The Persistence of Financial Distress*

Kartik Athreya† Jose Mustre-del-Rio‡ Juan M. Sánchez§

February 15, 2016

Preliminary and Incomplete

Abstract

How *persistent* is financial distress? We answer this question using data on the proximity to debt limits, household debt-income ratios, and the probability that given a past default, a household experiences repayment difficulties. We show that all of these measures indicate that household financial distress is an extremely persistent phenomenon. To what extent can standard theory, as represented by a basic incomplete-markets model in which consumers face state contingent borrowing limits, arising from default risk capture this observed persistence of financial distress? We show that the answer is "not well": None of a wide array of model variants is capable of capturing this aspect of consumer credit use. This is important, as these baseline models have informed policy discussions on how best to provide debt relief to mitigate consumer financial distress. We then show that a plausible extension of standard approach yields a better account for the persistence of financial distress [TBC].

JEL classification: D60, E21, E44.

Keywords: default, financial distress, consumer credit, credit card debt, persistence.

*The views expressed herein are those of the authors and should not be attributed to the FRB of St. Louis, Richmond or the Federal Reserve System.

†Federal Reserve Bank of Richmond; e-mail: kartik.athreya@rich.frb.org

‡Federal Reserve Bank of Kansas City; e-mail: jose.mustre-del-rio@kc.frb.org

§Federal Reserve Bank of St. Louis; e-mail: juan.m.sanchez78@gmail.com.

1 Introduction

How *persistent* is financial distress, and can it be accounted for with standard theory? The goal of this paper is to provide both evidence and theory regarding the evolution of household financial circumstances. We approach this question empirically by using several definitions that capture aspects of what might be regarded as financial distress: The proximity to debt limits, household debt-income ratios, and the probability that given a past default, a household experiences repayment difficulties. We then employ a standard model in which consumers face state-contingent borrowing limits to understand the extent to which theory can capture the persistence of financial distress.

Our main findings are twofold: First, as an empirical matter, all measures that we examine indicate that household financial distress is extremely persistent. Second, standard models are completely incapable of capturing this fact. We show that model can be constructed that produce either too much or too little persistent, but cannot closely approximate what we find in the data.

Our work contributes in two ways. First, to our knowledge, we are the first to focus on the empirical dynamics of consumer financial distress, which we broadly define to be situations in which the household remains susceptible to any deviation of income from its ex-ante expectation. In this sense, our measures are informed by the line of work emphasizing household insurance, particularly Kaplan and Violante (2010), and the "insurance coefficient" approach of Blundell, Pistaferri, and Preston (2008). Our emphasis, relative to the preceding work, is on direct measures of financial condition that have empirical counterparts.

Second, our work extracts a variety of previously unknown implications from the "standard default model." This class of environments has received significant attention in recent research: Adapting the classic work of Eaton and Gersovitz (1981) to the consumer context, work of Athreya (2002), Livshits, MacGee, and Tertilt (2007), and Chatterjee, Corbae, Nakajima, and Rios-Rull (2007) among others embedded the ability to default on loans in a setting where such decisions provide insurance, ex-post. We focus on the model's implications for the persistence of each empirical measure of financial distress, and show that a variety of variations on this model all fail to account for the data. This is important because this model class has been used to make statements on the desirability of changes to the stringency of consumer default regulation (such as personal bankruptcy law), and about the relevance of private information on consumer credit pricing and borrowing. The gross inability to match what we argue are reasonable metrics of financial difficulty that we uncover calls into question the policy implications of existing consumer debt models.

A consistent finding in work that uses the standard model is that debt relief, defined as the ability–through either formal or informal means–to not repay loans, makes credit both expensive, and so sensitive to borrower circumstances that the overall ability to smooth consumption is substantially worsened. A caveat is that once households are subject to the risk of sudden large shocks that demand that they consume, or spend, restore the ability of default to provide net benefits in an ex-ante sense. Such shocks are routinely modeled as extreme, i.i.d, events, and are motivated by appeal to sudden medical expenditures, legal judgements, or other such circumstances.

