Human Capital Risk, Contract Enforcement, and the Macroeconomy*

Tom Krebs

University of Mannheim†

Moritz Kuhn

University of Bonn

Mark L. J. Wright

FRB Chicago and NBER

August 2014

Abstract

We use data from the Survey of Consumer Finance and Survey of Income Program Participation to show that young households with children are under-insured against the risk that an adult member of the household dies. We develop a macroeconomic model with human capital risk, age-dependent returns to human capital investment, and endogenous borrowing constraints due to the limited pledgability of human capital (limited contract enforcement). We show analytically that, consistent with the life insurance data, in equilibrium young households are borrowing constrained and under-insured against human capital risk. A calibrated version of the model can quantitatively account for the life-cycle variation of life-insurance holdings, financial wealth, earnings, and consumption inequality observed in the US data. Our analysis implies that a reform that makes consumer bankruptcy more costly, like the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005, leads to a substantial increase in the volume of both 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. We especially thank the editor and three referees for many suggestions and insightful comments. Tom Krebs thanks the German Research Foundation for support under grant KR3564/2-1. 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: Federal Reserve Bank of Chicago, 230 S. LaSalle Street, Chicago, IL 60604.

1. Introduction

For many households, human capital comprises the largest component of their wealth. Human capital is also subject to hard-to-diversify risk which is, in the case of disability or mortality risk, potentially catastrophic in magnitude. In this paper, we use data from the SCF and SIPP and provide empirical evidence that young households are severely under-insured against one important type of human capital risk, namely the risk of the death of an adult household member, even though life insurance markets are fairly competitive. Specifically, the median young household with children buys life insurance that covers between 10 and 40 percent of the net present value loss associated with the death of an adult family member, whereas older households are close to fully insured (see figure 1).[1] We further argue that the observed pattern of under-insurance can be fully explained by borrowing constraints that emerge endogenously due to the non-pledgability of human capital and bankruptcy code that limits wage garnishment (limited contract enforcement). Finally, we demonstrate that our approach has quantitatively important macroeconomic implications.

Our argument for the under-insurance of young households proceeds intuitively as follows. As is well known, labor earnings increase rapidly from the level achieved upon entering the job market reaching a peak in late middle age (see figure 2). Following a long tradition in labor economics, we interpret the high earnings growth of young households as the result of investment in high-return post-schooling human capital (on-the-job training). Thus, young households have access to a risky investment opportunity with high expected returns, but they also desire to smooth consumption and have little assets beyond human capital. Consequently, young households have a strong motivation for borrowing, which is, however, tightly constrained by their lack of collateralizable assets and their inability to pledge future earnings as collateral on their debts. On the margin, young households prefer to either consume or invest in human capital rather than purchase insurance. Hence, young households under-invest in risky human capital *and* under-insure against human capital risk even if insurance is priced in an actuarially fair manner.

To develop the argument more formally, we present a theory of the life-cycle accumulation of human capital and financial capital, the allocation of financial capital across assets, and

[1] Our measure of income losses takes into account social security survivor benefits, progressive income taxes, and implicit insurance from the possibility of re-marriage. See Section 2 for details.

the decision to purchase insurance against human capital risk. The enforcement of credit contracts is imperfect in the sense that human capital is non-pledgable and wage garnishment in the case of default is limited. Households in the model are heterogeneous, differing in their age, labor market status, marital status, number of dependents, and holdings of their human capital and financial assets. Household preferences are logarithmic and, in the baseline model, the marginal utility of consumption is independent of family structure and in particular independent of the death event. In other words, there are no variations of preferences over the life-cycle that induce life-cycle variations in insurance demand.

In a first step, we show that the model is tractable in the sense that individual decision rules are linear in household wealth (financial plus human) and that the infinite-dimensional wealth distribution is not a relevant state variable. We exploit the tractability of the model in both the theoretical analysis and the quantitative analysis conducted in this paper. In the theoretical analysis, we confine attention to a simplified version of the model with mortality risk as the only source of human capital risk. We then show analytically that young married households are under-insured against the risk of the death of an adult household member in the sense that the life insurance coefficient increases with age, where we define the life insurance coefficient as the ratio of insurance payout in the case of the death of a spouse over the human capital loss associated with the death event.

In a second step we provide a quantitative assessment of the theory. To this end, we calibrate the model to match a number of features of the US data, as well as the main features of the US chapter 7 bankruptcy code. Specifically, the model is calibrated to match i) the empirical life-cycle profile of median household earnings, ii) estimates of labor market risk obtained by the empirical literature, iii) the empirical life-cycle profile of mortality risk and rates of demographic transition (such as marriage and childbirth), and iv) the human capital losses associated with the death of a spouse estimated in this paper. We emphasize three findings of our quantitative analysis.

First, in equilibrium young households are severely under-insured against human capital risk, whereas older households are almost fully insured. In other words, our quantitative analysis suggests that realistic life-cycle variations in human capital returns combined with the basic features of the US bankruptcy code generate substantial under-insurance of young households through endogenous borrowing constraints. We show that this under-insurance result holds for two different insurance measures. First, it holds when we focus on mortality risk and measure the degree of insurance by the life insurance coefficient, where we proxy the human capital loss associated with the death of a spouse by the net present value income loss. Specifically, the insurance coefficient for mortality risk increases from below 40 percent at age 23 to almost 100 percent at age 60. Second, the result also holds when we consider all human capital risk and measure insurance by comparing the consumption volatility of households with access to insurance markets to the consumption volatility of households without access to insurance markets.

Our second finding is that the calibrated model economy can quantitatively account for the empirical life-cycle pattern of a number of important economic variables. Most importantly, the model provides an accurate account of the empirical life-cycle pattern of life insurance holdings and the empirical life-cycle profile of under-insurance against the death of an adult family member. In other words, the under-insurance pattern observed in the life insurance data can be fully explained by borrowing constraints that emerge endogenously due to the non-pledgability of human capital and a bankruptcy code that limits wage garnishment (limited contract enforcement). We further show that the model produces life-cycle profiles of financial wealth, human wealth, and consumption inequality that are in line with the data. We take these results as corroboration of our theory since the model has not been calibrated to match the corresponding targets.

Our third finding is that limited contract enforcement has quantitatively important macroeconomic 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, such as the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005. The theoretical literature has argued that making it more costly to declare bankruptcy not only increases the volume of credit, but also the amount of insurance purchased by households. In addition, in our human capital model it further increases economic growth since it leads to more investment in the high-return asset. In this paper, we contribute to the debate by providing a quantitative analysis of the various channels in a macro model with physical capital and human capital. For the calibrated version of the model, we find that the growth-effect is moderate, but that the insurance-effect is substantial: improvements in insurance against all human capital risk (labor income risk plus mortality risk) lead to an overall welfare gain of the order of 0.5% of lifetime consumption.

In sum, in this paper we make an empirical contribution by providing evidence of under-insurance in the life insurance market, we make a theoretical contribution by showing analytically that a model with endogenous borrowing constraints due to limited contract enforcement can explain this empirical finding, and we make a contribution to the quantitative macro literature by demonstrating that a calibrated version of the model can account for the observed life-cycle patterns of insurance, earnings, financial wealth, and consumption. In addition to these substantive contributions, this paper also makes a methodological contribution by developing a tractable framework and demonstrating how to surmount the non-convexity in the household decision problem that is inherent in production economies with limited pledgability and risk-averse agents. The tractability of the model is indispensable for providing analytical results regarding the relationship between age (expected human capital returns) and insurance. Further, it is essential for our quantitative analysis since our calibration approach and the general equilibrium analysis require that the computation of the solution to the (highly complex) households decision problem is not too time-consuming.[2]

The rest of this paper is organized as follows. Following a review of the literature, Section 2 describes a number of aspects of the US data on life insurance and presents our measures of under-utilization of life insurance. Section 3 presents our theoretical model, and derives our underinsurance result analytically for a special case of the model. Section 4 describes our calibration while Section 5 describes our quantitative results. Section 6 establishes that the results of the model are robust to a number of changes to the model, while Section 7 concludes.

Literature Our paper is related to three strands of the literature. First, there is the literature on life insurance. In line with our approach, most papers in this literature have argued that the market for life insurance is fairly competitive and that issues of moral hazard or adverse selection are not of first-order importance – see Section 2.2 for a more detailed

[2] The quantitative macro literature on limited commitment/enforcement with labor market risk has so far not analyzed models with life-cycle heterogeneity or endogenous human capital accumulation. There is, of course, quantitative work on incomplete-market models with life-cycle heterogeneity and human capital investment, but the computation of optimal household decision rules is much less complex in the incomplete-market case since there a fewer assets and therefore fewer portfolio choices.

discussion. In contrast to our approach, most previous studies of life insurance demand have focused on variations in household preferences or health status to explain empirical regularities. For example, Hong and Rios-Rull (2012) provide a preference-based explanation of the observed life-cycle pattern of life insurance holdings and Koijen, Nieuwerburgh, and Yogo (2012) analyze the effect of health status on the life insurance demand of elderly men. One important exception is Hendel and Lizzeri (2003) who emphasize, as we do, commitment problems in the market for life insurance. However, Hendel and Lizzeri (2003) study the link between the inability of households to commit to long-term contracts and the structure of insurance contracts (i.e. front-loading), whereas our interest lies in the effect of endogenous borrowing constraints on the purchase of life insurance contracts.

Second, our paper relates to the literature on risk sharing in models with limited enforcement/commitment (Alvarez and Jermann, 2000, Kehoe and Levine, 1993, Kocherlakota, 1996, Thomas and Worrall, 1988). Our theoretical contribution to this literature is to develop a tractable human capital model and to show how to avoid the non-convexity problem in a class of limited enforcement production models.[3] We also make a substantive contribution to this literature. Specifically, we show that a calibrated macro model with physical capital and limited contract enforcement generates substantial lack of consumption insurance and fully explains the observed life-cycle pattern of life insurance holdings. In contrast, the previous literature did not consider life-cycle variations and found that consumption insurance is almost perfect in calibrated models with physical capital (Cordoba, 2004, and Krueger and Perri, 2006).[4] Broer (2014) provides a detailed discussion of the quantitative implications of limited commitment models without a life-cycle component. Finally, we share with Andolfatto and Gervais (2006) and Lochner and Monge (2011) the focus on models with human capital and endogenous borrowing constraints due to enforcement problems, but we go beyond their work by studying the interaction between borrowing constraints and insurance.

[3]Wright (2001) has shown how to circumvent the non-convexity issue in linear production models (AK-model) with limited commitment. See also Azariadis and Kaas (2008) for a contribution that exploits the linear production structure in limited commitment models. The model structure we use in this paper is based on the human capital model with incomplete markets analyzed in Krebs (2003).

[4]Krueger and Perri (2006) match the cross-sectional distribution of consumption fairly well, but the implied volatility of individual consumption is negligible in their model.

Third, our paper is related to the voluminous literature on macroeconomic models with incomplete markets, and in particular studies of human capital accumulation (Krebs, 2003, Guvenen, Kuruscu, and Ozkan, 2011, and Huggett, Ventura, and Yaron, 2011) and the life-cycle profile of consumption (Storesletten, Telmer, and Yaron, 2004a, and Kaplan and Violante, 2010). This literature has shown that the incomplete-market framework is a powerful tool for understanding human capital investment choices and the observed life-cycle behavior of income, consumption, and human capital. In this paper, we show that a model with one financial friction, limited contract enforcement, explains equally well the life-cycle pattern of income, consumption, and human capital. Further, we show that, at least in the context of our discussion of the reform of the US bankruptcy code, the policy implication of these two classes of models differ significantly.

2. Empirical Evidence

In this section, we discuss the empirical evidence that motivates this paper. Specifically, we present the life-cycle profile of life insurance (Section 2.1), the life-cycle profile of mortality risk and under-insurance against this risk (Section 2.3), and the life-cycle variation of the relative importance of human capital and financial capital (Section 2.4).

2.1 Life Insurance

Our primary source of data on life insurance holdings is the Survey of Consumer Finance (SCF). Our data are drawn from the 6 surveys of the SCF conducted between 1992 and 2007. The unit of observation is the "family", which corresponds to our concept of a household, and we measure the household's age as the age of the household head. We focus our attention on married households with at least one child since they constitute a group of households that we can identify in our data as a group with a clear motive for purchasing life insurance (see also the discussion and references in Inkmann and Michaelides, 2012). We construct life-cycle profiles by computing median household values using five-year age bins for each survey removing possible time effects using time dummies as in Huggett et al. (2011).[5] Further details on the data, definition of variables, and sample selection are provided in the Appendix.

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

Life insurance contracts can be approximately divided into Term Insurance and Whole Life Insurance.[6] Term insurance contracts only offer insurance against the death event, whereas whole life insurance contracts offer a combination of insurance and saving. We use the face value (amount of money paid in the case of death) of all insurance contracts, term insurance and whole life insurance, to construct the amount of insurance owned by a household, subtracting the savings component of whole life policies as reported in the SCF. The SCF presents total holdings for the household, and hence these data reflect the total payout from the death of both spouses (we present data on life insurance holdings by spouse from an alternative data source in Section 4 below).

Figure 3 shows the empirical life-cycle profile of life insurance purchases of married households with children. The first line plots the median across all such households and shows that households in their early 20's have roughly $15 thousand dollars of life insurance. This rises to about $150 thousand by the time these households reach their 40's, and decline to $50 thousand as households reach retirement age. The second line plots the median across only those households that have purchased some life insurance. Amongst these households, the young purchase around $85 thousand in life insurance, rising quickly to $200 thousand before declining slowly down to $75 thousand in their early 60's.

The inverted u-shaped pattern of both series seems to indicate that young households are under-insured. However, it could also mean that young households simply need less insurance. To establish whether households are underinsured against the risk of death of a spouse, we need to take a stand on the "appropriate" level of insurance. One approach would be to use a model to deduce optimal holdings in the absence of market frictions, and use this as a measure of full insurance. We return to this approach later in the paper after presenting our model. In Section 2.3 below we turn to an alternative approach in which we proxy the size of the loss of human wealth upon death of a spouse by the present value of income losses taking into account the implicit insurance that results from the possibility of re-marriage, social security survivor benefits, and progressive income taxation. These two approaches yield almost identical answers for our baseline model with marginal utility of consumption that is independent of the family state – compare figures 7 and 8.

[6] Universal life insurance is grouped with whole life insurance.

