Economic Quarterly—Volume 98, Number 4—Fourth Quarter 2012—Pages 255–307

Debt Default and the Insurance of Labor Income Risk

Kartik B. Athreya, Xuan S. Tam, and Eric R. Young

Recent research (e.g., Chatterjee et al. 2007, Livshits, MacGee, and Tertilt 2007) has found that allowing for debt default, such as through the relatively lenient U.S. bankruptcy code, is likely to improve *ex ante* welfare relative to more strict forms of debt forgiveness. The welfare gains come from improved consumption insurance provided by the option to not repay debt in some circumstances. Thus far, however, *all* instances where quantitative work finds a beneficial role for default have been ones with large and transitory shocks directly to household consumption *expenditures.* It is clear therefore that these "expense shocks" that lead to involuntary reductions in net worth are sufficient, given the specification of non-expense-related income risk in current models, to justify debt relief in forms resembling U.S. personal bankruptcy provisions.

The availability of bankruptcy, and more generally, default, will be reflected in the pricing on consumer debt, and so will affect households' ability to smooth consumption across dates and states of nature. It is therefore important to note that a significant amount of the risk to lifetime household resources may come from persistent shocks to *labor income* (Huggett, Ventura, and Yaron 2010). As a result, to the extent

Athreya is an economist at the Richmond Fed; Tam is affiliated with the University of Cambridge; Young is an economist at the University of Virginia. This article previously circulated under the title "Are Harsh Punishments for Default Really Better?" We would like to thank seminar and conference participants at UT-Austin, the Board of Governors, the Cleveland Fed, Georgetown University, the Philadelphia Fed, and Queen's University for comments on the earlier versions. We thank the EQ committee, especially Huberto Ennis, for detailed comments. Tam thanks the John Olin Foundation for financial support. The opinions expressed here do not reflect those of the Federal Reserve System or the Federal Reserve Bank of Richmond. All errors are the responsibility of the authors. E-mail: kartik.athreya@rich.frb.org.

that one might be able to locate other, more targeted, ways of insuring expense shocks, it is useful to better understand how effective debt forgiveness is for managing income risk in isolation.

In this article, we evaluate in detail the role of debt forgiveness in altering the impact of income risk in the absence of expense shocks. The experiments we present can be thought of as asking: "If we insure the out-of-pocket expenses that constitute expenditure shocks, is there still a role of debt relief as a form of insurance against 'pure labor income risk'?" We address this question by studying a range of specifications for households' attitudes toward the intra- and intertemporal properties of income, when expense shocks are not present. Our main finding is that, absent expenditure shocks, the ability to default very generally hinders the ability of households to protect themselves against labor income risk.

Despite the nature of our results, we stress that our work is not to be taken as a strong statement about the overall desirability of U.S. personal bankruptcy law, for two reasons. First, to the extent the expense shocks are a feature of reality, our model is missing a feature known to be capable of justifying bankruptcy protection. Second, *informal* default or "delinquency" whereby a borrower simply ceases making payments (and leaves themselves open to legally protected collections efforts) may simply increase if formal bankruptcy is made stricter or disallowed altogether. Indeed, in ongoing work (Athreya et al. 2013), we find that this channel is quantitatively relevant. These related, and coexisting, options to avoid debt repayment are not modeled here. Instead, our results apply more narrowly: They suggest that labor income risk alone may not provide a strong rationale for allowing households to default. In other words, our findings suggest that the scope of shocks that debt forgiveness is providing insurance against is limited, perhaps limited principally to relatively catastrophic outcomes.

It is interesting to note that similar results are now being located in the literature on sovereign debt. Namely, it has proved very difficult to find plausible circumstances in which the benefits to being able to repudiate debts (or perhaps more accurately, the costs of being unable to commit to repayment of sovereign debt) are positive. The reasons for the similarity of the results are natural. Most importantly, the models themselves are largely isomorphic in the optimization problems they lead to, and do not differ substantially enough in their quantitative specification of either preferences or risk. Moreover, even though sovereign debt models differ somewhat in the interpretation of the debt itself (i.e., that is public debt, not private), the standard assumption in that literature is that government is benevolent and seeks to borrow on behalf of households who themselves wish to smooth consumption.

This blurs the distinction between the path of public debt the government chose and that which households would have chosen.

Our results come from comparing allocations arising from two underlying trading environments. First, we study allocations arising from what we will refer to as the textbook, or "standard model" (SM), of consumption and saving in which households face uninsurable earnings risks with persistent and transitory components. In this model, households can only borrow using nondefaultable debt and also face liquidity constraints. Canonical examples of SM include those laid out in Deaton (1992, chapter 7) and Carroll (1997). To be consistent with the view that borrowing limits should be endogenously determined by repayment incentives, under SM, we investigate primarily the so-called "natural borrowing limit" case.[1]

The second trading arrangement we consider is one where, as before, households face life-cycle consumption/savings problems in which they encounter identical risks as in SM, but can issue defaultable debt. We will refer to this as the "default model" (DM). Benchmarks in this literature are Chatterjee et al. (2007) and Livshits, MacGee, and Tertilt (2007). Following these articles, default in the DM will be represented as a procedure whereby those with negative net worth can stop paying obligations, subject to any costs that may be present. The two trading arrangements we consider are thus clearly different. Nonetheless, they are related in a simple way: SM is the limiting case of DM as default becomes prohibitively costly.

To focus directly on the role of default in insuring labor income risk relative to the SM, we take two steps. First, as already noted, we deliberately set aside expenditure shocks. The presence of such shocks rules out the comparison of models with default against the standard model as budget sets would be empty for some dates and states were it not for the possibility of default. Second, we will examine a wider array of household preferences than has been done in the literature thus far. Specifically, we (i) separate risk aversion from the intertemporal willingness of households to substitute consumption, and (ii) evaluate the role of ambiguity aversion (or uncertainty aversion) when households are unsure of the stochastic environment they populate.

Both the separation of risk aversion from intertemporal elasticities and the possibility of ambiguity have been previously identified with a beneficial role for debt default. However, neither has been studied formally. The logic for suspecting that they may be important in delivering a welfare-enhancing role for default is as follows.

[1] See, e.g., Ljungqvist and Sargent (2004, p. 577).

First, the tradeoff between intertemporal and intratemporal smoothing was first suggested in Livshits, MacGee, and Tertilt (2007) in a life-cycle model of personal default. Assessing the relative importance of these motives therefore requires allowing for preferences in which the two attitudes can be distinct, irrespective of the uncertainty surrounding income. However, prior work has employed constant relative risk aversion (CRRA) preferences that conflate the two aspects of household preferences. In contrast, we employ Epstein-Zin recursive utility (Epstein and Zin 1989), which we select because of its tractability and demonstrated ability to improve the performance of asset pricing models, of which defaultable debt is a special case.

Second, with respect to the role of ambiguity in determining the value of an option to default, the legal and political history of bankruptcy law suggests that allowing for the release of debtors subject only to modest penalties is a policy that improves welfare if households are not perfectly sure of the probabilistic structure of income risk (see Jackson [2001] for one example).[2] This view is not confined to legal experts. As noted as early as Friedman (1957), agents will typically be unsure about the process that generates their labor income shocks, instead accepting that a family of potential distributions that may be difficult to distinguish are possible. Within this class of preferences, an agent who displays ambiguity aversion (Epstein and Schneider 2003) will solve a max-min problem—the agent will choose the member of the class that makes utility lowest and then choose consumption and savings in order to deliver the highest utility in this worst case.[3] It is precisely this feature of the problem that will allow for a more nuanced understanding of how penalties can be "excessive" and thereby welfare-reducing: Eliminating default through harsh penalties may leave the agent unwilling to borrow at all. As a result, such a policy could perversely inhibit both intertemporal and intratemporal consumption smoothing, despite "mechanically" alleviating the limited commitment problem that the young and poor face. U.S. bankruptcy law, for instance, appears directly predicated on the idea that penalties can indeed be excessive, in the sense that they may leave would-be borrowers unwilling to do so (see Jackson [2001]).

The potential role for ambiguity in altering the welfare implications of having defaultable debt is also suggested by the observation that, in

[2] Miao and Wang (2009) study the decision to exercise an option under ambiguity. Due to the presence of fixed costs, bankruptcy has option value. We focus on a related setting but are interested in the quantitative aspects associated with household consumption smoothing.

[3] These preferences are a special case of the more general ambiguity-averse preferences axiomatized by Klibanoff, Marinacci, and Mukerji (2009).

all extant work on consumer default, the relative gains seen in the SM relative to DM strongly depend on the "worst case" for household income. In particular, the large welfare losses in the DM relative to SM stem from the ability of young agents to borrow out to the natural debt limit. The natural debt limit is, however, extremely sensitive to small changes in the value of the worst-possible labor income realization, particularly for (i) young agents for whom the annuity value of future labor income is particularly high, and (ii) all agents when the risk-free borrowing rate is low.[4] This lower bound is difficult to estimate accurately (see Deaton [1992] or Pemberton [1998]) and the worst-case outcomes are the primary focus of ambiguity-averse agents; thus, it seems important to understand whether the superiority of SM hinges entirely on the lowest value of income.

Our main finding along these dimensions is that even in the presence of very high levels of uninsurable labor income risk, high risk aversion, an unwillingness to substitute intertemporally, and the presence of ambiguity, the ability of households to default on debt leads to allocations that all households prefer less than the outcome that arises when they retain full commitment to repay. The intuition for our welfare results involves the relationship between the current economic situation of the borrower and the price of debt. When short-term debt is used in a setting with household labor income risk that is persistent, limited commitment to debt repayment will make credit expensive anytime the household experiences a negative shock; pricing "moves against" the unlucky borrower. (In Athreya, Tam, and Young [2009], we argue that unsecured credit markets are not insurance markets for precisely this reason.) As a result, agents who most "need" debt to smooth consumption are exactly those that find themselves unable to obtain it, because they also pose the highest risk of default. Tam (2009) extends this result to longer-term arrangements; specifically, he finds that competitively priced longer-period debt (in which the pricing function is held fixed over a number of periods) is welfare-dominated by one-period debt.

In contrast, the possibility of welfare gains from lowering penalties by enough to yield default in equilibrium was first suggested by Dubey,

---

[4] Denoting by $y_{\min} > 0$ the lowest realization of potential labor income and r the risk-free interest rate on debt, the natural borrowing limit for an infinitely lived agent is given by $\underline{b}_{nat} \equiv -\frac{y_{\min}}{r}$, a function that asymptotes to $-\infty$ as interest rates go to zero. Assuming a credit card interest rate of 14 percent (the modal interest rate in Survey of Consumer Finances data in 1983 adjusted for a measure of realized inflation), the natural debt limit moves roughly seven times as much as the minimum income level. For good borrowers, for whom interest rate discounts have recently appeared (Furletti 2003; Livshits, MacGee, and Tertilt 2008), the natural debt limit will be even more sensitive.

Geanakoplos, and Shubik (2005). Theirs was a setting where borrowers of differential default risk were pooled together and thereby *did not* pay the individually actuarially fair price for their debt issuance. As a result of the stylized nature of their two-period model, it is not suitable for determining whether defaultable debt is welfare-improving in a more quantitatively oriented model economy. In some quantitative settings where pooling is imposed exogenously, Athreya (2002) and Mateos-Planas and Seccia (2006) find that welfare is higher in SM than DM. More recently, in a setting where private information allows for equilibrium pooling, the findings of Athreya, Tam, and Young (2009) suggest again that, as a quantitative matter, short-term defaultable debt is unlikely to be able to function as a form of insurance. Viewing these findings as a whole, they support the notion that the benefits of slacker borrowing constraints outweigh the costs of having no default option.

Lastly, with respect to political support for a policy allowing debt default, in addition to the welfare gains from having defaultable debt available in the presence of expense shocks, it seems possible that such provisions would enjoy support even in their absence. One obvious possibility is that the current regime may simply reflect objectives other than the maximization of the welfare of newborn agents. We therefore ask if *ex post* welfare can account for the evident political support enjoyed by proponents of relatively lax rules on default. Specifically, we ask whether model agents would choose to allow the option to default on debt in an economy where it was not already present (taking into account all changes resulting from the policy change). We find some support for such a change, but it falls well short of a majority. Support for the default option comes from relatively unlucky middle-aged college graduates: These are agents who borrowed a lot when young, in (rational) anticipation of higher income in middle age. When realized income did not materialize as expected, such households have significant debt as they approach retirement, and so will benefit from having debt obligations removed. Young agents, by contrast, are almost uniformly opposed to allowing defaultable debt, and even less-educated workers do not generally support it.

1. MODEL

Households in the model economy live for a maximum of $J < \infty$ periods. We assume that the economy is small and open, so that the risk-free interest rate is exogenous, while the wage rate is still determined by a factor price condition.[5] As a result, our welfare calculations will be biased toward finding a positive role for bankruptcy, since any lost resources arising from the implementation of default procedures like bankruptcy courts and legal costs will be ignored.

Households

Each household of age j has a probability $\psi_j < 1$ of surviving to age $j+1$ and has a pure time discount factor $\beta < 1$. Households value consumption per household member $\frac{c_j}{n_j}$ and attach a negative value $\lambda_{j,y}$ (in terms of a percentage of consumption) to all nonpecuniary costs of defaulting, which depend on type y to be defined below. Their preferences are represented by a recursive utility function $U\left(\left\{\frac{c_j}{n_j}\right\}_{j=1}^{J}\right)$ that we detail below. Households retire exogenously at age $j^* < J$.

We follow Chatterjee et al. (2007) in allowing for household-level costs from default that are primarily nonpecuniary in nature. The existence of nonpecuniary costs of default are also suggested by the calculations and evidence in Fay, Hurst, and White (1998) and Gross and Souleles (2002), respectively. The former article shows that a large measure of households would have "financially benefited" from debt default via personal bankruptcy but did not file for protection, while both articles document significant unexplained variability in the probability of default across households even after controlling for a large number of observables. These results suggest the presence of implicit unobserved collateral that is heterogeneous across households, including (but not limited to) any "stigma" associated with default along with any other costs that are not explicitly pecuniary in nature (as in Athreya [2004]). We will therefore sometimes refer to $\lambda_{j,y}$ as stigma in what follows, although we intend it to be more encompassing.