One motivation for our work, particularly our empirical efforts, is that the extreme events

of bankruptcy or outright repudiation play an important role in helping discipline the quantitative strength of limited commitment on allocations. Specifically, it is the observable event of personal bankruptcy that provides a main target for the calibration or parameterization of the models. Based on models that match salient observations on both consumer borrowing and default (especially bankruptcy), a variety of recent work has analyzed the implications of regulations (especially bankruptcy law) on outcomes. For example, recent reforms like BAPCPA, or the effects of competing social insurance policies on credit use have been studied through the lens of what is now a "standard default model" (e.g., Livshits, MacGee, and Tertilt (2007), Chatterjee, Corbae, Nakajima, and Rios-Rull (2007)). But as noted above, absent clear evidence that the baseline models used in these analyses captures well the time path of financial conditions, there is reason for concern.

The remainder of the paper is organized as follows. Section 2 defines our preferred measures of financial distress, and provides an empirical analysis of the behavior of these measures. Section 3 then lays out a battery of variants of a standard life-cycle model of consumption and savings. Section 4 provides the main comparisons of models with data, and illustrates the inability of any of these settings to capture what would appear to be critical features of observed financial distress. Section 5 studies potential variations that, while not standard, hold out promise for reconciling theory with data. Section 6 concludes.

2 The Empirics of Financial Distress

We use several data sets to characterize the persistence of financial distress. The first data source we employ is the Federal Reserve Bank of New York Consumer Credit Panel/Equifax. In particular, we use these data to construct two specific measures of *financial distress*, which we name *Default*, and *HLIM* (for "high credit usage relative to a person's credit limit"), respectively. The definitions are as follows:

- *Default*: An individual is in financial distress is he/she has any accounts categorized as 'severely derogatory'.
- *HLIM*: An individual is in financial distress if he/she is in the top 10% of the credit usage/credit limit distribution, where the distribution is defined over all people ages 25-65.

The results are presented in Figure 1.

The preceding definitions rely either on the realization of default events or the presence of credit usage that is relatively extreme in the population. However, neither makes direct use of wealth (or debt), or income as separate sources of information. To accomodate this, We also use data from the Survey of Income and Program Participation (SIPP) and the Panel Study of Income Dynamics (PSID) to consider another measure of financial distress. For this, we create the mty-ratio as in Kaplan, Violante, and Weidner (2014) measure of net liquid wealth/income. A person is under *mty distress=1* if her net liquid wealth/income is at or below the 10th percentile of the unconditional distribution of net liquid wealth/income. The results are presented in Figure 2.

Figure 1: Persistence of Financial Distress, default and HLIM

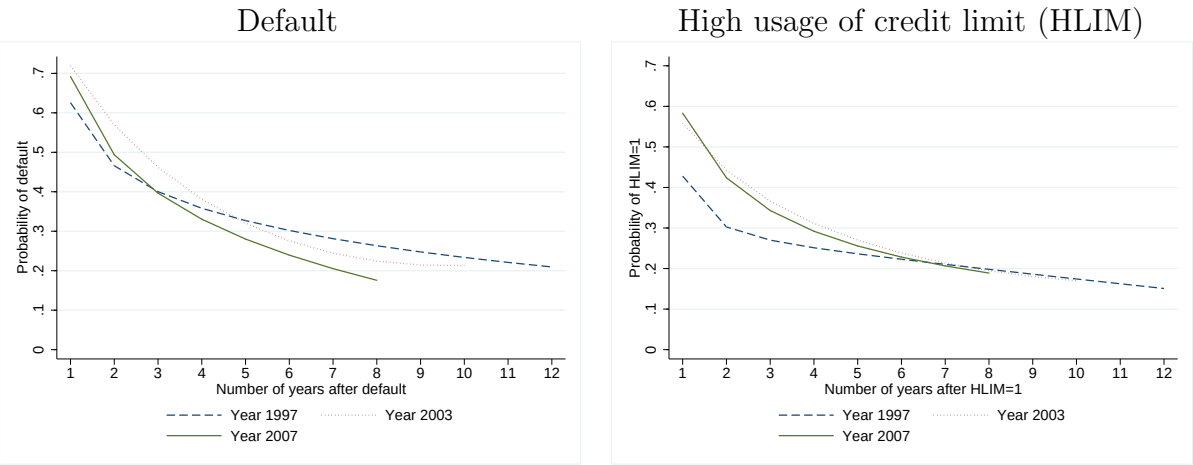
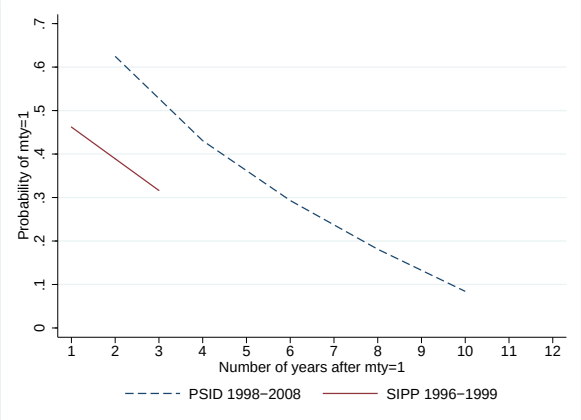


