Human Capital Risk, Contract Enforcement, and the Macroeconomy*

Tom Krebs

University of Mannheim†

Moritz Kuhn

University of Bonn

Mark L. J. Wright

FRB Chicago, UCLA, and NBER

December 2011

Abstract

We develop a macroeconomic model with physical and human capital, human capital risk, and limited contract enforcement. We show analytically that young (high-return) households are the most exposed to human capital risk and are also the least insured. We document this risk-insurance pattern in data on life-insurance drawn from the Survey of Consumer Finance. A calibrated version of the model can quantitatively account for the life-cycle variation of insurance observed in the US data and implies welfare costs of under-insurance for young households that are equivalent to a 4 percent reduction in lifetime consumption. A policy reform that makes consumer bankruptcy more costly leads to a substantial increase in the volume of credit and insurance.

Keywords: Human Capital Risk, Limited Enforcement, Insurance

JEL Codes: E21, E24, D52, J24

---

*We thank seminar participants at various institutions and conferences for useful comments. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

†Department of Economics, University of Mannheim, 68313 Mannheim, Germany. E-mail: tkrebs@econ.uni-mannheim.de. Moritz Kuhn: Department of Economics, Adenauer Alle 24-42, 53113 Bonn, Germany. Mark Wright: 230 S. LaSalle Street, Chicago, IL 60604.

1. Introduction

For many households, human capital constitutes the single most important component of their wealth. Empirical evidence suggests that human capital is distinguished by three characteristics.[1] First, ex ante returns to human capital investment vary greatly across the population. Second, human capital investment is very risky due to uncertainty about lifespan, health status, and labor market conditions. Third, human capital cannot be pledged as collateral. In this paper, we explore the macroeconomic implications of these special characteristics of human capital using a combination of theoretical, quantitative, and empirical methods. We emphasize five main findings.

First, we show theoretically that these three properties of human capital generate an interesting form of under-insurance, with the households that are most exposed to human capital risk also the least insured relative to their insurance needs. Second, we provide microeconomic evidence in support of this prediction by examining data on insurance against an important form of human capital risk: the death of a household member. Specifically, we use data on life-insurance contracts drawn from the Survey of Consumer Finance (SCF) and show that the extent of under-insurance is lowest for young households who also have the greatest share of their wealth invested in human capital. Third, we show that when we calibrate our model to life-cycle patterns in human capital returns, our model replicates the observed quantitative pattern of under-insurance. Fourth, we show that this under-insurance is important for welfare, with limited access to insurance against human capital risk reducing the welfare of young households by an amount equivalent to a 4 percent reduction in lifetime consumption. Fifth, we show that this under-insurance has important implications for policy: making consumer bankruptcy as costly as defaulting on student loans leads to a substantial increase in the volume of credit and insurance.

We begin our analysis by developing a macroeconomic model in which human capital

[1] The third characteristic is obvious. Section 6 discusses the empirical evidence on human capital risk and heterogeneity of human capital returns.

investments are risky, earn heterogeneous expected returns, and are not pledgeable as collateral. Households can buy insurance and borrow using unsecured debt, with their ability to borrow limited endogenously by the possibility that they might default. Default is modeled along the lines of Chapter 7 of the US bankruptcy code: in the case of default all debt is cancelled, all financial assets are seized, no future earnings are garnished, and access to financial markets is restricted for a period of time.

In a first step, we consider a simplified version of our model and show analytically that the equilibrium exhibits a negative relationship between risk and insurance: those households who are most exposed to human capital risk are the ones who have the least insurance in equilibrium. We show that this result holds for two different insurance measures, one that is defined as the ratio of insurance pay-out relative to the income loss and another that measures the reduction in consumption volatility due to insurance. Intuitively, households with high ex-ante human capital returns choose to invest the bulk of their wealth in human capital, that is, they are heavily exposed to human capital risk. In the absence of borrowing constraints, these households would like to borrow in order to invest even more in human capital and to buy insurance against human capital risk. However, the risk of bankruptcy, combined with the fact that these households mainly hold non-collaterizable assets, prevents them from borrowing and leaves them with little insurance in equilibrium.

In a next step, we turn to the quantitative analysis and study a calibrated version of the full model in which ex-ante returns to investment in human capital vary by age: younger households have higher human capital returns than older households. When we calibrate the model economy to match the US evidence on labor market and mortality risk and the life-cycle profile of median earnings growth, we find that in equilibrium a large number of younger households are severely under-insured, but older households are almost completely insured. Indeed, young households would gain around 4 percent of lifetime consumption from being fully insured, but they would have to borrow to achieve this welfare improvement. In short, the inability to pledge human capital generates large deviations from the full-insurance outcome for a large group of households.

Our analysis predicts that both human capital investment and the degree of under-insurance should decrease over the life-cycle. In this paper, we examine these two implications empirically using micro-level data drawn from the Survey of Consumer Finance (SCF). Specifically, we first measure the degree of under-insurance against an important form of human capital risks faced by a household: the death of a household member.[2] When we measure under-insurance using data on life-insurance purchases to estimate insurance pay-outs relative to our estimate of the present value of lost earnings of the household, we find that under-insurance is strongly decreasing in age. We also show that this result is robust to different sample selections and modifications in the definition of under-insurance. Second, when we measure the human capital choice of households in the data by computing the ratio of net financial wealth over labor income, we also find that the fraction of total wealth invested in human capital strongly decreases over the life-cycle. Moreover, the magnitude of the decline in under-insurance and human capital over the life-cycle is in line with predictions of our calibrated model economy. This provides additional corroboration of the theory since the model has not been calibrated to match these two targets.

Finally, we argue that our approach has important macro-economic implications. There has been a long-standing debate among academic scholars and policy makers with regard to the relative merits of alternative consumer bankruptcy codes. In the US, this debate has led to legislation making it more costly to declare bankruptcy. In this paper, we add to this debate by exploring a channel that has not been studied by the previous macro literature on consumer bankruptcy: making it more costly to declare bankruptcy not only increases the volume of credit, but also the amount of insurance purchased by households. In the human capital model analyzed here, it further increases economic growth since it leads to more

[2] One advantage of focusing on the market for life insurance is that other market imperfections, such as adverse selection, are likely to be less important. Further, pure life-insurance contracts (term life insurance) have a relatively simple structure and can in principle be purchased by most households.

investment in the high-return asset. For the calibrated version of the model, we show that these effects are quantitatively substantial, though the positive growth impact is dampened as a result of strong general equilibrium effects.

In addition to these substantive contributions, this paper also makes a methodological contribution by developing a tractable macro-economic model with risky human capital and limited contract enforcement. Both theoretical and applied work on heterogeneous-agent models with idiosyncratic risk and limited enforcement/commitment has struggled with two fundamental problems. First, the infinite-dimensional wealth distribution is in general a relevant state variable when computing recursive equilibria. Second, in models with investment, the choice of individual households is typically not convex, which calls into question the application of any first-order approach to the computation of equilibria. In contrast, for the class of models developed in this paper, we show that the maximization problem of individual households can be transformed into a convex problem and that the infinite-dimensional wealth distribution is not a relevant state variable.[3] This property allows us to show analytically our main result about risk and insurance, and is of great use in quantitative work dealing with higher-dimensional state variables.

2. Literature

This paper is most closely related to the large literature on risk sharing in models with limited commitment/enforcement. See, for example, Alvarez and Jermann (2000), Kehoe and Levine (1993), Kocherlakota (1996), Thomas and Worrall (1988) for contributions based on exchange models and Ligon, Thomas, and Worrall (2002), Kehoe and Perri (2002), Krueger and Perri (2006), and Wright (2001) for work on production models with capital. In addition to our methodological contribution, which shows how to deal with the non-convexity issue

[3] In this paper, we focus on logarithmic one-period utility functions, but it is easy to see that our characterization result holds more generally for CRRA utility functions. Indeed, our basic argument uses homotheticity of preferences, which means that the assumption of Epstein-Zin preferences is sufficient for the proof.

in a certain class of production models, we make three substantive contributions to this literature. First, we show that a calibrated macro model with physical capital and limited contract enforcement can generate substantial lack of consumption insurance once we introduce life-cycle considerations and human capital choices. In contrast, previous work in this literature, which has not considered ex-ante heterogeneity and human capital choice, has concluded that the effects of limited enforceability of contracts on risk sharing are small in calibrated macro models with physical capital and production.[4] Second, we show that our calibrated model economy provides a good quantitative account of the empirically observed life-cycle profile of human capital investment and consumption insurance. In particular, we show that the model can quantitatively explain the "under-insurance" puzzle in the life-insurance market. Finally, we introduce the human capital channel and show that our model has important implications for macro-economic policy analysis.[5]

This paper is also related the literature on macro-economic models with incomplete markets. Most work in this literature has taken the human capital of individuals as exogenous, but Krebs (2003), Guvenen, Kuruscu, and Ozkan (2011), and Huggett, Ventura, and Yaron (2011) are three contributions that have explicitly dealt with human capital investment when returns are uncertain and insurance markets are incomplete. Though these models provide useful insights into a number of important issues, they are necessarily silent about the underlying financial friction that explains the observed lack of insurance and the limits on borrowing. In particular, the standard incomplete-market model can in principle explain why self-insurance increases with age,[6] but it has nothing to say about the use of existing insurance markets over the life-cycle. Moreover, our analysis of the personal bankruptcy

[4]Krueger and Perri (2006) match the cross-sectional distribution of consumption fairly well, but the implied volatility of individual consumption growth is negligible in their model. A similar "almost full-insurance" result is obtained by Cordoba (2006).

[5]Andofatto and Gervais (2006) and Lochner and Monge (2011) analyze models with human capital investment and endogenous borrowing constraints due to enforcement problems, but they abstract from risk considerations and therefore cannot address the issues that take center stage in this paper.

[6]See, for example, Kaplan and Violante (2010).

law heavily relies on the endogeneity of borrowing constraints and the existence of some insurance markets.

Recent contributions by Chatterjee, Corbae, Nakajima, and Rios-Rull (2007), and Livshits, MacGee, and Tertilt (2007) analyze the consequences of reforming the consumer bankruptcy code based on models with equilibrium default and no insurance markets. In these papers, an increase in the cost of bankruptcy increases borrowing and reduces default, which leads to a reduction in risk sharing since default is a means towards smoothing consumption across states of nature. In contrast, in our model an increase in the cost of bankruptcy increases borrowing *and* improves risk sharing since households can take better advantage of existing insurance markets. Moreover, our quantitative work shows that the increase in equilibrium insurance is substantial. Clearly, neither our assumption of a complete set of insurance markets nor the assumption of no insurance markets is a correct representation of reality. Despite this caveat, our work makes a simple yet important point: any reform of the consumer bankruptcy law is likely to affect not only credit markets, but also insurance markets.

There is an extensive literature analyzing insurance markets based on models of adverse selection and moral hazard, and one basic implication of this approach is that households with higher risk exposure should buy more insurance (Chiappori and Salanie, 2000, and Chiappori, Jullien, Salanie, and Salanie, 2006). A number of empirical studies have found that this hypothesis is often rejected by the data (Chiappori and Salanie, 2000, and Bernheim, Forni, Gokhale, and Kotlikoff, 2003), and has dubbed this finding the "under-insurance puzzle". In this paper, we provide additional evidence supporting the findings of the previous literature and put forward an explanation of the puzzle for mortality risk in terms of limited contract enforcement. Of course, there are alternative explanations of a negative relationship between risk exposure and insurance based on adverse selection and preference heterogeneity (Chiappori et al.., 2006, and Cutler, Finkelstein, and McGarry, 2008), but we are not aware of any work in the macro literature that addresses this issue.

At a very general level, models of limited enforcement/commitment have one basic implication: in equilibrium, there is imperfect risk sharing if borrowing (short-sale) constraints are binding. We are not aware of any empirical study directly testing this joint hypothesis, but both the "perfect risk sharing hypothesis" and the "binding borrowing constraints hypothesis" have been tested separately. On risk sharing, almost all empirical studies using household level data have found that the full-insurance hypothesis is strongly rejected (Attanasio and Davis, 1996, Blundell, Pistaferri, and Peston, 2008, Cochrane, 1991, and Townsend, 1994). On borrowing constraints, Jappelli (1990) finds that a significant fraction of US households are credit constrained, and that these households are on average younger than the rest of the population. There is also an extensive empirical literature on credit constraints and college enrollment, which has reached somewhat mixed results. For example, Card (2001) concludes that borrowing constraints affect college enrollment decisions substantially, whereas Carneiro and Heckmann (2002) argue that only a small fraction of the population is affected. Lochner and Monge (2011) show that the number of affected individuals has increased significantly since the 1980s.

3. Model

In this section, we develop the model and define a stationary (balanced growth) recursive equilibrium.

3.1. Production

Time is discrete and open ended. There is no aggregate risk and we confine attention to stationary (balanced growth) equilibria. We assume that there is one all-purpose good that can be consumed, invested in physical capital, or invested in human capital. Production of this one good is undertaken by one representative firm (equivalently, a large number of identical firms) that rents physical capital and human capital in competitive markets and uses these input factors to produce output, Y, according to the aggregate production function $Y = F(K, H)$, where K and H denote the aggregate levels of physical capital and human capital, respectively. The production function, F, has constant-returns-to-scale, satisfies a Inada condition, and is continuous, concave, and strictly increasing in each argument. Given these assumptions on F, the derived intensive-form production function, $f(\tilde{K}) = F(\tilde{K}, 1)$, is continuous, strictly increasing, strictly concave, and satisfies a corresponding Inada condition, where we introduced the "capital-to-labor ratio" $\tilde{K} = K/H$. Given the assumption of perfectly competitive labor and capital markets, profit maximization implies

$$\begin{aligned} r_k &= f'(\tilde{K}) \\ r_h &= f(\tilde{K}) + f'(\tilde{K})\tilde{K} \ , \end{aligned} \tag{1}$$

where r_k is the rental rate of physical capital and r_h is the rental rate of human capital. Note that r_h is simply the wage rate per unit of human capital and that we dropped the time index because of our stationarity assumption. Clearly, (1) defines rental rates as functions of the capital to labor ratio: $r_k = r_k(\tilde{K})$ and $r_h = r_h(\tilde{K})$. Finally, physical capital depreciates at a constant rate, δ_k, so that the (risk-free) return to physical capital investment is $r_k - \delta_k$.

