information about one's type, it allows for intertemporal and inter-state transfers of purchasing power that agents would have chosen even under completely symmetric information. The game described above is a standard game of incomplete information, and as such, comes with a standard equilibrium concept: Perfect Bayesian Equilibrium (PBE).

Intuitively, a PBE is given by a set of beliefs and strategies for households and lenders, where the strategies for each are each optimal given the (common) beliefs they hold, and the beliefs they each hold are, in turn, derived rationally from the strategies they each (correctly) anticipate their opponents will use. In our model, optimal household strategy concerns the amount of debt to issue, and beliefs are over the prices they believe the loan market will charge them for all debt levels they may consider. In the case of lenders, optimal strategy concerns deciding what price to post for any debt request given their beliefs and the Bertrand game they each play with all other lenders. Lenders are required to use the decision rule for households' optimal debt issuance and Bayes' rule *wherever possible* to arrive at a posterior probability over a household's type, given the observables, as any other belief about the behavior of households will not survive in equilibrium.

A Perfect Bayesian Equilibrium (PBE) (see Mas-Colell, Whinston, Green, Definition 9.C.3, p. 285, and the additional requirement given on p. 452) is as follows. Denote the state space for households by $\Omega = \mathcal{B} \times \mathcal{Y} \times \mathcal{E} \times \mathcal{V} \times \mathcal{L} \times \mathcal{J} \times \{0, 1\} \times \mathcal{M} \subset \mathcal{R}^4 \times \mathcal{Z}_{++} \times \{0, 1\} \times [0, 1]$ and space of information as $\mathcal{I} \subset \mathcal{Y} \times \mathcal{E} \times \mathcal{V} \times \mathcal{L} \times \mathcal{J} \times \{0, 1\} \times \mathcal{M}$. Let the stationary joint distribution of households over the state be given by $\Gamma(\Omega)$. Let the stationary equilibrium

joint distribution of households over the state space Ω and loan requests b' be derived from the decision rules $\{b^*(\cdot), d^*(\cdot)\}$ and $\Gamma(\Omega)$, and be denoted by $\Psi^*(\Omega, b')$. Given $\Psi^*(\Omega, b')$, let $\psi^*(b')$ be the fraction of households (i.e., the marginal distribution of b') requesting a loan of size b' . Lastly, let the common beliefs of lenders on the household's state, Ω , given b' , be denoted by $\Upsilon^*(\Omega|b')$.⁸

Definition A PBE for the credit market game of incomplete information consists of (i) household strategies for borrowing $b^* : \Omega \rightarrow \mathcal{R}$ and default $d^* : \Omega \times \mathcal{E} \times \mathcal{V} \times \mathcal{L} \rightarrow \{0, 1\}$, (ii) lenders' strategies for loan pricing $q^* : \mathcal{B} \times \mathcal{I} \rightarrow [0, \frac{1}{1+r}]$ such that q^* is weakly decreasing in b' , and (iii) lenders' common beliefs about the borrower's state Ω given a loan request of size b' , $\Upsilon^*(\Omega|b')$, that satisfy the following:

1. **Households optimize:** Given lenders' strategies, as summarized in the locus of prices $q^*(b', I)$, decision rules $\{b^*(\cdot), d^*(\cdot)\}$ solve the household problem.
2. **Lenders optimize given their beliefs:** Given common beliefs $\Upsilon^*(\Omega|b')$, q^{I*} is the pure-strategy Nash equilibrium under one-shot simultaneous-offer loan-price competition.
3. **Beliefs are consistent with Bayes' rule wherever possible:** $\Upsilon^*(\Omega|b')$, is derived from $\Psi^*(\Omega, b')$ and household decision rules using Bayes rule whenever b is such that $\psi^*(b') > 0$.

2.3.4 Locating Equilibria

We now describe how we find outcomes that satisfy the conditions laid out above. A proof of the convergence of our algorithm can be found in the Appendix.

⁸Recall that the stationary distribution of households over the state space alone is given by $\Gamma(\cdot)$.

