Asset Pricing with Idiosyncratic Risk and Overlapping Generations[*]

Kjetil Storesletten[†], Christopher I. Telmer[‡] and Amir Yaron[§]

This Draft: January 2007
First Draft: June 1996

Abstract

What is the effect of non-tradeable idiosyncratic risk on asset-market risk premiums? Constantinides and Duffie (1996) and Mankiw (1986) have shown that risk premiums will increase if the idiosyncratic shocks become more volatile during economic contractions. We add two important ingredients to this relationship: (i) the life cycle, and (ii) capital accumulation. We show that in a realistically-calibrated life-cycle economy with production these ingredients mitigate the ability of idiosyncratic risk to account for the observed Sharpe ratio on U.S. equity. While the Constantinides-Duffie model can account for the U.S. value of 41% with a risk-aversion coefficient of 8, our model generates a Sharpe ratio of 33%, which is roughly half-way to the complete-markets value of 25%. Almost all of this reduction is due to capital accumulation. Life-cycle effects are important in our model — we demonstrate that idiosyncratic risk matters for asset pricing because it inhibits the intergenerational sharing of aggregate risk — but their net effect on the Sharpe ratio is small.

1   Introduction

The essence of Mehra and Prescott's (1985) equity premium puzzle is that investing in equity looks like too good of a deal; the stock market seems to reward risk-taking far more than than a representative agent would require. A large literature has asked if the representative-agent assumption lies at the heart of the puzzle. The idea is that individuals face idiosyncratic risks and are unable to insure against them, and that this affects the way they value financial assets. The plausibility of this story seems apparent. Non-financial wealth — human wealth in particular — is larger than financial wealth and is subject to substantial risks. Upon closer inspection, however, the story runs into difficulties. Idiosyncratic risks are, by definition, uncorrelated with aggregate risks. In contrast, asset pricing relies on dependence between sources of risk and asset returns in order to explain why some assets pay a higher expected return than others. The challenge for a theory of asset pricing driven by idiosyncratic risk, therefore, is to generate such dependence while still having the idiosyncratic shocks wash-out at the aggregate level.

An innovative response to this challenge is Mankiw (1986). In his model aggregate shocks and the *volatility* of idiosyncratic shocks are negatively related. He showed that these kinds of idiosyncratic shocks represent a source of aggregate risk in that they matter for the equity premium. We refer to this kind of aggregate risk as *countercyclical cross-sectional variance*, or *CCV risk*. Constantinides and Duffie (1996) formalized the pricing of CCV risk in a multiperiod setting and went on to derive a more general set of conditions under which it can resolve any aggregate-consumption-based asset-pricing puzzle.

Our paper adds two potentially-important ingredients which are absent in the Constantinides-Duffie model: (i) capital accumulation, and (ii) the life cycle. Why might these ingredients be important? First, regarding capital accumulation, a number of papers have shown that the degree of risk-sharing is increasing in the level of aggregate capital.[1] The asset-pricing effects of idiosyncratic risk, therefore, are likely to be exaggerated by the Constantinides-Duffie model.[2] Second, regarding the life cycle, the essence of their story is that non-tradeable idiosyncratic shocks to human wealth — the capitalized value of labor income — affect the valuation of financial wealth. The distribution of human wealth necessarily has a life-cycle dimension: the

---

[1]See, for example, Krueger and Perri (2006), Krusell and Smith (1997) and Storesletten, Telmer, and Yaron (2004a).

[2]To be specific, the Constantinides-Duffie model is an environment without physical capital accumulation *and* where aggregate financial capital must equal zero. The former simply means that the Constantinides-Duffie model is a Lucas tree economy. The latter is necessary for the construction of an autarkic equilibrium. It means that non-traded endowment income represents 100% of aggregate consumption, and that financial assets are zero-net supply claims on stochastic processes which are not permitted to be measurable with respect to the individual-specific information structure. Thus, there is no *aggregate capital*, in either a physical or a financial sense, which can serve as a buffer stock against adverse fluctuations in the value of human capital.

young have more than the old. The same is true, therefore, of idiosyncratic shocks. Surely this must matter for asset pricing? Consider, for example, retired people. They comprise roughly 20 percent of the adult population, they participate in equity markets at a much higher rate, yet they face little if any labor-market risk. If the solution to the equity-premium puzzle is that labor-market risk makes equity more risky, then why don't retirees hold all of the equity, thus resurrecting the puzzle?

We address these issues as follows. We begin with the life cycle. We construct two OLG versions of the Constantinides-Duffie model, one without retirement and one with retirement. We calibrate the models and confirm the above intuition: the existence of retirees reduces the Sharpe ratio on equity from 41% — which matches U.S. data — to 34%, which is roughly halfway to the complete-markets value of 25%. The expressions we derive provide a clear economic intuition for this, one which survives in richer environments. Retirees do not face non-tradeable CCV risk in human wealth. They, therefore, have a comparative advantage in bearing the aggregate risk inherent in financial wealth. The Constantinides-Duffie model is autarkic, so this cannot show up in portfolios. Where it shows up is consumption allocations. The consumption of retirees is more exposed to aggregate risk than that of workers. We interpret this as the intergenerational sharing of aggregate risk. This risk sharing is imperfect in the sense that the associated complete-markets allocation features uniform aggregate risk exposure across generations. Yet there *is* aggregate risk sharing in the sense that the young are endowed with more aggregate risk but, in equilibrium, bear less of it. This is a central intuition of our paper. Idiosyncratic risk matters for asset pricing because it inhibits the intergenerational sharing of aggregate risk. The more it does this, the larger will be the Sharpe ratio.

The remainder of our paper focuses on OLG economies with capital. The inclusion of capital makes the effect of retirement on the Sharpe ratio more complex. While the intergenerational risk sharing effect remains — which tends to reduce the Sharpe ratio — two additional effects arise which tend to offset each other. First, the *level* of capital enhances 'self-insurance:' the behavior of accumulating and decumulating buffer-stock savings in the face of good and bad idiosyncratic shocks. Self-insurance behavior mitigates the extent to which income shocks are manifest in consumption and tends to reduce the Sharpe ratio. Second, the *distribution* of capital can work in the opposite direction. The more that the distribution is skewed toward the old the more the young are exposed to CCV risk. This tends to increase the Sharpe ratio. In our calibration the distributional effect turns out to be quite strong. The Sharpe ratio in an economy without retirement is 31%, substantially less than the Constantinides-Duffie, no-retirement counterpart of 41%. The difference is due to the self-insurance characteristic of an economy with capital. When we introduce retirement, the Sharpe ratio actually increases slightly, to 33%. We provide quantitative evidence that this is driven by the distributional effect being slightly larger than the intergenerational risk-sharing effect.

Our model features non-degenerate trade in financial assets. Thus, we can say something about what kinds of portfolio rules support the imperfect aggregate risk sharing allocations described above. A useful benchmark is Bodie, Merton, and Samuelson (1992) (BMS). In their model wages are riskless and labor income is like a non-defaultable bond. This induces agents to reduce the fraction of financial wealth held in stocks as they age. In our model the same force is at work, but with risky wages. Wage variability has both an idiosyncratic and an aggregate component. The latter is less variable than stock returns. Therefore, just as in BMS, wage income can serve as a "hedge" for financial income, resulting in reduced stock holding with age. However, this only happens beyond a certain age. The youngest workers hold relatively little stock, resulting in hump-shaped portfolio rules.[3] This is driven by the fact that (i) wages are less variable than stock returns, (ii) wages are perfectly correlated with stock returns, and (iii) wages exhibit CCV.[4] Aggregate risk, therefore, is concentrated on a subset of the population — mostly retirees — thereby tending to increase the Sharpe ratio. Constantinides, Donaldson, and Mehra (2002) get big risk premiums via a similar outcome. Our mechanism, however, is different than theirs. In our model the young don't hold equity because they choose not to. In their model the young don't hold equity because they are not allowed to.

The remainder of our paper is organized as follows. Section 2 discusses related literature. Section 3 formulates a life-cycle version of the Constantinides and Duffie (1996) model and shows that retirement, to some extent, mitigates the model's ability to account for the equity premium puzzle. Section 4 formulates a model with capital and non-degenerate trade in financial assets and examines its quantitative properties. Section 5 provides an in depth analysis of some important economic properties of the model and Section 6 concludes.

2   Related Work

A number of papers examine the quantitative implications of the Constantinides and Duffie (1996) model for asset pricing. To understand how our paper fits in, it's important to understand the nature of Constantinides-Duffie's main result. They show that any given collection of asset price processes are consistent with a heterogeneous agent economy in which agents have 'standard' preferences and face idiosyncratic shocks with a particular volatility process. Their model's testable restrictions can be thought of in two ways. First, because the economy admits the construction of a representative agent, it restricts the joint behavior of aggregate consumption, as-

---

[3]Hump-shaped portfolio rules are (arguably) consistent with average portfolio behavior in the U.S. (*e.g.*, Amerkis and Zeldes (2000), Heaton and Lucas (2000)).

[4]Feature (ii) is inconsistent with high frequency movements in wages and stock returns. But at lower frequencies it is both natural (*i.e.*, it is a feature of most RBC models, like ours) as well as empirically valid (see, for instance, Benzoni, Collin-Dufresne, and Goldstein (2004)).

set returns and the cross-sectional variation in consumption. That is, conditional on knowledge of the cross-sectional variance, the model's first-order conditions can be tested without individual-level data. Papers by Balduzzi and Yao (2000), Brav, Constantinides, and Geczy (2002), Cogley (2002) Ramchand (1999) and Sarkissian (2003) investigate these restrictions and find mixed evidence. Second, if one asks what gives rise to the first-order conditions, the model restricts the joint behavior of individual labor income, asset returns, and individual consumption. Most critical is the requirement that labor income be a unit-root process with innovations which become more volatile during aggregate downturns. Our paper, and its companion paper Storesletten, Telmer, and Yaron (2004b), focus on these restrictions. The advantages to doing so are both related to data — income is certainly easier to measure than consumption — and the ability to understand how idiosyncratic risk interacts with asset pricing at a structural level.

Krusell and Smith (1997) laid much of the groundwork for our paper in studying the asset-pricing effects of idiosyncratic labor-market risk in models with capital. Our results on self-insurance with aggregate capital are basically life-cycle versions of results in their paper. They also demonstrate the limitations of the Constantinides-Duffie framework as it relates to the distribution of financial wealth, a result which is quite similar to our findings in a life-cycle context. An important distinction, however, is borrowing constraints. Krusell-Smith require extreme borrowing constraints (essentially zero) in order to generate risk premiums, whereas, due primarily to life cycle considerations and unit root shocks, borrowing constraints play no role in our study.

Finally, our paper builds on a large body of work on asset pricing with heterogeneous agent models, including Aiyagari (1994), Aiyagari and Gertler (1991), Alvarez and Jermann (2001), den Haan (1994), Gomes and Michaelides (2004), Guvenen (2005), Heaton and Lucas (1996), Huggett (1993), Lucas (1994), Mankiw (1986), Marcet and Singleton (1999), Ríos-Rull (1994), Telmer (1993), Weil (1992), and Zhang (1997). The stationary OLG framework we develop owes much to previous work by Ríos-Rull (1994), Huggett (1996) and Storesletten (2000).

3 An OLG Version of the Constantinides-Duffie Model

We begin with a life-cycle version of the Constantinides and Duffie (1996) model. There are two asset markets, a one-period riskless bond and an equity claim to a dividend process, D_t. The bond and equity prices are denoted q_t and p_t, respectively. Equilibrium will be autarkic, so limiting attention to two assets is without loss of generality.

The economy is populated by H overlapping generations of agents, indexed by $h = 1, 2, \ldots, H$, with a continuum of agents in each generation. Agents are born with

one unit of equity and zero units of bonds. Preferences are

$$U(c) = E_t \sum_{h=1}^{H} \beta^h (c_{it+h}^h)^{1-\gamma}/(1-\gamma) \ , \tag{1}$$

where c_{it}^h is the consumption of the i^{th} agent of age h at time t and β and γ denote the discount factor and risk aversion coefficients, respectively.