3.2. Households

There are a continuum of long-lived households of mass one. Households have an uncertain life-span and in the case of death they are replaced by new-born households. The exogenous state of an individual household in period t is denoted by s_t. We assume that the process of exogenous states, $\{s_t\}$, is Markov with stationary transition probabilities $\pi(s_{t+1}|s_t)$. Note that s_t can have several components (age, ability, mortality risk, labor market risk) and that we can incorporate ex-ante heterogeneity (age, ability) by assuming degenerate transition probabilities for certain components (see Sections 5 and 6 for particular applications). We denote by $s^t = (s_1, \ldots, s_t)$ the history of exogenous states up to period t (date-event, node) and let $\pi(s^t|s_0) = \pi(s_t|s_{t-1}) \ldots \pi(s_1|s_0)$ stand for the probability that s^t occurs given s_0. At time $t = 0$, the type of an individual household is characterized by his initial state, (k_0, h_0, s_0), where k_0 denotes the initial stock of physical capital and h_0 the initial stock of human capital (note that s_0 is not included in s^t). We take as given an initial distribution,

μ_0, of households over initial states (k_0, h_0, s_0), and a sequence of distributions, $\{\mu_{t,new}\}$, of new-born households over initial states.

Households are risk-averse and have identical preferences that allow for a time-additive expected utility representation with logarithmic one-period utility function and pure discount factor β. That is, for a household choosing the consumption plan $\{c_t\}$, expected life-time utility is given by

$$\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\nu(s_t)lnc_t(s^t)\pi(s^t|s_0) \tag{2}$$

where ν is a preference shifter that in the event of death of the household (family ceases to exists) is set to zero. Note that we have abstracted from the labor-leisure choice of households. Note also that with log-utility preferences, any deterministic change in household-size simply adds a constant to (2) without changing the optimal choice of households.

Each household can invest in human capital and buy and sell a complete set of financial assets (contracts) with state-contingent payoffs. More specifically, there is one asset (Arrow security) for each exogenous state s. We denote by $a_{t+1}(s_{t+1})$ the quantity bought (sold) in period t of the asset that pays off one unit of the good in the next period if s_{t+1} occurs in the next period. Given his initial state, (h_0, a_0, s_0), a household chooses a plan, $\{c_t, h_{t+1}, \vec{a}_{t+1}\}$, where the notation $\vec{a}$ indicates that in each period the household chooses a vector of asset holdings. Further, c_t stands for the function mapping partial histories, s^t, into consumption levels, $c_t(s^t)$, with similar notation used for the other choice variables. A budget-feasible plan has to satisfy the sequential budget constraint

$$\begin{aligned} r_h h_t + a_t(s_t) &= c_t + i_{ht} + \sum_{s_{t+1}} a_{t+1}(s_{t+1})q_t(s_{t+1}) \\ h_{t+1} &= (1-\delta_h(s_t))h_t + i_{ht} \\ 0 &\leq h_t + \sum_{s_{t+1}} a_{t+1}(s_{t+1})q_t(s_{t+1}) \\ c_t &\geq 0 \quad , \quad h_{t+1} \geq 0\ , \end{aligned} \tag{3}$$

where $q_t(s_{t+1})$ is the price of a financial contract in period t that pays off if s_{t+1} occurs in $t+1$. Note that in general prices depend on history and initial state, $q_t(s_{t+1}) =$

$q_t(s_{t+1}; s^t, a_0, h_0, s_0)$, though in our Markov setting the prices can be written as $q(s_{t+1}; s_t)$ (see below). In (3) i_{ht} is investment in human capital and $\delta_h(s_t)$ is the age- and shock-dependent depreciation rate of human capital. The term $\delta_h(s_t)$ captures all types of human capital risk as well as ex-ante heterogeneity in human capital returns (see Sections 5 and 6). Note that (3) has to hold in realization, that is, it has to hold for all t and all sequences $\{s_t\}$. Note also that the first inequality in (3) represents a debt constraint, which in our setting is equivalent to a no-Ponzi-scheme condition.

In addition to the standard budget constraint, each household has to satisfy a sequential enforcement (participation) constraint, which ensures that at no point in time individual households have an incentive to default on their financial obligations. More precisely, individual consumption plans have to satisfy

$$\sum_{n=0}^{\infty} \sum_{s^{t+n}|s^t} \beta^n \nu(s_{t+n})\, ln(c_{t+n}(s^{t+n}))\pi(s^{t+n}|s_t) \geq V_d(h_t, s_t) \ , \tag{4}$$

where V_d is the value function of a household who defaults. Note that (4) also has to hold in realization. Note further that the constraint set defined by (4) may not be convex since both the left-hand side and the right-hand side are concave functions of h.

The default value function, V_d, is defined by the following utility maximization problem. The consequences of default are designed to mimic Chapter 7 of the US bankruptcy code. Upon default, all debts of the household are cancelled and all financial assets seized so that $a_t(s_t) = 0$. Following default, a household is excluded from participation in financial markets for a period of time. For tractability, we assume exclusion continues until a stochastically determined future date that occurs with probability $(1 - p)$ in each period; that is, the probability of remaining in (financial) autarky is p. Following a default, households retain their human capital and continue to earn the wage rate $(1 - \tau_h)r_h$ per unit of human capital, where τ_h denotes the fraction of labor income that is garnished from households in default. In our baseline calibration, τ_h is set to zero as no wages are garnished under Chapter 7; later we analyze the effect of a reform of the bankruptcy code that allows for wages to be garnished (an increase in τ_h). After regaining access to financial markets, the households expected continuation value is $V^e(h, a, s)$, where (h, a, s) is the individual state at the time of regaining access. For the individual household the function V^e is taken as given, but we will close the model and determine this function endogenously by requiring that $V^e = V$, where V is the equilibrium value function associated with the maximization problem of a household who participates in financial markets.[7] In summary, a household who defaults in period t chooses a continuation plan, $\{c_{t+n}, h_{t+n}\}$, so as to maximize

$$\sum_{n=0}^{\infty} \sum_{s^{t+n}|s^t} p^n \beta^n \nu(s_{t+n})\, ln(c_{t+n}(s^{t+n}))\pi(s^{t+n}|s_t) + (1-p) \sum_{n=1}^{\infty} \sum_{s^{t+n}|s^t} p^{n-1} \beta^n\, V^e(h_{t+n}, s_{t+n})\pi(s^{t+n}|s_t) \ , \tag{5}$$

where $\{c_{t+n}, h_{t+n}\}$ has to solve the sequential budget constraint (2) with $a_t = 0$.

3.3 Equilibrium

In this paper, we confine attention to equilibria in which financial contracts are priced in a risk-neutral manner:

$$q(s_{t+1}; s_t) \quad = \quad \frac{\pi(s_{t+1}|s_t)}{1 + r_k - \delta_k} \ . \tag{6}$$

The pricing equation (6) can be interpreted as a zero-profit condition for financial intermediaries that can invest in physical capital at the risk-free rate of return $r_k - \delta_k$ and can fully diversify idiosyncratic risk for each insurance contract s_{t+1}.

Below we show that the optimal plan for individual households is recursive, that is, the optimal plan is generated by a policy function, g. This household policy function in conjunction with the transition probabilities, π, define a transition function over states, (h, a, s), in the canonical way. The transition function over individual states (h, a, s) conjunction with

[7] In other words, we assume rational expectations. The previous literature has usually assumed $p = 1$ (permanent autarky), and therefore did not have to deal with this issue. See, however, Krueger and Uhlig (2006) for a model with $p < 1$ following a similar approach to ours. Note also that the credit (default) history of an individual household is not a state variable affecting the expected value function, V^e. Thus, we assume that the credit (default) history of households is information that cannot be used for contracting purposes.

the initial distribution, μ_0, and sequence of distributions, $\{\mu_{t,new}\}$, induce a sequence of equilibrium distributions, $\{\mu_t\}$, of households over individual states, (h, a, s). Assuming a law of large numbers, aggregate variables can be found by taken the expectations with respect to the induced equilibrium distribution. For example, the aggregate stock of human capital held by all households in period t is given by $H_t = E[h_t] = \int h d\mu_t(h)$. A similar expression holds for the aggregate value of financial wealth. In equilibrium, human capital demanded by the firm must be equal to the corresponding aggregate stock of human capital supplied by households. Similarly, the physical capital demanded by the firm must equal the aggregate net financial wealth supplied by households. Because of the constant-returns-to-scale assumption, only the ratio of physical to human capital is pinned down by this market clearing condition. That is, in equilibrium we must have for all t

$$\tilde{K} = \frac{E[\sum_{s_{t+1}} q(s_{t+1}; s_t) a_{t+1}(s_{t+1})]}{E[h_t]} \,, \tag{7}$$

where $\tilde{K}$ is the capital-to-labor ratio chosen by the firm.

To sum up, we have the following equilibrium definition:

Definition A stationary recursive equilibrium is a collection of rental rates (r_k, r_h), an aggregate capital-to-labor ratio, $\tilde{K}$, a household value function, V, an expected household value function, V^e, a household policy function, g, and a sequence of distributions, $\{\mu_t\}$, of households over individual states, (h, a, s), so that

i) Utility maximization of households: for each initial state, (h_0, a_0, s_0), and given prices, the household policy function, g, generates a plan, $\{c_t, h_{t+1}, \vec{a}_{t+1}\}$, that maximizes expected lifetime utility (2) subject to the sequential budget constraint (3) and the sequential participation constraint (4).

ii) Profit maximization of firms: aggregate capital-to-labor ratio and rental rates satisfy the first-order conditions (1).

iii) Financial intermediation: financial contracts are priced according to (6)

iv) Aggregate law of motion: the sequence of distributions, $\{\mu_t\}$, is generated by g, π, μ_0,

and $\{\mu_{t,new}\}$.
v) Market clearing: equations (7) holds for all t when the expectation is taken with respect to the distribution μ_t.
vi) Rational expectations: $V^e = V$.

3.4 Discussion

The budget constraint (3) follows Jones and Manuelli (1990) and Rebelo (1991) by assuming that human capital and physical capital are produced using the same technology and that there are no diminishing returns to investment at the household level (competitive markets). In contrast, Lucas (1988) and Ben-Porath (1967) consider models with asymmetric production structures and diminishing returns at the household level. There are also differences with respect to the cost of human capital investment, where the former literature emphasizes direct costs and the latter indirect costs that arise when households have to allocate a fixed amount of time between work and human capital investment. To see the relationship of our approach to Ben-Porath (1967), note that a general formulation of the law of motion of human capital of a household would be $h_{t+1} = G(h_t, x_t, l_t, s_t)$, where l_t is time spent investing in human capital. Ben-Porath uses $G(h, l) = h + a(hl)^{\alpha}$ and Huggett et al. (2011) add human capital (depreciation) shocks: $G(h, l, s) = e^s\,(h + a(hl)^{\alpha})$, where s is normally distributed. Our formulation (3) assumes $G(h, x, s) = (1 - \delta_h(s))h + x$, but it is easy to see that our general equilibrium characterization result (propositions 1 and 2) goes through if $G(h, x, l, s) = g_1(l, s)h + g_2(l, s)x$, where g_1 and g_2 can be non-linear functions.[8]

We have chosen the specification (3) for two reasons. First, it keeps the model highly tractable, though a more general formulation of the type $G(h, x, l, s) = g_1(l, s)h + g_2(l, s)x$ would also deliver a tractability result. Second, it treats the production of physical and human capital fully symmetricly, which seems a useful abstraction given that our focus is on three properties of human capital that are *per se* unrelated to the production process,

[8]Formulation (3) makes another assumption that is very common in the literature, namely it lumps together general human capital (education, health) and specific human capital (on-the-job training).

namely that human capital is an asset that i) is risky, ii) has heterogeneous ex-ante returns, and iii) cannot be seized upon default.

The budget constraint (3) introduces risk and ex-ante heterogeneity in returns by assuming that the human capital depreciation rate depends on the exogenous state: $\delta_h = \delta_h(s)$. It is important to keep in mind that there is a formally equivalent formulation of the household problem in which risk and ex-ante heterogeneity of human capital returns arises because the productivity of human capital investment depends on s. More precisely, suppose that human capital evolves according to $h_{t+1} = (1 - \bar{\delta}_h) + z(s_t)x_t$, where z measures the productivity of human capital investment, that is, the number of goods needed to produce one more unit of human capital. It is straightforward to see that this formulation and formulation (3) lead to the same budget constraint if we set $z(s_t) = \left[(1 - \bar{\delta}_h)z(s_{t-1})\right] [r_h(1 - z(s_{t-1}) + (1 - \delta_h(s_t))]$.[9] Thus, our assumption in Section 6 that expected human capital returns are age-dependent does not literally mean that depreciation rates are age-dependent. Moreover, our choice of not imposing non-negativity constraint on human capital investment, which is essential for our tractability result, is much less severe than suggested by formulation (3). To see the last point, note that a non-negativity constraint on human capital investment in (3) means $h_{t+1}/h_t \geq 1 - \delta_h(s_t)$, whereas in the equivalent formulation with productivity differences it reads $h_{t+1}/h_t \geq 1 - \bar{\delta}_h$. Hence, if s has finite support, then for any solution to the household problem with budget constraint (3) we can find an equivalent formulation with $\bar{\delta}_h$ large enough so that the solution automatically satisfies the non-negativity constraint on human capital investment.

4. Equilibrium Characterization

In this section, we show that recursive equilibria can be found without knowledge of the

[9]For this equivalence result to hold, we have to change the definition of returns and total wealth accordingly, and the portfolio choices in the denominator of the market clearing condition (13) have to be multiplied by z.

endogenous wealth distribution and provide a convenient characterization of recursive equilibria as the solution to a finite-dimensional fixed-point problem. This characterization of recursive equilibria is then used for the subsequent analysis.

4.1. Household Problem

Denote total wealth (human plus financial) of a household at the beginning of the period by $x_t = h_t + \sum_s a_t(s)q(s)$. Further, denote the portfolio shares by $\theta_{ht} = h_t/x_t$ and $\theta_{at}(s_t) = a_t(s_t)/x_t$, and the total investment return by $1 + r_t = (1 + r_h - \delta_h(s_t))\theta_{ht} + \theta_{at}(s_t)$. Using this notation, the budget constraint (3) becomes

$$\begin{aligned}
x_{t+1} &= (1 + r(\theta_t, s_t))x_t - c_t \\
1 &= \theta_{h,t+1} + \sum_{s_{t+1}} q(s_{t+1}|s_t)\theta_{a,t+1}(s_{t+1}) \\
c_t &\geq 0 \ , \ x_{t+1} \geq 0 \ , \ \theta_{h,t+1} \geq 0 \ .
\end{aligned} \tag{8}$$

Clearly, (8) is the budget constraint corresponding to an inter-temporal portfolio choice problem with linear investment opportunities and no exogenous source of income. It also suggest that (x, θ, s) is the relevant state variable for the recursive formulation of the utility maximization problem. More specifically, the Bellman equation associated with the utility maximization problem of a household facing the budget constraint (8) and the sequential enforcement constraint (4) reads:

$$\begin{aligned}
V(x, \theta, s) \; &= \; \max_{c, x', \theta'} \left\{ ln\, c + \beta \sum_{s'} \nu(s')\, V(x', \theta', s')\, \pi(s'|s) \right\} \qquad (9) \\
s.t. \qquad & x' \; = \; (1 + r(\theta, s))x - c \\
& 1 \; = \; \theta'_h + \sum_{s'} q(s'|s)\theta'_a(s') \\
& c \geq 0 \ , \ x' \geq 0 \ , \ \theta'_h \geq 0 \\
& V(x', \theta, s') \; \geq \; V_d(x', \theta, s') \ ,
\end{aligned}$$

In contrast to the standard case without participation constraint, the Bellman equation (9) may have multiple solutions. However, in the Appendix we show that there is a maximal solution to (9), and this solution is also the value function of the corresponding utility maximization problem.