Figure 2: Persistence of Financial Distress, net liquid assets



Empirically, the behavior of all three notions sends a clear message: Financial distress is a persistent phenomenon. We now approach the second question: To what extent can we explain or account for these observations?

3 A General Model of Financial Distress

Environment. We study a life cycle Standard Incomplete Market (SIM) model close to that of Kaplan and Violante (2010). As they do, we model the choices of a household that lives up to T periods and works until age $W \leq T$. At any time a unit mass of households is born.

Households differ in their subjective discount factor β_j; denote the fraction of type-j households as p_j with $\sum p_j = 1$. Households survive to the next period with probability ϱ_n, which depends on age n. Preferences of households are given by the expected value of the discounted sum of momentary utility

$$\mathbb{E}_0 \left[\sum_{n=1}^{T} (\beta_j \varrho_n)^n u(c_n) \right],$$

where c_n is consumption at age n. The utility function u is strictly increasing, strictly concave, and twice differentiable.

At any period t a household i receives income y_t^i, which is independent of its preference type j. During working age, income has a fixed effect, a persistent component, a life cycle component, and an i.i.d component:

$$\log(y_t^i) = \alpha^i + l(n) + \varepsilon_t^i + z_t^i,$$

where α^i denotes the fixed effect, $l(n)$ denotes the life cycle component, ε_t^i is a transitory component, and z_t^i is a permanent component that follows a zero mean AR(1) process:

$$z_t^i = \rho_z z_{t-1}^i + e_t^i.$$

We assume α^i, ε_t^i, and e_t^i are normally distributed with variances σ_α^2, σ_ε^2, and σ_e^2, respectively. After retirement, the household receives a percentage of the last realization of the permanent component of its working-age income.

There are also expense shocks, x, which follow an AR(1) process with mean μ_x:

$$x_t^i = \mu_x + \rho_x x_{t-1}^i + \eta_t^i.$$

We assume η_t^i is normally distributed with variance σ_η^2.

Lenders can borrow at the risk-free interest rate r from the rest of the world. The price charged to a type (i, j) household of age n is $q_{i,j,n}(z_t, x_t, a_{t+1})$. This price depends on a_{t+1} because it determines the debt the household will have to pay back next period, which in turn affects her willingness to pay back the debt. It depends on x_t and z_t because these shocks affect next period's expenses and income, and thereby the probability of bankruptcy.

Household's problem. Hereafter, period-t variables will be expressed without any subscripts or superscripts, and period-$t+1$ variables will be represented with superscripts '′'. Households decide on consumption, c, and asset accumulation, a'. In addition, they decide whether to file for bankruptcy or to pay back debt.

Several assumptions determine the advantages and disadvantages of bankruptcy. The key advantage is the discharge of expenses and debts: current period expense obligations are

eliminated and in the period after bankruptcy, debt is set at zero. Thus, a household with too much debt or too large an expense shock may find it beneficial to file for bankruptcy. There are two disadvantages of doing so, however. In the period of bankruptcy, a proportion of income, τ, is lost.[1] Additionally, in that period, consumption equals income—neither saving nor borrowing are allowed. Finally, we also incorporate a utility cost of bankruptcy κ, which captures for example the time and psychic costs associated with filing.

