American Economic Journal: Macroeconomics 2012, 4(3): 153–183
http://dx.doi.org/10.1257/mac.4.3.153

A Quantitative Theory of Information and Unsecured Credit†

By Kartik Athreya, Xuan S. Tam, and Eric R. Young*

Important changes have occurred in unsecured credit markets over the past three decades. Most prominently, there have been large increases in aggregate consumer debt, the personal bankruptcy rate, the size of bankruptcies, the dispersion of interest rates paid by borrowers, and the relative discount received by those with good credit ratings. We find that improvements in information available to lenders on household-level costs of bankruptcy can account for a significant fraction of what has been observed. The ex ante welfare gains from better information are positive but small. (*JEL* D14, D82, G21)

For most of the postwar period, the unsecured market for credit has been small. Direct evidence from the Survey of Consumer Finances (SCF), as well as other sources (Ellis 1998), shows that unsecured credit did not appear in any significant amount in the United States until the late 1960s. However, over the past three decades there have been dramatic changes in this market. First, and perhaps the most well-known attribute of the unsecured credit market in the period we consider, has been the large increase in personal bankruptcy rates, from less than 0.1 percent of households filing annually in the 1970s to more than 1 percent annually since 2002. Sullivan, Warren, and Westbrook (2000) also notes that not only are bankruptcies more common now than before, they are also larger. As measured by ratio of median net worth-to-median US household income, the size of bankruptcies grew from 0.19 in 1981 to approximately 0.26 by 1997. More generally, the use of unsecured credit has intensified. It has nearly tripled, as measured by the ratio of aggregate negative net worth to aggregate income, from 0.30 percent in 1983, to 0.67 percent in 2001, to 0.80 percent in 2004 (as measured in the SCF).

Perhaps most dramatically, data from the SCF suggests that the distribution of interest rates for unsecured credit was highly concentrated in 1983 and very diffuse

*Athreya: Research Department, Federal Reserve Bank of Richmond, 701 East Byrd Street, Richmond, VA 23219 (e-mail: Kartik.Athreya@rich.frb.org); Tam: CFAP, Judge Business School, University of Cambridge, Trumpington Street, Cambridge CB1 2AG, UK (e-mail: xst20@cam.ac.uk); Young: Department of Economics, University of Virginia, Charlottesville, VA 22904 (e-mail: ey2d@virginia.edu). We thank Kosuke Aoki, V. V. Chari, Maxim Engers, Juan Carlos Hatchondo, Espen Henriksen, Mark Huggett, Narayana Kocherlakota, Dirk Krueger, Nathan Larson, Igor Livshits, Jim MacGee, Leo Martinez, Makoto Nakajima, Borghan Narajabad, Fabrizio Perri, Kevin Reffett, Víctor Ríos-Rull, Andy Skrzypacz, Nick Souleles, and numerous seminar participants. We are particularly grateful to two anonymous referees. We also thank Brian Gaines and Anne Davlin for excellent research assistance. Tam thanks the John Olin Fellowship and the Bankard Fund for Political Economy for financial support. Young thanks the Bankard Fund for Political Economy for financial support. The views expressed are those of the authors and are not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System. All errors are our responsibility.

† To comment on this article in the online discussion forum, or to view additional materials, visit the article page at http://dx.doi.org/10.1257/mac.4.3.153.

by 2004. Measured in terms of the variance of interest rates paid by those who report rolling over credit card debt, we find that as of 1983 the variance was 7.90 percentage points, but by 2004 this number had more than *tripled* to 26.63 percentage points. Interestingly, the change in the variance of rates has not been accompanied by large changes in the mean spread on unsecured credit relative to the risk-free rate (measured by the annualized 3-month T-bill rate). These spreads have fallen by only slightly more than 1 percentage point.

The change in dispersion of interest rates appears to be a consequence of more explicit pricing for borrower default risk. A variety of financial contracts, ranging from credit card lines, to auto loans, to insurance, began to exhibit terms that depended nontrivially on regularly updated measures of default risk, particularly a household's credit score and whether the household had a delinquent account.[1] Comparing data from 1983 and 2004, we find that the distribution of interest rates for delinquent households shifted significantly to the right of that of nondelinquent households, with the means of those with past delinquency being over 200 basis points greater than those with no such events on their records.[2]

The purpose of this paper is to measure the extent to which these changes may reflect improvements in the information directly available on borrowers' default risk. Unsecured credit markets are a particularly likely place for information to be relevant. Perfectly collateralized lending is, by definition, immune to changes in information. Private information is simply irrelevant. By contrast, in the case of unsecured debt, formal collateral is totally absent. As a result, changes in the asymmetry of information will likely have consequences for prices and allocations.

We emphasize that we do not think that information availability is the only thing likely changed over this period. Many other things germane to unsecured credit market statistics have changed over the past three decades (see Livshits, MacGee, and Tertilt 2010 for a detailed discussion of the various factors). Our aim is simply to better understand the role played by improved information by keeping the environment fixed in all other ways.

We develop a life-cycle model of consumption and savings in which borrowers vary in their willingness to repay debts as a result of shocks to both income and nonpecuniary costs of default. Importantly, lenders cannot always directly observe the state of a borrowing household. Nonetheless, agents' borrowing and repayment decisions may convey additional information about aspects of their condition that are otherwise private information. As a result, the gap between allocations under asymmetric information and their symmetric-information counterparts may be narrowed. One of the answers we provide here is a quantitative measure of this effect.

[1] At least two related findings stand out from the literature. First, the sensitivity of credit card loan rates to the conditional bankruptcy probability grew substantially after the mid-1990s (Edelberg 2006). Second, credit scores themselves became more informative. Furletti (2003), for example, finds that the spread between the rates paid by highest and lowest risk classifications grew from zero in 1992 to 800 basis points by 2002.

[2] All interest rate data is taken from the Survey of Consumer Finances. Specifically, the SCF question regarding late payments is "Now thinking of all the various types of debts, were all the payments made the way they were scheduled during the last year, or were payments on any of the loans sometimes made later or missed?" The variable for SCF1983 is "V930," where a value of "1" means "all paid as scheduled" and a value of "5" means "sometimes got behind or missed payments." For SCF2004, the variable is "X3004," where the values of "1" and "5" are the same as for 1983.

Our paper is most closely related to Sánchez (2010) and Livshits, MacGee, and Tertilt (2008). Sánchez (2010) posits that the increased size of the credit market can be attributed to declines in the cost of offering contracts that separate households by risk characteristics. Specifically, the era with contractual homogeneity and low levels of unsecured risky credit is characterized by prohibitively expensive costs of offering screening contracts. Relatedly, Livshits, MacGee, and Tertilt (2008) argue that in the past high "overhead" costs severely limited lenders from offering a wide menu of risky contracts. These papers therefore both suggest that screening could not, and did not, play an important role in allowing lenders to overcome the effect of asymmetric information on consumer default risk.[3]

The absence of screening as a route to escape the classical "lemons" problem does not, however, rule out the separation of borrowers. In particular, credit market *signaling* by borrowers, who are on the informed side of the market, still remains an option; borrowers will transmit information via the size of the loan they request. The equilibrium level of activity in the consumer credit market then depends on the ability (and desire) of relatively low-risk borrowers to use debt as a signal to separate themselves from those who pose higher default risk.

We proceed by studying two settings. First, we allow lenders to observe all relevant aspects of the state vector necessary to predict default risk. We intend this "full information" environment to represent the one currently prevalent (with outcomes compared against the 2004 SCF, among other sources). Second, we compare the preceding allocation to one where lenders are no longer able to observe all of these variables. We intend these "partial information" settings to be representative of periods prior to the mid-1980s (and here, compare outcomes against data from the earliest (1983) wave of the SCF, among other sources). The difference across these allocations is a quantitative measure of the effect of improved information about shocks in unsecured credit markets.

Our findings suggest that improvements in the ability of lenders to observe borrower characteristics can help account for a substantial proportion of most, though not all, of these observations. In terms of bankruptcy, our model suggests that relative to the current period, information by itself can account for approximately 46 percent of the total change in bankruptcy seen in the data. In terms of interest rate dispersion, our model suggests that the change in information captures 77 percent of the change in the variance of interest rates paid by households, and a similar amount (73 percent) of the change in what we term a "good borrower" discount. In terms of the change in overall indebtedness, the model predicts that information by itself would have led to an even greater change in borrowing than what was observed, at 148 percent.

Two somewhat broader messages of our analysis are as follows. First, we find that, in general, asymmetric information is damaging but not fatal to the functioning of the unsecured credit market. Second, a central aspect of our findings is that the

[3] Our work is also related to Drozd and Nosal (2007), which offers a theory of increased differentiation of borrowers based on declining contracting costs, and to Narajabad (2007) who uses improvements in the quality of symmetric information about borrowers to induce changes in the credit market. Both papers assume strong ex post commitment on the part of lenders, an assumption that is hard to square with the flexibility credit card issuers have to change the terms at will.

power of signaling is likely to be weaker when it is costs that are germane to default decisions that are *not* related to income, rather than the persistent component of income, that are unobservable. The model also suggests that information may not be important in the growth of the average size of bankruptcies, though it should be stated clearly that the model systematically fails to generate bankruptcies as large as the data. In what follows, we describe our model and parametrization scheme and then present results. The final sections provide more detail on related work and then present some conclusions.

I. The Model

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 life cycle 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 nonpecuniary 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. However, it also reflects a large number of other costs that are not explicitly pecuniary in nature (as in Athreya 2002), such as the added transactions costs associated with a bad credit history (additional difficulty in renting an apartment or obtaining a cell phone contract). We therefore model λ as a multiplicative factor that alters the value of consumption in the period in which a household files for bankruptcy. We will allow it to differ with human capital, to be persistent, and to potentially vary over time.

Household preferences are represented by the expected utility function

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

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 nonpecuniary 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 when it takes a low one.

The cost λ evolves stochastically and is persistant. As a result, there are at least two interactions between the household's expectations of future income and the cost of bankruptcy it faces. First, to the extent that consumption and expected future income are linked, the cost of bankruptcy varies with income. Second, and more fundamentally, because of the signaling content of debt those with low future income prospects *and* a low nonpecuniary cost of default will again find the cost of bankruptcy, income, and consumption to be related.

By comparing the utility from not filing for bankruptcy with that arising from filing, while holding consumption and family size fixed, our specification implies that the within-period purely nonpecuniary cost of bankruptcy is given by $c^{-\sigma}(\frac{1}{1-\sigma} - \frac{\lambda^{1-\sigma}}{1-\sigma})$. For $\sigma > 1$, (which is the case in our quantitative analysis), the cost of bankruptcy is decreasing in c.[4] Thus, comparing two households with different levels of planned consumption in the current period, the one with the higher level will face the smaller relative effective reduction in utility arising from bankruptcy. However, there are two main points to note here regarding how the costs of bankruptcy vary with consumption on the equilibrium path.

First, the decision to file for bankruptcy is made by comparing the value that can be attained by paying the costs of bankruptcy in the current period with the value attainable if a household chooses to repay fully. Critically, consumption in these two paths will in general *not* be the same, even in the current period. In particular, in equilibrium, households who pay the cost of bankruptcy will be those for whom the immediate relaxation in the budget constraint arising from the discharge of debts is most valuable. All else equal, these households will be ones in which current income is low and expected future income is low as well. In this sense, our specification does not make the net benefits of bankruptcy lower for the "poor" than for the "rich."

[4] We thank an anonymous referee for noting this point.

Second, to the extent that those households with relatively good future income prospects plan to consume more than others, starting in the current period onward they will actually face a higher penalty (in terms of the absolute level of expenditures that their penalty is equivalent to) than those planning to consume less. On the other hand, since the marginal value of consumption is lower for those planning to consume a relatively high amount, the bite of this penalty is partially offset.