In the applications in Sections 5 and 6, we consider cases in which the exogenous state has several components, $s_t = (s_{1t}, \ldots, s_{nt})$, and only the first component, s_{1t} (age, ability), exhibits serial correlation (predictive power), whereas the remaining components are independently distributed over time. With this application in mind, let us assume that $\pi(s_{t+1}|s_t) = \pi(s_{t+1}|s_{1t})$. We further assume that the expected value function is logarithmic. In this case, it is well-known that the default consumption policy function is linear in total wealth and that the default value function is logarithmic (see Appendix for details), that is, the optimal policy function is

$$\begin{aligned} c(x,\theta,s) &= (1-\beta)(1+r(\theta,s))x \\ \theta'(x,\theta,s) &= \theta'(s_1) \\ x'(x,\theta,s) &= \beta(1+r(\theta,s))x \ . \end{aligned} \tag{10}$$

and the corresponding value function is given by

$$V(x,\theta,s) = \tilde{V}(s_1) \ + \ \frac{1}{1-\beta}\left[\ln\, x + \ln\left(1+r(\theta,s)\right)\right] \ , \tag{11}$$

The intensive-form value function, $\tilde{V}$, and the optimal portfolio choices, θ', are the solution to

$$\tilde{V}(s_1) \ = \ ln(1-\beta) \ + \ \frac{\beta}{1-\beta} ln\beta \ + \ \frac{\beta}{1-\beta}\sum_{s'} ln(1+r(\theta'(s_1),s'))\pi(s'|s_1) + \beta\sum_{s'_1}\tilde{V}(s'_1)\pi(s'_1|s_1)$$

and

$$\begin{aligned} \theta'(s_1) \ &= \ arg \max_{\theta'\in\Gamma(s_1)} \sum_{s'} ln(1+r(\theta',s'))\pi(s'|s_1) \\ \Gamma(s_1) \ &\doteq \ \left\{\theta' \left| \theta'_h + \sum_{s'} \frac{\theta'_a(s')\pi(s'|s_1)}{1+r_k-\delta_k} = 1 \ , \ \ \theta'_h \geq 0 \ , \right.\right. \\ &\quad \left. \tilde{V}(s_1) + \frac{1}{1-\beta} ln(1+r(\theta'_h,\theta'_a(s'),s')) \geq \tilde{V}_d(s_1) + \frac{1}{1-\beta} ln(1+r_d(\theta'_h,s'))\right\} \ . \end{aligned} \tag{12}$$

Note that the intensive-form value function, $\tilde{V}$, and optimal portfolio choices, θ, only depend on the component of s that has predictive power (serial correlation), that is, $\tilde{V}$ and θ are independent of any i.i.d. component.

Proposition 1. Suppose that the expected value function, V^e, is logarithmic. Then the default value function, V_d, is logarithmic. Further, the value function, V, is logarithmic, that is, it has the functional form (11) and the optimal policy function is given by (10), where optimal portfolio choices and intensive-form value function are determined by the solution to (12).

Proof: See Appendix.

Remark 1 The participation constraint in the maximization problem (12) is linear since the investment return, r, is linear in the portfolio choice, θ. Thus, the choice set in the maximization problem in is convex, and the non-convexity problem alluded to in the introduction has been solved in the context of the current model.

4.2. Intensive-form equilibrium

Define the share of aggregate total wealth of households of age s_1 as

$$\Omega(s_{1t}) \doteq \frac{E\left[(1+r_t)x_t|s_{1t}\right]\pi(s_{1t})}{E[x_t]}$$

Note that $(1+r_t)x_t$ is total wealth of an individual household after assets have paid off (after production and depreciation has been taken into account). Note also that $\sum_{s_{1t}}\Omega(s_{1t}) = 1$. Further, Ω is finite-dimensional, whereas the set of distributions over (x, s) is infinite-dimensional. Using the definition of wealth shares and the property that portfolio choices are wealth-independent, in the Appendix we show that the market clearing condition (7) is equivalent to the intensive-form market clearing condition

$$\tilde{K} = \frac{\sum_{s_1}(1-\theta_h(s_1))\Omega(s_1)}{\sum_{s_1}\theta_h(s_1)\Omega(s_1)} \,. \tag{13}$$

Further, in the Appendix we also show that the stationarity condition for Ω is given by

$$\Omega(s_1') = \frac{\sum_{s_1}(1+\bar{r}(s_1,s_1'))\Omega(s_1)}{\sum_{s_1,s_1'}(1+\bar{r}(s_1,s_1'))\Omega(s_1)} , \tag{14}$$

where we defined the expected investment return conditional on current and future age, s_1 and s_1', as $\bar{r}(s_1,s_1') = \sum_{s_{-1}'} r(\theta(s_1),s')\pi(s'|s_1)$ with $s_{-1} = (s_2,\ldots,s_n)$. Note that $\pi(s'|s_1) = \pi(s'|s_1')\pi(s_1'|s_1)$, which is the expression used in most applications.

In sum, we have

Proposition 2. Suppose that $(\theta,\tilde{V},\tilde{K},\Omega)$ is a stationary intensive-form equilibrium, that is, the portfolio choice θ together with the intensive-form value function $\tilde{V}$ are the solution to (12), the intensive-form market clearing condition (13) holds, and Ω satisfies the stationarity condition (14). Then $(g,\tilde{V},\tilde{K},\{\mu_t\})$ is a stationary recursive (balanced growth) equilibrium, where g is the individual policy function defined by (10) and $\{\mu_t\}$ is the sequence of measures recursively defined by μ_0, g, and π.

Proof. See the Appendix.

Remark 2 Proposition 2 shows that the equilibrium can be found without knowledge of the infinite-dimensional wealth distribution – only the lower dimensional distribution Ω matters. Proposition 2 in conjunction with the characterization of the household problem stated in proposition 1 show that the model is highly tractable.

5. Example

5.1 Set-Up

In this section, we confine attention to an economy with exogenous state $s_t = (s_{1t}, s_{2t}, s_{3t})$, where the first component denotes the type of the household (age, ability), the second component represents human capital risk (health risk, labor market risk), and the third component determines whether the household is alive. Note that the type of mortality risk analyzed in Section 6, namely that a member of a multi-person household dies but the household continues to exist, amounts to a shock to the human capital stock of a household and therefore enters the household decision problem through the second component, s_{2t}. We assume that the first component can take on two values, $s_{1t} \in \{l, h\}$ (low and high human capital returns), and is fully persistent: $\pi(s_{1,t+1}|s_{1t}) = 1$ if $s_{1,t+1} = s_{1t}$ and $\pi(s_{1,t+1}|s_{1t}) = 0$ otherwise. We further assume that human capital risk is an i.i.d. random variable, $\pi(s_{2,t+1}|s_{2t}) = \pi(s_{2,t+1})$, with two-state support, $s_{2t} \in \{b, g\}$ (bad and good shock). The human capital depreciation rate of a household of type s_1 with shock s_2 is given by $\delta_h(s_1, s_2) = \bar{\delta}_h(s_1) + \eta(s_2)$. We assume that the mean depreciation rate for the low-return household is high, $\bar{\delta}_h(l) > \bar{\delta}_h(h)$, and that human capital shocks have mean 0: $\eta(b) > 0$ and $\eta(g) = -\eta(b)\pi(b)/\pi(g) < 0$. The third component takes on two values, $s_{3t} \in \{n, d\}$, corresponding to death, $s_{3t} = d$, or no-death, $s_{3t} = n$, of the household. We assume that death of the household is an absorbing state, $\pi(s_{3,t+1} = d|s_{3t} = d) = 1$, and denote the probability of death of a household by $p_d = \pi(s_{3,t+1} = d|s_{3t} = n)$, and normalize the utility in the death state to zero: $\nu(d) = 0$. Finally, we assume that defaulting households are not excluded form financial markets: $p = 1$, which rules out short positions in financial assets (see Appendix).

5.2 Consumption and Insurance

Using the policy function (10) of our equilibrium characterization result, we find that consumption growth is given by:

$$
\begin{aligned}
\frac{c_{t+1}}{c_t} &= \tilde{\beta}(1 + r(\theta(s_1), s_{2,t+1})) \\
&= \tilde{\beta}\left(\theta_h(s_1)\left(1 + r_h - \bar{\delta}_h(s_1) - \eta(s_{2,t+1})\right) + \theta_a(s_1, s_{2,t+1})\right)
\end{aligned} \tag{15}
$$

with an effective discount factor $\tilde{\beta} = \beta(1 - p_d)$. Consumption growth depends on human capital choice, $\theta_h(s_1)$, ex-ante human capital returns, $r_h - \bar{\delta}_h(s_1)$, ex-post shocks, $\eta(s_{2,t+1})$, and asset payoffs (insurance), $\theta_a(s_1, s_{2,t+1})$.

Consider now a bad human capital shock of size $\eta(b)$. Note that $\eta(b)$ is the percentage of human capital lost, which equals the percentage drop in permanent income. We define the consumption drop associated with this drop in permanent income as the difference between the percentage decline in consumption and the mean consumption growth rate (conditional on type). Using (15), we find

$$\textit{consumption drop} \;=\; \tilde{\beta}\left(\eta(b)\,\theta_h(s_1) - (\theta_a(s_1,b) - E[\theta_a|s_1])\right)\;, \tag{16}$$

where $E[\theta_a|s_1] = \pi(b)\theta_a(s_1,b) + \pi(g)\theta_a(s_1,g)$ is the mean holding of financial assets of a household of type s_1. Note that $\eta(b)\theta_h$ is the human capital loss as a fraction of total wealth, x, and $\theta_a(s_1,b) - E[\theta_a|s_1]$ is the insurance pay-out as a fraction of total wealth. When these two terms are equal, we have full insurance and the consumption drop is nil. When there is no insurance pay-out, the consumption drop is $\tilde{\beta}\eta(b)\theta_h(s_1)$, which is less than the original drop in permanent income, $\eta(b)\theta_h(s_1)$, as long as $\tilde{\beta} < 1$. In this sense, there is self-insurance in the model.

In this paper, we consider two measures of insurance. Both capture the degree to which households insure against human capital risk by purchasing insurance contracts.[10] Our first insurance measure is defined as the fraction of the income loss that is insured. More precisely, we define it as the ratio of the insurance pay-out in the case of a bad shock, $(\theta_a(s_1,b) - E[\theta_a|s_1])\,x$, to the associated human capital loss, $\eta(b)\,\theta_h(s_1)x$:

$$I_1(s_1) \;\doteq\; \frac{\theta_a(s_1,b) - E[\theta_a|s_1]}{\eta(b)\,\theta_h(s_1)}\;.$$

The insurance measure I_1 varies between 0 if $\theta_a(s_1,b) = E[\theta_a|s_1]$, in which case we have no insurance, and 1 if $\theta_a(s_1,b) - E[\theta_a|s_1] = \eta\,\theta_h(s_1)$.

Our second measure of insurance is based on the idea that insurance reduces consumption volatility, where volatility is measured by the standard deviation of consumption growth. More precisely, we define

$$I_2(s_1) \doteq 1 - \frac{\sigma\left[c_{t+1}/c_t|s_1\right]}{\sigma\left[c_{a,t+1}/c_{a,t}|s_1\right]}\;,$$

[10]Blundell et al (2008) introduce an insurance coefficient that measures the extent to which consumption responds to income shocks. Clearly, their measure captures consumption insurance through self-insurance and the explicit purchase of insurance contracts, whereas our approach confines attention to the latter channel.

where $\sigma\left[c_{t+1}/c_t|s_1\right]$ is the standard deviation of equilibrium consumption growth and $\sigma\left[c_{a,t+1}/c_{a,t}|s_1\right]$ is the standard deviation of consumption growth in financial autarky. Note that consumption growth in financial autarky is simply given by (15) with $\theta_s(s_1, s_{2,t+1}) = 0$. If we assume a symmetric shock distribution, $\pi(b) = \pi(g) = 1/2$, we can show that the insurance measure I_2 varies between 0 and 1.

Proposition 3. Consider the simple economy described above. In equilibrium, household with high ex-ante human capital returns, $s_1 = h$, invest more in human capital and have less insurance than low-return households, $s_1 = l$:

$$\begin{aligned} \theta_h(h) &\geq \theta_h(l) \\ I_1(h) &\leq I_1(l) \\ I_2(h) &\leq I_2(l) \end{aligned}$$

where the last inequality, $I_2(h) \leq I_2(l)$, holds under the additional assumption of a symmetric shock distribution. The inequalities are strict if in equilibrium there is some insurance, but not full insurance.

6. Quantitative Analysis

Section 6.1 lays out the framework used for the quantitative analysis and Section 6.2 discusses the data. In Section 6.3 we outline our calibration strategy and provide a survey of the relevant empirical literature. Section 6.4 briefly discusses our computational approach of equilibria, with most of the details are relegated to the Appendix. In Section 6.5 we present the main equilibrium implications, in particular the model's implications for the life-cycle profile of human capital investment and insurance. Section 6.6 considers an extension of the model with risk heterogeneity and Section 6.7 analyzes a version of the model without the life-cycle. Section 6.8 presents our policy experiment, namely a reform of the bankruptcy code. We also conducted an extensive sensitivity analysis with respect to the main parameters of the model, but do not report the results here because of space limitations – details are available on request.

6.1 Set-Up

In this section, we consider a version of the model with ex-ante heterogeneity in human capital returns due to age-differences. We further focus on two types of human capital risk: mortality risk and labor market risk. The mortality risk we have in mind is the risk that an adult member of a multi-person household dies and the household continues to exists, leading to a loss in human capital and labor income available to the household (measured in equivalence units that adjust for the change in household-size). Labor market risk refers, for example, to the loss of firm- or occupation-specific human capital in the case of job displacement. Internal promotions and upward movement in the labor market provide two examples of positive human capital "shocks" related to the labor market.