The household budget constraint during working age is given by

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

where q is an individual-specific bond price that depends on bond issuance b_j and a vector of individual characteristics I. Net worth *after* the current-period default decision is denoted a_j, and therefore satisfies $a_j = b_{j-1}$ if the household does not default and $a_j = 0$ otherwise; Δ is the pecuniary cost of filing for default. The last term is after-tax

[5] In our previous work we introduce a class of "special" agents who hold large amounts of capital for the purpose of endogenously obtaining a low, risk-free rate in the presence of low asset holdings for the median agent. Here we ignore the general equilibrium determination of returns and thus drop the special households from the model because their presence is irrelevant to the question at hand.

current labor income (τ is the tax rate). 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):

$$\log\left(e'\right) = \varsigma \log\left(e\right) + \epsilon', \tag{2}$$

and a purely transitory shock $\log(\nu)$. Both e and $\log(\nu)$ are independent mean zero normal random variables with variances that are y-dependent.[6] The budget constraint during retirement is

$$c_j + q\left(b_j, I\right) b_j \leq a_j + \Delta \mathbf{1}\left(d_j = 1\right) + \upsilon W \omega_{j^*-1,y} y e_{j^*-1} \nu_{j^*-1} + \Upsilon W, \tag{3}$$

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

The survival probabilities $\psi_{j,y}$ and the deterministic age-income terms $\omega_{j,y}$ differ according to the realization of the permanent shock. We interpret y as differentiating between non-high school, high school, and college education levels, as in Hubbard, Skinner, and Zeldes (1994), and the differences in these life-cycle parameters will generate different incentives to borrow across types. In particular, college workers will have higher survival rates and a steeper hump in earnings; the second is critically important as it generates a strong desire to borrow early in the life cycle. Less importantly, they also face slightly smaller shocks than the other two education groups. The life-cycle aspect of our model is key—in the data, defaults are skewed toward young households (who borrow at least in part for purely intertemporal reasons), particularly those who do not report medical expenses as a main contributor to their default.[7]

Nonpecuniary costs, λ, follow a two-state Markov chain with realizations $\{\lambda_{L,y}, \lambda_{H,y}\}$ that are independent across households, but serially dependent with transition matrix

$$\Pi_\lambda = \begin{bmatrix} \pi & 1-\pi \\ 1-\pi & \pi \end{bmatrix}.$$

Due to data limitations, we assume that the transition probability matrix is symmetric and type-invariant, so the only difference across types in terms of stigma costs are their realizations. Our parametrization is more flexible than we used in previous work (Athreya, Tam, and Young 2009, 2012) so that we can match the default rates across education groups. As we show in a subsequent section, the process is still not

[6] We approximate both e and ν with finite-state Markov chains. This approximation has the convenient property that income is bounded.

[7] See Sullivan, Warren, and Westbrook (2000).

flexible enough to match all the targets of interest, although it does a reasonable job. Households cannot borrow or save during the period in which they declare default; however, they face no restriction in any subsequent period.[8]

Loan Pricing

We focus throughout on competitive domestic 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 $\widehat{\pi} : b \times I \rightarrow [0, 1]$ denote the function that assigns a probability of default to a loan of size b given information I; $\widehat{\pi}(b, I)$ 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 $q(\cdot)$ satisfies

$$q_j(b, I) = \begin{cases} \frac{1}{1+r} & \text{if } b \geq 0 \\ \frac{(1-\widehat{\pi}(b,I))\psi_j}{1+r+\phi} & \text{if } b < 0 \end{cases} \tag{4}$$

given $\widehat{\pi}(b, I)$.

In terms of loan pricing, some remarks are in order. In earlier work, Athreya (2002) specified an exogenous credit limit and then limited the sensitivity of loan pricing by forcing all loans to be priced identically. This approach has the benefit of plausibly capturing the "optionality" of the typical unsecured debt contract, whereby households can count on being able to borrow at a predetermined interest rate up to a predetermined credit limit, i.e., a credit "line." A second benefit from this approach is that it might allow a shortcut to analyzing pooling outcomes that arise from private information on borrower characteristics. However, there are clear drawbacks to this approach as well. First, for the counterfactuals we are interested in, we desire a setting in which both the supply side of the credit market and prices jointly respond to changes in borrowing and repayment incentives. By contrast, in Athreya (2002), only prices responded. For large changes in default incentives, such as what we will examine, this is not a desirable limitation. More recently, Mateos-Planas and Seccia (2006) extended the approach of Athreya (2002) to allow for changes in credit limits, but both it and Athreya (2002) in the end employ a framework substantially different enough to make the comparison to the existing models described at the outset difficult. Second, from even a purely empirical perspective, there are reasons to avoid the use of pooling contracts. As

[8] That is, exclusion from credit markets beyond the initial period is not sustainable as a punishment.

documented in Livshits et al. (2012), and Athreya, Tam, and Young (2012), among others, the variation in unsecured credit terms is now large and appears sensitive to household-level conditions. Lastly, while not directly observable, it is plausible that while individual credit contracts are best characterized by a single interest rate and credit limit, the proper interpretation of credit in the model is the sum of all credit available to the household. In this case, then, the question is the extent to which the household would have to pay more, sooner or later, to acquire additional credit. Our chosen approach features pricing that responds to default in a manner that yields supply-side effects and makes the marginal cost of credit an increasing function.

Returning to the model, r is the exogenous risk-free saving rate and ϕ is a transaction cost for lending, so that $r+\phi$ is the risk-free borrowing rate; the pricing function takes into account the automatic default by those households that die at the end of the period.[9] We assume I contains the entire state vector for the household: $I=(a,y,e,\nu,\lambda,j)$. 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

$$\widehat{\pi}(b,I)=\sum\nolimits_{e',\nu',\lambda'}\pi_e(e'|e)\,\pi_\nu(\nu')\,\pi_\lambda(\lambda'|\lambda)\,d\big(b(a,y,e,\nu,\lambda,j),e',\nu',\lambda'\big). \tag{5}$$

Since $d\left(b,e',\nu',\lambda'\right)$ is the probability that the agent will default in state $\left(e',\nu',\lambda'\right)$ tomorrow at debt level b, integrating over all such events *tomorrow* produces the relevant default risk. This expression also makes clear that knowledge of the persistent component e is critical for predicting default probabilities; the more persistent e is, the more useful it becomes in assessing default risk.

Government

The only purpose of government in this model is to fund pension payments to retirees. The government budget constraint is

$$\tau W\int y\omega_{j,y}e\nu\Gamma(a,y,e,\nu,\lambda,j<j^*)=$$
$$W\int(\upsilon\omega_{j^*-1,y}ye_{j^*-1}\nu_{j^*-1}+\Upsilon)\,\Gamma(a,y,e,\nu,\lambda,j\geq j^*).$$

The left-hand side is the total revenue obtained by levy of a flat tax rate τ on all working agents, where the distribution of working

[9] We assume any savings of households who die is taxed at 100 percent and used to fund wasteful government spending.

households (those for whom $j < j^*$) over productivity levels and age is given by $\Gamma(\cdot)$. The right-hand side is the total expenditure on retirees (those for whom $j \geq j^*$). Recall that to provide a tractable representation of social security and retirement benefits, we assume that retirement income is composed of a fraction $\upsilon \in (0, 1)$ of income in the last period of working life plus a fraction Υ of average income (which is normalized to 1).

Price Determination

We assume that the risk-free rate r is exogenous and determined by the world market for credit. Given r, profit maximization by domestic production firms implies that

$$W = (1 - \alpha) \left(\frac{r}{\alpha}\right)^{\frac{\alpha}{\alpha - 1}},$$

where α is capital's share of income in a Cobb-Douglas aggregate production technology. Our assumption that the risk-free rate is exogenous deserves discussion. It is certainly reasonable to assume that the U.S. capital market is open, so empirically it is not implausible. Furthermore, if we close the economy we confront the high concentration of wealth puzzle directly—the median-wealth agent in the United States has little or no wealth and thus cares about default policy, since they may borrow in the future if unlucky, while the mean agent holds substantial wealth and is unlikely to be concerned with the default policy in place.[10] There is a caveat, however. Li and Sarte (2006) is an early article that establishes a role for general equilibrium feedback effects that overturn partial equilibrium implications. Though we suspect our findings are robust to the determination of the risk-free rate via general equilibrium restrictions, it is not known for sure whether this is the case.

Preferences

Here we present the recursive representations of the preferences we study.

[10] Chatterjee et al. (2007) calibrate their model to match the wealth distribution in the United States in a dynastic setting. As we have argued, life-cycle considerations are important for assessing the welfare effects of bankruptcy.

Constant Relative Risk Aversion

The agent's problem is standard under CRRA preferences, with the Bellman equation for a household of age j given by

$$v(a,y,e,\nu,\lambda,j) = \max_{b,d(e',\nu',\lambda')\in\{0,1\}} \left\{ \begin{array}{c} \frac{n_j}{\rho}\left(\frac{c_j}{n_j}\right)^{\rho} + \\ \beta\psi_{j,y}(EU) \end{array} \right\}$$

$$EU = \sum\nolimits_{e',\nu',\lambda'} \pi_e(e'|e) \times \pi_\nu(\nu')\pi_\lambda(\lambda'|\lambda) V\left(\begin{array}{c} b,y,e',\nu',\lambda',j+ \\ 1 \end{array} \right)$$

$$\begin{aligned} V(b,y,e',\nu',\lambda',j+1) = {} & (1-d(e',\nu',\lambda'))\, v(b,y,e',\nu',\lambda',j+1) + \\ & d(e',\nu',\lambda')\, v^D(0,y,e',\nu',\lambda',j+1), \end{aligned} \quad (6)$$

subject to the budget constraint given in (1) and (3), depending on their age.

The value function for a household that defaulted in the current period is given by

$$\begin{aligned} v^D(0,y,e,\nu,\lambda,j) &= \max\left\{\frac{n_j}{\rho}\left(\lambda\frac{c_j}{n_j}\right)^{\rho} + \beta\psi_{j,y}(EU)\right\} \\ EU &= \sum\nolimits_{e',\nu',\lambda'} \pi_e(e'|e) \times \\ & \quad \pi_\nu(\nu')\pi_\lambda(\lambda'|\lambda)\, v\left(\begin{array}{c} 0,y,e',\nu',\lambda',j+ \\ 1 \end{array} \right). \end{aligned} \quad (7)$$

$1-\rho \geq 0$ is the coefficient of relative risk aversion and also the inverse of the elasticity of intertemporal substitution. Given our assumptions, the budget constraints remain the same as for all other agent types, aside from current net worth being zero as a result of the default.

Epstein-Zin

Under Epstein-Zin preferences, a household of age j solves the dynamic programming problem

$$\begin{aligned}
v(a,y,e,\nu,\lambda,j) &= \max_{b,d(e',\nu',\lambda')\in\{0,1\}} \left\{ \begin{array}{c} n_j\left(\frac{c_j}{n_j}\right)^{\rho} + \\ \beta\psi_{j,y}(EU)^{\frac{\rho}{1-\sigma}} \end{array} \right\}^{\frac{1}{\rho}} \\
EU &= \sum\nolimits_{e',\nu',\lambda'} \pi_e(e'|e) \times \\
&\quad \pi_\nu(\nu')\pi_\lambda(\lambda'|\lambda) V(b,y,e',\nu',\lambda',j+1) \\
V(b,y,e',\nu',\lambda',j+1) &= (1-d(e',\nu',\lambda')) \times \\
v\left(\begin{array}{c} b,y,e',\nu',\lambda',j+ \\ 1 \end{array} \right)^{1-\sigma} + &\quad d(e',\nu',\lambda') \times \\
&\quad v^D(0,y,e',\nu',\lambda',j+1)^{1-\sigma}, \qquad (8)
\end{aligned}$$

subject to the usual budget constraints, and where

$$\begin{aligned}
v^D(0,y,e,\nu,\lambda,j) &= \max\left\{ n_j\left(\lambda\frac{c_j}{n_j}\right)^{\rho} + \beta\psi_{j,y}(EU)^{\frac{\rho}{1-\sigma}} \right\}^{\frac{1}{\rho}} \\
EU &= \sum\nolimits_{e',\nu',\lambda'} \pi_e(e'|e)\pi_\nu(\nu') \times \\
&\quad \pi_\lambda(\lambda'|\lambda) v\left(\begin{array}{c} 0,y,e',\nu',\lambda',j+ \\ 1 \end{array} \right) \qquad (9)
\end{aligned}$$

is the value of default. $\sigma \geq 0$ governs the household's aversion to fluctuations in utility across states of nature while $\rho \leq 1$ controls the substitutability between current and future utility; specifically, σ is the coefficient of relative risk aversion with respect to gambles over future consumption and $\frac{1}{1-\rho}$ is the elasticity of intertemporal substitution in consumption. When $\rho = 1 - \sigma$, these preferences generate the same ordering over stochastic streams of consumption as expected utility does.

2. RESULTS

The results are organized into two subsections. First, we study the roles played by aversion to fluctuations in consumption over time and across states-of-nature. We begin with expected utility preferences. We then relax this by employing Epstein-Zin preferences. Throughout this subsection, we consider parameter values that lie near the values implied by the benchmark calibration; these values ensure that model outcomes remain in congruence with cross-sectional facts on consumption and income inequality. We show that welfare under the default option is lower, at least *ex ante*. Second, based on this result, we ask the "inverse" question: Are there economies in which welfare in the standard model is worse? In this subsection, we no longer restrict ourselves to parameters dictated by U.S. data; rather, our goal is to understand whether any parameterizations within the parametric classes we study are capable of generating lax default as a welfare-improving policy. Specifically, we consider shocks with counterfactually large persistent and transitory components and preferences that display ambiguity aversion.

As noted at the outset, our approach throughout will shut down expense shocks in an otherwise standard consumption smoothing problem. A caveat is in order. While we have argued that this is informative about a case in which insurance is introduced where it was previously missing, it should be recognized that this is not necessarily identical to that case. In particular, the most direct route to addressing the question of whether default would remain useful if society located a way to insure what are presently uninsurable expenses is to explicitly model such an option. We opt for a simpler approach here in part because the form of such insurance, were it to become available, is not obvious a priori. This is primarily because it is unlikely to be provided privately, given that it has not emerged to this point. As a result, the form it takes will likely be as part of a tax-transfer scheme. Our model lacks the detail needed to address the associated incentive-related costs. Our approach is therefore similar to the thought-experiment of Lucas (1987) in which the costless removal of risks was employed as a benchmark for the gains from business cycle stabilization. Still, the reader should keep in mind the indirect nature of our approach and the limits it places on the interpretation of our results. In particular, our approach leads us to calibrate more than once, sometimes with only partial success, depending on the case under study, as opposed to calibrating once at the outset. We acknowledge this limitation and leave the alternate approach for future work.

Does Default Help Insure Labor Income Risk?