Our model allows us to study improvements in information arising from two places: improvements in lenders' ability to forecast the future *income* of borrowers, and improvements in the ability of lenders to assess factors *orthogonal* to predicted 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 differences in default risk among groups whose observables, perhaps including income, have already been forecasted as accurately by lenders as by the borrower. As mentioned at the outset, our model will suggest that improvements in such "softer" forms of information may indeed have played a role in generating the changes seen in unsecured credit markets. Because we wish to avoid overstating the power of adverse selection to unravel credit markets, we allow lenders to be perfectly informed about slow-moving components of the individual state vector, such as age and education, as these can likely be relatively easily inferred even if not directly observed.[5]

Given that we will study settings with partial information, creditors will generally want to track the history of default. We allow for tracking 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 current regulations requiring that bankruptcy filings disappear from one's credit score after ten years, and agents are prohibited from declaring Chapter 7 bankruptcy more than once every seven years. For tractability, we model the removal of the bankruptcy "flag" probabilistically. We denote by $\xi \in (0,1)$ the likelihood of the bad credit market flag disappearing tomorrow; having $m = 0$ does not prohibit the household from borrowing. This approach means that some households in our model will be able to declare bankruptcy more than once every seven years. However, since households in the US economy also have the option to declare Chapter 13 (once every nine months) , and typically have few non-exempt assets (the primary difference between Chapters 7 and 13 is the dispensation of assets), this abstraction is reasonable and avoids the need for a cumbersome state variable tracking "the number of periods since a filing."

Under partial information, the price charged to a household for issuing debt will generally depend on m, so that households with recent defaults will receive different credit terms than households with "clean" credit. When information is symmetric, this flag is useless, though it will, in general, be negatively correlated with debt (those with a documented past history of bankruptcy $m = 1$ will borrow less on average). This point also illustrates that inferring the extent to which bankruptcy

[5] We ignore regulations that require certain characteristics not be reflected in credit terms, such as those proscribed by the Equal Credit Opportunity Act.

affects future credit access is not clear cut; it can depend critically on whether the bankruptcy reveals information relevant for future default risk.

Under asymmetric information, we make an anonymous markets assumption; no past information about an individual (other than their current credit market status m) can be used to price credit. This assumption rules out the creation of a credit score that encodes past default behavior through a history of observed debt levels. Since income shocks are persistent, past borrowing would convey useful information, although it is an open question how much. Given the difficulties encountered by other researchers in dealing with dynamic credit scoring (see Chatterjee, Corbae, and Ríos-Rull 2011), we think it useful to consider an environment for which we can compute equilibria.

A. *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)

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

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 life-cycle 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 feature is critically important as it generates a strong desire to borrow early in the life cycle, 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.

All households then face a purely independently and indentically distributed shock to their expenditures, denoted $\chi \in \mathcal{X}$, that captures the effect of sudden changes in obligations that the household may not actively "choose" but nonetheless acquires. Examples include facing lawsuits, large out-of-pocket health risks, and unexpected changes to the economies of scale or the legal assignment of debts within the household, such as those arising from divorce (see Chatterjee et al. 2007 or Livshits, MacGee, and Tertilt 2010). Lastly, households are required to pay a proportional tax on labor earnings in each period, τ, to fund pension payments to retirees.

After receiving income and expenditure 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 when it is most needed. While this may not be an ideal abstraction, we make this assumption both for tractability and because it keeps our model close to related benchmarks of Chatterjee et al. (2007), and Livshits, MacGee, and Tertilt (2007). More substantively, the level of commitment on the part of lenders to not readjust 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.

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 (including the expense stemming from the expenditure shock). 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, it is 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

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

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

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

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]

[6] Exclusion 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.

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

B. *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 $I = (y, e, \nu, \chi, \lambda, j, m)$, 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 $\hat{\pi}^{b'}$, the break-even pricing function must satisfy

$$q(b', I) = \begin{cases} \dfrac{1}{1 + r} & \text{if } b \geq 0 \\ \dfrac{\left(1 - \hat{\pi}^{b'}\right)\psi_{j,y}}{1 + r + \phi} & \text{if } b < 0 \end{cases}. \tag{5}$$

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. Assuming a finite state space for all shocks, and Markov chains for e and λ, this implies that

$$\hat{\pi}^{b'} = \sum_{e', \nu', \lambda', \chi'} \pi_\chi(\chi')\,\pi_e(e'|\,e)\,\pi_\nu(\nu')\,\pi_\lambda(\lambda'|\lambda)\,d(b', e', \nu', \chi', \lambda'). \tag{6}$$

Since $d(b', e', \nu', \chi', \lambda')$ is the probability that the agent will default in state $(e', \nu', \chi', \lambda')$ 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.

Partial Information.—Partial information in our model refers to cases in which lenders cannot directly observe at least one of the current realization of the persistent component of earnings, e, or the nonpecuniary cost of bankruptcy, λ. The objects e and λ, because they are the only persistent components, are the two stochastic elements necessary and sufficient to perfectly forecast default risk for any given loan request b'. That is, knowledge of e and λ would, given b', collapse the model to the "full-information" case.

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 e and λ, even when these objects are not *themselves* useful for forecasting default risk. These variables are the current realization of the transitory shock ν, current net worth b, and total income.[8] The only exception to this limit on observability

[8] In our setting, if total income, $W\omega_{j,y}e\nu$, *and* the transitory shock, ν, were observable, e would be observable (since W and $\omega_{j,y}$ are always assumed observable). An alternative would be to allow total income to be observed,

is for the expenditure shocks, χ, which, given the interpretations it has in the literature, we allow to always be directly observable.

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 (PBE), 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 idependently and identically 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 independently and identically distributed 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 e, λ, or possibly both, as these are the two household-level state variables useful for forecasting the bankruptcy decision, one period hence, of a household that has requested a loan of b'. We will describe the inference problem for the case in which *neither* e nor λ is directly observable. The modifications for the case where only one of these objects is unobservable are obvious. If lenders cannot directly observe these items, however, they must resort to constructing an estimate of the value of the pair (e, λ) received by a household requesting b'. To make the best possible inference of (e, 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 $\Gamma(b, y, e, \nu, \chi, \lambda, j, m)$, and let the household decision rule for optimal borrowing be given by $b' = g(b, y, e, \nu, \chi, \lambda, j, m)$. For any *given* e and λ, not all values of the state vector are consistent with the observables, particularly the loan request b'. Therefore, let $S_{\{e,\lambda\}}$ denote the set of values for remaining household characteristics, b and ν, that, for a household with given e and λ, *could be* consistent with the observable vector (b', y, χ, j, m). This set is therefore the (possibly set-valued) pre-image of $g(\cdot)$

but not either of its stochastic components (e, ν). This arrangement endows lenders with more information than we allow them. We will show, however, that even in this case, additional information on income is actually not very powerful in altering allocations, and so would be even less relevant under the alternative in which lenders were allowed perfect observability of total income. Lastly, note that we *always* assume that lenders can observe age and education, and hence that lenders always know the expectation of a borrower's income conditioned on both these objects.

for a given e, λ. Notice that the state variables b and ν are both unobservable and useless for *directly* predicting e' and λ', but nevertheless contain useful information on current e and λ via the household's decision rule.

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

$$\Pr(e, \lambda | b', y, j, m, \chi) = \int_{S_{\{e,\lambda\}}} d\Gamma(b, y, e, \nu, \chi, \lambda, j, m). \tag{7}$$

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

$$\hat{\pi}^{b'} = \sum_{e,\lambda} \left[\sum_{e',\nu',\chi',\lambda'} \pi_{\chi}(\chi') \pi_e(e'|e) \pi_{\nu}(\nu') \pi_{\lambda}(\lambda'|\lambda) d(b', e', \nu', \chi', \lambda') \right] \Pr(e, \lambda \,|\, b', y, j, m, \chi). \tag{8}$$

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 interstate 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, PBE.

Intuitively, a PBE is given by a set of beliefs and strategies for households and lenders, where the strategies for each are 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.

A PBE (see Mas-Colell, Whinston, and Green (1995, 285) definition 9.C.3, and the additional requirement given on page 452) is as follows. Denote the state space for households by $\Omega = \mathcal{B} \times \mathcal{Y} \times \mathcal{E} \times \mathcal{V} \times \mathcal{L} \times \mathcal{J} \times \{0,1\} \subset \mathcal{R}^5 \times \mathcal{Z}_{++} \times \{0,1\}$ and space of information as $\mathcal{I} \subset \mathcal{Y} \times \mathcal{E} \times \mathcal{V} \times \mathcal{L} \times \mathcal{J} \times \{0,1\}$. 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 $\mu^*(b')$ 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 $\Upsilon^*(\Omega \mid b')$.[9]

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

- **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.

- **Lenders optimize given their beliefs:** Given common beliefs $\Upsilon^*(\Omega \mid b')$, q'^* is the pure-strategy Nash equilibrium under one-shot simultaneous-offer loan-price competition.

- **Beliefs are consistent with Bayes' rule wherever possible:** $\Upsilon^*(\Omega \mid b')$, is derived from $\Psi^*(\Omega, b')$ and household decision rules using Bayes rule whenever b is such that $\mu^*(b') > 0$.

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. 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.

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

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 in the online Appendix, 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.[10] 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. The algorithm is described in the online Appendix.

Discussion of Equilibrium Selection.—We now discuss the nature of the equilibrium that we compute. In terms of notation, we will use the abbreviation "PI(x)" to denote the partial information settings in which one or more objects x are not directly observable, and "FI" to denote the setting with full information. Thus, the case where neither the current e nor λ are visible is denoted "PI(e, λ);" the case where only e is not observable is denoted by "PI(e);" and lastly, the case where only λ is not visible is denoted by "PI(λ)."

We define a **pooling equilibrium** to be one in which the choice of b is the same for two observably equivalent but different households. That is, for PI (λ) we say an equilibrium pools borrowers if either both the low and high-λ individuals choose the same b given the same remaining state variables; or given the same λ two different a individuals choose the same b, given that they face the same q function. For PI (e, λ) the definition extends in the natural way. Inspecting the law of motion for net worth shows that pooling is possible, since a horizontal line at some b values intersects more than one decision rule, but the existence of possible pooling does not guarantee it will emerge on the equilibrium path.

Pooling is most likely to occur in the model when the pricing function is flat to the right of the equilibrium decision, and the indifference curve in (b, q) space is steep to the right of the equilibrium decision. In the model, almost all borrowers are "constrained" in the sense that their MRS is not equal to the slope of the pricing function (which is zero at all points of continuity and undefined elsewhere). That is, borrowers locate at the edge of "cliffs" in the pricing function that occurs when

[10] It 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).

an additional state tomorrow is added to the default set.[11] If the flat segment is sufficiently long just to the right of the equilibrium decision and the indifference curve is steep, the gain from reducing borrowing (in terms of improved interest rates) is too small, as it requires a very large drop in current consumption to obtain a small increase in q. As a result, multiple "types" might pool at a given cliff point.

Examining the pricing functions in Figure 1, we see two things. First, the pricing function is relatively steep, so pooling is an unlikely outcome. The relative steepness follows from standard results that a single asset is sufficient to make the utility cost of fluctuations small, meaning that households do not mind saving a little more. Forward induction arguments and/or the test of equilibrium dominance (see Cho and Kreps 1987) typically destroy pooling equilibria on the grounds that they are sustained by unreasonable beliefs about the identity of types who would choose certain off-equilibrium signals; indeed, with two types their Intuitive Criterion guarantees that agents arrive at the best separating allocation (the Riley equilibrium); and with three or more types, the D2 criterion of Banks and Sobel (1987) delivers the same. In the figure, we see that the PI pricing function closely tracks the pricing function for the low-λ type at low debt levels but the high-λ type for high debt levels. That is, the lender correctly interprets a low debt issuance level as likely coming from a relatively safe borrower. Thus, while we cannot prove that our algorithm selects equilibria according to any forward induction argument, it does seem to capture the flavor of these refinements. However, pooling does occur on the equilibrium path, so our procedure does not deliver the best separating equilibrium as D2 would.