In this environment, lifetime utility can be written as

$$G_{i,j,n}(z,x,\varepsilon,a) = \max\{\underbrace{V_{i,j,n}(z,x,\varepsilon,a)}_{\text{Pay}}, \underbrace{D_{i,j,n}(z,x,\varepsilon)}_{\text{Default}}\} \tag{1}$$

where V and D (defined below) are lifetime utilities for households paying back the debt and filing bankruptcy, respectively. This means that a household has the choice of filing bankruptcy. The policy function R indicates whether the household pays back the debt (repay) or not,

$$R_{i,j,n}(z,x,\varepsilon,a) = \begin{cases} 1 \text{ if } V_{i,j,n}(z,x,\varepsilon,a) \geq D_{i,j,n}(z,x,\varepsilon), \\ 0 \text{ otherwise.} \end{cases}$$

Suppose the household chooses to file for bankruptcy. Then, lifetime utility is

$$D_{i,j,n}(z,x,\varepsilon) = u(y_{i,n}(z,\varepsilon)(1-\tau)) - \kappa + \varrho\beta_j \mathbb{E}\left[G_{i,j,n+1}(z',x',\varepsilon',0)|z,x\right]. \tag{2}$$

The household's consumption equals net income (labor income minus the proportion lost due to bankruptcy). In the period after bankruptcy, the household will have no debt.

Now suppose the household decides to pay back the debt. Then she faces the debt price $q_{i,j,n}(z,x,a')$ and her lifetime utility is

$$\begin{aligned} &V_{i,j,n}(z,x,\varepsilon,a) = \max_{\{a',c\}} u(c) + \varrho\beta_j \mathbb{E}\left[G_{i,j,n+1}(z',x',\varepsilon',a')|z,x\right], \\ &\text{subject to} \\ &\qquad c + a' q_{i,j,n}(z,x,a') = a + y_{i,n}(z,\varepsilon) - x, \\ &\qquad c \geq 0. \end{aligned} \tag{3}$$

Debt prices. How are prices determined? They must imply zero-expected profits. In general, a price function $q_{i,j,n}(z,x,a')$ implies zero profits if the following equation is satisfied.

$$q_{i,j,n}(z,x,a') = \frac{1}{1+r}\varrho\mathbb{E}\left[R_{i,j,n}(z',x',\varepsilon',a')|z,x\right]. \tag{4}$$

Looking at this equation it is very clear why the price function (or interest rates) depends on (a',z,x). It depend on a' because it affects the bankruptcy decision, R, in each possible state. It depends on z because it determines the transition probability to each z' and therefore next period's earned income, y. Lastly, it depends on x because it determines the transition probability to next period's expenditure shock z', which in turn determines tomorrow's disposable income.

[1] Chatterjee, Corbae, and Rios-Rull (2008) build a model where no punishment is required after filing bankruptcy. There, asymmetric information is crucial to create incentives for debt repayment, because households signal their type by paying back their debt.

Equilibrium. An equilibrium in this economy is a set of value functions, optimal decision rules for the consumer, default probabilities, and bond prices, such that equations (1) to (3) are satisfied and prices satisfy the zero-profit condition (4).

4 Accounting for Financial Distress? Models vs. Data

4.1 Calibration

We consider three variants of the general framework described in the previous section. The *baseline model* assumes individuals do not differ in their subjective discount factor ($\beta_j = \beta$ for all j), and are not subject to persistent expenditure shocks x. The *expenditure model* assumes individuals do not differ in their subjective discount factor ($\beta_j = \beta$ for all j), but are subject to persistent expenditure shocks x. Lastly, the *beta-het model* assumes individuals differ in their subjective discount factor, but are not subject to persistent expenditure shocks x.