Off-Equilibrium-Path Beliefs From the household’s perspective, however, it is essential that they know the price they will face at *any* debt level they might contemplate; it is only then that they can solve a well-posed optimization problem. So what is a household to expect that a lender will infer about them should they contemplate issuing a debt level no one is expected to choose in a proposed equilibrium? A trivial example in our context is the PBE in which lenders believe that all borrowers will default with probability one on any debt, and no one borrows. Given the ability of such pessimism to be self-fulfilling in signaling models, we structure our iterative procedure to avoid limiting borrowing in such a manner. Since our process for locating equilibrium price functions is one that generates a monotone sequence of functions, points at which pricing functions reach zero will remain there in all subsequent iterations.

The component of our equilibrium for which off equilibrium-path beliefs are germane is the pricing function q^* . This function is derived as the fixed point of a mapping that we describe further below. The locus q^* describes for each agent type (where “type” denotes all directly observable characteristics of a borrower) what pricing they can expect, and what default risk lenders expect, at all debt levels. However, this pricing function, even if taken as given by all participants, will not necessarily lead to all debt levels being chosen by any given agent type. Thus, in each of these places, any value taken by the function q^* necessarily reflects some off-equilibrium-path beliefs on the part of participants.

The other study of credit scoring, Chatterjee, Corbae, and Ríos-Rull (2011), assumes that agents randomly make mistakes and thus take all feasible actions with positive probability. Because they have only four possible asset levels, all actions are feasible for all agents at all times, and thus they never need impose any off-equilibrium-path beliefs. Clearly, such assumptions are inconsistent with the data.

Finding Equilibrium Loan Prices The only tractable way to locate equilibria is iteratively. Any iterative procedure, in turn, necessarily requires, and ultimately influences, a set of off equilibrium-path beliefs, and we now describe how our approach does so. We begin our iterations, for reasons described in more detail below, by positing that lenders hold the most optimistic views they could have about a borrower: they literally ignore the possibility of default until it becomes a certainty. That is, the implied beliefs for the intermediary are such that default is predicted to never occur except when it must (to allow the household to obtain positive consumption) in every state of the world. As a result, credit availability at the outset of our iterations, defined in terms of the interest rate for any loan size that one might obtain, is maximized under this specification.⁹ This is because debt levels in excess of what is observed under this initial pricing function will not be chosen when prices are (weakly) higher, which they will be in any equilibrium. While we make no formal qualitative claims about the extent of credit availability in equilibrium, the manner in which update pricing for debt helps us preserve credit access along the path to, as well as under, the equilibrium pricing function.

Given a pricing function, households make decisions about debt. For all debt levels selected by positive fraction of households, we can then compute the likelihood of default as a function of their observable type. These are Steps 1-5 in the algorithm. Next, for debt levels that lie in any range not chosen by any households, we proceed as follows. Let k be the index of iteration, let $q^{(k)}$ and be the current price function, and $\psi^{(k)}$ be the associated distribution of debt requests induced by agent decisions, given $q^{(k)}$. For any given debt request b' such that $\psi^{(k)}(b') = 0$, define $b' < b^- < 0$ as the nearest lower debt level relative to b' at which $\psi^{(k)}(b') > 0$. b^- is therefore the upper end (or right-endpoint) of a segment of

⁹It is useful to compare our initial pricing function with the natural borrowing limit, the limit implied by requiring consumption to be positive with probability 1 in the absence of default. Our initial debt limit is larger than the natural borrowing limit, as agents can use default to keep consumption positive in some states of the world; we only require that they not need to do this in *every* state of the world. This point is also made in Chatterjee *et al.* (2007).

debt levels no household requests, given $q^{(k)}$. Define $\hat{q}(b^-, I)$ to be the actuarially fair price at b^- ; this is where we use optimal inference by lenders to construct $\hat{\pi}^{b'}$.

Next, define $b^+ < b' < 0$ to be the nearest higher debt level (i.e. a more negative value for assets) relative to b' at which $\psi^{(k)}(b^+) > 0$. b^+ is therefore the lower end (or left-endpoint) of the same segment of debt levels, again given $q^{(k)}$. Denote by $\hat{q}(b^+, I)$ the actuarially fair price at b^+ . Thus, at any b , $\hat{q}(b, I)$ is the actuarially fair price at any debt level that is observed under $\psi^{(k)}$, and equal to $\hat{q}(b^-, I) \forall b \in (b^+, b^-]$. The collection of these segments is then denoted $\hat{q}(\cdot, I)$.¹⁰ With $\hat{q}$ in hand, as well as the current iterate for the pricing function $q^{(k)}$, we construct the next iterate of the pricing function, $q^{(k+1)}$ as a convex combination of $q^{(k)}$ and $\hat{q}(\cdot, I)$ that places weight Ξ on $q^{(k)}$. The preceding is Step 6.