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

$$\tau W\left(\sum_{y,e,\nu,\chi,\lambda,j,m}\int_b \omega_{j,y} e\nu d\Gamma(b,y,e,\nu,\chi,\lambda,j < j^*,m)\right) \tag{9}$$

$$= W\sum_{y,e,\nu,\chi,\lambda,j,m}\int_b (\theta\omega_{j^*-1,y} e_{j^*-1}\nu_{j^*-1} + \Theta)\, d\Gamma(b,y,e,\nu,\chi,\lambda,j \geq j^*,m).$$

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 households 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.

[11] This property is not generic. Because indifference curves are U-shaped, eventually, it could be the case that households choose values for b on the interior of a flat segment of q. We do not find any such outcomes in our model, however.

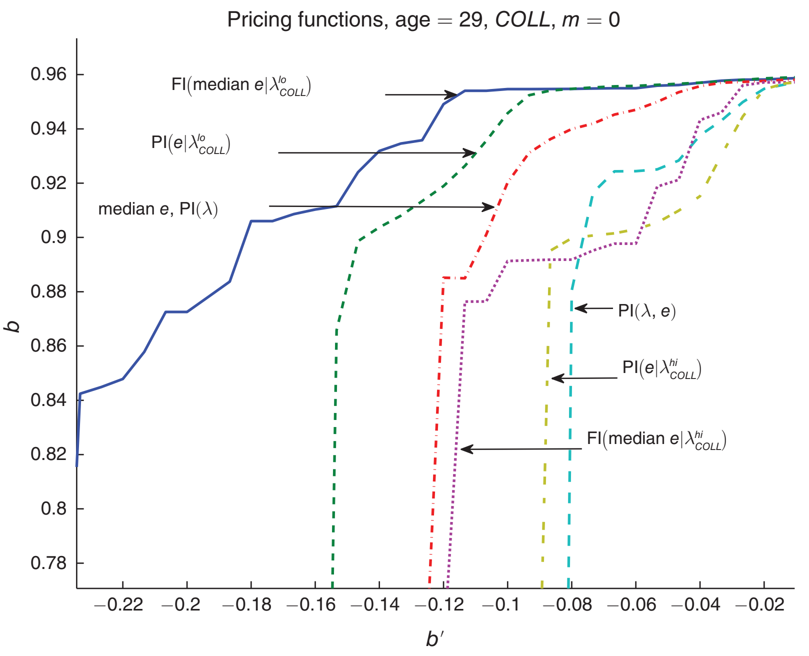

Figure 1. Information and Credit Supply

II. Calibration

Our calibration assigns numerical values to model parameters under the maintained assumption that the current setting is one of full information. This assumption requires explanation. It is clear that there is asymmetric information between unsecured borrowers and lenders in the data due to the regulatory constraints we mentioned earlier. Our approach is driven primarily by two considerations. First, it is computationally intractable to calibrate our asymmetric information model.[12] Second, it turns out that the FI outcomes are actually close to PI outcomes when the only information that is not directly observable is the current persistent component of income: FI and PI(e) outcomes are close. In this sense, one could consider PI(λ) as a benchmark. Anticipating the results, this finding of our model suggests that it is improvements in the ability of lenders to forecast or assess household characteristics relevant to bankruptcy decisions that are *not* summarized by income alone which have been important in driving the changes seen in unsecured credit markets.

Our targets are the ratio of median debt discharged in bankruptcy to median US household income, the total fraction of US households with negative net worth, the conditional mean of debt for those who hold debt of each of the three educational groups, and the personal bankruptcy rate for each of the

[12] The PI cases simply take too long for us to explore much of the parameter space (one equilibrium takes 10 days even when efficiently parallelized on a 32-processor cluster). Moreover, modifications intended to make the convergence more rapid have generally led to equilibria in which credit is nearly shut down. Given that the beliefs for lenders required to sustain those equilibria are somewhat unreasonable, we are stuck with slow methods.

three educational groups. We therefore have a total of eight targets, so we calibrate eight model parameters. These are: the discount factor β, two values for the nonpecuniary cost of bankruptcy for each educational level (6 parameters—$\lambda^{lo}_{NHS}, \lambda^{hi}_{NHS}, \lambda^{lo}_{HS}, \lambda^{hi}_{HS}, \lambda^{lo}_{COLL}, \lambda^{hi}_{NHS}$), and the (common) persistence parameter for this cost, ρ_λ.

The statistics on default rates are based on the measures from Sullivan, Warren, and Westbrook (2000). The debt targets are all from the SCF (2004), and use exactly the same definitions employed by Chatterjee et al. (2007) and Sánchez (2010), presented in Table 1. The maintained assumption within the model is that households have a single asset with which to smooth consumption. Given our life-cycle setting, this assumption is not likely to be crucial for the question we pose. As noted above, in the data (as well as in the model), defaults are largely the province of the young (Sullivan, Warren, and Westbrook 2000). Young households also have few gross assets, implying that negative net worth and unsecured debt largely coincide (see also table 2.4 in Sullivan, Warren, and Westbrook 2000). As in Chatterjee et al. (2007) and Sánchez (2010), we identify debt with negative net worth, and the specific value we target for aggregate-debt-to-income in the FI setting is 0.67 percent, the same as theirs. For the measure of borrowers, we choose a target of 12.5 percent, as it lies in the middle of the interval defined by the estimate of 6.7 percent in Chatterjee et al. (2007) and 17.6 percent in Wolff (2006).

To parameterize "expense" shocks, χ, we follow Livshits, MacGee, and Tertilt (2007), and specify an idependently and identically distributed random variable that is allowed to take three values $\{0, \chi_l, \chi_h\}$. The shock χ_l is set to be 7 percent of mean income, and χ_h is set at 27 percent of mean income. The relative likelihoods are $\Pr(\chi = 0) = 0.9244$, $\Pr(\chi = \chi_h) = 0.0710$, and $\Pr(\chi = \chi_l) = 0.0046$. As a result, only a minority of households receive shocks, and of these, most do not receive a catastrophic one. Our calibration target for the model's bankruptcy rate is 1.2 percent, in line with the average over the period 2000–2006. The median discharge to median US household income ratio that we use takes the ratio of median debt discharged in Chapter 7 bankruptcy (taken from Sullivan, Warren, and Westbrook 2000, table 2.4) and compares this relative to median US household income from the US Census Bureau, and is therefore set to 0.27.[13]

For earnings risk, we do not calibrate any parameters. We instead use values consistent with those previously obtained in the literature. We set $\theta = 0.35$ at an exogenous retirement (model) age of 45 and $\Theta = 0.2$, yielding an overall replacement rate for retirement earnings of approximately 55 percent. The income process is of the "restricted income profiles (RIP)" type and is taken from Hubbard, Skinner, and Zeldes (1994), which estimates separate processes for non-high school, high school, and college-educated workers for the period 1982–1986.[14]

[13] Alternatively, since we identify defaultable debt with negative net worth in the model, we could use net worth among filers as our measure of debt discharged. Sullivan, Warren, and Westbrook (2000, 72) report median net worth among filers in 1997 data as $-\$10,542$. Comparing this number with US median household income yields a very similar target. The near-equivalence of gross unsecured debt and negative net worth reflects the fact that nearly all uncollateralized debts are dischargeable for nearly all filers, because few have significant positive asset positions (including home equity).

[14] Figure 2 in the online Appendix displays the relative incomes over the life cycle for all three human capital groups.

Table 1—Calibration Targets and Model Performance

Target	Model	Data	Source
Med(Discharge)/Med(US HH Income)	0.1329	0.2688	Sullivan et al. (2000)
Fraction(Net Worth $<$ 0)	0.1720	0.1250	Chatterjee et al. (2007) & Wolff (2006)
Agg. NW(NW <0)/Agg. Income | NHS	0.1432	0.0800	the 2004 SCF
Agg. NW(NW <0)/Agg. Income | HS	0.1229	0.1100	the 2004 SCF
Agg. NW(NW <0)/Agg. Income | COLL	0.0966	0.1500	the 2004 SCF
Default rate | NHS	1.237%	1.228%	Sullivan et al. (2000)
Default rate | HS	1.301%	1.314%	Sullivan et al. (2000)
Default rate | Coll	0.769%	0.819%	Sullivan et al. (2000)

In Athreya, Tam, and Young (2009), we study the effect of the rise in the volatility of labor income in the US and find it to be quantitatively unimportant. Therefore, we use the process estimated on the earlier data even though we compute the FI case assuming it applies to 2004. The process for $\omega_{j,y}$ displays a more pronounced hump for college types than the others. Details are available in the online Appendix. The shocks are discretized with 15 points for e and 3 points for ν. The resulting processes have a common $\rho = 0.95$, with $= \sigma_{\epsilon}^2$ (0.033, 0.025, 0.016) and $\sigma_{\nu}^2 = (0.04, 0.021, 0.014)$ for non-high-school, high school, and college agents. The measures of the three groups are 16, 59, and 25 percent, respectively. The risk-free rate on savings is set to 1 percent to reflect the assumption that households have access to a risk-free and liquid savings instrument. For lending costs, we set $\phi = 0.03$ to generate a 3 percent spread between rates of return on broader measures of capital and the risk-free borrowing rate, consistent with transactions costs measured by Evans and Schmalensee (1999). Δ is set equal to 0.03. If one unit of model output is interpreted as \$40,000—roughly median income in the US—then the filing cost is equal to \$1,200, and is an estimate inclusive of filing fees, lawyer costs, and the value of time. Finally, $\xi = 0.2589$ implies that 95 percent of households that do not file for bankruptcy again will have clean credit after 10 years.

The calibration generates a stochastic process for the nonpecuniary cost λ with the following properties. We specify that it follows a Markov chain, and the calibration assigns this process very high persistence. As a result, it appears that such costs are likely partly in the nature of a household-level "type," as they will remain generally unchanged for any given household over time. However, the variance of the cost is nontrivial, and, as a result, households will differ substantially from each other in their willingness to repay debt, all else fixed. Lastly, we set risk aversion/inverse elasticity of intertemporal substitution at $\sigma = 2$ for all households. The calibrated parameter values are collected in Table 2.

III. Results

We evaluate the implications of changes in information on six aggregates of interest. These are the overall bankruptcy rate, measured by dividing the American Bankruptcy Institute's aggregate Chapter 7 filing rate by the number of US households in 2004 (FI case) and in 1983 (PI cases); the variance of interest rates on unsecured debt in the 2004 SCF (FI case) and 1983 SCF (PI cases); the difference in average credit-card-to-risk-free rate spreads paid by those with a past bankruptcy

TABLE 2—CALIBRATED PARAMETERS

Parameter	Interpretation	Value
β	discount rate	0.9628
ρ_λ	persist. of non-pec. cost	0.9636
λ_{NHS}^{lo}	low val. of λ for Non-High-School	0.7675
λ_{NHS}^{hi}	high val. of λ for Non-High-School	0.9088
λ_{HS}^{lo}	low val. of λ for High-School	0.7310
λ_{HS}^{hi}	high val. of λ for High-School	0.9320
λ_{NHCOLL}^{lo}	low val. of λ for College	0.7831
λ_{COLL}^{hi}	high val. of λ for College	0.9071

on their record relative to those without one, measured as the difference between the average real credit card interest rate on which balances were paid and the ex-post real (using PCE deflator) annualized 3-month T-bill rate; the mean interest rate spread paid by borrowers under PI and FI; the average debt in the economy, as measured by the ratio of total negative net worth to total economywide income (from the SCF 1983 and 2004); and the median unsecured debt (as measured by Sullivan, Warren, and Westbrook 2000, table 2.4) at the time of bankruptcy relative to US median household income from US Census data from 2004 (under FI) and 1983 (under PI).[15]

We now address our main question: how do the six aggregates change when information is systematically withdrawn from lenders? These changes measure the role better information plays in determining credit market outcomes.