In this subsection, we evaluate the implications of default relative to the standard model for a variety of empirically plausible values for agent attitudes toward intra- and intertemporal consumption smoothing. Before evaluating these alternatives, we present our argument for why default regimes must be a matter of policy rather than an endogenous outcome of decentralized trading arrangements. The most prevalent form of explicitly unsecured credit is that arising from the open-ended revolving debt plan offered by credit card lenders. Credit card lending, in turn, has been (certainly since the mid-1990s) extremely competitive.[11] The relevance of the competitiveness of the U.S. unsecured lending industry is that the credit market cannot be punitive in its treatment of those who default. That is, no single firm would be willing to treat an individual borrower any worse than the current assessment of their state would justify. As a result, a household contemplating default in such a setting can safely rule out being "punished" for it. In the case where default conveys no additional information to a lender than what it was able to observe *ex ante*, there is literally no change in terms that are "caused" by the act of reneging on a payment obligation. Conversely, when default does reveal information, the change in terms is again not "punitive" in nature, but instead reflects an updated assessment of default risk. As a result, "high" *ex post* interest rates following default are implausibly ascribed to deadweight loss-inducing penalties. In the symmetric-information and competitive setting we study, punishments that are ex post inefficient will not be sustainable. Even if any single lender could withhold credit after default, the presence of other lenders would undermine the possibility of anything purely punitive. As a result, default costs capable of sustaining unsecured credit markets are likely to require intervention by policymaking authorities.[12] Thus, in the market for unsecured consumer debt, it is likely that *any costs of default filing that are in any way punitive have to be policies.*[13]

At the outset, we noted that for plausible parameterizations of preferences that admit an expected utility representation, the

[11] The average interest rate on credit card balances is high—currently 14 percent—relative to more secured forms of debt. As Evans and Schmalensee (2005) have pointed out, however, it is straightforward to account for the interest rate after funding costs, transactions costs, and, most crucially, default costs are taken into account, without relying on market power distortions.

[12] Most dynamic contracting models of limited borrower commitment, for example, currently use implicit or explicit appeals to public institutions with commitment to punish, in order to motivate penalties for the value of autarky. In recent work, Krueger and Uhlig (2006) show that the inability of the supply side of the credit market to commit to punishments can have severe implications for the existence of the market itself. In the "normal" case, Krueger and Uhlig (2006) show that competition in fact collapses credit and insurance markets completely even without informational frictions.

[13] We want to be clear that what we call "penalties" differs from the usage in Ausubel and Dawsey (2008), where rates imposed after late or missed payments are labeled punitive. They attribute the high values of such rates to a common agency problem. Modeling the bilateral contracting problem that would arise in the presence of noncompetitive intermediation is well beyond the goals for this article. We are pursuing the endogenous determination of interest rate hikes for delinquent borrowers in other work.

standard model typically maximizes welfare. Our first step is to understand whether this argument against default obtains only because of the restriction to expected utility or is a more fundamental property of models of life-cycle consumption smoothing. To collapse the model to the standard model, the specific quantitative experiment we consider is the imposition of a cost of default Δ that is large enough to eliminate all default on the equilibrium path.[14] Before proceeding, we note the following property of our model.

Proposition 1 *For each* (a, y, e, ν, j) *there exists* Δ *large enough that* $\widehat{\pi}(b) = 0$.

This result relies on the nonnegativity condition for consumption—if Δ exceeds the labor income of the household in the current period, default cannot occur since consumption would have to be negative. Given that total labor income is bounded (by assumption) and borrowing is proscribed in the period of default, we can always impose a cost of filing sufficient to generate zero default along the equilibrium path. We then compute the change in lifetime utility for each individual given a Δ that exceeds the maximum required; in the absence of general equilibrium effects, we can compute these changes for each individual, rather than simply for newborns, without the need to track transitional dynamics. We will focus in general on *ex ante* welfare of newborns.

Calibration

We consider a benchmark case of expected utility, where $\rho = 1 - \sigma = -1$. We choose $(\beta, \lambda_{L,y}, \lambda_{H,y}, \pi)$ to match the default rates of each type y, the measure of negative net worth as a fraction of gross domestic product for each type y, the fraction of borrowers, and the discharge ratio (mean debt removed via default divided by mean income at time of filing). Table 1 contains the constellation of parameters that fits best (when viewed as exactly identified generalized method of moments with an identity weighting matrix). Other parameters are identical to those in Athreya, Tam, and Young (2009)—these include the resource cost of default Δ, the income processes faced by each type, the measure of each type, and the parameters of the retirement system (θ, Θ).[15]

[14] Similar results would obtain if the government could impose "shame" on households by choosing values for λ, provided it could make λ large enough to guarantee zero default on the equilibrium path. In our model, the Inada condition on consumption implies that such a λ always exists.

[15] Specifically, we set $\upsilon = 0.35$, $\Upsilon = 0.2$, $\phi = 0.03$, $\Delta = 0.03$, $\varsigma = 0.95$, $\sigma^2_{n,\epsilon} = 0.033$, $\sigma^2_{n,\nu} = 0.04$, $\sigma^2_{h,\epsilon} = 0.025$, $\sigma^2_{h,\nu} = 0.021$, $\sigma^2_{c,\epsilon} = 0.016$, and $\sigma^2_{c,\nu} = 0.014$.

Table 1 Calibration

Case	$\rho = -1, \sigma = 2$		$\rho = -0.5, \sigma = 2$		$\rho = -1, \sigma = 5$	
Parameter,Target	**Parameter**	**Outcome**	**Parameter**	**Outcome**	**Parameter**	**Outcome**
$\lambda^h_{nhs}, \pi_{nhs} = 1.03\%$	0.8972	0.31%	0.8668	1.24%	0.9376	0.51%
$\lambda^l_{nhs}, E(\frac{b}{y}\|b<0)_{nhs} = 0.1552$	0.7624	0.2071	0.6929	0.2104	0.7538	0.1561
$\lambda^h_{hs}, \pi_{hs} = 1.11\%$	0.8832	0.97%	0.8064	1.29%	0.8872	1.31%
$\lambda^l_{hs}, E(\frac{b}{y}\|b<0)_{hs} = 0.5801$	0.7135	0.1835	0.6933	0.1825	0.6236	0.2553
$\lambda^h_{coll}, \pi_{coll} = 0.57\%$	0.7067	0.63%	0.7136	0.79%	0.7055	0.76%
$\lambda^l_{coll}, E(\frac{b}{y}\|b<0)_{coll} = 0.7251$	0.5698	0.1504	0.6352	0.1506	0.4205	0.2194
$\beta, \Pr(b<0) = 12.5\%$	0.9765	17.5%	0.9895	13.3%	0.9532	12.5%
$\rho_\lambda, \frac{E(b\|d=1)}{E(y\|d=1)} = 0.56$	0.8597	0.3986	0.6655	0.4073	0.7658	0.4630

Our model is not capable of exactly matching the entire set of moments—for example, we underpredict default rates and discharge, generally underpredict debt-to-income ratios, and overpredict the measure of borrowers. This inability arises because the model actually places very tight links between some variables, restricting the minimization routine's ability to independently vary them.[16] In the end, one either accepts that expense shocks do indeed play a very dominant role in default data, or one is left with a puzzle relative to standard consumption-savings models. Still, we note that the qualitative findings from our analysis do not appear to depend on our specification of the stochastic process for λ.[17]

Expected Utility and Ex Ante Welfare

We consider two environments—one with the calibrated value for Δ and one with a cost Δ sufficient to eliminate default on the equilibrium path. Table 2 contains the welfare gain from the standard model in which it is infeasible for any household to declare default. Consistent with our previous work, we find that welfare is higher in the standard model *ex ante* for every newborn (independent of type). College types benefit the most from the change, and their welfare gain is substantial (1.2 percent of lifetime consumption). To aid the discussion in subsequent sections where we alter preference parameters, we quickly summarize the reasons for the welfare gains here.

In the standard model, the loss of resources generated by the filing cost is not present. Since we do not impose an economy-wide resource constraint, these lost resources are not important. Instead, the welfare gain is driven by an improved allocation of consumption. By the law of total variance, the variance of consumption over the life cycle can be decomposed into two components:

$$Var\left(\log\left(c\right)\right) = Var\left(E\left[\log\left(c\right)|\text{age}\right]\right) + E\left[Var\left(\log\left(c\right)|\text{age}\right)\right].$$

[16] Consider an attempt to improve the model's prediction for the measure of borrowers by increasing β. Holding all other parameters constant, this reduces default rates and debt-to-income ratios for all types (and these variables are generally already too small). To counteract this effect, one might then move λ for each type and each state. Consider first increasing both λ_i^H and λ_i^L for one type i. While this change would increase the default rate—default becomes less costly—it would via a supply-side effect tend to reduce debt levels (see Athreya [2004]). By contrast, suppose we increase λ_i^H and decrease λ_i^L; this change has countervailing effects on both default rates and debt levels and default rates could rise because it becomes cheaper for H types, but fall as it becomes more expensive for L types. A similar tension exists for debt-to-income ratios—driving it up for one type tends to drive it down for the other.

[17] In the real world, "stigma" may also be a function of aggregate default rates (an agent cares less about default if everyone else is defaulting), in which case this invariance may break. To analyze this case would be of interest, but it poses some challenges with respect to calibration. We therefore defer it to future work.

Table 2 Welfare Gains (without Recalibration)

$\sigma = 2$ & $EIS = 0.5$	Coll	HS	NHS
$DM \rightarrow SM$	1.21%	0.54%	0.52%
$\sigma = 2$ & $EIS = 0.67$	Coll	HS	NHS
$DM \rightarrow SM$	0.58%	0.21%	0.13%
$\sigma = 5$ & $EIS = 0.5$	Coll	HS	NHS
$DM \rightarrow SM$	0.47%	0.16%	0.13%

We label the first term the "intertemporal" component of consumption smoothing; it represents how expected consumption differs across time periods. The second term is the "intratemporal" component; it measures how much consumption varies across agents of a given age. Roughly speaking, how costly the first component is in terms of welfare depends on the elasticity of intertemporal substitution, because it measures the deterministic variance of consumption over time, whereas the welfare cost of the second part is governed by static risk aversion. In Figure 1 we see that the standard model, or "no-default" case (SM), improves intertemporal smoothing (the curve gets flatter) because all lending becomes risk-free. Thus, as we noted in the introduction, the only debt limit that is relevant is the natural debt limit, which is very large in our model for newborn agents. Turning to the intratemporal component, in Figure 2 we see that the SM improves this as well, restating the analysis in Athreya, Tam, and Young (2009) that unsecured credit markets do not provide insurance. Here, bad shocks trigger tightening of credit constraints, making consumption smoothing across states of nature more difficult. As a result, young agents are unable to respond effectively to bad income realizations when they can default, causing their consumption to be highly volatile. Under the SM, the natural debt limit is sufficient to protect them against adverse shocks; by middle age, default has ceased to be relevant and thus the two cases largely coincide.[18]

The differences in outcomes across the DM and SM cases are given in Figures 3, 4, and 5, and are driven by changes in the pricing functions agents face. In Figure 3 we show the pricing functions in the low costs of default environment facing a young college agent across realizations of the persistent shock e. The initial flat segment is driven by Δ and is increasing in the current realization of the persistent shock e. As debt

[18] The figures are drawn for the aggregate, since the results are the same for each type qualitatively. Figures decomposed by type are available from the authors upon request.

Figure 1 Intertemporal Consumption Smoothing, Expected Utility

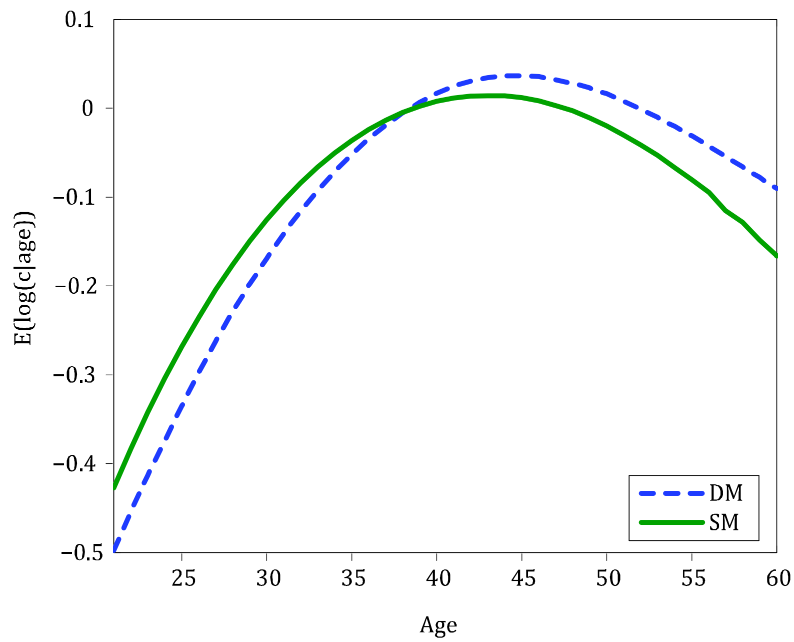

increases, more realizations of e' would trigger default, causing q to decline until it reaches zero; looking across e values we see that higher e realizations permit more borrowing. Of course, higher e realizations in our model are typically associated with less, not more, borrowing, so these increased debt limits are not particularly valuable; instead, the tightening of credit limits when e is low generates substantial costs for poor agents. In contrast, under SM pricing is flat out to the natural debt limit. Crucially, transitory shocks do not impact pricing; because ν' cannot be predicted using ν, the current transitory shock has no effect on the default decision tomorrow conditional on b (b is changed by the transitory shock, however).