We want to point out here that we do *not* confront the agent with the function $\hat{q}$; we cannot guarantee that the sequence of $\hat{q}$ functions is monotone. Instead, we update the pricing functions extremely slowly using a weight Ξ on the current iterate $q^{(k)}$ close to one (we use $\Xi = 0.985$), and thereby obtain a monotone sequence of pricing functions (see the proof of convergence in the Appendix where this property is used). This procedure constitutes Steps 7 and 8 in the algorithm. Repeating this procedure to convergence, we obtain the equilibrium pricing function and loan request distribution. In addition, note that the segments starting at b^- will, in subsequent iterations, lead those there to potentially lower their borrowing slightly in the next iteration, as they receive discontinuously better pricing for doing so – even under the heavily convexified pricing function we will face them with at that point.¹¹ As a result, some pooling can occur in equilibrium as agents of different types are inframarginal at any discrete jumps in loan pricing.

Given the definition of equilibrium for the game, a stationary equilibrium for the overall model is standard: we simply require that the pricing functions and distribution of house-

¹⁰Note that $\hat{q}$ is a (lower semicontinuous) step-function.

¹¹We thank a referee for bringing this to our attention.

holds over the state space are invariant under the optimal decision-making described above, and that the tax rate τ allows the government to meet its budget constraint.

Finding Equilibrium Credit Scores Having obtained q^* , we can now assess the agents' perceptions about how their credit score will evolve. The law of motion for μ is

$$\mu' = \hat{\mu} \Pr(\lambda' = \lambda^{hi} | \lambda = \lambda^{hi}) + (1 - \hat{\mu}) \Pr(\lambda' = \lambda^{hi} | \lambda = \lambda^{lo}). \quad (10)$$

$\hat{\mu}$ is the posterior belief that the household is the bad type given $(b', e, \nu, j, m, y, \mu)$, which equals

$$\hat{\mu} = \frac{\Pr(\lambda = \lambda^{hi} | b', e, \nu, j, m, y, \mu)}{\Pr(\lambda = \lambda^{hi} | b', e, \nu, j, m, y, \mu) + \Pr(\lambda = \lambda^{lo} | b', e, \nu, j, m, y, \mu)}.$$

To compute $\hat{\mu}$ we use the stationary distribution of households to construct the probability of observing a given vector for an individual who is actually a bad type (in the same way we use it to construct q^*), then update the prior μ to the posterior $\hat{\mu}$, which then feeds into the law of motion for μ .

3 Results

We first present results that compare the 'FI' allocation to the 'PI' allocation where both m and μ are observed by lenders, and λ is not. We focus first on the role of information for default rates and bond prices, then discuss the dynamic behavior of μ . Finally, we consider allocations where m is absent, where μ is absent, and where m is permanently attached to an individual.

3.1 Aggregate Results

Table 2 compares the outcomes under FI and PI (with credit scoring) for three facts about the credit market (which we use for calibration): the bankruptcy rate for each education type, the median amount of discharged debt relative to the median income of a filer, and the median income of a filer relative to the median income of borrowers.

Table 2: Aggregate Statistics

	FI	PI
BK Rate NHS	0.64%	0.53%
BK Rate HS	0.68%	0.54%
BK Rate College	0.47%	0.39%
Median Discharge / Median Income of Filers	0.283	0.266
Median Income of Filers / Median Income of Borrowers	0.690	0.655

From the table we see that improved information leads to a rise in bankruptcy rates across the board, with the largest absolute increase being for HS-educated households. We also see that discharge increases relative to income for filers, so they are in worse shape under full information when they file, but their income relative to other borrowers also increases so it is not the "worst" borrowers who are filing.

Table 3 shows the interest rate spreads that a type without the bankruptcy flag receives

on average.