A. *Full Information*

In the full information case, lenders form forecasts of default risk that coincide with the borrowing household's own conditional expectation. For brevity, we omit any detailed discussion of the Full Information setting here (the reader is directed to Athreya, Tam, and Young 2009 for a complete presentation). As a brief summary, Table 1 presents the aggregate unsecured credit market statistics from the model under full information.[16] The calibration procedure matches closely the targets we set in terms of the fraction of those with negative net worth, the overall level of unsecured debt relative to income for each education group, and the bankruptcy rate for each education group. However, the benchmark model is not able to fully capture the ratio of median discharge to median income. Part of the difference between the model and the data is that the latter employ a substantially broader

[15] We use the SCF sample weights to obtain a representative measure. An alternative debt-to-income target is the debt-to-income ratio for households where a member files. However, while debt at the time of filing is well measured, household income is not. Instead, for a large fraction of filings, only the filer's income, and not the *household's* income, is recorded. We therefore normalize by US household income instead. While we do not report it here, the results are very similar when we target these other ratios. In particular, as we will show, the model does well along a number of dimensions, but fails to capture the size of bankruptcies seen in the data. To measure the variance of interest rates, we use only those accounts on which interest is paid, and we do not weight observations by debt levels.

[16] Figure 5 in the online Appendix contains more detail on loan pricing for this case.

measure by including the large debts associated with small business failures (see Sullivan, Warren, and Westbrook 2000), a feature that is absent from the model. Unfortunately, the data that clearly separate discharge according to the nature of the filing (purely personal or small business) do not exist.

B. *The Effect of Asymmetric Information*

The central goal of our paper is to assess the extent to which improvements in information held by creditors can help account for the large changes seen in unsecured credit aggregates over the past three decades. Our results suggest that along several, but not all, dimensions, the answer is "a substantial amount." We focus our attention on the behavior of the six unsecured credit market aggregates laid out earlier, and measure the extent to which moving from the FI case to a partial information setting generates changes in these aggregates. We take the benchmark partial information environment to be PI(e, λ).

Information and Aggregates.—The first two columns of Table 3 compare the levels of each of the six items of interest in the data as of 1983 against the predictions of the PI(e, λ) model, while the third and fourth columns represent recent (2004) data and the predictions of the model under FI. Given that we have calibrated the model to match the 2004 data as closely as possible, what is relevant is the extent to which the change in information accounts for differences seen between 1983 and 2004. For ease of exposition, the bottom half of the table displays changes in the levels of the variables of interest, and the fraction of these changes generated by our model when we move from an FI setting to the PI(e, λ) environment.

In terms of bankruptcy, our model suggests that relative to the current period, information by itself can account for approximately 46 percent of the total change in bankruptcy seen in the data. In terms of interest rate dispersion, our model suggests that the change in information captures 77 percent of the change in the variance of interest rates paid by households and a similar amount (73 percent) of the change in the "good borrower" discount. The model also captures the observed change in the spread of mean unsecured interest rates relative to the risk-free rate reasonably well (97 percent). In terms of the change in overall indebtedness, the model predicts that information by itself would have led to an even greater change in borrowing than what was observed, at 148 percent.

Finally, we turn to the size of bankruptcies. As stated earlier, we measure the model's predictions for the ratio of median debt discharged in bankruptcy to median US household income. Given this normalization, it is clear that bankruptcies have grown larger over time. As seen in Table 1, the benchmark calibration is unable to match the size of the median bankruptcy, and instead generates bankruptcies of approximately half the size. Moreover, the model predicts that improved information by itself should have left the size of equilibrium debt discharge essentially time invariant. One mechanism potentially at play in limiting the size of debts discharged, even when aggregate borrowing increases, is seen in Table 2 in the online Appendix. Nearly 40 percent of bankruptcy filers are those with no expense shock ("Low χ"), and a high risk of default ("High λ"). Under FI, a switch in λ from low to high

Table 3—Model Performance for Aggregates

	1983		2004	
Levels	Data	Model	Data	Model
BK rate	0.20%	0.72%	1.21%	1.18%
Rate variation	7.90%	15.53%	26.63%	29.01%
Good borrower discount	−0.22%	1.21%	2.23%	3.01%
Mean rate spread	10.08%	8.21%	11.07%	9.19%
Agg. NW(NW $<$ 0)	0.40%	0.42%	0.80%	1.04%
Med(Discharge)/Med(US HH income)	0.192	0.134	0.269	0.133
Changes(1983–2004)	Data	Model	Explain change%	
BK rate	−1.01%	−0.46%	45.54	
Rate variation	−18.73%	−14.48%	77.31	
Good borrower discount	−2.45%	−1.80%	73.47	
Mean rate spread	−0.99%	−0.96%	96.97	
Agg. NW(NW $<$ 0)	−0.40%	−0.62%	147.62	
Med(Discharge)/Med(US HH income)	−0.077	0.001	−1.35	

will lead to an increase in the cost of borrowing, and thereby limit the size of debt that an agent can accumulate. In a PI setting, however, the borrower can continue to accumulate debts for some time before defaulting. The preceding statements are about the *sensitivity* of borrowing terms to personal circumstances. But the average cost of borrowing also matters, and works in an offsetting direction with respect to the size of bankruptcies. The FI setting is one in which, in general, credit is cheaper and more households can borrow at any given rate, relative to a given PI setting. As a result, more households are prone to having bankruptcies generated via expense shocks, simply because more households can borrow (especially when comparing PI(λ), where only λ is unobservable, to FI).[17]

One part of the model's inability to generate large enough bankruptcies likely stems from our use of relatively flexible one-period debt. As mentioned earlier, not only can such a setting hinder smoothing to a greater extent than one in which contracts take the form of credit lines that lenders are implicitly committed to honoring (perhaps due to reputational concerns), it may also—for the very same reason—be part of why bankruptcies are small. The difference can certainly be regarded as a shortcoming of the benchmark model with respect to its ability to generate bankruptcies as large as in the data.

There are two other asymmetric information economies one could consider within the framework of our model. These are the cases in which only persistent income shocks are unobservable, PI(e), and the one in which only nonpecuniary costs are unobservable, PI(λ). Table 4 displays the outcomes under these two information regimes. Two results are worth noting. First, the ability to observe income is actually not vital to allocations. Comparing the first two columns of Table 4 shows that allocations are largely similar across these two information regimes. One reason for this can be seen in Figure 1. Denote by PI($e \,|\, \lambda^{lo}_{COLL}$) the pricing in the PI($e$) case for a college-educated household with nonpecuniary cost $\lambda = \lambda^{lo}_{COLL}$. Define

[17] We thank an anonymous referee for highlighting this mechanism.

Table 4—Unsecured Credit Market Aggregates

	FI	PI(e)	PI(λ)	PI(λ, e)
Aggregate debt	0.0104	0.0094	0.0053	0.0042
Med(discharge)/med(US HH income)	0.1329	0.1351	0.1371	0.1342
Fraction of borrowers	0.1720	0.1709	0.1711	0.1493
Debt/income ratio | NHS	0.1432	0.1339	0.1206	0.1203
Debt/income ratio | HS	0.1229	0.1182	0.0964	0.9063
Debt/income ratio | COLL	0.0966	0.0944	0.0863	0.0768
Default rate | NHS	1.237%	1.018%	0.809%	0.778%
Default rate | HS	1.301%	1.197%	0.819%	0.789%
Default rate | COLL	0.769%	0.728%	0.638%	0.463%

PI($e \mid \lambda^{hi}_{COLL}$) analogously. We see that under either value for the observable λ, the pricing function confronting a household with the modal persistent income shock (which is the median e) in a FI setting is quite similar to the *single* price function that households face under PI. In other words, for any given observable value of λ, the pricing most households receive will not differ substantially between FI and PI(e).

The second finding is that the ability to assess a household's cost of bankruptcy *is* important. Table 4 shows that when λ is the only unobservable, allocations change quite substantially relative to FI. While the size of bankruptcies and the measure of borrowers remains stable, we see that debt/income ratios and bankruptcy rates fall very close to the benchmark private information setting. As a result, the model suggests that not all asymmetric information is equally important, and in particular that asymmetry of information on income by itself does not lead to large reductions in credit supply. Our findings also are consistent with the fact that credit card lenders rarely, if ever, follow up on cardholder income over time, and typically ask for only self-reported income at the time of application.[18]

How Does Our Approach Matter for the Outcomes Selected?—Both our initial choice of pricing function and our procedure for "filling in the holes" at intermediate iterates of the pricing function are "optimistic" in the sense that lenders believe that borrowers are no riskier than anyone with the next *lower* observed debt level. As a result, within the interval defined by any two *observed* debt levels, borrowing costs will not rise with debt, even though higher debts will generally be associated with greater default risk.[19] Thus, if there are "holes" in the equilibrium set of borrowing levels, it will be despite the fact that we allowed households to select borrowing levels that were "subsidized" from their perspective at iterations leading up to the converged pricing function (though of course, not in equilibrium).

Our algorithm also gives relatively low-risk (e.g., high current income) agents relatively higher incentives to borrow, since our initial pricing function will affect the behavior of the low-risk agents more than high-risk households. This is because high-risk agents are willing to pay a premium to borrow that the low-risk agents

[18] As reported on a credit card industry website, current practice is to use credit history and self-reported income at the time of application to estimate income, *without* verification of income (http://www.creditcards.com/credit-card-news/credit-card-application-income-check-1282.php).

[19] Of course, in equilibrium, lenders will earn zero profits, so they are indifferent across different specifications of these beliefs. Our approach just aims to sustain as much lending as possible.

would not. As a result, our updating is set to work against going to a much lower pricing function q and remaining there (that is, protected against much higher interest rates for any given debt issuance).

A second point to cover is the role played by the off-equilibrium-path beliefs that are induced by our algorithm and reflected in the converged price function $q^*(b)$ facing any particular type of agent. First, one can see from Figure 1 that the pricing functions are monotone in the face value of debt. Therefore, one immediate restriction on off-equilibrium beliefs is that households never believe that lenders would view them as less risky if they increased the size of the loan they request. This restriction plays a role in the convergence of the algorithm, and is shown by applying Theorem 6 from Chatterjee et al. (2007).

Next, consider the bottom-right panel in Figure 3, which is the PI(λ, e) case, for a young (age 29) college-educated household. This type is representative of a group that has the desire, all else equal, to borrow. For this group, we see that the distribution of debts for this class of households contains a region where no one borrows (i.e., debt levels in excess of approximately -0.08 units). Looking at the pricing for these regions indicates that the pricing never improves with debt, and also inherits the property from the updating procedure that it does not get worse until one hits the next higher level of debt that attracts a positive measure of borrowers.