The potential tradeoff between the two components of smoothing motivated the life-cycle analysis of Livshits, MacGee, and Tertilt (2007) and Athreya (2008), so why doesn't default generate this tradeoff? As discussed in Athreya, Tam, and Young (2009), default can either help or hinder intratemporal smoothing, depending on which agent you ask. An agent facing an income process with low intertemporal variance

Figure 2 Intratemporal Consumption Smoothing, Expected Utility

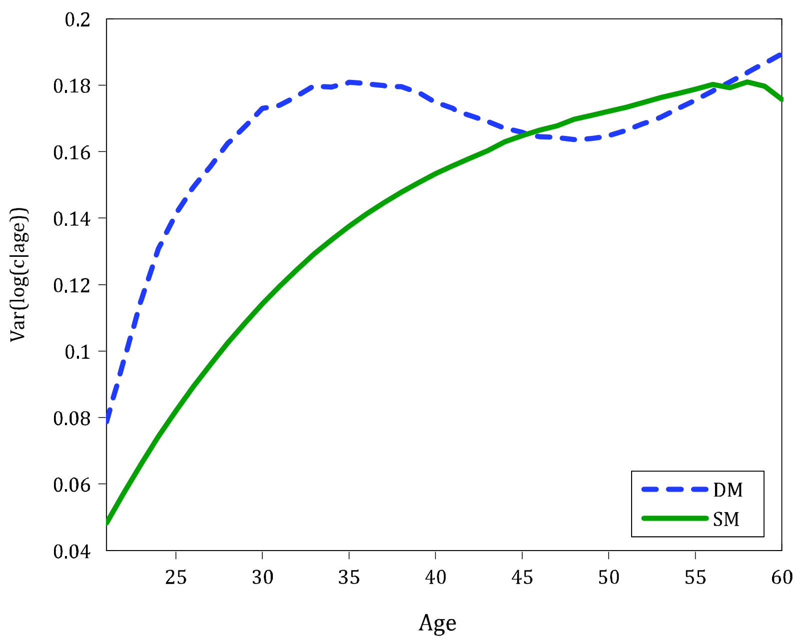

but high intratemporal variance—that is, tomorrow's expected income is close to current income but tomorrow's income has substantial risk—may benefit from default; the intertemporal distortion is minimal while the potential to truncate the consumption distribution at the low end conveys significant benefits (even once pricing is taken into account). In contrast, an agent facing the opposite process—income that grows over time and is relatively safe—generally does not benefit; default is not used because pricing prevents it and the intertemporal distortion is substantial, leading to significant welfare losses. In our model, a young agent is of the second type, especially a college-educated one, while older households are members of the first type.

Ex Post Welfare—Voting over Default Policy

Because we study a small open-economy model in which the risk-free rate is fixed, but also allow all pricing to be individualized, there are

Figure 3 Pricing, Expected Utility

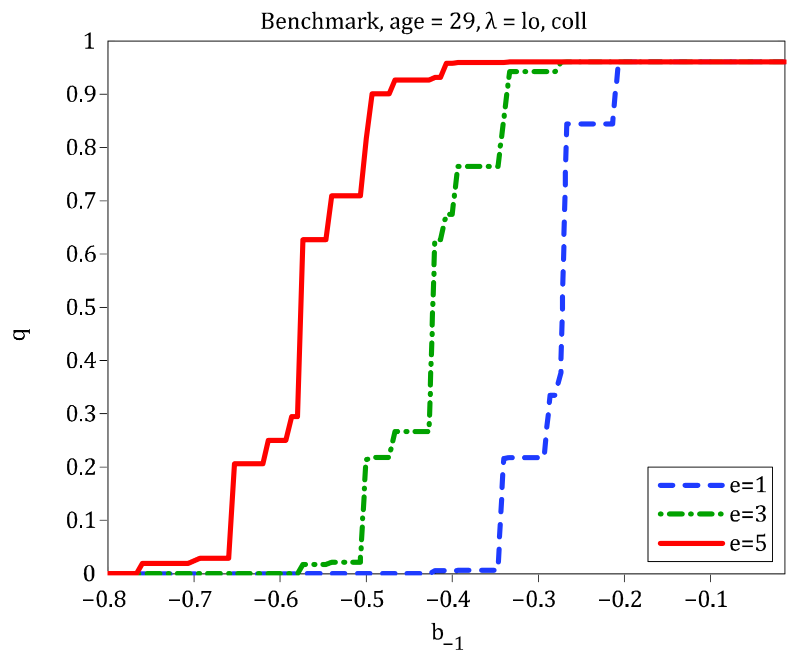

no "pecuniary" externalities. We can therefore compute the welfare consequences of policy changes for any agent at any point in the state space; since the distribution plays no role in pricing (and therefore no role in welfare), we do not need to calculate the transitional dynamics of the model to get the welfare changes. We ask agents of a given age and type whether, conditional on their current state, they would be in favor of eliminating the option to declare default. Figure 6 displays the measure of each type, conditional on age, that would support retaining default with the calibrated Δ. A substantial portion of college types oppose elimination, but they are all middle-aged and have experienced histories of bad shocks; the peak in opposition occurs earlier for high-school types and later for non-high-school types, with correspondingly fewer such households opposing overall. For the convenience of the reader, Table 3 presents the aggregate measures of each type that oppose eliminating default (the column labeled "DM Regime"); they are small for each education group. Furthermore, as is clear from the figures, almost no newborns oppose eliminating the option.

Figure 4 Pricing, Expected Utility

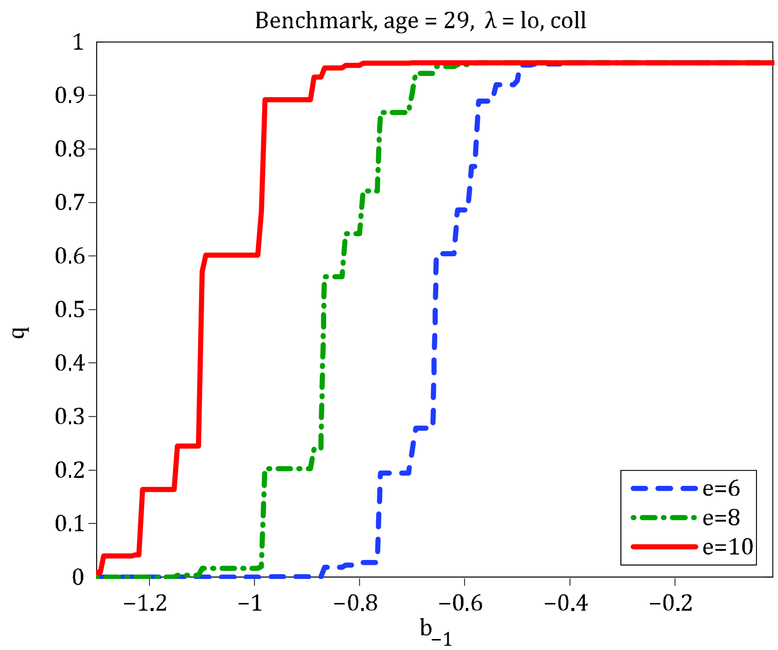

We now consider the inverse of the preceding experiment: Agents of different ages and types are asked if they would prefer to *introduce* default (again, with Δ set to its calibrated value) into a setting in which it is currently prohibited. As seen in Figure 7, a nontrivial fraction of agents would like to introduce default. The intuition here is that the no-default case allows significant borrowing at the risk-free rate. As a result, many households, especially the college-educated, borrow when young in anticipation of higher earnings. The relatively unlucky among them then find themselves indebted by middle age and thereby will benefit from the discharge of debts. Moreover, by virtue of being middle-aged, these households place relatively low value on being able to access the cheap unsecured debt later in life. This effect is especially strong for the college-educated, for whom purely intertemporal consumption smoothing motives dictate a strong effort to save for retirement beyond middle age. As a result, a substantial proportion of high-school- and college-educated household groups would support the introduction of default when they reach middle age. In contrast,

Figure 5 Pricing, Expected Utility

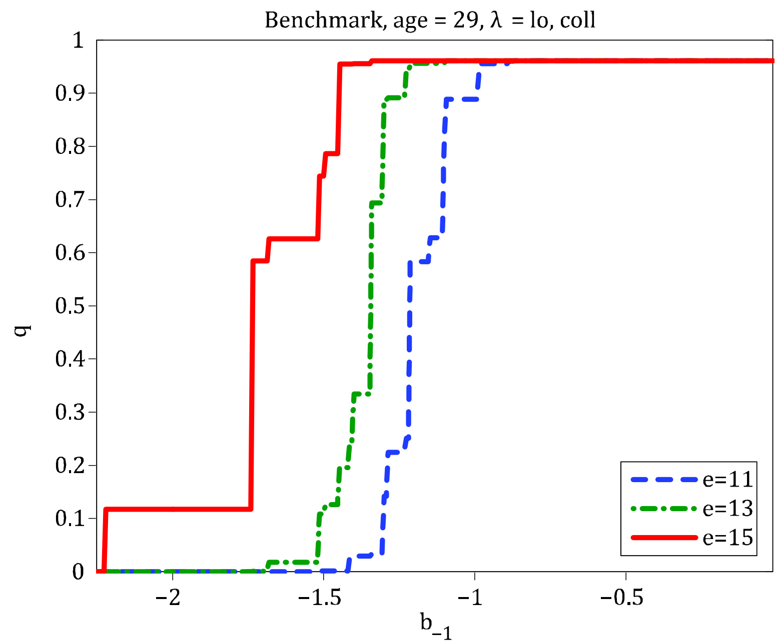

those who have not completed high school support the introduction of default only late in working life, when the subsequent increase in borrowing costs is not long-lasting. However, as Table 3 shows (the column "SM Regime"), the aggregate number of agents who vote in favor of introduction falls well short of majority status.

Separating Risk Aversion from Intertemporal Substitution

As discussed above, the two pieces of the variance decomposition have welfare costs that depend (mainly) on different aspects of preferences. Our benchmark case using CRRA expected utility restricted these two aspects of preferences to be reciprocals of each other. Here, we relax that requirement by using the Epstein-Zin preference structure, and consider two particular deviations. First, we make households more tolerant of intertemporal variance than in the expected utility benchmark by employing a high value for ρ. Second, the default option

Figure 6 Fraction Supporting Bankruptcy, BK Regime

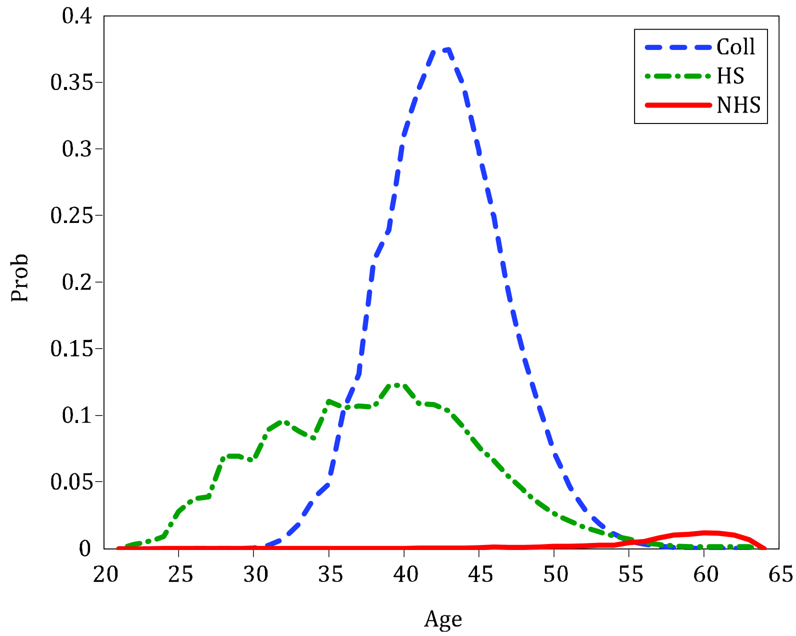

may shrink the volatility of intratemporal consumption, at least for some ages. Given this, making intratemporal variance more painful to households may help us explain the presence of low default costs. We therefore select a relatively high value for σ. It is important to note that this particular combination of insensitivity to the timing of consumption and sensitivity to the income state in which it occurs is the arrangement that gives default its best chance of improving *ex ante* welfare and does not lie within the class of expected utility preferences.

The specific experiments we investigate involve changing ρ and σ without recalibrating the entire model. This type of change generates two effects—an effect conditional on borrowing (which we call the price effect) and an effect caused by changes in the number of borrowers (the extensive effect). We then compare the results with cases where the model is recalibrated (to the extent that is possible) in an attempt to isolate the two effects.

We first consider changes in ρ. To understand how this change affects welfare, it is helpful to first consider the extreme case of $\rho = 1$,

Table 3 Measure of Agents in Favor of Bankruptcy

Education	DM Regime	SM Regime
College	6.45%	4.09%
High School	4.05%	3.26%
Non-High School	0.16%	0.24%
Total	4.05%	2.98%

making the household infinitely willing to move consumption deterministically through time. As $\rho \to 1$, the Bellman equation converges to the form

$$v(a, y, e, \nu, \lambda, j) = \max\left(c_j + \beta\psi_{j,y}\left(\begin{array}{c}\sum_{e',\nu',\lambda'} \pi_e(e'|e)\,\pi_\nu(\nu') \times \\ \pi_\lambda(\lambda'|\lambda)\,V(b, y, e', \nu', \lambda', j+1)\end{array}\right)^{\frac{1}{1-\sigma}}\right).$$

Here, the household will either completely frontload or backload consumption, depending on the relationship between the discount factor and the interest rate. For the parametrization we use, the effective discount factor (β times the survival probability) lies between the risk-free saving and borrowing rates for almost every age, meaning that households don't wish to borrow and, critically, do not value the default option *no matter how risk averse* they are. For some older households, whose survival probabilities are relatively low, the effective discount factor is sufficiently low that they want to borrow and "frontload" their consumption; the option to default makes borrowing expensive enough to render complete frontloading impossible. This, in turn, reduces the welfare of these households—since they face no uncertainty, default is either probability zero or one and pricing therefore eliminates it. Obviously such extreme consumption behavior is inconsistent with U.S. cross-sectional facts; in particular, the model with $\rho = 1$ would miss very badly on the life-cycle pattern of consumption inequality, which in the data is substantially smaller than income inequality.

