Does Rising Income Risk Lead to Better Risk Sharing?*

Kartik Athreya[†] Federal Reserve Bank of Richmond

Xuan S. Tam[‡] University of Virginia

Eric R. Young[§] University of Virginia

April 6, 2008

Abstract

This paper studies the extent to which the rising use of unsecured credit and personal bankruptcy may have altered the transmission of increased income risk to consumption variability over the past several decades. We find that changes in income risk do not drive changes in default rates and bankruptcy and that changes in information do not alter the consequences of rising income risk for consumption volatility. If risk sharing has indeed improved over this period, the reasons do not lie in the unsecured credit market.

Keywords: Risk Sharing, Asymmetric Information, Bankruptcy, Default

JEL Classification Codes: D91, D82, D52, G22

---

*We thank helpful discussions with Leo Martinez, Erwin Quintin, and Pierre Sarte, as well as Brian Gaines for research assistance. We also thank participants at the Iowa Alumni Workshop for comments. Young acknowledges the financial support of the Bankard Fund for Political Economy. All errors are the responsibility of the authors. This paper does not reflect the opinion of the Federal Reserve Bank of Richmond or the Federal Reserve System.

[†]Research Department, Federal Reserve Bank of Richmond, kartik.athreya@rich.frb.org.

[‡]Department of Economics, University of Virginia, xst2q@virginia.edu.

[§]Corresponding author. Department of Economics, University of Virginia, ey2d@virginia.edu.

1. Introduction

Labor earnings have become more dispersed in the cross-section since 1980, a fact that has been established by a number of authors.[1] This increase reflects two changes in the dispersion of earnings. First, the variance of earnings between non-random, clearly-delineated groups, such as education, race, and gender, has grown, and second, the variance of growth in earnings *within* similarly-defined groups has increased. It is the second observation that suggests that earnings have become *riskier*. To this point, Krueger and Perri (2006) argues that the rise in inequality is heavily driven by bigger idiosyncratic shocks experienced after entering the workforce, attributing nearly half of the rise in inequality to an increase in the volatility of innovations to a highly-persistent component of the earnings process and one-fourth to an increase in the volatility of a purely-transitory component.[2] Primiceri and van Rens (2006) also find that an increase in the volatility of the highly-persistent component is a major contributor to rising earnings inequality.

The idiosyncratic component of income risk dwarfs other components, making the study of its effect on consumption variability a topic of enduring interest. Implicit in the assumption that idiosyncratic income risk may actually matter for consumption is the presumption that markets are incomplete. In particular, ever since Friedman (1957), the canonical 'income fluctuation problem' has been that of a household facing idiosyncratic risk that it may smooth *only* through the use of a non-contingent bond.[3] Based on lifecycle data on consumption and income, Storesletten, Telmer, and Yaron (2004) argues that models with only non-contingent debt can account for a variety of features of life-cycle smoothing. In contrast, Krueger and Perri (2006) argue that such models dramatically understate the insurance opportunities available to households and argue for a model based on a large number of contingent claims, but where limits to commitment prevent full risk-sharing.

In this paper, we evaluate how recent increases in idiosyncratic income risk are likely to have contributed to consumption variability for US households, holding other conditions fixed. To do this, we modify the "standard" income fluctuation problem in two ways: (1) we allow for default on debt and (2) we model improvements over time in the information held by lenders on the default

[1] We list only a few contributions to this vast literature: Moffitt and Gottschalk (2002), Bowlus and Robin (2004), Krueger and Perri (2006), Primiceri and van Rens (2006), Cunha and Heckman (2007), and Blundell, Pistaferri, and Preston (2008).

[2] The residual increase is attributed to an increase in a permanent shock realized upon entry into the workforce (education, race, or gender). In particular, the rising inequality between educational groups has received a significant amount of attention – see Acemoglu (2002) or Hornstein, Krusell, and Violante (2005) for surveys.

[3] Notable contributions include Schectman (1976), Bewley (1977), and Chamberlain and Wilson (2000).

risk of borrowers. Our paper is therefore part of the research agenda advocated by Deaton and Paxson (1994) and echoed in Blundell, Pistaferri, and Preston (2008): it is a study of a model that lies between the polar extremes of risk-free bond economies and models with complete sets of contingent claims. We are motivated to model bankruptcy and information improvements because of two observations regarding the unsecured loan market: (i) the rapid growth in use of unsecured debt and (ii) the rapid growth in default rates on these debts. Specifically, in precisely the period in which income risk appears to have grown, the relative size of unsecured revolving debts more than quadrupled, from 2 percent of GDP to approximately 9 percent of GDP (see Figure 1), while the annual US personal bankruptcy filing rate increased five-fold, from a rate of $300,000$ in 1980 to 1.6 million annually since 2000 (see Figure 2).[4]

Both the accumulation of debt and the use of bankruptcy reflect the attempts of households to smooth consumption in response to shocks. With respect to debt, Gross and Souleles (2002) and Sullivan (2008) argue that in the data unsecured credit is used to smooth consumption – households whose income prospects decline tend to accumulate more debt.[5] With respect to bankruptcy, Sullivan, Warren, and Westbrook (2000) report that at least 80 percent of filers cite income interruption – particularly employment disruption – as one of their reasons for filing.[6] Holding income risk fixed, one might expect that consumption variability should fall as a result of the greatly expanded use of credit and default. Furthermore, one might expect that the increase in income risk documented at the outset led households to both borrow more and default more frequently. We show that neither prediction is likely to be correct.

The fundamental reason for the breakdown of the basic intuition given above is that the supply side of the credit market will not allow such changes, all else equal. In particular, the increased incentives to borrow and default created by growth in earnings risk will trigger the repricing of debt, choking off opportunities to use both debt and bankruptcy as insurance mechanisms.[7] Therefore,

[4]The expansion of unsecured debt does not seem to be merely a replacement of other, more collateralized forms of credit; rather, households now appear to be able to borrow more against future income than two decades ago. Specifically, Bird, Hagstrom and Wild (1995), and Narajabad (2007) both argue that overall access to credit expanded for households; with the former study showing this expansion was largest for the riskiest populations. Similar trends appear in UK data; Attanasio *et al.* (2002) and Blundell and Preston (1998) document a rise in earnings risk and Power and Young (2008) show a sharp rise in bankruptcy filings. In the US, the rise in filings was large enough that approximately one in six US households filed for bankruptcy in the period 1980-2000

[5]Calem, Gorday, and Mester (2006) argue that credit card usage has high persistence. This result seems consistent with the use of debt to smooth both life-cycle changes in income, as well as highly-persistent income shocks.

[6]White (2007) argues that this survey data is not reliable; Fay, Hurst, and White (1998) in fact argue that there is no evidence that bankrupt households have experienced negative economic shocks. However, it is also evident in other data (SCF, PSID) that bankrupt households are generally in poor economic shape; see Budría Rodríguez *et al.* (2002), Livshits, MacGee, and Tertilt (2006), and Chatterjee *et al.* (2007).

[7]It is legitimate to question whether unsecured debt markets work this way, given the observation that many

the growth in debt and default, especially in the face of rising income risk, demands an explanation as to why households were able to borrow and default more than ever before. A growing empirical literature – Chandler and Parker (1989), Barron and Staten (2003), Furletti (2003), and Edelberg (2006), to name but a few – has assembled evidence that the supply side of the unsecured credit has expanded in response to improvements in the ability of lenders to distinguish between borrowers' risk characteristics. Athreya, Tam, and Young (2007) shows that a quantitative equilibrium model that allows for asymmetric information can indeed generate facts consistent with the view that information held by lenders about borrowers has improved; Narajabad (2007) finds similar results using a model of improved, but symmetric, signal precision.[8] We use the model from Athreya, Tam, and Young (2007) to examine the link between income risk, default, and consumption smoothing.

We address three narrowly focused questions. First, in a class of models capable of generating unsecured borrowing and equilibrium bankruptcy, how do changes in income risk lead to changes in consumption? Second, how does the answer to the preceding question depend on the amount of information possessed by lenders? Third, what are the welfare implications of the rise in income risk and improved information when viewed through our preferred model of credit markets?

We show that in general, increases in income risk will lead to increases in consumption risk, *whether or not* lenders are fully informed on the default risk posed by a given household. With respect to overall consumption variability, as measured by the variance of log consumption over the life-cycle, we show that surprisingly it is *higher* under a "full-information" regime than under a partial information one. This result obtains because when information is very limited, no young households can borrow – as a result, intertemporal smoothing is hindered, but as households enter middle age few have large debts that they must repay which then keeps consumption variability low. In contrast, when information improves lenders begin to price default risk, and the only households who borrow are those with relatively good future earnings prospects. Intertemporal smoothing improves for these households; their intertemporal needs are high, and the premium they pay for borrowing is low. *Ex post*, households who are unlucky first allow consumption to fall, and if they are very unlucky finally file for bankruptcy. In the interim, consumption variability between the unlucky and lucky diverges. With respect to income risk and welfare, we find that holding fixed the ability to default, higher income risk lowers welfare, but the welfare loss of increased income

credit cards appear to offer lines (fixed rates and limits); Mateos-Planas (2007) and Mateos-Planas and Ríos-Rull (2007) are two attempts to model lines. Since credit card agreements are explicit about the ability of lenders to reprice at will, we believe that our arrangement captures key features of the data.

[8] See also work by Livshits, MacGee, and Tertilt (2007), Sánchez (2007), and Drozd and Nosal (2007).

risk can be mitigated by harshly penalizing default.

Our punchline, therefore, is that changes in household use of an *a priori* plausible set of markets and institutions (unsecured debt and bankruptcy) are not the mechanisms that seem to have decoupled income and consumption volatility. Our work therefore appears to stand in contrast to Krueger and Perri (2006). The reason for the difference in outcomes lies in the way in which income risk matters for limited commitment. Krueger and Perri (2006) assumes that failure to deliver on financial obligations is met with permanent autarky, which immediately means that increased income risk actually mitigates the limited commitment problem. Given the markets are otherwise assumed complete, consumption smoothing improves. By sharp contrast, in our model markets are not assumed to be complete and increases in income risk worsen the limited commitment problem created by non-waivable personal bankruptcy. As a result, the young and the unlucky have difficulty accessing unsecured credit and face greater difficulty smoothing consumption.

Despite the preceding implication, our approach is not inconsistent with the basic mechanism in Krueger and Perri (2006), precisely because we do not allow for improvements in insurance arrangements that may have occurred as a result of heightened income risk. In fact, our findings suggest that one must assign even more importance to changes in the direct insurability of income risk, given that credit markets are not likely to have helped. Speaking to growth in the importance of "pure" insurance markets, Primiceri and van Rens (2006) argue that much of the increase in income risk was anticipated and therefore insurable, while Heathcote, Storesletten, and Violante (2006) argue that while much of the increase in income volatility represents genuinely increased *risk* to households it has been in components that are insurable. However, it has not been possible to subject these alternatives to the kind of test we have performed here, forcing the stylized models to confront observable trends in the specific markets that 'decentralize' a given story. Thus, a useful aspect of our model that differentiates it from complete-markets approaches is that there is a natural link from the assets purchased and sold by agents in the theoretical model to observable transactions in credit markets and bankruptcy court.

Our final comments in this section are intended to spell out further precisely what we do and do not do here. Our paper is about the *consequences* of changes in income inequality, not the causes. With respect to these consequences, the innovations in our work are twofold. First, and most importantly, ours is the first detailed analysis of how consumption changes with income risk in the presence of unsecured debt and bankruptcy. Second, we are the first to accommodate the role played by income risk in conjunction with reductions in the asymmetry of information between lenders and borrowers suggested by recent empirical work.

Lastly, while our model makes predictions for a variety of conditional moments of consumption and for changes in these moments in response to changes in income risk, we do not attempt to directly confront these predictions with the data. Our reason for this decision is that the evolution of consumption inequality in the data is at present in dispute. Krueger and Perri (2004, 2006) argue that the rising risk in earnings has not been fully transmitted to consumption, leading to a much smaller increase in consumption volatility. Davis (2004) argues against the characterization of consumption inequality of Krueger and Perri (2004, 2006). Specifically, Davis (2004) notes that Slesnick (1992) and Battistin (2003) find serious discrepancies between aggregate measures constructed from the CEX and from the National Income and Product Accounts, while Attanasio (2002) and Battistin (2003) find discrepancies between the Interview Survey and the Diary Survey of the CEX; the Diary Survey, which is a more comprehensive survey intended to mitigate measurement error, shows rising consumption inequality. Fisher and Johnson (2006) note that a comprehensive measure of consumption expenditures from the PSID also display more inequality. More recently, Attanasio, Battistin, and Ichimura (2007) finds that nondurable consumption risk did not remain flat, but rather *rose* nontrivially in the 1990's.

The paper is organized as follows. We first review the implications of increased income risk on consumption volatility in several models of consumption smoothing, *none* of which allows for equilibrium default on an otherwise risk-free debt. We also lay out a simple two-period model of default to highlight some key links between default and consumption risk. We then lay out our benchmark model of defaultable debt, and assess its implications for the impact of income risk, both for measured bankruptcy rates, but more crucially, for consumption smoothing. A final section relaxes the symmetric information assumption and compares a low-risk partial information setting with a high-risk full information one. We then conclude with some proposals for how to take our model's implications to the data.

2. Smoothing in Models without Equilibrium Default

It is instructive to review what is known about the mapping from income risk to consumption volatility in models in which default cannot occur on the equilibrium path; our discussion will be based on comparing steady states for four models – Aiyagari (1994) (Model **A**), Zhang (1997) (Model **Z**), Krueger and Perri (2006) (Model **KP**), and Mendoza, Quadrini, and Ríos-Rull (2007) (Model **MQR**) – with different assumptions about the span of assets traded and the determination of holding limits on those assets. Krueger and Perri (2006) conduct similar exercises in these models, but they concentrate on transitional dynamics rather than simple comparisons across steady states. For the sake of comparison with our main theoretical model – for which transitional dynamics cannot readily be computed – we choose the simpler (albeit somewhat less informative) route. The appendix details explicitly the models that we use for this exercise; we do not attempt to calibrate these models as closely to the data as we do our main model, but rather use them to demonstrate general properties. Our results are not sensitive to the exact parameter choices made in each model.

We also use this section to discuss a simple two-period model with default to illustrate precisely how the option to renege on payments in some states of the world alters the nature of consumption smoothing. The extant literature on two-period default models – such as Zame (1993) and Dubey, Geanakoplos, and Shubik (2006) – focuses on the existence and efficiency of equilibrium; they are silent on the precise manner in which default is used to alter the consumption allocation state-by-state. Our discussion highlights two key possibilities – (i) default can improve consumption smoothing by introducing contingency and therefore mitigating precautionary motives and (ii) default can worsen consumption smoothing by restricting borrowing. Our two-period examples are intended to starkly illustrate how households in differing circumstances may view the option to default differently; our quantitative model is intended to sort out these competing effects.

In Figure (3) we plot the cross-sectional standard deviation of labor earnings on the horizontal axis and the cross-sectional standard deviation of consumption on the vertical; these values are obtained by feeding in a sequence of progressively more volatile labor income processes into each model.[9] The figure shows clearly that the models with uncontingent claims – Models **A** and **Z** – tie consumption volatility tightly to income volatility, while the model with a full set of contingent claims – Model **KP** – produces no relationship between the two statistics. The intermediate case – Model **MQR** – produces only a mild positive connection between income risk and consumption risk. We now discuss the properties of each model that generates this graph.

In model **A**, agents can self-insure only through buffer stocks; that is, they increase their savings a bit in order to reduce the probability that a string of bad shocks will cause them to run their assets down to the debt limit. They do not reduce this probability to zero, however, because to

[9]Because each of these models satisfies the mixing condition from Aiyagari and Alvarez (2001), the cross-sectional standard deviation is also the time series standard deviation of a single agent. Thus, we can report cross-sectional statistics as measuring volatility. This property will not hold in our main model.

do so requires infinite assets.[10] When income risk rises, the conditional probability of a sequence that exhausts a household's assets rises, leading to more precautionary savings; but since this precautionary savings does not eliminate the increased risk, consumption becomes more volatile. General equilibrium considerations play a role here, as the increased demand for assets in the presence of higher risk reduces the return on assets; in an endowment setting (such as Huggett 1993) the rise in the price of the asset would entirely wipe out the increased asset holdings, but with capital accumulation the decline in the return only eliminates some of the increased asset supply.[11]

In Model **Z** the increase in precautionary savings is mitigated by the fact that the debt limit increases.[12] When debt limits are determined by the value of exclusion from credit markets, they naturally increase as the value of that state declines; with risk averse agents, increased variance translates immediately into lower utility in autarchy and therefore larger debt positions can be sustained. As Figure (3) clearly shows, the link between consumption risk and income risk is almost unchanged when debt limits adjust in this fashion. This result is quantitative – we find that increased income risk does not move the default-free debt limits much.