2.2 Life Insurance – Two Issues

Life insurance contracts can be divided into insurance that households purchase directly from insurance companies and insurance that is obtained through employment or membership in organizations (group insurance). If the amount of group insurance offered by the employer exceeds the amount households want to hold, then these households are "involuntarily" over-insured and the insurance holdings observed in the data are not the outcome of the optimal insurance choice by households. Clearly, the phenomenon of involuntary over-insurance can only occur for households who have not purchased any individual life insurance from insurance companies. Although the SCF does not distinguish between group insurance and insurance purchased individually, data on employer provided life insurance are available from the Survey of Income and Program Participation (SIPP). Based on data drawn from the SIPP, we find that for each age between 23 and 60, the median household with kids holds substantially more life insurance than the amount of insurance provided by the employer. Further, for the median household with children the amount of employer-provided life insurance is roughly constant over the life-cycle and the shape of the life-cycle profile of total (group plus individual) life insurance holdings is therefore not much affected by the presence of group life insurance. See the Appendix for more details. Thus, we conclude that the consideration of insurance purchases as voluntary is appropriate to a first approximation. Hong and Rios-Rull (2012) come to a similar conclusion after analyzing data drawn from the International Survey of Consumer Financial Decisions.

In line with most of the previous literature on life insurance (Hendel and Lizzeri 2003, Hong and Rios-Rull 2012, and Koijen et al 2012), we model the market for life insurance as a competitive market with actuarially fair pricing. This seems reasonable given the large number of competing providers and the lack of regulatory inference, and has found support in the data.[7] We also follow the bulk of this literature in abstracting from considerations of asymmetric information. We argue that this is reasonable given that moral hazard problems appear small, and that adverse selection is limited by the requirement of a medical exam and the provision of a medical history with the risk that a policy will be voided if health information is not fully disclosed. Further, the available empirical evidence suggests that adverse selection is not of first-order importance in the market for life insurance (see, for

[7] For example, Winter (1981) finds little evidence of discrimination in the pricing of life insurance policies.

example, Cawley and Philipson 1999, Hendel and Lizzeri, 2003, and Koijen et al. 2012).[8]

2.3 Human Capital Risk and Underinsurance

Human capital is subject to a significant amount of idiosyncratic risk. In this paper, we divide these risks into labor market risk and demographic risk, with a particular focus on the mortality risk of spouses. We view labor market risk as that risk which affects observed labor earnings, which includes job displacement risk and some forms of disability risk. We follow a substantial literature (e.g. Huggett et al. 2011, Krebs 2003) and set the parameters describing human capital risk so that the implied labor market earnings process is consistent with estimates of permanent labor market risk obtained by the empirical literature (Carroll and Samwick 1997, Meghir and Pistaferri 2004, Storesletten et al. 2004), and defer a discussion of the details until we calibrate the model in Section 4.

Demographic risk captures the effects of marriage, divorce, childbirth, and death of a spouse on household earnings. Rates of marriage and childbirth are calibrated using data from the SIPP. As a result of the small numbers of young widows and widowers in the SIPP, we cannot reliably estimate a life-cycle profile of re-marriage rates for widows. We therefore follow the macro literature on life insurance (Hong and Rios-Rull, 2012) and use re-marriage rates of divorcees as a first proxy for the re-marriage rates of widows/widowers, but in contrast to the previous literature we introduce an adjustment factor to take into account what is known about re-marriage rates of widows and widowers from the economics and sociology literatures.[9] Section 4.2 on calibration of the model discusses in detail our approach. Mortality risk is chosen to match the year-to-year average survival rates for the period 1991-2000 from the US life-tables for the respective group.

Computation of the size of the human capital loss associated with the death of a spouse is also inhibited by the relative paucity of data on young widows and widowers. We focus on households with median earnings and approximate the amount of household human capital lost in the case of death by the (expected) present value of the after-tax earnings differential

---

[8]By contrast, there is considerable evidence of adverse selection in the market for annuities, where a medical exam is not required (Brugiavini 1990 and Friedman and Warshawski 1990). See also Society of Actuaries (2012) for details on the range of data collected by life insurers.

[9]Both widows and widowers have lower re-marriage rates than divorcees for each age group (Norton and Miller, 1990, and Wilson and Clarke, 1992).

between married households and the corresponding single household after including Social Security survivor benefits. We use the income tax and survivor benefit schedules from the year 2000, the mid point of our sample. See Section D of the Appendix for a detailed description of our approach and Section H of the Appendix for a robustness analysis.

Figure 4 plots the ratio of human capital losses in the event of death of a spouse to household labor earnings for the sample of married households with kids. To allow comparability with the life insurance data described above (which is aggregated for one household) the loss associated with the death of the household head is added to the loss from the death of their spouse. The first line depicts the loss of labor pre-tax earnings without allowing for the possibility that a widow or widower can remarry and shows that young households with children lose roughly 30 years of earnings following the death of a spouse. The second line includes the effect of taxes and social security survivor benefits on lost earnings, but does not allow for re-marriage, and shows that the government provides a substantial amount of insurance against the death of a spouse: for young households with children, the loss has declined to 15 years of earnings after taxes and transfers. Finally, the third line, which also allows for remarriage as a kind of informal insurance against loss of a spouse, shows that the resulting income loss is reduced to between 8 and 9 years of annual earnings for young households, with smaller reductions for older households who face lower remarriage rates.[10] Overall, our results suggest that income losses in the case of death of a spouse are substantial, but much less than a simple calculation that does not take into account non-linear taxes, social security survivor benefits, and remarriage, would suggest. Further, human capital losses expressed as a fraction of household human capital decline with age, suggesting that younger households should purchase more life insurance than older households.

Although other sources of informal insurance are possible, we argue that, with one exception, they are likely to be insignificant. For example, it is possible that the surviving spouse increases their own labor supply. However, the substantial empirical literature examining the responses of spouses labor force participation and hours worked to shocks in earnings or disability typically finds little or no effect (Gallipoli and Turner 2011, Heckman and Macurdy

[10] Our computed income losses without re-marriage are in line with the results in the literature (for example, Burkhauser et al. (2005) and Weaver 2010). This literature, however, has not computed effective income losses taking into account re-marriage.

1980, Gruber and Cullen 1996). Similarly, private transfers from outside of the household appear insignificant following the disability of a spouse (Gallipoli and Turner 2009) and we argue are also likely insignificant following death. In Section I of the Appendix we provide a detailed analysis that shows that the degree of informal insurance is relatively small for the affected group of households. One possibility that is plausibly significant is that the cost of living for a family is reduced following the death of a spouse, and we return to this issue in Section 4 below.

With our measure of the human capital loss in hand, we can now present our estimates of underinsurance. To this end, in this paper we introduce the life insurance coefficient defined as the ratio of life insurance holdings (face value) over the human capital loss in the case of death of an adult family member. Note that the size of the human capital loss in the case of death is equal to the amount necessary to achieve full insurance against mortality risk in our baseline model. Note further that our insurance coefficient differs from the insurance coefficient introduced by Blundell at al. (2008) in three important dimensions. First, Blundell et al. (2008) consider the consumption response to all income shocks due to events that are not further specified, whereas we confine attention to one clearly defined event, namely the death of a spouse. Second, Blundell et al (2008) separate income shocks into a transitory and a permanent component and then separately define insurance coefficients for each component. In contrast, we focus on human capital losses that are permanent and therefore only consider an insurance coefficient with respect to permanent income shocks. Finally, the insurance coefficient estimated by Blundell et al. (2008) lump together different mechanisms of consumption insurance: insurance through friends and family, self-insurance through own saving, and insurance through the purchase of insurance contracts. In contrast, we focus on one type of insurance mechanism: insurance through the purchase of insurance contracts.

Figure 1 plots the insurance coefficient for married households with children. The first line plots holdings for all households with kids and shows that the median household is insured against only one-tenth of the of the loss expected from the death of a spouse. This rises to roughly 50% by middle age, and to 75% by retirement. The second line depicts the same data for the sample of married households with kids that purchase some life insurance. This figure begins at roughly 30% and rises to close to 100% only as households reach their late

50's. This is our main empirical result: there exists a positive correlation between age and the degree to which households insure against mortality risk by purchasing life insurance. Further, young households are severely under-insured, whereas older households are almost fully insured.[11]

2.4 Human Capital and Financial Capital

We now turn to a discussion of the relative importance of human capital and financial capital, and how this importance varies over the life-cycle. To this end, we use earnings (labor income) as a proxy for human capital and construct the life-cycle of the ratio of financial wealth to earnings. Data on earnings and financial wealth are also drawn from the 6 surveys of the Survey of Consumer Finance (SCF) conducted between 1992 and 2007 and life-cycle profiles are constructed in the way described in Section 2.1. We continue to focus on married households with children. The variable "financial wealth" is defined as "net worth" in the SCF, which is the value of all assets (including housing and excluding human capital) minus the value of all debt (including mortgage debt). Labor earnings are defined as wages and salaries plus two-third of business and farm income.

Figure 5 plots the median ratio of net worth to labor earnings for married households with children of each age starting at 23 and ending at 60. As shown in the figure, households in their 20's and early 30's hold almost all of their wealth in human capital with the stock of net financial assets (including housing) less than one years flow of income from their human capital (i.e. labor earnings). By age 45, household net worth is roughly twice labor earnings, and it is not until households reach their 50's that net worth exceeds three time annual labor earnings.

The pattern in figure 5 is driven by the rapid accumulation of net financial assets, as labor earnings are also increasing in the early part of the life cycle. To illustrate this, figure 2 plots the lifetime profile of labor earnings derived from our data for households from age 23, when many households have left college, to age 60.[12] As is well-known, labor earnings

[11] The fact that households are underinsured even after conditioning only on those households that purchase life insurance suggests that fixed costs in the purchase of life insurance are not the only reason for the observed underinsurance.

[12] The age-earnings profile in figure 2 is computed from cross-sectional data, but a very similar concave life-cycle pattern emerges in studies that use panel data drawn from the PSID (Heathcote et al., 2010).

rapidly increase until age 35-40, after which they grow more modestly reaching a peak about age 50, and then declining as households approach retirement.

We follow a long tradition by interpreting these earnings profiles as the result of human capital accumulation decisions motivated by high returns to post-college education and on-the-job training (e.g. Becker 1964 and Ben-Porath 1967). There is a large literature estimating the returns to college-education, and some work on the returns to post-college education. Overall, the literature suggests a rate of return in the range of $8\% - 10\%$ (Krueger and Lindhal, 2001), though individual estimates vary considerably and there is a large amount of heterogeneity due to differences in ability (Cunha, Heckman, and Navarro, 2005, and Taber, 2001). Estimates of wage gains from on-the-job training imply rate of returns that are even higher than 10% (Blundell et al., 1999, and Mincer, 1994). Clearly, this evidence suggests that for many young households there is a strong incentive to invest in human capital. Moreover, this incentive exists regardless of whether returns are higher for the young because of decreasing returns to human capital accumulation, as in Ben-Porath (1967) and Huggett et al. (2011), or because human capital investment is less productive for older households, as in the model we describe next.

3. Model

In this section, we develop a general version of the model and discuss two theoretical results. The model structure is similar to the incomplete-market model with human capital developed in Krebs (2003), but adds life-cycle considerations and limited contract enforcement. Our first theoretical result is a convenient characterization of equilibria (proposition 1 and proposition 2) that highlights the tractability of the model. Wright (2001) provides a characterization result similar to propositions 1 and 2 for a class of linear production models (AK models) with limited commitment.[13] Proposition 3 is an analytical result and shows that, for a special case of the general model, age and insurance are positively correlated in

Further, this concave pattern is also observed for the earnings or wages of individual workers, though the household earnings profile lies, of course, strictly above the individual earnings profile (Heathcote et al., 2010, Huggett et al., 2011).

[13] Angeletos (2007) and Moll (2012) develop tractable models of entrepreneurial activity in which individual consumption/saving policies are also linear in wealth. In all these approaches, tractability is achieved through the assumption that individual investment returns are independent of household wealth.

equilibrium. Proofs are relegated to the Appendix.

3.1. Goods 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 good that can be consumed or invested in physical 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. This constant-returns-to-scale assumption in conjunction with the assumption that human capital is produced under constant-returns-to-scale (see below) implies that the model exhibits endogenous growth.

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 households of mass one. Households are indexed by their age j, the exogenous state (shock) s_j, their human capital, h_j, and their asset holdings, a_j. In our quantitative analysis, the exogenous state has two components, $s_j = (s_{1j}, s_{2j})$, where s_{1j} refers to the family state of the household and s_{2j} describes idiosyncratic labor market risk. The family state s_{1j} is defined by the marital status (married, widowed, single-not-widowed), the number of kids, and the gender in the case of a single household, for a total of 17 different states. Note that mortality risk corresponds to the transition from married household to widowed household. The process $\{s_j\}$ is Markov with stationary transition probabilities $\pi_j(s_{j+1}|s_j)$. We denote by $s^j = (s_1, \ldots, s_j)$ the history of exogenous states up to age j and let $\pi_j(s^j|s_0) = \pi_j(s_j|s_{j-1}) \ldots \pi_0(s_1|s_0)$ stand for the probability that s^j occurs given s_0. At age $j = 0$, a household begins life in the initial state (a_0, h_0, s_0).

The life of a household is divided into three phases. The first phase runs from age $j = 0, \ldots, J$ and is the focus of our analysis. In the quantitative application, we identify $j = 0$ with age 23 and $j = J$ with age 60. Thus, with 17 different family states and 38 different age-groups, we have $17 * 38 = 646$ different household types in the first phase of life. In this phase, households are working and married households face the (age-dependent) risk that an adult member of the household dies (mortality risk). For simplicity, we assume that the event that both adults die simultaneously has zero probability and, given our focus on married households with children, that single households do not face mortality risk. Married households also face the (age-dependent) risk of divorce. Single households meet with age-dependent probability to form a married household. Some households have children and the number of children in a household can increase or decrease by one. All transition probabilities over family states are exogenous. See Section 4 and the Appendix for more details about the specification of these transition probabilities.

The second phase of life, $j = J + 1$, is the pre-retirement stage. This phase is similar to the first phase, but now households do not age. However, they retire stochastically with fixed probability. The third and final phase of life is retirement. In the retirement phase, households receive no labor income and can only invest in a risk-free asset. Retired households die with constant probability and are then replaced by a household age $j = 0$ (age 23 in Section 4). Given that the focus of our analysis is on married households with kids in the first phase of life (i.e. with a household head not older than 60), we relegate to the Appendix a more detailed discussion of the decision problem in the second and third phase of life.

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. For a household choosing the consumption plan $\{c_j\}$, expected life-time utility is given by

$$\sum_{j=0}^{J} \beta^j \sum_{s^j} u(c_j, s_j) \pi_j(s^j|s_0) \\ + \beta^{J+1} \sum_{s^{J+1}} V_{J+1}(h_{J+1}(s^J), a_{J+1}(s^{J+1}), s_{J+1}) \pi_{J+1}(s^{J+1}|s_0) \tag{2}$$

where $u(c_j, s_j) = \gamma_0(s_j) + \gamma_1(s_j)\, lnc_j$ is the one-period utility function of the household, V_{J+1} is the value function in the pre-retirement stage and γ_0 and γ_1 are preference shocks that depend on the family state. For our baseline quantitative model we use $u(c, s) = ln(c)$ so that $\gamma_1(s_j) = 1$ (state-independent marginal utility of consumption). We then relax this assumption in Section 6. Note that we have abstracted from the labor-leisure choice of households.