As to the question of how much pooling the model generates, we proceed in two steps. First, as a general matter, pooling will occur under asymmetric information when the relatively low-risk agent in a given pool finds it not very valuable to lower their requested debt amount in return for a lower interest rate (higher q). As a way of assessing the overall extent of pooling, we study a representative set of decision rules across the unobserved components of the household's state (b, λ, e). In Figure 2, we plot the law of motion b (which, in a slight abuse of notation, nests the decision rule for today's debt issuance *and* the effect of declaring bankruptcy today, which forces $b' = 0$). Note the close proximity of the decision rules to each other across values of λ for many debt and savings levels. However, notice also that these households clearly differ in their default incentives. In particular, for the household with the median income shock, the level of debt that triggers default (seen in the jump to $b' = 0$) is substantially lower for households with $\lambda = \lambda^{hi}_{COLL}$ than for those with $\lambda = \lambda^{lo}_{COLL}$. The pattern in Figure 2 is representative of what is seen across education levels for most households in the early part of the life cycle (e.g., below age 40). Given the distribution of initial wealth b (more on this below), we conclude that a nontrivial proportion of households with differing default risk sends essentially the same signals in equilibrium. Of course, when income is low ("lo e") we see that decisions under either value of λ are similar, as households are driven to default more directly by low income than by the nonpecuniary cost that bankruptcy carries for them. By definition, however, such realizations of income are rare, and hence do not carry as much weight in determining aggregate outcomes. The proximity of decision rules across households with differing λ at relatively low initial debt levels, and their divergence at higher debt levels, is precisely why the PI(λ) and PI(λ, e) environments look substantially different from either the PI(e) or FI cases.

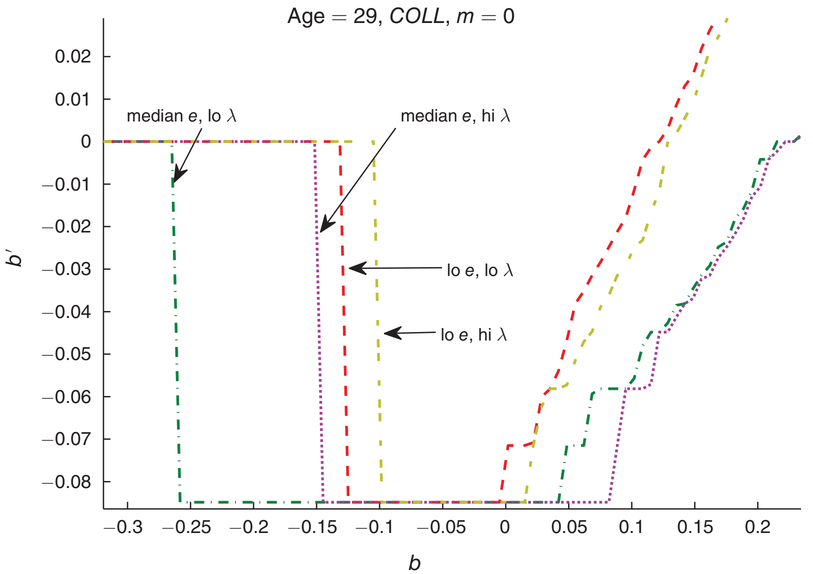

Figure 2. Evidence of Pooling in Decision Rules

The fact that decision rules overlap in many places, however, may still be misleading because it does not display the fraction of households located at points of the state space that lead them to send similar signals. To address this issue, we now return to Figure 3. Throughout this figure we hold λ fixed at λ_{COLL}^{hi}. Beginning with the top-right panel, we see that there are indeed several combinations of initial wealth b and persistent shock e that lead households to send identical signals. The horizontal dotted line in the top-right panel shows one particular signal b' that is sent by multiple types of households that are observably indistinguishable, and the top-left panel takes this level of debt (on the y-axis) and then displays on the x-axis the equilibrium price q for this particular debt level.

The next question is how many households exist at this debt level. To answer this question we move to the bottom-right panel, which displays the distribution of households over initial debt (net worth) levels b, and see that households closer to the median e (here, the line labeled e_6) are relatively more populous (unsurprisingly) amongst those who borrow b' than are the other classes of agents who issue the same level of b'. The latter are, naturally, those households whose current persistent shock realization is not as good (e.g., e_1), but whose initial wealth b was relatively high. These households wish to smooth consumption, and so will choose to issue the same debt as their luckier (in terms of e) counterparts.

Lastly, the bottom-left panel displays the fraction of households paying this particular bond price for their issuances. Note that this price reflects the information we just discussed: lenders understand that a given signal b' is not equally likely to have come from the various combinations of b, e (and λ in the PI(λ) and PI(λ, e) cases). Instead, they will use the distribution plotted in the bottom-right panel to assess the odds.

Two final points are worth emphasizing. First, there is only one systematic way to assess the role of the algorithm in equilibrium selection and allocations, which is

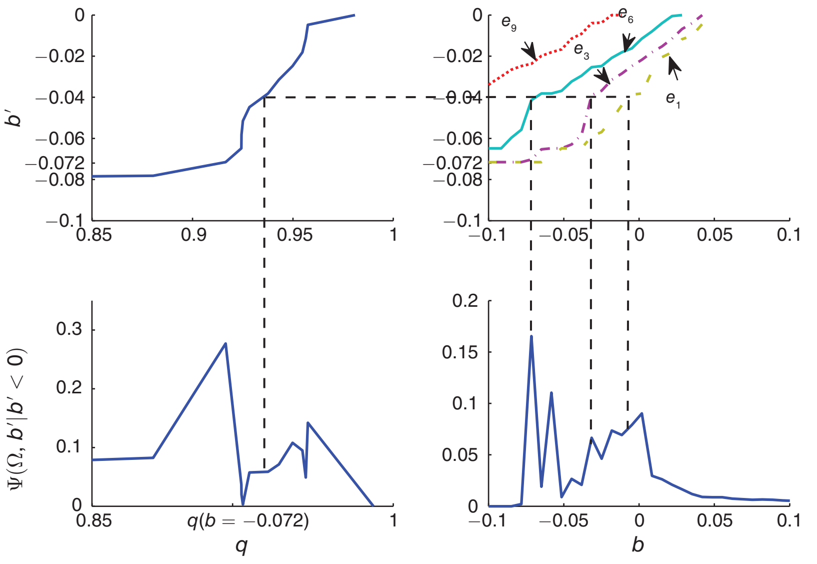

Figure 3. Measuring Pooling

to check its performance against a PI setting in which one already has a direct method for solving for an equilibrium, and then check if iterations lead to the same outcomes. This test is intractable as a general matter of course, and is why we employ the numerical approach we do in the first place. Second, there is still one case where we *can* check the algorithm when all unobserved shocks are independently and identically distributed. In this case, we know already what the PI outcome is. It is the FI outcome because the private information concerns variables that the lender has no interest in calculating. Therefore, one test that our algorithm should pass is that allocations and prices converge to the FI outcomes. We find that our algorithm passes this test.

Information and Credit Supply.—We turn now to evaluating the role that asymmetric information plays on the supply side of the unsecured credit market. This role is measured by comparing the pricing of unsecured credit arising from a given asymmetric information setting with others, or the most natural counterpart from the full information setting. Figure 1 contains at least three messages. First, quite naturally, private information of every variety considered here constrains credit supply; ceteris paribus, loan prices are high (q is lower) for a given borrowing level under private information. For example, as we move from FI to PI(e), notice that q under either value for the nonpecuniary cost of bankruptcy shifts credit supply inward. Second, under all information regimes, unsecured debt in excess of about 5 percent of economy-wide average income begins to carry default risk. The loan pricing function begins at this point to fall, and implies an interest rate higher than the risk-free borrowing rate. Third, as emphasized already, the model suggests that private information about bankruptcy costs is more important than private information about income, for a relatively natural reason. In an environment where only the persistent shock was unobservable, agents have a way of separating themselves. An agent with a poor current realization will be willing to pay more for a large loan to smooth consumption than one who has received a good shock. As a result, those asking for large loans in this case will be "correctly" identified as relatively high-risk borrowers, who, because of currently low income, will be willing to pay a price that is roughly actuarially fair. As a result, aggregate allocations under PI(e) are relatively close to those under FI.[20]

In sharp contrast to the ability of the unsecured market to function relatively smoothly in the absence of direct information on the current component of persistent income risk, when nonpecuniary costs are unobservable, outcomes are more sharply restricted relative to FI. Looking again at Figure 1, and holding fixed the current persistent shock at its median value, we see that a household faces a much higher interest rate for any given debt level they might issue. If further compounded by unobservable income risk, credit supply shrinks substantially. In the latter case, this contraction is because an interaction between the desire to borrow when faced with a particularly bad persistent shock and a low nonpecuniary cost of default makes it very attractive to borrow, if possible. The strength of adverse selection incentives can be seen by looking at the FI case. Here, we can see that the difference in loan pricing for those with high and low λ is very substantial, which further explains why, when combined with asymmetric information on income e, credit is relatively restricted.

Information, Interest Rates, and Dispersion.—As seen in Table 5, our model makes predictions for the effect of improved information on the mean and variance of interest rates paid by borrowers who have an indicator of past bankruptcy ($m = 1$) on their record. As documented in Edelberg (2006) and Furletti (2003), past default appears substantially more correlated with credit terms now than in past decades. In particular, the positive correlation of interest rates with past defaults may be seen as a form of "punitive" sanctions imposed by creditors. However, under competitive lending and full information, such penalties will not be viable. Nonetheless, given the persistence of shocks, the income events that trigger default may well persist, and therefore justify risk premia on lending. Indeed, we will argue that this view is a plausible interpretation of the data.

Note first that even under FI, there will generally be a nonzero correlation between the terms of credit and m; even though the flag is irrelevant under FI, it turns out that many who file do so when they have been hit by a bad persistent shock. As a result the pricing they face under FI is worse than that faced by the average household. Table 5 shows how the flag is related to credit pricing under two PI cases relative to FI.[21] The most useful thing to note is that in either case, the entries for $\text{Var}(r \mid m = 1)$ and $\text{Var}(r \mid m = 0)$ show that the model does fairly well at replicating the *level* of the dispersion under FI, and also the fact that dispersion is higher for those with the flag $m = 1$. However, we also see that the reduction in information to either the PI(λ, e) case or the PI(λ) case leads to a substantially smaller reduction in the variance than

[20] We are indebted to an anonymous referee for clarifying this mechanism to us.
[21] See also Table 3 in the online Appendix.

Table 5—Dispersion and Credit Sensitivity: PI and FI

	1983			2004	
	Data	PI(λ, e)	PI(λ)	Data	Model
Levels					
$E(r - r^f)$	10.08	8.21	8.24	11.07	9.17
$E(r - r^f \mid m = 1)$	9.86	9.41	9.47	12.85	12.18
$E(r - r^f \mid m = 0)$	10.08	8.20	8.23	10.68	9.15
$\mathrm{Var}(r)$	7.90	15.53	15.51	26.63	29.01
$\mathrm{Var}(r \mid m = 1)$	8.68	17.07	17.41	33.88	34.55
$\mathrm{Var}(r \mid m = 0)$	7.53	15.51	15.49	25.60	29.18
Changes					
$E(r - r^f \mid 1983) - E(r - r^f \mid 2004)$				−0.99	−0.96
$E(r \mid m = 1) - E(r \mid m = 0)$	−0.22	1.21	1.24	2.23	3.03
$\mathrm{Var}(r \mid m = 1) - \mathrm{Var}(r \mid m = 0)$	1.15	1.56	1.92	7.28	5.37
$\mathrm{Var}(r \mid 1983) - \mathrm{Var}(r \mid 2004)$				18.73	14.48
$\mathrm{Var}(r \mid m = 1, 1983) - \mathrm{Var}(r \mid m = 1, 2004)$				25.20	17.48
$\mathrm{Var}(r \mid m = 0, 1983) - \mathrm{Var}(r \mid m = 0, 2004)$				18.07	13.67

occurred in the data. For instance, we see that the reduction in variance predicted by the model (taking those with $m = 1$), for example, is from 34.55 to 17.07, while in the data, the reduction was almost twice as large, going from 33.88 to 8.68. A related prediction of the model also suggests that information is not the whole story. The model fails to explain why the difference between the mean interest rates paid by those with and without a bad flag is close to zero (-22 basis points), instead predicting that the difference should be a positive spread of 121 basis points.[22] The model thus suggests that information can help explain the difference in dispersion seen amongst those with a bad flag and those without. This finding is of interest because interest rate dispersion was never targeted.