In contrast, Model **KP** permits the agents to trade a full set of contingent assets. When income risk increases, agents can fully diversify away this risk by trading, provided that they do not run afoul of the debt limits. As in Model **Z**, these debt limits increase when income risk increases, permitting households to take larger negative positions in these assets; naturally, the effect is to reduce the link between income risk and consumption inequality. The key result for the perfect insurance outcome is shown in Ábrahám and Cárceles-Poveda (2007) and Córdoba (2007): the amount of capital held in the steady state is sufficient to eliminate the effect of the debt constraints on consumption risk (for reasonable parameter values). The combination divorces consumption risk from income risk in this model.[13] We note here that the endogeneity of the

[10]Consider the Euler equation for a consumer in Model **A** with $r - \delta \leq \beta^{-1} - 1$. If the inequality holds strictly, agents will never accumulate an infinite amount of assets and therefore are exposed to consumption risk given a bad enough sequence of shocks. If $r - \delta = \beta^{-1} - 1$, then marginal utility is a nonnegative supermartingale and consumption must converge to zero almost surely unless assets are infinite.

[11]In fact, in a Huggett (1993) economy the relationship between income risk and consumption risk is nearly 1 to 1. It is not precisely 1 to 1 because higher moments of the distribution of asset holdings shift.

[12]The interactions between debt limits and income risk are not necessarily straightforward in this model, since the wage does not remain constant. Young (2007) contains a discussion of the general equilibrium effects in a related model with elastic search effort.

[13]Krueger and Perri (2006) do not obtain full risk sharing in their economy because they do not consider the stationary limit. Ábrahám and Cárceles-Poveda (2007) note that if households are sufficiently impatient or the elasticity of output with respect to capital is small enough full risk-sharing may not obtain in the limit; a related result is provided in Córdoba (2007) who notes that the shock process plays a role as well.

debt limits is not particularly important – even if we set the debt limits to zero exogenously, the relationship between consumption and income volatility remains essentially zero.[14] The model based on Mendoza, Quadrini, and Ríos-Rull (2007) (Model **MQR**) – contingent claims trades limited by 'no diversion' and 'limited liability' restrictions – leads to intermediate levels of risk sharing.[15]

We have some reservations about Model **KP**, despite its apparent success at replicating the data. First, survey data strongly suggest that agents default when they receive bad shocks, while in Model **KP** agents have an incentive to default when they have received bad shocks in the past but a good shock in the current period.[16] Agents who get low shocks expect to receive transfers through the risk-sharing arrangement and thus have no incentive to renege; in contrast, agents with high debts who receive a good shock must begin to pay that debt down, so by defaulting they rid themselves of the obligation and can enjoy the persistent high income realizations without making any payments. Second, as noted in Kehoe and Levine (2007) one can price a defaultable security in the presence of a full set of contingent claims; one can therefore generate any default rate in equilibrium (and defaulting is economically meaningless). The viewpoint that default is meaningless does not seem to be shared by policymakers (the rising default rate certainly seems to have played a role in the recent change in bankruptcy rules). Thus, we feel justified in arguing that a model which has default occurring on the equilibrium path is needed to confront the trends in income risk and consumption inequality; we present this model in the next section.

2.1. Smoothing with Default: A Two-Period Model

In this section we present two examples that clarify the role that bankruptcy plays in determining the extent of risk-sharing and intertemporal smoothing. Consider a model with two period, $t = 1, 2$. In the first period, labor income w_1 is certain. Income in the second period is stochastic and may take two values $\left\{w_2^H, w_2^L\right\}$ with $w_2^H > w_2^L$; the probabilities of these two events is given by π_H and $1 - \pi_H$. The household maximizes the expectation of a time-separable utility function and discounts the future at rate β. We consider two cases. In the first case, the income process is set such that households benefit in terms of *ex ante* welfare from the option to default, while in the

[14] The complete elimination of a link between income risk and consumption risk is a quantitative result; the overall weak link in Model **KP** is not. A similar result can be obtained using the model from Córdoba (2007) – the only difference is that debt limits are set to imply limited collateralization of contingent debts.

[15] Model **MQR** spans the space of models without default via one parameter, which governs the loss of resources in the event of diversion; the appendix contains a simple presentation. We plot only one representative case.

[16] See our previous discussion of the evidence on the economic shape of bankrupt households.

second case we consider the opposite situation. Preferences will remain fixed across the two cases.

Without bankruptcy households can obtain the risk-free rate on debt, since they must repay debt in all income states in period 2. When default is allowed, – subject to a non-pecuniary cost – borrowers can decide not to repay debts acquired in period 1 but must promise to pay more in non-default states at the time of borrowing. The statement of the household's problem in each of these cases is presented fully in the appendix; in what follows, we present the results of two examples designed to highlight the role bankruptcy plays in both consumption smoothing over time and across states.

In the first case, the parameter values are set to capture three features: (i) to have a low-income state that is bad enough that the level of debt that can be *feasibly* repaid in that state is very low, (ii) to have the agent default only if the low income is realized, and (iii) preferences and income are such that the household would borrow under complete markets. In Figure (4), consumption under complete markets is labelled C and is given by the pair $\{c_1^{CM},\ c_2^{CM}\}$. Turning next to a "standard incomplete-markets model" in which only non-defaultable borrowing is allowed, the optimal choice is given by first-period consumption c_1^A, and because full insurance is no longer possible period 2 consumption is a *state-contingent* pair $\{c_{2L}^A,\ c_{2H}^A\}$. In this case, the household's inability to insure against uncertain second period income actually leads it to save, even though under complete markets it would have borrowed; this outcome is the usual precautionary savings effect found in models of uninsurable idiosyncratic risk (our two-period model is essentially Aiyagari 1994). The preceding two cases are polar, in the sense that the first offers complete contingency of assets and the second none. An impetus for allowing bankruptcy and default is to aid risk-sharing by making repayment contingent on circumstances that prevail in the future. The consumption allocation that results from permitting default is given by $(c_1^D, \{c_{2L}^D,\ c_{2H}^D\})$. The household again borrows in the first period, as it did under complete markets, but borrows an amount that it will repay only if $w_2 = w_2^H$, i.e. only in the good period 2 income state. The fact that it will default in one state is what makes the household's period 2 consumption uncertain, but as seen in the figure the ex post utility attained in each state is higher than that under standard incomplete markets. *Ex ante* expected utility is therefore also higher when default is allowed.

The key intuition for this case is that when households face an income process that is likely to generate higher future income, but features some rare and bad states, they will in general wish to borrow but the maximal risk-free borrowing they can access will be small. In our two-period example, the household cannot guarantee nonnegative second-period consumption unless it borrows less than w_2^L; if small, this limitation will hamper purely intertemporal smoothing. The reason that bankruptcy (subject to penalties) is able to improve on this outcome is that the household may now borrow more, knowing that it can default and pay the penalty if the income realization is w_2^L. It is here that the probability of low income comes into play; if $1 - \pi_H$ is low, then the premium that must be paid on risky loans will remain small and essentially allow the household to borrow against future income. In doing so, they achieve an allocation closer to complete markets than is possible with uncontingent debt. Thus, for households in situations where rare bad states are possible, default can not only help share risk *ex post* but also can facilitate borrowing to improve *inter*temporal smoothing.

We turn next to a case where the non-waivable right to bankruptcy actually harms both intertemporal and intratemporal smoothing (see Figure 5). The key feature of this example is an income process that is expected to grow between the first and second periods: both period 2 income states are higher than period 1 income. As before, consumption under complete markets is labelled C and is given by the pair $\{c_1^{CM}, c_2^{CM}\}$. When non-defaultable borrowing is allowed, the optimal choice is the triple $(c_1, \{c_{2L}^A, c_{2H}^A\})$, and the consumption allocation when default is permitted is given by $(c_1^D, \{c_{2L}^D, c_{2H}^D\})$. The example is set so that the household will (i) default in both states of income for any level debt taken in period 1, and (ii) even given the resulting inability to borrow chooses not to save as a precautionary measure in the first period. Therefore, $c_{2L}^D = w_2^H$, $c_{2H}^D = w_2^L$. Thus, the household now is no longer able to smooth intertemporally at all – the outcome is unambiguously worse than the outcome with pure non-contingent debt. The intuition for the harm done by the bankruptcy option is that households cannot commit to repaying debt and therefore cannot borrow. But given that they face very little utility risk in the second period – endowments in both states are higher than current income and are fairly close together – bankruptcy limits intertemporal smoothing without helping risk sharing.

In summary, the relationship between bankruptcy, consumption smoothing, and welfare is not obvious and is likely to vary nontrivially across households that differ in income prospects. As is well known, there is substantial heterogeneity in the slope and variability of household income across educational attainment; young, poorly-educated households in the US face relatively flat earnings profiles with substantial risk, while their college-educated counterparts face much steeper earnings growth when young and much smaller risks for much of their working lives. Therefore, the trade-offs created by bankruptcy in the first example correspond to those facing less-educated households, while the trade-offs faced by college-educated groups may be better described by the second example. However, these statements must also be conditioned on age; older households who are close to their peak in earnings may look more like the first example than the second. By modelling precisely this source of heterogeneity, we hope to provide reasonable conclusions on the way defaultable debt helps households manage changes in income and income risk.

3. A Quantitative Model of Unsecured Debt

In order to provide a quantitatively serious evaluation of the roles played by bankruptcy and unsecured borrowing in modifying the effect of income risk on consumption, we use the framework laid out in Athreya, Tam, and Young (2007). The salient features of that setting are: (1) households live for a maximum of $J < \infty$ periods; (2) households save and borrow using defaultable debt; and (3) households belong to two classes differing in terms of (i) access to financial intermediation and (ii) exposure to idiosyncratic risk.

3.1. Normal Households

If alive at age j, a normal household faces probability $\psi_j < 1$ of surviving to age $j+1$. Households value consumption, discount the future at rate $\beta < 1$, and attach a negative value λ_j (in utility terms) to all nonpecuniary costs of bankruptcy; we assume λ evolves stochastically.[17] Heterogeneous default costs appear relevant, as argued in Fay, Hurst, and White (1998), and permit us to match the data on both borrowing and default.[18] Households maximize expected utility:

$$\sum\nolimits_{j=1}^{J} \sum_{s^j} \left(\prod\nolimits_{i=0}^{j} \beta\psi_{j,y}\right) \Pi\left(s^j\right) \left[\frac{n_j}{1-\sigma}\left(\frac{c_j}{n_j}\right)^{1-\sigma} - \lambda_j \mathbf{1}\left(m_{j-1}=0, m_j=1\right)\right], \tag{3.1}$$

Households may differ in their "credit" status (described in detail further below) denoted by the state variable m_j. Let $\Pi\left(s^j\right)$ be the probability of a given history of events s^j through age j. $\sigma \geq 0$ is the coefficient of relative risk aversion. Households retire exogenously at age $j^* < J$, and n_j is the number of household members at age j.[19]

[17] This cost is intended to be a parsimonious method for capturing all of the factors and complications associated with bad credit other than ones associated with the direct filing cost or the terms of credit, as in Athreya (2002). As Gross and Souleles (2002) document, there is significant unexplained variability in the probability of default across households, even after controlling for a large variety of observables.

[18] Chatterjee *et al.* (2007) permit shocks to discount factors; these shocks are troublesome because they induce simultaneously a desire to borrow and to default. In contrast, shocks to stigma costs only alter the incentive to default (directly).

[19] We assume that labor is supplied inelastically. Many studies of rising inequality assert that changes in labor hours are a critical source of insurance (Storesletten, Telmer, and Yaron 2004 or Heathcote, Storesletten, and Violante

In each period a typical normal household makes a consumption-savings decision (c, b), where b is the amount of borrowing/saving and c is consumption. The household also selects a rule for default on any accumulated debt as a function of next period's shock. The cost of default is a non-pecuniary cost λ_j borne only in the period in which default occurs, a pecuniary cost Δ also paid only in the period of default, and a subsequent marker of 'bad' credit, $m_j = 1$. The parameter $\xi \in (0, 1)$ governs the likelihood of subsequently changing credit status. Under partial information, the price charged a household for issuing debt will in general depend on m, so that households with recent defaults will receive different credit terms than households with 'clean' credit. The probabilistic removal of this marker represents current regulations forcing the removal of bankruptcy from one's credit score after a fixed period (10 years in the US).

During working age, the household's budget constraint d is given by

$$c_j + q\left(b_j, I\right) b_j + \Delta \mathbf{1}\left(m_{j-1} = 0, m_j = 1\right) \leq a_j + (1 - \tau)\, W \omega_{j,y} y e \nu, \tag{3.2}$$

The function $q(\cdot)$ denotes an individual-specific bond price that depends on bond issuance b_j and a vector of individual characteristics I, a is net worth, Δ is the pecuniary cost of filing for bankruptcy, and the last term is after-tax current income. Log labor income is the sum of five terms: the aggregate wage index W, a permanent shock y realized prior to entry into the labor market, a deterministic age term $\omega_{j,y}$, a persistent shock e that evolves as an AR(1)

$$\log\left(e'\right) = \rho \log\left(e\right) + \epsilon', \tag{3.3}$$

and a purely transitory shock $\log(\nu)$. The shocks to labor earnings, ϵ and $\log(\nu)$, are both independent mean zero normal random variables with variances that are y-dependent. The budget constraint during retirement is

$$c_j + q\left(b_j, I\right) b_j \leq a + \theta W \omega_{j^*-1,y} y e_{j^*-1} \nu_{j^*-1} + \Theta W, \tag{3.4}$$

where for simplicity we assume that pension benefits are a composed of a fraction $\theta \in (0, 1)$ of income in the last period of working life plus a fraction Θ of average income (which is normalized

2006) and abstracting away from it will bias the outcomes in the model. While we do not disagree with the idea that agents can vary their labor supply in response to shocks, we suspect that the frictionless divisible labor model overstates the extent to which this variation is feasible. Since our interest is not in providing a complete accounting of risk sharing, we abstract from this issue.

1). Because bankruptcy is not a retiree phenomenon, we deliberately keep the specification of retirees simple. There do not exist markets for insurance against income or survival risk and we abstract from any sources of long-run growth.

The permanent shock in the model is intended to capture the large differences in lifetime income associated with different human capital level; y allows us to differentiate between non-high school, high school, and college education levels, as in Hubbard, Skinner, and Zeldes (1994). We abstract away from any connection between y and family size, but we permit the survival probabilities $\psi_{j,y}$ and the deterministic age-income terms $\omega_{j,y}$ to differ according to the realization of the permanent shock.[20] The differences in these life-cycle parameters will generate different incentives to borrow across types; in particular, college workers will have a steeper hump in earnings, which is critically important as it generates empirically relevant demand for borrowing early in the life cycle when observed default rates are highest.

We model the nonpecuniary cost process λ as a two-state Markov chain with realizations $\{\lambda_1, \lambda_2\}$ and transition matrix $[\pi_{ij}]_{i,j=1}^{2}$. We will discipline these costs by requiring that the benchmark model replicate several moments summarizing the behavior of unsecured credit markets.

3.1.1. Recursive Formulation

Let $d\left(b, e', \nu', \lambda'\right)$ denote the (indicator) decision rule governing bankruptcy if the household declares bankruptcy in the event that next period's shocks are $\left(e', \nu', \lambda'\right)$ and 0 otherwise. For the convenience of the reader, let $X = (e, \nu, \lambda)$ denote the exogenous components of the household state vector. In recursive terms, a household of age j with good credit $m = 0$ solves

$$v\left(a, y, X, j, m=0\right) = \max_{b, d(X') \in \{0,1\}} \left\{ \frac{n_j}{1-\sigma} \left(\frac{c_j}{n_j}\right)^{1-\sigma} + \beta \psi_{j,y} E_{X'} \left[\begin{array}{c} \left(1 - d\left(b, X'\right)\right) v\left(b, y, X', j+1, m'=0\right) + \\ d\left(b, X'\right) v^D\left(0, y, X', j+1, m'=1\right) \end{array} \right] \right\} \tag{3.5}$$

where

$$v^D\left(0, y, X', j+1, m'=1\right) = \left\{ \frac{n_j}{1-\sigma} \left(\frac{c_j}{n_j}\right)^{1-\sigma} - \lambda + \beta \psi_{j,y} E_{X'} \left[\begin{array}{c} \xi v\left(a', y, X', j+1, m'=0\right) + \\ \left(1-\xi\right) v\left(a', y, X', j+1, m'=1\right) \end{array} \right] \right\} \tag{3.6}$$

[20]Stochastic mortality and age-dependent family size do not play an important role in our model, but we keep them so as not to bias our estimates of the parameters of the model in systematic ways.

denotes the value from defaulting on previously accumulated debts; notice that households cannot borrow or save during the period in which they declare bankruptcy.[21] Households who currently have a bad credit indicator (i.e. $m = 1$) solve

$$v\left(a, y, X, j, m=1\right) = \max_{b, d(X') \in \{0,1\}} \left\{ \begin{array}{c} \frac{n_j}{1-\sigma}\left(\frac{c_j}{n_j}\right)^{1-\sigma} + \\ \beta\psi_{j,y} E_{X'} \left[\begin{array}{c} \xi \left[\begin{array}{c} \left(1 - d\left(b, X'\right)\right) v\left(b, y, X', j+1, m'=0\right) + \\ d\left(b, X'\right) v^D\left(0, y, X', j+1, m'=1\right) \end{array} \right] + \\ (1-\xi) \left[\begin{array}{c} \left(1 - d\left(b, X'\right)\right) v\left(b, y, X', j+1, m'=1\right) + \\ d\left(b, X'\right) v^D\left(0, y, X', j+1, m'=1\right) \end{array} \right] \end{array} \right] \end{array} \right\}. \tag{3.7}$$

3.1.2. Loan Market

We focus throughout on competitive lending. There exists a competitive market of intermediaries who offer one-period debt contracts and utilize available information to offer individualized credit pricing. Let I denote the information set for a lender and $\widehat{\pi}^b$: $b_{|I} \to [0,1]$ denote the function that assigns a probability of default to a loan of size b, given information I. $\widehat{\pi}^b$ is identically zero for positive levels of net worth and is equal to 1 for some sufficiently large debt level. The break-even pricing function must satisfy

$$q\left(b, I\right) = \begin{cases} \frac{1}{1+r} & \text{if } b \geq 0 \\ \frac{\left(1-\widehat{\pi}^b\right)\psi_j}{1+r+\phi} & \text{if } b < 0 \end{cases} \tag{3.8}$$

given $\widehat{\pi}^b$. The risk-free saving rate is given by r and ϕ represents transactions costs for lending, so that $r + \phi$ is the risk-free borrowing rate.[22] With full information, a variety of pricing arrangements will lead to the same price function. However, as is well known (e.g. Hellwig 1990), under asymmetric information settings outcomes often depend on the particular "microstructure" being used to model the interaction of lenders and borrowers. Under full information, our approach is completely standard (see Chatterjee *et al.* 2007 or Livshits, MacGee, and Tertilt 2006), as we seek a setting that delivers to households a function $q\left(b, y, e, \nu, \lambda, j, m\right) : b \to \left[0, \frac{1}{1+r}\right]$ that they can

[21] Under current U.S. bankruptcy code, a filing is likely to be judged fraudulent if it is immediately followed by significant asset accumulation; since filings normally take several months to resolve, it seems that a short period of exclusion is imposed as a punishment by law. Longer periods of exclusion that are imposed by the market are not consistent with competitive behavior.