Households can invest in physical capital as well as human capital and they can 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 $a_{j+1}(s_{j+1})$ the quantity bought (sold) at age (in period) j of the asset that pays off one unit of the good if s_{j+1} occurs at age $j+1$ (in the next period). Given an initial state, (h_0, a_0, s_0), a household chooses a plan, $\{c_j, h_{j+1}, \vec{a}_{j+1}\}$, where the notation $\vec{a}$ indicates that in each period the household chooses a vector of asset holdings. Further, c_j stands for the function mapping partial histories, s^j, into consumption levels, $c_j(s^j)$, with similar notation used for the other choice variables. A budget-feasible plan has to satisfy the sequential budget constraint, human capital evolution equation, and non-negativity constraints on total wealth (financial plus human), consumption and human capital

$$\begin{aligned} i) \quad && z(s_j) r_h h_j + a_j(s_j) &= c_j + x_{hj} + \sum_{s_{j+1}} a_{j+1}(s_{j+1}) q_j(s_{j+1}) \\ ii) \quad && h_{j+1} &= (1 - \delta_h + \eta_j(s_j)) h_j + \varphi_j(s_j) h_j + \phi x_{hj} \\ iii) \quad && h_j + \sum_{s_{j+1}} a_{j+1}(s_{j+1}) q_j(s_{j+1}) &\geq 0 \\ iv) \quad && c_j &\geq 0 \quad , \quad h_{j+1} \geq 0 \ , \end{aligned} \tag{3}$$

where $q_j(s_{j+1})$ is the price of a financial contract in period j that pays off if s_{j+1} occurs in $j+1$, which in our Markovian setting only depends on asset type s_{j+1} and current state s_j.

Note that the equations in (3) have to hold in realization; that is, they hold for all j and all sequences s^j. Note also that (3iii) represents a debt constraint, and that (3iv) requires the stock of human capital, h_{j+1}, to be non-negative, which prevents elderly workers from shorting their human capital.

The variable x_{hj} captures the resource cost of (active) human capital investment measured in consumption units and ϕ is a parameter describing the productivity of this investment. The term z_j in equation (3ii) is a labor productivity shock that captures transitory movements in earnings and we normalize its mean to one: $E[z_j] = 1$. In the human capital evolution equation, the term η_j measures the loss/gain of household human capital when there is a transition from married household to single households and vice versa (death of spouse, divorce, marriage).[14] The term $\varphi_j h_j$ represents increases in human capital that do not require an active input of resources, including returns to experience and learning-by-doing one's job, which are often referred to as experience capital. Note that this term has a random component so that φ_j also captures any labor market risk that is not part of the productivity shock z_j. The randomness in φ_j might describe variations in the return to experience that often occur when workers switch their employer and/or occupation. For our quantitative analysis, we assume $\varphi_j(s_j) = \bar{\varphi}_j + \hat{\varphi}(s_j)$, where the first component describes age-dependent learning-by-doing effects and the second component captures labor market risk. Finally, δ_h is the depreciation rate of human capital.

Note that the budget constraint (3) assumes that physical capital and human capital are produced using similar technologies in the sense that one unit of physical capital can be transformed into ϕ units of human capital. Thus, we assume constant returns to scale at the household level. This assumption, also made in Krebs (2003), implies that the household decision problem displays a certain linearity with respect to physical capital investment and human capital investment in the sense that goods invested in either human capital or physical capital generate returns that are independent of household size, where size is measured by total wealth (see below). For simplicity, we also assume that the time cost of human capital investment is exogenous, but this assumption is not essential for our tractability result.

The assumptions we make in (3) have the advantage that they keep the model highly

[14] For notational ease, we expand the family state, s_{1j}, to include last period's marital status (married, single widowed, single not widowed) so that the η-shock only depends on the current s_{1j} and not on $s_{1,j-1}$.

tractable, which, as we argued before, is essential for the theoretical and quantitative analysis conducted in this paper. Tractability in the general case requires that we do not impose a restriction on the ability of households to decumulate human capital. However, in the calibrated model economy used for our quantitative analysis, the restriction $h_{j+1} \geq (1-\delta_{hj})h_j$ is always satisfied in equilibrium; that is, it holds for all household types at all ages and all realizations of uncertainty. Similarly, the restriction that human capital investment is non-negative, $\varphi_j h_j + \phi x_{hj} \geq 0$, is always satisfied in equilibrium. Thus, imposing these restrictions in (3) would not change the conclusions drawn in the quantitative analysis. [15]

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

$$\begin{aligned} & \sum_{n=0}^{J-j} \beta^n \sum_{s^{j+n}|s^j} u(c_{j+n}, s_{j+n}) \pi_j(s^{j+n}|s_j) \\ + \quad & \beta^{J+1-j} \sum_{s^{J+1}|s^j} V_{J+1}(h_{J+1}(s^J), a_{J+1}(s^{J+1}), s_{J+1}) \pi_{J+1}(s^{J+1}|s_j) \\ \geq \quad & V_d(h_j(s^{j-1}), a_j(s^j), s_j) \end{aligned} \tag{4}$$

for all j and s^j, where V_d is the value function of a household who defaults. Note 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. This is the non-convexity issue alluded to in the introduction; in proposition 1 we show how this problem is surmounted in the current setting.

The default value function V_d, is defined by the household decision problem in default. In this paper, we allow for different specifications of this default problem. In the baseline version of the model, we model default along the lines of Chapter 7 of the US bankruptcy code. More precisely, we assume that upon default all debts of the household are canceled and all financial assets seized so that $a_j(s_j) = 0$. Following default, households are excluded

[15] Note that in (3) we have explicitly imposed a non-negativity constraint on the stock of human capital, and our general characterization of the household decision rule (proposition 1) holds with this constraint imposed. Of course, for a certain range of parameter values this constraint binds in equilibrium, but for the parameter values used in our quantitative analysis this constraints never binds (does not bind for all households types and uncertainty states).

from purchasing insurance contracts and borrowing (going short), but they can still save in a risk-free asset. We assume that 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 r_h per unit of human capital. 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.[16] Details are found in the Appendix.

3.3 Equilibrium

We assume that insurance markets (financial markets) are perfectly competitive and abstract from transactions costs. Thus, insurance contracts (financial contracts) are priced in an actuarially fair manner (risk neutral pricing):

$$q_j(s_{j+1}; s_j) = \frac{\pi_j(s_{j+1}|s_j)}{1 + r_f} \ . \tag{5}$$

The pricing equation (5) can be interpreted as a zero-profit condition for financial intermediaries that can invest in physical capital at the risk-free rate of return $r_f = r_k - \delta_k$ and can fully diversify idiosyncratic risk for each insurance contract s_{j+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. This transition function together with the initial distribution, μ_0, and sequence of distributions for new-born households, $\{\mu_{t,new}\}$, induce a sequence of equilibrium distributions, $\{\mu_t\}$, of households over individual states. We assume that the

[16] The previous literature has usually assumed $p = 1$ (permanent autarky). See, however, Krueger and Uhlig (2006) for a model with $p = 0$ 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.

financial capital of households who die is inherited by new-born households, which imposes a restriction on the mean of the marginal distribution $\mu^m_{t,new}$ over a. Note that we allow the distribution $\{\mu_{t,new}\}$ to have an explicit time-dependence since in our endogenous growth model the mean value of h and a will growth over time, and a stationary distribution over intensive-form or growth-adjusted variables can only be obtained if the mean value of the extensive-form variables also grows for new-born households. In our quantitative analysis, we directly specify the distribution of new-born households over growth-adjusted states.

Assuming a law of large numbers, aggregate variables can be found by taking expectation 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 $H_t = \int \sum_j h_j d\mu_{tj}(h_j)$. 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

$$\tilde{K} = \frac{\int \sum_j \sum_{s_{j+1}} q_j(s_{j+1}; s_j) a_{j+1}(s_{j+1}) d\mu_{jt}(h_j)}{\int \sum_j h_j d\mu_{tj}(h_j)} \ , \qquad (6)$$

where $\tilde{K}$ is the capital-to-labor ratio chosen by the firm that determines the equilibrium rental rates of physical capital and human capital, r_k and r_h, through (1).

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), such 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 household plan 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 (5)
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 (6) holds for all t when the expectation is taken with respect to the distribution μ_t.
vi) Expected household value function is identical to the household value function: $V^e = V$.

Note that the equilibrium value of $\tilde{K}$ determines the equilibrium growth rate of the economy (see Appendix for details). Note also that in equilibrium the goods market clearing condition (aggregate resource constraint) automatically holds:

$$C_t + K_{t+1} + \frac{1}{\phi}H_{t+1} = (1-\delta_k)K_t + \frac{1}{\phi}H_t + \frac{1}{\phi}\int\sum_j(\eta_j(s_j)+\varphi_j(s_j)-\delta_h)h_j d\mu_{tj}(h_j) + F(K_t, H_t) \tag{7}$$

3.4. Characterization of Household Problem

We next show that optimal consumption choices are linear in total wealth (human plus financial) and portfolio choices are independent of wealth. This property of the optimal policy function allows us to solve the quantitative model, which has a large number of household types and uncertainty states, without using approximation methods. The property also implies that the household decision problem is convex and the first-order approach can be utilized.

To state the characterization result, denote total wealth (human plus financial) of a household of age j at the beginning of the year by $w_j = h_j/\phi + \sum_{s_j} a_j(s_j)q_{j-1}(s_j)$. Note that ϕ measures the productivity of goods investment in human capital and $1/\phi$ is is the shadow price of one unit of human capital in terms of the consumption/capital good. Denote the portfolio shares by $\theta_{hj} = h_j/(\phi w_j)$ and $\theta_{a,j}(s_j) = a_j(s_j)/w_j$. The sequential budget constraint (3) then reads:

$$\begin{aligned} w_{j+1} &= (1 + r_j(\theta_j, s_j))w_j - c_j \\ 1 &= \theta_{h,j+1} + \sum_{s_{j+1}} q_j(s_{j+1}|s_j)\theta_{a,j+1}(s_{j+1}) \\ c_j &\geq 0 \ , \ w_{j+1} \geq 0 \ , \ \theta_{j+1} \geq 0 \end{aligned} \tag{8}$$

with

$$1 + r_j(\theta_j, s_j) \doteq [1 + \phi z(s_j) r_h - \delta_h + \eta_j(s_j) + \varphi_j(s_j)]\, \theta_{hj} + \theta_{aj}(s_j)$$

Clearly, this is the budget constraint corresponding to an inter-temporal portfolio choice problem with linear investment opportunities and no exogenous source of income. It also shows that (w, θ, s) can be used as individual state variable for the recursive formulation of the utility maximization problem. Using this notation, we have the following result:

Proposition 1. The value function and the optimal policy function are given by

$$\begin{aligned} V_j(w_j, \theta_j, s_j) &= \tilde{V}_{0j}(s_j) + \tilde{V}_{1j}(s_j)\,[\ln w_j + \ln(1 + r_j(\theta_j, s_j))] \\ c_j(w_j, \theta_j, s_j) &= \tilde{c}_j(s_j)\,(1 + r_j(\theta_j, s_j))\, w_j \\ \theta_{j+1}(w_j, \theta_j, s_j) &= \theta^*_{j+1} \\ w_{j+1}(w_j, \theta_j, s_j) &= (1 - \tilde{c}_j(s_j))\,(1 + r_j(\theta_j, s_j))\, w_j \end{aligned} \tag{9}$$

where the value function coefficients, $\tilde{V}_{0j}(s_j)$, $\tilde{V}_{d,0j}(s_j)$, and $\tilde{V}_{1j}(s_j)$ as well as the optimal consumption-to-wealth ratio, $\tilde{c}$, and the optimal portfolio choice, θ^*_{j+1} are the solution to a maximization problem with linear constraints (see the Appendix).[17]

Proof: See Appendix.

Proposition 1 provides a convenient characterization of the solution to the household decision problem for given investment returns (partial equilibrium). We next turn to the determination of investment returns (general equilibrium).

3.5. Equilibrium Characterization

Define the share of aggregate total wealth of households of age j and state s_j as

$$\Omega_j(s_j) \doteq \frac{E\,[(1 + r_j)w_j|s_j]\,\pi_j(s_j)}{\sum_j \sum_{s_j} E\,[(1 + r_j)w_j|s_j]\,\pi_j(s_j)}$$

Note that $(1+r_j)w_j$ is total wealth of an individual household after assets have paid off (after production and depreciation has been taken into account). Note also that $\sum_j \sum_{s_{1j}} \Omega(s_{1j}) =$

[17]The Appendix also contains the corresponding expressions for the default value function and default consumption policy.

1. Further, Ω is finite-dimensional, whereas the set of distributions over (w, s) is infinite-dimensional. Using the definition of wealth shares and the property that portfolio choices are wealth-independent, in the Appendix we show the following result:

Proposition 2. Suppose that $(\theta, \tilde{c}, \tilde{V}, \tilde{K}, \Omega)$ solves the fixed-point problem defined by the equations (A4), (A10), and (A11) in the Appendix. Then $(g, \tilde{V}, \tilde{K}, \{\mu_t\})$ is a stationary (balanced growth) equilibrium, where g is the individual policy function induced by $(\tilde{c}, \theta)$ and $\{\mu_t\}$ is the sequence of measures induced by the policy function g, the initial measure, μ_0, and the transition matrix over demographic and labor market states, π.

Proof. See the Appendix.

Proposition 2 shows that the stationary equilibrium can be found without knowledge of the infinite-dimensional wealth distribution – only the lower dimensional distribution Ω matters. Proposition 2 facilitates our quantitative analysis significantly since it implies that there is no need to keep track of the entire wealth distribution when computing equilibria.

3.6. Analytical Results

We now derive analytical results for a special case of the model. We use these results to discuss the main determinants of individual consumption, and to prove that in equilibrium there is a positive relationship between age and insurance.

We focus on the first phase of life, $j = 1, \ldots, J$, and on households with two adult members (married households). We consider the case with only mortality risk so that $s_j \in \{d, n\}$, where $s_j = d$ is the event that the death of an adult household member occurs and $s_j = n$ is the event that death does not occur. Note that after the event "death of an adult household member" the household continues to exist. Mortality risk is an i.i.d. random variable, η, with age-independent probability π that death occurs and age-independent human capital loss $\eta(d)$ in the death event. We normalize the mean of η to zero: $\pi\eta(d) + (1 - \pi)\eta(n) = 0$. Note that $\eta(d)$ is the fraction of household human capital that is lost in the event that an adult member of the household dies. We also assume constant labor productivity $z(s_j) = 1$ and no human capital risk beyond mortality risk: $\varphi_j(s_j) = \bar{\varphi}_j$.