Information and Welfare.—Having found that changes in information may have been important in explaining changes in the behavior of the US unsecured credit market, we now address the normative question of whether these changes are likely to have improved household welfare. The reason that it is not obvious that ex ante welfare will rise with more precise information is that some households may receive a subsidy from others by successfully pooling with them under partial information. If this subsidy is large for households who draw an initial type that will be worse off than average, then the suppression of information helps the credit market provide a form of insurance.[23]

Given our model specification, there are five ways in which the information held by lenders can improve. Starting from the PI(λ, e) case, lenders can move to one of three information settings: PI(λ), PI(e), and FI. And starting from PI(λ) and PI(e), lenders' information can improve to FI. In addition, because the economy is one with a fixed risk-free rate (and a tax rate that only depends on the invariant age

[22] The online Appendix contains additional results on the relationship between mean interest rates paid by borrowers by education-type when $m = 1$ and $m = 0$.

[23] The preceding is not merely a theoretical curiosity, it is the reasoning behind recent changes in statutory restrictions on the information that lenders may explicitly use, such as the Equal Credit Opportunity Act. The literature refers to this outcome as the Hirshleifer effect.

distribution), the welfare gains are additive. There are no spillovers across agents as a result of the differential credit constraints they face across informational regimes. For example, the welfare gain from moving from the PI(λ, e) case to FI is the sum of three items: the welfare gain/loss from a move from PI(λ, e) to PI(λ) plus the gain/loss from a move from PI(λ) to PI(e), and lastly from PI(e) to FI. For brevity, therefore, we present only the minimal number of cases needed to examine all five welfare improvements. The results are given in Table 6 and report how much welfare would change if we moved according to each specific improvement in information. Given the additivity, one can directly add the row entries in any given column to see the gains from particular improvement in information.

In terms of average gains or losses from better information, we find that, ex ante, households are always better off under full information.[24] In fact, summing the entries in the 'NHS' column of Table 6, we find that these households gain the most from improvements in information, while the college educated gain the least.[25] Part of the reason that the welfare gains are not huge, and benefit the more skilled by less, is that the PI environments still allow for considerable amounts of credit, and the relative contraction in credit supply (see Figure 4, for a picture that displays pricing function for the NHS households, which can be compared with Figure 1) is smaller for the skilled than the unskilled. From Figure 1, we see that, holding λ fixed, a move from FI to PI(λ, e), college educated households can still borrow roughly 10 percent of income at nearly the risk-free rate, while NHS households, as seen in Figure 4, will pay a default premium for essentially any borrowing they might attempt. Table 6 also assesses the welfare role played by the nonobservability of nonpecuniary costs λ alone, by comparing welfare in the move from PI(λ, e) to PI(e) and FI (for the particular value of λ under consideration). The punchline to the preceding results is that the gains in ex ante household welfare from improvements in information are positive but small.

We also present some welfare calculations involving the elimination of bankruptcy. As noted in Athreya, Tam, and Young (2009), the asymmetric information problem disappears if borrowers can commit to repayment. After all, if the borrower pays back in all states tomorrow, the lender has no incentive to estimate the current state. With expenditure shocks, a no-bankruptcy regime may be infeasible, so we present two related but feasible cases. We permit bankruptcy only if the household receives an expenditure shock, or we permit bankruptcy only if consumption would otherwise be negative (or zero). It is obvious that solving the asymmetric information problem is far less important than fixing the commitment problem for borrowers, for the reasons discussed in our earlier paper.

IV. Concluding Remarks

In this paper, we have shown that improved information held by unsecured creditors on factors relevant for the prediction of default can account for a nontrivial

[24] These welfare gains are computed after the newborn household learns y but before e, ν, χ, λ are observed.

[25] These results run in the opposite way from the gains of eliminating default. Athreya, Tam, and Young (2009) show that banning default benefits the college types more than the NHS types.

Table 6—Information and Ex Ante Welfare

	COLL	HS	NHS
PI $(\lambda, e) \rightarrow$ PI (λ)	0.029%	0.052%	0.097%
PI $(\lambda) \rightarrow$ PI (e)	0.023%	0.030%	0.047%
PI $(e) \rightarrow$ FI	0.034%	0.035%	0.065%
FI $\rightarrow$ NBK (if $\chi = 0$)	1.403%	1.205%	1.103%
NBK (if $\chi = 0$) $\rightarrow$ NBK (if $c > 0$)	3.242%	3.432%	5.045%

Notes: $\chi = 0$ represents no adverse expenditure shock. NBK (if $\chi = 0$): Can only file for bankruptcy if $\chi > 0$. NBK (if $c > 0$): Can only file for bankruptcy if $c \leq 0$.

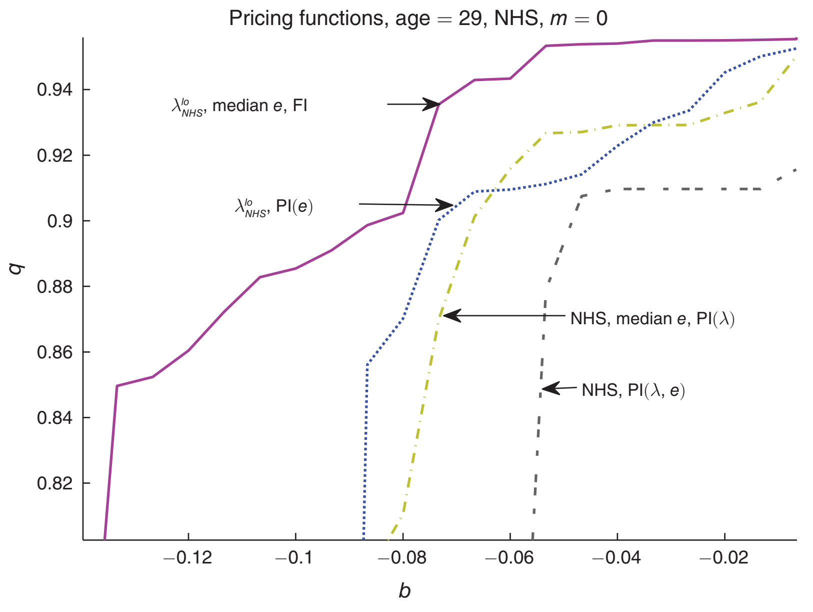

Figure 4. Information, Education, and Credit Supply

portion of the changes seen in unsecured credit markets. Given the prohibitive costs of overcoming private information via rich arrays of screening contracts, we show that the remaining route of borrower signaling could have been effective at separating borrowers according to risk characteristics. Quantitatively, a central aspect of our findings is that the power of signaling is likely to be weaker when it is nonpecuniary costs, rather than the persistent component of income, that are unobservable. A technical contribution of our work is an algorithm to compute equilibria with individualized pricing and asymmetric information; the algorithm is general and could easily be applied to alternative settings.

As mentioned in the paper, a feature of recent work on consumer default, including the present paper, is that it imposes a type of debt product that may not mimic all the features of a standard unsecured contract offered by real-world credit markets. Given that credit conditions are typically only adjusted by two events—default or entering the market to either purchase more credit or to retire existing lines—credit lines may be a more appropriate abstraction. This issue is tackled in symmetric information settings by Mateos-Planas (2007) and Tam (2009).

Lastly, the "RIP" income process we use is standard, but future work on bankruptcy will likely benefit from allowing for learning about income. Both consumption dynamics and usefulness of bankruptcy ultimately hinge on the persistence of risk faced by households. Using the "HIP" specification of Guvenen (2007), where estimated persistence is lower, is therefore a natural next step in assessing the role of consumer default and bankruptcy, and the role of information in altering unsecured credit use.

REFERENCES

Athreya, Kartik B. 2002. "Welfare Implications of the Bankruptcy Reform Act of 1999." *Journal of Monetary Economics* 49 (8): 1567–95.

Athreya, Kartik. 2004. "Shame as It Ever Was: Stigma and Personal Bankruptcy." *Federal Reserve Bank of Richmond Economic Quarterly* 90 (2): 1–19.

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.

Athreya, Kartik, Xuan S. Tam, and Eric R. Young. 2012. "A Quantitative Theory of Information and Unsecured Credit: Dataset." *American Economic Journal: Macroeconomics*. http://dx.doi.org/10.1257/mac.4.3.153.

Banks, Jeffrey S., and Joel Sobel. 1987. "Equilibrium Selection in Signaling Games." *Econometrica* 55 (3): 647–61.

Chatterjee, Satyajit, Dean Corbae, Makoto Nakajima, and José-Víctor Ríos-Rull. 2007. "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default." *Econometrica* 75 (6): 1525–89.

Chatterjee, Satyajit, Dean Corbae, and José-Víctor Ríos-Rull. 2008. "A Theory of Credit Scoring and Competitive Pricing of Default Risk." http://sites.google.com/site/deancorbae/research/cs040811paper.pdf?attredirects=0.

Cho, In-Koo, and David M. Kreps. 1987. "Signaling Games and Stable Equilibria." *Quarterly Journal of Economics* 102 (2): 179–221.

Drozd, Lukasz A., and Jaromir B. Nosal. 2007. "Competing for Customers: A Search Model of the Market for Unsecured Credit." http://finance.wharton.upenn.edu/~ldrozd/my_files/OLDCC.pdf.

Edelberg, Wendy. 2006. "Risk-Based Pricing of Interest Rates for Consumer Loans." *Journal of Monetary Economics* 53 (8): 2283–98.

Ellis, Diane. 1998. "The Effect of Consumer Interest Rate Deregulation on Credit Card Volumes, Charge-Offs, and the Personal Bankruptcy Rate." *FDIC: Bank Trends: Analysis of Emerging Risks in Banking* 98–05.

Evans, David S., and Richard Schmalensee. 1999. *Paying with Plastic: The Digital Revolution in Buying and Borrowing*. Cambridge, MA: MIT Press.

Fay, Scott A., Erik Hurst, and Michelle J. White. 1998. "The Bankruptcy Decision: Does Stigma Matter?" University of Michigan Working Paper 98–01.

Furletti, Mark. 2003. "Credit Card Pricing Developments and Their Disclosure." Federal Reserve Bank of Philadelphia, Payment Cards Center Discussion Paper 03–02.

Gross, David B., and Nicholas S. Souleles. 2002. "An Empirical Analysis of Personal Bankruptcy and Delinquency." *Review of Financial Studies* 15 (1): 319–47.

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

Hellwig, Martin. 1989. "Asymmetric Information, Financial Markets, and Financial Institutions: Where Are We Currently Going?" *European Economic Review* 33 (2/3): 277–85.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1994. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." *Carnegie-Rochester Conference Series on Public Policy* 40: 59–125.

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

Livshits, Igor, James MacGee, and Michèle Tertilt. 2008. "Costly Contracts and Consumer Credit." http://www.nber.org/papers/w17448.pdf.

Livshits, Igor, James MacGee, and Michèle Tertilt. 2010. "Accounting for the Rise in Consumer Bankruptcies." *American Economic Journal: Macroeconomics* 2 (2): 165–93.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory*. New York: Oxford University Press.

Mateos-Planas, **Xavier.** 2007. "A Model of Credit Limits and Bankruptcy with Applications to Welfare and Indebtedness." http://webspace.qmul.ac.uk/xmateos-planas/research/credlim.pdf.

Narajabad, **Borghan.** 2007. "Information Technology and the Rise of Household Bankruptcy." http://www.ruf.rice.edu/~bn2/InformationTechnologyAndBankruptcy.pdf.

Sánchez, **Juan M.** 2010. "The IT Revolution and the Unsecured Credit Market." Federal Reserve Bank of St. Louis Working Paper 2010–022.