[22] The pricing function takes into account the automatic default by those households that die at the end of the period.

take parametrically when optimizing; the compactness of the range for q implies that the household problem has a compact opportunity set and therefore possesses a solution.

The microstructure that underlies our pricing function is modelled as a two-stage signalling game between borrowers and lenders. In the first stage, borrowers name a level of debt b that they wish to issue. 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 price competition for borrowers. Thus, households view the pricing functions as schedules and understand how changes in their desired borrowing will alter the terms of credit (that is, they know $D_b q(b, I)$) because they compute the locus of Nash equilibria under price competition. Exactly how the pricing function depends on the components in I will be specified next. A formal statement of equilibrium for the full information case is omitted since it is entirely standard.

In the full information case, I includes all components of the household state vector: $I = (y, e, \nu, \lambda, j, m)$. Zero profit for the intermediary requires that the probability of default used to price debt must be consistent with that observed in the stationary equilibrium, implying that

$$\widehat{\pi}^b = \sum\nolimits_{e', \nu', \lambda'} \pi_e (e'|e) \, \pi_\nu (\nu') \, \pi_\lambda (\lambda'|\lambda) \, d \left(b(a, y, e, \nu, \lambda, j, m), e', \nu', \lambda'\right). \tag{3.9}$$

Since $d(b, e', \nu', \lambda')$ is the probability that the agent will default in state (e', ν', λ') tomorrow at debt level b, integrating over all such events *tomorrow* is the relevant default risk. With partial information, we will need to integrate over current states as well as future ones, since pieces of the state vector will not be observable. We defer a formal discussion of the partial information case until later in the paper.

3.2. Special Households

As mentioned above (and as discussed in more detail in Athreya, Tam, and Young 2007), we posit a measure μ_s of households who face neither idiosyncratic risk nor financial market frictions and therefore earn a higher rate of return on savings. These households face a present-value budget constraint

$$\sum_{j=1}^{J} \left(\frac{1}{1 + MPK - \delta}\right)^{j-1} c_j = A_1 \tag{3.10}$$

with lifetime wealth given by

$$A_1 = k_1 + \sum_{j=1}^{j^*-1} \left(\frac{1}{1+MPK-\delta} \right)^{j-1} (1-\tau) \, wy\omega_{j,y} + \left(\frac{1}{1+MPK-\delta} \right)^{j^*-1} \sum_{j=j^*}^{J} \left(\frac{1}{1+MPK-\delta} \right)^{j-j^*+1} (\theta wy\omega_{j^*-1,y} + \Theta) \, ; \tag{3.11}$$

Special households face the mean earnings profiles of the college type, and for convenience we assume logarithmic preferences. As a result, we can obtain closed-form representations for their decisions:

$$c_1 = \frac{1-\beta_s}{1-\beta_s^J} A_1 \tag{3.12}$$

$$c_j = \beta_s^j \, (1+MPK-\delta)^{j-1} \, c_1. \tag{3.13}$$

Given the decision rule for consumption, household capital stocks evolve as follows:

$$k_{j+1} = (1+MPK-\delta) \, k_j + (1-\tau) \, wy\omega_{j,y} - \frac{1-\beta_s}{1-\beta_s^J} \beta^{j-1} \, (1+MPK-\delta)^j \, A_1. \tag{3.14}$$

The purpose of the special households is to permit sufficient flexibility to match both the average amount of wealth held in the economy – which determines the marginal product of capital – and the concentration of wealth in the hands of a small minority of households for whom the option to default is irrelevant. An alternative to our approach is used in Chatterjee *et al.* (2007) and Sánchez (2007) – they calibrate the earnings process to match the US distributions of income and wealth. For our overlapping generations model the burden of this approach is immense as it forces us to extend the range of permissible asset positions out to very high positive levels. Given that we are not interested in the behavior of agents with positive net worth, the cost seems too much to bear.

3.3. Government

Government in this model exist simply to fund pension payments to retirees, which implies the following government budget constraint:

$$\tau(1-\mu_s)W\int y\omega_{j,y}e\nu\Gamma(a,y,e,\nu,\lambda,j<j^*,m)+\tau\frac{\mu_s}{J}W\sum_{j=1}^{j^*-1}y\omega_{j,y} \tag{3.15}$$

$$=\quad(1-\mu_s)W\int(\theta\omega_{j^*-1,y}ye_{j^*-1}\nu_{j^*-1}+\Theta)\Gamma(a,y,e,\nu,\lambda,j\geq j^*,m)+\frac{\mu_s}{J}W\sum_{j=1}^{j^*-1}(\theta\omega_{j^*-1,y}y+\Theta).$$

3.4. Market Clearing

Loan markets clear when the risk-free saving rate equals the marginal product of capital net of depreciation δ:

$$r=\alpha K^{\alpha-1}N^{1-\alpha}-\delta-\vartheta.$$

The parameter ϑ denotes intermediation costs that apply to both saving and borrowing. Given the resources destroyed in intermediation, total capital solves the equation

$$\begin{aligned}K\quad=\quad&\frac{1-\mu_s}{1+\alpha K^{\alpha-1}N^{1-\alpha}-\delta-\vartheta}\int b(a,y,e,\nu,\lambda,j,m)\Gamma(a>0,y,e,\nu,\lambda,j,m)+\\&\frac{1-\mu_s}{1+\alpha K^{\alpha-1}N^{1-\alpha}-\delta-\vartheta+\phi}\int\left(1-\pi^b(a,y,e,\nu,\lambda,j,m)\right)b(a,y,e,\nu,\lambda,j,m)\Gamma(a<0,y,e,\nu,\lambda,j,m)+\\&\frac{\mu_s}{J}\sum_{j=1}^{J}k_j.\end{aligned} \tag{3.16}$$

Total labor input is the weighted sum of the labor inputs of the two classes of households, and so is given by

$$N=(1-\mu_s)\int y\omega_{j,y}e\nu\Gamma(a,y,e,\nu,\lambda,j,m)+\frac{\mu_s}{J}\sum_{j=1}^{j^*-1}y\omega_{j,y} \tag{3.17}$$

where $\alpha\in(0,1)$ is the elasticity of output with respect to capital. The labor market clears when the aggregate wage index equals the marginal product of labor:

$$W=(1-\alpha)K^{\alpha}N^{-\alpha}. \tag{3.18}$$

Goods market clearing occurs when total consumption plus total transactions costs equals total output less depreciation:

$$C = (1-\mu_s)\int c\Gamma(a,y,e,\nu,\lambda,j,m) + \frac{\mu_s}{J}\sum_{j=1}^{J} c_j \tag{3.19}$$

$$T = (1-\mu_s)\int \left[\phi\mathbf{1}(b<0)\,b + \Delta\mathbf{1}\big(m=0,m'=1\big) + \vartheta b\right]\Gamma(a,y,e,\nu,\lambda,j,m) \tag{3.20}$$

$$C+T = K^{\alpha}N^{1-\alpha} - \delta K, \tag{3.21}$$

where $\mathbf{1}(A)$ is the indicator function for set A.

4. Calibration

We set $\sigma = 2$. We set $\theta = 0.35$ at an exogenous retirement (model) age of 45 and $\Theta = 0.2$, yielding an overall replacement rate around 55 percent. The income process 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. Figure (6) displays the path for $\omega_{j,y}$ for each type; the large hump present in the profile for college-educated workers implies that they will want to borrow early in life to a greater degree than the other types will (despite their effective discount factor being somewhat higher due to higher survival probabilities). We hold the age-dependent component fixed and assume that the persistence parameter is unchanged across periods as well. For each period the process is discretized with 15 points for e and 7 points for ν. The resulting processes for the low risk period are

$$\log(e') = 0.95\log(e) + \epsilon'$$
$$\epsilon \sim N(0, 0.033)$$
$$\log(\nu) \sim N(0, 0.04)$$

for non-high school agents,

$$\log(e') = 0.95\log(e) + \epsilon'$$
$$\epsilon \sim N(0, 0.025)$$
$$\log(\nu) \sim N(0, 0.021)$$

for high school agents, and

$$\begin{aligned}\log(e') &= 0.95\log(e)+\epsilon' \\ \epsilon &\sim N(0,0.016) \\ \log(\nu) &\sim N(0,0.014)\end{aligned}$$

for college agents; the measures of the three groups are 15, 58, and 22 percent, respectively, and the measure of special agents is $\mu_s = 0.05$.[23] We set $\phi = \vartheta = 0.03$ to generate a 6 percent spread between risk-free saving rates received by normal households and the risk-free borrowing rate.[24] 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 \$1200.[25] Finally, $\xi = 0.25887$ implies that 95 percent of households who do not file for bankruptcy again will have clean credit after 10 years.

We introduce changes in the income process consistent with the measurements in Krueger and Perri (2006): the variance of the permanent component, the persistent component, and the transitory component contribute 36 percent, 40 percent, and 24 percent to the overall rise in cross-sectional variance, respectively, with the total unconditional variance of lifetime income assumed to rise by 20 percent. Under the high-risk regime we calibrate the other parameters (β,λ_1,λ_2,β_s,α,δ) to match six targets: a measure of borrowers equal to 12.5 percent, an aggregate negative net worth to GDP ratio of 0.6 percent, a bankruptcy filing rate of 1.2 percent, a risk-free saving rate of 1 percent, a 70 percent labor income share, and an annual depreciation rate of 10 percent.[26] Our target for bankruptcy is consistent with a model in which only income is uncertain; that is, there are no shocks to expenses. Expense shocks create involuntary creditors that allow households to suddenly acquire very large debts with no corresponding change in measured consumption. The difficulties in measuring rare "catastrophic" shocks, their true "uninsurability" (for example, Medicare and emergency rooms are always available to deal with medical shocks), and their persistence can lead to a serious mismeasurement of the role of credit use and bankruptcy for managing income risk.

[23] Recall that all special agents are college-educated types, so the total measure of college agents is 20 percent. We normalize units of measurement such that $N = 1$.

[24] The spread between saving rates and capital returns is thus equal to 3 percent, consistent with transactions costs measured by Evans and Schmalensee (1999).

[25] This cost is an estimate inclusive of filing fees, lawyer costs, and the value of time.

[26] The value for the ratio of negative net worth to GDP is taken from Budría Rodríguez *et al.* (2002); we exclude the self-employed and consider both credit cards and installment loans as unsecured debt. For the measure of borrowers, we roughly average the numbers from Chatterjee *et al.* (2007) – 6.7 percent – and Wolff (2004) – 17.6 percent to get a ballpark measure of 12.5 percent.

We therefore calibrate the model's bankruptcy target to be net of this measure. Sullivan, Warren, and Westbrook (2000) report that one-fifth of filers report that health played *some* role in their decision to file. Using an overall filing rate of 1.5%, our target becomes 1.2%.

We calibrate the high-risk period and then reduce the variances to match the earlier period; we argue that the earlier period is best interpreted as suffering from asymmetric information and that model is too computationally demanding to be calibrated effectively. The resulting calibrated parameter values are listed in Table (1), along with all other parameters for the convenience of the reader; we do reasonably well at matching the targeted moments, but not perfectly.

5. Results

We present our results in three steps. First, we investigate the role of bankruptcy in the transmission of income risk to debt accumulation and consumption risk under full (and symmetric) information (FI) under the relatively high risk (H) income process corresponding to recent (2004) data; here, we compare outcomes with and without the bankruptcy option. This case represents our baseline calibration and is denoted (FI-H). Given our benchmark, we then isolate the role of income risk by studying allocations obtaining under full information under the relatively lower-risk (L) income processes faced by households in the early 1980's. We denote this case (FI-L). Third, we evaluate the role that information changes play in this transmission; this discussion is deferred until after we describe the model with asymmetric information. Throughout our analysis we restrict attention to comparisons of steady states.[27]

To understand how unsecured credit affects unconditional consumption volatility, it is critical to recognize that both intertemporal and intratemporal smoothing will be affected by changes in bankruptcy policy and credit availability (as we demonstrated in the two-period examples). Therefore, in what follows, we will repeatedly focus on a useful decomposition for the variance of (log) consumption:

$$V\left(\log\left(c\right)\right)=E\left[V\left(\log\left(c\right)|j\right)\right]+V\left(E\left[\log\left(c\right)|j\right]\right); \tag{5.1}$$

the total variance of consumption is the mean of the variances of consumption conditional on being age j plus the variance of mean consumption conditional on being age j. The first term yields the effectiveness of any given market arrangement on *intra*temporal smoothing – it is the average

[27] Krueger and Perri (2006) investigate an infinite-horizon model and are therefore able to explicitly compute transitional dynamics. For our OLG model the memory requirements for even a 10-year transition path exceed the capacity of a 64-bit addressing system.

dispersion of consumption occurring within households of any given age and so provides a natural measure of "risk sharing". The second term measures *inter*temporal smoothing by capturing the extent to which mean consumption – which is precisely what would obtain for *all* households under complete insurance markets – evolves over the lifecycle. Once bankruptcy is a possibility and income risk is even partially uninsured, credit supply may tighten and lead to an increase in the contribution of intertemporal movements to overall consumption volatility.[28]

5.1. What Effect Does Bankruptcy Have on Consumption Smoothing?

Our first step in studying the role that bankruptcy plays in smoothing consumption is to compare our benchmark model to an otherwise identical one where bankruptcy is prohibited. Figure (7) plots the variance of log consumption over the lifecycle for our benchmark model with and without the bankruptcy option (under the high-risk regime).[29] The figure suggests a trade-off with respect to intratemporal smoothing – when bankruptcy is an option, households experience higher consumption variance when young but lower consumption variance when old. The intuition is similar to that discussed in Athreya (2008). Without default, households can borrow at the risk-free rate up to the natural debt limit and can therefore achieve good intertemporal smoothing. *Ex post*, any borrowing done early in life must be repaid, leaving unlucky households relatively exposed to bad shocks later in their lifecycle. In contrast, when bankruptcy is permitted borrowing does not have to be repaid if the household gets bad enough outcomes; thus, intratemporal smoothing is improved later in life.

A second cause of this reduction in variance beyond middle-age is a borrowing rate that moves against unlucky households, leading to a curtailment of borrowing when young and a decline in the ability to respond to shocks early in life. Figure (8) plots the price functions for an age 29 household over various different values of the shocks (for brevity we plot only the college type with the low stigma cost; the top panel is for the lowest realizations of e, the middle panel is the middle realizations, and the bottom panel is the highest realizations of e, and we ignore the transitory shock since it does not alter the pricing functions in a quantitatively-important way). With full information about borrowers, any change in circumstances that increases default risk – such as a low realization of e – will increase the borrowing rate and therefore inhibit the ability to smooth when

[28] In any infinite horizon setting without aggregate risk – such as Chatterjee *et al.* (2007) – the second term is zero.
[29] The picture for the low-risk regime looks quite similar and is therefore omitted.

young (since only the young borrow it is mainly their consumption smoothing that is affected).[30]

Table (2) presents our consumption-smoothing decomposition for the two cases; overall, consumption volatility is higher when bankruptcy is permitted. The presence of bankruptcy seems to create a trade-off between intertemporal and intratemporal smoothing. In particular, with respect to intratemporal smoothing, HS and NHS households both benefit from the presence of the bankruptcy option, while the college educated see a slight increase in the mean variability of consumption over the life-cycle. By contrast, intertemporal smoothing is uniformly better when bankruptcy is ruled out. Livshits, MacGee, and Tertilt (2006) was the first paper to note that bankruptcy may induce a tradeoff whereby intertemporal smoothing is restricted but intratemporal smoothing can be improved. However, this trade-off is simply an artifact of aggregating the increased variance experienced by older households who borrowed when young with the reduced variance of the young; in fact, bankruptcy in the calibrated model reduces both components of consumption smoothing for young households, so there is no trade-off for households who end up using the option to default in equilibrium. Figure (9) displays default rates over the lifecycle; since default is an activity for the young, the higher volatility experienced by the young shows that bankruptcy does not assist in smoothing consumption.