Table 3: Mean Spread for Good Flag Borrowers

	λ^{hi}			λ^{lo}		
	NHS	HS	College	NHS	HS	College
Price Getting	6.79%	5.47%	4.29%	7.37%	6.52%	4.76%
Price Supposed to Get	6.53%	5.32%	4.24%	7.83%	6.76%	4.91%
Difference	0.26%	0.15%	0.05%	-0.46%	-0.24%	-0.15%

From this table we see two facts. First, equilibrium rates differ across the types by less than "what they should get" based on their individual default probabilities. In other words, the asymmetric information environment is providing some insurance against being the bad type.

Second, the gap between the interest rate they get and the interest rate they should get is positive for the high λ type (the bad type) and negative for the low λ type, and this gap (in absolute value) is smaller for high-educated types. This result is to be expected given standard adverse selection models, the bad type benefits at the expense of the good type, unless the market unravels completely.¹²

¹²As we noted in our previous paper, complete unraveling ($q = 0$) is an equilibrium. However, it is not one consistent with notions of forward induction like those developed in Cho and Kreps (1987) or Banks and Sobel (1987).

In contrast, Table 4 shows the spreads for borrowers carrying the bankruptcy flag.

Table 4: Mean Spread for Bad Flag Borrowers

	λ^{hi}			λ^{lo}		
	NHS	HS	College	NHS	HS	College
Price Getting	8.28%	6.29%	4.72%	8.98%	7.01%	5.08%
Price Supposed to Get	7.95%	6.12%	4.66%	9.05%	7.26%	5.20%
Difference	0.33%	0.17%	0.06%	-0.07%	-0.15%	-0.12%

The gaps for the bad type are larger and smaller for the good type. The gap is also not monotonic in education for the good type, as HS types get the largest decrease.

3.2 Credit Scores

Figures (2)-(5) present the dynamic behavior of μ . For each panel, we plot the average dynamics of the credit score over time, conditional on an individual being a borrower at time $t = 1$. In Figure 1 we see that bad types (high λ) without bankruptcy flags tend to have their credit scores improve (mean reversion) and then deteriorate, while good types without bankruptcy flags have the opposite pattern. Bad types with bankruptcy flags in general see declines in their credit scores, and good types with bankruptcy flags can expect improvements. It is clear from these figures that μ and m reflect different information about the borrower.

Musto (2004) presents evidence that the removal of a bankruptcy flag significantly improves the terms at which a household can borrow. In Figure (6), we present the interest rate charged in the periods leading up to a flag removal as well as those periods immediately following the removal, averaged over all borrowers. In all cases the terms improve dramatically, even when additional information regarding the default probability (namely μ) is still

present. Figure (7) shows that the credit score of these households also improves.

3.3 Unobserved Income

Now suppose that e is unobservable. Given new debt issuance b' , default decision $\mathcal{D}$, and current μ , as well as invariant distribution $\Gamma(b, e, \nu, \lambda, j, m, \mu)$, we have that

$$\mu' = \psi(b', \mathcal{D}, b, \nu, j, m, \mu) = \sum_e \tilde{\Gamma}(e, \lambda_1 | b', d, b, \nu, j, m, \mu),$$

which gives the fraction of households with $\lambda = \lambda^{hi}$ who would choose b' and $\mathcal{D}$ given they have (b, ν, j, m, μ) , integrated over the unknown e ; obviously if $\mathcal{D} = 1$ then the only consistent asset choice is $b' = 0$. We use this section to examine whether our results regarding the relative unimportance of private information on income still hold when credit scoring is permitted.

3.4 Changing Information

We next compare allocations that differ with respect to the amount of information present in the credit market. First, we shut off μ , resulting in the PI equilibrium from Athreya, Tam, and Young (2012) without expenditure shocks. Second, we shut off m , so that past defaults are not observed. Third, we make m permanent, so that any past default is recorded permanently.