Sullivan, **Teresa A.**, **Elizabeth Warren**, **and Jay Lawrence Westbrook.** 2000. *The Fragile Middle Class: Americans in Debt*. New Haven: Yale University Press.

Tam, **Xuan S.** 2009. "Long-Term Contracts in Unsecured Credit Markets." http://www.laef.uscb.edu/pages/conferences/cdab10/papers/tam.pdf.

Wolff, Edward N., ed. 2006. "Changes in Household Wealth in the 1980s and 1990s in the United States." In *International Perspectives on Household Wealth*, 107–50. Northampton, MA: Elgar Publishing Ltd.

This article has been cited by:

1. Lukasz A. Drozd1 and Ricardo Serrano-Padial2 1Research Department, Federal Reserve Bank of Philadelphia, Ten Independence Mall, Philadelphia, PA 19106 (e-mail: lukaszadrozd@gmail.com) 2School of Economics, Lebow College of Business, Drexel University, 3220 Market Street, Philadelphia, PA 19104 (e-mail: rspadial@gmail.com) . 2017. Modeling the Revolving Revolution: The Debt Collection Channel. *American Economic Review* **107**:3, 897-930. [Abstract] [View PDF article] [PDF with links]
2. Thomas Hintermaier, Winfried Koeniger. 2016. Debt Portfolios and Homestead Exemptions. *American Economic Journal: Macroeconomics* **8**:4, 103-141. [Abstract] [View PDF article] [PDF with links]

A Quantitative Theory of Information and Unsecured Credit: Additional Appendix

Kartik Athreya*
Federal Reserve Bank of Richmond

Xuan S. Tam†
University of Cambridge

Eric R. Young‡
University of Virginia

January 3, 2012

1 Appendix: Not For Publication

1.1 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

* Corresponding author. Research Department, Federal Reserve Bank of Richmond, kartik.athreya@rich.frb.org.

† Centre for Financial Analysis & Policy, Judge Business School, University of Cambridge, xst20@cam.ac.uk.

‡ Department of Economics, University of Virginia, ey2d@virginia.edu.

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, j, m, χ)

2. Guess the initial function $q^{(0)}(b', y, j, m, \chi)$ discussed above

3. Solve household problem to obtain $b' = g(b, y, e, \nu, \chi, \lambda, j, m)$, $S_{\{e,\lambda\}}$, and $d(b', e', \nu', \chi', \lambda')$, and $\Gamma(\cdot)$.

4. For all b' observed, use $S_{\{e,\lambda\}}$ and $\Gamma(\cdot)$ to compute $\Pr(e, \lambda|b', y, j, m, \chi)$ the probability that an individual is in (e, λ) given observed (b', y, j, m, χ). Knowledge of $\Pr(e, \lambda|b', y, j, m, \chi)$ and the distribution of households over the remaining observable state variables implies $\Upsilon(\Omega|b')$.

5. Compute

$$\widehat{\pi}^{b}(b', y, j, m, \chi) = \sum\nolimits_{e} \sum\nolimits_{\nu} \sum\nolimits_{\lambda} \sum\nolimits_{\chi'} \pi_{\chi}(\chi') \pi^{b'}(b', y, e, \nu, \chi, \lambda, j, m) \Pr(e, \lambda|b', y, j, m, \chi), \tag{1}$$

the expected probability of default for an individual in observed state (b', y, j, m, χ);

6. Fill in the "holes" in $\widehat{\pi}^{b'}$ for all b' not observed by applying the interim off-equilibrium beliefs as described in the main text.

7. Compute an "intermediate" price function $\widehat{q}$ for all b', that is actuarially fair (competitive) given the preceding estimate $\widehat{\pi}^{b'}$:

$$\widehat{q}(b', y, j, m, \chi) = \frac{\left(1 - \widehat{\pi}^{b'}(b', y, j, m, \chi)\right)\psi_j}{1 + r + \phi} \quad \text{for all } b \geq b_{\min}(b, y, j, m, \chi); \tag{2}$$

8. Set

$$q^{(1)}\left(b', y, j, m, \chi\right) = \Xi q^{(0)}\left(b', y, j, m, \chi\right) + (1 - \Xi)\, \widehat{q}\left(b', y, j, m, \chi\right)$$

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 $\mu^{(k)}$ be the associated distribution of debt requests induced by agent decisions, given $q^{(k)}$. For any given debt request b' such that $\mu^{(k)}\left(b'\right) = 0$, define $b' < b^{-} < 0$ as the nearest lower debt level relative to b' at which $\mu^{(k)}\left(b'\right) > 0$. b^{-} is therefore the upper end (or right-endpoint) of a segment of debt levels no household requests, given $q^{(k)}$. Define $\widehat{q}\left(b^{-}, I\right)$ to be the actuarially fair price at b^{-}; this is where we use optimal inference by lenders to construct $\widehat{\pi}^{b'}$.

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 $\mu^{(k)}(b^{+}) > 0$. b^{+} is therefore the lower end (or left-endpoint) of the same segment of debt levels, again given $q^{(k)}$. Denote by $\widehat{q}(b^{+}, I)$ the actuarially fair price at b^{+}. Thus, at any b, $\widehat{q}\left(b, I\right)$ is the actuarially fair price at any debt level that is observed under $\mu^{(k)}$, and equal to $\widehat{q}(b^{-}, I) \forall b \in (b^{+}, b^{-}]$. The collection of these segments is then denoted $\widehat{q}(\cdot, I)$.[1] With $\widehat{q}$ 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 $\widehat{q}(\cdot, I)$ that places weight Ξ on $q^{(k)}$. The preceding is Step 6.

We want to point our here that we do *not* confront the agent with the function $\widehat{q}$; we cannot guarantee that the sequence of $\widehat{q}$ functions is monotone. Instead, we update the

[1]Note that $\widehat{q}$ is a (lower semicontinuous) step-function.

pricing functions extremely slowly using a weight Ξ on the current iterate $q^{(k)}$ close to one (we use Ξ =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.[2] As a result, some pooling can occur in equilibrium as agents of different types are inframarginal at any discrete jumps in loan pricing. We discuss measures of pooling in the main text extensively.

Because the household value function is continuous but not differentiable or concave, we solve the household problem on a finite grid for b', using linear interpolation to evaluate the value function at points off the grid. Similarly, we use linear interpolation to evaluate q at points off the grid for b'. To compute the optimal savings behavior we use golden section search (see Press *et al.* 1993 for details of the golden section algorithm) after bracketing with a coarse grid search; we occasionally adjust the brackets of the golden section search to avoid the local maximum generated by the nonconcave region of the value function. To calibrate the model we use a derivative-free method to minimize the sum of squared deviations from the targets; the entire program is implemented using OpenMPI instructions in Fortran95 on a 32-processor Mac cluster. We make no claims about the program working for parameter values we have not explored. The computational cost of the PI equilibria is very large (roughly ten days of computing time per equilibrium), and our attempts to speed it up have led to equilibria with substantially less borrowing.

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

[2] We thank a referee for bringing this to our attention.

not familiar to all readers, we provide some basic definitions as well.

1.2 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 ($f \precsim g$ 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 $\Phi : Q \to Q$ 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 $\Phi : Q \to Q$ as the mapping defined in the algorithm; heuristically, Φ takes a given loan pricing function q^0, constructs the break-even pricing function $\widehat{q}$ 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:

$$\Phi\left(q^0\right) = \Xi\widehat{q}\left(q^0\right) + (1 - \Xi)\, q^0$$

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

Lemma 1.1 $(Q, \precsim)$ *is a complete lattice.*

Proof $(Q, \precsim)$ 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) $(Q, \precsim)$ is a complete lattice.

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

Theorem 1.2 *(Tarski-Kantorovitch) Let* $(Q, \precsim)$ *be a partially-ordered set and* $\Phi : Q \to Q$ *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 1.3 $\Pr(d = 1 | b)$ *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 1.4 $\Pr(d = 1 | q)$ *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 1.5 Φ *has a maximal fixed point* $q^* = \Phi(q^*)$*. Furthermore,* $\{\Phi^n(\mathbf{1})\} \to q^*$*.*

Proof Under Conjecture 1.4, 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 $\{\Phi^n(\mathbf{1})\} \to q^*$, 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 $\Lambda > 0$; in that case, no default will occur at $b > -\Lambda$ 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.

1.3 The Roles of λ and Δ Under Full Information

The non-pecuniary cost of bankruptcy, λ, plays an important role in allocations. It is meant to capture various aspects of deadweight costs borne by households in bankruptcy. There are two key aspects to this process – the coefficient of variation and the persistence. If one forces λ to remain constant and uniform across households when it is chosen to match the observed filing rate, the model produces a too-small discharge-income ratio and can no longer capture the heterogeneity in default costs implied by the estimates of Fay, Hurst, and White (1998). The calibrated value of λ in this case is also too small, in the sense that the model generates counterfactually-small bankruptcies, and as a result will understate the welfare costs of frequent default. If λ differs across households but is iid, discharge rates remain too low as the average λ needs to be small. Without persistence, no household's implicit collateral is expected to be particularly valuable in the next period and thus cannot support large debts. Thus, our calibration allows for the high cross-sectional dispersion and high persistence in λ needed in order to jointly support (i) large risky debts on which default premia are paid, (ii) frequent default, and (iii) relatively large discharges. If we change the average λ in the economy, the effect is to move default rates and discharge levels in opposite directions (see Athreya 2004). Furthermore, changing merely the average λ delivers little change in the dispersion in the terms. Thus, changes in stigma can be dismissed as the force driving all of the changes in the unsecured credit market.

Lastly, we discuss the roles played by the two main transactions costs, Δ and ϕ. As noted in Livshits, MacGee, and Tertilt (2007) and Athreya (2004), dropping transactions costs can potentially deliver the trends in the default rates and debts observed in the data, so these changes are worth examining as competitor stories. No household in the model would default on any debt less than this cost (i.e, when $b > -\Delta$), so higher values of Δ can support larger debts in general. Changes in Δ only alter the length of the initial flat segment where risk-free borrowing is sustained (Figure 1 plots price q as a function of borrowing b). Changes in ϕ only shift the pricing functions up and down (see Figure 1), altering the cost of issuing any given amount of debt. Thus, neither change will alter the variance of interest rates that agents receive, as they affect all agents symmetrically. To get a change in the distribution of interest rates, one needs to generate changes in the slope of the pricing functions. As a result, stories that place falling transactions costs at the heart of the changes in the unsecured credit market cannot account for the homogeneity observed in the earlier period. More details on experiments with (λ, Δ, ϕ) are available upon request.

Having given a flavor of how pricing works in the model, we turn now to credit availability under full information; the 'supply side' of the credit market is seen most clearly in the pricing of debt facing households in varying states of income. For questions regarding unsecured credit, the young are the most relevant population, and we therefore focus on their access to credit. Figure 1 in the main text displays the pricing functions for college types at age 29 given both the low and high value of λ; as would be expected the higher the realization of e the more credit is available (at any given interest rate). For low realizations of e the pricing functions look like credit lines – borrowing can occur at a fixed rate (in this case, the risk-free rate) up to some specified level of debt, after which the interest rate goes to ∞. For higher realizations the increase in the interest rate is more gradual, meaning that some risky borrowing will occur in equilibrium; for some borrowers, the marginal gain from issuing debt is sufficiently high that they are willing to pay a default premium to do it. The pricing functions for noncollege types look similar but involve higher interest rates at any given level of debt. Similar pictures arise for older agents – they are weakly decreasing in debt with more gradual increases in interest rates for luckier agents. Middle-aged agents (say, age 45) can borrow significantly more than their younger counterparts, although they choose not to do so in equilibrium because they are saving for retirement. Given a high value of λ agents can borrow a lot more (as would be expected).