5.2. How Does Bankruptcy Transmit Rising Income Risk to Consumption?

We now discuss how the presence or absence of bankruptcy alters the transmission of rising income risk into consumption. The main result of this section is that the option to declare bankruptcy, by itself, does not mitigate the transmission of increased income risk into increased consumption volatility. Figure (10) plots the variance of log consumption over the life cycle under low and high income risk when bankruptcy is prohibited, while Figure (11) does the same when bankruptcy is permitted.

In Figure (10), the difference in consumption inequality tracks the difference in income inequality (see Figure 12). With bankruptcy and full information, the gap between the curves is nearly identical. If one views our model as a reasonable instrument for measuring the effects of bankruptcy and default, one should conclude that the ability to default is unlikely to have played a major role in reducing consumption risk. Our findings here are consistent with a similar experiment conducted

[30] One might argue that credit card arrangements look less like bonds and more like lines (a fixed pair of rates and limits). As we noted above, actual arrangements involve an option by the lender to reprice at will. Matteos-Planas (2007) and Matteos-Planas and Ríos-Rull (2007) are preliminary attempts to study credit lines in the presence of default; see also Athreya (2002).

in Livshits, MacGee, and Tertilt (2006).

Why does bankruptcy fail to buffer household consumption against increased income risk? Two factors are at work. First, the increase in income risk modelled here posits growth in the volatility of the persistent component of income, in addition to transitory factors. As a result, the cross-sectional variance of income grows steadily with age and therefore so does the difference between income volatility in the two regimes. Figure (12) shows that early in the lifecycle – when unsecured debts are most useful and default is most prevalent – the two income regimes generate relatively small differences in the variance of cross-sectional income. This fact sharply limits the ability of bankruptcy to blunt increases in income risk.

The second factor making bankruptcy a poor form of social insurance against increased income risk is that the conditional mean future expected income of a household with a bad current shock is even lower than before. This force is especially powerful when the increase is driven by growth in the persistence of income shocks, As a result, the choice facing an indebted household is between rolling over increasingly expensive debt for another period, in hopes of a better draw of income, or defaulting. By middle age, expected future income is declining at a rate which makes further borrowing for the unfortunate expensive, as it would trigger default at high rates; the result is that intermediaries do not extend credit to such households. Thus, increased income risk does not markedly change overall borrowing; as seen in Figure (9) it is the only low-risk college-educated workers who default at appreciable rates beyond very young ages. In summary, increased income risk does not generate enough income dispersion early in life, full information makes high-risk borrowing more expensive among the young, and bankruptcy is irrelevant for older households. Thus, under full information the ability to declare bankruptcy does not meaningfully limit the transmission of default risk to consumption risk.

Table (5), which appears later in the paper, explicitly computes the amount of consumption that households would pay to live under different regimes of income risk and information. Most relevant for this section, agents will pay more to live in a low risk regime when bankruptcy is prohibited, despite the reductions in variance that eliminating bankruptcy generates. One must therefore be careful to not interpret changes in consumption volatility as being indicative of changes in welfare.[31]

[31] This point is obviously well known when leisure is valued.

6. A Model of Default with Partial Information

As argued at the outset, a variety of empirical work, and the quantitative model of Athreya, Tam, and Young (2007) suggests that over the past several decades there has been a systematic decline in the asymmetry of information between borrowers and lenders. We now model an increase in income risk when it *coincides* with an improvement in the information lenders possess about borrowers. To accommodate asymmetric information, we use the signalling model developed in our previous paper. We assume that "earlier period" of the model corresponding to the early 1980's is one where information is asymmetric with respect to the stochastic components of household's income (including total income), current stigma cost, and current net worth. We maintain the assumption throughout that age and education are observable.[32] Therefore, we have $I = (y, j, m)$ (with (a, e, ν, λ) not observed).[33]

The first concern for solving the partial information economy is that lenders must form beliefs over the probability of an individual being in a particular state (e, ν, λ) given *whatever* is observed, knowing also that what is observed is a function of lenders' *a priori* beliefs; that is, beliefs must satisfy a fixed point condition. Let $\Pr(e, \nu, \lambda | b, y, j, m)$ denote the probability that an individual's shock vector in any period takes a given value (e, v), conditional on observing the size of borrowing, the permanent shock, age, and credit status. Given this assessment, the lender can compute the likelihood of default on a loan of size b:

$$\widehat{\pi}^b = \sum_e \sum_\nu \sum_\lambda \left[\sum_{e'} \sum_{\nu'} \sum_{\lambda'} \pi_e \left(e'|e\right) \pi_\nu \left(\nu'\right) \pi_\lambda \left(\lambda'|\lambda\right) d \left(b, e', \nu', \lambda'\right) \right] \Pr\left(e, \nu, \lambda | b, y, j, m\right). \tag{6.1}$$

In a stable environment with a small number of creditors, or one with an efficient technology for information sharing, intermediaries must form beliefs that incorporate everything they either know or can infer from observables; competitors who exploit this information may be able to 'cream-skim' the best borrowers away from those who form beliefs in any other way.[34] In equilibrium, if this information exists it must be incorporated by all intermediaries into their belief functions; we view this arrangement as a natural analogue to the conditions that prevailed in the early 1980s.

[32] Regulatory restrictions prohibit the use of age in determining credit terms, at least in the unsecured credit market, along with race and gender. We study the possibility that types are unobserved in a companion paper, which focuses on estimating the costs of such regulations.

[33] We separate b from I even though b is observable because the borrower takes the derivative of q with respect to b and it is therefore more convenient to make it a separate argument.

[34] This point is related to the extensive survival literature, which investigates whether agents who form beliefs that deviate from rational expectations can survive in asset markets.

$\Pr(e, \nu, \lambda | b, y, j, m)$ assigns a probability to each type that borrows b in equilibrium.

In the partial information environment the calculation of $\Pr(e, \nu, \lambda | b, y, j, m)$ is nontrivial, because it involves the distribution of endogenous variables. First, let the invariant distribution over states be denoted by $\Gamma(a, y, e, \nu, \lambda, j, m)$. In a stationary equilibrium the joint conditional probability density over shock pair (e, v, λ) must be given by

$$\Pr(e, \nu, \lambda | b, y, j, m) = \int_a \Gamma(a = f(b, y, e, \nu, \lambda, j, m), y, e, \nu, \lambda, j, m), \tag{6.2}$$

where f is the inverse of g with respect to the first argument wherever $\Gamma(a, y, e, \nu, \lambda, j, m) > 0$; that is,

$$a = f(b, y, e, \nu, \lambda, j, m)$$

and

$$b = g(a, y, e, \nu, \lambda, j, m).$$

Thus, the decision rule of the household under a given pricing scheme is inverted to infer the state conditional on borrowing. Using this function the intermediary then integrates over the stationary distribution of net worth, conditional on observables, and uses this probability to formulate beliefs.

It is possible that intermediaries in the partial information world would find it profitable to offer a menu of contracts and separate types (meaning agents with different realizations of the shocks (e, ν, λ)) in this manner. We restrict attention to the pure signalling model, which is not only tractable but also consistent with the relative homogeneity of unsecured loan contracts prior to 1990 (see Furletti 2003 or Edelberg 2006); permitting lenders to offer separating contracts would not be consistent with this observation.[35] However, as described in detail in Athreya, Tam, and Young (2007), a central complication arising under signalling with asymmetric information is the need for a rule specifying how lenders' assign beliefs about a household's state for values *not* observed in equilibrium; that is, how should lenders assign beliefs regarding repayment by households where $\Gamma(a, y, e, \nu, \lambda, j, m) = 0$? This issue matters because a household's decision on the equilibrium path depends on its understanding of lender behavior at *all* feasible points in the state space, including those that never arise. While zero profits must obtain on the equilibrium path, off-equilibrium pricing depends on off-equilibrium beliefs which are, in general, not easy to restrict as Bayes' rule

[35] Why the intermediaries did not use these contracts in the earlier period is a question beyond the scope of this paper.

does not provide guidance about how to set these beliefs. We therefore assign off-equilibrium beliefs to deal with the proliferation of equilibria typically present in signaling models, without ruling out plausible outcomes *a priori*.[36] The assignment of off-equilibrium beliefs will inform the algorithm we use to compute equilibria. Our algorithm is iterative – we guess pricing functions, compute implied default rates, recompute pricing functions based on the new default rates, and iterate to convergence. Athreya, Tam, and Young (2007) provides a detailed description of the algorithm and the main computational details for handling the asymmetric-information environment.

6.0.1. Equilibrium

We now formally define an equilibrium for the game between borrowers and lenders. We 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 $\mathcal{I} \subset \mathcal{Y} \times \mathcal{E} \times \mathcal{V} \times \mathcal{L} \times \mathcal{J} \times \{0,1\}$.

Definition 1. *A Perfect Bayesian Equilibrium for the model consists of (i) household strategies for borrowing $b^* : \Omega \rightarrow \mathcal{R}$ and default $d^* : \Omega \times \mathcal{E} \times \mathcal{V} \times \mathcal{L} \rightarrow \{0,1\}$ and intermediary strategies for lending $q^* : \mathcal{R} \times \mathcal{I} \rightarrow \left[0, \frac{1}{1+r}\right]$ and (ii) beliefs about the borrower state Ω given borrowing $\mu^* (\Omega|b)$, that satisfy*

1. ***Lenders optimize**: Given beliefs $\mu^* (\Omega|b)$, q^* is the pure-strategy Nash equilibrium under price competition.*

2. ***Households optimize**: Given prices $q^* (b, I)$, b^* solves the household problem.*

3. ***Beliefs are consistent with Bayes' rule**: The stationary joint density of Ω and b, $\Gamma_{\mu^*} (\Omega, b)$, that is induced by (i) lender beliefs and the resultant optimal pricing, (ii) household optimal borrowing strategies, and (iii) the exogenous process for earnings shocks and mortality, is such that the associated conditional distribution of Ω given b, denoted $\Gamma^b_{\mu^*} (b)$, is $\mu^* (\Omega|b)$.*

4. ***Off-Equilibrium Beliefs**: $q^* (b, I) = 0 \; \forall b$ such that $\Gamma^b_{\mu^*} (b) = 0$.*

[36] It turns out that modelling the game as signalling rather than screening is significantly easier. In a screening game the lenders would move first, and then we would need to check deviations in the infinite-dimensional space of alternative pricing functions. Here households move first and we only need to check deviations in the space of borrowing levels, which is implicit in our use of the pricing function as a schedule confronting the borrower. There is a connection between our procedure and the Intuitive Criterion from Cho and Kreps (1987). Sánchez (2007) employs a screening model.

Since our shocks are continuous random variables, the debt levels that get zero weight in the stationary distribution are those above and below any levels that get positive weight (Γ has a connected support). Obviously, for default decisions the upper limit is irrelevant; thus, as noted above, we are imposing a belief about the behavior of agents who borrow more than *any* agent would in equilibrium, no matter what their current unobserved state. Given that q is weakly-decreasing in b, the natural assumption is that this agent intends to default with probability one. Again, we direct the reader to Athreya, Tam, and Young (2007) for a complete discussion.

For the convenience of a reader who does not want to read two papers, Table (3) replicates the summary results of Athreya, Tam, and Young (2007) by comparing how selected aggregates change when the economy moves from a low-risk partial information setting to a high-risk full information one. Four key aggregates – discharge to income ratio, debt to income ratio, measure of borrowers, and measure of defaulters – show marked increases when the economy moves from a situation of asymmetric and partial information to one with full information. In addition, although not shown here, the experiment also generates growth in the dispersion in credit card loan rates rise, further evidence that changes related to the information lenders have about borrowers are critical to explaining the data. For a complete discussion of the underlying mechanisms we direct the reader to our earlier work.

6.1. Bankruptcy and Consumption Smoothing Under Partial Information

To isolate the role of better information, we first study the effect of improvements in information on consumption smoothing under a regime where income risk is held constant and default is permitted (we consider the high-risk setting).[37] The first result, seen in Figure (13), is that the variance of log consumption across age groups is *smaller* under PI than FI. Second, intertemporal smoothing is better; for college-educated households, there is a clear flattening of the profile of mean consumption by age as seen in Figure (14); Table (4) presents the average numbers for the convenience of the reader. That is, a (perhaps surprising) lesson is that improvements information will tend to *increase* the measured volatility of consumption in a lifecycle model.

Why does information work this way? Two factors are key: (1) a well-defined hump-shaped profile of mean income over the life-cycle, and (2) the presence of the bankruptcy/default option (which, to repeat, is the only reason information matters for credit supply at all). What Figure

[37] Again, the low-risk comparison is quite similar.

(8) shows is that under full information, the ability to default allows those with good future income prospects to borrow while limiting the ability of those who are doing badly (in the sense of having received a bad current shock) to borrow and smooth. The three panels of the pricing function make clear that those with the lowest income, conditional on age, must pay substantially more to borrow than their currently lucky counterparts. Therefore, a source of the welfare loss in this case relative to a world in which bankruptcy is prohibited is precisely that the latter allows all households – even the currently unlucky – to borrow at the risk-free rate.

The hump-shape of expected income over the lifecycle creates significant purely-intertemporal smoothing needs for households and so would generate low net worth positions early in life even in the absence of income risk. However, the presence of income risk has a significant impact of the consequences of debt accrued in early life on consumption smoothing later in life. More specifically, Figure (8) shows that under FI, households with income near and above the conditional mean for their age will be able to borrow at relatively low interest rates, while those who present more substantial default risk will not be able to do so. As a result, by mid-life many households have accrued debts and now must decide on whether to repay them. And most households *will*: *ex post* repayment rates will be high because under FI they would not otherwise have been able to borrow in the first place. As a result, consumption inequality will grow with age.

In contrast, the central finding of Atheya, Tam, and Young (2007) (which we do not detail here) is that partial information (PI) equilibria are capable of sustaining only very small amounts of unsecured borrowing. All households, even those with bright income prospects, are effectively rationed to those debt levels that carry little or no default risk. Thus, under PI, many young can be said to be meaningfully "credit constrained" since their default risk is overestimated due to pooling. In addition to these constraints directly precluding borrowing today, the inability to borrow in the future creates a stronger precautionary savings motive under PI than under FI. The combination of an inability to borrow and a desire to accumulate a buffer stock of assets – even when young – allows households later in life to effectively smooth consumption. The mechanisms at work here are the product of both credit supply and credit demand – an infinite-horizon setting would not generate any "intertemporally" motivated debts and so would not properly reflect the role that improved information, in conjunction with default, is likely to have played in altering consumption smoothing.

The intuition above is supported by Figure (13). Under FI, peak borrowing periods are early in life which then implies that by early middle age (approximately age thirty five) many households have accumulated significant debts – which, as argued earlier, most will choose to repay. Comparing this outcome to that obtaining under PI, we see that the *difference* in log consumption variance is highest at this age. Thus, the big difference created by improved information is that by early middle-age the household who could not borrow (PI) is *ex post* better able to smooth. This *ex post* advantage is transitory however, because once households have entered the age in which their primary motivation is to save for retirement differences in information on default risk are simply irrelevant. Moreover, since households save similarly after middle-age under both information regimes, consumption smoothing after middle age also becomes increasingly similar.

Turning to the transmission of increases in income risk to consumption, we find that the information setting is not critical – the changes in consumption risk are very similar if we assume that information is full or partial and held fixed. The counterpart to Figure (11) under partial information is very similar and is omitted. The implication of this result is that changes in consumption risk do not identify changes in the information set of lenders, given changes in income risk.

6.2. Changes in Income Risk and Information

We turn next to our preferred interpretation of events: that US households experienced an increase in income risk, while lenders experienced improvements in the ability to measure household default risk. Figure (15) shows the variance of log consumption over the life cycle under the two empirically-relevant settings – a low-risk, partial information setting (PI-L) versus a high-risk, full information setting (FI-H). We see from Figure (15) that the results are similar to the case where income risk was held at current levels and information improved. Viewed as a whole, our experiments have two implications: (i) changes in income risk do not play an important role in driving changes in default rates and bankruptcy (as found in Livshits, MacGee, and Tertilt 2006) and (ii) the effect of increased income risk does not depend the amount of information possessed by lenders. In other words, the increased credit access afforded by better information is not likely to have altered the transmission of income risk to consumption volatility.

6.3. Welfare

As seen in our experiments, measured consumption inequality rises with improvements in information; however, Table (5) shows that *ex ante* welfare does as well. The first section of Table (5) shows that eliminating bankruptcy under full information is desirable for all newborns; the gains are largest for college-educated workers since their ability to smooth intertemporally improves the most. Conversely, the smallest gains accrue to the non-high-school types because their intertemporal motives are weakest. Rising income risk harms all groups, with the largest costs borne by the non-high-school group. The losses are uniformly smaller if bankruptcy is permitted. Roughly speaking, distortions are larger in the presence of bankruptcy, so the marginal cost of additional income risk is smaller when bankruptcy is an option. The information assumption is not critical here; rising income risk always harms all households.