We calculate the welfare effects of information by asking how much a newborn agent would pay to live in the FI economy rather than any of the other economies. As in our previous work (Athreya, Tam, and Young 2009), we interpret our welfare calculations through a simple decomposition of lifetime consumption variance:

$$\text{var}(\log(C)) = E[\text{var}(\log(C) | j)] + \text{var}(E[\log(C) | j]); \quad (11)$$

that is, the variance of (log) consumption over an agent’s lifetime is equal to the average of the variance of consumption over ages plus the variance of average consumption over the lifecycle. We refer to these components as the ‘intratemporal’ and ‘intertemporal’ components, respectively.¹³

4 Conclusion

This paper has investigated the value of dynamic record-keeping in unsecured credit markets. We are currently working on versions that incorporate regulatory restrictions on information – such as the Equal Credit Opportunity Act and the Fair Credit Reporting Act – that limit the conditioning information that lenders can use.

References

- [1] **Aiyagari, S. Rao, and Ellen R. McGrattan.** 1998. "The Optimum Quantity of Debt." *Journal of Monetary Economics* 42(3): 447-69.
- [2] **Athreya, Kartik, Juan M. Sánchez, Xuan S. Tam, and Eric R. Young.** 2012. "Bankruptcy and Delinquency in a Model of Unsecured Debt." Working Paper 2012-042A, Federal Reserve Bank of St. Louis.
- [3] **Athreya, Kartik, Xuan S. Tam, and Eric R. Young.** 2009. "Unsecured Credit Markets Are Not Insurance Markets." *Journal of Monetary Economics* 56(1): 83-103.
- [4] **Athreya, Kartik, Xuan S. Tam, and Eric R. Young.** 2012. "A Quantitative Theory of Information and Unsecured Credit." *American Economic Journal: Macroeconomics* 4(3): 153-83.

¹³Obviously these two moments do not exhaustively determine the welfare consequences of changing information; we simply use them as a useful ‘rule of thumb’.

- [5] **Banks, Jeffrey S., and Joel Sobel.** 1987. Equilibrium Selection in Signalling Games. *Econometrica* 55(3), 647-662.
- [6] **Chatterjee, Satyajit, P. Dean Corbae, Makoto Nakajima, and José-Víctor Ríos-Rull.** 2007. "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default." *Econometrica* 75(6): 1525-90.
- [7] **Chatterjee, Satyajit, P. Dean Corbae, and José-Víctor Ríos-Rull.** 2010. "A Theory of Credit Scoring and Competitive Pricing of Default Risk." Mimeo, Federal Reserve Bank of Philadelphia, University of Texas-Austin, and University of Minnesota.
- [8] **Cho, In-Koo, and David M. Kreps.** 1987. Signalling Games and Stable Equilibria. *Quarterly Journal of Economics* 102(2), 179-221.
- [9] **Davey, Brian A., and Hilary A. Priestley.** 2002. *Introduction to Lattices and Order, Second Edition.* Cambridge Press.
- [10] **D'Erasmus, Pablo.** 2011. "Government Reputation and Debt Repayment." Mimeo, University of Maryland.
- [11] **Fay, Scott A., Erik Hurst, and Michelle J. White.** 1998. "The Bankruptcy Decision: Does Stigma Matter?" University of Michigan Working Paper 98-01.
- [12] **Granas, Andrzej, and James Dugundji.** 1986. *Fixed Point Theory.* Springer Publishing.
- [13] **Gross, David B. and Nicholas S. Souleles.** 2002. "An Empirical Analysis of Personal Bankruptcy and Delinquency." *Review of Financial Studies* 15(1): 319-47.
- [14] **Hellwig, Martin.** 1989. "Asymmetric Information, Financial Markets, and Financial Institutions: Where Are We Currently Going?" *European Economic Review* 33(2-3): 277-85.

- [15] **Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes.** 1994. "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving." *Carnegie-Rochester Conference Series on Public Policy* 40: 59-126.
- [16] **Ionescu, Felicia, and Nicole Simpson.** 2011. "Credit Scores and College Investment: Implications for Student Loan Policies." Mimeo, Colgate University.
- [17] **Livshits, Igor, James MacGee, and Michèle Tertilt.** 2011. "Costly Contracts and Consumer Credit." Mimeo, University of Western Ontario and Stanford University.
- [18] **Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green.** 1995. *Microeconomic Theory*. Oxford University Press.
- [19] **Musto, David K.** 2004. "What Happens When Information Leaves a Market? Evidence from Postbankruptcy Consumers." *Journal of Business* 77(4): 725-48.
- [20] **Narajabad, Borghan.** 2012. "Information Technology and the Rise of Household Bankruptcy." *Review of Economic Dynamics* 15(4): 526-50.
- [21] **Sánchez, Juan M.** 2010. "The IT Revolution and the Unsecured Credit Market." Working Paper 2010-022, Federal Reserve Bank of St. Louis.
- [22] **Tam, Xuan S.** 2009. "Long-Term Contracts in Unsecured Credit Markets." Mimeo, University of Cambridge.