We assume that young households have a higher rate at which they gain work experience on the job, $\bar{\varphi}_j > \bar{\varphi}_{j+1}$, which implies that expected human capital returns of young households are larger than the returns of older households. We choose state-independent preference parameters $\gamma_0(s_j) = \gamma_1(s_j) = 1$. Finally, we assume that defaulting households are not excluded form financial markets, $p = 0$, which rules out short positions in financial assets (see Appendix).

Using the policy function (9) of our equilibrium characterization result, we find that in this example consumption growth is given by:

$$\begin{aligned} \frac{c_{j+1}}{c_j} &= \beta(1 + r_{j+1}(\theta_{j+1}, s_{j+1})) \\ &= \beta\left\{(1 + \phi r_h - \delta_h + \varphi_{j+1} + \eta(s_{j+1}))\,\theta_{h,j+1} + \theta_{a,j+1}(s_{j+1})\right\} \end{aligned} \tag{10}$$

Consumption growth depends on human capital choice, $\theta_{h,j+1}$, ex-ante human capital returns, $\phi r_h - \delta_h + \varphi_{j+1}$, ex-post shocks, $\eta(s_{j+1})$, and asset payoffs (insurance), $\theta_{a,j+1}(s_{j+1})$. From (10) we immediately conclude that consumption is independent of mortality shocks if $\theta_{a,j+1}(d) - E[\theta_{a,j+1}] = \eta(d)\,\theta_{h,j+1}$, where $E[\theta_{a,j+1}] = \pi\theta_{a,j+1}(d) + (1-\pi)\theta_{a,j+1}(n)$ is the fraction of total wealth the household is holding as financial wealth. This is intuitive since $(\theta_{a,j+1}(d) - E[\theta_{a,j+1}])\,w_{j+1}$ is the insurance pay-out in the case of death and $\eta(d)\,\theta_{h,j+1}w_{j+1}$ is the human capital loss in the case of death, and when the two are equal we have full insurance and therefore deterministic consumption growth.

The above discussion demonstrates that in our model of mortality risk without "preference shocks" full insurance is achieved when the actual insurance payout, $(\theta_{a,j+1}(d) - E[\theta_{a,j+1}])\,w_{j+1}$, is equal to the human capital loss in the case of death, $\eta(d)\,\theta_{h,j+1}w_{j+1}$. Thus, we can define the life insurance coefficient as the ratio:

$$I_{j+1} = \frac{\theta_{a,j+1}(d) - E[\theta_{a,j+1}]}{\eta(d)\,\theta_{h,j+1}} \tag{11}$$

Clearly, in any model without "preference shocks" the life insurance coefficient I varies between 0 (no insurance) and 1 (full insurance). For our example economy we have the following result:

Proposition 3. Suppose the economy is as described above. In equilibrium, young households are less insured than old households and a larger fraction of their total wealth is invested in human capital:

$$\begin{aligned} \theta_{h,j} &\geq \theta_{h,j+1} \\ I_j &\leq I_{j+1}\ , \end{aligned}$$

where the inequalities are strict if in equilibrium there is some insurance, but not full insurance.

Proof. See the Appendix.

Next we establish this result in a calibrated version of our general model.

4. Calibrating the Model

We now turn to the quantitative analysis. Section 4.1 lays out the model specification. Section 4.2 discusses our calibration strategy and the relevant empirical literature, while Section 4.3 discusses computation.

4.1 Model Specification

We set the length of a time period to one year and let $j = 23, \ldots, 60$ for the first phase of life. As in Huggett et al. (2011), we restrict attention to households up to age 60 for the following three reasons. First, the number of households with positive labor income in each age group drops rapidly in our SCF sample after age 60. Second, labor force participation falls near the traditional retirement age for reasons that are not modeled 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 only up to age 60 and introducing a transition-group of households with stochastic retirement, we can ensure that for the calibrated model economy the rate of decumulation of human capital is bounded by the rate of depreciation over the entire life-cycle.

For our baseline model, we assume a one-period household utility function $u(c, s) = ln(c)$. In other words, we assume a constant marginal utility of household consumption and normalize the constant to one: $\gamma_1(s) = 1$. Note that for this preference specification we have $\tilde{c}_j = 1 - \beta$ in (9). Note also that with this specification utility is defined over household consumption without any adjustment for household size. In order to take into account the insurance effect of any reduction in living costs resulting from a smaller family size, below we scale the income losses associated with the death of the husband/wife according to a cost of living adjustment (see Section 4.2 below). We choose this simple preference specification to focus attention on our basic mechanism. Specifically, this preference specification ensures that the life cycle pattern of under-insurance generated by the model is solely due to the interaction of human capital risk with endogenous borrowing constraints. We discuss the possibility of life cycle variation in preferences in Section 6.3.

We assume that the exogenous state variable has two components. The first component describes the family state that is a pair (m_j, k_j), where m_j is the marital state and k_j denotes the number of kids. We assume $k_j \in \{0, 1, 2, 3\}$ and $m_j \in \{ma, fw, fn, mw, mn\}$ corresponding to married (ma), female widowed (fw), female not widowed (fn), male widowed (mw), and male not widowed (mn). Thus, we have in total 17 family states (we do not have non-widowed single male households with children since we assume that after divorce children live with their mother). The transition matrix over family states is discussed in more detail in the Appendix. Transitions across marital states, m_j, lead to changes in the stock of human capital denoted by η_j. In particular, η_j captures the human capital losses in the case of the death of a spouse (transition from $s_{1j} = ma$ to $s_{1,j+1} = fw$ or $s_{1,j+1} = mw$). We parameterize these human capital losses as follows.

We assume that for a fraction π_j of households we have $\eta_j = \alpha\tilde{\eta}_j$ and for a fraction $(1-\pi_j)$ of households we have $\eta_j = 0$, where α is distributed according to a generalized Pareto distribution with cumulative distribution function $F(\alpha) = 1 - \left(1 + \frac{\psi(\alpha-\mu)}{\sigma_\alpha}\right)^{-\frac{1}{\psi}}$. Households with $\eta_j > 0$ buy life insurance and households with $\eta_j = 0$ do not. Thus, π_j is the participation rate in the life insurance market. Below we calibrate $\tilde{\eta}_j$ to the match the present value income losses of married households with median earnings as discussed in Section 2.3, and we use α to capture the heterogeneity of life insurance holdings conditional on age j for those households who have purchased life insurance. We assume that households with $\eta_j > 0$ keep their value until retirement and that households with $\eta_j = 0$ draw a new value in the next year $j+1$. This new value is drawn from the distribution of $\alpha\tilde{\eta}_{j+1}$ with probability q_{j+1} and is equal to zero (no participation) with probability $(1 - q_{j+1})$. Thus, the participation rate evolves according to $\pi_{j+1} = \pi_j + q_j(1 - \pi_j)$. We assume that probability q_{j+1} has a linear trend: $q_{j+1} = (1 - d)q_j$.

The second component of the exogenous state describes labor market risk specified by the two variables z_j and $\varphi_j = \bar{\varphi}_j + \hat{\varphi}$. We assume that $\bar{\varphi}_j$ is deterministic and that productivity shocks, z_j, and human capital shocks, $\hat{\varphi}_j$, are i.i.d. with a finite, symmetric distribution that approximates a normal distribution. The assumption that human capital shocks are independently and (approximately) normally distributed is also made by Huggett et al.

(2011) and Krebs (2003). We assume that z_j has mean 1 and $\hat{\varphi}_j$ has mean 0 and denote variances of these two random variables by σ_z^2 and σ_φ^2, respectively.

Households in pre-retirement age (age $J = 61$) work and the duration of this phase of life ends stochastically with retirement. Households age $J = 61$ solve a recursive version of the household decision problem described in Section 3 (see also the Appendix). Upon retirement, the human capital of households becomes unproductive.[18] Retired households can save in a risk-free asset. Households age $j = 23, \ldots, 61$ do not die and retired households die stochastically, in which case they are replaced by a new-born household of age 23. The financial capital of deceased households is passed on to new-born households. New-born households also receive an initial endowment of human capital. The distribution of new-born households over human capital, physical capital, and family states is discussed in Section 4.2 below.

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

4.2 Model Calibration

4.2.1 Mortality Risk

Mortality risk is captured in the model by the transition from the marital state $m_j = ma$ to the marital state $m_{j+1} = mw$ (female spouse dies, producing a widower) or $m_{j+1} = fw$ (male spouse dies). We choose the probability that a male or female spouse dies to match the year-to-year average survival rates for the period 1991-2000 for the US life-tables for the respective group (see figure A5 in the Appendix). We use the re-marriage rates of divorcees from the SIPP as a proxy for the re-marriage rates of widows/widowers, but introduce an adjustment to take into account of the evidence that indicates lower re-marriage rates for widows and widowers. Specifically, we compute the life-cycle profile of re-marriage rates of female/male divorcees from the SIPP and then scale down this life-cycle profile so that the

[18]Formally, we assume that the labor productivity of retired households drops to zero. To avoid that households sell their human capital upon retirement, we also assume 100 percent depreciation of existing human capital in the first period of retirement. Of course, in equilibrium households are almost fully insured against this retirement-shock. Alternatively, we can assume that upon retirement households receive a lump-sum payment from the government.

average marriage rate corresponds to the average re-marriage rate of widows/widowers in the SIPP data. The result is depicted in figure A1 in the Appendix, and is in line with the evidence of re-marriage rates of widows and widowers presented in Norton and Miller (1990) and Wilson and Clarke (1992).

The size of the negative human capital shock in the case of the death of an adult household member, η_j, is calibrated as follows. Recall that we assume $\eta_j = \alpha\tilde{\eta}_j$ with probability π_j, where π_j is the fraction of households who participate in the life insurance market and α is distributed according to a generalized Pareto distribution. For married households of age j with number of kids k_j, we calibrate $\tilde{\eta}_j$ to the match the present value income losses as discussed in Section 2.3, but here we compute the losses separately for the case of death of the husband and death of the wife. The losses we compute are the losses for a married household with median earnings and they take into account the implicit insurance that arises from social security survivor benefits, progressive income taxation, and re-marriage. In addition, we make an adjustment to the human capital losses to take into account the insurance effect of any reduction in living costs resulting from a smaller family size, which we have ignored so far. We account for this effect by scaling the income losses according to the cost of living adjustment suggested by the consumption equivalence scale of Ruggles (1990). The resulting life-cycle profiles of human capital losses in case of death of the husband, respectively wife, are shown in figures A7 and A8 in the Appendix.[19]

We calibrate the remaining parameters ψ, μ, σ_α, π_{23}, q_{23} and d determining mortality risk as follows. We choose the values of the generalized Pareto distribution, ψ, μ, and σ_α, to match three moments of the empirical distribution (not conditional on age) of life insurance holdings for married households with children who have purchased some insurance: the mean, median, and the upper 10 percentile. The resulting parameter values are $\psi = 0.4714$, $\mu = 0.4876$, and $\sigma_\alpha = 0.6250$. We set the value of π_{23} equal to the empirical participation rate of 23-year old married households with children and choose q_{23} and d so as to match the

[19]Note that our approach to estimating the human capital loss, η, takes into account the implicit insurance provided by the possibility of re-marriage since re-marriage probabilities enter into our calculation of the expected income losses depicted in figures A7 and A8. Consequently, we do not incorporate into the model a positive human capital shock upon re-marriage to avoid double-counting. This approach also ensures that monetary gains from re-marriage are treated in the same way as monetary gains from the social security survivor benefits – both enter into the model only through the size of the human capital shock η.

life cycle profile of the observed participation rates for this group (we minimize the squared distance between model and data).

4.2.2 Divorce Risk

Divorce risk is captured in the model by the transition from the marital state $m_j = ma$ to the marital state $m_{j+1} = fn$ and $m_{j+1} = mn$. We choose the age-dependent probability of divorce so that we match the corresponding separation rates in the SIPP data (see figure A2 in the Appendix). After divorce, the new single-female household receives x percent of the total household human capital and the new single-male household y percent of the total household human capital. The number x is the ratio of median earnings of a single-female household over median earnings of a married household. We assume that after divorce the financial wealth is split equally between the man and the woman.

4.2.3 Labor Market Risk

We calibrate the two parameters σ_z^2 and σ_φ^2 as follows. Conditional on the family structure, our human capital model implies a labor income process that is consistent with the specification used by the empirical literature on labor market risk. Indeed, there is a one-to-one mapping between the two parameters σ_z^2 and σ_φ^2 and the labor income risk parameters estimated by the empirical literature. Specifically, a large literature (Carroll and Samwick 1997, Meghir and Pistaferri 2004, Storesletten et al. 2004b) has estimated transitory and permanent labor income risk as follows. Observed labor income, y_j, is decomposed into a transitory component, y_j^T, and a permanent/persistent component, y_j^p, with $\ln y_j = \ln y_j^T + \ln y_j^p$, where the transitory component is an i.i.d. process with $\ln y_j^T \sim N(0, \sigma_T^2)$ and the permanent component is a logarithmic random walk with innovation term $\epsilon \sim N(0, \sigma_\epsilon^2)$. The two variances σ_T^2 and σ_ϵ^2 can then be estimated separately using various moment conditions (Meghir and Pistaferri, 2004). We next show that our model specification leads to a labor income process with an i.i.d. component and a random walk component, and that estimates of σ_T^2 and σ_ϵ^2 therefore provide us with estimates for σ_z^2 and σ_φ^2.

Labor income in the model is given by $y_j = z_j r_h h_j$ so that $\ln y_j = \ln z_j + \ln r_h + \ln h_j$. In equilibrium $\ln h_j$ follows a random walk (see below) so that it can be identified with the permanent component, $\ln y_j^p$. Since z_j is i.i.d. it can be identified with the transitory component y_j^T ($\ln r_h$ is a constant) and we therefore have $\sigma_z^2 = \sigma_T^2$. Given a value for σ_z^2, a value for σ_φ^2 can be found as follows. Consider the evolution of the human capital stock h_j

$$\begin{aligned} \ln h_{j+1} - \ln h_j &= \ln \theta_{h,j+1} - \ln \theta_{hj} + \ln \beta + \ln(1 + r_j(\theta_h, s_j)) \\ &\approx \ln \theta_{h,j+1} - \ln \theta_{hj} + \ln \beta + z(s_{2j})\phi r_h - \delta_h + \eta_j(s_{1j}) + \bar{\varphi}_j + \hat{\varphi}(s_{2j}) \end{aligned} \tag{12}$$

where we used the equilibrium policy for human capital and the approximation $ln(1+x) \approx x$. Equation (12) shows that conditional on family structure, human capital in the model follow a random walk with age-dependent drift and innovation term that is approximately normally distributed with variance $\sigma_\epsilon^2 = \phi^2 r_h^2 \sigma_z^2 + \sigma_\varphi^2$. Given values for ϕ, r_h, σ_ϵ^2, and σ_z^2 we find a value for σ_φ^2 using $\sigma_\varphi^2 = \sigma_\epsilon^2 - \phi^2 r_h^2 \sigma_z^2$.