Each agent receives nontradeable endowment income of y_{it}^h,

$$y_{it}^h = G_t \exp(z_{it}^h) - D_t \ , \quad h = 1, 2, \ldots, (H-1) \tag{2}$$
$$y_{it}^H = G_t \exp(z_{it}^H) - (p_t + D_t) \ , \tag{3}$$

where G_t is an aggregate shock (defined more explicitly below) and the idiosyncratic shocks, z_{it}^h, follow a unit root process with heteroskedastic innovations,

$$z_{it}^h = z_{i,t-1}^{h-1} + \eta_{it} \tag{4}$$
$$z_{i,t}^0 = 0 \tag{5}$$
$$\eta_{it} \sim \mathrm{N}(-\sigma_t^2/2, \, \sigma_t^2) \tag{6}$$
$$\sigma_t^2 = a + b \log(G_t/G_{t-1}) \ . \tag{7}$$

This structure is essentially identical to the Constantinides-Duffie formulation, the only exception being that in the last period of life the amount $p_t + D_t$ is subtracted from income, instead of just D_t. In either case the implication is that the amount of aggregate financial wealth is zero. This property is critical for the construction of an autarkic equilibrium. In Section 4 we relax this condition and allow for trade. The incorporation of positive financial wealth will turn out to be a driving force in our results.

The equilibrium of this model is autarky with individual consumption $c_{it}^h = G_t \exp(z_{it}^h)$. Bond and equity prices satisfy

$$q_t = \beta^* E_t \lambda_{t+1}^{-\gamma^*} \tag{8}$$
$$p_t = \beta^* E_t \lambda_{t+1}^{-\gamma^*} (p_{t+1} + D_{t+1}) \ , \tag{9}$$

where $\lambda_{t+1} = G_{t+1}/G_t$, $\beta^* = \beta \exp(\gamma(1+\gamma)a/2)$ and $\gamma^* = \gamma - b\gamma(1+\gamma)/2$ (see Constantinides and Duffie (1996) for derivations). A cross-sectional law of large numbers implies that the variable G_t, and therefore the growth rate λ_t, coincides with per-capita consumption, which we denote C_t (the reason for making a potential distinction will become apparent in the next section),

$$C_t = \frac{1}{H} \tilde{E}_t \sum_{h=1}^{H} G_t \exp(z_{it}^h) = G_t \ ,$$

where $\tilde{E}_t$ is a cross-sectional expectations operator which conditions on time t aggregate information. Since $C_t = G_t$, the pricing equations (8) and (9) represent a representative agent equilibrium where the agent's preference parameters (β^*, γ^*) are amalgamations of actual preference parameters (β, γ) and technological parameters (a, b). The main idea behind the Constantinides-Duffie model is that (i) because $\beta^* > \beta$, the model may resolve the 'risk-free rate puzzle,' and (ii) if $b < 0$ (*i.e.*, the volatility of idiosyncratic shocks is countercyclical) then 'effective' risk aversion exceeds actual risk aversion $(\gamma^* > \gamma)$, and the model may resolve the equity premium puzzle.

3.1   Calibration

We now ask if the values of a and b implied by labor market data satisfy the above requirements and help the model account for the equity premium. We use estimates from Storesletten, Telmer, and Yaron (2004b) which are based on annual PSID data, 1969-1992. They show that (a) idiosyncratic shocks are highly persistent and that a unit root is plausible, (b) the conditional standard deviation of idiosyncratic shocks is large, averaging 17%, and (c) the conditional standard deviation is countercyclical, increasing by roughly 68% from expansion to contraction (from 12.5% to 21.1%). In Appendix A we show that these estimates map into values $a = 0.0143$ and $b = -0.1652$.

We use a stochastic process for λ_t which is essentially the same as that of Mehra and Prescott's (1985): a two-state Markov chain with mean, standard deviation and autocorrelation of aggregate consumption growth of 0.018, 0.033, and -0.14, respectively (we use a slightly lower value for the standard deviation which matches our dataset). We choose the 'effective' discount factor, β^*, to match the average U.S. riskfree interest rate, and the effective risk aversion coefficient, γ^*, to match either the U.S. Sharpe ratio or the unlevered U.S. equity premium. Table 1 reports the implications for the 'actual' risk aversion coefficient, γ. To match the Sharpe ratio, a value of $\gamma^* = 13.6$ is required. This corresponds to an actual risk aversion coefficient of $\gamma = 7.8$. To match the equity premium $\gamma^* = 15.42$ is required, which corresponds to $\gamma = 8.6$.[5] To facilitate a direct comparison with the numerical results (shown below) for an economy with trade, we also report the results for the case in which actual risk aversion is 8. Time preference is characterized by $\beta^* = 1.140$ $(\beta = 0.69)$, $\beta^* = 1.148$ $(\beta = 0.64)$, and $\beta^* = 1.142$ $(\beta = 0.68)$ respectively.

The Constantinides-Duffie model, then, is successful at what it sets out to do; given a realistic parameterization for idiosyncratic risk, it accounts for the equity

---

[5]Cogley (2002) formulates an asset-pricing model with idiosyncratic risk (to individuals' consumption) and uses the empirical time-varying cross-sectional moments of consumption growth from the Survey of Consumer Expenditures (CEX) to ask what level of risk aversion would be required to account for the empirical equity premium. Interestingly, this approach delivers a risk aversion of 8 (assuming a plausible level of measurement error).

premium without resorting to extreme values for risk aversion and/or negative time preference. Along other dimensions, of course, the model is counterfactual. It generates excessive volatility in both risky and riskless asset returns and cannot account for the ubiquitous rejections of Euler equation tests based on (8) and (9) (*i.e.*, such tests typically reject for all values of β^* and γ^*). Constantinides and Duffie (1996) prove that this can be rectified with an alternative process for the conditional variance σ_t^2 from equation (6). The remainder of our paper, however, focuses on a more fundamental set of the model's restrictions, those which involve age and risk sharing.

3.2   The Implications of Retirement

We now introduce retirees and ask to what extent they mitigate the model's success. There are two senses in which the process (2)–(7) does not capture retirement. First, agents face idiosyncratic income shocks in all periods of life. Second, agents receive income each period until death, thus obviating the need to save for retirement. We begin by incorporating the first feature, which can be analyzed in the no-trade environment. The second requires trade and is incorporated in Section 4.

We define a retired agent as one who does not receive an idiosyncratic shock beyond some retirement age so that, for retirees, $a = b = 0$. Given this, equations (8) and (9) no longer describe autarkic equilibrium prices. Marginal rates of substitution (at autarky) are

$$\text{workers:} \qquad \beta E_t (\frac{G_{t+1}}{G_t})^{-\gamma} e^{\gamma(1+\gamma)a/2} (\frac{G_{t+1}}{G_t})^{\gamma b(1+\gamma)/2} \ , \qquad (10)$$

$$\text{retirees:} \qquad \beta E_t (\frac{G_{t+1}}{G_t})^{-\gamma} \ . \qquad (11)$$

Retirees differ from workers in two ways. First, with $a > 0$ the exponential term in equation (10) is positive, implying that retirees discount future consumption more than workers. Intuitively, the absence of idiosyncratic risk reduces their demand for precautionary savings and they assign a lower price to a riskfree bond. Second, if $b < 0$, retirees appear less risk averse than workers, assigning a relatively high value to risky assets or, equivalently, demanding a relatively small risk premium. By removing the countercyclical volatility from the retiree's endowments we have effectively given them a greater capacity to bear aggregate risk.

We can now do one of two things to characterize an equilibrium. We can allow trade and solve for market clearing prices to replace equations (8) and (9). This would involve a substitution of consumption from retirement toward the working years, and an increased exposure to aggregate risk for retired individuals. Alternatively, we can follow Constantinides and Duffie (1996) and characterize endowments which give rise to a no-trade equilibrium, but subject to the constraint that retirees do not receive

idiosyncratic shocks. The difference between these endowments and those in equations (2) and (3) will be suggestive of what will characterize an equilibrium with trade.

A three-generation example, $H = 3$, will make the point. Generations 1 and 2 receive endowments according to equations (2)–(7). Generation 3 — the old agents — receive

$$y_{it}^3 \;\;=\;\; f_t \, G_t \exp(z_{it}^3) - (p_t + D_t) \;\;, \tag{12}$$

but with $z_{it}^3 = z_{it}^2$ (*i.e.*, the innovation in equation (4) equals zero), and

$$f_t = e^{-a(1+\gamma)/2} \left(\frac{G_t}{G_{t-1}} \right)^{-b(1+\gamma)/2} \;\;.$$

Given the endowment (12), the prices (8) and (9) once again support an autarkic equilibrium. Relative to the original endowment, the old now receive less goods (on average) with more aggregate risk, just as the above intuition suggests. What has changed, however, is aggregate consumption. Assigning a population weight of 20 percent to the old generation (corresponding to the U.S. population), aggregate consumption is

$$
\begin{aligned}
C_t \;\;&=\;\; \tilde{E}_t \left(0.8[G_t \exp(z_{it}^1) + G_t \exp(z_{it}^2)] + 0.2 f_t G_t \exp(z_{it}^3) \right) \\
&=\;\; G_t \left(0.8 + 0.2 \, e^{-a(1+\gamma)/2} \left(\frac{G_t}{G_{t-1}} \right)^{-b(1+\gamma)/2} \right) \;\;, \tag{13}
\end{aligned}
$$

which, because we've added aggregate risk to the endowment of the old, can be substantially more variable than G_t.

The prices (8) and (9) are now valid, but only in an economy with more variability in aggregate consumption growth than the original. The above calibration (which underlies Table 1) is therefore invalid. Aggregate consumption growth, as implied by equation (13), now has a standard deviation of 4.2 percent, 27% larger (0.9 percentage points) than the benchmark volatility of consumption growth. In this sense, adding retirees implies that, without changing preferences, the model can only account for asset prices with an unrealistically high amount of aggregate variability.

An alternative is to re-calibrate the process G_t/G_{t-1} so that aggregate consumption growth, C_t/C_{t-1} from equation (13), has mean, standard deviation and autocorrelation which match the U.S. data. Results are given in the 6th to 8th rows of Table 1. Holding risk aversion fixed (row 6), we find that the required reduction in the variability of aggregate consumption growth causes the model's Sharpe ratio to fall from 41.2 percent to 34.4 percent. The equity premium falls from percent to 3.4 percent to 2.3 percent. For the alternative calibration (row 8), the Sharpe ratio and equity premium fall from 45.9 percent to 38.6 percent and 4.1 percent to 2.9 percent, respectively.

To summarize, retirement has the effect one might expect. Because retirees do not face countercyclically heteroskedastic shocks — the driving force in the Constantinides-Duffie model — they are less averse to bearing aggregate risk. An autarkic allocation must therefore skew the aggregate risk toward the old, who are content to hold it in return for a relatively low expected return. In this sense, the incorporation of retirement resurrects the equity premium puzzle.

4   Models With Trade

The previous section emphasized the importance of how idiosyncratic shocks are distributed over the life cycle. Equally important is the distribution of what is being shocked: the human wealth represented by the flow of income, y_{it}^h. Human wealth typically accounts for a large fraction of total wealth for young people and a small fraction for older people. Given the nature of our question — How do shocks to human wealth affect the valuation of financial wealth? — incorporating this seems of first-order importance. It may also overturn the implication of the previous section, which was driven by older agents bearing the lion's share of the aggregate risk. If a realistic human/financial wealth distribution reverses this, making the younger agents who face the idiosyncratic risk instrumental in pricing the aggregate risk, the incorporation of retirement may actually help the model to account for the equity premium.

The major cost of incorporating a life-cycle wealth distribution is that, necessarily, we must allow for trade (*i.e.*, if nontradeable income is zero after retirement, the young must save and the old must dissave). With several exceptions — Gertler (1999) for example — this means using computational methods to analyze the model. The benefits, however, are numerous. First, we can make the model more realistic along certain dimensions which are important for calibration (*e.g.*, the demographic structure). Second, a more realistic life-cycle distribution of human wealth will necessarily imply a more realistic life-cycle distribution of financial wealth (recall that in the Constantinides-Duffie model financial wealth equals zero). This is important because financial wealth is the means with which agents accomplish buffer-stock savings and self-insurance. Finally, the model will display partial risk-sharing behavior, even with unit root idiosyncratic shocks. Partial risk sharing is an undeniable aspect of U.S. data on income and consumption.

With all this in mind, we make the following changes to the framework of Section 3.

Financial markets.