Turning to the welfare implications of better information, we see that improved information benefits all groups. The largest gains accrue to the college types again, for the same reason as above – their ability of borrow intertemporally is substantially improved under full information. Non-high-school workers don't really care which information regime they live in, since they can't borrow either way. When we combine the changes, college workers are better off under the high risk full information regime than under partial information with low risk, whereas high school workers are *worse off* (but still better off than if information had not improved); again, the non-high-school workers are essentially indifferent.

The result that some households are worse off has implications for the concerns that motivate regulations designed to ensure access to credit by blocking the use of some information.[38] There is a "second-best" aspect to the idea that one may want to protect some households from increases in income risk by restricting information; our model does not support this concern because the welfare losses are driven by the heightened income risk, not the change in information. In fact, ridding the economy of the bankruptcy option generates large gains in welfare even when income risk is rising; recent bankruptcy reform is a move towards the welfare-maximizing regime of $NBK - H$.[39]

7. Concluding Remarks

In this paper, we studied the extent to which unsecured credit and personal bankruptcy have altered the transmission of increased income risk to consumption variability over the past several decades. Our central result is that while plausible improvements in information that coincided with rising income risk greatly expanded the supply of credit, their effect on risk-sharing is not unambiguous. In particular, we find that these changes are likely to have improved intertemporal consumption smoothing for the lucky, but have made credit access harder for the unlucky, all else equal. As a

[38]The Equal Credit Opportunity Act in the U.S. is one such regulation.

[39]If bankruptcy is so damaging, why it exists at all is an open question. One study that attempts to integrate bankruptcy into an efficient risk sharing arrangement is Grochulski (2004).

result, if risk-sharing has improved in the face of rising income risk, as suggested by some recent work (Krueger and Perri 2006), it has occurred independently of these changes. Our approach should be viewed as a complement to complete-markets models, such as Krueger and Perri (2006) – since changes in credit markets are not likely to have improved risk-sharing, any such improvements must have come from other (more explicit) insurance arrangements.

Our model makes a host of predictions about the use of bankruptcy and changes in income risk. While we have calibrated the model using unconditional moments from the data, we have not confronted the model with the full range of conditional implications, particularly for the response of agents over the education and age distributions. One reason we eschew this comparison was mentioned in the Introduction: the literature does not agree on the temporal changes that have occurred in the consumption distribution. Until this debate is settled, comparisons of the model to the data are premature (since we might be comparing the model to the 'wrong' changes). Another reason is that we have not attempted a complete accounting of the dynamics of default; we intentionally abstract from many shocks that households experience in the data (especially shocks to household size and composition and medical and health care expenditures) in order to highlight the role that income risk plays in default.[40] We do not apologize for this decision; instead, we believe that our results generate a rich set of testable implications which we hope will prove useful for empirical analysis, provided the data is sufficiently good.

With respect to the role of changes in information, in the actual economy, not all improvements need be due to technological progress. In particular, many countries place a host of regulations that prevent *any* observed equilibrium from being a full information outcome; as we referenced above, the Equal Credit Opportunity Act (ECOA) bans lenders from conditioning terms on age, gender, and race.[41] Furthermore, the ECOA also bans even the *appearance* of such conditioning; correlations between proscribed characteristics and lending terms can trigger penalties even if the protected characteristics are not directly used (one such practice is redlining based on zip code data). Preliminary investigation of the costs of banning such information suggests that they may be quite large – in particular, pooling over age is very costly as it limits the ability of young college workers to borrow against their high expected future income. Athreya, Tam, and Young (2008)

[40] Our stigma shocks are designed to soak up some of the implications of these other shocks on average, but conditional moments are clearly not going to be captured. In fact, it is not clear that we have modelled the changes in the income distribution properly either; Jensen and Shore (2007) argue for changes in volatility that are concentrated mainly among the initially high variance groups.

[41] The Data Protection Directives in the E.U. and several acts in the U.K. (the Race Relations Act and Sex Discrimination Act) are similar regulations.

represents ongoing work aimed at a thorough analysis of the welfare and distributional consequences of such statutory limits on information use.

References

[1] Ábrahám, Árpád and Eva Cárceles-Poveda (2007), "Risk Sharing Under Limited Commitment," mimeo, University of Rochester and SUNY-Stony Brook.

[2] Acemoglu, Daron (2002), "Technical Change, Inequality, and the Labor Market," *Journal of Economic Literature* **40(1)**, pp. 7-72.

[3] Aiyagari, S. Rao (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics* **109(3)**, pp. 659-84.

[4] Aiyagari, S. Rao and Fernando Alvarez (2001), "Efficient Dynamic Monitoring of Unemployment Insurance Claims," mimeo, University of Chicago.

[5] Athreya, Kartik (2002), "Welfare Implications of the Bankruptcy Reform Act of 1999," *Journal of Monetary Economics* **49(8)**, pp. 1567-95.

[6] Athreya, Kartik (2008), "Default, Insurance, and Debt over the Life-Cycle," mimeo, Federal Reserve Bank of Richmond.

[7] Athreya, Kartik, Xuan S. Tam, and Eric R. Young (2007), "A Quantitative Model of Information and Unsecured Lending," mimeo, Federal Reserve Bank of Richmond and University of Virginia.

[8] Athreya, Kartik, Xuan S. Tam, and Eric R. Young (2008), "The Equal Credit Opportunity Act: A Welfare Analysis," unpublished notes, University of Virginia.

[9] Attanasio, Orazio (2002), "Consumption and Income Inequality: What We Know and What We Can Learn From It," Plenary Lecture, Society of Economic Dynamics Conference.

[10] Attanasio, Orazio, Erich Battistin, and Hidehiko Ichimura (2007), "What Really Happened to Consumption Inequality in the U.S.?" in Berndt, Ernst R. and Charles R. Hulten (eds.), *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, University of Chicago Press.

[11] Attanasio, Orazio, Gabriella Berloffa, Richard Blundell, and Ian Preston (2002), "From Earnings Inequality to Consumption Inequality," *Economic Journal* **112(478)**, pp. C52-9.

[12] Barron, John M. and Michael Staten (2003), "The Value of Comprehensive Credit Reports," in Miller, Margaret J. (ed.), *Credit Reporting Systems and the International Economy*, MIT Press, pp. 273-310.

[13] Battistin, Erich (2003), "Errors in Survey Reports of Consumption Expenditures," Working Paper **03/07**, Institute for Fiscal Studies.

[14] Bewley, Truman F. (1977), "The Permanent Income Hypothesis: A Theoretical Formulation," *Journal of Economic Theory* **16(2)**, pp. 252-92.

[15] Bird, Edward J., Paul A. Hagstrom, and Robert Wild (1997), "Credit Cards and the Poor," Institute for Research on Poverty Discussion Paper No. **1148-97**.

[16] Blundell, Richard, Luigi Pistaferri, and Ian Preston (2008), "Consumption Inequality and Partial Insurance," forthcoming, *American Economic Review.*

[17] Bowlus, Audra J. and Jean-Marc Robin (2004), "Twenty Years of Rising Inequality in US Lifetime Labor Values," *Review of Economic Studies* **71(3)**, pp. 709-43.

[18] Budría Rodríguez, Santiago, Javier Díaz-Giménez, Vincenzo Quadrini, and José-Víctor Ríos-Rull (2002), "Updated Facts on the U.S. Distributions of Earnings, Income, and Wealth," *Federal Reserve Bank of Minneapolis Quarterly Review* **26(3)**, pp. 2-35.

[19] Calem, Paul S., Michael B. Gorday, and Loretta J. Mester (2006), "Switching Costs and Adverse Selection in the Market for Credit Cards: New Evidence," *Journal of Banking and Finance* **30(6)**, pp. 1653-85.

[20] Chamberlain, Gary and Charles A. Wilson (2000), "Optimal Intertemporal Consumption Under Uncertainty," *Review of Economic Dynamics* **3(3)**, pp. 365-95.

[21] Chandler, Gary G. and Lee F. Parker (1989), "Predictive Value of Credit Reports," *Journal of Retail Banking* **11(4)**, pp. 47-54.

[22] Chatterjee, Satyajit, P. Dean Corbae, Makoto Nakajima, and José-Víctor Ríos-Rull (2007), "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," forthcoming *Econometrica.*

[23] Chatterjee, Satyajit, P. Dean Corbae, and José-Víctor Ríos-Rull (2006a), "A Finite-Life Private-Information Theory of Unsecured Debt," forthcoming *Journal of Economic Theory.*

[24] Chatterjee, Satyajit, P. Dean Corbae, and José-Víctor Ríos-Rull (2006b), "A Competitive Theory of Credit Scoring," mimeo, Federal Reserve Bank of Philadelphia, University of Texas-Austin, and University of Minnesota.

[25] Cho, In-Koo and David M. Kreps (1987), "Signalling Games and Stable Equilibria," *Quarterly Journal of Economics* **102**, pp. 179-221.

[26] Córdoba, Juan Carlos (2007), "U.S. Inequality: Debt-Constraints or Incomplete Asset Markets," forthcoming, *Journal of Monetary Economics.*

[27] Cunha, Flávio and James J. Heckman (2007), "The Evolution of Inequality, Heterogeneity, and Uncertainty in Labor Earnings in the U.S. Economy," NBER Working Paper **13526**.

[28] Davis, Steven J. (2004), "Comments on 'The Welfare Consequences of the Increase in Inequality in the United States'," *NBER Macroeconomics Annual 2004*, pp. 121-32.

[29] Deaton, Angus and Christina Paxson (1994), "Intertemporal Choice and Inequality," *Journal of Political Economy* **102(3)**, pp. 437-67.

[30] Dubey, Pradeep, John Geanakoplos, and Martin Shubik (2005), "Default and Punishment in General Equilibrium," *Econometrica* **73(1)**, pp. 1-37.

[31] Drozd, Lukasz A. and Jaromir B. Nosal (2007), "Competing for Customers: A Search Model of the Market for Unsecured Credit," mimeo, University of Wisconsin and University of Minnesota.

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

[33] Evans, David S. and Richard L. Schmalensee (1999), *Paying with Plastic: The Digitial Revolution in Buying and Borrowing*, MIT Press.

[34] Fay, Scott A., Erik Hurst, and Michelle J. White (1998), "The Bankruptcy Decision: Does Stigma Matter?" University of Michigan Working Paper **98-01**.

[35] Fisher, Jonathan D. and David S. Johnson (2006), "Consumption Mobility in the United States: Evidence from Two Panel Data Sets," *Topics in Economic Analysis and Policy* **6(1)**.

[36] Friedman, Milton (1957), *A Theory of the Consumption Function*, Princeton: National Bureau of Economic Research.

[37] Furletti, Mark (2003), "Credit Card Pricing Developments and Their Disclosure," Federal Reserve Bank of Philadelphia Payment Cards Center Discussion Paper **03-02**.

[38] Grochulski, Borys (2004), "Optimal Personal Bankruptcy Design: A Mirrlees Approach," mimeo, Federal Reserve Bank of Richmond.

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

[40] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2007), "Insurance and Opportunities: The Welfare Implications of Rising Wage Dispersion," mimeo, Georgetown University, University of Oslo, and New York University.

[41] Hellwig, Martin (1990), "Asymmetric Information, Financial Markets, and Financial Institutions: Where Are We Currently Going?" *European Economic Review* **33(2-3)**, pp. 277-85.

[42] Hornstein, Andreas, Per Krusell, and Giovanni L. Violante (2005), "The Effects of Technical Change on Labor Market Inequalities," in Aghion, Phillipe and Stephen Durlauf (eds.) *Handbook of Economic Growth*, Elsevier-North Holland.

[43] Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes (1994), "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving," *Carnegie-Rochester Conference Series on Public Policy* **40**, pp. 59-126.

[44] Huggett, Mark (1993), "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control* **17(5-6)**, pp. 953-69.

[45] Jensen, Shane T. and Stephen H. Shore (2007), "Changes in the Distribution of Income Volatility," mimeo, Wharton School and Johns Hopkins University.

[46] Kehoe, Timothy J. and David K. Levine, "Bankruptcy and Collateral in Debt Constrained Markets," forthcoming, *Macroeconomic Dynamics*.

[47] Krueger, Dirk and Fabrizio Perri (2004), "On the Welfare Consequences of the Increase in Inequality in the U.S." *NBER Macroeconomics Annual 2004*, pp. 83-121.

[48] Krueger, Dirk and Fabrizio Perri (2006), ”Does Income Inequality Lead to Consumption Inequality? Evidence and Theory,” *Review of Economic Studies* **73(1)**, pp. 163-93.

[49] Livshits, Igor, James MacGee, and Michèle Tertilt (2006), ”Accounting for the Rise in Consumer Bankruptcies,” mimeo, University of Western Ontario and Stanford University.

[50] Livshits, Igor, James MacGee, and Michèle Tertilt (2007a), ”Consumer Bankruptcy: A Fresh Start,” *American Economic Review* **97(1)**, pp. 402-18.

[51] Livshits, Igor, James MacGee, and Michèle Tertilt (2007b), ”Costly Contracts and Consumer Credit,” mimeo, University of Western Ontario and Stanford University.

[52] Matteos-Planas, Xavier (2007), ”A Model of Credit Limits and Bankruptcy with Applications to Welfare and Indebtedness,” mimeo, SUNY-Stony Brook.

[53] Matteos-Planas, Xavier and José-Víctor Ríos-Rull (2007), ”Credit Lines,” mimeo, SUNY-Stony Brook and University of Minnesota.

[54] Moffitt, Robert A. and Peter Gottschalk (2002), ”Trends in the Transitory Variance of Earnings in the United States,” *Economic Journal* **112(478)**, pp. C68-73.

[55] Mendoza, Enrique G., Vincenzo Quadrini, and José-Víctor Ríos-Rull (2007), ”Financial Integration, Financial Deepness, and Global Imbalances,” NBER Working Paper **12909**.

[56] Narajabad, Borghan (2007), ”Information Technology and the Rise of Household Bankruptcy,” mimeo, Rice University.

[57] Power, John and Garry Young (2007), ”Characteristics of the Insolvent and Those in Financial Distress,” mimeo, Bank of England.

[58] Primiceri, Giorgio E. and Thijs van Rens (2006), ”Predictable Life-Cycle Shocks, Income Risk, and Consumption Inequality,” mimeo, Northwestern University and Universitat Pompeu Fabra.

[59] Sánchez, Juan M. (2007), ”The Role of Information in Consumer Debt and Bankruptcy,” mimeo, University of Rochester.

[60] Schectman, Jack (1976), ”An Income Fluctuation Problem,” *Journal of Economic Theory* **12(2)**, pp. 218-41.

[61] Slesnick, Daniel T. (1992), "Aggregate Consumption and Saving in the Postwar United States," *Review of Economics and Statistics* **74(4)**, pp. 585-97.

[62] Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron (2004), "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy* **112(3)**, pp. 695-717.

[63] Sullivan, James X. (2008), "Borrowing During Unemployment: Unsecured Debt as a Safety Net," forthcoming, *Journal of Human Resources*.

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

[65] White, Michelle J. (2007), "Bankruptcy Reform and Credit Cards," *Journal of Economic Perspectives* **21(4)**, pp. 175-99.

[66] Wolff, Edward N. (2004), "Changes in Household Wealth in the 1980s and 1990s in the U.S." in Wolff, Edward N. (Ed.), *International Perspectives on Household Wealth*, Elgar Publishing Ltd.

[67] Young, Eric R. (2007), "Crowding Out or Crowding In? Unemployment Insurance and Private Insurance," mimeo, University of Virginia.

[68] Zame, William R. (1993), "Efficiency and the Role of Default When Security Markets are Incomplete," *American Economic Review* **83(5)**, pp. 1142-64.

[69] Zhang, Harold H. (1997), "Endogenous Borrowing Constraints with Incomplete Markets," *Journal of Finance* **52(5)**, pp. 2187-209.

8. Appendix – Models without Default

In this appendix we detail briefly Models **A**, **Z**, **KP**, and **MQR**. We abstract from many features that we use in the main model – overlapping generations, multiple sources of earnings risk, preference shocks, transactions costs – in order to cleanly represent the relationship between earnings risk and consumption inequality. In all cases the single source of uncertainty is a shock to earnings that follows an AR(1) process:

$$\log\left(e'\right) = \rho \log\left(e\right) + \nu'$$

where ν is white noise with variance σ_ν^2.

8.1. Model A

This subsection lays out the model based on Aiyagari (1994). The key features of this model are that agents can trade in only one uncontingent asset, that holdings of this asset are exogenously limited by a constant floor, and that default is prohibited. The household problem is then

$$v^A\left(a,e\right) = \max_{a' \geq \underline{a}} \left\{\log\left(c\right) + \beta E\left[v^A\left(a',e'\right)|e\right]\right\}$$

subject to the budget constraint

$$c + a' \leq \left(r + 1 - \delta\right) a + we.$$

In general equilibrium it must be the case that total asset holdings equals the aggregate capital stock and total labor input equals the sum over efficiency units of labor. Defining the stationary distribution of agents over assets and earnings realizations by $\Gamma\left(a,e\right)$, we require that (r,w) be such that

$$\begin{aligned} \int a\Gamma\left(a,e\right) &= K \\ \int e\Gamma\left(a,e\right) &= N. \end{aligned}$$

The goods market clearing condition then follows,

$$\int c\left(a,e\right)\Gamma\left(a,e\right) + \delta K = K^{\alpha} N^{1-\alpha}.$$

The parameter values are standard: $\beta = 0.99$, $\delta = 0.025$, $\alpha = 0.36$, $\underline{a} = 0$, and $\rho = 0.95$.