Figure 1: FICO Scores and Default Rates




Figure 2: Dynamics of Credit Score




Figure 3: Dynamics of Credit Score




Figure 4: Dynamics of Credit Score




Figure 5: Dynamics of Credit Score




Figure 6: Mean Pricing Before and After Flag Removal





Figure 7: Mean Credit Score Before and After Flag Removal



5 Appendix

5.1 Computing Partial Information Equilibria

The imposition of conditions on beliefs off-the-equilibrium path makes the computational algorithm we employ relevant for outcomes, so we discuss in some detail our algorithm for computing partial information competitive equilibria. The computation of the full information equilibrium is straightforward using backward induction; since the default probabilities are determined by the value function in the next period, we can solve for the entire equilibrium, including pricing functions, with one pass. The partial information equilibrium is not as simple, since the lender beliefs regarding the state of borrowers influence decisions and are in turn determined by them; an iterative approach is therefore needed.

1. Fix an agent type by observables (y, e, ν, j, m, μ)
2. Guess the initial function $q^{(0)}(b', y, e, \nu, j, m, \mu)$ discussed above
3. Solve household problem to obtain $b' = g(b, y, e, \nu, \mu, \lambda, j, m)$, $S_{\{\lambda\}}$, and $d(b', e', \nu', \mu', \lambda')$, and $\Gamma(\cdot)$.
4. For all b' observed, use $S_{\{\lambda\}}$ and $\Gamma(\cdot)$ to compute $\Pr(\lambda|b', y, e, \nu, j, m, \mu)$ the probability that an individual has stigma cost λ given observed $(b', y, e, \nu, j, m, \mu)$. Knowledge of $\Pr(\lambda|b', y, e, \nu, j, m, \mu)$ and the distribution of households over the remaining observable state variables implies $\Upsilon(\Omega|b')$.
5. Compute

$$\hat{\pi}^b(b', y, e, \nu, j, m, \mu) = \sum_{\lambda} \pi^{b'}(b', y, e, \nu, \mu, \lambda, j, m) \Pr(\lambda|b', y, e, \nu, j, m, \mu), \quad (12)$$

the expected probability of default for an individual in observed state $(b', y, e, \nu, j, m, \mu)$;

6. Fill in the “holes” in $\hat{\pi}^{b'}$ for all b' not observed by applying the interim off-equilibrium beliefs as described in the main text.
7. Compute an “intermediate” price function $\hat{q}$ for all b' , that is actuarially fair (competitive) given the preceding estimate $\hat{\pi}^{b'}$:

$$\hat{q}(b', y, e, \nu, j, m, \mu) = \frac{\left(1 - \hat{\pi}^{b'}(b', y, e, \nu, j, m, \mu)\right) \psi_j}{1 + r + \phi} \quad \text{for all } b \geq b_{\min}(b, y, e, \nu, j, m, \mu); \quad (13)$$

8. Set

$$q^{(1)}(b', y, e, \nu, j, m, \mu) = \Xi q^{(0)}(b', y, e, \nu, j, m, \mu) + (1 - \Xi) \hat{q}(b', y, e, \nu, j, m, \mu)$$

where Ξ is set very close to 1 (we use 0.985), return to Step 1 and repeat until the pricing function converges.