With trade, the menu of assets is no longer innocuous. We now *limit* asset trade to a riskless and a risky asset. The latter takes the form of ownership of an aggregate

9

production technology. Agents rent capital and labor to constant-return-to-scale firm which then splits its output between the two. Labor is supplied inelastically and, in aggregate, is fixed at N. Denoting aggregate consumption, output and capital as Y_t, C_t and K_t respectively, the production technology is

$$
\begin{align}
Y_t &= Z_t K_t^\theta N^{1-\theta} \tag{14}\\
K_{t+1} &= Y_t - C_t + (1-\delta_t)K_t \tag{15}\\
r_t &= \theta Z_t K_t^{\theta-1} N^{1-\theta} - \delta_t \tag{16}\\
w_t &= (1-\theta) Z_t K_t^\theta N^{-\theta}, \tag{17}
\end{align}
$$

where r_t is the return on capital (the risky asset), w_t is the wage rate, θ is capital's share of output, Z_t is an aggregate shock and δ_t is the depreciation rate on capital. The depreciation rate is stochastic:

$$
\delta_t = \delta + (1-Z_t)\frac{s}{Std(Z_t)} \quad , \tag{18}
$$

where δ controls the average and s is, approximately, the standard deviation of r_t.[6]

This production process delivers four key ingredients: (i) the model is tractable (solving the analogous endowment economy is substantially more difficult), (ii) the volatility of the return on equity can be calibrated realistically, (iii) the volatility of aggregate consumption growth can be calibrated realistically, (iv) the return on human capital — essentially the wage rate — can be substantially less volatile than the return on equity. Each ingredient is critical for our question. The first two are obvious.[7] The third ensures that the aggregate part of the asset-pricing Euler equations is realistic (*i.e.*, see equations (8) and (9)), which is essential if we are to isolate the incremental impact of idiosyncratic risk. The fourth is instrumental for life-cycle portfolio choice. It determines which age cohorts hold equity in equilibrium and, consequently, whether or not idiosyncratic risk is priced.

Endowments.

The endowment processes (2)–(7) are of a special form required to support an autarkic outcome. Since this is no longer required, and because of the incorporation

---

[6]Greenwood, Hercowitz, and Huffman (1988) and Greenwood, Hercowitz, and Krusell (1997) have used a similar production technology in a business cycle context. Boldrin, Christiano, and Fisher (2001) have done so in an asset pricing context. Our technology is essentially a reduced-form representation of, for instance, Greenwood, Hercowitz, and Krusell (1997), equation (B3).

[7]The tractability afforded by the production economy is due to the fact that the return on the risky asset can be written as a simple function of the market-clearing level of aggregate capital. In an endowment economy, in contrast, an separate stochastic process for the risky asset return must be characterized. den Haan (1997) was, to our knowledge, the first to do the latter in a model with two assets and a large number of agents.

of production, we reformulate them as follows. First, to capture the fact that young people have relatively little financial wealth relative to human wealth, we endow all newborn agents with zero units of equity and zero units of bonds. Next, the non-tradeable endowment now takes the form of labor efficiency units, not units of the consumption good.[8] At time t the ith working agent of age h is endowed with n_{it}^h units of labor which they supply inelastically. Retirees are agents for whom h exceeds a retirement age $\bar{H}$. They receive $n_{it}^h = 0$. For workers,

$$\log n_{it}^h = \kappa_h + z_{i,t}^h \quad , \tag{19}$$

where κ_h is used to characterize the cross-sectional distribution of mean income across ages, and

$$z_{it}^h = \rho z_{i,t-1}^{h-1} + \eta_{it} \quad , \quad \eta_{it} \sim \mathrm{N}(0, \sigma_t^2) \quad ,$$

with $z_{it}^0 = 0$. We use a two-state specification for σ_t^2:

$$\begin{aligned} \sigma_t^2 &= \sigma_E^2 \ \ \text{if} \ \ Z \geq E(Z) \\ &= \sigma_C^2 \ \ \text{if} \ \ Z < E(Z) \ . \end{aligned}$$

Individual labor income now becomes the product of labor supplied and the wage rate: $y_{it}^h = w_t n_{it}^h$.[9]

With $\rho = 1$ this process is analogous to the Constantinides-Duffie process, (2)–(7). The exceptions are that (i) income is now a share of the aggregate wage bill instead of aggregate consumption, (ii) financial income is no longer 'taxed' at 100 percent as in (2)–(7), thereby implying that aggregate financial wealth is zero, and (iii) the variance of the innovations to z_{it}^h is now a discrete function of the technological shock Z, not a continuous function of aggregate consumption growth.

4.1   Equilibrium

The state of the economy is a pair, (Z, μ), where μ is a measure defined over an appropriate family of subsets of $S = (\mathcal{H} \times \mathcal{Z} \times \mathcal{A})$, $\mathcal{H}$ is the set of ages, $\mathcal{H} = \{1, 2, \ldots, H\}$, $\mathcal{Z}$ is the product space of all possible idiosyncratic shocks, and $\mathcal{A}$ is the set of all possible beginning-of-period wealth realizations. In words, μ is simply a distribution

---

[8]Strictly speaking, this is inconsistent with the empirical approach of Storesletten, Telmer, and Yaron (2004b) which measured idiosyncratic risk using labor *income*, not hours worked. To reconcile the two, we have generated simulated data on labor income from our model and estimated a labor income process identical to that in Storesletten, Telmer, and Yaron (2004b). Owing in large part to relatively low variability in the wage rate, w_t, the results were very similar. In this sense, the population moments for labor income in our model have been calibrated to sample moments on non-financial income from the PSID.

[9]Our model assumes that bequests are zero. This provides focus on our main point: the effect of intergenerational dispersion in the ratio of human to total wealth.

of agents across ages, idiosyncratic shocks and ex-post wealth. The existence of aggregate shocks implies that μ evolve stochastically over time (*i.e.*, μ belongs to some family of distributions over which there is defined yet another probability measure). We use G to denote the law of motion of μ,

$$\mu' = G(\mu, Z, Z') \ .$$

The bond price and the return on equity can now be written as time-invariant functions $q(\mu, Z)$ and $r(\mu, Z)$. The wage rate is $w(\mu, Z)$. Omitting the (now redundant) time t and individual i notation, the budget constraint for an agent of age h is,

$$
\begin{aligned}
c_h + k'_{h+1} + b'_{h+1} q(\mu, Z) &\leq a_h + n_h w(\mu, Z) &\qquad (20) \\
a_h &= k_h r(\mu, Z) + b_h \\
k'_{H+1} &\geq 0 \\
b'_{H+1} &\geq 0
\end{aligned}
$$

where a_h denotes beginning-of-period wealth, k_h and b_h are beginning-of-period capital and bond holdings, and k'_{h+1} and b'_{h+1} are end-of-period holdings. We do not impose any portfolio restrictions over and above restricting terminal wealth to be non-negative (the third and fourth restriction).

Denoting the value function of an agent of age h as V_h, the choice problem can be represented as,

$$V_h(\mu, Z, z_h, a_h) = \max_{k'_{h+1}, b'_{h+1}} \left\{ u(c_h) + \right.$$

$$\left. \beta E \, V'_{h+1}\Big(G(\mu, Z, Z'), Z', z'_{h+1}, k'_{h+1} r(G(\mu, Z, Z'), Z') + b'_{h+1} \Big) \right\} \ , \qquad (21)$$

subject to equations (20). An equilibrium is defined as stationary price functions, $q(\mu, Z)$, $r(\mu, Z)$ and $w(\mu, Z)$, a set of cohort-specific value functions and decision rules, $\{V_h, k'_{h+1}, b'_{h+1}\}_{h=1}^H$, and a law of motion for μ, $\mu' = G(\mu, Z, Z')$, such that r and w satisfy equations (16) and (17), the bond market clears,

$$\int_S b'(\mu, Z, z_h, a_h) \, d\mu = 0 \ ,$$

aggregate quantities result from individual decisions,

$$
\begin{aligned}
K(\mu, Z) &= \int_S k_h(\mu, Z, z_h, a_h) \, d\mu \\
N &= \int_S n_h \, d\mu \ ,
\end{aligned}
$$

agents' optimization problems are satisfied given the law of motion for (μ, Z) (so that $\{V_h, k'_{h+1}, b'_{h+1}\}_{h=1}^H$ satisfy problem (21)), and the law of motion, G, is consistent with individual behavior. We characterize this equilibrium and solve the model using the computational methods developed by Krusell and Smith (1997) and described further in Appendix B.

4.2 Quantitative Properties

Our model has three motives for trade which are absent in the Constantinides-Duffie framework. First, retirees don't face any idiosyncratic shocks but workers do. Second, retirees don't receive any income. Therefore, they must save for retirement by accumulating financial assets while working and then sell these assets to younger agents as they age. Third, if $\rho < 1$ working-age agents will self-insure against mean-reverting idiosyncratic shocks by trading in financial markets. In what follows, we eliminate the latter motive for trade and set $\rho = 1$. The reasons are that we'd like to emphasize the first two motives — the life-cycle motives — and we'd like to maintain some comparability with the Constantinides-Duffie model. Computational tractability is also more manageable with unit-root shocks.[10]

We calibrate our economy as follows. The most important issues are listed here, with additional details relegated to Appendix A.

1. Idiosyncratic risk, captured by equation (19), follows a unit-root process with a regime-switching conditional variance function chosen to match the estimates in Storesletten, Telmer, and Yaron (2004b). Their estimate of ρ is 0.952. We scale down the variances in our model so that, with $\rho = 1$, the unconditional variance over the life-cycle matches that implied by their $\rho = 0.952$ estimates. This results in $\sigma_E = 0.0768$ and $\sigma_C = 0.1298$.

2. The discount rate β is chosen to ensure the capital-to-output ratio is set to 3.3.

3. The magnitude of the depreciation shocks in equation (18) is set so that the standard deviation of aggregate consumption growth is 3.3 percent. We choose this (as opposed to matching the variability of equity returns) because, just as in representative agent models, realistic properties for aggregate consumption are the primary disciplinary force on asset-pricing models with heterogeneity. Equations (8) and (9) make this clear. The resulting implications for the standard deviation of equity returns is reported in Table 2. The volatility of the theoretical equity premium is 6.8%, 3.2 percentage points less than the U.S. sample value.

4. We examine economies with and without retirees. In both cases agents are born with zero financial assets. In economies with retirees — those of age $\bar{H}$ or greater — retired agents receive zero labor income and comprise 20 percent of the population.

---

[10]A previous version of the paper did examine mean-reverting shocks and found the asset-pricing implications to be qualitatively similar for $\rho = 0.92$. This isn't terribly surprising given that our model features finite lives.

Our main results are in Table 2. The first row reports the Sharpe ratio and the mean and standard deviation of the risk-free and risky rates of return for an economy *without* retirement. This economy is analogous to the no-retirement Constantinides-Duffie economy described in Table 1. The main difference, however, is that aggregate financial capital is positive in our economy but zero in the Constantinides-Duffie economy. The impact is substantial. The Sharpe ratio falls from 42.1 to 30.9. Why? Because positive aggregate capital permits self-insurance behavior: the accumulation and decumulation of a stock of precautionary savings in the face of good and bad idiosyncratic shocks, implying that consumption responds less than one-for-one to an earnings shock. Unlike the Constantinides-Duffie model, agents in our model exhibit self-insurance behavior even with unit-root shocks. This is described further in Section 5.2. For now, what's important is that self-insurance behavior mitigates exposure to idiosyncratic shocks, thereby mitigating exposure to CCV risk and reducing the Sharpe ratio.

The second row of Table 2 shows what happens when retirement in introduced. Relative to the no-retirement economy, the Sharpe ratio *increases* slightly, from 30.9 to 32.6. This is surprising in light of Section 3.2, where retirement resulted in a substantial *decrease* in the Sharpe ratio. What's going on is, again, driven by the existence of aggregate capital. To understand this, see Figure 1 which plots the age-profile of financial wealth. The figure shows that young agents in the economy with retirement accumulate wealth more slowly than those in the economy without retirement. That is, retirement induces a *distributional effect* which shifts capital-holdings toward older agents. This tends to increase the Sharpe ratio since it reduces the ability of young agents to self-insure and increases their exposure to idiosyncratic shocks and CCV risk. We substantiate this interpretation further in Section 5.