8.2. Model Z

This subsection lays out the model based on Zhang (1997). Relative to Model **A**, Model **Z** determines the borrowing limit $\underline{a}$ endogenously; this limit is specified as the value of assets that prevents default in every state of the world tomorrow. The household problem is

$$v^Z(a,e) = \max_{a' \geq a^d} \left\{ \log(c) + \beta E\left[v^Z(a',e') \,|e \right] \right\}$$

subject to the budget constraint

$$c + a' \leq (1 + r - \delta) a + we.$$

The level of assets that prevents default in the event that e^* is realized tomorrow, denoted $a^d(e^*)$, is the unique solution to

$$v^Z\left(a^d(e^*), e^*\right) = v^D(e^*),$$

where the value of default is specified to be the autarkic outcome

$$v^D(e) = \log(we) + \beta E\left[v^D(e') \,|e \right].$$

The borrowing limit is then specified to be the largest such value a^d (meaning the maximum amount of borrowing that prevents default in every state of the world tomorrow):

$$a^d = \max_{e^*} \left\{ a^d(e^*) \right\}.$$

In general equilibrium it must be the case that total asset holdings equals the aggregate capital stock and total labor input equals the sum over efficiency units of labor. Defining the stationary distribution of agents over assets and earnings realizations by $\Gamma(a,e)$, we require that (r,w) be such that

$$\begin{aligned} \int a\Gamma(a,e) &= K \\ \int e\Gamma(a,e) &= N. \end{aligned}$$

The goods market clearing condition then follows,

$$\int c(a,e)\,\Gamma(a,e) + \delta K = K^{\alpha} N^{1-\alpha}.$$

The parameter values are the same as in Model **A** with the obvious exception of $\underline{a}$.

8.3. Model KP

Here we lay out Model **KP**, which has a full set of contingent claims and endogenous debt limits (based on Krueger and Perri 2006). The household problem is then

$$v^{KP}(a,e) = \max_{a'(e') \geq \underline{a}(e')} \left\{ \log(c) + \beta E\left[v^{KP}\left(a'\left(e'\right), e'\right) | e \right] \right\}$$

subject to the budget constraint

$$c + \int a'\left(e'\right) f\left(e' | e\right) \leq (1 + r - \delta)\, a + we.$$

The (state-contingent) level of assets that prevents default in the event that e^* is realized tomorrow (the debt limit for that contingent claim) is the unique solution to

$$v^{KP}\left(\underline{a}(e^*), e^*\right) = v^{D}(e^*),$$

where the value of default is specified to be the autarkic outcome

$$v^{D}(e) = \log(we) + \beta E\left[v^{D}\left(e'\right) \right].$$

In general equilibrium it must be the case that total asset holdings equals the aggregate capital stock and total labor input equals the sum over efficiency units of labor. Defining the stationary distribution of agents over assets and earnings realizations by $\Gamma(a,e)$, we require that (r,w) be such that

$$\begin{aligned} \int a\Gamma(a,e) &= K \\ \int e\Gamma(a,e) &= N. \end{aligned}$$

The goods market clearing condition then follows,

$$\int c\left(a,e\right)\Gamma\left(a,e\right)+\delta K=K^{\alpha}N^{1-\alpha}.$$

The parameter values are the same as Model **Z**.

8.4. Model MQR

Here we lay out Model **MQR**, which has a full set of contingent claims but limits trade in these assets through two restrictions – a 'no diversion' constraint and a 'limited liability' constraint (see Mendoza, Quadrini, and Ríos-Rull 2007). The household problem is

$$v^{KP}\left(a,e\right)=\max_{k}\left\{\log\left(c\right)+\beta E\left[v^{KP}\left(a'\left(e'\right),e'\right)|e\right]\right\}$$

subject to the budget constraint

$$\begin{aligned}c+k &\leq a\\ a'\left(e'\right) &= \left(1+r-\delta\right)k+we'-\phi\int f\left(e'|e\right)we'\\ a'\left(e'\right) &\geq 0.\end{aligned}$$

The second equation is the incentive constraint – it imposes the condition that agents do not divert their resources (at a proportional cost of ϕ) and eat them. The third equation requires that end of period wealth must be nonnegative and embodies a form of limited liability.

In general equilibrium it must be the case that total asset holdings equals the aggregate capital stock and total labor input equals the sum over efficiency units of labor. Defining the stationary distribution of agents over assets and earnings realizations by $\Gamma\left(a,e\right)$, we require that (r,w) be such that

$$\begin{aligned}\int k\Gamma\left(a,e\right) &= K\\ \int e\Gamma\left(a,e\right) &= N.\end{aligned}$$

The goods market clearing condition then follows,

$$\int c\left(a,e\right)\Gamma\left(a,e\right)+\delta K=K^{\alpha}N^{1-\alpha}.$$

The parameter values are the same as Model **Z**, and we set $\phi = 0.4$ as an intermediate case. $\phi = 0$ eliminates trade in contingent claims completely, while $\phi \geq 1$ implements complete markets.

9. Two-Period Model of Default

In our two-period model of default, filing for bankruptcy in the second period entails a utility penalty (as in Zame 1993). The household problem can be written as

$$u = \max_{c_1, c_2(w_2), d(w_2) \in \{0,1\}} \left\{ \log(c_1) + \beta \int \left[(1 - d(w_2)) \log(c_2(w_2)) + d(w_2)(\log(w_2) - \lambda) \right] f(w_2 | w_1)\, dw_2' \right\}$$

subject to the period budget constraints

$$\begin{aligned} c_1 &\leq a + w_1 - a' \\ c_2(w_2) &\leq a'(1 + r) + w_2. \end{aligned}$$

For the case where default is valuable, we choose $\beta = 0.5$, $w_1 = 1$, w_2 is uniform over the two-point set $\{0.25, 1.75\}$, and $\lambda = 0.25$. $r = 0.1$ for the default-free environment and $r = 0.2$ when the agents can default (the risk-free rate adjusted for the fact that default occurs in 50 percent of the states in period 2). For the case where default is not preferred, we set $w_1 = 0.1$ and assume w_2 is uniform over the two-point set $\{1.25, 1.75\}$; now the equilibrium interest rate in the default environment is infinite. Clearly, these parameters are rigged to deliver the outcome we desire in each case. The maximization is computed using FFSQP from AEM Design (a feasible sequential quadratic programming approach with both active set and interior point methods to handle the constraints).

10. Computation

The computational method used to solve the no-default models is closely related to the method used in Young (2007). We use cubic splines to parameterize the value functions, which we solve using value iteration with Howard acceleration. We then iterate to find the invariant distributions using the lottery method from Young (2007). For Model **MQR** we solve the "modified model" (see the computational appendix in Mendoza, Quadrini, and Ríos-Rull 2007) that converts the contingent claims into a risk-free asset and a portfolio of pure insurance claims that can be solved out analytically. The resulting one-dimensional maximization problems are all solved using a combination of Brent's method and Newton-Raphson iterations, and the equilibrium interest rates are also located using Brent's method. For Model **KP** (and the related Model **C**) we use the FFSQP routine to do the high-dimensional maximization over contingent asset portfolios; due to zeros in the state transition equation driven by the highly-persistent nature of labor earnings shocks, we prohibit the purchase of any contingent claim for states that have a conditional probability of less than 10^{-8} (this ban keeps the constraint set numerically well-behaved).

The computational method for the main model is detailed in Athreya, Tam, and Young (2007); it is similar to the method used for the default-free models, but uses linear splines instead of cubic splines because the value function is not smooth everywhere. We use golden section search to compute the optimal asset position tomorrow, using a coarse grid search to avoid the local optima. Due to the significant burden this model poses, particularly the partial information setting, we implement it using OpenMP interfacing across 8 Pentium 4 processors. To calibrate the model we set up a system of five equations in five unknowns (the calibration targets plus the market clearing condition for capital); we then minimize the sum of squared deviations using a variety of methods (Nelder-Mead, Hooke-Jeeves, and quasi-Newton), giving higher weight to the equilibrium condition for the capital market than the calibration targets. We found that the Nelder-Mead method produced the most robust (albeit slow) convergence to the calibration targets. Overall, the calibration procedure took several days to converge. Each FI equilibrium took less than one day, as it involves only one nonlinear equation in one unknown (the capital market clearing condition); we solve this equation using Brent's method. The PI economy is extremely costly to compute; it took several days to compute an allocation for a given interest rate, and computing equilibria take weeks. Since the interest rate does not change much relative to FI (the presence of the special households ensures that the interest rate is almost pinned down by agents who don't care about default), we chose to fix r in the PI settings and save ourselves weeks of computation.

Table 1

Parameters/Calibration

Parameters				
$\sigma = 2.0000$	$\vartheta = 0.0300$	$\Delta = 0.0300$	$\lambda_2 = 0.0358$	$\Theta = 0.2000$
$\beta = 0.9823$	$\beta_s = 0.9948$	$\xi = 0.2589$	$\pi_{HH} = \pi_{LL} = 0.8668$	$\alpha = 0.3000$
$\phi = 0.0300$	$\mu_s = 0.0500$	$\lambda_1 = 1.6715$	$\theta = 0.3500$	$\delta = 0.1000$

Calibration	Model	Target
Discharge/Income Ratio	0.2830	0.5000
Fraction of Borrowers	0.1255	0.1250
Debt/GDP Ratio	0.0227	0.0068
Default Rate	1.4356	1.2000

Table 2

Bankruptcy and Consumption Smoothing

	FI-H			NBK-H		
	Coll	HS	NHS	Coll	HS	NHS
Intra	0.1291	0.1645	0.2109	0.1226	0.1664	0.2144
Inter	0.0238	0.0178	0.0246	0.0070	0.0082	0.0182
Total	0.1529	0.1823	0.2355	0.1296	0.1746	0.2326

Table 3

Unsecured Credit Market Aggregates

	FI-H	PI-L
Discharge/Income Ratio	0.2830	0.1080
Fraction of Borrowers	0.1255	0.0499
Debt/GDP Ratio	0.0227	0.0014
Default Rate	1.4356	0.0010

Table 4

Information and Consumption Smoothing

	FI-H			PI-H		
	Coll	HS	NHS	Coll	HS	NHS
Intra	0.1291	0.1645	0.2109	0.1064	0.1450	0.1910
Inter	0.0238	0.0178	0.0246	0.0314	0.0237	0.0216
Total	0.1529	0.1823	0.2355	0.1378	0.1687	0.2126

Table 5

Change in Welfare

	Coll	HS	NHS
$FI^L \rightarrow NBK^L$	2.15%	0.63%	0.60%
$FI^H \rightarrow NBK^H$	1.64%	0.25%	0.06%
$FI^L \rightarrow FI^H$	−0.20%	−0.33%	−0.48%
$NBK^L \rightarrow NBK^H$	−0.70%	−0.70%	−1.01%
$PI^L \rightarrow FI^L$	0.57%	0.06%	−0.14%
$PI^H \rightarrow FI^H$	0.74%	0.23%	0.09%
$PI^L \rightarrow PI^H$	−0.37%	−0.50%	−0.70%
$PI^L \rightarrow FI^H$	0.37%	−0.27%	0.06%
$PI^L \rightarrow NBK^H$	2.74%	0.68%	0.46%