Returning to less extreme values for ρ, Figure 8 displays the pricing function across several different values of ρ and demonstrates the effect on loan prices. As ρ increases, the pricing function shifts downward because at any given level of debt an agent with a higher ρ is more willing to default. The intuition for this result is not straightforward. When ρ increases, the household is more willing to accept variability in consumption across time. If a household enters the current period with some debt and wishes not to lower debt, they have two options: (i) borrow more if possible or (ii) default and void those obligations. Borrowing more is only feasible if there is a reasonable commitment to repay. But since a bad shock would lead to low mean consumption,

Figure 7 Fraction Supporting Bankruptcy, NBK Regime

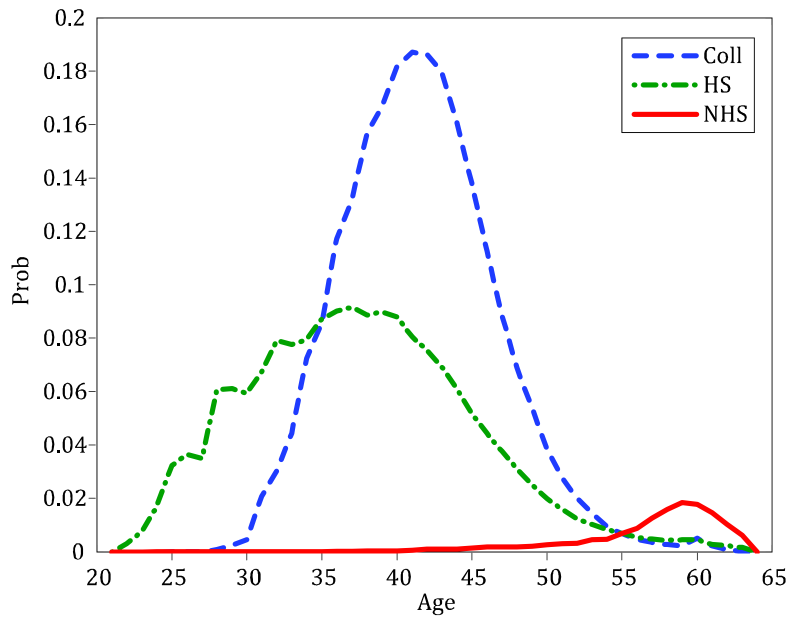

default becomes attractive, and households lack strong commitment to repay debt. As a result, they cannot borrow easily. For the cases with "intermediate" values for ρ, the creation of strong default incentives makes intertemporal smoothing more costly, but the latter is relatively unimportant.

Consider next an experiment where σ, the risk aversion with respect to gambles over future utility, is increased. Again, turning first to the polar case, let $\sigma \to \infty$, so that the household becomes infinitely risk averse. In this case, the limiting household Bellman equation takes the form

$$v(a, y, e, \nu, \lambda, j) = \max \left\{ \begin{array}{c} n_j \left(\frac{c_j}{n_j}\right)^\rho + \\ \beta\psi_{j,y} \min_{e',\nu',\lambda'} \left\{V\left(b, y, e', \nu', \lambda', j+1\right)\right\}^\rho \end{array} \right\}^{\frac{1}{\rho}},$$

subject to the usual budget constraints seen earlier in equations (1) and (3).

When households are infinitely risk averse, they choose not to borrow for the reasons outlined in Athreya, Tam, and Young (2009)—

Figure 8 Pricing, Epstein-Zin with Different EIS

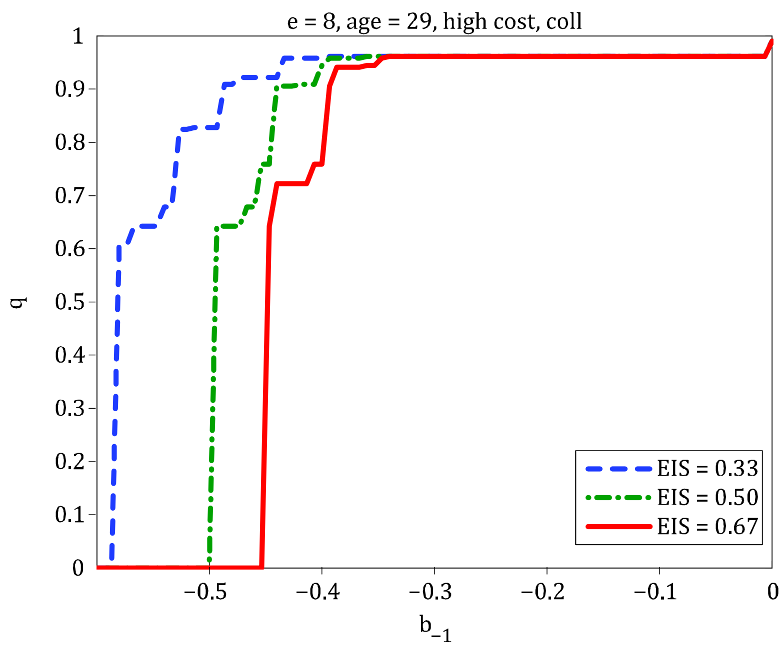

unsecured credit markets do not provide insurance and thus agents will be unwilling to pay the transaction cost to borrow. As a result, there is a welfare gain to living in the standard model, as no household has negative net worth. Again, extreme preferences render the model grossly inconsistent with cross-sectional facts; here, consumption inequality would be essentially zero over all ages.

Returning again to more intermediate cases, we see that changes in risk aversion generates two effects. The extensive margin effect is similar to increasing ρ, but for different reasons. When σ is large, households have a strong demand for precautionary savings; for $\sigma = 5$, for example, we see a clear decline in the measure of total borrowers, again making default overall less damaging. The pricing effect is also similar; by increasing risk aversion, we make the household less willing to have consumption differ across states of the world tomorrow. Conditional on borrowing, the pricing functions reveal a stronger desire to default—for any given b, the price of debt is decreasing in σ (see Figure 9). As above, there are only two options for a household with

Figure 9 Pricing, Epstein-Zin with Different Risk Aversion

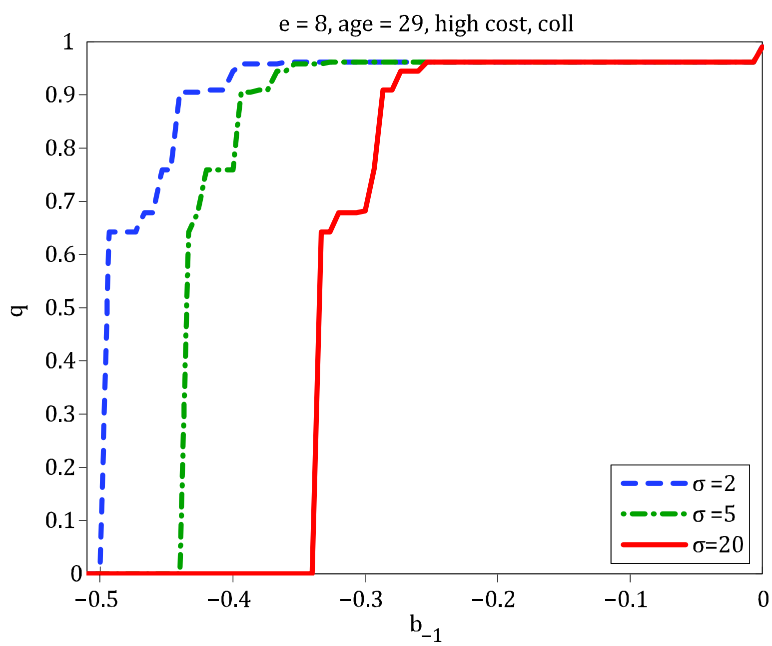

debt; since even a moderately bad outcome will cause a highly risk-averse agent to default, commitment is not possible, leaving default as the only option for smoothing consumption across states.[19] Combining these results into one statement, we see that no combination of (ρ, σ) leads to default being a welfare-improving policy, although for extreme cases it will be nearly innocuous.

Table 2 shows that welfare is higher (for newborns) in the standard model (SM), but that the gains from (imposing the high Δ) decline with risk aversion and elasticity of intertemporal substitution (EIS). $\rho > 1 - \sigma$, which is satisfied when either parameter increases, implies the household has a preference for early resolution of uncertainty; thus, default appears to be least damaging when households prefer to resolve their risk early rather than late.

[19] Our model satisfies the conditions noted in Chatterjee et al. (2007) that imply default occurs only if current debt cannot be rolled over: If $d(\epsilon', \nu', \lambda') > 0$ for some $\epsilon', \nu', \lambda'$, then there does not exist b such that $a + y - q(b, Y)b > 0$ for total income Y.

Table 4 Welfare Gains (with Recalibration)

$\sigma = 2$ & $EIS = 0.5$	College	High School	Non-High School
$DM \rightarrow SM$	1.21%	0.54%	0.52%
$\sigma = 2$ & $EIS = 0.67$	College	High School	Non-High School
$DM \rightarrow SM$	0.28%	0.05%	0.04%
$\sigma = 5$ & $EIS = 0.5$	College	High School	Non-High School
$DM \rightarrow SM$	1.28%	0.57%	0.56%

All of these results are obtained without recalibrating the model. To ensure that our findings are not particularly sensitive to this strategy, we also recalibrate the model for different values of ρ and σ, to the extent that this recalibration is possible; Table 1 contains the new parameter values that best fit the targets under alternative settings. By doing so, we attempt to shut off the extensive margin, although we are not completely successful. When we recalibrate, we find that with high EIS all welfare gains from eliminating default are substantially reduced, with both noncollege types now barely benefiting at all (see Table 4), while for high risk aversion the welfare gains increase slightly. As noted above, this welfare change is entirely due to the shifts in the pricing function that higher EIS and/or higher risk aversion engender. Thus, for no parameter combination that we consider do we observe welfare gains from retaining the default option.

A summary of findings thus far is that default significantly worsens allocations for income risk and preference parameters that are empirically plausible for U.S. data, as well as for more extreme values of preference parameters within the class of Epstein-Zin non-expected utility preferences. We turn now to the question of whether such policies continue to remain desirable under two additional (and more substantial) departures from the settings studied so far.

Is the Standard Model Ever Worse?

We begin this section by allowing for the underlying volatility of income to be driven by relatively more and less persistent income shocks. For this experiment, we hold the unconditional variance of labor income fixed and vary the relative contributions of the persistent component e and the transitory component ν. We then ask whether a relaxation in the household's understanding of the probabilistic structure of earnings risk can open the door for welfare-improving default. For this experiment, we allow for households to display ambiguity aversion in the sense of Klibanoff, Marinacci, and Mukerji (2009).[20]

The Roles of Persistent and Transitory Income Risk

It has long been known that self-insurance, and therefore also the benefit of insurance markets, hinges critically on the persistence of the risks facing households. As a general rule, the more persistent are shocks, the more difficult they are to deal with via the accumulation of assets in good times and decumulation and borrowing in bad times. In contrast, purely transitory income shocks can typically be smoothed effectively. In a pure life-cycle model, however, there are additional impediments to self-insurance: Young households are born with no wealth and often face incentives to borrow arising from purely intertemporal considerations. In particular, those with relatively high levels of human capital, especially the college-educated, can expect age-earnings profiles with a significant upward slope into late middle age. As a result, such households would like to borrow even in the absence of any shocks to income, often substantially, against their growing expected future income. In contrast, those households with low human capital face a far less income-rich future, and as a result borrow primarily to deal with transitory income risk.

In order to understand the role that the persistence of income risk plays in the welfare gains or losses arising from U.S.-style bankruptcy and delinquency, we now evaluate the effects of changes in the persistent component of household income risk for all three classes of households. However, in order to avoid conflating persistence and overall income volatility, we adjust the variance of transitory income volatility such that the overall variance of log labor income remains constant.[21] Figure 10 and Tables 5 and 6 present the welfare and consumption smoothing implications of the standard model under varying income shock persistence. The first column of each table documents the fraction of total variance contributed by the persistent component.

Normatively, three findings are noteworthy. First, and perhaps most importantly, the standard model displays higher welfare irrespective of the nature of shocks accounting for observed income

[20] There are connections between ambiguity aversion and the concept of Knightian uncertainty from Bewley (2002), although the latter concept does not permit preferences to be represented by a utility function and is therefore hard to analyze quantitatively. There are also connections between ambiguity aversion and robust decisionmaking as defined by Hansen and Sargent (2007).

[21] Athreya, Tam, and Young (2009) are primarily concerned with the role of income variance in models of default.

Figure 10 Welfare Gains from Eliminating Bankruptcy

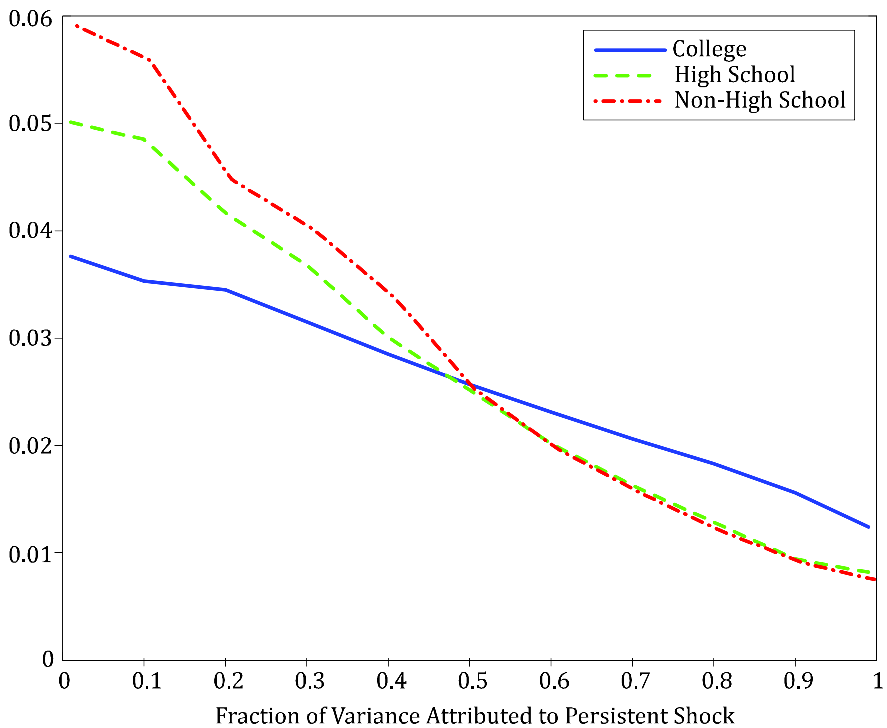

volatility. This result strengthens our findings thus far, and it further suggests that defaultable debt is simply unlikely to be useful to households. It is also a particularly important form of robustness, given both the general importance of persistence for the efficacy of self-insurance and borrowing and because estimates of income shock persistence vary dramatically—see Guvenen (2007), Hryshko (2008), or Guvenen and Smith (2009) for discussions of the debate between so-called "restricted income profiles" (RIP), in which all households draw earnings from a single stochastic process, and "heterogeneous income profiles" (HIP), in which households vary in the processes from which they derive earnings. This debate has implications for models like ours because these two models differ, sometimes strongly, in the persistence of earnings shocks their structure implies. Most recent work now suggests that income-process parameters vary over the life cycle as well (Karahan and Ozkan 2009).