The discussion shows how estimates of the transitory and permanent component of labor market risk, σ_T^2 and σ_ϵ^2, provide us with estimates of σ_z^2 and σ_φ^2 (for given values of ϕ and r_h). Estimates of σ_T and σ_ϵ vary considerably, with a midpoint of around 0.3 and 0.15, respectively. For σ_T^2 we choose the midpoint of 0.30. For the standard deviation of the permanent component we follow Huggett et al (2011) and choose a somewhat lower value, namely 0.0123. This choice is motivated by the fact that estimates of permanent labor income risk will overstate the true value of the variance if there is earnings profile heterogeneity in addition to stochastic shocks with a permanent component (Baker and Solon 2003 and Guvenen 2007).

4.2.4 Investment Returns

We calibrate an annual risk-free rate of $r_f = 3\%$, in line with Kaplan and Violante (2010) and roughly in line with Huggett et al. (2011) and Krueger and Perri (2006) who use a 4% annual risk-free rate, but also deduct capital income taxes.

The depreciation rate of human capital, δ_h, is chosen to match a target value for the ratio of human capital investment to GDP (see below). For given value of δ_h, we choose the age-dependent learning-by-doing parameters $\bar{\varphi}_j$ to match the life-cycle profile of earnings of the median household in our sample. Specifically, we assume that the age-dependence is described by an exponential function, $\bar{\varphi}_j = A + Be^{-Cj}$, and choose the coefficients A, B, and C in order to minimize the distance (L2-norm) between the empirical life-cycle of median earnings from age 23 to age 60 and the corresponding model prediction. The implied life-cycle profile of $\bar{\varphi}_j$ is shown in figure A9 in the Appendix.

4.2.5 Bankruptcy Code

We calibrate the costs of default to match features of the U.S. bankruptcy code. Specifically, we assume that households forfeit all financial assets, experience no garnishment of labor income, and are unable to borrow or buy insurance products for an average length of 7 years, so that the probability of re-establishing full financial market access is $(1-p) = 1/7$.[20] Households in default may save in the risk-free asset, and may continue to rent their human capital to firms.

4.2.6 Preferences, Endowment, and Production

We follow Huggett et al. (2011) and assume a capital share in output, α, of .32, and target an aggregate capital to output ratio of 2.94. We also target an aggregate ratio of physical capital to human capital, $\tilde{K}$, of 0.4. This value lies somewhere in the middle of the range of estimates by the empirical literature, which suggests that the aggregate stock of human capital is $2-4$ times larger than the aggregate stock of physical capital (Jorgenson and Fraumeni, 1993). Together with the interest rate target of 3 percent, these requirements pin down the parameter values $\delta_k = 0.0785$ and $A = 0.1818$.

The retirement probability of households is chosen so that retirement occurs on average at age 65 and the death probability of retired households is chosen so that the expected age of death is 85. We choose an annual discount factor $\beta = 0.95$ and the human capital productivity parameter ϕ to match the average of the ratio of financial wealth to labor income, which yields $\phi = 0.675$. We choose the frequency distribution of newborn households (households age 23) over family types s_1 in the model equal to the empirical distribution.

Newborn households inherit the financial capital of deceased retired households and receive an initial endowment of human capital. We assume that all newborn households of a particular family type s_1 have the same endowment of financial and human capital, and assign initial endowments of human capital to different family types so that we match the empirical distribution of earnings across family type at age 23.

[20] Our parameterization is bracketed by Krueger and Perri (2006), who assume $(1-p) = 0$, Chatterjee et al. (2007), who use $(1-p) = 1/10$, and Livshits et al. (2007), who use $(1-p) = 1$ following the first period of default. The degree of variation in the parameter p reflects the fact that, as in our model, these papers abstract from a number of the costs of consumer default, and hence the calibration of the parameter p in part is a proxy for other default costs. In light of this, some authors have argued that the parameters governing the cost of default should be calibrated to match some aspect of the data, as in the choice of the level of wage garnishment in Livshits et al. (2007).

The calibration procedure described so far leaves free two parameter values: The human capital depreciation rate and the total human capital endowment of 23-year old households. We choose the values of these two parameters to match a value of 0.4 for the ratio $\tilde{K}$ (see above) and a value of 6 percent for the aggregate ratio of human capital investment to output, X_h/Y. The value of 6 percent for X_h/Y is based on the work by Mincer about the cost of investment in on-the-job training. Specifically, using information on time spent in training and wages, Mincer (1989) estimates that the total volume of job training investment by US companies amounted to 296 billion in 1987, which results in a value of 6 percent for X_h/Y. Mincer (1989) also argues that this number is broadly consistent with a number of alternative empirical studies about on-the-job training.

The human capital depreciation rate, δ_h, implied by our calibration is 4.29 percent annually. This value lies within the range of values used in the literature. Specifcally, Mincer (1989) finds an annual depreciation rates for human capital accumulated on the job of about 4 percent. The literature on education usually finds somewhat lower values – see Browning, Hansen, and Heckman (1999) for a survey. However, macro work on human capital accumulation has often used calibrated models with annual human capital depreciation rates of 6 percent (Manuelli and Jones, 1990, and Krebs, 2003).

4.3 Computation of Equilibrium

The computation of equilibria is based on propositions 1 and 2. More specifically, 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 (proposition 1). Given the solution to the household problem, we compute a stationary relative wealth distribution, Ω, using the law of motion described in the Appendix (proposition 2). We use this Ω to compute a new $\tilde{K}$ and iterate over $\tilde{K}$ until the clearing holds. A detailed description of our solution method can be found in the Appendix.

Table 1 shows the parameter values of the calibrated model economy together with the targets used to calibrate the model. The model also has implications for additional aggregate macro statistics. Specifcally, the calibrated model implies an annual aggregate consumption growth rate of 0.67 percent and an investment-to-output ratio, X_k/Y of 23 percent. These values are in line with the corresponding values for the US economy.

5. Results

We next present the results of our quantitative analysis of the model. Section 5.1 compares the models implications for insurance to the data, beginning with life insurance holdings and concluding with a description of consumption insurance. Section 5.2 discusses the models implications for wealth and the cross sectional variation in consumption levels. Section 5.3 presents the welfare consequences of mortality risk. Section 5.4 discusses our policy experiment: a reform of the bankruptcy code.

5.1 Insurance

5.1.1 Insurance of Median Household

In this subsection we assess the ability of the model to reproduce the life-cycle pattern of life insurance holdings. We begin the analysis on the median household conditional on age and the purchase of life insurance (intensive margin), and then turn to a discussion of heterogeneity within age groups in Section 5.1.2 below. Figure 6 shows the life-cycle profile of median life insurance holdings of married households with children who have purchased life insurance (intensive margin) in the data and as predicted by the model. As discussed before, the data display an inverted u-shape pattern: the median young married household holds around $85 thousand in life insurance, rising quickly to $200 thousand before declining slowly down to $75 thousand in their early 60's. Figure 6 shows that the model is able to match these data both qualitatively and quantitatively.

We next evaluate the extent of underinsurance for the group of households depicted in figure 6 by analyzing the corresponding life insurance coefficient. Recall that in this paper we define the life insurance coefficient as the ratio of insurance payout in the case of death (actual insurance) over the human capital loss in the case of death (insurance need). In figure 7 we depict the life cycle profile of the life insurance coefficient when we proxy the human capital loss by the present value losses in the case of death, computed as described in Section 2.3. As in figure 6, the model provides an excellent quantitative account of the data, which is expected given that the fit between model and data in figure 6 is very good. Figure 7 shows that under-insurance for young households is severe and that insurance is strongly correlated with age: both the data and the model, young households are insured against roughly 30% of their potential loss, with the figure rising close to 100% only as households reach their late 50's.

The measure of underinsurance in figure 7 is based on estimated present value income losses. Alternatively, we can compute life insurance coefficients using the model-based measure of human capital loss in the denominator of the insurance coefficient. The result of this computation is depicted in figure 8 and confirms the result already shown in figure 7: both in the data and in the model, there exists a strong correlation between age and the degree to which households purchase insurance against mortality risk. Indeed, the correlation between age and life insurance coefficient is roughly the same in both figures; the only difference is that young households are insured against 30% of the losses implied by the present value of income losses and against 40% of the losses according to the model-based calculation. This result is expected given that the human capital losses in the calibrated model economy are in line with the estimated net present value losses.

In the model, the human capital losses in the case of a husband's death are different from the human capital losses in the case of a wife's death. Consequently, the life insurance holdings for the two events differ. The SCF does not provide information about the split of insurance between husband and wife, but the SIPP data provide information about this split. In figure 9 we plot life insurance holdings separately for husband and wife, where we again focus on married households with children who have purchased some life insurance. The data show that in both cases there is an inverted u-shape, but this inverted u-shape is much more pronounced for insurance against the husband's death. Further, life insurance against the husband's death is about twice as much as life insurance against the wife's death. Figure 9 also shows that the model provides a good quantitative account of both life-cycle profiles, though on average the model slightly under-predicts holdings of insurance against the death of the husband, and slightly over-predicts the holdings of insurance against the death of the wife.

Figures 6-9 analyze to what extent households insure against mortality risk. Households in the model also face labor market risk and demographic risk in addition to mortality risk, and we next investigate the extent of insurance against all types of risk. To this end, we consider the model implication for the life-cycle variation of consumption insurance measured as lack of consumption volatility. More precisely, we define an insurance measure $1-\sigma_c/\sigma_{a,c}$, where σ_c is the standard deviation of equilibrium consumption growth for households with full access to financial markets and $\sigma_{a,c}$ is the standard deviation of equilibrium consumption growth for households with no access to credit and insurance markets (but the possibility of saving in a risk-free asset). Figure 10 shows that consumption insurance increases substantially with age. For example, the value of our insurance measure begins at 0.24 for households age 23 and increases to 0.81 for households age 60.

5.1.2 Heterogeneity in Insurance

In this subsection we move beyond the focus on the median married households with children who has purchased some life insurance and discuss additional moments of life insurance holdings. We begin with a discussion of the extensive and intensive margin. Figure 11 depicts the life-cycle profile of the median life insurance holdings for all married households with children and those married households with children who have purchased some life insurance. Figure 11 shows that this version of the model provides a good account of the empirical life-cycle profile of both the extensive margin and the intensive margin. This is confirmed by figure A10 in the Appendix, which shows the life-cycle profile of the participation rate for married households with children. From figure A10 we conclude that the model provides a good match of the data on participation rates in the life insurance market.

Figure 12 depicts the life-cycle profile of life insurance for married households with children who have purchased some life insurance and face the mean level of mortality risk and figure 13 shows the corresponding life cycle profile for households facing the highest level of mortality risk (upper 10-percentile). Figures 12 and 13 show that the model matches well the empirical life-cycle profile of these two moments, namely mean and upper 10-percentile. Note that our calibration approach allows us to match by construction the mean and upper 10-percentile of life insurance holdings averaged over the life-cycle, but leaves no degree of freedom for matching the entire life-cycle profile. Note also that figures 11, 12 and 13 show that the shape of the life cycle profile of life insurance holdings is very similar for all three moments (median, mean, upper 10-percentile), both in the data as well as according to the model.

Figures A11 and A12 in the Appendix show the life-cycle profile of the life insurance coefficient for the median, mean, and upper 10-percentile of the distribution – the corresponding plot for the median is shown in figure 7. The figures show that for all three moments there exists a strong correlation between age and insurance, and that the correlation is roughly the same for all three moments. In sum, this paper's main quantitative result regarding the under-insurance of young households holds for all three moments both in the data and in the model.

5.2 Wealth and Consumption

An essential feature of our mechanism generating under-insurance of young households is that young household have little financial wealth relative to their human wealth. In our model, the portfolio mix between human and financial capital is measured by θ_h, the fraction of total wealth invested in human capital. Empirically, we construct a measure of portfolio holdings by taking the ratio of (net) financial wealth to labor earnings, and compare this to the model generated analog which is given by $\frac{1-\theta_h}{\phi r_h \theta_h}$. Figure 14 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 for young households, and matches the observed increase in financial wealth relative to human wealth through age 50, although it over-predicts wealth holdings for the oldest households. We view this as a success of the model, as 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.

Another important dimension of the data is the consumption dispersion over the life-cycle. Figure 15 compares the variance of log adult-equivalent consumption in the US, estimated using data from the Consumer Expenditure Survey, from three different studies—Aguiar and Hurst (2008), Deaton and Paxson (1994), and Primciceri and van Rens (2009)—to 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).

Note that the consumption implications of our model differ substantially from the results obtained by the previous literature on limited enforcement (Krueger and Perri, 2006, and Boer, 2014) due to the life-cycle component and human capital investment. For example, in Section 5.1 we have shown that young households insure less than 40 percent of their human capital losses upon death of a spouse (see figure 7), which implies that these households experience a substantial (and permanent) drop in consumption levels in the event of the death of an adult family member relative to their expected consumption path. Further, in our calibrated model economy the participation constraints bind frequently for our youngest households as they have an intense desire to borrow to invest in human capital and to smooth consumption over their lifecycle. This is why our measure of overall consumption insurance in Figure 10 shows that less than 25% of the volatility in the consumption of the young is insured in our benchmark model.

Though our model provides a better account of the consumption data than previous limited enforcement models analyzed in the literature, our baseline model still fails to match some dimensions of the consumption data. Specifically, the distribution of equilibrium consumption in the baseline model is rightward skewed since consumption growth is bounded from below by $\beta(1 + r_f)$. This property of the standard limited commitment model is well-known. For example, Broer (2014) argues that it is at odds with the CEX data. However, the distribution of equilibrium consumption growth becomes richer once we introduce state-dependent marginal utility of consumption, $\gamma_1 = \gamma_1(s_1)$, endogenous leisure choice with non-separable preferences, or unobserved heterogeneity in β. A deeper study of this interesting issue is beyond the scope of the current paper and constitutes an important topic for future research.

5.3 Welfare Cost of Mortality Risk

In this section, we turn to an analysis of the welfare effect of mortality risk and the welfare cost of lack of insurance against this type risk. To this end, we compute the welfare cost of mortality risk as the difference between welfare of a 23-year old married household with children who is fully insured against mortality risk and welfare of the same household in a situation with no access to insurance against mortality risk, where the no-insurance situation is defined by the absence of both private life insurance markets and public insurance through the social security system. To facilitate the interpretation of our results, we report all welfare changes in terms of the dollar value of the equivalent variation in the present value of consumption.[21] The welfare cost is defined in an ex-ante sense before households know

[21] Welfare for households who are fully insured is computed based on the equilibrium for a model with no η-shocks for all ages and family types. Welfare without insurance against mortality risk is computed using η-shocks that do not take into account social security survivor benefits and restrict asset payoffs (portfolios)

their mortality risk parameter η and their number of children, that is, it is the social welfare cost for the group of 23-year old married households with children. Table 2 summarizes the results of our welfare calculations and shows that the welfare cost of mortality risk is $\$10,350$, which is a substantial amount. This number hides a great deal of heterogeneity. The welfare cost for a 23-year old married household with median level of mortality risk is $\$2,385$, whereas this welfare cost is $\$7,501$ for the group of married households with average (mean) levels of mortality risk, and this cost increases to $\$91,653$ for the 10 percent of 23-year old married households who face the highest levels of mortality risk.[22] The welfare cost of mortality risk has two parts: the cost of moving from full insurance to imperfect insurance and the cost of moving from imperfect insurance to no insurance. This decomposition of the welfare cost of mortality risk is shown in columns 2 and 3. A comparison of columns 2 and 3 shows that both components of the welfare costs are substantial, but the larger part of this welfare cost is the welfare difference between no insurance and imperfect insurance.