Common parameters across models. As in Kaplan and Violante (2010), a period in the model refers to a year; households enter the model at age 25, retire at age 60, and die no later than at age 82. Survival rates are obtained from Kaplan and Violante (2010). With a retirement income replacement ratio of 75 percent, we replicate the mean income after retirement in the data. The parameters $\sigma_e, \sigma_\varepsilon$ and the life cycle component of the income process are calibrated following Kaplan and Violante (2010). As in Storesletten, Telmer, and Yaron (2004), the fixed effect takes two values. Table 1 presents these values.

Table 1: Parameters common across models

Parameter	Value
σ_α, Standard deviation of permanent component	0.205
σ_e, Standard deviation of persistent shocks	0.089
ρ_z, Autocorrelation of persistent component	0.980
σ_ϵ, Standard deviation of transitory component	0.242
σ, Coefficient of relative risk aversion	2.000
r, Risk free interest rate	3.000%

Baseline model. For this model the parameters to be determined are the subjective discount factor, the fraction of earnings lost when households default, and the utility cost associated with default (β, τ, κ). These three parameters are jointly calibrated to match the average default rate, debt/income ratio, and the fraction of wealth held by the second quintile of net worth. We take values for the first two targets from Livshits, MacGee, and Tertilt (2007). However, similar to Chatterjee, Corbae, and Rios-Rull (2007) we interpret the reasons for default in this baseline model as stemming only from job loss or credit misuse and hence divide the default rate by a half. Information on net worth is taken the Survey of Consumer Finances (SCF) as reported in Chatterjee, Corbae, and Rios-Rull (2007). Table 2 summarizes this calibration.

Table 2: Statistics and Parameter Values, Baseline Model

Statistic	Target	Model	Parameter	Value
prorated default rate (%)	0.42	0.39	β	0.920
debt/income (%)	8.40	0.43	τ	110.306
% of wealth in 2nd quintile	1.3	1.3	κ	0.019

Expenditure model. This model has as additional parameters those characterizing the expenditure process $(\mu_x, \rho_x, \sigma_\eta^2)$. We take values for (ρ_x, σ_η^2) from the estimated expenditure process in Hubbard, Skinner, and Zeldes (1994). Next we jointly calibrate $(\mu_x, \beta, \tau, \kappa)$ to match the average default rate, debt/income ratio, and the fraction of wealth held by the second and third quintiles of net worth. Because of the inclusion of expenditure shocks we target all bankruptcies as reported in Livshits, MacGee, and Tertilt (2007). Table 3 summarizes this calibration.

Table 3: Statistics and Parameter Values, Expenditure Model

Statistic	Target	Model	Parameter	Value
default rate (%)	0.84	0.84	β	0.922
debt/income (%)	8.40	1.11	τ	86.135
% of wealth in 2nd quintile	1.3	0.94	κ	0.044
% of wealth in 3rd quintile	5.0	7.67	μ_x	-2.552
Hubbard, Skinner, and Zeldes (1994)			ρ_x	0.9
Hubbard, Skinner, and Zeldes (1994)			σ_ν	0.15

Beta-het model. For this model we calibrate (τ, κ, β_l, β_h , p_l) to match the average default rate, debt/income ratio, and the second, third and fourth quintiles of net worth. Table 4 presents this calibration.