We let the length of a time period be one year and consider a version of the general framework with exogenous state $s_t = (s_{1t}, s_{2t}, s_{3t})$. The first component of s_t denotes age, the second component represents mortality risk discussed above, and the third component subsumes all of labor market risk. We assume that the second and third component, s_2 and s_3, are independently distributed over time, but allow for an age-dependence of the distribution: $\pi(s_{2,t+1}, s_{3,t+1}|s_{2,t}, s_{3,t}) = \pi(s_{2,t+1}, s_{3,t+1})$ and $\pi(s_{2t}, s_{3t}|s_{1t}) = \pi(s_{2t}|s_{1t}) * \pi(s_{3t}|s_{1t}) \neq \pi(s_{2t}) * \pi(s_{3t})$. The age-component can take on the values $s_{1t} \in \{23, \ldots, 60, pre-retirement, retirement, death\}$. From age 23 to 60, the household is working and the transition from one age-group to the next is deterministic: $\pi(j+1|j) = 1$ for $j = 23, \ldots, 60$. Households in pre-retirement age also work, but the duration of this phase of life ends stochastically with retirement. The retirement probability is chosen so that retirement occurs on average at age 65. Finally, retired households die stochastically, in which case they have reached the absorbing state $s_{1t} = death$ and are replaced by a new-born household of age 23. For the preference shifter in (2) we assume $\nu(s_1 = death) = 0$ and $\nu(s_1) = 1$ for all $s_1 \neq death$.

We assume that the human capital depreciation rate can be decomposed as follows: $\delta_h(s_{1t}, s_{2t}, s_{3t}) = \bar{\delta}_h(s_{1t}) + \eta(s_1, s_{2t}) + \xi(s_{3t})$. The age-dependent mean, $\bar{\delta}_h(s_{1t})$, determines the expected human capital return of a household of age s_1 through $r_h - \bar{\delta}_h(s_1)$. We use a parsimonious specification for the life-cycle schedule and assume that the function $\bar{\delta}_h(.)$ is a fourth-order polynomial. We assume that the second term can take on two values, $s_{2t} \in \{b, g\}$, where $s_{2t} = b$ denotes the bad shock that a member of the household dies and $s_{2t} = g$ denotes the good shock that the death-event does not occur. We assume that the size of the human capital loss if $s_{2t} = b$ is independent of age: $\eta(s_1, b) = \eta(b)$. However, we allow the death probabilities to be age-dependent, and choose the realizations $\eta(s_1, g) < 0$ so that η is a random variable with mean zero.

Finally, $\xi(s_3)$ represents labor market risk. We assume that the human capital shocks due to labor market risk are log-normally distributed, $\ln(1+r_h-\bar{\delta}_h(s_1)+\xi) \sim N(\mu(s_1)-\sigma^2/2, \sigma^2)$. The assumption that human capital shocks are independently and log-normally distributed is also made by Huggett et al. (2011). In our setting, it has the advantage that it leads to a stochastic process of earnings that is consistent with the specification of a large number of empirical papers on labor market risk (see below). Note that the mean of human capital returns is increasing in μ and independent of σ, whereas the variance of human capital returns is independent of μ and increasing in σ.[11]

Our choice to match the basic life-cycle facts only up to age 60 follows Huggett et al. (2011) and is motivated by several considerations. First, the number of households for each age-group in our SCF-sample drops rapidly after age 60. Second, labor force participation falls near the traditional retirement age for reasons that are not modelled here. Third, the closer we get to the traditional retirement age, the more important non-negativity constraints on human capital investment become. By fitting the empirical life-cycle of earnings and wealth only up to age 60 and introducing a transition-group of households with stochastic

[11]More specifically, we have $1 + r_h - \bar{\delta}_h(s_1) = e^{\mu(s_1)}$ and $var[r_h - \bar{\delta}_h(s_1) + \xi | s_1] = var[\xi | s_1] = e^{2\mu(s_1)+\sigma^2}\left(e^{\sigma^2} - 1\right)$.

retirement, we can ensure that for the calibrated model economy the optimal choice of human capital investment is non-negative over the entire life-cycle.

We assume a Cobb-Douglas aggregate production function, $f(\tilde{k}) = A\tilde{K}^{\alpha}$. The computation of equilibria is based on the characterization results in proposition 1 and proposition 2. See the Appendix for more details on our computational approach.

6.2 Data

Data on earnings, financial wealth, and life-insurance are drawn from the 6 surveys of the Survey of Consumer Finance (SCF) conducted between 1992 and 2007. In the Appendix, we discuss in more detail the data, definition of variables, and sample selection. Here we only mention that the survey provides information about "families" corresponding to our concept of a household, and that we include single-person households as well as multi-person households in our basic sample. However, we also considered the the sub-sample of multi-person households, but the results for the empirical life-cycle profile of earnings, earnings growth, and wealth-to-earnings ratio were almost unchanged. For the case of life-insurance, we discuss below (Section 6.5.4) the effect of sample selection criteria on the empirical life-cycle profile. Household age refers to the age of the household head. The model variable "financial wealth" is associated with the variable "net worth" in the SCF, which is the value of all assets (excluding human capital) minus the value of all debt.

Our life-cycle profiles of earnings and earnings growth in Figures 1 and 2 are constructed as follows. We first compute median household earnings for each age group and survey (calendar time) using a centered 5-year age bin, and then remove possible time effects using time dummies as in Huggett et al. (2011).[12] This gives us a life-cycle profile of median earnings, which we smooth using a third-order polynomial. Finally, we compute from this smoothed life-cycle profile of median earnings a life-cycle profile of earnings growth rates. For the life-cycle of ratio variables plotted in Figure 3 (wealth-to-earnings ratio) and Figure

[12] We have also used cohort-dummies, with similar results.

7 (ratio of insurance face value to income loss), we compute the median of the ratio variable, but the results are very similar when computing the ratio of the median of the respective variables.

6.3 Calibration

We now discuss the targets used to calibrate the model and the resulting parameter values.

6.3.1 Bankruptcy Code

We assume that bankruptcy leads to an exclusion from financial markets for 7 years, but no loss of earnings, $p = 1/7$ and $\tau_h = 0$. In comparison, Chatterjee et al. (2007) use $p = 1/10$ and $\tau_h = 0$, and Livshits et al. (2007) choose no exclusion period beyond the period of default, but also introduce additional cost of bankruptcy. In all these papers, there are no insurance markets, and therefore no cost of bankruptcy that relates to the loss of access to insurance markets. In contrast, we assume that households in default are excluded from participation in credit *and* insurance markets, and in this sense we assume a harsher punishment of default than the previous literature. Krueger and Perri (2006) assume $p = +\infty$ for technical reasons, but also allow defaulting households to save.[13]

6.3.2 Mortality Risk

We calibrate the first component of human capital risk, η, as follows. We choose the probability that an adult member of the household dies so that we match the year-to-year average survival rates for the period 1991-2000 for the US life-tables for the respective age-group. The size of the negative human capital shock in the case of the death of a household member is set to $\eta(b) = 0.20$, that is, 20 percent of the human capital of a household is lost.

A value of $\eta(b) = 0.20$ is in line with the evidence presented in Weaver (2010), who

[13]Our model abstracts from a number of aspects of consumer default, and calibration of the parameter p is therefore subject to a fair amount of uncertainty. In this sense, one could argue that p, respectively τ_h, should be treated as a free parameter that is chosen to match a particular dimension of the data, an approach taken by Livshits et al. (2007).

reports projections based on a micro-simulation model. More specifically, Weaver (2010) reports that the median of the of the ratio of income of a widow to income of the couple before widowhood is expected to be between 0.61 and 0.63 depending on the birth cohort, where income includes social security payments but excludes any asset income (see table 6 in Weaver (2010)).[14] Given that the current paper deals with insurance provided by private markets, it seems reasonable to use as target the income loss after social security payments (public transfer payments) have been taken into account. Based on the equivalence scale used for the official US poverty threshold, a household of size $n-1$ needs 78% of the income of an n-person household independent of household size n. Thus, an income value of .62 translates into an equivalent income of $0.62/0.78 = 0.80$ of the income before widowhood, which amounts to an effective income drop of 20 percent.

There are two reasons why the value of 20 percent overstates the effective income loss for the typical household. First, Weaver (2010) and related studies focus on a group of households for which the surviving adult satisfies the eligibility requirements for Social Security widow benefits, a group that is more at risk than the typical US household. Second, by using the Official US Equivalence Scale, we have used equivalence weights that lead to the largest decline in living standard for the same drop in family income. For example, if we use the OECD Equivalence Weights, the implied drop in effective income is significantly smaller that 20 percent.

6.3.3 Labor Market Risk

In our baseline calibration, we choose the value of the variance parameter, σ^2, of the second component of human capital risk, ξ, based on the estimation results reported in Huggett et al. (2011), who find $\sigma^2 = 0.0123$ (a standard deviation of .11). To see how the model's earnings process with $\sigma^2 = 0.0123$ relates to the empirical findings of the extensive literature on labor market risk, note that labor income is given by $y_{ht} = r_h h_t$. Thus, earnings growth

[14] Holden and Brand (2003) find an income drop of very similar magnitude based on a sample of widowed women in the PSID data.

rates are equal to human capital growth rates: $y_{h,t+1}/y_{h1} = h_{t+1}/h_t$. Changes in human capital, in turn, can be computed using the optimal policy (10) leading to

$$\begin{aligned} \frac{h_{t+1}}{h_t} &= \frac{\theta_{h,t+1}}{\theta_{h,t}}\frac{x_{t+1}}{x_t} \\ &= \frac{\theta_{h,t+1}}{\theta_{h,t}}\beta(1+r(\theta_t,s_t)) \end{aligned} \tag{17}$$

If θ_h is close to one, which is the case for young households, then we can use the approximation $1+r(\theta,s) \approx 1+r_h-\delta_h(s)$. In this case, we find:

$$\ln y_{h,t+1} - lny_{ht} = g_t + \epsilon_t \tag{18}$$

Thus, log-labor income follows a random walk with age-dependent drift, $g_t = \ln\frac{\theta_{h,t+1}}{\theta_{h,t}} + \ln\beta + \mu_t - \frac{\sigma^2}{2}$, and stochastic innovation term $\epsilon_t \sim N(0,\sigma^2)$.[15] The random walk specification is often used by the empirical literature to model the permanent component of labor income risk with the additional assumption that the innovation term is normally distributed. Estimates of the variance of this innovation term, σ^2, usually range from .0225 (Carroll and Samwick (1997)) to .0361 (Meghir and Pistaferri (2004)) up to .0625 (Storesletten et al. (2004)), where we averaged over age-groups and, if applicable, over business cycle conditions. However, these values will overstate the true value of σ^2 if there is earnings profile heterogeneity in addition to stochastic shocks with a permanent component, which is the reason why we follow Huggett et al. (2011) and use the smaller value of $\sigma^2 = 0.0123$.[16]

Inference about the parameters governing labor market risk crucially depends on the

[15]We have ϵ_t instead of ϵ_{t+1} in equation (18), and the latter is the common specification for a random walk. However, this is not a problem if the econometrician observes the idiosyncratic depreciation shocks with a one-period lag. In this case, (18) is the correct equation from the household's point of view, but a modified version of (18) with ϵ_{t+1} replacing ϵ_t is the specification estimated by the econometrician.

[16]There are also reasons why the model might understate the true amount of human capital risk. For example, if we use $\ln(1+r(\theta,s))$ instead of $\ln(1+r_h-\delta_{ht})$ in (18), we would need to choose a higher value for σ^2 to match the same variance of labor income changes (see also Krebs, 2003). Further, the assumption that earnings innovations are (log)-normally distributed is likely to understate the true amount of human capital risk if, as the evidence indicates (Geweke and Keane, 2000), the actual distribution of earnings innovations has a fat lower tail.

degree to which life-cycle profiles of earnings are heterogeneous. More specifcally, the empirical literature on labor income risk can be broadly divided into two strands: one that, after controlling for observable characteristics, assumes that income profiles are homogeneous (MacCurdy (1982)) and one that assumes heterogeneity of income profiles (Lillard and Weiss (1979)). The first strand usually finds a large random walk component or at least a highly persistent component close to a random walk, whereas the second strand often finds that the estimated persistence parameter significantly differs from the random walk specification. For example, Guvenen (2007) finds that the estimated auto-correlation coefficient drops from 0.988 to 0.821 after income heterogeneity has been taken into account. However, based on Monte Carlo simulations, Hryshko (2009) finds that the random walk hypothesis cannot be rejected. Meghir and Pistaferri (2010) suggest that these two theories might not be mutually exclusive: when Baker and Solon (2003) estimate the parameters of a generalized earnings process that allows for profile heterogeneity, a random walk component, and a transitory component modeled as an AR(1), they find that the variance of the random walk component is precisely estimated and substantial, though smaller than the estimate when no earnings profile heterogeneity is allowed.

6.3.4 Expected Returns

We use an annual risk-free rate of $r_f = 3\%$. In comparison, Kaplan and Violante (2010) choose the same value and Huggett et al. (2011) and Krueger and Perri (2006) choose an annual risk-free rate of four percent, but also allow for capital income taxation.[17] We use the observed life-cycle profile of mean earnings to infer the life-cycle profile of expected human capital returns as follows.

Taking the conditional expectations in (17), we find

$$E\left[\frac{h_{t+1}}{h_t}|s_{1t}, s_{1,t-1}\right] = \frac{\theta_h(s_{1t})}{\theta_h(s_{1,t-1})}\beta\left(1 + \theta_h(s_{1,t-1})(r_h - \bar{\delta}_h(s_{1,t-1})) + (1 - \theta_h(s_{1,t-1}))r_f\right) \quad (19)$$

[17]Using real financial returns as a proxy for physical capital returns, anything between 1% (T-bills) and 7 % (stocks) seems defensible and has been used in the literature.

For given portfolio choice, (19) shows that an increase in expected human capital return, $r_h - \bar{\delta}_h(s_1)$, increases the earnings growth rate. For our calibrated model economy, this relationship still holds once we take into account the endogenous response of the portfolio choice. Thus, we can use (19) in conjunction with the life-cycle profile of earnings growth to identify the life-cycle profile of expected human capital returns, which is the approach taken here. More precisely, we choose the coefficients of the fourth-order polynomial $\bar{\delta}_h(.)$ in order to minimize the distance (L2-norm) between the empirical life-cycle of median earnings growth from age 23 to age 60 and the corresponding model prediction.[18] Figures 1 and 2 show the life-cycle profile of median earnings and median earnings growth in the data and according to the model, where the match between data and model is almost perfect even though we restricted the depreciation schedule to a fourth-order polynomial. Note that our life-cycle profile of median earnings is very similar to the life-cycle reported in Huggett et al. (2011), who use individual earnings data drawn from the PSID.

By drawing inference about human capital returns from the life-cycle profile of median earnings, we follow a long tradition in human capital theory. In most of the previous work, for example Porath (1967) and more recently Huggett et al. (2011), human capital returns decline with age because there are diminishing returns at the micro-level and the level of human capital is increasing with age. In contrast, in our approach human capital returns are decreasing with age because human capital investment is less productive for older households. In either approach, however, young households have a very strong incentive to invest in human capital since the expected returns to this investment are high.