The above numbers provide the cost of lacking both private and public insurance against mortality risk. The welfare cost of lacking only the private insurance market for a 23-year old married household is $\$2,773$, see column 4 of table 2. Further, this welfare cost is \$370 for a 23-year old married household with median level of mortality risk, $\$1,009$ for the group of married households with average (mean) levels of mortality risk, and $\$14,721$ for the 10 percent of 23-year old married households who face the highest levels of mortality risk. Columns 5 and 6 show the decomposition of these welfare cost into the two parts corresponding to moving from full insurance to imperfect insurance and form imperfect insurance to no insurance. As in the previous case, the largest part of this welfare cost is the difference between no insurance and imperfect insurance.

5.4 Reform of the Consumer Bankruptcy Code

to be constant across death states and no-death states. In all cases we are careful to make the appropriate mean adjustments in η so that we only consider mean-preserving spreads. We compute welfare cost as the product of the equivalent percentage change in consumption times the present value of consumption of a 23-year old married households with children and median wealth.

[22]Here median, mean, and upper 10-percentile refer to the moments of the distribution of mortality risk for households with $\alpha > 0$. Thus, these welfare cost are the costs for those households who have purchased some life insurance and whose life insurance demand is depicted in figures 6, 15, and 16. In our computation of the social welfare cost we average over all values of α including $\alpha = 0$ and therefore take into account households who have not purchased life insurance.

There has been a long-standing debate among academic scholars and policy makers as to the relative merits of alternative consumer bankruptcy codes. For example, 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, an effect that has also been studied in the theoretical contributions of Attanasio and Rios-Rull (2000) and Krueger and Perri (2010). In this section, we provide a quantitative evaluation of a change in the consumer bankruptcy code on household borrowing, insurance, and welfare.

Our experiment is motivated by the Bankruptcy Abuse Prevention and Consumer Protection Act (BAPCPA) of 2005, which made it more difficult to file for bankruptcy under chapter 7—where debt is only repaid out of existing assets—and therefore forced more households to file under chapter 13 of the US bankruptcy code, where debts are repaid out of current earnings over a period of 3 to 5 years. See, for example, White (2007) for a detailed account of the US bankruptcy code before and after the reform. We implement this experiment by assuming that after implementation of the BAPCPA, in the event of bankruptcy, 30% of households are randomly assigned to file under chapter 13.[23] In line with the code, we model the consequence of filing for bankruptcy under chapter 13 as an exclusion from borrowing and insurance markets for an average of 4 years and a 25 percent garnishment of labor income during the period of exclusion. The BAPCPA also increased bankruptcy filing costs significantly, and we incorporate this change in legislation by introducing a one-off cost that is paid in the year of filing for bankruptcy by all defaulting households. We assume this cost to be proportional to the wealth of households and set the cost parameter so that it amounts

[23]An important change the BAPCPA introduced was the "means test". This means test restricted filing under chapter 7 to those households with income below median income adjusted for family type, which suggest that after the reform 50 percent of all households are forced to file bankruptcy under chapter 13. However, defaulting households differ from non-defaulting households, and we take account of this fact by assuming that only 30 percent defaulting households are forced into chapter 13 after the reform. The number of 30 percent corresponds to the fraction of defaulting households in our SCF 2004 sample who have above median income.

to \$2,000 for the median wealth household. Finally, the BAPCPA increased the minimum number of years that have to pass until a consumer can file a second time under chapter 7 from 7 years to 8 years, and we incorporate this change in legislation by assuming that households filing for bankruptcy under chapter 7 are excluded for 8 years after the reform (instead of 7 years before the reform).

We compute the welfare consequences of the reform by comparing the lifetime utility of new-born households (households age 23) in the two economies, before and after the reform. For this comparison, we compute the welfare change $\Delta(s_1)$ as the equivalent variation of the bankruptcy reform, measured in units of lifetime consumption, and then average over family states s_1 using the fixed stationary distribution over s_1 (this distribution over exogenous shocks is not affected by the policy experiment). Note that the welfare change $\Delta(s_1)$ is independent of the initial wealth level of a household so that in our case there is no need to average over wealth using an endogenous wealth distribution. Note further that we conduct a steady state comparison in the sense that we do not take into account the transition path of the aggregate capital-to-labor ratio $\tilde{K}$ (and the corresponding transition path of investment returns).

We compute three different measures of the welfare effect of the reform of the bankruptcy code for new-born households. First, we keep human capital investment fixed and ask how much young households gain when additional borrowing enables them to buy more insurance. This welfare gain from better insurance is 0.5 percent of lifetime consumption. A welfare gain of half percent of lifetime consumption through improvements in risk sharing is substantially larger than any gain that the model of Krueger and Perri (2006) would predict, where households are almost fully insured even before the reform. Second, we allow households to adjust human capital choices but keep investment returns fixed (partial equilibrium with endogenous human capital accumulation). The welfare gain for households age 23 is 0.56 percent of lifetime consumption. Finally, we consider endogenous human capital accumulation in general equilibrium. In this case, the rate of return on human capital investment goes down and the return to physical capital investment goes up, which introduces a negative welfare effect for young households who are almost fully invested in human capital. The net welfare gain for newborn households is 0.25 percent when we take into account the general equilibrium adjustments of investment returns.

In figure A13 in the Appendix we plot the life-cycle profile of the ratio of net worth over labor income before and after the reform and show in figure A14 the insurance measure $1 - \sigma_c/\sigma_{a,c}$ before and after the reform. Figure A10 confirms that households borrow more after the reform – the youngest households have positive net worth before the reform and negative net worth after the reform. Figure A14 supports the idea that the reform improves insurance – the reform increases the insurance coefficient $1 - \sigma_c/\sigma_{a,c}$. Note that the reform improves insurance against all types of risk, including labor market risk. Hence, the welfare gain of 0.43% is attributed to a gain in risk sharing in all insurance markets.

In general equilibrium the reform of the bankruptcy code increases $\tilde{K}$ since households invest more in human capital after the reform. We find that the reform of the bankruptcy code increases $\tilde{K}$ by 1.2 percent of its initial value of 0.4. In our endogenous growth model, any change in $\tilde{K}$ also changes the aggregate growth rate of the economy. Specifically, the equilibrium value $\tilde{K}$ is in general lower than the value of $\tilde{K}$ that maximizes aggregate growth, and an increase in $\tilde{K}$ increases aggregate growth, an effect that is discussed in more detail in Krebs (2003). In our calibrated model economy, the growth gain is relatively modest: the annual growth rate increases by 0.02 percentage point.

6. Extensions and Robustness

In this section, we discuss several extensions of the baseline model that help us understand additional dimensions of the data. We also perform a battery of robustness checks with respect to our data analysis and changes in the parameter values of the calibrated model economy. Due to space limitation many of the details of the analysis and all figures are relegated to the Appendix.

6.1 Model Extensions

The baseline model used in Section 5 assumes that the marginal utility of consumption is independent of family structure and in particular independent of the death event. Motivated by two recent contributions (Koijen et al., 2012, and Hong and Rios-Rull, 2012), we have considered an extension of the baseline model that incorporates household preferences that depend on the number of children and the health status of the household. Details of the model extension and the analysis can be found in Section G.1 of the Appendix. The results can be summarized as follows. First, with the changes added to the model, the basic facts about life insurance and other asset holdings over the life-cycle for all married households with children are unchanged – see in particular figure A11 in the Appendix. Second, this extension improves the match between model and data in the sense that the extended model replicates additional cross-sectional facts. Third, if we interpret the change in marginal utility following the death of a parent as reflecting the consequent change in the cost of living, the resulting changes are relatively modest and increase in the number of kids.

In the baseline model, prior to retirement all agents can buy a complete set of insurance products, including both life insurance and annuities. However, we constrain retirees to save in a risk-free asset with any wealth remaining at their death distributed to newborn households. In Section G.2 of the Appendix we also discuss an extension of the baseline model in which retirees have a bequest motive and can purchase annuities. We conclude that the restriction on retirees in the baseline model is relatively innocent.

6.2 Sensitivity Analysis

We have conducted an extensive sensitivity analysis varying the main parameters of interest within a range of empirically plausible values. In particular, we considered realistic changes in the parameters controlling mortality risk and contract enforcement. Overall, the main quantitative results of this paper have shown to be quite robust to these variations in parameter values (targets). Details of our sensitivity analysis can be found in the Section H of the Appendix.

6.3 Empirical Robustness

We also conducted an extensive robustness analysis of our empirical results. The details of this analysis can be found in Section I of the Appendix. The results of this analysis can be summarized as follows.

We analyzed to what extent our approach to estimating human capital losses upon the death of a spouse leads us to overestimate or perhaps underestimate the true losses. In particular, we used panel data drawn from the SIPP to investigate if there is evidence of selection bias when we estimate earnings losses by comparing the earnings of married households with the earnings of single households. Further, we considered if there are additional channels of insurance through inter vivo transfers using data drawn from the SIPP and the

PSID. Our conclusion is that there is no evidence that our estimates of earnings losses have a substantial bias.

We also investigated — one paragraph about life insurance holding by wealth in the data ...

7. Conclusion

In this paper, we provided empirical evidence of under-insurance in the market for life insurance. We then developed a tractable macroeconomic model with risky human capital and used the model to provide an explanation of our empirical finding in terms of endogenous borrowing constraints due to limited enforcement. We also used the framework to analyze the possible consequences of under-insurance. The results of this paper suggest at least two lines of future research.

First, in the paper we restricted attention to insurance against one form of human capital risk, the death of a family member, and insurance against this type of risk that can be purchased in the market for life insurance. A promising avenue for future research is to investigate the extent to which limited contract enforcement helps us understand empirical patterns in other insurance markets. For example, one basic implication of the adverse selection and moral hazard approach to insurance is that households with higher risk exposure should buy more insurance (Chiappori and Salanie, 2000), and a number of empirical studies have found that this hypothesis is rejected in the data (Chiappori and Salanie, 2000, and Bernheim, Forni, Gokhale, and Kotlikoff, 2003). In contrast, the limited enforcement approach suggests a negative correlation between risk exposure and insurance if households with high risk exposure are also the households with the largest share of human capital in total wealth. Thus, theories of limited contract enforcement have the potential to explain the empirical findings of Chiappori and Salanie (2000) and Bernheim, Forni, Gokhale, and Kotlikoff (2003).[24]

A second line of research would broaden the set of assets available to households. The most important alternative asset is housing, which is also risky and which is, to varying

[24]Note that adverse selection in conjunction with unobserved preference heterogeneity provides an alternative explanation of the observed inverse relationship between risk exposure and insurance (Cutler, Finkelstein, and McGarry, 2008).

degrees, 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*

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

Angeletos, M. (2007) "Uninsured Idiosyncratic Investment Risk and Aggregate Saving," *Review of Economic Dynamics* 10: 1-30.

Attanasio, O., and V. Rios-Rull (2000) "Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare?", *European Economic Review*: 44: 1225-1258.

Azariadis, C., and L. Kaas (2008) "Credit and Growth under Limited Commitment," *Macroeconomic Dynamics* 12: 20-30.

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.

Becker, G. (1964) *Human Capital: A Theoretical and Empirical Analysis*, University of Chicago Press.

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., Dearden, L., Meghir, C., and B. Sianesi (1999) "Human Capital Investment: The Returns from Education and Training to the Individual, the Firm, and the Economy," *Fiscal Studies* 20: 1-23.

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

Broer, T. (2012) "The Wrong Shape of Insurance? What Cross-sectional Distributions Tell Us About Models of Consumption Smoothing," *Working Paper*

Browning, M., Hanse, L., Heckman, J (1999) "Micro Data and General Equilibrium Models," in: Taylor, J., and Woodford, M (eds) *Handbook of Macroeconomics* vol 15: 543-663, Amsterdam, New York and Oxford.

Brugiavini, A. (1993) "Uncertainty Resolution and the Timing of Annuity Purchases," *Journal of Public Economics* 50: 31-62

Burkhauser, R., Giles, P., Lillard, D., and J. Schwarze (2005) "Until Death Do US Part: An Analysis of the Economic Well-Being of Widows in Four Countries," *Journal of Geren-*

tology 60: 238-246.

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

Cawley, J., and T. Philipson (1999) "An Empirical Examination of Information Barriers to Trade in Insurance," *The American Economic Review* 89: 827-46

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.

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

Cullen, JB., J. Gruber (2000) "Does Unemployment Insurance Crowd out Spousal Labor Supply?" *Journal of Labor Economics* 18: 546-72

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., and 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.

Friedman, B., and M. Warshawsky (1990) "The Cost of Annuities: Implications for Saving Behavior and Bequests," *The Quarterly Journal of Economics* 105: 135-54.

Gallipoli, G., and L. Turner L (2009) "Disability in Canada: A Longitudinal Household Analysis," Working Paper, University of British Columbia.

Gallipoli, G., and L. Turner (2011) "Household Responses to Individual Shocks: Disability and Labor Supply," Working Paper, University of British Columbia.

Gustman, A., Mitchell, O., Samwick, A., and T. Steinmeier (1997) "Pension and Social Security Wealth in the Health and Retirement Study," NBER Working Paper 5912.

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.

Heckman, J., and T. Macurdy (1980) "A Life Cycle Model of Female Labour Supply," *The Review of Economic Studies* 47: 47-74

Hendel, I., and A. Lizzeri (2003) "The Role of Commitment in Dynamic Contracts:

Evidence From Life Insurance," *Quarterly Journal of Economcis* 118: 299-327.

Hong, J., and V. Rios-Rull (2012) "Life Insurance and Household Consumption," *American Economic review* 102: 3701-3730.

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

Inkmann, J., and A. Michaelides (2012) "Can the Life Insurance Market Provide Evidence for a Bequest Motive?", *The Journal of Risk and Insurance* 79: 671-695.

Johnson, R., Burman, L., and D. Kobes (2004) "Annuitized Wealth at Older Ages: Evidence from the Health and Retirement Study," Final Report to the Employee Benefits Security Administration U.S. Department of Labor

Jorgenson, D., and B. Fraumeni (1993) "The Accumulation of Human and Non-Human Capital, 1948-84," In *The Measurement of Saving, Investment, and Wealth*. Eds. R. Lipsey and H. Tice *NBER Studies in Income and Wealth* 52: 227-82.