Why do young agents save less when they know they will receive less income when they are old? At first blush, this seems a contradiction. The answer is related to how we calibrate our models. In order to make sensible comparisons, we insist that, in all models, the capital-output ratio is equal to 3.3. This means that in an economy with retirees there must be less aggregate capital. The reason is that retirement implies less labor supply. Since holding fixed the capital-output ratio is equivalent to holding fixed the capital-labor ratio, retirement must also imply less aggregate capital. We accomplish this by lowering the discount factor from 0.98 to 0.80 (see Table 1). The answer to the question, then — Why do young agents save less? — is driven by a general equilibrium effect. We want to understand the effects of retirement by comparing economies with similar amounts of capital relative to output or, equivalently, similar rates of return. *Ceteris paribus*, introducing retirement will decrease the rate of return. Therefore, we undo this by lowering the discount factor. This has the desired effect on aggregate capital and the rate of return. It also has an interesting distributional effect in that more of the ownership of aggregate capital shifts towards the old.

14

5 Explaining the economic forces at work

The above interpretations of our results might seem like story telling. This section attempts to do what all good computational economics should do: substantiate the stories by describing other aspects of the solution as well as supplementary experiments. We begin by describing consumption allocations, followed by the portfolio rules which support them.

There are three main economic effects at work:

1. *Self-insurance behavior.*

2. *Intergenerational sharing of aggregate risk.*

3. *Life-cycle distribution of aggregate capital*

The first two decrease the Sharpe ratio whereas the third increases it. In our calibration the first dominates and the second and the third are basically offsetting.

More specifically, the first effect echoes the discussion in Section 4.2. That is, a key feature of our model is the existence of aggregate capital. Unlike the autarkic Constantinides-Duffie economy of Section 3.2, aggregate capital induces imperfect risk sharing, even with unit-root shocks. This mitigates the effect of idiosyncratic risk and reduces the Sharpe ratio. A number of previous papers have reached similar conclusions — most notable for our setup is Krusell and Smith (1997). Section 5.2, below, provides further details on the imperfect-risk-sharing properties of our model and compares them to U.S. data.

To understand the second effect, consider Figure 2 which plots an age-dependent measure of 'aggregate risk bearing:'

$$Cov\left(R_{t+1}\,,\,\frac{c_{h+1,t+1}}{c_{h,t}}\right)\ \ .$$

The essential point of Section 3.2 was that retirement causes this measure to increase with age. That is, non-traded CCV risk for workers means that retired agents have a comparative advantage in aggregate risk-bearing. In equilibrium, therefore, their consumption covaries more with aggregate variables such as the equity return. A simple calculation confirms this. For $\gamma = 8$, using the calibration from Section 3.2, the above covariance is larger for retirees than for workers by a factor of $-b(1+\gamma)/2 = 1.75$.

Figure 2 shows that a similar mechanism is at work in our model with capital. The three lines in the figure correspond to an economy without retirement, with retirement, and with retirement and complete markets. Consider first the distinction between complete and incomplete markets, with retirement. With complete markets

aggregate risk-bearing is uniform across age. With incomplete markets it increases, just as in the Constantinides-Duffie model with retirement. This, then, makes precise what we mean by 'idiosyncratic risk inhibiting the intergenerational sharing of aggregate risk.' Imperfectly-pooled idiosyncratic shocks cause the sharing of aggregate shocks to depart from the first-best outcome. The resulting asset-pricing effects are the central point of our paper. Table 2 shows that the Sharpe ratio in the complete-markets economy is 25% compared to 33% in our benchmark economy.

Consider next the effect of retirement, with incomplete markets. Figure 2 shows that the allocation with retirement represents a shift of aggregate risk-bearing from the young, toward the old. This mirrors Section 3.2. There, the effect on the Sharpe ratio was unambiguously negative. Here, however, Table 2 shows that the Sharpe ratio increases slightly. Why? The answer is driven by the last of the above three economic forces, the life-cycle distribution of aggregate capital.

The third effect is described by Figure 3. We plot the age-dependent *variability* of consumption growth, a measure of self-insurance behavior which is very important for asset pricing. The graph shows that, with retirement, the young face substantially higher consumption variability than without. This is because the reduction in aggregate capital held by the young — recall Figure 1 — mitigates the extent to which the young can self-insure. This is what we referred to above as a 'distributional effect.' Retirement induces less capital-holding among the young, which increases their exposure to CCV risk, which tends to increase the Sharpe ratio. The small net increase in the Sharpe ratio in Table 2 indicates that this distributional effect dominates the effect of enhanced intergenerational risk sharing documented in Figure 2, but only slightly. The two effects are basically offsetting.

5.1 Portfolio behavior

Now that we understand the distribution of aggregate risk in consumption, we can explain the portfolio rules which support it. Figures 4 and 5 plot the portfolio levels and weights, respectively, for both complete and incomplete market economies. We also find it useful to report some results for homoskedastic economies: incomplete-market economies with idiosyncratic risk but without CCV risk.

We start with complete markets. The distinguishing feature is that consumption growth rates across agents are equated with the aggregate consumption growth rate. However, the supporting portfolio policies differ because agents of different ages have different amounts of human and financial capital. Consider first the retirees. They hold diversified portfolios of stocks and bonds. The average share of stocks for a 70-year old, for instance, is about 1/2. The reason is that stock returns are much more volatile than aggregate consumption growth (7.4% versus 3.3%). Retirees have zero labor income, so, in order to replicate aggregate consumption risk, they hold lots of their wealth in bonds.

16

Consider next the relatively old workers. They are different from retirees in that they still have labor income. They also hold more of their financial wealth in stocks. The reason is that, similar to the Bodie, Merton, and Samuelson (1992) (BMS) model, labor income has bond-like properties. In BMS this means that labor income is deterministic. Not so in our model. Labor income is risky. In fact, it is perfectly correlated with stock returns (like in most RBC-style models). However, it is also a lot less volatile because stock returns bear the brunt of the depreciation-rate shocks from equation (18). The upshot is the same mechanism as in BMS; older workers hold more stock than retirees because their labor income serves as a partial hedge against their stock portfolio. See Figures 4 and 5

Things change, however, for young workers. They have negative financial wealth and, as a result, the youngest agents in the complete-market economy actually short-sell stocks. The result is a *hump-shaped* portfolio profile over the entire life-cycle. What drives this, vis-a-vis BMS, is a combination of negative financial wealth and risky wages. Average wage rates are risky and perfectly correlated with stock returns. With negative financial wealth it is as if the agent is levered in aggregate risk.[11] Hence, by shorting stocks, young agents reduce their exposure to aggregate risk, thus implementing the complete-market allocation.

We can now describe behavior in the incomplete market, CCV economy. The key to understanding it is understanding how its distribution of aggregate risk differs from the complete-market economy, and how differences in portfolio rules implement this. Recall that in Figure 2 aggregate risk in the CCV economy is concentrated on the old. In the complete-market economy it is uniformly distributed. Consequently, going from complete markets to CCV, we should see stock holdings shifted from the young to the old. This implication is clearly borne out in Figure 4 where the stock profile of the CCV economy is, for the young, shifted to the right. All agents younger then 55, in the CCV economy, hold less stock and more bonds than their complete-market counterparts (after age 55 the relationship is reversed). Similarly, Figure 5 shows that after 55, the share of stocks is higher in the CCV economy than in the complete-market economy.

The resulting hump-shaped pattern in equity ownership in the CCV economy (see Figure 5) is broadly consistent with U.S. data and has been the focus of recent work by Amerkis and Zeldes (2000) and Heaton and Lucas (2000). Brown (1990) shows that non-tradeable labor income can generate hump-shaped portfolio rules in age, and Amerkis and Zeldes (2000) discuss a similar phenomenon.

---

[11]Consider, for example, a worker who maintains a large debt, invested in bonds. If consumption equals wages net of the deterministic interest rate payments, consumption will be *more* volatile than wages. With sufficiently large debt, the agent's consumption will be more volatile than aggregate consumption, so she would want to reduce her exposure to aggregate risk. Consequently, the agent would short stocks as an insurance against aggregate risk: shorting stocks implies that in good (bad) times, when earnings growth is large (small), stock repayment is large (small).

5.2 Risk Sharing

A counterfactual implication of the Constantinides and Duffie (1996) model is that the equilibrium features no risk sharing while the bulk of the existing evidence suggests that partial risk sharing is a better characterization of reality. This seems important for the question at hand, which essentially asks how idiosyncratic consumption risk affects the market price of risk. Surely the *magnitude* of the consumption risk which agents face — a synonym for the degree of partial risk sharing — is relevant for this question?

An advantage of the life-cycle model is that, even with unit root shocks, allocations exhibit partial risk sharing. The reason involves the way in which the life-cycle savings interacts with 'buffer-stock savings:' the savings reaction to an unexpected shock. In our model, provided that financial wealth is positive, the marginal propensity to save out of current income is increasing in the level of current income but decreasing in the level of wealth. The implication is that, in spite of being characterized by unit-root shocks, agents in our economy are able to partially smooth consumption (*i.e.*, achieve some self-insurance).

One measure of self-insurance behavior is the volatility of consumption relative to that of income. In our model, with the exception of the youngest, the cross-sectional variance in consumption is less than that of income. Averaged over age, consumption is roughly 10 percent less variable (in terms of the standard deviation). In U.S. data, this value is at least 35 percent (see Deaton and Paxson (1994), Storesletten, Telmer, and Yaron (2004a), or Heathcote, Storesletten, and Violante (2005)), so, by this metric, our model exhibits too little risk sharing. An alternative — one which is more directly related to the essence of risk sharing — is the cross-sectional volatility of consumption *growth*, reported in Figure 3. In this case, we see a larger difference between our model and the autarkic outcome. As shown, autarky implies that, for workers, the graph is flat at 0.107. Our model features a monotonically decreasing graph, starting at roughly autarky and falling to near zero. The main reason is what we've emphasized above: a decreasing ratio of human to total capital and the resulting mitigation of the impact of idiosyncratic shocks. Risk sharing behavior is yet another dimension of our model for which this ratio is the main economic force at work.

5.3 Effect of Risk Aversion

Lowering risk aversion to 3 obviously lowers the Sharpe ratio and also the contribution of CCV-risk to the Sharpe ratio (over and above the Sharpe ratio with complete markets). However, the qualitative findings discussed above and documented in the figures (which all pertain to economies with a risk aversion of 8), remain unchanged with a risk aversion of 3.

6 Conclusions

Our main question is whether idiosyncratic labor-market risk matters for the pricing of aggregate risk. We emphasize the importance of the life cycle and of capital accumulation. A priori, it seemed that both would *reduce* the ability of idiosyncratic risk to account for the equity premium puzzle. We've learned that this isn't quite right. Although the CCV effect first pointed out by Mankiw becomes diminished, it remains important. The Sharpe ratio in our model is halfway between the observed U.S. value of 41% and the complete-markets value of 25%. What turns out to be critical is a realistic calibration of the level of aggregate capital. We make this clear by showing that in an economy *without* aggregate capital — our OLG version of the Constantinides-Duffie model — life-cycle effects unambiguously reduce the Sharpe ratio. Capital changes this. Depending on the amount of it, life-cycle effects can either increase or decrease the Sharpe ratio. Our calibration results in a marginal increase. Getting the amount of aggregate capital right makes a big difference in quantitatively assessing the impact of idiosyncratic risk.

What else have we learned about the life-cycle and idiosyncratic risk? Mankiw's main idea was that there is an important interaction between idiosyncratic and aggregate *shocks*. We show that life-cycle effects generate a further (endogenous) interaction which relates to *allocations*. This is most evident in risk-sharing allocations. Idiosyncratic risk is concentrated on the young and is difficult to transfer across generations. Aggregate risk is not. But the former can interfere with the latter. The existence of idiosyncratic risk can inhibit the intergenerational sharing of aggregate risk. We demonstrate this by comparing the consumption allocations between our incomplete-markets model and its complete-markets counterpart. With complete markets, aggregate risk sharing across generations is uniform. With incomplete markets old agents bear more aggregate risk than young agents. To borrow from the title of Mankiw's paper, our model features a "concentration of aggregate risk" on a subset of the population. We show that this increases the model's Sharpe ratio.