Table 4: Statistics and Parameter Values, Beta-het Model

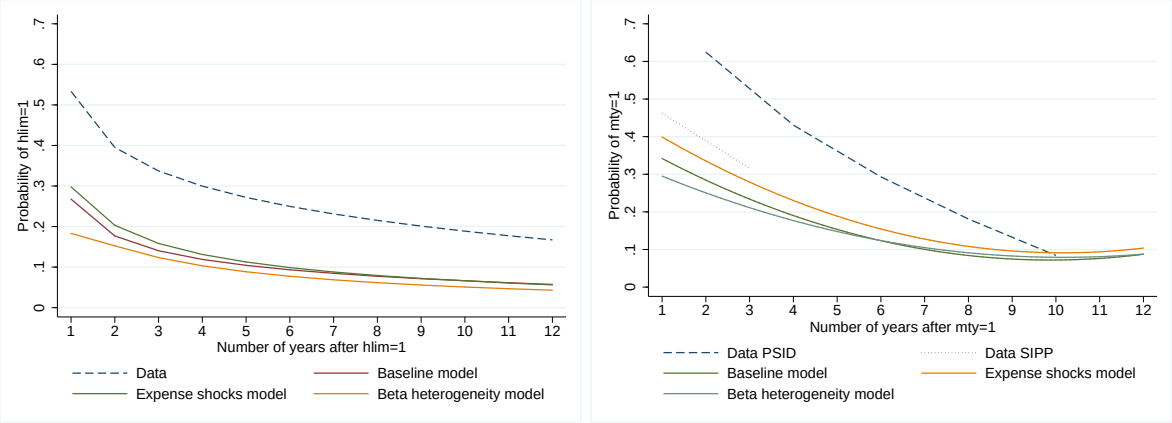
Statistic	Target	Model	Parameter	Value
default rate (%)	0.84	0.86	τ	112.334
debt/income (%)	8.40	0.06	κ	0.003
% of wealth in 2nd quintile	1.3	0.92	β_l	0.897
% of wealth in 3rd quintile	5.0	5.0	β_h	0.999
% of wealth in 4th quintile	12.2	18.4	p_l	0.858

4.2 Persistence of financial distress

In the model, borrowing limits are calculated by looking at the distribution of borrowing by *(age ,permanent income type, persistent income shock, persistent expenditure shock, discount*

factor type) and taking the borrowing limit as the bottom 10th percentile of this distribution. Additionally, in the model usage-percentiles are calculated using all ages. We should adjust Equifax accordingly.