Table 5 Consumption Smoothing (DM)

	Intra			Inter			Total		
	Coll	**HS**	**NHS**	**Coll**	**HS**	**NHS**	**Coll**	**HS**	**NHS**
1.0%	0.0306	0.0462	0.0575	0.0359	0.0364	0.0386	0.0665	0.0826	0.0961
10.0%	0.0377	0.0561	0.0872	0.0343	0.0367	0.0357	0.0720	0.0938	01229
20.0%	0.0459	0.0807	0.1092	0.0336	0.0347	0.0325	0.0795	0.1154	0.1417
30.0%	0.0538	0.0884	0.1367	0.0327	0.0327	0.0297	0.0865	0.1211	0.1664
40.0%	0.0619	0.1013	0.1472	0.0316	0.0301	0.0280	0.0925	0.1314	0.1752
50.0%	0.0700	0.1146	0.1613	0.0305	0.0284	0.0263	0.1005	0.1430	0.1876
60.0%	0.0779	0.1280	0.1797	0.0294	0.0264	0.0241	0.1065	0.1544	0.2038
70.0%	0.0859	0.1413	0.1992	0.0283	0.0247	0.0224	0.1141	0.1660	0.2216
80.0%	0.0946	0.1543	0.2182	0.0272	0.0231	0.0211	0.1218	0.1774	0.2393
90.0%	0.1053	0.1681	0.2368	0.0258	0.0212	0.0199	0.1311	0.1893	0.2567
99.0%	0.1248	0.1863	0.2566	0.0235	0.0187	0.0180	0.1483	0.2050	0.2680

Table 6 Consumption Smoothing (SM)

		Intra			Inter			Total	
	Coll	**HS**	**NHS**	**Coll**	**HS**	**NHS**	**Coll**	**HS**	**NHS**
1.0%	0.0196	0.0307	0.0474	0.0318	0.0314	0.0120	0.0514	0.0621	0.0594
10.0%	0.0271	0.0397	0.0577	0.0315	0.0298	0.0124	0.0586	0.0695	0.0801
20.0%	0.0360	0.0541	0.0771	0.0311	0.0290	0.0131	0.0671	0.0831	0.0902
30.0%	0.0444	0.0683	0.0971	0.0306	0.0284	0.0137	0.0750	0.0967	0.1108
40.0%	0.0524	0.0820	0.1173	0.0300	0.0277	0.0144	0.0824	0.1097	0.1317
50.0%	0.0600	0.0951	0.1364	0.0295	0.0271	0.0151	0.0895	0.1222	0.1515
60.0%	0.0673	0.1076	0.1550	0.0291	0.0267	0.0158	0.0964	0.1343	0.1708
70.0%	0.0743	0.1197	0.1729	0.0288	0.0262	0.0164	0.1031	0.1495	0.1893
80.0%	0.0811	0.1314	0.1903	0.0285	0.0258	0.0171	0.1096	0.1627	0.2075
90.0%	0.0878	0.1428	0.2072	0.0282	0.0255	0.0178	0.1160	0.1638	0.2250
99.0%	0.0935	0.1528	0.2218	0.0280	0.0253	0.0182	0.1215	0.1781	0.2400

Second, the effect of the contribution of persistent shocks to income volatility depends on the education level of households. In particular, when volatility is driven primarily by persistent shocks, the relatively well-educated benefit from the elimination of default substantially more than their less-educated counterparts. Conversely, when most income variability is driven by large but transitory shocks, it is the relatively less-educated who benefit most from the elimination of the default option. The intuition for this result comes from the nature of borrowing: College types borrow primarily to use future expected income today while less-educated types borrow to smooth shocks.

Third, within each educational class, the welfare losses from default decline monotonically as the relative contribution of the persistence of the shock grows; default on debt is least (most) useful when income volatility is driven primarily by shocks that are transitory (persistent). What is surprising, but in keeping with the main theme of our results, is that in *no* case is it true that U.S.-style default is *ex ante* more desirable than allocations obtaining under the standard model. Moreover, even in the case where essentially all income risk is delivered in the form of persistent shocks where credit markets are least useful in dealing with income risk, outcomes that allow for default are worse for agents than those arising in the standard model. The welfare in the standard model is non-trivially higher, at up to 1.24 percent of consumption for college-educated households (as seen in Figure 10).

In Figures 11 and 12 we display the measure of borrowers at each age and the conditional mean of debt among those who borrow for two levels of the importance of persistent income risk.[22] The fact that the losses from allowing default rise for all agent types with the importance of transitory shocks is a consequence of the increased usefulness of credit in dealing with transitory income risk. Conversely, when shocks are primarily persistent, a negative realization requires more frequent borrowing and leads, all else equal, to more debt in middle age; the combination is ultimately unable to stem the transfer of income risk to consumption volatility. In Tables 5 and 6, we see that, irrespective of default policy, persistence translates into higher consumption volatility, and that the presence of lax default policy seen in Table 6 does little to stem the flow of income risk into consumption risk (echoing our previous result in Athreya, Tam, and Young [2009]).

We turn next to the relationship between shock persistence and equilibrium default rates, displayed in Figure 13. Default is “U-shaped,” with high default rates at both ends. To understand this shape, con-

[22] From the perspective of a newborn, the measure of borrowers of a given age equals the probability of the newborn borrowing at that age.

Figure 11 Fraction of Borrowers

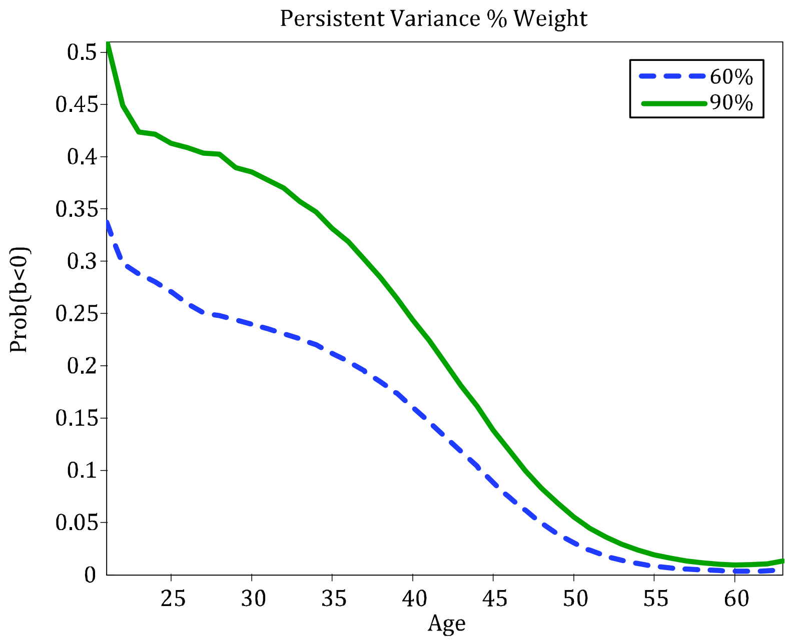

sider first the case where the labor income shocks are nearly all transitory (the left side of the graph). Here, agents can generally manage their risk effectively via saving and dissaving, but they choose to augment the self-insurance mechanism with default at higher rates than they do in the benchmark setting. The reason they do so is that risk-based pricing is not effective here, because there is no useful information contained in the current labor income of the borrower that would identify future bad risks. In contrast, in the case where labor income is driven entirely by the persistent component (the right side of the graph), high default is the result of agents being generally unable to smooth consumption; persistent shocks are hard to smooth using assets alone (and if permanent are in fact impossible). As a result, despite the pricing effects, borrowers will use default relatively often (and pay the costs to do so). The middle parts of the graph, where default is lowest, balance these two effects.

Intuitively, in the standard model, borrowers realize that debt must be repaid, and under high persistence, heavy borrowing in response

Figure 12 Mean Debt of Borrowers

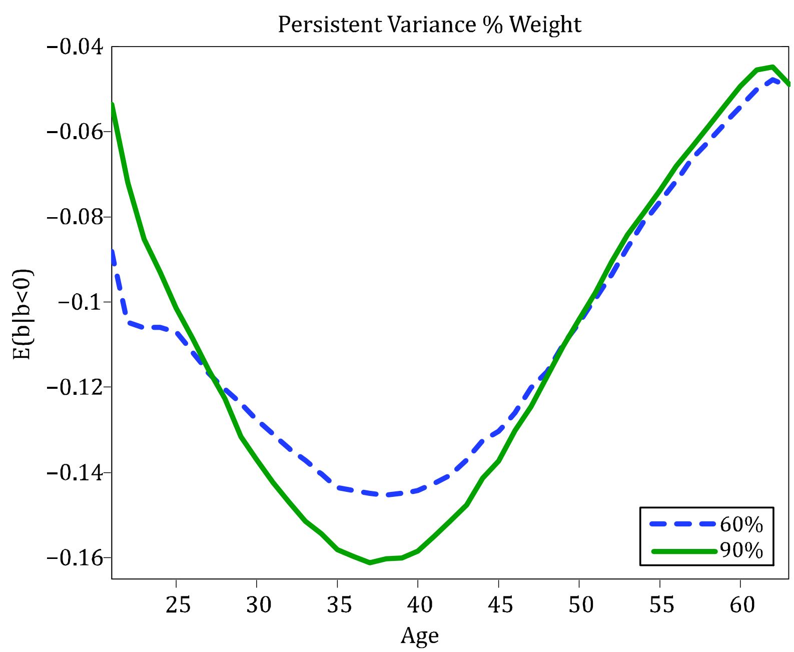

to a negative shock makes low future consumption relatively likely. Nonetheless, credit markets are willing to lend to such households at the risk-free rate (adjusted for any transactions costs of intermediation), making total debt rise. When default is available, borrowing today to deal with persistent income risk does not expose the borrower to severe consumption risk in the long term as default offers an "escape valve," but it does expose lenders to severe credit risk in the near term. Creditors then price debt accordingly; as seen in Figure 14, when shocks are primarily persistent, as the current shock deteriorates so do the terms at which borrowers can access credit. Moreover, under a bad current realization of income, households facing persistent risk see a disproportionate decline in the price of any debt they may issue, while the reverse occurs in the event of a good current realization of income; the pricing functions essentially "switch places." Yet, despite the increased sensitivity of loan pricing to the borrower's current income state under relatively high persistence, the welfare gains under the SM, though still positive, fall. This result obtains because of the reduction

Figure 13 Default Rates

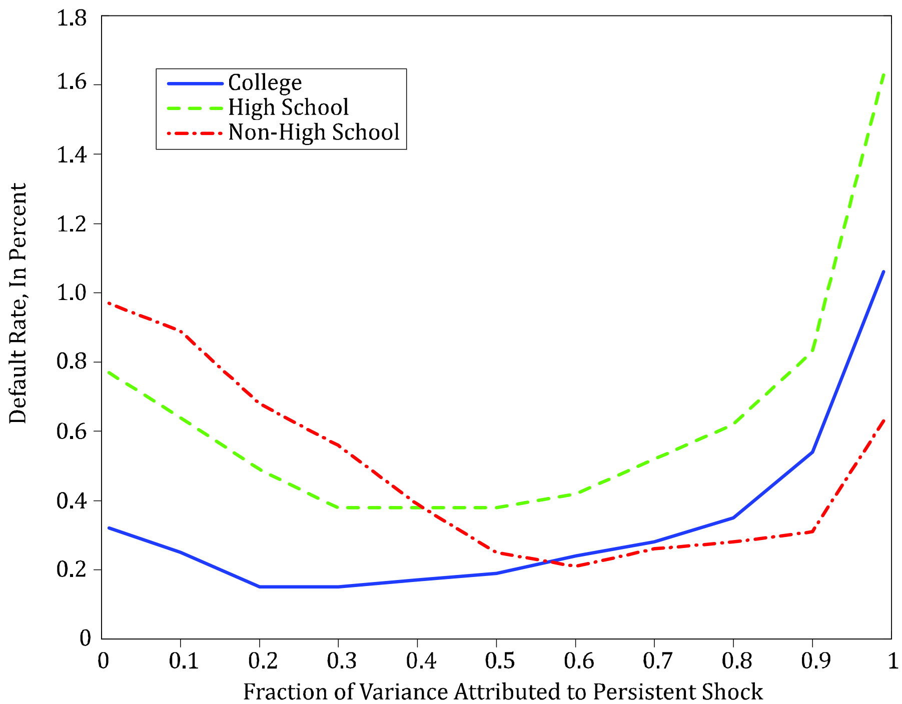

in the ability of self-insurance, inclusive of borrowing, to prevent income fluctuations from affecting consumption. To sum up, income risk is quantitatively relevant in governing the losses conferred by default, but irrelevant for altering the qualitative welfare property that, in the absence of expense shocks, the default option lowers welfare.

Ambiguity Aversion

We turn next to the question of whether default can improve outcomes when households are not perfectly certain about the probabilistic structure of income risk. Households that face ambiguity are uncertain about the probability process for their incomes; if ambiguity-averse, these households behave pessimistically and therefore adopt views about their income that would, for example, imply that it would mean-revert more slowly from low realizations. In such a situation, borrowing to smooth away temporary falls may not be optimal, since asset