These risk-sharing allocations are supported by portfolio allocations. Like Bodie, Merton, and Samuelson (1992), portfolio choice in our model is driven by life-cycle variation in the ratio of human wealth to financial wealth. What's different, however, is idiosyncratic labor-market risk which affects the former but not the latter. There are two main forces at work. First, as an agent ages, idiosyncratic risk becomes less important to them. This happens both because they face fewer (persistent) shocks in the future and because human wealth declines as a fraction of total wealth. The CCV effect, therefore, becomes less important with age and tolerance for equity-holding increases. Second, because equity returns are substantially more volatile than the wage rate, age also brings with it an increased exposure to aggregate shocks, because an increasing share of an agent's income derives from financial assets instead of human wealth. This effect eventually counteracts the first effect and, late in the working life, tolerance for equity-holding begins to decrease with age. Taken together, the two

19

effects imply hump-shaped portfolio rules. Young agents hold zero equity, retired agents hold diversified portfolios of equity and bonds, and middle-aged agents hold levered equity, issuing bonds to both the young and the old. In our model, hump-shaped portfolio rules represent the (imperfect) intergenerational sharing of aggregate risk.

Constantinides, Donaldson, and Mehra (2002) (CDM) also stress the importance of life-cycle effects for the equity premium. Like us, an important feature of their model is that young agents hold zero equity, thereby concentrating aggregate risk on older agents. The reasons, however, are fundamentally different than in our framework, which gives rise to stark, testable restrictions between the two. Our model is distinguished by idiosyncratic risk *within* generations. A young agent's choice to avoid equity is a portfolio allocation decision: equity is too risky, so they choose not to hold any. In the CDM framework, where heterogeneity only exists across generations, the driving force is consumption smoothing and how it interacts with borrowing constraints. Young agents receive a relatively meager endowment, cannot borrow or short sell equity, and therefore choose not to hold any assets whatsoever. The two models, therefore, offer starkly different interpretations of why one might see a young household choose not to hold equity. The testable restrictions are related to overall savings behavior and how important the precautionary motive is. In our model the average young household is a net saver during the first third of their lives. That is, the precautionary motive dominates the life cycle motive, and the decision to avoid equity is driven by *risk*, in our case an avoidance of CCV risk. The CDM framework is consistent with the same average, young household not accumulating any assets but, in contrast, viewing equity (in a shadow value sense) as an attractive investment. Which of these interpretations is more important — it seems clear to us that the world features aspects of each of them — is something we leave to future work.

References

Aiyagari, S. R., (1994), Uninsured idiosyncratic risk and aggregate saving, *Quarterly Journal of Economics* 109, 659–684.

Aiyagari, S. R. and M. Gertler, (1991), Asset returns with transactions costs and uninsured individual risk, *Journal of Monetary Economics* 27, 311–331.

Alvarez, F. and U. Jermann, (2001), Quantitative asset pricing implications of endogeneous solvency constraints, *Review of Financial Studies* 14, 1117–1152.

Amerkis, J. and S. P. Zeldes, (2000), How do household portfolio shares vary with age?, Unpublished manuscript, Columbia University.

Balduzzi, P. and T. Yao, (2000), Does heterogeneity matter for asset pricing, Unpublished manuscript, Boston College.

Benzoni, L., P. Collin-Dufresne, and R. Goldstein, (2004), Portfolio choice over the life-cycle in the presence of 'trickle down' labor income, Working paper, University of Minnesota.

Bodie, Z., R. C. Merton, and W. F. Samuelson, (1992), Labor supply flexibility and portfolio choice in a life cycle model, *Journal of Economic Dynamics and Control* 16, 427–49.

Boldrin, M., L. J. Christiano, and J. D. Fisher, (2001), Habit persistence, asset returns and the business cycle, *American Economic Review* 91, 149–166.

Brav, A., G. M. Constantinides, and C. Geczy, (2002), Asset pricing with heterogeneous consumers and limited participation: Empirical evidence, *Journal of Political Economy* 110, 793–824.

Brown, D. P., (1990), Age clienteles induced by liquidity constraints, *International Economic Review* 31, 891–911.

Cogley, T., (2002), Idiosyncratic risk and the equity premium: Evidence from the consumer expenditure survey, *Journal of Monetary Economics* 29, 309–334.

Constantinides, G. M., J. B. Donaldson, and R. Mehra, (2002), Junior can't borrow: a new perspective on the equity premium puzzle, *Quarterly Journal of Economics* 117, 269–296.

Constantinides, G. M. and D. Duffie, (1996), Asset pricing with heterogeneous consumers, *Journal of Political Economy* 104, 219–240.

Deaton, A. and C. Paxson, (1994), Intertepmoral choice and inequality, *Journal of Political Economy* 102, 437–467.

den Haan, W., (1994), Heterogeneity, aggregate uncertainty and the short term interest rate: a case study of two solution techniques, Working paper, University of California at San Diego.

den Haan, W., (1997), Solving dynamic models with aggregate shocks and heterogeneous agents, *Macroeconomic Dynamics* 1, 355–386.

Gertler, M., (1999), Government debt and social security in a life cycle economy, *Carnegie Rochester Conferance Series on Public Policy* 50, 61–110.

Gomes, F. and A. Michaelides, (2004), Asset pricing with limited risk sharing and heterogeneous agents, Working paper, LBS.

Graham, J. R., (2000), How big are the tax benefits of debt?, *Journal of Finance* 55, 1901–1941.

Greenwood, J., Z. Hercowitz, and G. W. Huffman, (1988), Investment, capacity utilization and the business cycle, *American Economic Review* 78, 402–18.

Greenwood, J., Z. Hercowitz, and P. Krusell, (1997), Long-run implications of investment-specific technological change, *American Economic Review* 87, 342–362.

Guvenen, F., (2005), A parsimonious macroeconomic model for asset pricing: Habit formation or cross-sectional heterogeneity?, Working Paper, University of Rochester.

Heathcote, J., K. Storesletten, and G. L. Violante, (2005), Two views of inequality over the life-cycle, *Journal of the European Economic Association (Papers and Proceedings)* 3 (2-3), 543–552.

Heaton, J. and D. J. Lucas, (1996), Evaluating the effects of incomplete markets on risk sharing and asset pricing, *Journal of Political Economy* 104, 443–487.

Heaton, J. and D. J. Lucas, (2000), Portfolio choice and asset prices; the importance of entrepreneurial risk, *Journal of Finance* 55, 1163–1198.

Huggett, M., (1993), The risk free rate in heterogeneous-agents, incomplete insurance economies, *Journal of Economic Dynamics and Control* 17, 953–969.

Huggett, M., (1996), Wealth distribution in life-cycle economies, *Journal of Monetary Economics* 38, 469–494.

Krueger, D. and F. Perri, (2006), Does income inequality lead to consumption inequality? Evidence and theory, *Review of Financial Studies* 73, 163–193.

Krusell, P. and A. A. Smith, (1997), Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns, *Macroeconomic Dynamics* 1, 387–422.

22

Krusell, P. and A. A. Smith, (1998), Income and wealth heterogeneity in the macro-economy, *Journal of Political Economy* 106, 867–896.

Lucas, D. J., (1994), Asset pricing with undiversifiable risk and short sales constraints: Deepening the equity premium puzzle, *Journal of Monetary Economics* 34, 325–341.

Mankiw, N. G., (1986), The equity premium and the concentration of aggregate shocks, *Journal of Financial Economics* 17, 211–219.

Marcet, A. and K. J. Singleton, (1999), Equilibrium assets prices and savings of heterogeneous agents in the presence of portfolio constraints, *Macroeconomic Dynamics* 3, 243–277.

Mehra, R. and E. Prescott, (1985), The equity puzzle, *Journal of Monetary Economics* 15, 145–61.

Ramchand, L., (1999), Asset pricing in international markets in the context of agent heterogeneity and market incompleteness, *Journal of International Money and Finance* 18, 871–890.

Ríos-Rull, J. V., (1994), On the quantitative importance of market completeness, *Journal of Monetary Economics* 34, 463–496.

Sarkissian, S., (2003), Incomplete consumption risk sharing and currency risk premiums, *Review of Financial Studies* 16, 983–1005.

Storesletten, K., (2000), Sustaining fiscal policy through immigration, *Journal of Political Economy* 108, 300–323.

Storesletten, K., C. I. Telmer, and A. Yaron, (2004a), Consumption and risk sharing over the life cycle, *Journal of Monetary Economics* 59(3), 609–633.

Storesletten, K., C. I. Telmer, and A. Yaron, (2004b), Cyclical dynamics in idiosyncratic labor market risk, *Journal of Political Economy* 112(3), 695–717.

Storesletten, K., C. I. Telmer, and A. Yaron, (2007), *Asset Prices and Intergenerational Risk Sharing: The Role of Idiosyncratic Earnings Shocks*, forthcoming, *Handbook of Investments: The Equity Risk Premium*, editor Rajnish Mehra, North Holland, Amsterdam.

Telmer, C. I., (1993), Asset pricing puzzles and incomplete markets, *Journal of Finance* 48, 1803–1832.

Weil, P., (1992), Equilibrium asset prices with undiversifiable labor income risk, *Journal of Economic Dynamics and Control* 16, 769–790.

Zhang, H., (1997), Endogenous borrowing constraints with incomplete markets, *Journal of Finance* 52, 2187–2209.

A Calibration Appendix

This appendix first describes the calibration of the no-trade (Constantinides and Duffie (1996)) economies in Section 3 and Table 1, and then goes on to describe the calibration of the economies with trade, presented in Section 4 and Table 2. It also demonstrates the sense in which our specification for countercyclical volatility — heteroskedasticity in the innovations to the idiosyncratic component of log income — is consistent with the approach used by previous authors (*e.g.*, Heaton and Lucas (1996), Constantinides and Duffie (1996)). In each case the cross sectional variance which matters turns out to be the variance of the *change* in the log of an individual's share of income and/or consumption.

Calibration of No-Trade Economies

Aggregate consumption growth follows an *i.i.d* two-state Markov chain, with a mean growth of 1.8% and standard deviation of 3.3%. This is essentially the process used in Mehra and Prescott (1985) with slightly more conservative volatility. The Constantinides and Duffie (1996) model is then 'calibrated' via a re-interpretation of the preference parameters of the Mehra and Prescott (1985) representative agent. Recall that we use β and γ to denote an *individual* agent's utility discount factor and risk aversion parameters, respectively. Constantinides and Duffie (1996) construct a representative agent (their equation (16)) whose rate of time preference and coefficient of relative risk aversion are (using our notation),

$$ -\log \beta^* = -\log(\beta) - \frac{\gamma(\gamma + 1)}{2} a \ \ , \tag{22} $$

and

$$ \gamma^* = \gamma - \frac{\gamma(\gamma + 1)}{2} b \ \ , \tag{23} $$

respectively. In these formulae, the parameters a and b relate the cross sectional variance in the change of the log of individual i's share of aggregate consumption (y_{t+1}^2, using Constantinides-Duffie's notation) to the growth rate of aggregate consumption:

$$ Var(\log \frac{c_{i,t+1}/c_{t+1}}{c_{it}/c_t}) = a + b \log \frac{c_{t+1}}{c_t} \ \ . \tag{24} $$

All that we require, therefore, are the numerical values for a and b which are implied by our PSID-based estimates in Table 1 of Storesletten, Telmer, and Yaron (2004b).

Our estimates are based on income, y_{it}. Because the Constantinides-Duffie model is autarkic, we can interpret these estimates as pertaining to individual consumption, c_{it}. Balduzzi and Yao (2000), Brav, Constantinides, and Geczy (2002), and Cogley (2002) take the alternative route and use microeconomic consumption data. While

their results are generally supportive of the model, they each point out serious data problems associated with using consumption data. Income data is advantageous is this sense. In addition, our objective is just as much relative as it is absolute. That is, consumption is endogenous in the model of Section 4, driven by risk sharing behavior and the exogenous process for idiosyncratic income risk. What Table 1 asks is, "what would the Constantinides-Duffie economy look like, were its agents to be endowed with idiosyncratic risk of a similar magnitude?" Also, "how does our model measure up, in spite of its non-degenerate (and more realistic) risk sharing technology?" Using income data seems appropriate in this context. For the remainder of this appendix we set $c_{it} = y_{it}$.