Figure 1: Revolving Debt/Disposable Income

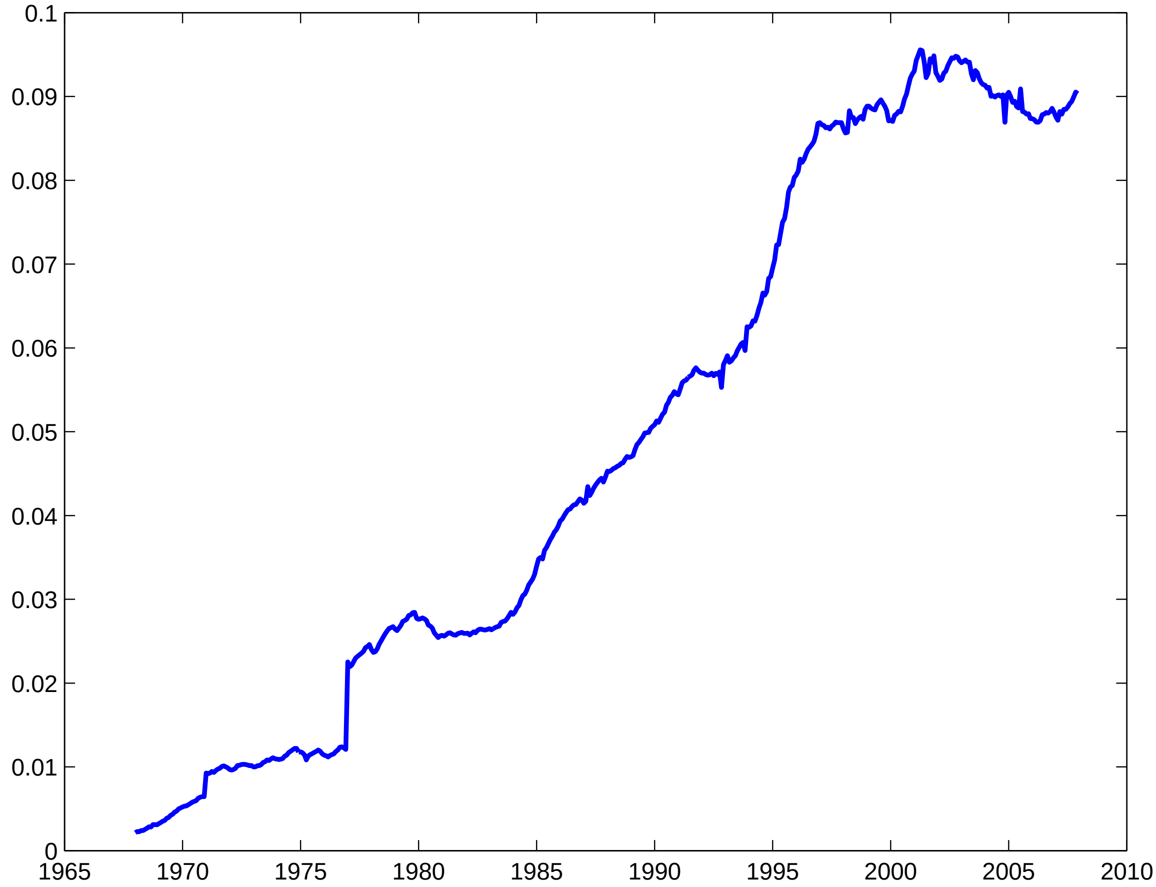

Figure 2: Chapter 7 Filings Per Capita

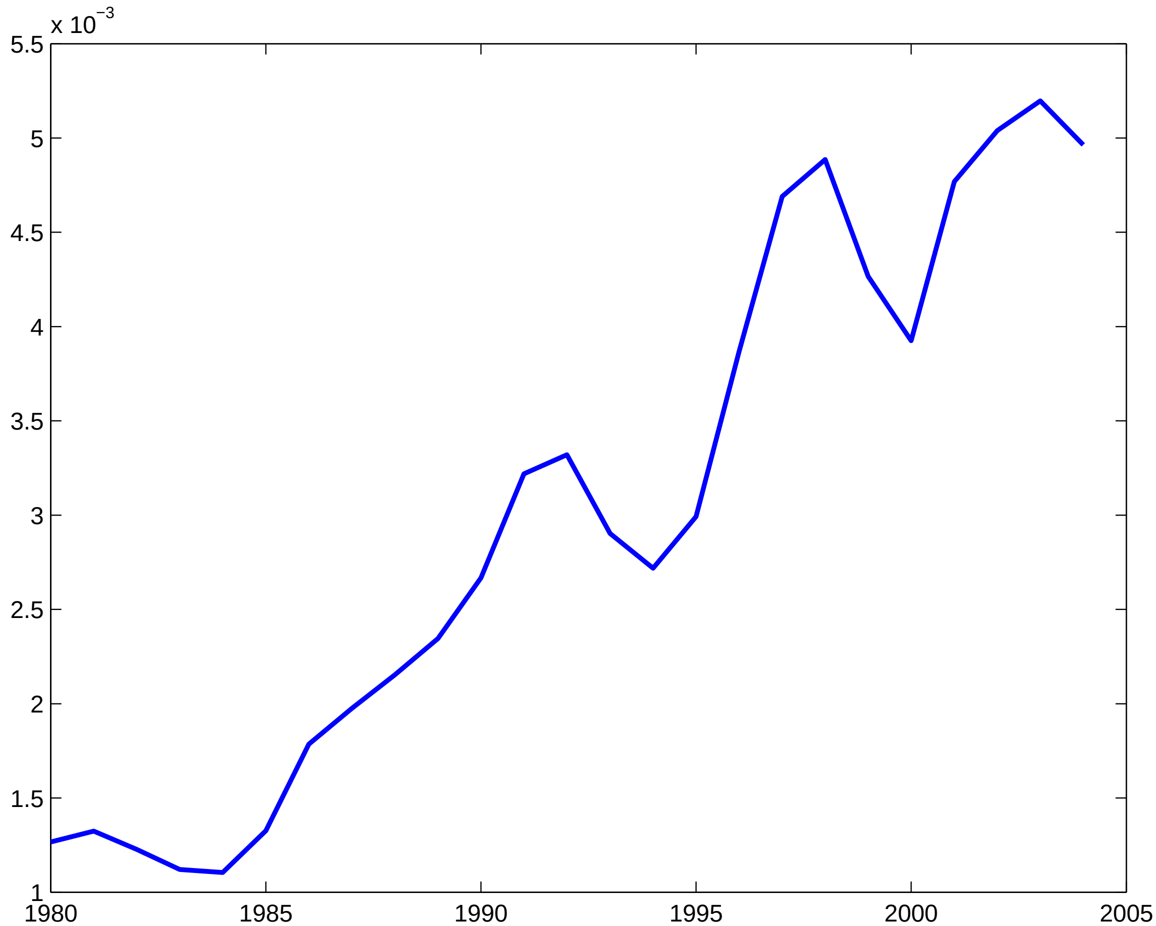

Figure 3: Income vs. Consumption Volatility

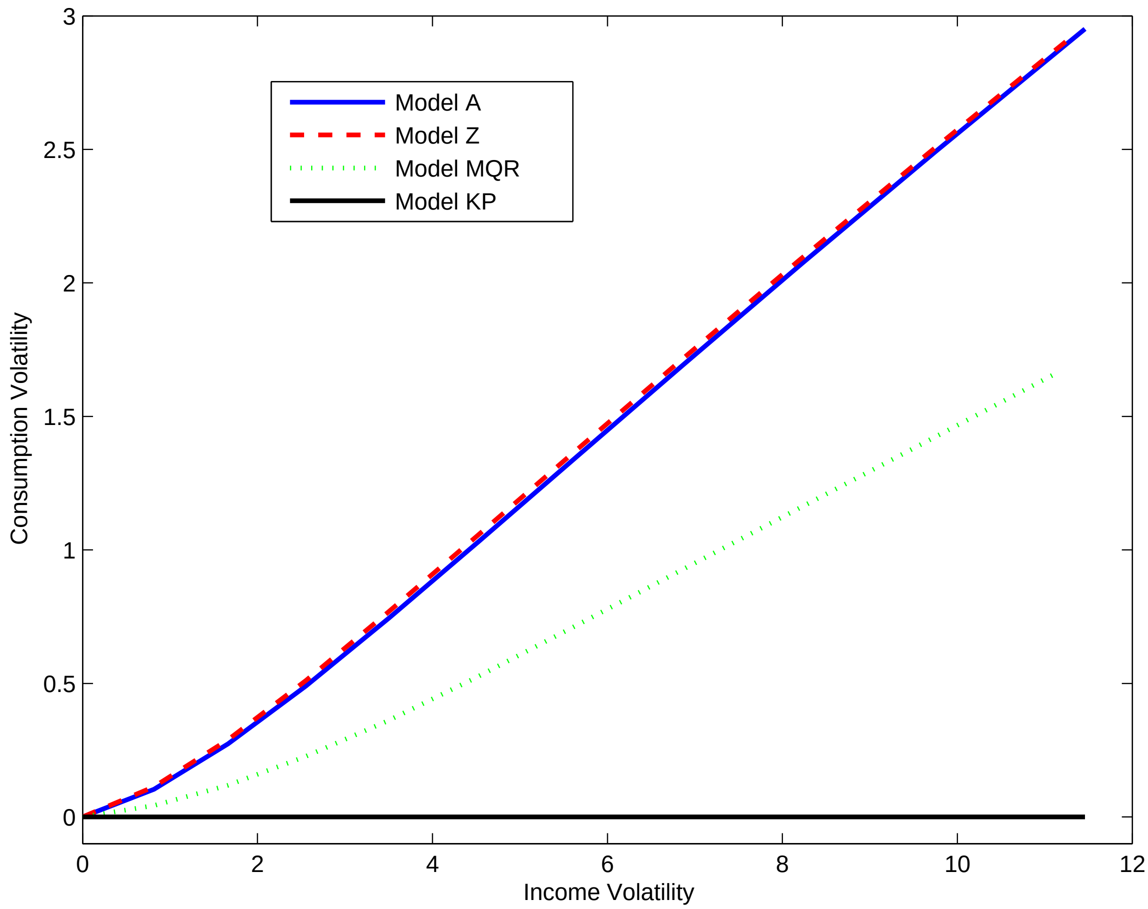

Figure 4: Bankruptcy Improves Smoothing

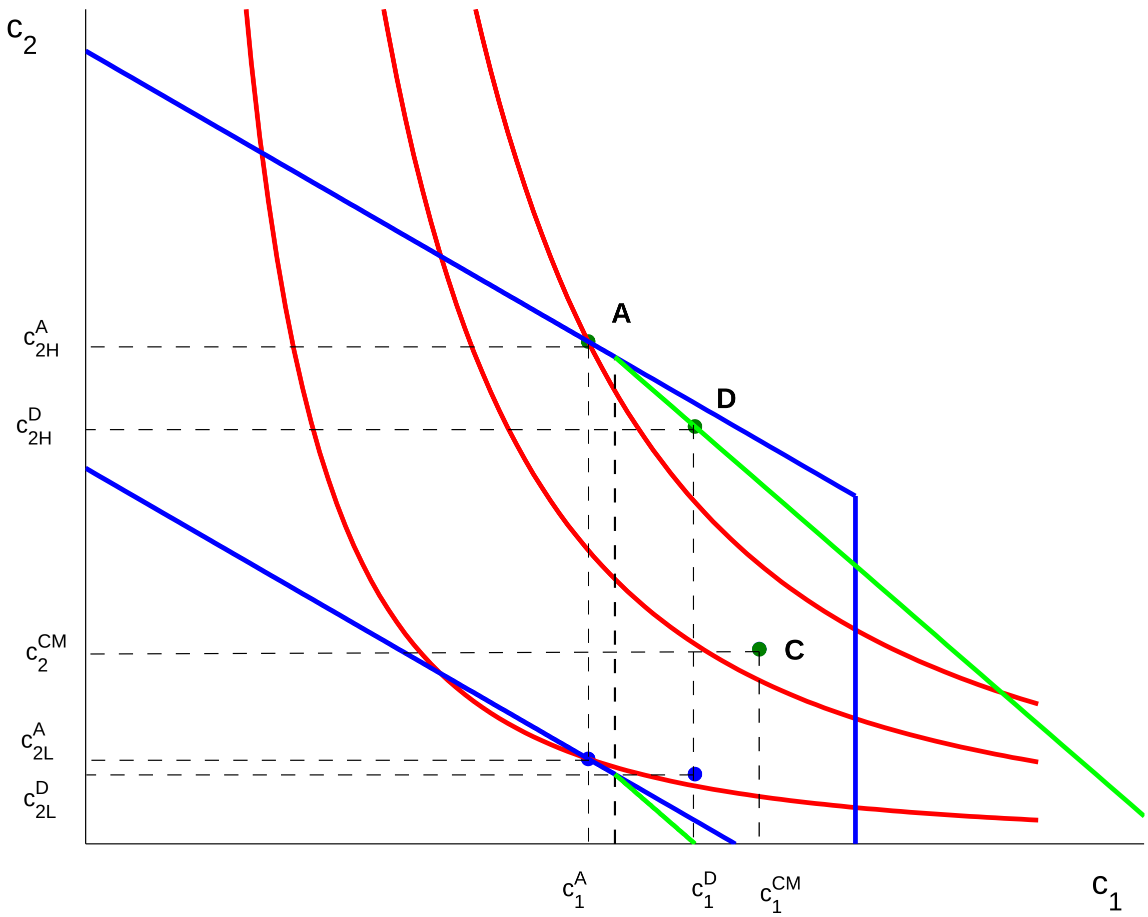

Figure 5: Bankruptcy Hinders Smoothing

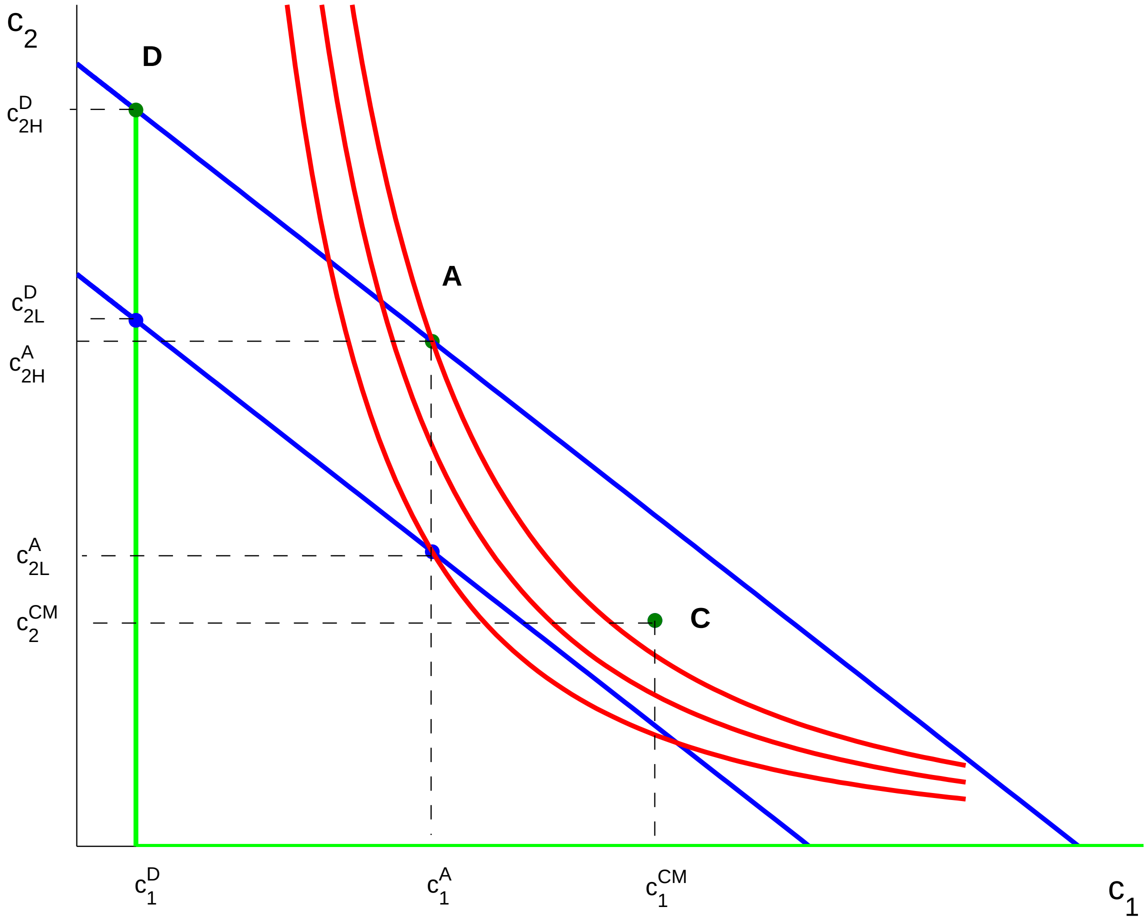

Figure 6: Efficiency Units

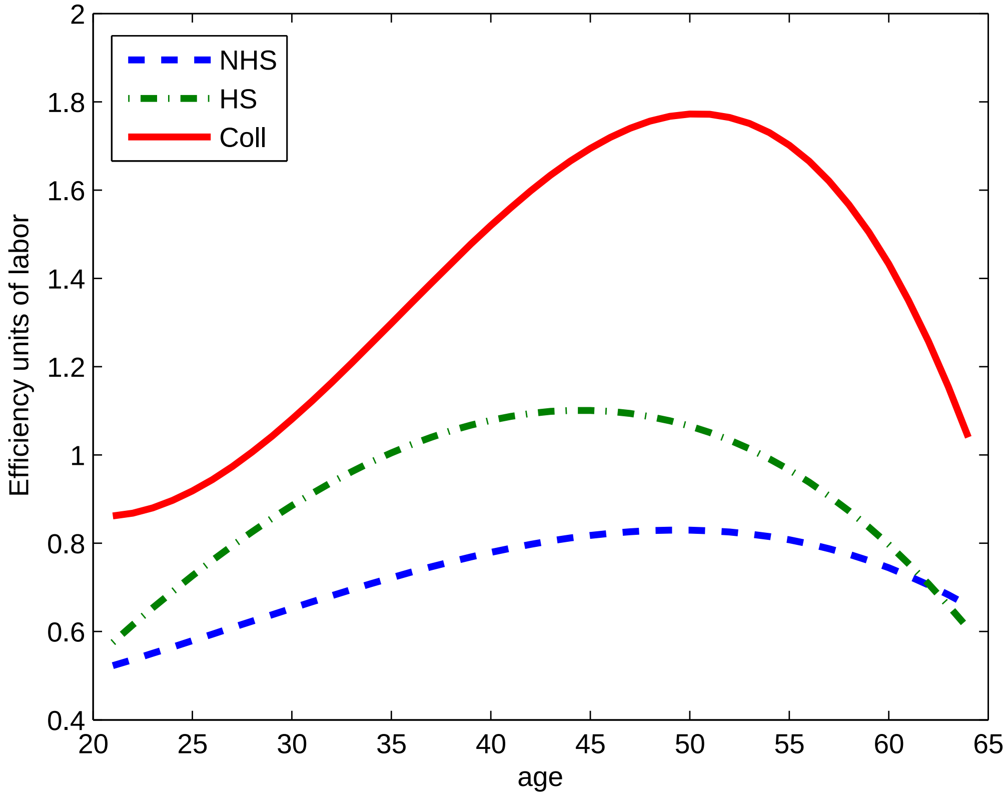

Figure 7: Effect of Bankruptcy Option, High Risk

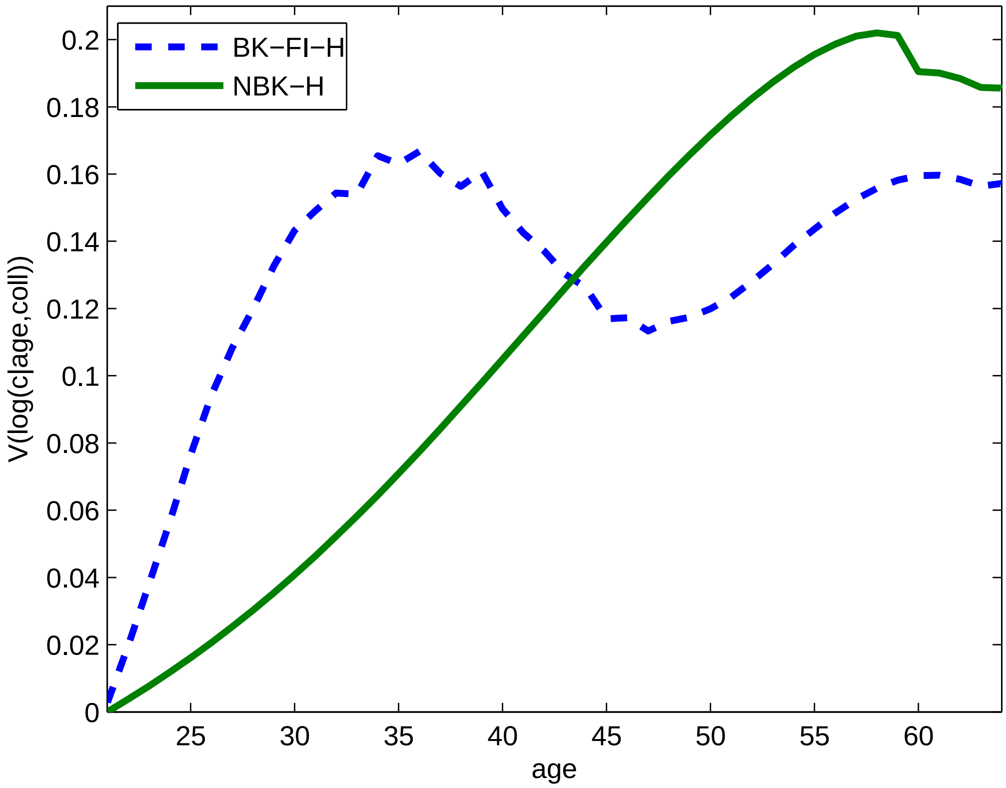

Figure 8: Pricing Functions