Given a pricing function, households make decisions about debt. For all debt levels selected by positive fraction of households, we can then compute the likelihood of default as a function of their observable type. These are Steps 1-5 in the algorithm. Next, for debt levels that lie in any range not chosen by any households, we proceed as follows. Let k be the index of iteration, let $q^{(k)}$ and be the current price function, and $\eta^{(k)}$ be the associated distribution of debt requests induced by agent decisions, given $q^{(k)}$. For any given debt request b' such that $\eta^{(k)}(b') = 0$, define $b' < b^- < 0$ as the nearest lower debt level relative to b' at which $\eta^{(k)}(b') > 0$. b^- is therefore the upper end (or right-endpoint) of a segment of debt levels no household requests, given $q^{(k)}$. Define $\hat{q}(b^-, I)$ to be the actuarially fair price at b^- ; this is where we use optimal inference by lenders to construct $\hat{\pi}^{b'}$.

Next, define $b^+ < b' < 0$ to be the nearest higher debt level (i.e. a more negative value for assets) relative to b' , at which $\eta^{(k)}(b^+) > 0$. b^+ is therefore the lower end (or left-endpoint)

of the same segment of debt levels, again given $q^{(k)}$. Denote by $\hat{q}(b^+, I)$ the actuarially fair price at b^+ . Thus, at any b , $\hat{q}(b, I)$ is the actuarially fair price at any debt level that is observed under $\eta^{(k)}$, and equal to $\hat{q}(b^-, I) \forall b \in (b^+, b^-]$. The collection of these segments is then denoted $\hat{q}(\cdot, I)$.¹⁴ With $\hat{q}$ in hand, as well as the current iterate for the pricing function $q^{(k)}$, we construct the next iterate of the pricing function, $q^{(k+1)}$ as a convex combination of $q^{(k)}$ and $\hat{q}(\cdot, I)$ that places weight Ξ on $q^{(k)}$. The preceding is Step 6.

We want to point out here that we do *not* confront the agent with the function $\hat{q}$; we cannot guarantee that the sequence of $\hat{q}$ functions is monotone. Instead, we update the pricing functions extremely slowly using a weight Ξ on the current iterate $q^{(k)}$ close to one (we use $\Xi = 0.985$), and thereby obtain a monotone sequence of pricing functions (see the proof of convergence below). This procedure constitutes Steps 7 and 8 in the algorithm. Repeating this procedure to convergence, we obtain the equilibrium pricing function and loan request distribution. In addition, note that the segments starting at b^- will, in subsequent iterations, lead those households at b^- to potentially lower their borrowing slightly in the next iteration, as they receive discontinuously better pricing for doing so, even under the heavily convexified pricing function we will face them with at that point.¹⁵ As a result, some pooling can occur in equilibrium as agents of different types are inframarginal at any discrete jumps in loan pricing.

We next provide a proof that the numerical procedure has a maximal fixed point, and that this fixed point is the one we obtain. Because the machinery used for this proof may not be familiar to all readers, we provide some basic definitions as well.

¹⁴Note that $\hat{q}$ is a (lower semicontinuous) step-function.

¹⁵We thank a referee for bringing this to our attention.

5.2 Proof of Convergence

Define Q to be a collection of nondecreasing step functions defined over a finite set of points $\mathcal{D} \subset \mathcal{R}$; the number of steps is therefore necessarily finite, so Q itself is finite. Let Q contain a maximal element $\mathbf{1}$ and a minimal element $\mathbf{0}$, using the pointwise ordering of functions ($f \preceq g$ if and only if $f(b) \leq g(b) \forall b \in \mathcal{D}$); suppose all such members can be ordered. Endow Q with the Scott topology, which defines order-continuity: a function $\Phi : Q \rightarrow Q$ is **order-continuous** (or **Scott-continuous**) if $\sup \{\Phi(R)\} = \Phi \{\sup(R)\}$ for each $R \subset Q$. Since Q is finite, order-continuity reduces to pointwise convergence in $\mathcal{R}^n$.

Define $\Phi : Q \rightarrow Q$ as the mapping defined in the algorithm; heuristically, Φ takes a given loan pricing function q^0 , constructs the break-even pricing function $\hat{q}$ implied by behavior of agents confronted with q^0 , and then produces a new pricing function q^1 as a convex combination of the two functions:

$$\Phi(q^0) = \Xi \hat{q}(q^0) + (1 - \Xi) q^0$$

for some $\Xi \in [0, 1]$. Define $\Phi^n(q^0)$ as Φ composed n times.