We need to establish the relationship between our specification for idiosyncratic shocks and the log-shares of aggregate consumption in equation (24). Denote individual $i's$ share at time t as γ_{it}, so that,

$$\log \gamma_{it} \equiv \log c_{it} - \log \tilde{E}_t c_{it} \ \ ,$$

where the notation $\tilde{E}_t(\cdot)$ denotes the cross-sectional mean at date t, so that $\tilde{E}_t c_{it}$ is date t, per-capita aggregate consumption. The empirical specification in Storesletten, Telmer, and Yaron (2004b) identifies an idiosyncratic shock as the residual from a log regression with year-dummy variables:

$$z_{it} = \log c_{it} - \tilde{E}_t \log c_{it} \ \ ,$$

which have a cross-sectional mean of zero, by construction, and a sample mean of zero, by least squares. The difference between our specification and the log-share specification is, therefore,

$$
\begin{aligned}
\log \gamma_{it} - z_{it} &= \tilde{E}_t \log c_{it} - \log \tilde{E}_t c_{it} \\
&= \tilde{E}_t \log \gamma_{it} - \log \tilde{E}_t \gamma_{it} \ \ .
\end{aligned}
$$

The share, γ_{it}, is *defined* so that its cross-sectional mean is always unity. The second term is therefore zero. For the first term, note that in both our economy and the statistical model underlying our estimates, the cross sectional distribution is log normal, *conditional* on knowledge of current and past aggregate shocks. If some random variable x is log normal and $E(x) = 1$, then $E(\log x) = -Var(\log x)/2$. As a result,

$$\log \gamma_{it} - z_{it} = -\frac{1}{2} \tilde{V}_t(\log \gamma_{it}) \ \ ,$$

where $\tilde{V}_t$ denotes the cross-sectional variance operator. Because lives are finite in our model, and because we interpret data as being generated by finite processes, this cross-sectional variance will always be well defined, irrespective whether or not the shocks are unit root processes.

The quantity of interest in equation (24) can now be written as,

$$
\begin{aligned}
\log \frac{c_{i,t+1}/c_{t+1}}{c_{it}/c_t} &\equiv \log \gamma_{i,t+1} - \log \gamma_{it} \\
&= z_{i,t+1} - z_{it} - \frac{1}{2}\left(\tilde{V}_{t+1}(\log \gamma_{i,t+1}) - \tilde{V}_t(\log \gamma_{it}) \right) \qquad (25)
\end{aligned}
$$

The term in parentheses — the difference in the variances — does not vary in the cross section. Consequently, application of the cross-sectional variance operator to both sides of equation (25) implies,

$$
\tilde{V}_{t+1}\left(\log \frac{c_{i,t+1}/c_{t+1}}{c_{it}/c_t} \right) = \tilde{V}_{t+1}\left(z_{i,t+1} - z_{it} \right) \quad .
$$

The process underlying our estimates is

$$
z_{i,t+1} - z_{it} = (1 - \rho)z_{it} + \eta_{i,t+1} \quad ,
$$

where the variance of $\eta_{i,t+1}$ depends on the aggregate shock. For values of ρ close to unity the variance of changes in z_{it} is approximately equal to the variance of $\eta_{i,t+1}$. The left side of equation (24) is, therefore, approximately equal to the variance of innovations, $\eta_{i,t+1}$,

$$
\tilde{V}_{t+1}\left(\log \frac{c_{i,t+1}/c_{t+1}}{c_{it}/c_t} \right) \approx \tilde{V}_{t+1}\left(\eta_{i,t+1} \right) \quad .
$$

For unit root shocks — which we assume for most of Section 4, this holds exactly. The estimates of σ_E and σ_C in Storesletten, Telmer, and Yaron (2004b), Table 1, are therefore sufficient to calibrate the Constantinides-Duffie model.

All that remains is to map our estimates into numerical values for a and b from equation (24). Since aggregate consumption growth is calibrated to be an $i.i.d$ process with a mean and standard deviation of 1.8% and 3.3% respectively, — aggregate consumption growth, the variable on the right hand side of equation (24), takes on only two values, 5.1% and -1.5%. Computing the parameters a and b, then simply involves two linear equations:

$$
\begin{aligned}
\sigma_E^2 &= a + 0.051b \\
\sigma_C^2 &= a - 0.015b \quad ,
\end{aligned}
$$

Storesletten, Telmer, and Yaron's (2004b) estimates are $\sigma_E^2 = 0.0156$ and $\sigma_C^2 = 0.0445$. These estimates, however, are associated with $\rho = .952$. For our unit root economies, we scale them down so as to maintain the same average unconditional variance (across age). This results in $\sigma_E^2 = 0.0059$ and $\sigma_C^2 = 0.0168$. The resulting values for a and b are $a = 0.0143$ and $b = -0.1652$.

27

Calibration of Models with Trade

The models in Section 4 are calibrated as follows. A period is interpreted as one year. The aggregate shock in equation (16) follows a first-order Markov chain with values $Z \in \{0.98, 1.02\}$. The unconditional probabilities are 0.5 and the transition probabilities are such that the probability of remaining in the current state is 2/3 (so that the expected duration of a 'business cycle' is 6 years). Capital's share of output, θ from equation (16), is set to 0.40, and the average annual depreciation rate, δ, is set to match the average riskfree rate of 1.3 percent. This results in $\delta = 8.8\%$. The parameter s is chosen so that the standard deviation of the risky return, r_t, is 10 percent.

Risk aversion is set to 8 so that the no-trade economy matches the Sharpe ratio (see Table 1). The discount factor β is chosen so that the average capital-output ratio is equal to 3.3.

The demographic structure is calibrated to correspond to several simple properties of the U.S. work force. The annual population growth is 1%. Agents are 'born' at age 22, retire at age 65 and are dead by age 85. 'Retirement' is defined as having one's labor income drop to zero and having to finance consumption from an existing stock of assets.

The following Table illustrates the aggregate properties of our economy. As in Mehra and Prescott (1985), the sample starts in 1929.

Panel A: Population Moments of HP-filtered data, Theoretical Economy

	Std Dev	Autocorrelation	Correlation with Output
Output	0.037	0.84	1.00
Investment	0.056	0.69	0.79
Consumption	0.037	0.64	0.90

Panel B: Sample Moments of HP-filtered data, U.S. Economy, 1929-2005

	Std Dev	Autocorrelation	Correlation with Output
Output	0.073	0.616	1.000
Investment	0.298	0.451	0.815
Consumption	0.036	0.697	0.887

U.S. sample moments are based on annual NIPA data, 1929-2005. Theoretical moments are computed as sample averages of a 20,000 periods long simulated time series. Both

the empirical and the simulated data are filtered using an HP-smoothing parameter of 400.

The calibration of our model is focused on asset pricing, targeting the variability of consumption and the variability of the return on equity. However, as is clear from this table, the production side of our economy is not too unrealistic. This holds true even if we were to focus on growth rates. Note that the large empirical volatility of investments is due to the sample incorporating the Great Depression.

B Computational Appendix

Our general solution strategy follows the work of den Haan (1994), den Haan (1997) and, in particular, Krusell and Smith (1997) and Krusell and Smith (1998). The crucial step is the specification of a finite dimensional vector to represent the law of motion for μ. Given this, each individual faces a finite-horizon dynamic programming problem. The essence of the fixed point problem is the consistency of the law of motion for μ with the law of motion implied by individual decisions. More specifically, our algorithm involves the following steps.

Algorithm

1. Approximate the distribution of agents, μ, with a finite number of moments or statistics, μ_m. The idea is to capture the information relevant for portfolio decisions in an efficient way as possible. Natural candidates are various moments of individual wealth and bond holdings. Instead, we use aggregate capital and the conditional expected equity premium ξ_t as moments.[12] Note that ξ_t is in an agent's period t information set. The seemingly unconventional state variable – a conditional price – captures in an efficient way the price information subsumed in a range of equity and bond-holding moments.

2. To solve agent's dynamic programming problem it is necessary to forecast both μ'_m and ξ'. We approximate the agents' expectations for the law of motion of μ_m and ξ by

$$(\mu'_m, \xi') = \hat{G}(\mu_m, \xi, Z, Z') = A(Z, Z') \times (\mu_m, \xi) \qquad (26)$$

where $A(Z, Z')$ is an $m \times (m + 1)$ matrix (conditional on Z and Z'). The aggregate shock Z can take on two values, $Z \in \{\underline{Z}, \bar{Z}\}$, so each element in the

---

[12]The conditional expected equity premium is defined as $\xi_t \equiv E_t\{R_{t+1}\} - q_t^{-1}$, where R_{t+1} is the return on equity in period $t + 1$ and q_t is the period t price of a claim that pays one unit of the consumption good in period $t + 1$. Note that, given ξ_t and conditional expectations over the future states of the world, the implicit bond price is $q_t = (E_t\{R_{t+1}\} - \xi_t)^{-1}$. We use ξ because it fluctuates substantially less than q, which implies that our approximation of the decision rules become more accurate.

matrix $A(Z, Z')$ above can take on four different values. Assume a particular set of values for $A(Z, Z') \ \forall Z, Z' \in \{\underline{Z}, \bar{Z}\}$.

3. Using the specification above, we solve the following modified version of (21):

$$\hat{V}_h(\xi, \mu_m, Z, z, \epsilon, a) = \max_{b'_{h+1}, k'_{h+1}} \{u(c_h) +$$

$$\beta E \left[\hat{V}'_{h+1}(\hat{G}(\mu, Z, Z'), Z', z', \epsilon', k'_{h+1} \cdot R(\hat{G}(\mu, Z, Z'), Z') + b'_{h+1}) \right] \} (27)$$

subject to (20).[13] The implementation of this is described below.

4. Assume an initial distribution of a large, but finite, number of agents, μ, across wealth, idiosyncratic shocks and age (we use 1000 agents in each age cohort). Using the decision rules obtained in (27), simulate a long sequence of the economy (20100 periods) and discard the first 100 periods from this sequence. Note that, for each period in time, ξ must be set so that the bond market clears. That is, find a ξ^* such that $\int b'_h(\mu_m, \xi^*, Z, z, \epsilon, a) d\mu = 0$. This is the sense in which ξ is an 'endogenous moment.'

5. Update $\hat{G}$ by running a linear regression of (ξ', μ'_m) on μ_m and ξ from the realized sequence in Step 4. If the coefficients change, use the updated $\hat{G}$ and return to Step 3. Continue this process until convergence.

6. Evaluate the ability of $\hat{G}$ to forecast μ'_m and ξ'. If the goodness of fit is not satisfactory, return to Step 1 and increase the number of moments or change the functional form of $\hat{G}$.

Moments of μ_m and accuracy

Following Krusell and Smith (1997), we began with just the first moment, aggregate capital, $\mu_1 = \log(\bar{k})$. This variable has strong predictive power on $\log(\bar{k}')$ (R^2 of 0.9998), but less predictive power on ξ' (see Table B1).

Next, we ask what other moment(s) matter for forecasting ξ'. To this end, we collected long time series of 18 additional moments of the distribution of agents (see Table B1 for details). Of these, the moments with the largest marginal improvement of forecast accuracy of ξ' (over the forecast including only $\log(\bar{k}')$) are the wealth of workers and the fraction of agents constrained in the bond market, which each improve the R^2 with on average 0.05 and 0.03, respectively. Including all the 18 moments (together with $\log(\bar{k})$) increase the forecast accuracy to 0.994. Finally, we regressed ξ' on ξ and $\log(\bar{k})$, and found that the R^2 increased to 0.9992, with a

---

[13]Note that in order to ensure that the bond market clears each period, ξ is included as an argument in the value function (Krusell and Smith (1997) use the bond price). One difference from their approach is that, as ξ enters the value function for all age groups, it does not simplify our computations to exclude ξ' from next period value functions and rely on "approximate" future market clearing.

standard deviation of forecast error less than 0.02% of ξ. Hence, ξ provides a better forecast than all the 18 moments together. Our interpretation of this finding is that the the conditional expected equity premium or, equivalently, the bond price, capture a large amount of information relevant for future bond prices. Formally, we could not reject the hypothesis that the residuals from predicting (μ'_m, ξ') by $\hat{G}(\mu_m, \xi, Z, Z')$ are uncorrelated with the 18 variables described above for the simulations we employ.

In summary, we include ξ as an "endogenous" moment in μ_m in order to improve the forecast of ξ'. Note that, as the value functions explicitly incorporate ξ as a parameter, including ξ in the forecast of ξ' come at zero computational cost.

Dynamic Programming Problem

We now describe how the dynamic programming problem in (27) is solved.

1. First, we choose a grid for the continuous variables in the state space. That is, we pick a set of values for $\bar{k}$, ξ, and a. The grid points are typically chosen to lie in the stationary region of the state variables and in addition, for wealth, near the borrowing constraint and far in excess of the maximum observed wealth holdings (conditional on age). We pick 13 points for aggregate capital, 13 points for the conditional expected equity premium, and 29 points for individual wealth at each age.