6.4.5 Preferences and Production

We follow Huggett et al. (2011) and assume a capital share in output, α, of .32. We choose the remaining parameters β (annual discount factor), A (productivity), δ_k (capital depreciation rate), and p_d (probability of death of retired household) so that we match given

[18] The depreciation rate for households in the transition period, $\bar{\delta}_h(pre-retirement)$, is chosen to be equal to the implied depreciation rate for a households age 62.5.

values of i) the aggregate capital-to-output ratio, ii) the real interest rate, iii) the life-cycle average of the median of the ratio of financial wealth to labor income for the age-group $23-60$, and iv) the average human capital return for the age group $23-60$. For the aggregate capital-to-output ratio, we follow Huggett et al. (2011) and use a value of 2.95. The average of the median of the ratio of financial wealth to labor income in the data is equal to 2.5 for the age-group $23-60$. For the real interest rate we use 3 percent and for the average human capital return a value of 6 percent.

6.4 Computation

The computation of equilibria is based on proposition 2, that is, we compute intensive-form stationary recursive equilibria. To do this, we start with an aggregate capital-to-labor ratio, $\tilde{K}$, which defines the rental rates r_k and r_h, and solve the intensive-form household problem (12). Given the solution to the household problem, we compute a stationary relative wealth distribution, Ω, using (14). We use this Ω to compute a new $\tilde{K}$ and iterate over $\tilde{K}$ until the market clearing condition (13) holds. A detailed description of our solution method for solving the household problem (12) can be found in the Appendix.

6.5 Results: Life-Cycle Implications

6.5.1 Portfolio Choice

The portfolio mix between human and financial capital is measured by θ_h, the fraction of total wealth invested in human capital. Empirically, we can measure (net) financial wealth and the payoff to human capital holdings, namely labor income. Thus, we use the ratio of financial wealth to labor income, which in the model is given by $\frac{1-\theta_h}{r_h\theta_h}$, as our empirical measure of the portfolio choice of a households. Figure 3 shows the life-cycle profile of this ratio in the SCF data and according to the model. Clearly, the model provides a very good account of the this dimension of the data. In particular, the model matches well the observed increase in financial wealth relative to human wealth over the life-cycle, even though it has not been calibrated to match this target. In other words, one basic prediction of the theory, namely that households with high expected human capital returns should be heavily invested in human capital, is qualitatively and quantitatively supported by the empirical evidence.

6.5.2 Consumption Insurance

We first consider the model's implication for the life-cycle variation of consumption insurance based on the insurance measure $I_2 = 1 - \sigma_c/\sigma_{c,a}$ (see our definition in Section 5). The figure shows that insurance increases substantially with age. For example, the value of this insurance measure begins at 0.24 for households age 23 and increases to 0.81 for households age 60. Further, the lack of consumption insurance for young households is significant in welfare terms. To see this, we plot in Figure 5 the welfare consequences of this lack of insurance. More precisely, for each age group we compute the welfare gain of removing/insuring all risk assuming that the mean level of consumption remains the same, that is, we keep the portfolio choices fixed. The welfare gains shown in Figure 5 are expressed in percentage of lifetime consumption (consumption equivalent variation). Clearly, these gains are substantial. For example, for a household age 23 the welfare gain from perfect consumption insurance is about 4% of lifetime consumption. In comparison, Lucas (2003) finds, for the same preferences, welfare cost of aggregate consumption fluctuations that are less than 0.1 percent of lifetime consumption. In other words, the welfare cost of under-insurance of idiosyncratic human capital risk in our current heterogeneous-agent model are much larger than the cost of business cycle fluctuations in a corresponding representative-agent model.

6.5.3 Consumption Inequality

Here we compare the model's implication for consumption dispersion over the life-cycle with the pattern found in the US data. Figure 6 plots the variance of log adult-equivalent consumption in the US data (Consumer Expenditure Survey) from three studies, Aguiar and Hurst (2008), Deaton and Paxson (1994), and Primciceri and van Rens (2009), and the corresponding variance implied by the model. The figure shows that the model captures the increase in consumption dispersion observed in the data. Indeed, the model matches quite well the estimates of consumption dispersion reported by Aguiar and Hurst (2008), in particular the concave shape of the life-cycle profile of consumption dispersion. Note that these estimates are very similar to the one found in Heathcote et al. (2010).

The theoretical life-cycle profile of consumption dispersion depicted in Figure 6 is very similar to the life-cycle profile implied by the incomplete-market model analyzed in Huggett et al. (2010). In other words, the two human capital models, one with exogenous incomplete markets and one with complete markets and endogenous borrowing constraints, make very similar prediction with respect to this particular dimension of the data. We now turn to the analysis of another dimension of the data that is more informative for distinguishing between these two approaches, namely the purchase of life insurance contracts. By construction,the standard incomplete market model is silent about this dimension of the data, whereas the limited enforcement model makes a sharp prediction about the relationship between risk exposure and purchase of insurance contracts.

6.5.4 Life Insurance

Consider the event "death of a household member". As in the example of Section 5, we define a measure of insurance as the ratio of insurance pay-out to income loss, I_1. However, in the current version of the model, we also have labor market risk, and we therefore average over human capital shocks due to labor market risk by using $E[\theta_a|s_1, s_2 = b] - E[\theta_a|s_1]$ in the numerator of I_1. In order to construct an empirical proxy for I_1, note that $(E[\theta_a|s_1, s_2 = b] - E[\theta_a|s_1])\, x$ is the value of insurance pay-outs in the case of death, and $\eta(b)\theta_h(s_1)x$ is the value of the human capital loss. In the SCF data, we have detailed information about the size of the life insurance holdings and the corresponding insurance pay-out in the case of death (the face value of the contract), which provides us with an estimate of $(E[\theta_a|s_1, s_2 = b] - E[\theta_a|s_1])\, x$. For the human capital loss, $\eta(b)\theta_h(s_1)x$, we compute a measure of the present value of labor income lost as the product of current labor income, times a present value factor, times our value for $\eta(b)$. The present value factor, in turn, is defined as the present value of one dollar that grows according to the life-cycle profile of median earnings depicted in Figure 1, where we use the inter-temporal marginal rate of substitution, $\tilde{\beta}\,(c_{t+1}/c_t)^{-1}$, of the relevant age-group to discount future earnings.

In Figure 7 we show the life-cycle variation of the insurance measure I_1 in the data and according to the model. For the data, we show the insurance measure for two different samples: one sample including all households and one sample that is restricted to households that have purchased some life insurance. Clearly, in the data this insurance measure increases with age, where the increase in much more pronounced for the sample of all households. The model also generates a substantial increase of insurance with age and fits the data almost perfectly when we restrict ourselves to the sample of households with some life-insurance. However, for the sample of all households, the slope is much steeper than the slope predicted by the model. In other words, insurance purchase has an extensive margin and an intensive margin, and the data indicate that both margins are important for understanding life insurance. The model predicts very well the intensive margin, but it misses the extensive margin since in the equilibrium of the model all households purchase a positive amount of life insurance.

We have considered various changes in our sample selection criteria to analyze the robustness of our empirical result. For example, we restricted the sample to multi-person households and multi-person households with children, and also conditioned on education. Figure 8 shows the results and reveals the robustness of two important features of the data: insurance increases with age and this increase is much stronger if we do not condition on participation in the insurance market. Further, the slopes of the life-cycle plots are surprisingly similar for the sample of all households, the sample of multi-person households, and the sample of multi-person households with children. In other words, multi-person households with children are better insured than single-person households, but the increase of insurance with age is not so different from the increase observed for single-person households. The results for education are very similar and not shown here. Finally, we used different ways of computing the present value of lost earnings. For example, we used different values of $\eta(b)$

and also the risk-free rate to discount future earnings. The former had mainly the effect of shifting the age-insurance profile and the ladder of making the profile steeper.

To better understand the extensive margin in the data, we plot in Figure 9 the life-cycle pattern of the participation rate in the life-insurance market, that is, the fraction of households who have purchased some life-insurance. Figure 9 shows that participation rates increase until around age 45 and and remain roughly constant. Further, the initial increase in the participation rate is substantial. In the next section, we consider an extension of the model that incorporates this extensive margin and replicates the pattern of insurance for both the sample of all households and the sample of households that hold some life-insurance.

6.6 Extension: Heterogeneity in Mortality Risk

Our previous analysis of mortality risk allows for age-dependent death probabilities, but assumes that the size of the income loss, $\eta(b)$, is the same for all households independent of age or other characteristics. In this section, we introduce heterogeneity in these income losses. For example, effective income losses in the event "death of a household member" depend on family structure, and that family structure may change. More formally, we assume that that of all the households with a head aged 23, a fraction $(1 - \pi(23))$ have a value $\eta(b) = 0$ and the remaining $\pi(23)$ households have a value $\eta(b)$ that is drawn from a uniform distribution with support $[\eta_1, \eta_2]$, where $\eta_1 > 0$. Households with $\eta(b) > 0$ buy life-insurance, and $\pi(23)$ is therefore the participation rate of households age 23. This age-dependent participation rate, $\pi(age)$, increases with age for the following reason. Households who have $\eta(b) > 0$ will keep their value until retirement, but for households with $\eta(b) = 0$, there is a positive (and generally age-dependent) probability in each period that they draw a new value $\eta(b) > 0$ from the uniform distribution with support $[\eta_1, \eta_2]$, and then keep that new value forever.

The heterogeneity in mortality risk introduces an additional source of heterogeneity in household choice. More precisely, the current portfolio choice of an individual household, θ, now depends on current age and the current value of $\eta(b)$. Keeping in mind that there are now two components of the exogenous state that have predictive power, the definition of stationary recursive equilibrium of Section 4 can be applied without modification.[19]

We calibrate the underlying parameters of the distribution of mortality risk as follows. The median value of $\eta(b)$ for the sample of households who participate in the insurance market is simply the median value of the given uniform distribution: $\eta_{m,p} = .5 * (\eta_2 - \eta_1) + \eta_1$. We require this median value to be 0.2 as in the previous calibration. The median value of $\eta(b)$ for the entire sample of all households is given by $\eta_{m,all}(age) = (\pi(age) - 0.5) * (\eta_2 - \eta_1) + \eta_1$. Note that the former median value is independent of age, whereas the latter increases with age. The function $\pi = \pi(age)$ can be estimated from our data on participation rates in the life-insurance market (see Figure 9). Using this function and the restriction $.5 * (\eta_2 - \eta_1) + \eta_1 = 0.2$, we find the following formula: $\eta_{m,all}(age) = .4 * \pi(age) + \bar{\eta}$, where $\bar{\eta}$ is a constant that depends on the underlying distributional parameters. We choose $\bar{\eta}$ to match the life insurance of households age 25 in the data. All other calibration targets are chosen as in the previous analysis.

The results can be summarized as follows. The choices of individual households with respect to their portfolio holdings and insurance of labor market risk are only mildly affected by differences in $\eta(b)$. As a result, the implications of the calibrated model economy with additional heterogeneity in mortality risk for the main macroeconomic variables and the life-cycle profile of human capital investment and consumption insurance (Figures 3-6) are almost identical to the implications of the benchmark model discussed before. However, the implications for the life insurance market differ substantially. More specifically, the extended model implies the same age-insurance relationship for households participating in the life insurance market, but makes a very different prediction for the entire sample of all households.

Figure 10 plots the life-cycle profile of the median of life-insurance coverage according

[19] Formally, we can think of $s_1 = (s_{11}, s_{12})$ in our general formulation of Section 4, with s_{11} denoting age and s_{12} indexing mortality risk.

to the model and in the data for the two samples of households. It shows that the model is successful in matching the data for both sets of households, those who participate and those who do not. Put differently, the extended model matches both the intensive and the extensive margin of observed choices. Note that our insurance measure has the shock size, $\eta(b)$, in the denominator, and that for the one plot (data and model) we use $\eta_{m,all}(b) = .2$ and for the other plot (again data and model) we use $\eta_{m,p}(age)$.

6.7 The Importance of the Life-Cycle

Here we present the equilibrium implications of a model without life-cycle variations in expected human capital returns: $\bar{\delta}_h(s_1) = \bar{\delta}_h$. We keep all other assumptions unchanged and choose the calibration targets as follows. We choose the human capitals shocks, η and ξ, as before. We choose the parameter value $\bar{\delta}_h$ to be equal to the mean of $\bar{\delta}_h(s_1)$ in our previous calibration, which ensures that we have "on average" the same technology for producing human capital. We choose the remaining parameters β, A, and δ_k to match i) a real interest rate of 3%, ii) an aggregate capital-to-labor ratio of 2.95, and iii) an earnings growth rate that equals the average growth rate of households age 23 to 60 implied by our previous analysis.

For the calibrated model economy without age-dependent human capital returns, we find that consumption insurance is close to perfect: $I_1 = .96$ and $I_2 = .78$. Moreover, the welfare gain from removing all risk, mortality risk and labor market risk, is 0.4 percent of lifetime consumption. This result underscores the importance of ex-ante heterogeneity for our analysis. Without this heterogeneity, the consumption and welfare effects of limited contract enforcement are substantially smaller.

6.8 Reform of Consumer Bankruptcy Regulation

In this section, we analyze the consequences of a policy reform that makes bankruptcy more costly. Specifically, we consider an experiment in which the bankruptcy code is changed to make the consequences for a consumer of declaring bankruptcy similar to those for defaults on student loans, which allow for the garnishment of up to 15% of labor income. That is, as a result of the policy change, τ_h is increased from zero to 0.15. Other authors have conducted related policy experiments in models of consumer bankrupcty where access to insurance markets has been exogenously prohibited and hence where bankrupcty is the only form of insurance available (for example, Chatterjee et al. 2007 and Livshits et al. 2007). In these worlds, tougher sanctions against consumers declaring bankrupcty restricts insurance possibilities. By contrast, our focus is on the extent to which access to insurance will endogenously increase following the imposition of tougher sanctions on consumers in bankrupcty.

Garnishment of wages in default allows households to borrow more to invest in human capital, *and* to buy more insurance against human capital risk. As a result of allowing garnishment of 15 percent in labor income, there is an increases in aggregate human capital investment resulting in a very small decline in the returns to human capital investment and a forty basis point increase in the risk free rate. Everything else equal, this rise in the risk free rate benefits older households at the expense of the highly leveraged young households. The growth rate of the economy increases slightly from 1.1 to 1.2 per-cent per year.

Despite the modest aggregate effects of allowing a relatively modest level of wage garnishing, the microeconomic effects of the policy change are often large. Prior to the reform, the average level of debt to income in our economy averaged 34% for households aged between 23 and 60, which is somewhat larger than the corresponding value of unsecured consumer debt in the US data, but less than the ratio of consumer debt to income once durable goods and housing are included. After the policy reform, debt to income levels more than double, with particularly large increases observed for the youngest households.