Figure 3: Persistence of Financial Distress, Data vs. Model



4.3 Life-cycle profile of persistence of financial distress

Figure 4: Life-cycle profile of persistence, HLIM

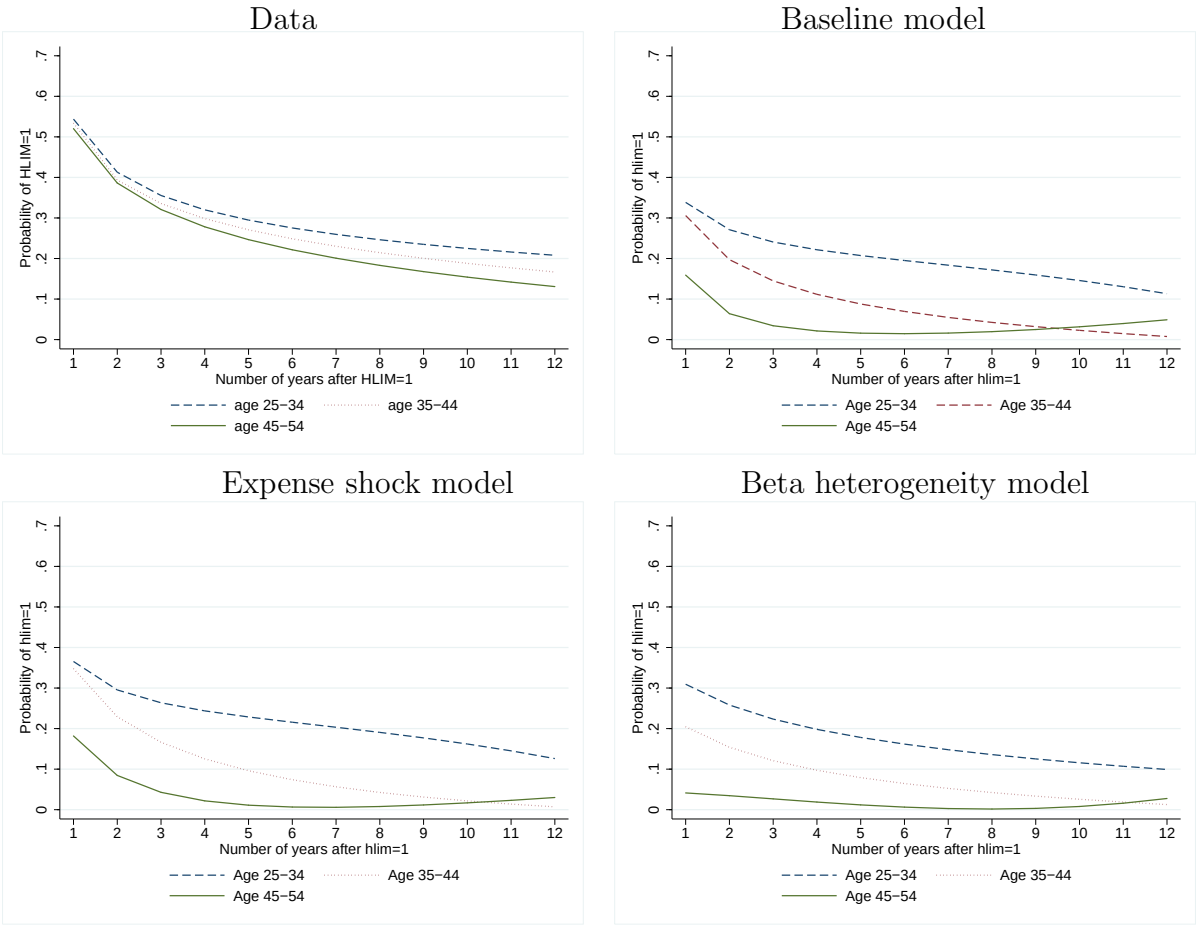
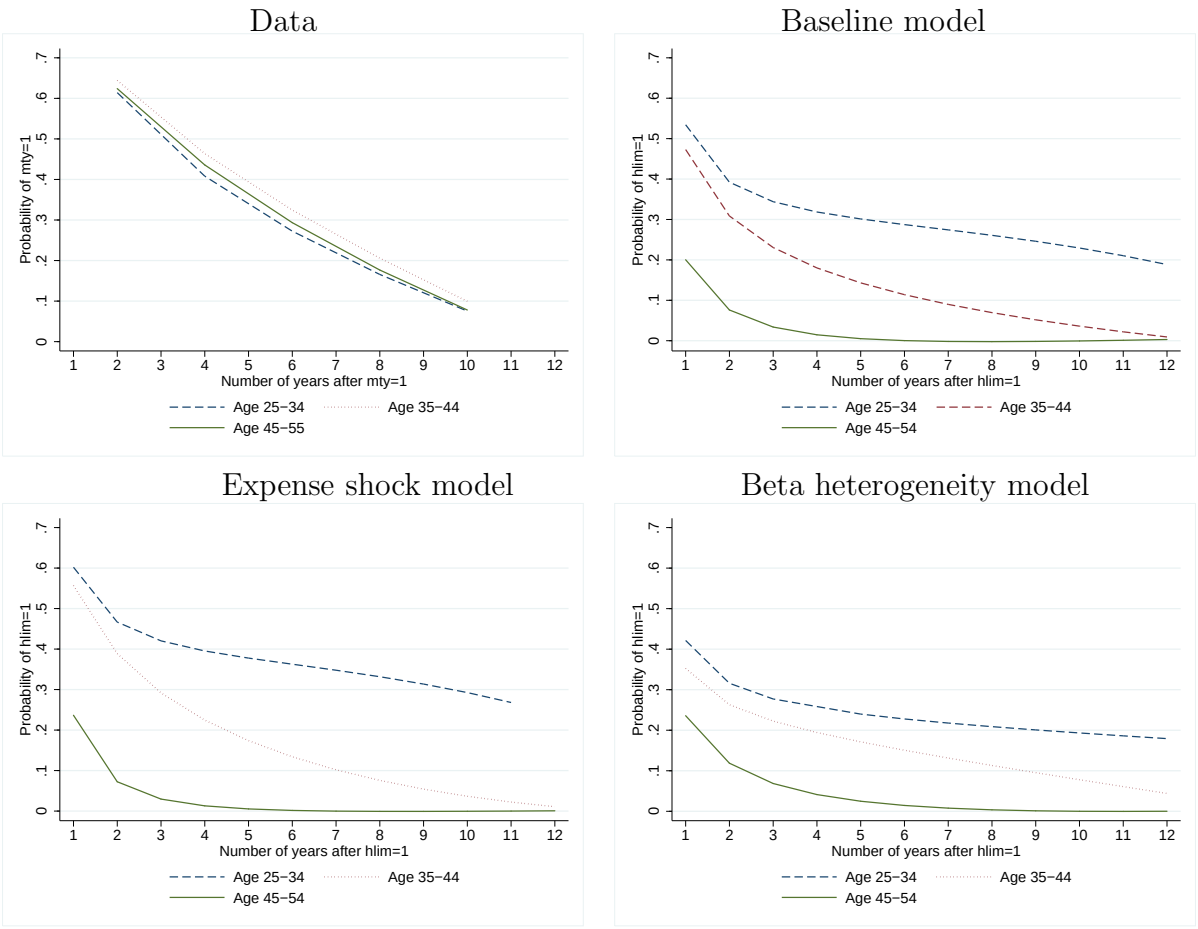