2. Second, we make piecewise linear approximations to the decision rules by solving for portfolio holdings on the grid and iterating on the Euler equations.

   This is done in the following way. Given the terminal condition associated with (20), the decision rules of the oldest agents (H years old) must be $b'_{H+1} = k'_{H+1} = 0$, in any state of the world. That is, the agent consumes all their wealth.

   Knowing c_H, we can in turn compute b'_H and k'_H at each grid point using Euler equations of an $H-1$ year old agent:

$$
\begin{aligned}
u'_{H-1}(c_{H-1}) &= E\{u'_H(c'_H)R' \mid \mu_m, Z, z, \epsilon\} \\
qu'_{H-1}(c_{H-1}) &= E\{u'_H(c'_H) \mid \mu_m, Z, z, \epsilon\}
\end{aligned}
\tag{28}
$$

   Knowing b'_H and k'_H at each grid point, we then obtain a piecewise linear approximation of the decision rules by linear interpolation (outside the grid we do linear extrapolation). Computing c_{H-1} is then straightforward, and this procedure is repeated for $H-2$ year old agents and iterated backwards until $h = 1$. Note that no further iterations are needed; given the (imperfect) expectations $\hat{G}$ and the decision rules for $h+1$ years old agents, the piecewise approximations are found in one single step for h years old agents.

Table B1
Predictability

Regressors	R^2			
$(Z_t = , Z_{t+1} =)$	(1,1)	(1,2)	(2,1)	(2,2)
$\log(\bar{k}), \xi$	.999	.999	.999	.999
$\log(\bar{k})$	.740	.907	.865	.922
$\log(\bar{k}), \{X_j\}_{j=1}^6$	.842	.931	.909	.941
$\log(\bar{k}), \{X_j\}_{j=1}^6, \{X_j^*\}_{j=1}^6$	.984	.997	.996	.996
$\log(\bar{k}), \{X_j\}_{j=1}^6, \{X_j^*\}_{j=1}^6, \{B_j\}_{j=2}^6$	.986	.998	.997	.997
$\log(\bar{k}), \{X_j\}_{j=1}^6, \{X_j^*\}_{j=1}^6, \{B_j\}_{j=2}^6, fr_{bt}, fr_{st}$	.987	.998	.997	.997
$\log(\bar{k}), X_1^*$	.803	.951	.922	.962
$\log(\bar{k}), fr_{bt}$	.802	.923	.898	.935

$X_j \equiv E[(a_{it} - \bar{a}_t)^j]$, a_{it} denotes the wealth of agent i at time t and $\bar{a}_t$ the average wealth at time t, $X_j^* \equiv E[(a_{it}^* - \bar{a}_t^*)^j]$, a_{it}^* denotes the wealth of *working* agent i and $\bar{a}_t^*$ the average wealth of workers at time t, $B_j \equiv E[(b_{it})^j]$, and b_{it} denotes the bond holdings of agent i at time t. fr_{bt} and fr_{st} denote the fraction of agents that are constrained in time t at the bond and equity market respectively. The bottom two rows represent the top two individual regressors among all regressors other than $\log(\bar{k})$ and ξ. The reported R^2 are of regressions of $\xi' = a(Z, Z') + X(Z, Z') * B(Z, Z')$

C Relationship to Companion Papers

This paper originally circulated as part of the working paper "Persistent Idiosyncratic Shocks and Incomplete Markets." It was subsequently split into two (distinct) companion papers, this paper and Storesletten, Telmer, and Yaron (2004b). The latter is an empirical study of CCV using PSID data. A third paper — Storesletten, Telmer, and Yaron (2007) — borrows the basic OLG model of Section 4 from this paper and uses it to focus upon the impact of having or not having CCV shocks on equity risk premia and portfolio rules. The two papers overlap somewhat in the discussion of portfolio behavior, but the basic questions and punchline are quite different. This paper asks (i) if retirement matters in an autarkic endowment economy and (ii) if the results change when we allow for capital accumulation and do a careful job of calibration. Storesletten, Telmer, and Yaron (2007) does not ask the first question. While it does use the issue of retirement to motivate the overall use of an OLG model for idiosyncratic risk and asset pricing, it does not conduct the retirement versus no-retirement experiments and does not make the central point about aggregate capital and self-insurance behavior with unit-root shocks. Its focus is on something absent in this paper: the marginal impact of CCV risk in a life-cycle context.

Table 1
Asset Pricing Properties – No Trade Economies

	Risk Aversion	Riskfree Rate Mean	Riskfree Rate Std Dev	Equity Premium Mean	Equity Premium Std Dev	Sharpe Ratio
U.S. data		1.30	1.88	6.85	16.64	41.17
U.S. data, unlevered		1.30	1.88	4.11	10.00	41.17
Models Without Trade (Constantinides-Duffie):						
No Retirement (match SR)	7.8	1.30	5.87	3.41	8.28	41.2
No Retirement	8.0	1.30	6.01	3.55	8.43	42.1
No Retirement (match EP)	8.6	1.30	6.55	4.11	8.95	45.9
Retirement (SR)	7.8	1.30	5.00	2.35	6.83	34.4
Retirement	8.0	1.30	5.02	2.42	6.89	35.1
Retirement (EP)	8.6	1.30	5.62	2.87	7.43	38.6

'Models Without Trade' correspond to a calibration of the Constantinides and Duffie (1996) model using the idiosyncratic risk estimates from Storesletten, Telmer, and Yaron (2004b), Table 1, and the aggregate consumption moments from Mehra and Prescott (1985). Details are given in Appendix A. In rows labeled 'match SR' and 'match EP,' risk aversion is chosen to match the U.S. Sharpe ratio and the mean equity premium, respectively. Rows labeled 'Retirement' hold risk aversion at these levels and then incorporate retirement, defined as old agents not receiving any idiosyncratic shocks (Section 3.2).

U.S. sample moments are computed using non-overlapping annual returns, (end of) January-over-January, 1956-1996. Estimates of means and standard deviations are qualitatively similar using annual data beginning from 1927, or a monthly series of overlapping annual returns. Equity data correspond to the annual return on the CRSP value weighted index, inclusive of distributions. Riskfree returns are based on the one month U.S. treasury bill. Nominal returns are deflated using the GDP deflator. All returns are expressed as annual percentages. Unlevered equity returns are computed using a debt to firm value ratio of 40 percent, which is taken from Graham (2000).

Table 2
Asset Pricing Properties – Economies with Trade

|	Risk Aversion	β	K/Y	σ_C^2	σ_E^2	Riskfree Rate Mean	Equity Premium Mean	Equity Premium Std Dev	Sharpe Ratio
Complete Markets	3	0.965	3.3	0	0	4.4	0.63	7.0	9.0
Incomplete Markets (with Retirement)	3	0.948	3.3	0.0168	0.0059	2.3	0.86	6.7	12.8
Complete Markets	8	0.959	3.3	0	0	4.8	2.28	9.1	25.0
Incomplete Markets (with Retirement)	8	0.8	3.3	0.0168	0.0059	1.6	2.23	6.8	32.6
Incomplete Markets (Without Retirement)	8	0.977	3.3	0.0168	0.0059	1.3	2.31	7.5	30.9