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

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.

Koijen, R., Nieuwerburgh, S., and M. Yogo (2012) "Health and Mortality Delta: Assessing the Welfare Cost of Household Insurance Choice," *NBER Working Paper*.

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.

Krueger, D., and F. Perri (2010) "Public vs Private Risk Sharing" *Journal of Economic Theory* 146: 920-956.

Krueger, A., and M. Lindhal (2001) "Education for Growth: Why and for Whom?", *Journal of Economic Literature* 39: 1101-1136.

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.

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

Mincer, J. (1994) "Investment in US Education and Training," *NBER Working Paper*.

Moll, B. (2012) "Productivity Losses From Financial Frictions: Can Self-Financing Undo Capital Misallocation?", Working Paper.

Norton, A. (1990) and L. Miller (1990) "Remarriage Among Women in te United States: 1985," in *Studies in Household and Family Formation*.

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

Ruggles, P. 1990. *Drawing the Line: Alternative Poverty Measures and their Implications for Public Policy*. Washington, D.C.: The Urban Institute.

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

Society of Actuaries (2012) "Report of the Society of Actuaries Preferred Underwriting Structures Survey Subcommittee: A Review of Current (December 2010) Preferred Underwriting Criteria," Society of Actuaries.

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

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

Taber, C. (2001) "The Rising College Premium in the Eighties: Return to College or Return to Unobserved Ability,"*Review of Economic Studies* 68: 665-691.

Weaver, D. (2010) "Widows and Social Security,"*Social Security Bulletin* 70: 89-109.

White, M. (2007) "Bankruptcy Reform and Credit Cards *Journal of Economic Perspectives* 21: 175-199.

Wilson, B., and S. Clarke (1992) "Remarriages: A Demographic Profile," *Journal of Family Issues* 13: 123-141.

Winter, R. (1981) "On the Rate Structure of the American Life Insurance Market," *The Journal of Finance* 36: 81-96.

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

Table 2. Welfare Cost of Mortality Risk

	Δ	Δ_1	Δ_2	Δ^p	Δ^p_1	Δ^p_2
social	$\$10,350$	$\$6,246$	$\$4,104$	$\$2,773$	$\$1,760$	$\$1,013$
median	$\$2,385$	$\$1,696$	$\$689$	$\$370$	$\$320$	$\$50$
mean	$\$7,501$	$\$4,948$	$\$2,553$	$\$1,009$	$\$754$	$\$255$
10-percentile	$\$91,653$	$\$55,039$	$\$36,614$	$\$14,721$	$\$8,999$	$\$5,722$

Note: All welfare changes are expressed as the dollar value of the equivalent variation in the present value of consumption for a 23-year old married houshold with children. Δ is the welfare cost of no life insurance computed as the welfare difference between full insurance against mortality risk and no insurance against mortaility risk. Δ_1 is the welfare cost of moving from full insurance to partial insurance against mortality risk. Δ_2 is the welfare cost of moving from partial insurance to no insurance against mortality risk. Δ^p is the welfare cost of no private life insurance computed as the welfare difference between full insurance against mortaility risk and no private insurance against mortality risk. Δ^p_1 is the welfare cost of moving from full insurance to partial private insurance against mortality risk. Δ^p_2 is the welfare cost of moving from partial private insurance to no private insurance against mortality risk. "Social" refers to the social welfare cost, that is, the welfare cost before the level of mortality risk is known. "Median", "mean", and "10-percentile" refer to the welfare cost for households with median, mean, and highest level of mortality risk, respectively.

Figure 1: Life Insurance Coefficient

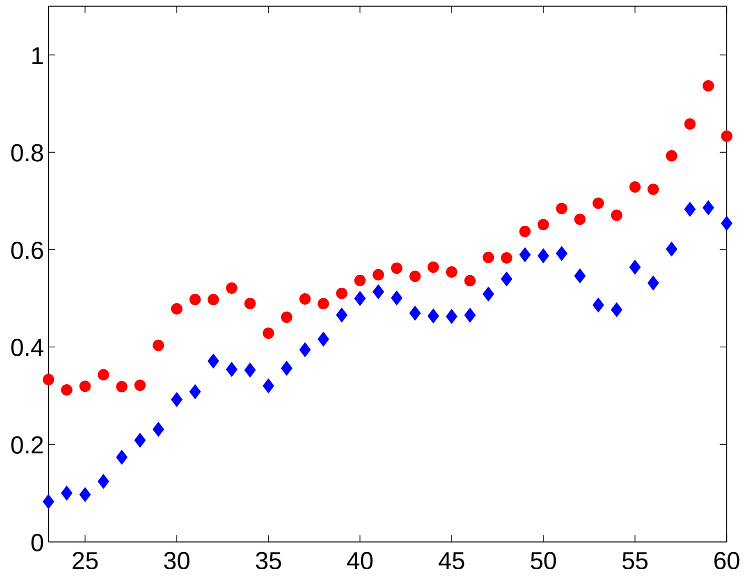

Notes: Life-cycle profile of life insurance coefficient. Life insurance coefficient is the ratio of life-insurance holdings to present value income loss in case of death. Red dots are for married households with children that have purchased life insurance. Blue diamonds are for all married households with children. All data are for households age 23 - 60 from the SCF, surveys 1992 - 2007 and show medians of the respective groups. See appendix for calculation of present value loss.

Figure 2: Labor income

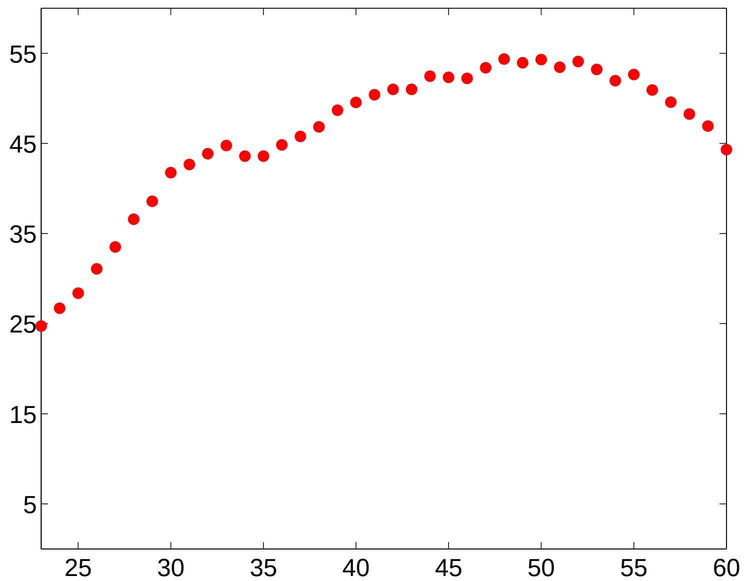

Notes: Life-cycle profile of median labor income for married households age 23 - 60 with children from the SCF, surveys 1992 - 2007 (thousands of year 2000 dollars).

Figure 3: Life Insurance

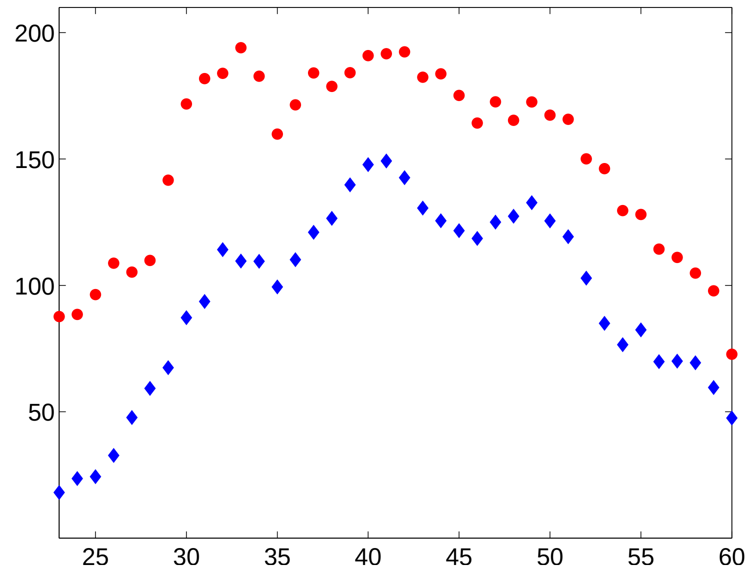

Notes: Life-cycle profile of face value of all life insurance contracts (thousands of year 2000 dollars). Red dots are for married households with children that have purchased life insurance. Blue diamonds are for all married households with children. All data are for households age 23 - 60 from the SCF, surveys 1992 - 2007 and show medians of the respective groups.

Figure 4: Human capital loss

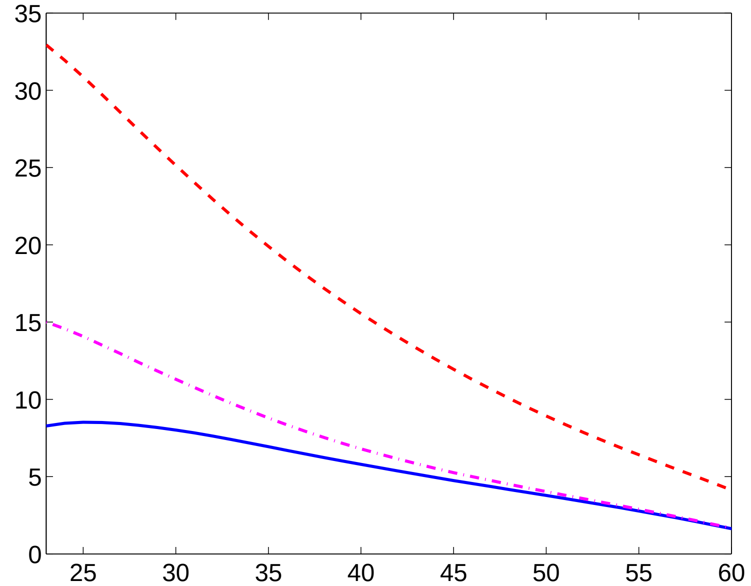

Notes: Life-cycle profile of sum of expected human capital loss in case of husbands and wifes death for all married households with children. Human capital loss is ratio of present value income loss over current labor income. Red dashed line: loss before transfers and taxes with zero probability to remarry. Pink dashed-dotted line: loss after transfers and taxes and zero probability to remarry. Blue solid line: loss after transfers and taxes and empirical remarriage rates. All data are for households age 23 - 60 from the SCF, surveys 1992 - 2007. See appendix for further details.

Figure 5: Networth to labor income ratio

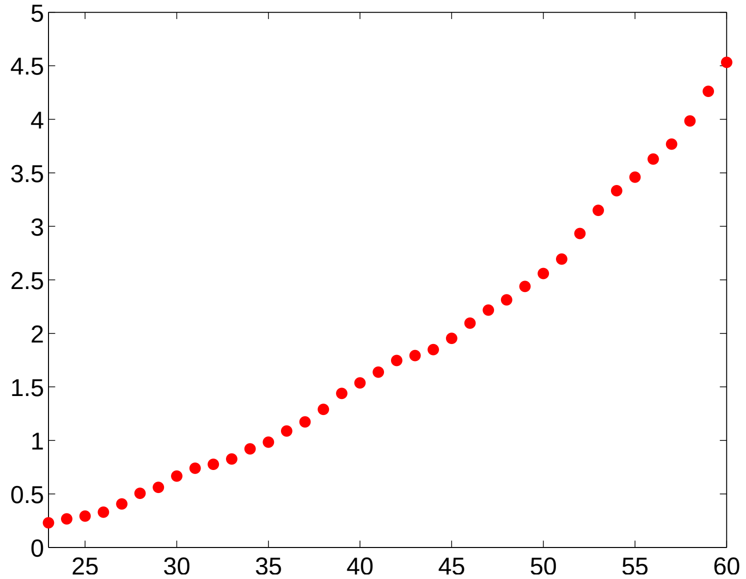

Notes: Life-cycle profile of the median ratio of networth to labor income for married households age 23 - 60 with children from the SCF, surveys 1992 - 2007.

Figure 6: Life insurance

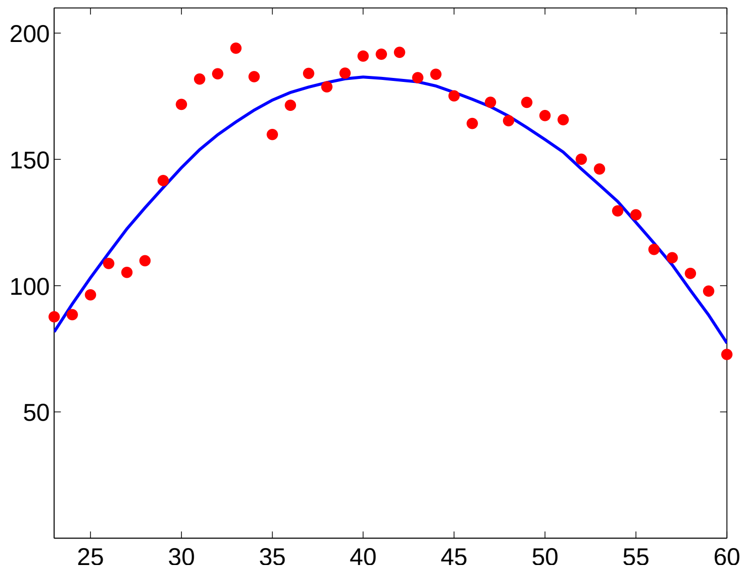

Notes: Life-cycle profile of face value of life insurance contracts (thousands of year 2000 dollars) for married households age 23 - 60 with children. Blue solid line shows model. Red dots show median of all households that have purchased life insurance from the SCF data.

Figure 7: Life Insurance Coefficient I

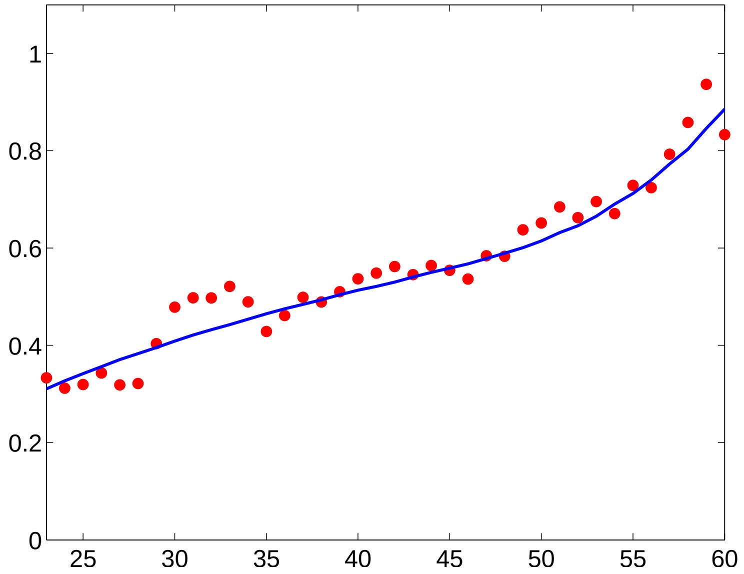

Notes: Life-cycle profile of life insurance coefficient. Life insurance coefficient is the ratio of life-insurance holdings to present value income loss for married households age 23 - 60 with children. Blue solid line shows model. Red dots show data for all households that have purchased life insurance.