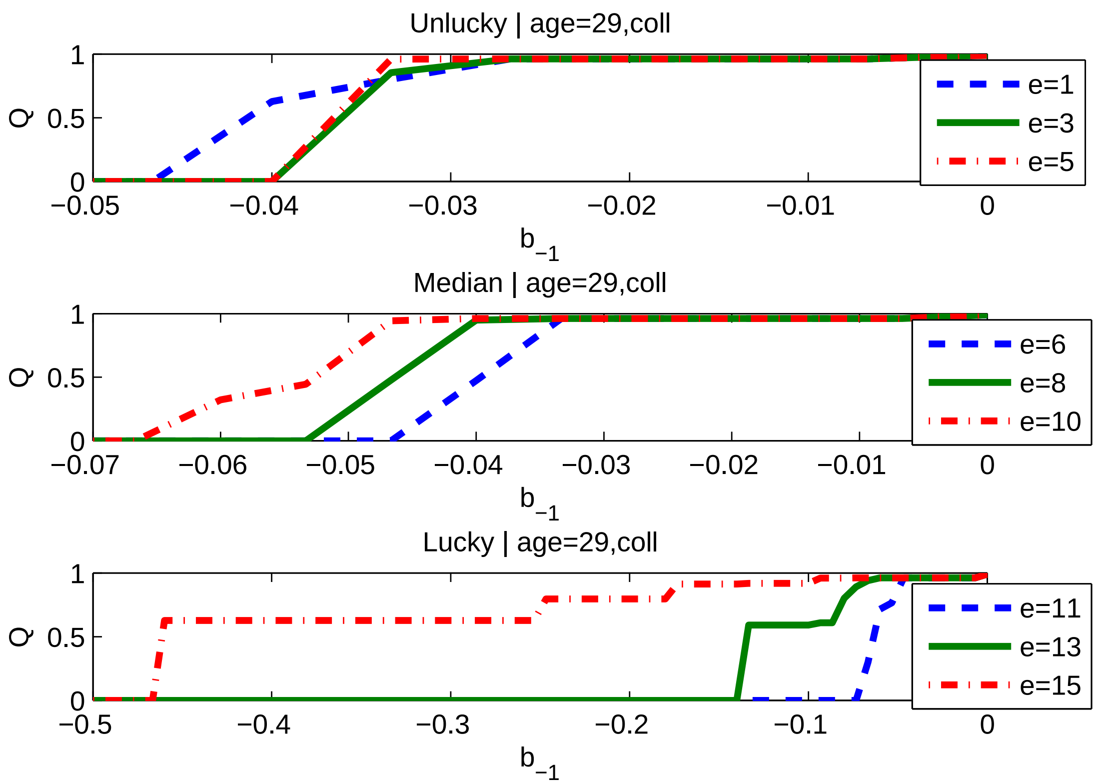

Figure 9: Default Over Lifecycle

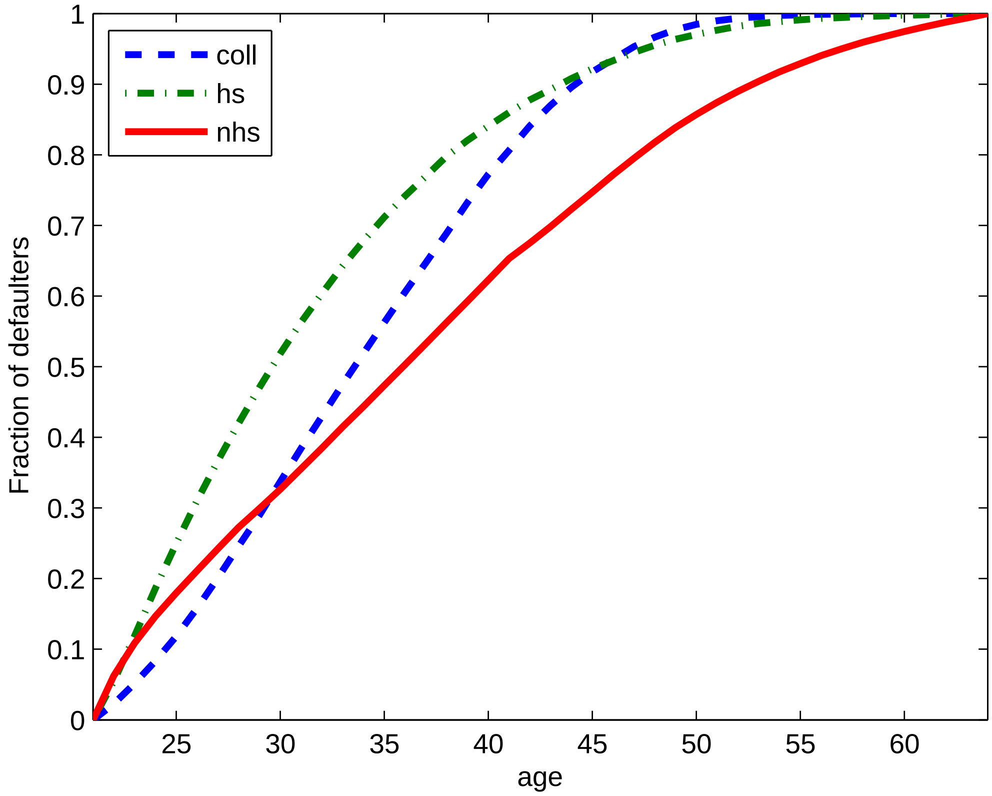

Figure 10: Increased Income Risk, No Bankruptcy

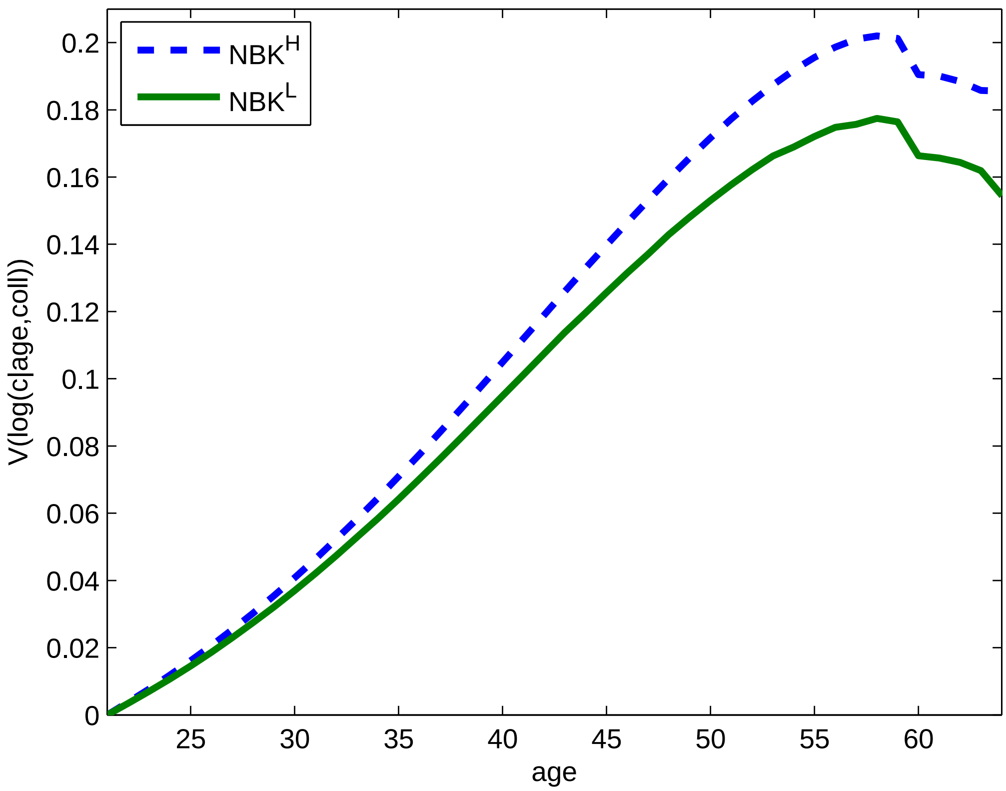

Figure 11: Increased Income Risk, Bankruptcy

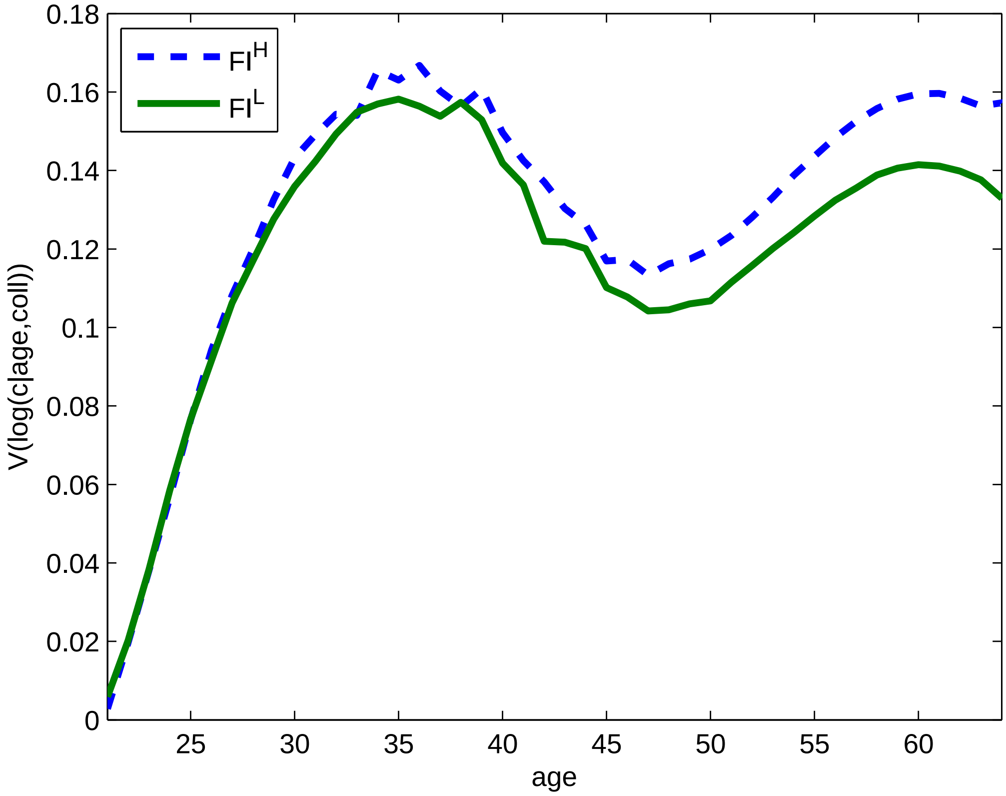

Figure 12: Δ Income Dispersion over Lifecycle

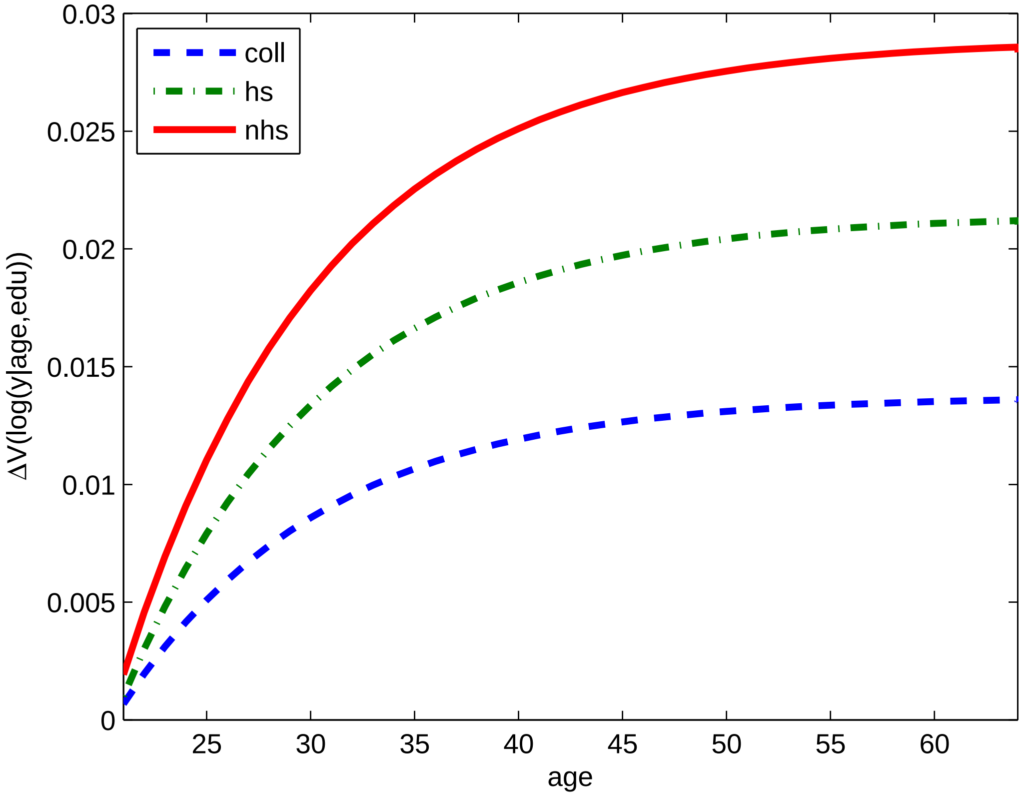

Figure 13: Role of Information for Consumption Volatility

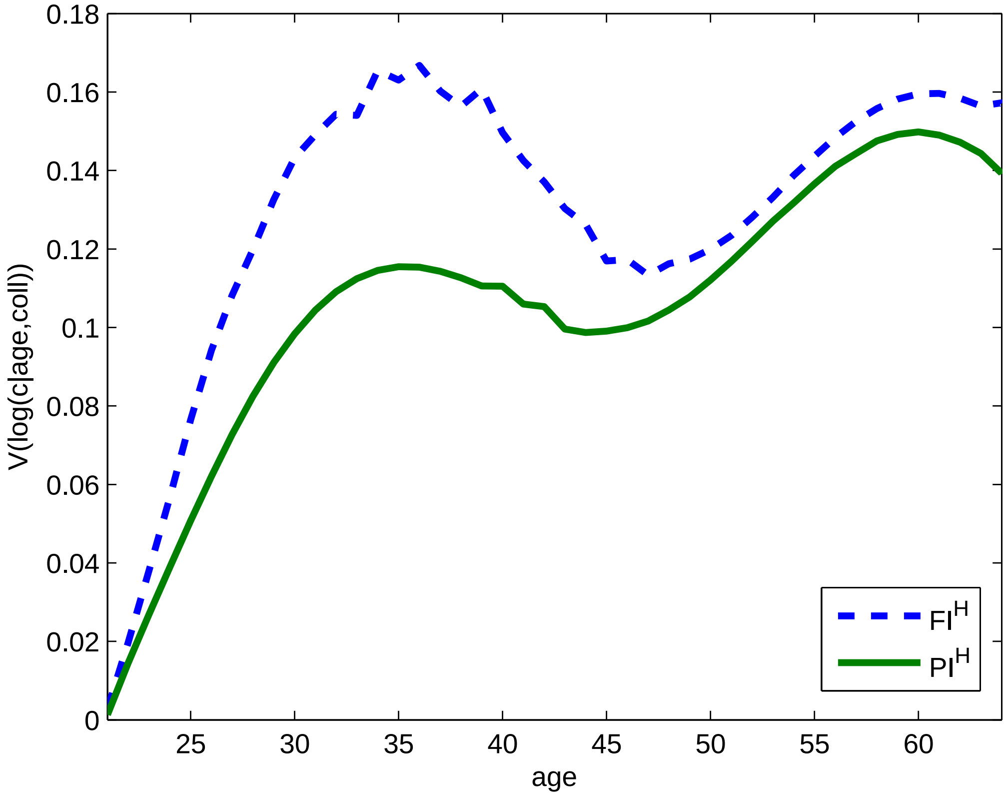

Figure 14: Role of Information for Mean Consumption

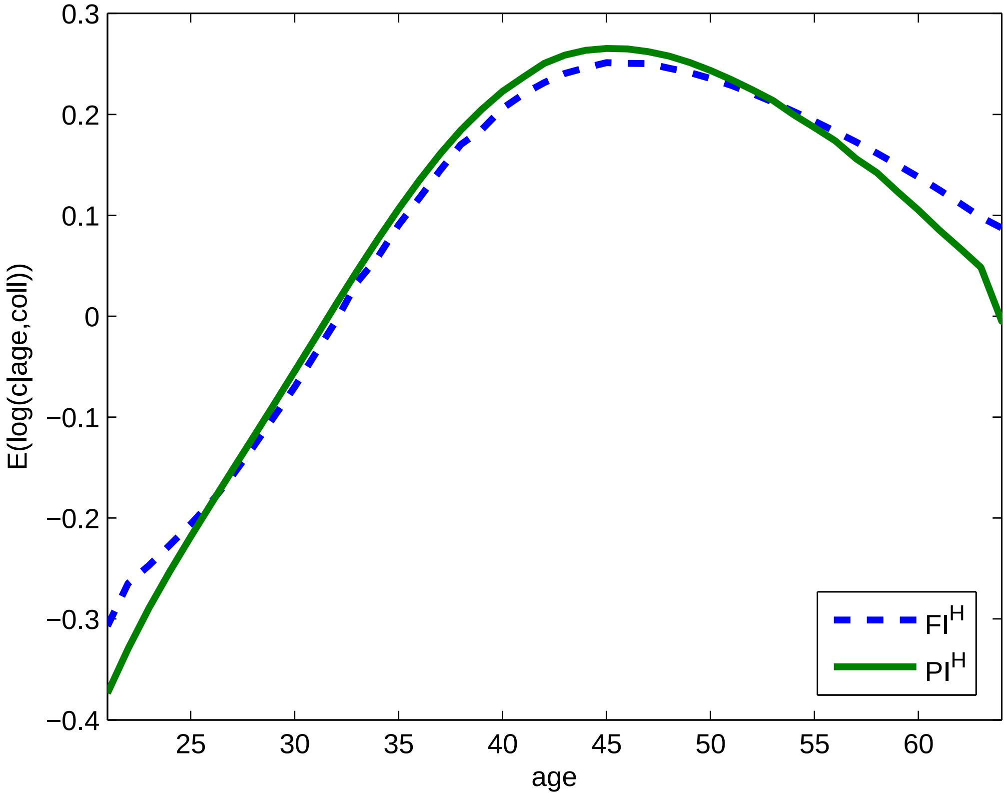

Figure 15: Consumption Volatility