Population moments from the models described in Section 4. The calibration procedure is discussed in the text and appendix. β is the discount factor, K/Y is the average capital/output ratio, and σ_C^2 and σ_E^2 are the variances of the permanent innovations to earnings in recessions (contractions) and booms (expansions), respectively. All economies are calibrated so that aggregate consumption volatility is 3.3%. In economies without retirement, agents work and live until age 85. In economies with retirement, agents work until age 65 and live until age 85, receiving zero labor income during the retirement years. Idiosyncratic shocks are calibrated so the unit root economy has the same average volatility as that in an economy based on the estimates of Storesletten, Telmer, and Yaron (2004b).
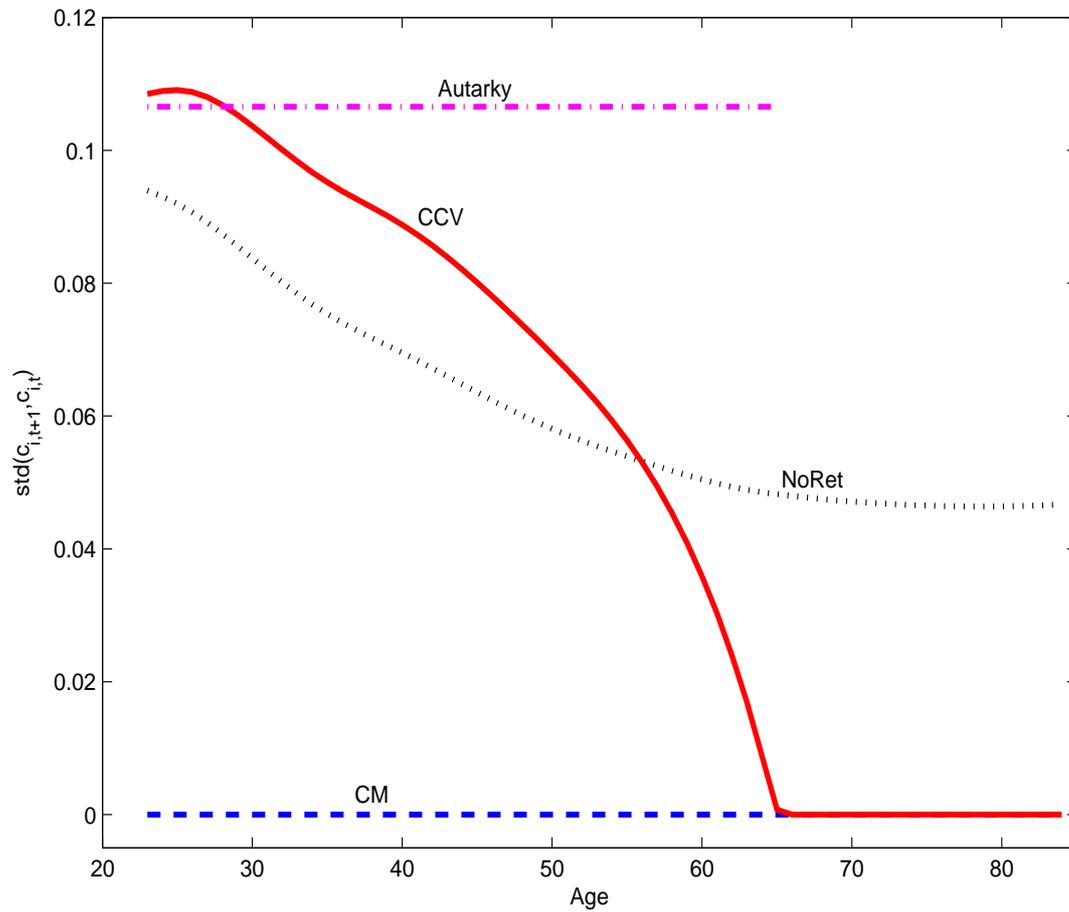
Figure 1
Financial Wealth by Age
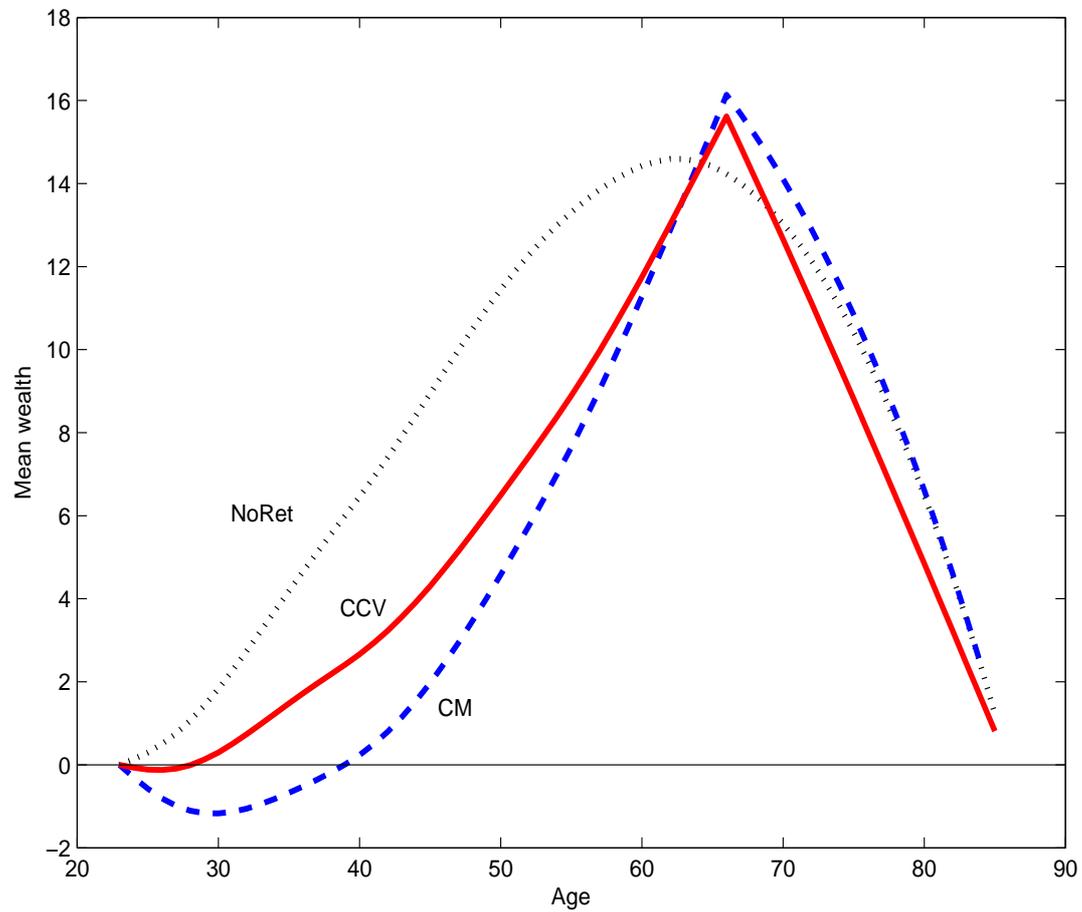
Age profile of average financial wealth (across agents of the same age) for three different economies: complete markets (CM), incomplete markets with retirement (CCV), and incomplete markets without retirement (NoRet).
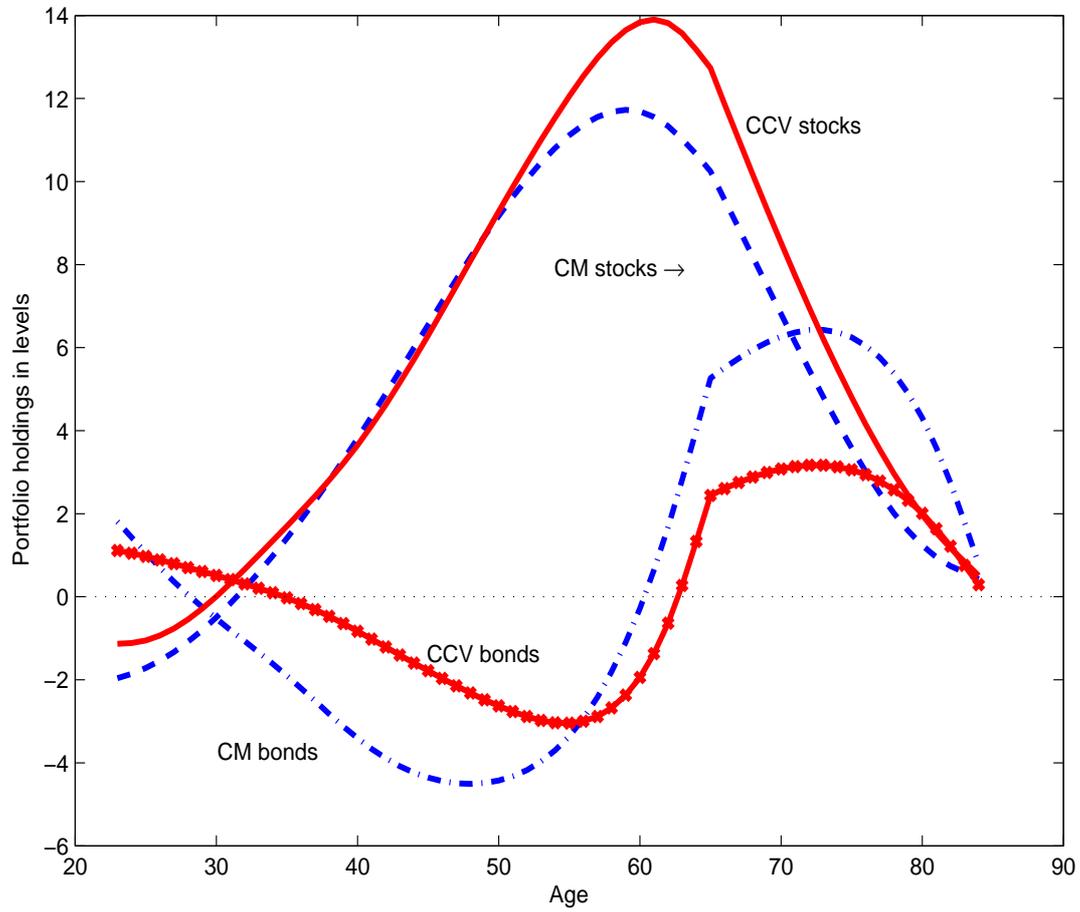
Figure 2
Aggregate Risk Bearing by Age
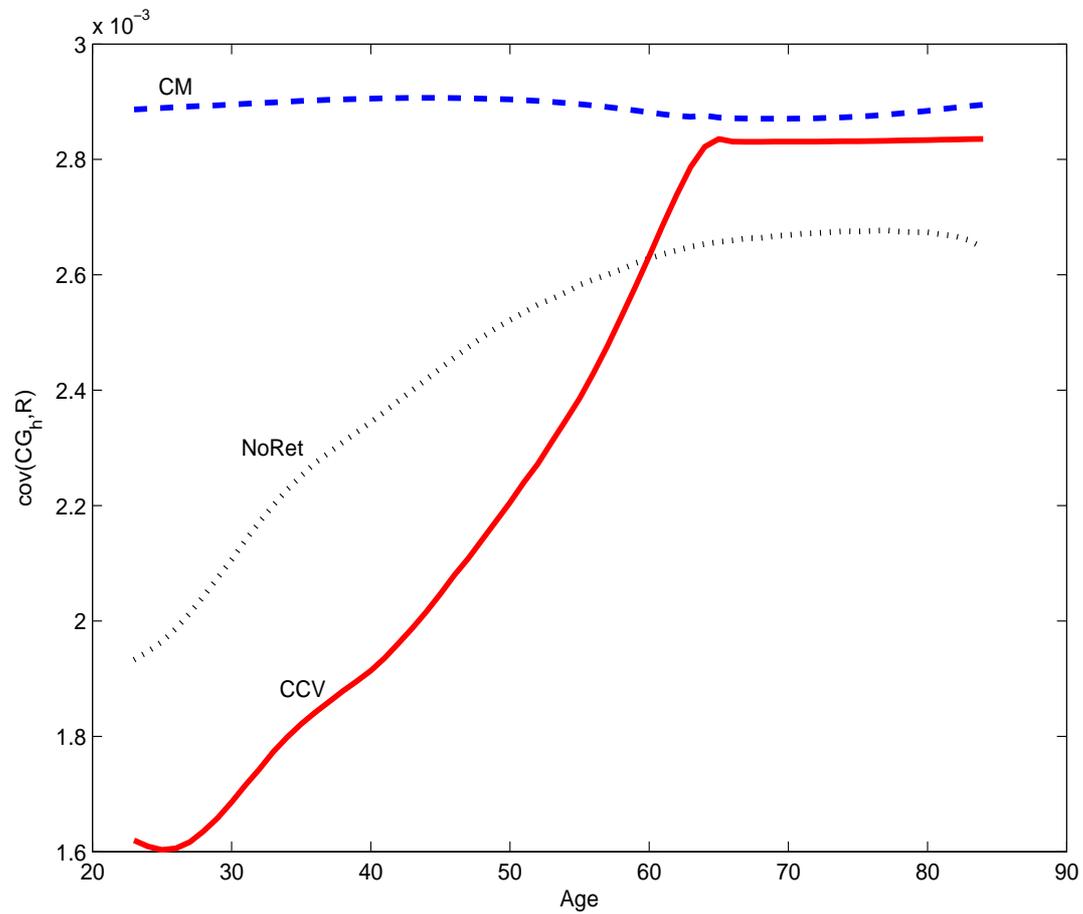
The measure of 'aggregate risk bearing' described in Section 5 of the text: the age-specific covariance between individual consumption growth and the return on equity. The lines correspond to complete markets (CM), incomplete markets with retirement (CCV), and incomplete markets without retirement (NoRet).

Figure 3
Standard Deviation of Consumption Growth by Age

The lines correspond to complete markets (CM), incomplete markets with retirement (CCV), and incomplete markets without retirement (NoRet).

Figure 4
Quantity of Bonds and Stocks, by Age

The lines labeled 'CM' correspond to the complete-markets economy. Those labeled 'CCV' correspond to the incomplete-market economy, with retirement.

Figure 5
Bond and Stock Portfolio Shares, by Age
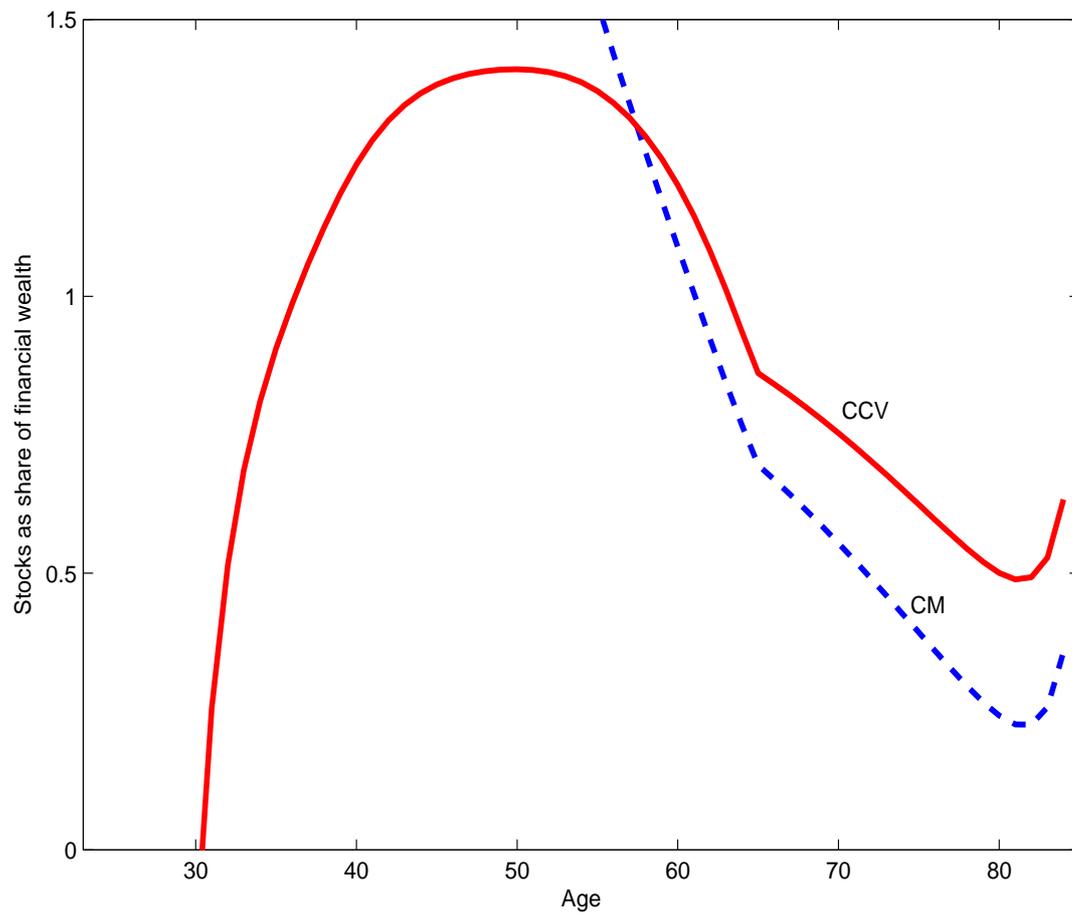
The line labeled 'CM' corresponds to the complete-markets economy. Those labeled 'CCV' corresponds to the incomplete-market economy, with retirement.