Figure 8: Life Insurance Coefficient II

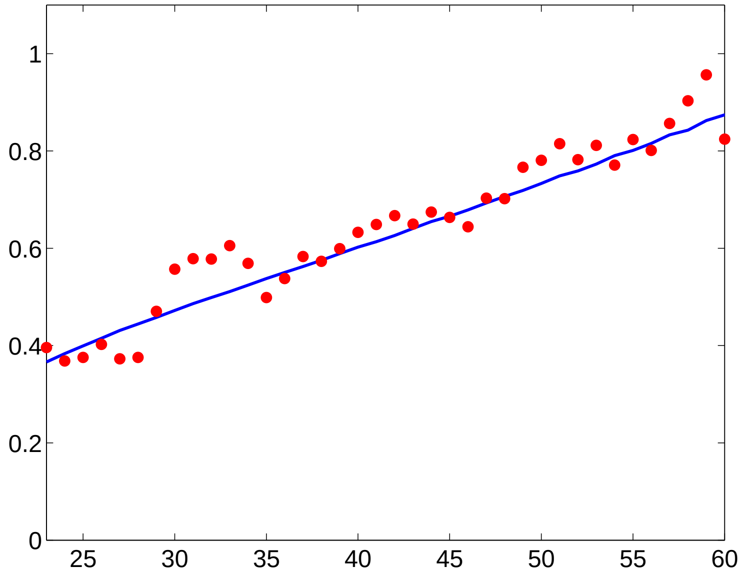

Notes: Life-cycle profile of life insurance coefficient. Life insurance coefficient is the ratio of life insurance holdings over full insurance for married households age 23 - 60 with children. Blue solid line shows model. Red dots show data for all households that have purchased life insurance.

Figure 9: Life insurance for husband and wife

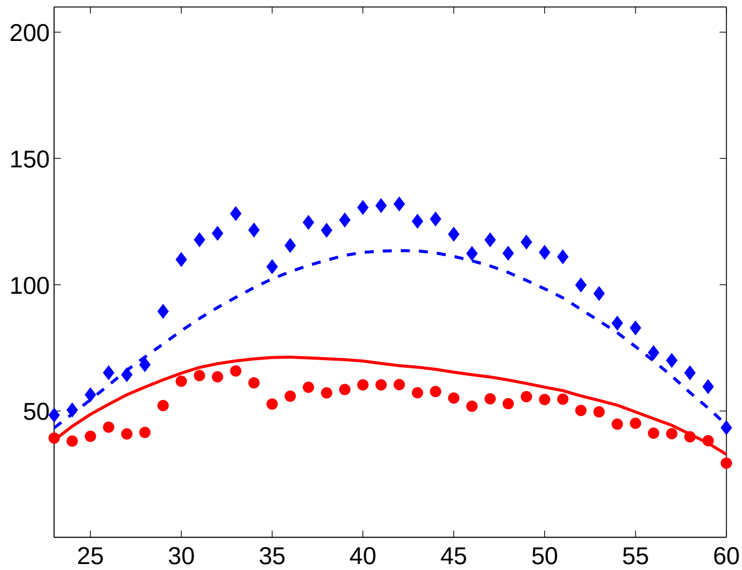

Notes: Life-cycle profile of face value of life-insurance contracts (thousands of year 2000 dollars) for married households age 23 - 60 with children. Red solid line shows face value of life insurance for wifes death from model. Red dots show face value of life insurance for wifes death from data. Blue dashed line shows face value of life insurance for husbands death from model. Blue diamonds show face value of life insurance for husbands death from data. Data are from the SCF and the SIPP. See appendix for details of the construction of data profiles.

Figure 10: Consumption insurance

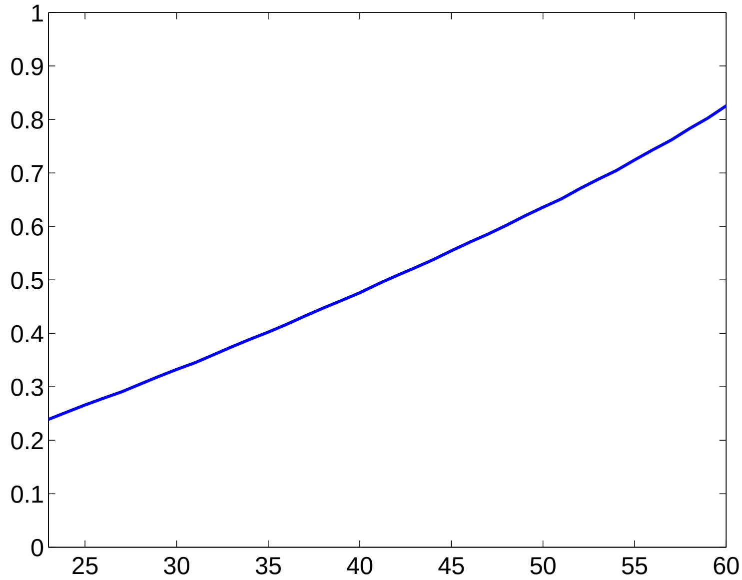

Notes: Consumption insurance in the model for married households age 23 - 60 with children. The insurance measure is one minus the ratio of the standard deviation of consumption in equilibrium relative to the standard deviation of consumption in financial autarky.

Figure 11: Life Insurance: Extensive and Intensive Margin

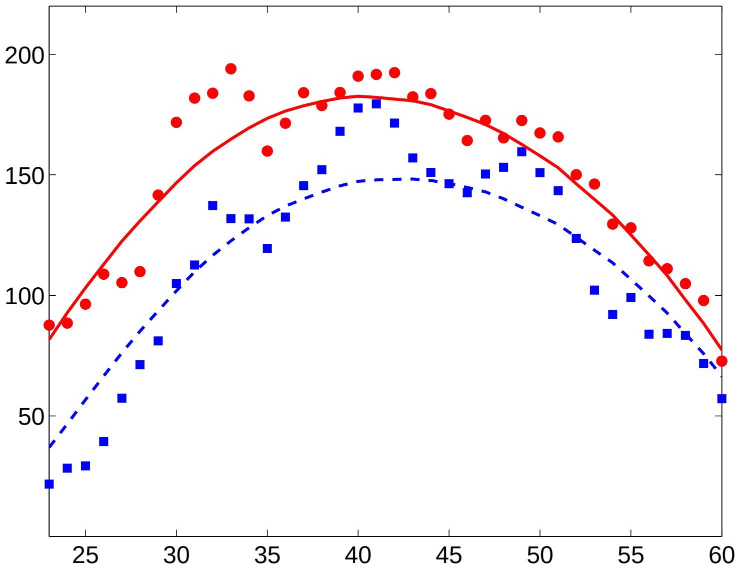

Notes: Life-cycle profile of face value of life insurance contracts (thousands of year 2000 dollars) for married households age 23 - 60 with children. Red solid line shows model prediction for all households that have purchased life-insurance. Red dots show all households that have purchased life insurance from the data. Blue dashed line shows model prediction for all households. Blue squares show all households from the data. All data are for households age 23 - 60 from the SCF, surveys 1992 - 2007 and show medians of the respective groups.

Figure 12: Life Insurance (mean mortality risk)

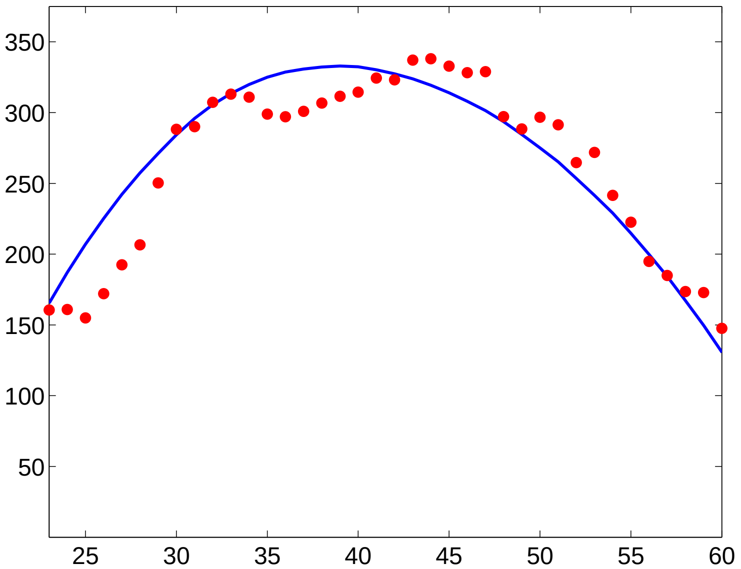

Notes: Life-cycle profile of face value of life insurance contracts (thousands of year 2000 dollars) for married households age 23 - 60 with children and mean mortality risk. Blue solid line shows model. Red dots show mean of all households that have purchased life insurance from the SCF, surveys 1992 - 2007.

Figure 13: Life Insurance (highest 10 % mortality risk)

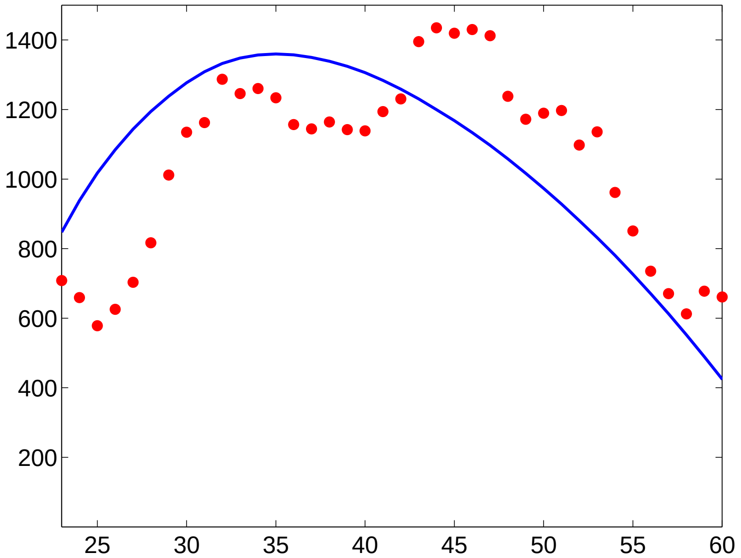

Notes: Life-cycle profile of face value of life insurance contracts (thousands of year 2000 dollars) for married households age 23 - 60 with children with the 10 % highest mortality risk. Blue solid line shows model. Red dots show mean of households with the highest 10 % life-insurance holdings that have purchased life insurance from the SCF, surveys 1992 - 2007.

Figure 14: Net Worth to Labor Income

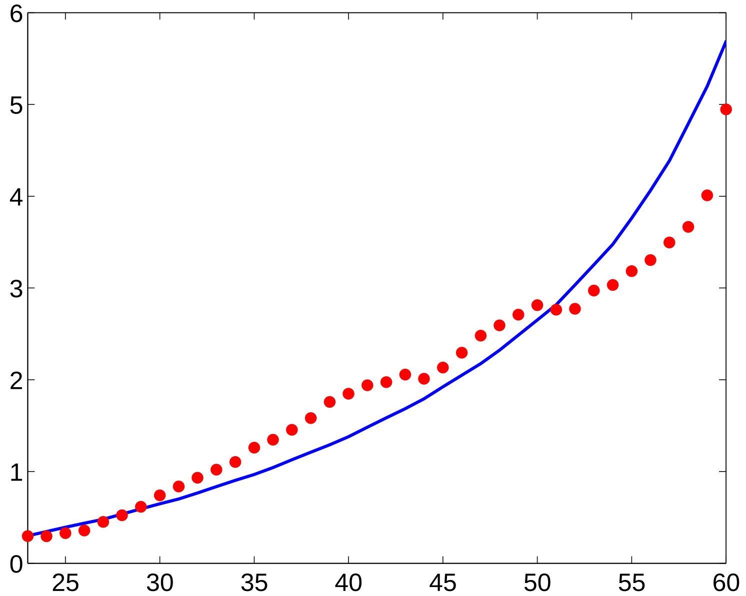

Notes: Life-cycle profile of the median ratio of networth to labor income for married households age 23 - 60 with children. Blue solid line shows model and red dots SCF data.

Figure 15: Consumption Inequality

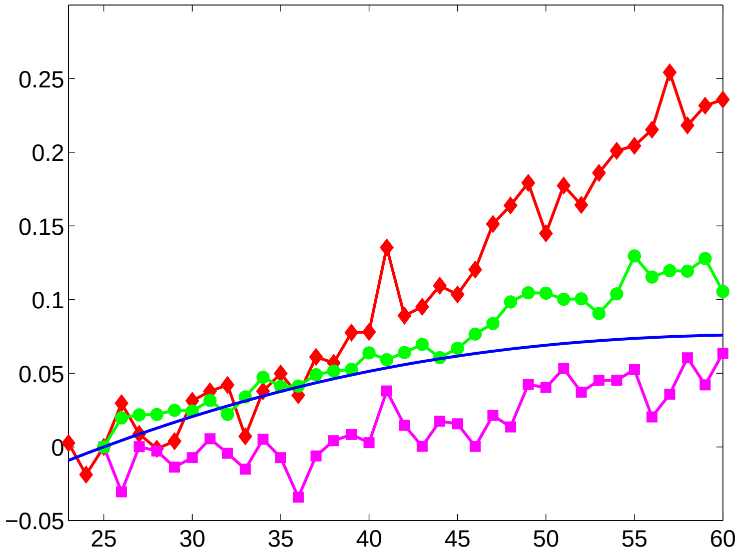

Notes: Life-cycle profile of the cross-sectional variance of consumption. The blue solid line shows the model prediction. The red diamonds show the profile estimated by Deaton and Paxson (1994), the green dots are the estimates of Aguiar and Hurst (2008), and the pink squares are the estimates of Primiceri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

Table 1: Calibration

parameter	value	description
β	0.95	discount factor
ϕ	0.6749	(inverse of) price of human capital
p_{ret}	0.2	probability of retiring
p_{death}	0.05	probability of dying
σ_η	0.0913	standard deviation of permanent shocks
σ_z	0.3	standard deviation of transitory shocks
p	0.8571	probability of remaining in financial autarky
α	0.32	capital share in output
δ_k	0.0785	physical capital depreciation rate
δ_h	0.0429	human capital depreciation rate
A	0.1818	total factor productivity
λ	1.6461	human capital endowment of young households