Lemma 5.1 *$(Q, \preceq)$ is a complete lattice.*

Proof $(Q, \preceq)$ is clearly a lattice, as it is a collection of lower-semicontinuous step functions that are nonincreasing and ordered pointwise. To see that it is a complete lattice, note that an arbitrary collection of lower-semicontinuous functions has a pointwise supremum, and $\mathbf{0}$ bounds all collections from below. Therefore, following Davey and Priestley (2002) $(Q, \preceq)$ is a complete lattice.

We will apply the following theorem (see Granas and Dugundji 1986).

Theorem 5.2 *(Tarski-Kantorovitch) Let $(Q, \preceq)$ be a partially-ordered set and $\Phi : Q \rightarrow Q$ be order-continuous. Assume there exists $q \in Q$ such that (i) $q \geq \Phi(q)$ and (ii) every*

countable chain in $\{x|x \leq q\}$ has an infimum. Then the set of fixed points of Φ is not empty. Furthermore, $q^* = \inf_n \{\Phi^n(q)\}$ is a fixed point and q^* is the maximum of the set of fixed points of Φ in $\{x|x \leq q\}$.

Any complete lattice is a partially-ordered set.

Let $d = 1$ denote the decision to default next period, and $\Pr(d = 1|b)$ denote the probability of default next period conditional on issuing debt b today.

Proposition 5.3 $\Pr(d = 1|b)$ is nondecreasing given $q(b)$. That is, the probability of default weakly rises with the amount of debt, holding fixed pricing.

Proof See Chatterjee *et al.* (2007), Theorem 6.

Conjecture 5.4 $\Pr(d = 1|q)$ is nonincreasing at each b . That is, the probability of default on a given amount of debt is weakly falling in the price of that debt.

This conjecture rules out the possibility that two iterates q^n and q^{n+1} "cross" each other, in that $q^n(b) > q^{n+1}(b)$ for some b and $q^n(b) < q^{n+1}(b)$ for other b ; that is, it implies that $\{\widehat{q}(q^n)\}$ is a monotone chain. The economic content of the conjecture is that the pool of borrowers who choose a particular b level in equilibrium cannot improve as q falls, implying that the equilibrium default rate on that debt level must weakly rise. The conjecture does not appear to be provable in general, but it is satisfied by the numerical procedure we use. Fortunately, the proof of convergence below does not require $\{\widehat{q}(q^n)\}$ to be a monotone chain, only $\{q^n\}$; with careful choice of Ξ we have been able to guarantee monotonicity of $\{q^n\}$ in all cases we examined.

Theorem 5.5 Φ has a maximal fixed point $q^* = \Phi(q^*)$. Furthermore, $\{\Phi^n(\mathbf{1})\} \rightarrow q^*$.

Proof Under Conjecture 5.4, Φ is a monotone nonincreasing mapping in the pointwise ordering ($q^n \succeq q^{n+1}$). Φ is order-continuous because the sequence $\{q^n(b)\}$ is monotone for

each b and confined to a compact set $[0, 1]$. By the Tarski-Kantorovitch theorem, the set of fixed points is nonempty and the chain $\{\Phi^n(\mathbf{1})\} \rightarrow q^*$, the maximal element of the set of fixed points.

Uniqueness is not generally assured, since $q = 0$ is a fixed point; uniqueness therefore only obtains when there does not exist any fixed point with $q \geq 0$. A sufficient condition for $q^* \neq 0$ is that $\Delta > 0$; in that case, no default will occur at $b > -\Delta$ so $q = 0$ is never the maximal fixed point. A separate sufficient condition is $\lambda > 0$, since again there will exist a small enough debt level that will never be defaulted on, although it is not possible to characterize analytically where this debt level lies. Necessary conditions for $q^* \neq 0$ are unknown.

Any equilibrium q must be a fixed point of Φ . Since our program converges to the maximal fixed point, it converges to the competitive equilibrium with the lowest interest rate functions. This equilibrium has the property that, with an exogenous risk-free rate r , budget sets are largest under full information (and therefore consumer welfare is maximized). Under asymmetric information the first statement still holds (budget sets are larger when interest rates are lower), but utility may not be maximized for all individuals due to the potential for pooling; nevertheless, we think that this equilibrium is the natural one to study.