The consequences of the policy reform for insurance against human capital risk also vary significantly across households. As shown in Figure 4, prior to the reform, our measure of consumption insurance varied from roughly 25% for households with a head aged 23 to more than 80% for households with heads aged 60. Averaging across these households, our consumption insurance measure was 51.3%. Following the policy reform this number rises by more than ten percent to 57%. Households with a head aged 61 and above were already highly insured against consumption fluctuations, and so looking at the average across the entire population aged 23 and above the policy reform produced a rise in the level of consumption insurance of just under ten per-cent. Similar effects are observed for life insurance.

As a measure of welfare, we compute the equivalent variation of the policy reform measured in units of lifetime consumption. An unweighted average of these estimates across all households aged 23 to 60 reveals that welfare increased by 0.63% of lifetime consumption. Almost the exact same number is found if we compute the (common) change in all households lifetimes consumption required to generate an equivalent level of population weighted utilitarian social welfare to that generated under the policy change. Almost all of this gain in welfare is due to the direct insurance effect, that is, the increases in the growth rate experiences by most households is so small that its impact on welfare is relatively small.

7. Conclusion

In this paper, we developed a tractable macroeconomic model in which households accumulate human capital that is both idiosyncratically risky and non-pledgeable against consumer debt. We used the framework to analyze the possible causes and consequences of under-insurance. The results of this paper suggest three lines of future research.

The first concerns the measurement of the extent of insurance against the various forms of human capital risk. In the paper, we restricted the attention to insurance against one form of human capital risk – the death of a family member – for three reasons. First, it is one of the most important – in the sense of a large shock size – forms of human capital risk. Second, it is a readily quantifiable risk for which other market imperfections, such as adverse selection, are likely to be less important. Third and finally, it is a type of human capital risk for which an insurance market exists that has a relatively simple structure. Future research on the observed lack of insurance against human capital risk needs to quantify the extent of under-insurance coming from different sources.

A second line of research concerns the extent of unobserved heterogeneity in returns to human capital across the population. Heterogeneity of human capital returns due to ability differences has been central to the work by, among others, Guvenen et al. (2011), Hugget et al (2011), and Cunha, Heckman, and Navarro (2005). In the current paper, we restricted attention to differences in returns by age, and argued that this dimension of heterogeneity can go a long way towards explaining a number of empirical facts about human capital choice and under-insurance. An important task for future research is to determine the extent to which additional heterogeneity is important in explaining additional empirical facts about human capital choice, borrowing, and insurance.

Finally, a third line of research would broaden the the set of assets available to households. The most important alternative asset is housing, which is also risky and which is, to varying degrees, partially collateralizable. All else equal, the perceived (utility) rates of return to housing investment are large, so that access to this asset will further strengthen the results of this paper: households would like to borrow to invest in housing and human capital, and these investment opportunities will compete with the need to purchase insurance. To what extent this effect is offset by the fact that some housing wealth can be used as collateral against borrowing remains an open quantitative question.

References

Aguiar, M. and E. Hurst (2008) "De-constructing Lifecycle Expenditures," *Working Paper*

Alvarez, F., and U. Jermann (2000) "Efficiency, Equilibrium, and Asset Pricing With Risk of Default," *Econometrica* 68: 775-797.

Andolfatto, D., and M. Gervais (2006) "Human Capital Investment and Debt Constraints," *Review of Economic Dynamics* 9: 52-67.

Attanasio, O., and S. Davis (1996) "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy* 104: 1227-62.

Baker, M., and G. Solon (2003) "Earnings Dynamics and Inequality Among Canadian Men, 1976-1992: Evidence from Longitudinal Income Tax Records," *Journal of Labor Economics* 21: 267-288.

Ben-Porath, Y. "The Production of Human capital and the Life-Cycle of Earnings," *Journal of Political Economy* 75: 352-365.

Bernheim, D., Forni, L., Gokhale, J., and L. Kotlikoff (2003) "The Mismatch Between Life Insurance Holdings and Financial Vulnerabilities: Adequacy of Life Insurance: Evidence from the Health and Retirement Survey," *American Economic Review* 93: 354-365.

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

Card, D. (2009) "Estimating the Returns to Schooling: Progress on Some Persistent Econometric Problems," *Econometrica* 69: 1127-60.

Carneiro, P., and J. Heckman (2002) "The Evidence on Credit Constraints in Post-Secondary Schooling," *Economic Journal* 112: 705-734.

Carroll, C., and A. Samwick (1997) "The Nature of Precautionary Wealth," *Journal of Monetary Economics* 40: 41-72.

Chatterjee, S., D. Corbae., M. Nakajima, and V. Rios-Rull (2007) "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica* 75: 1525-1589.

Chiappori, P., and B. Salanie (2000) "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy* 108: 56-78.

Cochrane, J. (1981) "A Simple Test of Consumption Insurance," *Journal of Political Economy* 99: 957-976.

Cordoba, J. (2008) "US Inequality: Debt Constraints or Incomplete Markets?" *Journal of Monetary Economics* 55: 350-364.

Cunha, F., Heckman, J., and S. Navarro (2005) "Separating Uncertainty from Heterogeneity in Life Cycle Earnings," *Oxford Economic Papers*, 57: 191-261.

Cutler, D., Finkelstein, A., K. McGarry (2008) "Preference Heterogeneity and Insurance Markets: Explaining a Puzzle of Insurance," *American Economic Review* 98: 157-162.

Deaton, A., and C. Paxson (1994) "Intertemporal Choice and Inequality," *Journal of Political Economy* 102: 437-67.

Geweke, J., and M. Keane (2000) "An Empirical Analysis of Male Income Dynamics in the PSID: 1968-1989," *Journal of Econometrics* 96: 293-356.

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

Guvenen, F., Kuruscu, B., and S. Ozkan (2011) "Taxation of Human Capital and Wage Inequality: A Cross-Country Analysis," Working Paper.

Holden, K., and J. Brand (2003) "Income Change and Distribution upon Widowhood: Comparison of Britain, U.S., and Germany," in E. Overbye and P. Kemp (eds.): *Pensions: Challenges and Reform.* Aldershot, Ashgate 2003.

Hryshko, D. (2009) "RIP to HIP: The Data Reject Heterogeneous Labor Income Profiles," *Working Paper*.

Huggett, M., G. Ventura, and A. Yaron (2011) "Sources of Lifetime Inequality, " *American Economic Review* 101: 2923-54.

Jappelli, T. (1990) "Who is Credit Constrained in the U.S. Economy?" *Quarterly Journal of Economics* 219-234.

Jones, L., and R. Manuelli (1990) "A Convex Model of Equilibrium Growth: Theory and Policy Implications," *Journal of Political Economy* 98: 1008-1038.

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

Kehoe, T., and D. Levine (1993) "Debt-Constrained Asset Markets," *Review of Economic Studies* 60: 868-888.

Kehoe, P., and F. Perri (2002) "International Business Cycles with Endogenous Incomplete Markets," *Econometrica*, 70: 907-28.

Kennickell, A., and M. Starr-McCluer (1994) "Changes in Family Finances from 1989 to 1992: Evidence from the Survey of Consumer Finances, " *Federal Reserve Bulletin* 80: 861-82.

Kocherlakota, N. (1996) "Implications of Efficient Risk Sharing Without Commitment," *Review of Economic Studies* 63: 595-609.

Krebs, T. (2003) "Human Capital Risk and Economic Growth," *The Quarterly Journal of Economics* 118: 709-744.

Krueger, D., and F. Perri (2006) "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory" *Review of Economics Studies*, 73: 163-193.

Krueger, D., and H. Uhlig (2006) "Competitive Risk Sharing Contracts with One-Sided Limited Commitment," *Journal of Monetary Economics* 53: 1661-1691.

Ligon, E., J. Thomas, and T. Worrall (2002) "Informal Insurance: Arrangements with Limited Commitment: Theory and Evidence from Village Economies," *Review of Economic Studies* 69: 209-244.

Liliard, L., and Y. Weiss (1979) "Components of Variations in Panel Earnings Data: American Scientists 1960-70," *Econometrica* 47: 437-454.

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

Lochner, L., and A. Monge-Naranjo (2011) "The Nature of Credit Constraints and Human Capital," *American Economic Review* 101: 2487-2529.

Lucas, R. (2003) "Macroeconomic Priorities," *American Economic Review* 93: 1-14.

Lucas, R. (1988) "On the Mechanics of Economic Development," *Journal of Monetary Economics* 22: 3-42.

MaCurdy, T. (1982) "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis," *Journal of Econometrics* 18: 83-114.

Meghir, C. and L. Pistaferri (2004) "Income Variance Dynamics and Heterogeneity," *Econometrica* 72: 1-32.

Meghir, C., and L. Pistaferri (2010) "Earnings, Consumption, and Life-Cycle Choices," NBER Working Paper.

Primiceri, G. and T. van Rens (2009) "Heterogeneous Life-Cycle Profiles, Income Risk and Consumption Inequality," *Journal of Monetary Economics* 56: 20-39.

Rebelo, S. (1991) "Long-Run Policy Analysis and Long-Run Growth," *Journal of Political Economy* 99: 126-150.

Rustichini, A. (1998) "Dynamic Programming Solution of Incentive Constrained Problems," *Journal of Economic Theory* 78: 329-354.

Storesletten, K., Telmer, C., and A. Yaron (2004) "Cyclical Dynamics of Labor Market Risk," *Journal of Political Economy*

Thomas, J., and T. Worrall (1988) "Self-Enforcing Wage Contracts," *Review of Economic Studies* 55: 541-554.

Townsend, R. (1994) "Risk and Insurance in Village India," *Econometrica* 62: 539-591.

Wright, M. (2001) "Investment and Growth with Limited Commitment," Working Paper, Stanford University.

Figure 1: Life-cycle profile of log labor income

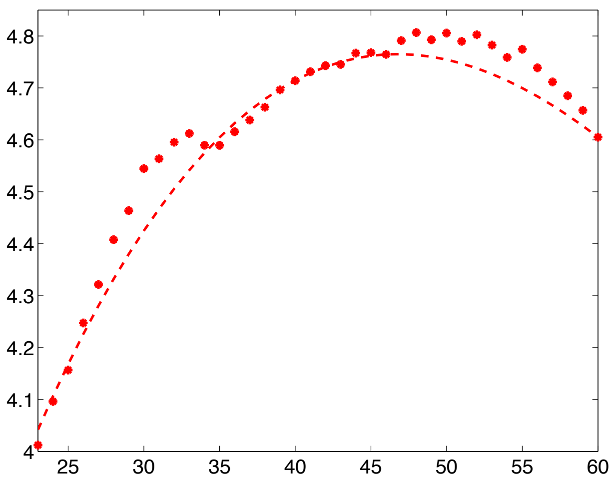

Notes: Life-cycle profile of median of log labor income for households age 23-60 from the SCF, surveys 1992-2007. The red points show the data and the red dashed line a polynomial fit of the data points.

Figure 2: Life-cycle profile of labor income growth

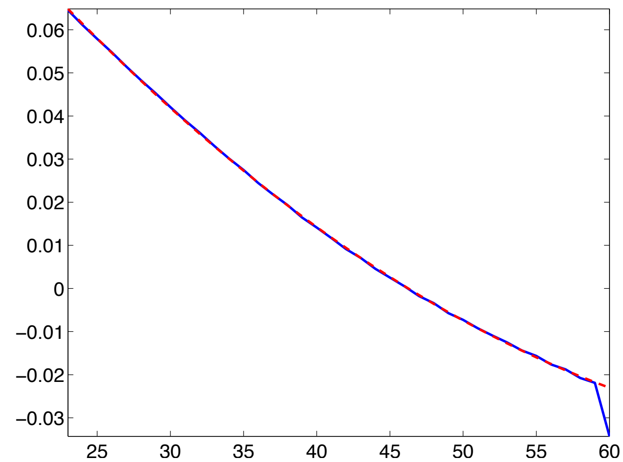

Notes: Life-cycle profile of the growth rate of labor income for households age 23-60. The red dashed line shows the smoothed data profile derived from figure 1 and the solid blue line the model fit.

Figure 3: Life-cycle profile of portfolio choice

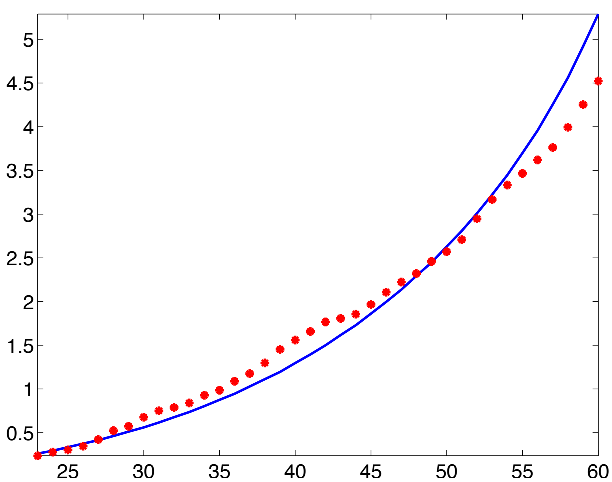

Notes: Life-cycle profile of the median of the ratio of financial wealth (value of all assets minus value of all debt) relative to labor income for households age 23-60. The red points show the data and the blue solid line the model prediction.

Figure 4: Life-cycle profile of consumption insurance

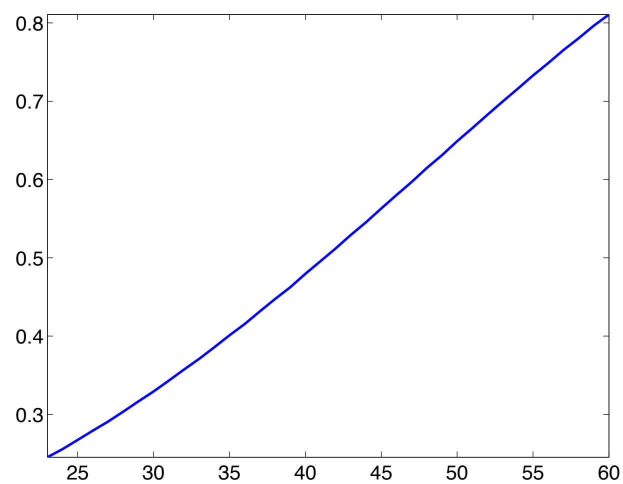

Notes: Life-cycle profile of consumption insurance in the model. The insurance measure is one minus the ratio of the standard deviation of consumption in equilibrium relative to the standard deviation of consumption in financial autarky.

