Asset Prices and Intergenerational Risk Sharing – the Role of Idiosyncratic Earnings Shocks*

Kjetil Storesletten[†], Chris Telmer[‡], and Amir Yaron[§]

October 2005

Abstract

Constantinides and Duffie (1996) show that if idiosyncratic labor income shocks (i) are highly persistent, and (ii) become more volatile during economic contractions, the impact on asset prices can be substantial. We argue that life-cycle effects also play a fundamental role. We use a stationary overlapping-generations model to show that life-cycle effects can mitigate the equity premium because aggregate risk can be shifted to retirees who do not face labor risk. However, in spite of this channel, our model, with realistic life cycle features, can still account for about 75% of the average equity premium and the Sharpe ratio observed on the U.S. stock market.

---

*Prepared for the *Handbook of Investments: Equity Risk Premium*, edited by Rajnish Mehra. This chapter builds on the work in Storesletten, Telmer and Yaron (2001a).

[†]University of Oslo and CEPR.

[‡]Tepper School of Business, Carnegie Mellon University.

[§]The Wharton School, University of Pennsylvania and NBER.

1 Introduction

This chapter analyzes the channels by which idiosyncratic earnings shocks affect asset prices. We show that the asset pricing effects of idiosyncratic shocks depend on the relative magnitude of human capital shocks to financial risk. The exposure to this dimension changes dramatically over the life cycle – making life cycle consideration an important element of our framework. More specifically, we show that individuals' portfolio choices are sensitive to idiosyncratic earning shocks when (i) they display persistence and counter-cyclical volatility – features documented using earnings data, and that (ii) the shocks to human capital are large relative to financial risk when young. Young agents are more exposed to human capital shocks since they hold little financial capital and because the shocks are persistent they have a large impact on discounted earnings (i.e., human capital). As agents age, the exposure to financial (human) capital increases (declines). Agents with no labor income – retirees – generally prefer to reduce their exposure to financial risk and thus to aggregate risk. However, with counter-cyclical volatility of earnings shocks, agents who are fully exposed to human capital risk – the youngest agents – would also like to reduce their exposure to aggregate risk. As a consequence, our model displays average portfolio rules which are hump-shaped in age. That is, agents choose to hold very little equity when young, levered equity positions when middle-aged, and relatively small equity positions when retired.

We show that in equilibrium these portfolio choices manifest themselves to asset prices that depend on how much intergenerational risk sharing takes place across cohorts. Without idiosyncratic shocks and with constant aggregate wages (and thus return on human capital), the old would prefer to issue debt while the young would prefer to be levered in equity. Such intergenerational risk sharing would tend, *ceteris paribus*, to lower the equity premium. However, the presence of idiosyncratic shocks dissuades the young from holding equity while the relative volatile risky equity dissuades the retirees from taking levered positions in equity, and thus curtails intergenerational risk sharing.

Following the work in Storesletten, Telmer, and Yaron (2001a) we show that quantitatively the effects of idiosyncratic risk are large. Idiosyncratic risk inhibits intergenerational risk sharing, imposing a disproportionate share of aggregate risk on the wealthy middle-aged cohorts who demand an equity premium for doing so. We use a stationary overlapping-generations model to show how life-cycle portfolio choices interact with intergenerational risk sharing to accentuate the equity premium. For a risk aversion of 8 our model is able to account for about 75% of the average equity premium and the Sharpe ratio observed on the U.S. stock market.

It is important to note that the driving force in our model is a concentration of equity ownership on middle-aged agents. The model of Constantinides, Donaldson, and Mehra (2002) shares a similar feature. Where their model is driven by portfolio constraints, however, ours is driven by portfolio choices made in light of how nontradeable and tradeable risks interact. We discuss these issues further below.

An advantage of our model relates to risk-sharing behavior. U.S. data on income and consumption indicate that, while complete markets may not characterize the world, neither does a distinguishing feature of the Constantinides and Duffie (1996) framework: autarky. The cross-sectional standard deviation of U.S. consumption, for instance, is roughly 35 percent smaller than that of non-financial earnings.[1] However, our framework shows that in spite of the autarky dimension and the lack of realistic life cycle features, the Constantinides-Duffie model is still able to provide useful quantitative asset pricing results. Papers by Balduzzi and Yao (2000), Brav, Constantinides, and Geczy (2002), Cogley (2002) Ramchand (1999) and Sarkissian (2001) investigate the Constantinides-Duffie model's consumption Euler equation restrictions directly and find mixed evidence. Our approach, emphasizing endogenous asset pricing, seem to be consistent with the channels put forth in their paper.