Figure 14 Pricing Functions

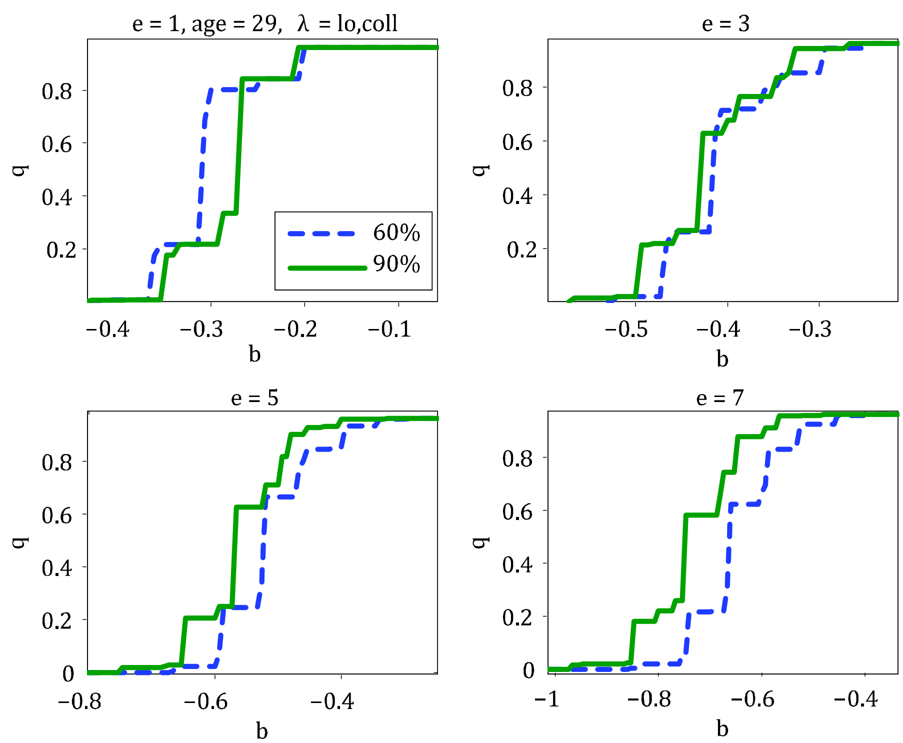

decumulation is not effective against permanent shocks, and therefore in the absence of a default option households may be unwilling to do so. In contrast, if default is an option, the household may be willing to borrow since, even if their pessimism is validated, consumption can be protected via discharge. We formalize this idea, as in Klibanoff, Marinacci, and Mukerji (2009), by assuming agents are averse to ambiguity. In this formulation, a household of age j solves the dynamic programming problem

$$v(a,y,e,\nu,\lambda,j) = \max_{b,d(e',\nu',\lambda')\in\{0,1\}} \left\{ \begin{array}{c} \frac{n_j}{1-\sigma}\left(\frac{c_j}{n_j}\right)^{1-\sigma} + \\ \beta\psi_{j,y} \sum_{e',\nu'} p(e',\nu'|e,\nu) \times \\ \Phi(EU) \end{array} \right\}$$

$$EU = \sum_{e',\nu',\lambda'} \pi_e(e'|e)\,\pi_\nu(\nu')\,\pi_\lambda(\lambda'|\lambda) \times V\left(b,y,e',\nu',\lambda',j+1\right), \tag{10}$$

subject to budget constraints, (1) and (3), where $\Phi(\cdot)$ is given as follows:

$$\Phi(x) = \begin{cases} \frac{1-\exp(-\eta x)}{1-\exp(-\eta)} & \text{if } \eta > 0 \\ x & \text{if } \eta = 0 \end{cases}$$

determines preferences over ambiguity. $\eta \geq 0$ controls the attitude toward ambiguity; as η increases, the household becomes more averse to ambiguous stochastic processes. The restrictions on the choices of $p(e',\nu'|e,\nu)$ are that they must sum to 1 for each (e,ν) and every element must lie in some set $\mathcal{P} \subset [0,1]$; we nest the standard model by setting the $\mathcal{P}$ to be an arbitrarily small interval around the objective probabilities.[23] We use π to denote objective probabilities and p to denote subjective ones; note that households are assumed to be uncertain only about the distribution of income shocks, not the process for λ.

Because we are interested in these preferences only to the extent that they may provide an environment in which relatively low-cost default and debt discharge are welfare-enhancing, we will deliberately take the most extreme case of $\eta = \infty$, yielding the max-min specification from Epstein and Schneider (2003):

$$\begin{aligned} v(a,y,e,\nu,\lambda,j) &= \max_{b,d(e',\nu',\lambda')\in\{0,1\}} \left\{ \begin{array}{c} \frac{n_j}{1-\sigma}\left(\frac{c_j}{n_j}\right)^{1-\sigma} + \\ \beta\psi_{j,y} \min_{p(e',\nu'|e,\nu)} EU \end{array} \right\} \\ EU &= \sum_{e',\nu',\lambda'} p\left(e',\nu'|e,\nu\right) \times \\ &\quad \pi_\lambda\left(\lambda'|\lambda\right) V\left(b,y,e',\nu',\lambda',j+1\right) \\ V\left(b,y,e',\nu',\lambda',j+1\right) &= \left(1-d\left(e',\nu',\lambda'\right)\right) v\left(b,y,e',\nu',\lambda',j+1\right) + \\ &\quad d\left(e',\nu',\lambda'\right) v^D\left(0,y,e',\nu',\lambda',j+1\right), \end{aligned} \tag{11}$$

[23] We do not require that the household assume that the probabilities of the independent events are independent in every distribution that is considered. That is, the household may be concerned that the independence property is misspecified and therefore select a worst-case distribution in which the events are correlated.

where

$$v^D(0,y,e,\nu,\lambda,j+1) = \max\left\{ \begin{array}{c} \frac{n_j}{1-\sigma}\left(\lambda\frac{c_j}{n_j}\right)^{1-\sigma} + \\ \beta\psi_{j,y}\min_{p(e',\nu'|e,\nu)} EU \end{array} \right\}$$

$$EU = \sum\nolimits_{e',\nu',\lambda'} p(e',\nu'|e,\nu) \times \pi_\lambda(\lambda'|\lambda)\, v(0,y,e',\nu',\lambda',j+1) \qquad (12)$$

is the value of default.

The min operator that appears in front of the summation reflects the agent's aversion to uncertainty; as shown by Epstein and Schneider (2003), a household who is infinitely uncertainty-averse chooses the subjective distribution of future events that is least favorable and then makes their decisions based on that subjective distribution. The size of the set of possible processes $\mathcal{P}$ measures the amount of ambiguity agents face; a typical p_{ij} element lies in the interval $\left[\mathbf{p}_1^{ij}, \mathbf{p}_2^{ij}\right] \subset [0,1]$.[24]

Standard ambiguity aversion models imply that households will learn over time and reject stochastic processes that are inconsistent with observed data (for example, a household who initially entertains the possibility of permanently receiving the worst possible income level forever will dismiss this process as soon as one non-worst realization occurs). For simplicity, we will focus our attention on a special case of extreme ambiguity aversion in which this learning does not occur; if default is not useful in this environment, it is likely of less use to households than when they face less uncertainty over time. The intuition is that the income process we buffet agents with is a non-unit process. To the extent that households would realize by a certain age that the data they've received makes unit-root earnings unlikely, they would be able to rule out such a persistent process and thereby smooth more effectively, and as a result, may not value default as much as someone viewing shocks as permanent.

Given the qualifications and considerations discussed above, we now evaluate outcomes in the standard model in the case where $\mathcal{P} = [0,1]$, the most extreme case possible (households behave as if the minimum income draw will be realized with probability 1 next period). The intuition is that such a case offers the possibility, discussed at the outset, that lax penalties for default might actually encourage the use of credit for consumption in a setting where the agent's aversion to ambiguity would otherwise preclude becoming indebted. And in fact, we *do* find that this case delivers default as welfare-improving for some

[24] Hansen and Sargent (2007) provide an interpretation of $\mathcal{P}$ in terms of detection probabilities.

Table 7 Welfare Effects Under Ambiguity Aversion

$\mathcal{P} = [0, 1]$	Non-High School	High School	College
DM $\rightarrow$ SM	0.215%	0.189%	−0.185%
$\mathcal{P} = \min(1, \pi + 0.5)$	Non-High School	High School	College
DM $\rightarrow$ SM	0.296%	0.219%	0.044%

agents (see Table 7). However, this finding is very limited: Benchmark default costs improve welfare for only the college type and the welfare gain is tiny (under 0.2 percent of consumption). As a result, unconditional *ex ante* welfare is negative since college types are not a large enough group to overcome the losses to the remainder of the population. It is interesting to see, however, that the welfare changes from allowing default are now reversed—the largest gains are experienced by the most educated, while the least educated suffer more. Part of the intuition for this result is that it is the best educated who face the steepest mean age-earnings profiles. Therefore, these agents would have the strongest purely intertemporal motives to borrow, absent any ambiguity. Low default costs mitigate the effect of ambiguity and allow for states in which a temporarily unlucky college-educated agent would find borrowing desirable.

Pricing is presented in Figures 15 and 16. Notice that for the low realization of e, the pricing function under ambiguity aversion is everywhere below the baseline expected utility case, but for the higher realization they switch places; ambiguity-averse agents with high income actually pose *less* of a default risk. The difference in pricing stems only from a difference in the households' willingness to default next period for a given b. Since default has a fixed cost component (Δ), households want to time their usage of default; in particular, households must balance the gains from defaulting tomorrow from those arising from waiting until additional shocks have been realized. This fact places the expectations of income in periods after tomorrow at the heart of the timing of default decisions, and here households who face ambiguity about the income process act quite differently from those in the benchmark economy.[25]

[25] The exposition is simpler if we refer to the expectations of the households facing ambiguity as coinciding with the choice of p, because the ambiguity-averse agents act *as if* those probabilities were the objective ones. Of course, if one were to ask ambiguity-averse agents about their forecasts of future income, they would use the true objective probabilities; they just do not use these probabilities for decisions. The proper phrasing of our statement "ambiguity-averse agents expect low future income" would be the more cumbersome "ambiguity-averse agents act as if they expect low future income." We

Figure 15 Pricing, Ambiguity Aversion, Low e

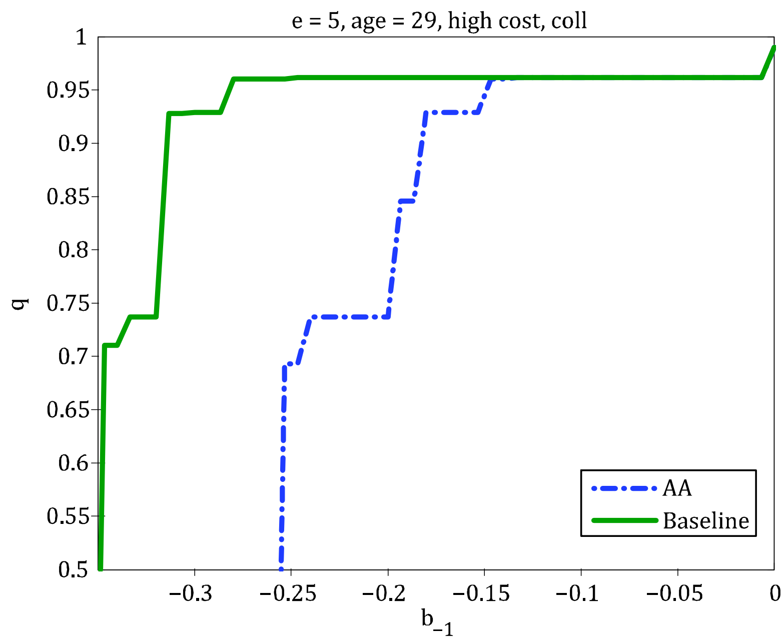

Take first the household with low e. For a "rational expectations" household, income in the distant future is expected to be better than whatever is realized tomorrow, as e is persistent but mean-reverting; for the household facing ambiguity, however, income is actually expected to be no better, or even worse, than tomorrow's realization. Since ambiguity-averse households do not think the future will be better, they may as well default next period if the realization of income is bad; lenders must therefore offer them higher interest rates to break even. In contrast, the ambiguity-averse household with higher e views a realization near the mean for next period as unexpectedly *good*, but does not expect better times in the more-distant future. Default in the next period is therefore not as valuable as waiting for a future period when those bad states are expected to occur. In contrast, without ambiguity a bad realization will induce the household to substantially

abuse the notion of expectation slightly as a result, and beg for the reader's indulgence on this matter.

Figure 16 Pricing, Ambiguity Aversion, High e

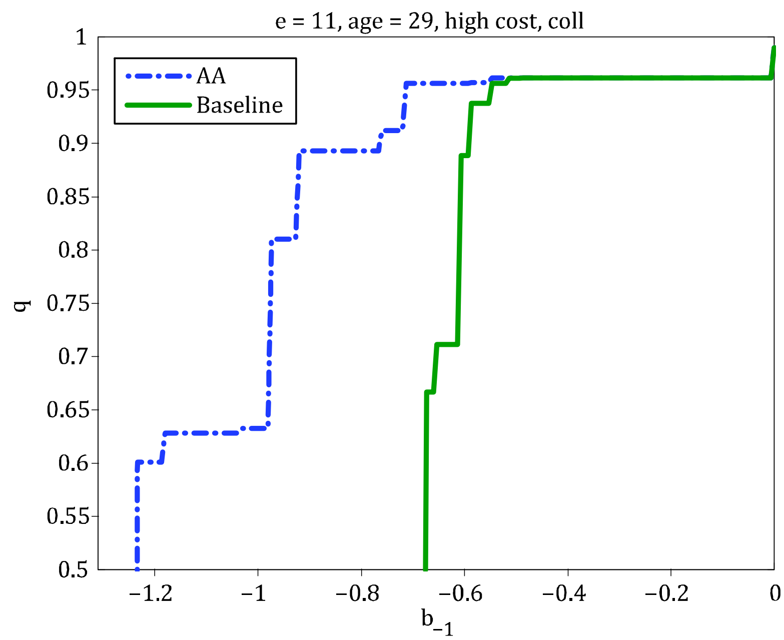

revise their future expectations downward, making default today more attractive (the decline in future income makes the fixed cost of default worth paying).[26] The result is that ambiguity-averse households with high current income obtain better terms.

Is such extreme ambiguity aversion "reasonable?" It seems highly unlikely that households entertain a stochastic process in which they receive the worst possible outcome forever with probability one as reasonable, at least not for long—after all, they need only observe the fact that their income is occasionally higher than the lower bound to discard this process empirically. As we noted above, we could introduce this learning into the model—since the households are simply learning about an exogenous process, it can be done "offline"—but it is computationally quite burdensome to condition the set of

[26] The median e has the pricing functions crossing, so that agents who face ambiguity are more likely to default on small debts but less likely to default on large ones.

permissive stochastic processes on the history of observations.[27] It is also the case that this extreme ambiguity leads to a discrepancy between model and data in terms of borrowing patterns; there is far too little debt, which lessens our interest in making this economy "more realistic." If we consider smaller limits for $\mathcal{P}$, such as 10 percent above or below the objective value, we find that default is welfare-reducing for all education levels. Thus, while ambiguity aversion provides a theoretical foundation for default options, it does not appear to provide an empirically tenable one.

3. CONCLUDING REMARKS

We have studied the efficacy of default in helping households better insure labor income risk in a large range of settings in which risk aversion, intertemporal smoothing motives, income risk, and uncertainty—and attitudes to uncertainty—over income risk itself were all varied. Our findings here suggest that within the broad class of models used thus far to develop quantitative theory for unsecured consumer credit and default, relatively generous U.S.-style default does not appear to be capable of providing protection against labor income risk.

Despite the fact that we find that labor income risk is not well hedged from the ex ante perspective, we also show that there are *ex post* beneficiaries from allowing default as it currently is; specifically, we show that the standard model generates a positive measure of agents *ex post* who would vote to introduce default. Our calibrated model predicts that these agents do not constitute a majority, though, since they are primarily college-educated middle-aged households who have been unlucky enough to still have significant debt. This result warrants further investigation since it may help explain why default penalties are becoming less stringent over time (with the exception of some aspects of the most recent reform).