Figure 5: Life-cycle profile of welfare cost of under-insurance

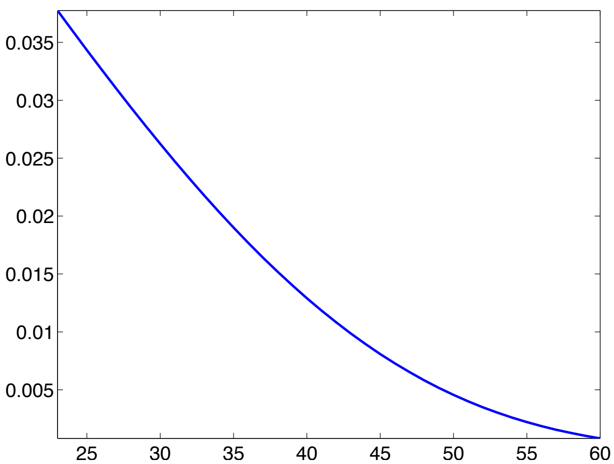

Notes: Welfare cost of under-insurance expressed in percentage of life-time consumption.

Figure 6: Life-cycle profile of consumption inequality

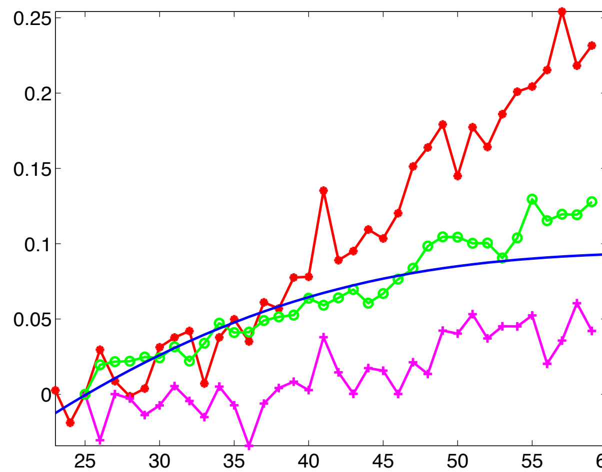

Notes: Life-cycle profile of the cross-sectional variance of consumption. The blue solid line shows the model prediction. The red points show the profile estimated by Deaton and Paxson (1994), the green circles are the estimates of Aguiar and Hurst (2008), and the crosses are the estimates of Primiceri and van Rens (2009). The data have been normalized to 0 at age 25.

Figure 7: Life-cycle profile of life insurance

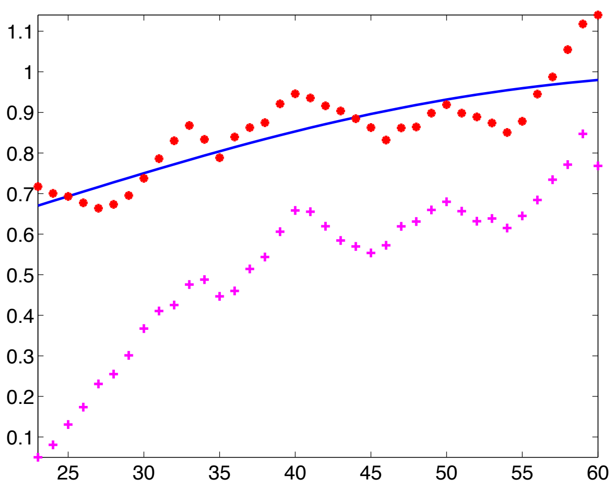

Notes: Life-cycle profile of the ratio of the median value of life insurance (insurance pay-out) to the median value of the permanent income loss. The median permanent income loss in the data is the product of median current income times present value factor times the fraction of income lost by the median household, which is assumed to be constant at 20%. The pink crosses show the data for the entire sample and the red points for the subsample of households that have purchased life insurance. The solid blue line is the model prediction.

Figure 8: Life-cycle profile of life insurance for different sample selections

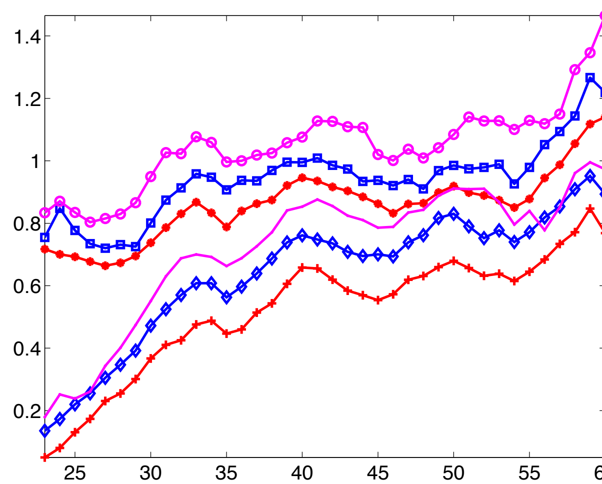

Notes: Life-cycle profile of the ratio of the median value of life insurance (insurance pay-out) to the median value of the permanent income loss. The median permanent income loss is the product of median current income times present value factor times the fraction of income lost by the median household. The red crosses show all households and the red points show the subsample of households that have purchased life insurance. The blue squares and diamonds show the same statistics for the subsample of households that are married or living with a partner. The pink circles and dots show the same statistics for the subsample of households that are married or living with a partner and have kids living in the household. The fraction of income lost by the median household is constant at 20%.

Figure 9: Participation rate in life-insurance market

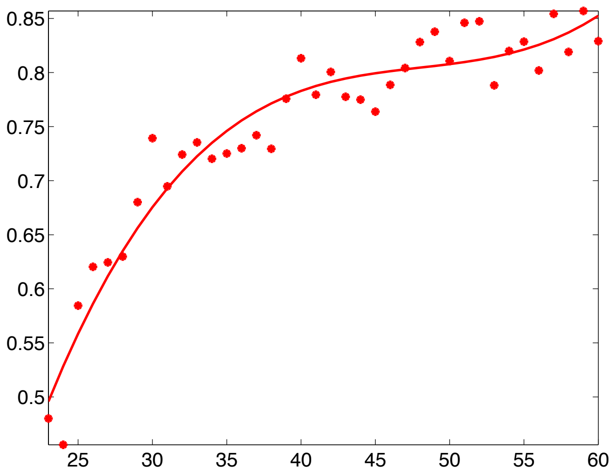

Notes: Life-cycle profile of the participation rate in the life-insurance market for households age 23-60. For each age the red star shows the share of households in the SCF that have positive holdings of life-insurance. The red solid line shows a polynomial fit to the data.

Figure 10: Life-cycle profile of life insurance (extended model)

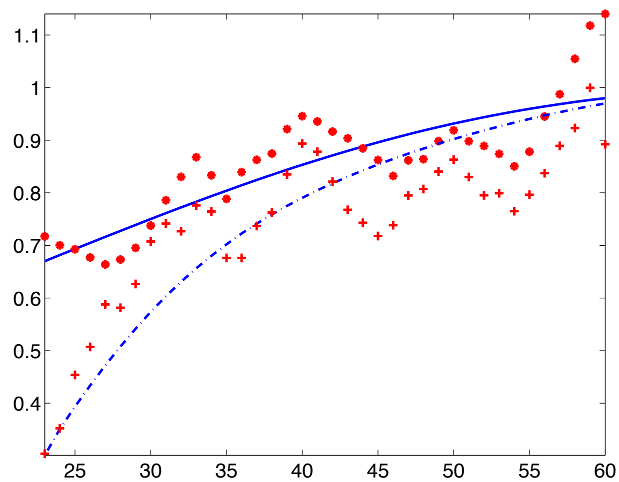

Notes: Life-cycle profile of the ratio of the median value of life insurance (insurance pay-out) to the median value of the permanent income loss. The median permanent income loss is the product of median current income times present value factor times the fraction of income lost by the median household. The red crosses show the data for the entire sample and the red points for the subsample of households that have purchased life insurance. The blue dashed dotted line is the model prediction for the entire sample. The blue solid line is the model prediction for the subsample of households that have purchased life insurance. The median fraction of income lost is different for the profiles. For the red points and the blue solid line the fraction is constant at 20%. For the red crosses and the blue dashed dotted line the fraction is increasing in age as described in the main text.

Appendix

A.1 Default value function

Denote by $r_d(\theta_h, s) = (1+(1-\tau_h)r_h - \delta_h(s))\theta_h$ the investment return of a defaulting household, which is simply the human capital return times the fraction of wealth invested in human capital. After the default period, we have $\theta_h = 1$, but in the period of default we have in general $\theta_h \neq 1$. The budget constraint of a household in default is $x' = (1 + r_d(s'))x - c$. In the period of re-gaining access to financial market, a household in default has no financial assets, and we still have $\theta_h = 1$. Suppose that the expected value function has the functional form

$$V^e(x, s) = \tilde{V}^e(s_1) \; + \; \frac{1}{1-\beta}\left[ln\, x \; + \; ln(1 + r_d(1, s))\right] \; . \qquad (A1)$$

Given (A1), it is straightforward to show that the autarky consumption policy function is linear in wealth, $c_d = (1-\beta)(1 + r_d(\theta_h, s))x$, and that the autarky value function has the functional form:

$$V_d(x, \theta_h, s) \; = \; \tilde{V}_d(s_1) + \frac{1}{1-\beta}[lnx \; + ln(1 + r_d(\theta_h, s))] \; . \qquad (A2)$$

The intensive-form default value function, $\tilde{V}_d$, is defined by the recursion

$$\begin{aligned}\tilde{V}_d(s_1) \;\; = \;\; & \ln(1-\beta) \; + \; \frac{\beta}{1-\beta}ln\beta \; + \; \frac{\beta}{1-\beta}\sum_{s'} ln(1 + r_d(1, s'))\pi(s'|s_1) \qquad (A3)\\ & + \beta p \sum_{s_1'} \tilde{V}_d(s_1')\pi(s_1'|s_1) + \; \beta(1-p)\sum_{s_1'} \tilde{V}^e(s_1')\pi(s_1'|s_1)\end{aligned}$$

For given $\tilde{V}^e$, equation (A3) determines uniquely the function $\tilde{V}_d$, which in turn pins down the default value function through (A2).

A2. Proof of proposition 1.

Let T be the operator associated with the Bellman equation (9). Adapting the argument made in Rusticchini (1998), the following result can be shown to hold in our setting:

Lemma Suppose that V_d and V^e are continuous functions. Suppose further that there is a unique continuous solution, V_0, to the Bellman equation without participation constraint. Let T stand for the operator associated with the Bellman equation. Consider the set of continuous functions B_W that are bounded in the weighted sup-norm $||V|| \doteq sup_x|V(x)|/W(x)$, where the weighting function W is given by $W(x) = |L(x)| + |U(x)|$ with U an upper bound and L a lower bound, and endow this function space with the corresponding metric.* Then

---

*Thus, B_W is the set of all functions, V, with $L(x) \leq V(x) \leq U(x)$ for all $x \in X$. For each particular

i) $\lim_{n\to\infty} T^n V_0 = V_\infty$ exists and is the maximal solution to the Bellman equation (9)

ii) V_∞ is the value function, V, of the sequential household maximization problem.

Suppose now that V^e has the functional form (A1), which implies that V_d has the functional form (A2). Clearly, both functions are continuous. Further, it is straightforward to show that in this case the Bellman equation without participation constraint has a unique continuous solution, V_0, and that this solution has the functional form (11). We now show that if $V_n = T^n V_0$ has the functional form, then the same is true for $V_{n+1} = TV_n$. To see this, note that V_{n+1} is defined as

$$\begin{aligned}
V_{n+1}(x,\theta,s) &= TV_n(x,\theta,s) \qquad\qquad (A4)\\
&= \max_{c,x',\theta'} \left\{ lnc \;+\; \beta \sum_{s_1'} \tilde{V}_n(s_1')\pi(s_1'|s_1) + \frac{\beta}{1-\beta}\left[\sum_{s'} ln(1+r(\theta',s'))\pi(s'|s_1) \;+\; ln\, x'\right]\right\}\\
s.t.\quad & x' = (1+r(\theta,s))x - c\\
& 1 \;=\; \theta_h' + \sum_{s'} \frac{\theta_a'(s')\pi(s'|s_1)}{1+r_f}\\
& \sum_{s'} \frac{\pi(s'|s_1)\theta_a'(s')}{1+r_f} \geq -D\\
& x' \geq 0 \;\;,\;\; c \geq 0 \;\;,\;\; \theta_h' \geq 0\\
& \tilde{V}_n(s_1) + \frac{1}{1-\beta} ln(1+r(\theta_h',\theta_a'(s'),s')) \geq \tilde{V}_d(s_1) + \frac{1}{1-\beta} ln(1+r_d(\theta_h',s'))
\end{aligned}$$

Clearly, the solution to the maximization problem defined by the right-hand-side of (A4) has the form

$$\begin{aligned}
x_{n+1}' &= \beta(1+r(\theta_{n+1},s))x \qquad\qquad (A5)\\
c_{n+1} &= (1-\beta)(1+r(\theta_{n+1},s))x\ ,
\end{aligned}$$

where the subscript $n+1$ indicates step $n+1$ in the iteration. Substituting this policy function into the right-hand-side of (A4) shows that V_{n+1} has the desired property.

From the lemma we know that $V_\infty = \lim_{n\to\infty} T^n V_0$ exists and that it is the maximal solution to the Bellman equation (9) as well as the value function of the corresponding sequential maximization problem. Since the set of functions with this functional form is a

application of the lemma, it has to be shown that this definition of the set of candidate value functions is without loss of generality for certain lower bound, L, and upper bound, U. In our case, the construction of the lower and upper bound is straightforward.

closed subset of the set of continuous functions, we know that V_∞ has the functional form. This proves proposition 1.

A.3 Proof of Proposition 2

From proposition 1 we know that individual households maximize utility subject to the budget constraint and participation constraint. Thus, it remains to show that the intensive-from market clearing condition (13) is equivalent to the market clearing conditions (7) and the stationary version of the law of motion for Ω is (14).

Let $\tilde{x}_t = (1 + r_t)x_t$ be the aggregate total wealth in period t after production and depreciation has taken place. The aggregate stock of human capital in period $t + 1$ is

$$\begin{aligned} H_{t+1} &= E[\theta_{h,t+1}x_{t+1}] \qquad\qquad (A6) \\ &= \beta E[\theta_{h,t+1}(1 + r_t)x_t] \\ &= \beta \sum_{s_{1t}} E[\theta_{h,t+1}\tilde{x}_t|s_{1t}]\pi(s_{1t}) \\ &= \beta \sum_{s_{1t}} \theta_h(s_{1t})E[\tilde{x}_t|s_{1t}]\pi(s_{1t}) \\ &= \beta E[\tilde{x}_t] \sum_{s_{1t}} \theta_h(s_{1t})\Omega(s_{1t}) \ . \end{aligned}$$

The second line in (A6) uses the equilibrium law of motion for the individual state variable x, the third line is simply the law of iterated expectations, the fourth line follows from the fact that the portfolio choices only depend on s_1, and the last line is a direct implication of the definition of Ω. A similar expression holds for the aggregate stock of physical capital, K_{t+1}. Dividing the two expressions proves the equivalence between (7) and (13).