Figure 5: Life-cycle profile of persistence, MTY



5 Alternative Models

Here we consider Guvenen (2007) and explore a model in which households are learning their earning.

6 Concluding Remarks

References

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

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887–1921.

Chatterjee, S., D. Corbae, M. Nakajima, and J. V. Rios-Rull (2007): "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica*, 75(6), 1525–1591.

Chatterjee, S., D. Corbae, and J. V. Rios-Rull (2007): "A Theory of Credit Scoring and Competitive Pricing of Default Risk," Manuscript.

——— (2008): "A Finite-Life Private-Information Theory of Unsecured Consumer Debt," *Journal of Economic Theory*, 142, 149–177.

Eaton, J., and M. Gersovitz (1981): "Debt with Potential Repudiation: Theoretical and Empirical Analysis," *Review of Economic Studies*, 48, 289–309.

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

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1994): "The importance of precautionary motives in explaining individual and aggregate saving," *Carnegie-Rochester Conference Series on Public Policy*, 40(1), 59–125.

Kaplan, G., G. Violante, and J. Weidner (2014): "The Wealthy-Hand-to-Mouth," *Brookings Papers on Economic Activity*, Spring, 2051–2079.

Kaplan, G., and G. L. Violante (2010): "How Much Consumption Insurance Beyond Self-Insurance?," *American Economic Journal: Macroeconomics*, 2, 53–87.

Livshits, I., J. MacGee, and M. Tertilt (2007): "Consumer Bankruptcy: A Fresh Start," *American Economic Review*, 97(1), 402–418.

Storesletten, K., C. I. Telmer, and A. Yaron (2004): "Consumption and Risk Sharing over the Life Cycle," *Journal of Monetary Economics*, 51, 609–633.

7 Appendix

7.1 Constructing net liquid wealth-to-income ratios

We use the Survey of Income and Program Participation (SIPP) to measure the persistence of net liquid wealth to labor income across households. The SIPP is a survey of about 20,000 households designed to be representative of the U.S. population. Households are interviewed every four months. At each interview, information on work experience is collected for the three preceding as well as most recent month. Once a year households are also asked a special set of questions regarding assets and liabilities. Our analysis uses the 1996 panel of the SIPP which covers 1996 through 1999 allowing us to observe up to four wealth measurements per household.

We restrict our sample to households with complete wealth observations. Following Kaplan et al (2014), we restrict the sample to households whose head is between the ages of 22 and 79 and where annual income is positive. Household income includes earnings from all members, household-level means-tested cash transfers, and other income (which captures other government transfers). For consistency purposes, we aggregate these measures to an annual frequency whenever necessary.

Liquid wealth captures individual and joint checking and savings accounts (interest bearing or not), interest earning assets at other institutions, equity in stocks and mutual fund shares, and savings bonds. Liquid debt is own and joint credit card debt.

Illiquid wealth includes home equity (and principal owed), equity from other real estate, IRAs and thrift savings accounts, vehicle equity, and the current value of life insurance policies. Illiquid debt includes all mortgages and other debts and loans (medical bills, educational loans, etc.)

In general, the wealth aggregates from SIPP will not match other datasets like the Panel Study of Income Dynamics (PSID), or the Survey of Consumer Finances (SCF) because the SIPP does not record assets like equity in defined-benefit pension plans. Additionally, the SIPP does not over-sample the very wealthy.