Our results also suggest that "expense" shocks or catastrophic movements in net worth are likely to be essential to justify the view of default as a welfare-improving social institution. To the extent that uninsured, catastrophically large, and "involuntary" expenditures are indeed a feature of the data, a natural question is whether consumer default is the best way to deal with such events. Given the nature of resource transfers created by default and the constraints that it imposes

[27] Since this learning is not Bayesian, it can be quite difficult to write recursively, and, in any case, learning about discrete processes generally involves a large number of states. Campanale (2008) investigates non-Bayesian learning in a two-state model where the approach taken introduces only one additional state.

on the young, who disproportionately account for *both* the income-poor and uninsured, this statement seems unlikely.

With respect to future work, it is worth stressing that since expense shocks and their absence seem so important to the implications of the class of models considered here, the value of purely empirical work better documenting the nature of expense shocks, and their (a priori plausible) connection to income shocks (for example, job loss leading to insurance loss, which in turn exposes households to out of pocket expenditures), is high. Relatedly, the pivotal role played by borrowing costs "moving against" unlucky borrowers seems important to independently substantiate. In the absence of such work, it remains a possibility that the welfare findings of this article (and essentially all others) hinges too much on an institutional arrangement for borrowing that is inaccurate. Use of detailed household level credit card pricing and income information seems productive.

In addition to the preceding, in light of the findings of this article and the larger quantitative theory of consumer default, two directions seem particularly useful. First, a more "normative" approach that asks if observed default procedures can arise an optimal arrangement under plausible frictions, may yield different conclusions. One interesting example of the latter approach is the theoretical work of Grochulski (2010), where default is shown to be one method for decentralizing a constrained Pareto optimum in the presence of private information. Quantifying models with default and endogenously derived asset market structures may lead to better understanding of policy choices in this area (such as why Europe has chosen to make default available under very strict conditions, and social insurance generous, while the United States has chosen the opposite).

Second, with respect to the experiments we studied, we were led to allow for two specific preference extensions beyond CRRA expected utility in order to accurately assess the particular tradeoffs created by default. While we emphatically did not attempt to turn the article into a survey of any larger variety of non-expected utility preferences, some further extensions seem potentially important: disappointment aversion (Gul [1991] or Routledge and Zin [2010]), deviations from geometric discounting (Laibson 1997), habit formation (Constantinides 1990), and loss aversion (Barberis, Huang, and Santos 2001). Why these preferences specifically? In each case, the more general preference structure breaks the link between risk aversion and intertemporal substitution (and generally makes risk aversion state-dependent), and some (such as nongeometric discounting and loss aversion) provide arguments for government intervention; there is also extensive empirical work supporting many of them. A recent contribution to this literature is Nakajima (2012), who investigates whether the temptation preferences of Gul and Pesendorfer (2001) alter the consequences of default reform.[28] We suspect other work will follow.

APPENDIX: COMPUTATIONAL CONSIDERATIONS

We make some brief points here regarding the computation of the model. The model is burdensome to calibrate, and all programs are implemented using Fortran95 with OpenMP messaging.

In all the models we study, the objective function (the right-hand side of the Bellman equation) is not globally concave, since the discrete nature of the bankruptcy decision introduces convex segments around the point where the default option is exercised (we find that, as in Chatterjee et al. [2007], the default decision encompasses an interval and in our case it extends to $b = -\infty$ as Δ is smaller than even the worst income realization). The nonconcavity poses a problem for local optimization routines, so we approach it using a global strategy. We use linear splines to extend the value function to the real line and a golden section search to find the optimum, with some adjustments to guarantee that we bracket the global solution rather than the local one. It is straightforward to detect whether we have converged to the local maximum at any point in the state space, as the resulting price function will typically have an upward jump.

For the ambiguity aversion case we have a saddlepoint problem to solve. By the saddlepoint theorem we can do the maximization and minimization in any order; the minimization (conditional on b and d) is a linear program that we solve using a standard simplex method conditional on some b (as in Routledge and Zin [2009]). We then nest this minimization within our golden section search, again with adjustments to deal with the presence of the local maximum. For our model, this linear program turns out to be extremely simple to solve—the household puts as much weight as allowed on the worst possible outcome, then as much weight as allowed on the next worst, and so on.

To impose boundedness on the realizations of income, we approximate both e and ν by Markov chains using the approach in Flodén (2008). Having income be bounded above is convenient since it implies

[28] Nakajima (2009) finds that increasing borrowing constraints in a model with quasi-geometric discounting is not always welfare-improving, similar to Obiols-Homs (2011).

Figure 17 Optimal Choice of b given q

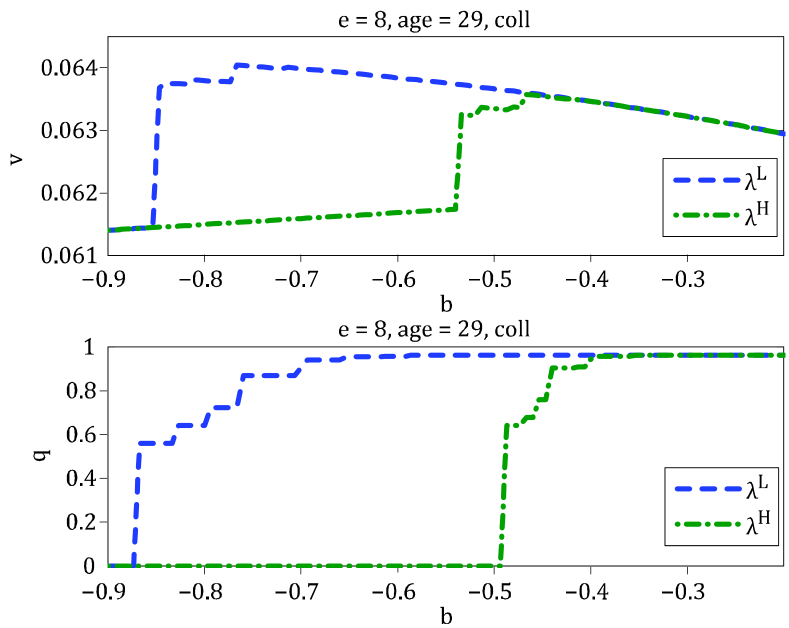

that there always exists a cost of default Δ such that bankruptcy is completely eliminated because it becomes infeasible. Quite naturally, bankruptcy is also likely not to occur when Δ is high enough even if filing is feasible for some types; in general, households with high income are not interested in the default option in our model.[29]

Figure 17 shows a typical objective function for a household in our benchmark case (expected utility with $\sigma = \rho^{-1} = 2$). The objective function has three distinct segments. The first segment is at the far right, where the values for both the low- and high-cost types coincide. In this region, default is suboptimal because borrowing either does not or barely exceeds Δ. The second segment is at the other end, where $q(b) = 0$; although impossible to see in the picture, the low-cost de-

[29] Households with high income realizations do not want to pay the stigma cost (which is proportionally higher for them) even if they are currently carrying a large amount of debt (which is very rare due to persistence). Thus, our model does not predict any "strategic" default, which can arise in models that rely on exclusion as a punishment for bankruptcy.

fault experiences slightly more utility in this region since default is less painful. The action is all in the middle segment. For this particular individual, the high-cost type (λ_L) borrows significantly more than the low-cost type; this extra borrowing reflects primarily the pricing function (as seen in the lower panel) and not any particular desire to borrow. High-cost types have more implicit collateral and are less likely to default at any given debt level, so they face lower interest rates. As a result, high-type borrowers today who become low-type borrowers tomorrow are a main source of default in our model—they both have debts and are not particularly averse to disposing of those debts through the legal system. Since type is persistent, low-type borrowers today will not generally make the same choice—the supply side of their credit market will contract.

REFERENCES

Athreya, Kartik. 2002. "Welfare Implications of the Bankruptcy Reform Act of 1999." *Journal of Monetary Economics* 49 (November): 1,567–95.

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

Athreya, Kartik. 2008. "Default, Insurance, and Debt over the Life-Cycle." *Journal of Monetary Economics* 55 (May): 752–74.

Athreya, Kartik, Juan M. Sánchez, Xuan Tam, and Eric R. Young. 2013. "Bankruptcy and Delinquency in a Model of Unsecured Debt." Federal Reserve Bank of St. Louis Working Paper 13-2.

Athreya, Kartik, Xuan S. Tam, and Eric R. Young. 2009. "Unsecured Credit Markets Are Not Insurance Markets." *Journal of Monetary Economics* 56 (January): 83–103.

Athreya, Kartik, Xuan S. Tam, and Eric R. Young. 2012. "A Quantitative Theory of Information and Unsecured Credit." *American Economic Journal: Macroeconomics* 4 (July): 153–83.

Ausubel, Lawrence M., and Amanda E. Dawsey. 2008. "Penalty Interest Rates, Universal Default, and the Common Pool Problem of Credit Card Debt." Unpublished manuscript, University of Maryland and University of Montana.

Barberis, Nicholas, Ming Huang, and Tano Santos. 2001. "Prospect Theory and Asset Prices." *Quarterly Journal of Economics* 116 (February): 1–53.

Bewley, Truman F. 2002. "Knightian Decision Theory, Part I." *Decisions in Economics and Finance* 25 (November): 79–110.

Campanale, Claudio. 2008. "Learning, Ambiguity, and Life-Cycle Portfolio Allocation." Mimeo, Universidad de Alicante.

Carroll, Christopher D. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *The Quarterly Journal of Economics* 112 (February): 1–55.

Chatterjee, Satyajit, P. Dean Corbae, Makoto Nakajima, and José-Víctor Ríos-Rull. 2007. "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default." *Econometrica* 75 (November): 1,525–90.

Constantinides, George M. 1990. "Habit Formation: A Resolution of the Equity Premium Puzzle." *Journal of Political Economy* 98 (June): 519–43.

Deaton, Angus. 1992. *Understanding Consumption*. New York: Oxford University Press.

Dubey, Pradeep, John Geanakoplos, and Martin Shubik. 2005. "Default and Punishment in General Equilibrium." *Econometrica* 73 (January): 1–37.

Epstein, Larry G., and Martin Schneider. 2003. "Recursive Multiple-Priors." *Journal of Economic Theory* 113 (November): 1–31.

Epstein, Larry G., and Stanley E. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption Growth and Asset Returns I: A Theoretical Framework." *Econometrica* 57 (July): 937–69.

Evans, David S., and Richard L. Schmalensee. 2005. *Paying with Plastic: The Digital Revolution in Buying and Borrowing.* Cambridge, Mass.: MIT Press.

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

Flodén, Martin. 2008. "A Note on the Accuracy of Markov-Chain Approximations to Highly Persistent AR(1) Processes." *Economics Letters* 99 (June): 516–20.

Friedman, Milton. 1957. *A Theory of the Consumption Function.* Princeton, N. J.: Princeton University Press.

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

Grochulski, Borys. 2010. "Optimal Personal Bankruptcy Design under Moral Hazard." *Review of Economic Dynamics* 13 (April): 350–78.

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

Gul, Faruk. 1991. "A Theory of Disappointment Aversion." *Econometrica* 59 (May): 667–86.

Gul, Faruk, and Wolfgang Pesendorfer. 2001. "Temptation and Self-Control." *Econometrica* 69 (November): 1,403–35.

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

Guvenen, Fatih, and Anthony A. Smith, Jr. 2009. "Inferring Labor Income Risk from Economic Choices: An Indirect Inference Approach." Mimeo, University of Minnesota and Yale University.

Hansen, Lars Peter, and Thomas J. Sargent. 2007. *Robustness.* Princeton, N. J.: Princeton University Press.

Hryshko, Dmytro. 2008. "RIP to HIP: The Data Reject Heterogeneous Labor Income Profiles." Mimeo, University of Alberta.

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

Huggett, Mark, Gustavo Ventura, and Amir Yaron. 2011. "Sources of Lifetime Inequality." *American Economic Review* 101 (December): 2,923–54.

Jackson, Thomas H. 2001. *The Logic and Limits of Bankruptcy Law.* Hopkins, Minn.: Beard Books.

Karahan, Fatih, and Serdar Ozkan. 2009. "On the Persistence of Income Shocks over the Life Cycle: Evidence and Implications." Penn Institute for Economic Research PIER Working Paper 09-05 (December).

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2009. "Recursive Smooth Ambiguity Preferences." *Journal of Economic Theory* 144 (May): 930–76.

Krueger, Dirk, and Harald Uhlig. 2006. "Competitive Risk Sharing Contracts with One-Sided Commitment." *Journal of Monetary Economics* 53 (October): 1,661–91.

Laibson, David A. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (May): 443–77.

Li, Wenli, and Pierre-Daniel Sarte. 2006. "U.S. Consumer Bankruptcy Choice: The Importance of General Equilibrium Effects." *Journal of Monetary Economics* 53 (April): 613–31.

Livshits, Igor, James MacGee, and Michèle Tertilt. 2011. "Costly Contracts and Consumer Credit." Working Paper 17448. Cambridge, Mass.: National Bureau of Economic Research (September).

Ljungqvist, Lars, and Thomas J. Sargent. 2004. *Recursive Macroeconomic Theory.* Cambridge, Mass.: MIT Press.

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

Livshits, Igor, James MacGee, and Michèle Tertilt. 2008. "Costly Contracts and Consumer Credit." Mimeo, University of Western Ontario and Stanford University.

Lucas, Robert E., Jr. 1987. *Models of Business Cycles.* Oxford: Basil-Blackwell Ltd.

Mateos-Planas, Xavier, and Giulio Seccia. 2006. "Welfare Implications of Endogenous Credit Limits with Bankruptcy." *Journal of Economic Dynamics and Control* 30 (November): 2,081–115.

Miao, Jianjun, and Neng Wang. 2009. "Risk, Uncertainty, and Option Exercise." Manuscript, Boston University and Columbia University.

Nakajima, Makoto. 2012. "Rising Indebtedness and Temptation: A Welfare Analysis." *Quantitative Economics* 3 (July): 257–88.

Obiols-Homs, Francesc. 2011. "On Borrowing Limits and Welfare." *Review of Economic Dynamics* 14 (April): 279–94.

Pemberton, James. 1998. "Income Catastrophes and Precautionary Saving." Mimeo, University of Reading.

Routledge, Bryan R., and Stanley E. Zin. 2009. "Model Uncertainty and Liquidity." *Review of Economic Dynamics* 12 (October): 543–66.

Routledge, Bryan R., and Stanley E. Zin. 2010. "Generalized Disappointment Aversion and Asset Prices." *Journal of Finance* 65 (August): 1,303–32.

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

Tam, Xuan S. 2009. "Long-Term Contracts in Unsecured Credit Markets." Mimeo, University of Virginia.