Define the expected investment return conditional on age, $\bar{r}(s_{1t}, s_{1,t+1})$, as in (14). The law of motion for Ω can be found as:

$$\begin{aligned} \Omega_{t+1}(s_{1,t+1}) &= \frac{E[\tilde{x}_{t+1}|s_{1,t+1}]\pi(s_{1,t+1})}{E[\tilde{x}_{t+1}]} \qquad\qquad (A7) \\ &= \frac{E[(1 + r_{t+1})\tilde{x}_t|s_{1,t+1}]\pi(s_{1,t+1})}{E[(1 + r_{t+1})\tilde{x}_t]} \\ &= \frac{\sum_{s_{1t}} E[(1 + r_{t+1})\tilde{x}_t|s_{1t}, s_{1,t+1}]\pi(s_{1t}|s_{1,t+1})\pi(s_{1,t+1})}{\sum_{s_{1t},s_{1,t+1}} E[(1 + r_{t+1})\tilde{x}_t|s_{1t}, s_{1,t+1}]\pi(s_{1t}, s_{1,t+1})} \\ &= \frac{\sum_{s_1}(1 + \bar{r}(s_{1t}, s_{1,t+1}))E[\tilde{x}_t|s_{1t}]\pi(s_{1t})}{\sum_{s_{1t},s_{1,t+1}}(1 + \bar{r}(s_{1t}, s_{1,t+1}))E[\tilde{x}_t|s_{1t}]\pi(s_{1t})} \\ &= \frac{\sum_{s_{1t}}(1 + \bar{r}(s_{1t}, s_{1,t+1}))\Omega(s_{1t})}{\sum_{s_{1t},s_{1,t+1}}(1 + \bar{r}(s_{1t}, s_{1,t+1}))\Omega(s_{1t})} \end{aligned}$$

where the second line uses the equilibrium law of motion for the individual state variable x, the third line is simply the law of iterated expectations, the fourth line follows from the fact that portfolio choices only depend on s_1 in conjunction with the definition of $\bar{r}$, and the last line is a direct implication of the definition of Ω. Imposing stationarity, $\Omega_{t+1} = \Omega_t$, yields equation (14). This completes the proof of proposition 2 .

A.4. Proof of proposition 3

For each household type $s_1 \in \{l, h\}$, the solution of the household maximization problem (12) determines the optimal portfolio choice $\theta(s_1) = (\theta_h(s_1), \theta_a(s_1, .))$. Without loss of generality, assume that both households have some insurance in equilibrium, but not full insurance: $\theta_a(s_1, b) \neq \theta_a(s_1, g)$ and $\eta(b)\,\theta_h(s_1) \neq (\theta_a(s_1, b) - E[\theta_a|s_1])$. In this case, for both types s_1 the participation constraint binds if $s_2 = g$ and does not bind if $s_2 = b$. If the participation does not bind, the consumption growth rate must be equal to $1 + r_f$ with log-utility, which given the consumption rule (10) implies that the portfolio return in the bad state is equal to the risk-free rate. Adding the budget constraint, we find that the optimal portfolio choice, $(\theta_h(s_1), \theta_a(s_1, .))$, is determined by the following three equations:

$$
\begin{aligned}
\theta_h(s_1)\left(1 + r_h - \bar{\delta}_h(s_1) - \eta(b)\right) + \theta_a(s_1, b) &= 1 + r_f \qquad (A8)\\
\theta_h(s_1)\left(1 + r_h - \bar{\delta}_h(s_1) - \eta(g)\right) + \theta_a(s_1, g) &= e^{-(1-\beta)(\tilde{V}-\tilde{V}_d)}\theta_h(s_1)\left(1 + r_h\bar{\delta}_h(s_1) - \eta(g)\right)\\
\theta_h(s_1) + \frac{\pi(b)\theta_a(s_1, b)}{1 + r_f} + \frac{\pi(g)\theta_a(s_1, g)}{1 + r_f} &= 1\ .
\end{aligned}
$$

Suppose now that defaulting households keep access to financial markets: $p = 0$. In this case, we have $\tilde{V} = \tilde{V}_d$, and from the third equation in (A8) it follows that $\theta_a(s_1, g) = 0$. Further, solving for θ_h using $\theta_a(s_1, g) = 0$ yields:

$$
\theta_h(s_1) = \frac{\pi(g)}{1 - \frac{\pi(b)}{1+r_f}(1 + r_h - \bar{\delta}_h(s_1) - \eta(b))} \qquad (A9)
$$

Clearly, equation (A9) shows that $\theta_h(h) > \theta_h(l)$ if $\bar{\delta}_h(h) < \bar{\delta}_h(l)$. It further follows from equation ((A8) that the insurance pay-out is given by:

$$
\theta_a(s_1, b) - E[\theta_a|s_1] = \pi(g)\left(1 + r_f - \theta_h(s_1)(1 + r_h - \bar{\delta}_h(s_1) - \eta(b))\right)\ . \qquad (A10)
$$

Using $\theta_h(h) > \theta_h(l)$, it follows that $\theta_a(h, b) - E[\theta_a|s_1 = h] < \theta_a(l, b) - E[\theta_a|s_1 = l]$. This proves the first part of the proposition. A similar argument proves the second part of proposition 3.

A.5. Computation

Here we discuss the solution to the household decision problem. With a slight abuse of notation, denote the age by $s_1 = j = 23, \dots, 61, r$ and human capital risk by $(s_2, s_3) = s$, where we interpret the age 61 as the transition period before retirement and the age r as retirement. In this case, the equation system (12) determining the intensive-form value function and the optimal portfolio choice becomes

$$\tilde{V}_j = ln(1-\beta) + \frac{\beta}{1-\beta} ln\beta + \frac{\beta}{1-\beta} \sum_{s'} ln(1 + r_{j+1}(\theta_{j+1}, s'))\pi_{j+1}(s') + \beta \tilde{V}_{j+1} \qquad (A11)$$

for $j = 23, \dots, 60$ and for $j = 61$:

$$\begin{aligned} \tilde{V}_{61} &= ln(1-\beta) + \frac{\beta}{1-\beta} ln\beta + \frac{\beta}{1-\beta}(1-p_r) \sum_{s'} ln(1 + r_{61}(\theta_{61}, s'))\pi(s') + p_r \sum_{s'} ln(1 + r_r(\theta_{61}, s'))\pi_r(s') \\ &\quad + \beta(1-p_r)\tilde{V}_{61} + \beta p_r \tilde{V}_r \end{aligned}$$

with

$$\begin{aligned} \theta_{j+1} &= arg \max_{\theta' \in \Gamma_{j+1}} \sum_{s'} ln(1 + r_{j+1}(\theta', s'))\pi_{j+1}(s') \\ \Gamma_{j+1} &\doteq \Bigg\{ \theta' \Bigg| \theta'_h + \sum_{s'} \frac{\theta'_a(s')\pi_{j+1}(s')}{1 + r_f} = 1 \;\; , \;\; \theta'_h \geq 0 \;\; , \\ &\qquad \tilde{V}_{j+1} - \tilde{V}_{d,j+1} \geq \frac{1}{1-\beta} \left[ln(1 + r_{d,j+1}(\theta'_h, s')) - ln(1 + r_{j+1}(\theta'_h, \theta'_a(s'), s')) \right] \Bigg\} \;. \end{aligned}$$

for $j = 23, \dots, 60$ and for $j = 61$

$$\begin{aligned} \theta_{61} &= arg \max_{\theta' \in \Gamma_{61}} \left[(1-p_r) \sum_{s'} ln(1 + r_{61}(\theta', s'))\pi_{61}(s') + p_r \sum_{s'} ln(1 + r_r(\theta', s'))\pi_r(s') \right] \\ \Gamma_{61} &\doteq \Bigg\{ \theta' \Bigg| \theta'_h + \sum_{s'} \frac{\theta'_{a1}(s')\pi(s')}{1 + r_k - \delta_k} + \sum_{s'} \frac{\theta'_{a2}(s')\pi(s')}{1 + r_k - \delta_k} = 1 \;\; , \;\; \theta'_h \geq 0 \;\; , \\ &\qquad \tilde{V}_{61} - \tilde{V}_{d,61} \geq \frac{1}{1-\beta} \left[ln(1 + r_{d,61}(\theta'_h, s')) - ln(1 + r_{61}(\theta'_h, \theta'_{a1}(s'), s')) \right] \Bigg\} \;. \\ &\qquad \tilde{V}_r - \tilde{V}_{d,r} \geq \frac{1}{1-\beta} \left[ln(1 + r_{d,r}(\theta'_h, s')) - ln(1 + r_r(\theta'_h, \theta'_{a2}(s'), s')) \right] \Bigg\} \;. \end{aligned}$$

The default intensive-form value function is given by

$$\tilde{V}_{d,j} = \ln(1-\beta) + \frac{\beta}{1-\beta} ln\beta + \frac{\beta}{1-\beta} \sum_{s'} ln(1 + r_{d,j+1}(s'))\pi_{j+1}(s') + \beta p \tilde{V}_{d,j+1} + \beta(1-p)\tilde{V}_{j+1} \qquad (A12)$$

for $j = 23, \ldots, 60$ and for $j = 61$:

$$\begin{aligned}\tilde{V}_{d,61} &= \ln(1-\beta) + \frac{\beta}{1-\beta}\ln\beta + \frac{\beta}{1-\beta}\sum_{s'}\ln(1+r_{d,61}(s'))\pi_6 1(s') \\ &\quad + (1-p_r)\left(\beta p\tilde{V}_{d,61} + \beta(1-p)\tilde{V}_{61}\right) + p_r\left(\beta p\tilde{V}_{d,r} + \beta(1-p)\tilde{V}_r\right) \ ,\end{aligned}$$

where we imposed rational expectations: $\tilde{V}^e = \tilde{V}$. Finally, since retired households only hold physical capital and die with a constant probability, p_r, the intensive-form value function for them is

$$\tilde{V}_r = \frac{1}{1-\beta(1-p_r)}\left(\ln(1-(1-p_r)\beta) + \frac{\beta(1-p_r)}{1-\beta(1-p_r)}\ln\beta(1-p_r) + \frac{\beta(1-p_r)}{1-\beta(1-p_r)}\ln(1+r_f)\right) \quad (A13)$$

For $j = 61$, the Bellman equation (A11-A13) define a fixed point problem, which we solve by iterating over $\tilde{V}_{61}$ and $\tilde{V}_{d,61}$ until convergence. For all ages $j = 23, \ldots, 60$, we solve the Bellman equation backwards starting at $j = 60$. The portfolio choice problem for given $\tilde{V}_j$ and $\tilde{V}_{d,j}$ is solved in two steps. In a first step, we order the set of possible shocks, s', according to shocks size: $S = \{s_1, \ldots, \bar{s}\}$, where $\bar{s}$ is the shock for which the human capital shock is largest. We then fix the number of shocks for which the participation constraint is binding, that is, we conjecture that the participation constraint is binding for the first k shocks and not binding for the remaining $N - k + 1$ shocks. Given this conjecture, the portfolio choice, θ_j, is the solution to a linear equation system that is defined by the following equalities

$$\begin{array}{rccr} i) & \textit{participation constraint holds with equality} & \forall s' \in S_1 & (A14) \\ ii) & \textit{human capital return equals } r_f & \forall s' \in S_2 & \\ iii) & \textit{budget constraint holds} & & \end{array}$$

where S_1 is the set of shocks for which the participation constraint is binding and S_2 the set of shocks for which it is not binding. Note that equations i)-iii) are always necessary conditions for an optimum. They are also sufficient if two additional conditions, namely that a) the solution to (A14) satisfies the participation constraint for all shock realizations and b) the set S_2 has maximal size. We find the solution to (A14) satisfying these two additional requirements as follows.

We begin with $S_2 = S$ and checks if the solution to (A14) satisfies condition a), that is, satisfies the participation constraint for all $s' \in S$. If not, we choose $S_2 = S \setminus \{\bar{s}\}$, where $\bar{s}$ is the maximal element of the ordered set S, and check if the new solution to (A14) satisfies the participation constraint for all s'. We continue this process of eliminating states from the set S until we reached a point at which the solution to (A14) satisfies the participation constraint for all shocks $s' \in S$.

A.5. Data

The data are for the years 1992, 1995, 1998, 2001, 2004, and 2007 drawn from the Survey of Consumer Finances (SCF) provided by the Federal Reserve Board. The Survey collects information on a number of economic and financial variables of individual families through triennial interviews, where the definition of a "family" in the SCF comes close to the concept of a "household" used by the U.S. Census Bureau. See Kennickell and Starr-McCluer (1994) for details about the SCF.

For the sample selection, we follow as closely as possible Heathcote et al. (2010)[†]. We restrict the sample to households where the household head is between 23 and 60 years of age. We drop the wealthiest 1.46% and the poorest 0.5% of households in each year. Heathcote et al. (2010) show that this step makes the sample more comparable to the PSID or CEX data. We drop all households that report negative labor income or that report positive hours worked but have missing labor income or that report positive labor income but zero or negative hours worked. We compute the average wage by dividing labor income by total hours worked, and drop in each year households with a wage that is below half the minimum wage of the respective year. For the data on life-insurance, we restrict the sample further to households that are married or live with a partner.

For the definition of variables we follow Kennickell and Starr-McCluer (1994). We only depart from their variable definitions when considering labor income, where we follow Heathcote et al.(2010) and add two-thirds of the farm and business income as additional labor income. As common in the literature, we associate financial wealth in the model with net worth in the SCF. Households' net worth includes the cash value of life-insurance as in Kennickell and Starr-McCluer (1994), but does not include the face value of insurance contracts. We associate life-insurance in the model with the face value of life-insurance from the data. All data has been deflated using the BLS consumer price index for urban consumers (CPI-U-RS). A detailed description of the relevant variables is as follows:

- **Assets** are the sum of financial and non-financial assets. The main categories of non-financial assets are cars, housing, real estate, and the net value of businesses where the household holds an active interest. Except for businesses all values are

[†] We use their Sample B for our analysis.

gross positions, i.e. before outstanding debt. The main categories of financial assets are liquid assets, CD, mutual funds, stocks, bonds, cash value of life-insurance, other managed investment, and assets in retirement accounts (e.g. IRAs, thrift accounts, and pensions accumulated in accounts.)

- **Debt** is the sum of housing debt (e.g. mortgages, home equity loans, home equity lines of credit), credit card debt, installment loans (e.g. cars, education, others), other residential debt, and other debt (e.g. pension loans).
- **Net-worth** is the sum of all assets minus all debt.
- **Labor income** is wages and salaries plus 2/3 of business and farm income.
- **Life-insurance** is the face value of all term life policies and the face value of all policies that build up a cash value. The cash value is not part of the life-insurance, but is part of the financial assets of an household.