Lastly, a reason for our focus on improved information as a candidate explanation for the facts is the increase in dispersion of credit terms observed (and so paid in equilibrium). Figure 3 displays model's implications for the evolution of the variance of equilibrium borrowing rates over the life-cycle across the two main information regimes we consider. These rates are not weighted by the amount of debt, they are direct measures of dispersion computed identically to what we measure from the SCF. The dispersion in interest rates is fairly flat over the life-cycle, and the variance is systematically higher for the less educated groups, since those groups are the ones who borrow and default on the equilibrium path. In the model, agents are willing to pay fair premiums for the option to default and do so. Furthermore, since all pricing is actuarially-fair with respect to default risk, agents who pose higher risk will pay higher prices to borrow. We also note that in general, the less well educated pay higher interest rates throughout life than their better educated counterparts. Given the earnings process, this is not surprising. However, what is interesting here is that the model suggests that partial information may lead to *lower* average interest rates, or alternatively, that improvements in information may well coincide with the observation of more households borrowing at higher interest rates. The reason is intuitive, and reflects the pricing functions displayed earlier. In essence, loan interest rates under partial information are frequently high enough to discourage borrowing.

1.3.1 Information and the "Causes" of Bankruptcy

Here, we detail the joint role played by expenditure shocks and non-pecuniary costs of bankruptcy. The results are given in Table 1. Each cell in which the information regime is held fixed presents the joint distribution of expenditure shocks and non-pecuniary costs of those who have filed for bankruptcy. What is clear is that under all information regimes the bulk of filers have low non-pecuniary costs (high λ). More interestingly, defaulting households typically have *not* received the largest expenditure shock. Specifically, the median expenditure shock is zero, and yet households with this realization account for more than half of all filers. Perhaps the most interesting finding here is that as that information becomes more limited, we see that the fraction of households in bankruptcy who have a low non-pecuniary cost of filing *and* have received no expenditure shock grow dramatically. This result suggests that "strategic bankruptcy" is certainly a real possibility under PI regimes in a way that is more severely restricted under FI. In particular, while the latter group accounts for 46 percent under FI, they are 76 percent of all filers under PI(λ, e). With respect to income, Figure 4 shows that most defaulting households have not experienced extremely bad transitory shocks.

Despite the possibility that the joint distribution of expenditure shocks and non-pecuniary costs suggests some strategic filing, what is still true is that the high expenditure shock is rare, and so may not be seen often among filers for that reason alone. Nonetheless, such a shock may well "push" someone into bankruptcy when it occurs. Table 2 presents the conditional probability of bankruptcy, given a particular constellation of the two variables of interest. It is clear here that the high shock does indeed "cause" a disproportionate amount of bankruptcy, relative to its unconditional likelihood. Moreover, Table 2 makes clear the power of this shock: conditional on getting this bad expenditure shock, the probability of bankruptcy is not highly sensitive to either the non-pecuniary cost or the information regime that prevails. By contrast, filing probabilities vary much more substantially across information and non-pecuniary costs when the expenditure shocks are smaller.

1.4 Tables

Table 1: Expenditure Shocks and Non-Pecuniary Costs Among Filers

	FI		**PI (e)**		**PI (λ)**		**PI (λ, e)**	
	High λ	Low λ	High λ	Low λ	High λ	Low λ	High λ	Low λ
Low χ	0.3776	0.0118	0.1664	0.0065	0.0405	0.0000	0.0496	0.0000
Median χ	0.4621	0.0623	0.6299	0.0755	0.7793	0.0393	0.7608	0.0476
High χ	0.0515	0.0346	0.0725	0.0492	0.0849	0.0561	0.0853	0.0567

Table 2: Expenditure Shocks and Non-Pecuniary Costs as Triggers

	FI		**PI** (e)		**PI** (λ)		**PI** (λ, e)	
	High λ	Low λ	High λ	Low λ	High λ	Low λ	High λ	Low λ
Low χ	0.964%	0.030%	0.385%	0.015%	0.068%	0.000%	0.077%	0.000%
Median χ	15.380%	2.071%	18.986%	2.276%	17.123%	0.863%	15.430%	0.965%
High χ	26.422%	17.751%	33.728%	22.889%	28.792%	19.025%	26.702%	17.750%

Table 3: Credit Sensitivity

	$b < 0, m = 0$		$b < 0, m = 1$	
Mean	b	r	b	r
COLL	0.0718	7.33%	0.0688	9.74%
HS	0.0402	9.58%	0.0341	12.55%
NHS	0.0310	12.16%	0.0285	13.71%

1.5 Figures

References

[1] **Davey, Brian A. and Hilary A. Priestley**. 2002. *Introduction to Lattices and Order, Second Edition*. Cambridge Press.

[2] **Granas, Andrzej and James Dugundji**. 1986. *Fixed Point Theory*. Springer Publishing.

[3] **Press, William H., Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery**. 1993. *Numerical Recipes in FORTRAN: The Art of Scientific Computing*. Cambridge University Press.

Figure 1: The Role of Transactions Costs

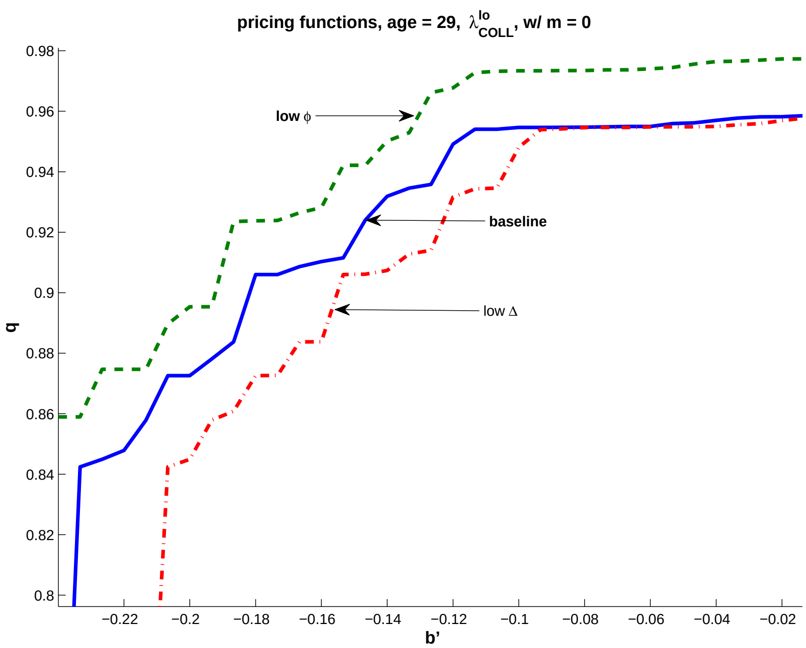

Figure 2: Labor Productivity over the Life-cycle

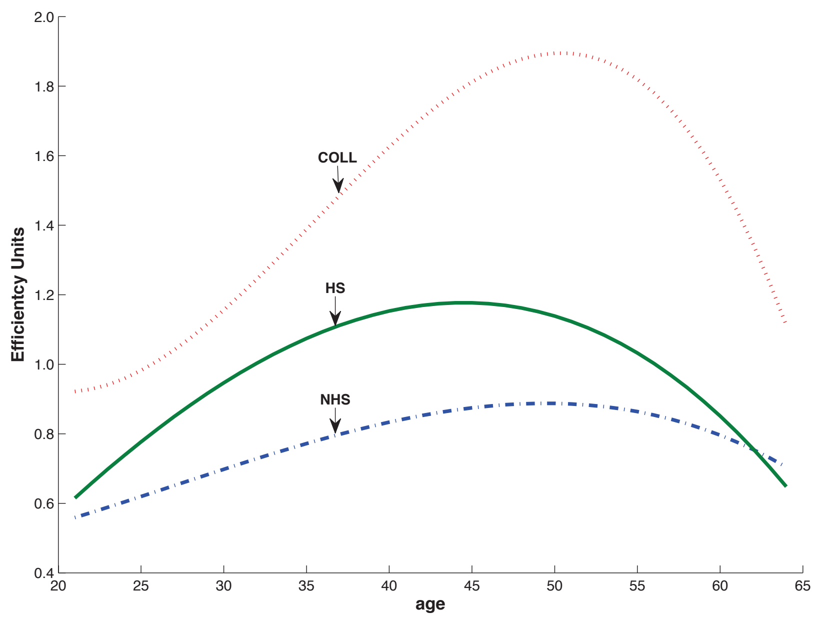

Figure 3: Variance of Equilibrium Borrowing Rates over the Life-Cycle

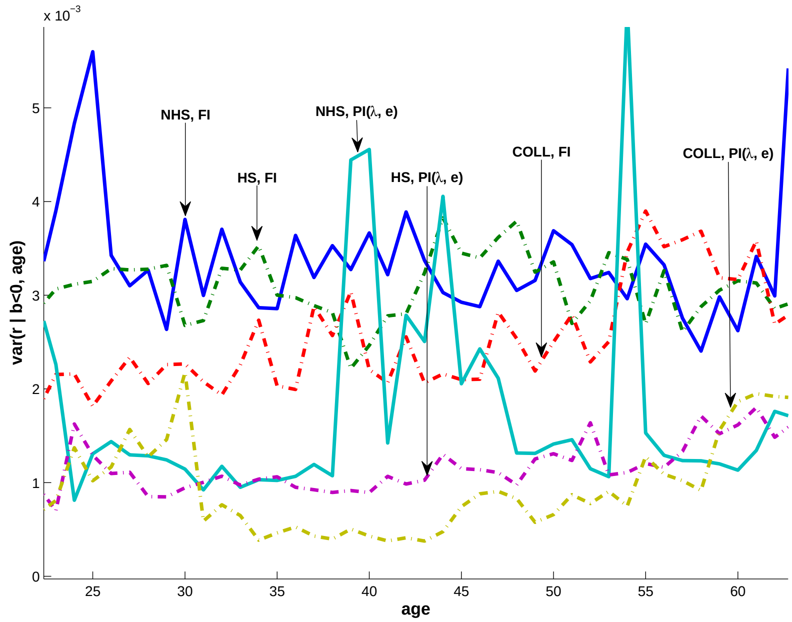

Figure 4: The Income Shocks of Defaulters

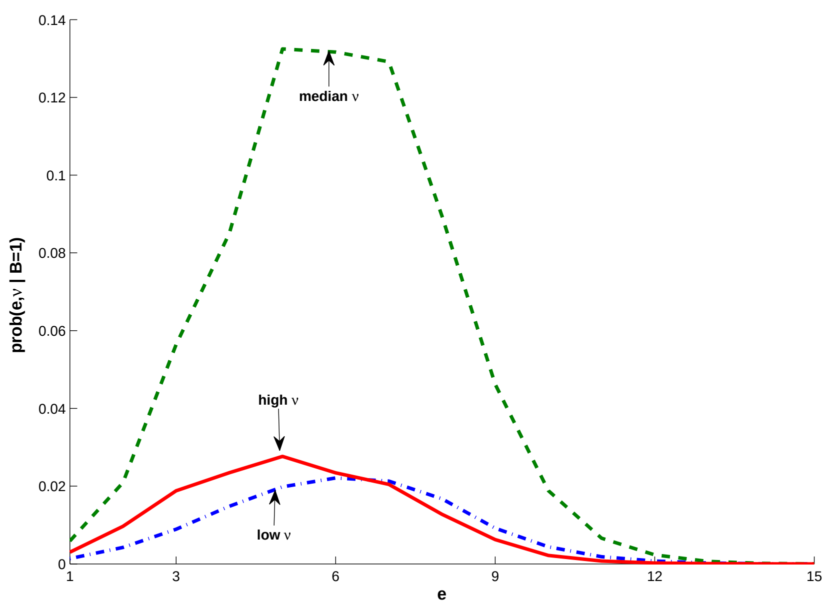

Figure 5: Credit Supply Under Full Information