The idea that market incompleteness may contribute to the equity premium is not new, and by and large most of the quantitative findings have been 'negative' in terms of the ability of the proposed models to generate a viable equity premium. The agents in these models tend to be 'very efficient' in insuring themselves against the relatively transitory income shocks they face; thus the resulting equity premium is essentially the one derived in Mehra and Prescott (1985) (e.g., Telmer (1993), Heaton and Lucas (1996)). To generate a sizeable equity premium such models usually had to resort to large transaction costs and/or tight borrowing constraints. The distinguishing aspect of our paper, life-cycle, has also important implications for the ability to model persistent (almost unit root) processes for individual earnings shocks while maintaining aggregate quantities that are characterized by relatively transitory processes. A number of studies have examined more specifically the quantitative implications of the Constantinides and Duffie (1996) model. The closest to our work is Krusell and Smith (1997) and Gomes and Michaelides (2004). The latter use a similar life cycle model, but focus on the role of fix entry costs for matching the relative size of stockholders. Others include, Aiyagari (1994), Aiyagari and Gertler (1991), Alvarez and Jermann (2001), den Haan (1994), 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). More recent examples using life cycle economies to asses portfolio choice and equity returns include Olovsson (2004) and Benzoni, Colin-Dufresne, and Goldstein (2004).

[1] A large number of papers, including Altonji, Hayashi, and Kotlikoff (1991), Attanasio and Davis (1996), Attanasio and Weber (1992), Cochrane (1991) Deaton and Paxson (1994) Mace (1991) and Storesletten, Telmer, and Yaron (2004a) provide evidence which is suggestive of imperfect risk sharing. Altug and Miller (1990) find opposing evidence. The numerical value cited here is based upon evidence in both Deaton and Paxson (1994) and Storesletten, Telmer, and Yaron (2004a).

The remainder of this chapter is organized as follows. In Section 2 we formulate a life-cycle version of the Constantinides and Duffie (1996) model, calibrates it, and examine its quantitative asset pricing properties. We then introduce retirement in two steps. First, we assume that the oldest 20 percent of the population do not receive idiosyncratic shocks, but do receive retirement income. This formulation can be studied within the Constantinides-Duffie no-trade framework. Second, we introduce life-cycle savings by assuming that retirement income is equal to zero. Equivalently, we introduce a more realistic distribution of human to financial wealth. This is done in a class of computational economies, *with* trade, presented in Section 2.2. Section 3 concludes.

2 Model(s)

In this section we describe the life cycle model we use for our quantitative analysis. We start, however, with a three-period OLG example that is designed to highlight the role of risk sharing within a life cycle model. We then proceed to describe our calibrated model. This model, although admittedly not parsimonious, allows us to analyze the role life cycle and intergenerational risk sharing play in determining the equity premium.

2.1 An OLG Version of the Constantinides-Duffie 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,

$$z_{it}^h = z_{i,t-1}^{h-1} + \eta_{it} \quad (4)$$

$$\eta_{it} \sim \mathrm{N}(-\sigma_t^2/2\,,\,\sigma_t^2) \quad (5)$$

$$\sigma_t^2 = a + b\log(G_t/G_{t-1}) \quad (6)$$

$$z_{i,t}^0 = 0\ . \quad (7)$$

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. Equilibrium is autarkic with individual consumption $c_{it}^h = G_t \exp(z_{it}^h)$. Bond and equity prices satisfy

$$q_t = \beta^* E_t \lambda_{t+1}^{-\gamma^*} \quad (8)$$

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

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.

2.1.1 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 12%, and (c) the conditional standard deviation is countercyclical, increasing by roughly 67% from expansion to contraction (from 8.7% to 14.6%). In Appendix A we show that these estimates map into values $a = 0.0181$ and $b = -0.2067$.

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 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 $\gamma = 7.3$. To match the equity premium $\gamma^* = 15.42$ is required, which corresponds to $\gamma = 8.0$. Time preference is characterized by $\beta^* = 1.140$ ($\beta = 0.65$) and $\beta^* = 1.148$ ($\beta = 0.59$), 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 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.

2.1.2 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 2.2.

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

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

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

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

$$y^3_{it} \quad = \quad f_t \, G_t \exp(z^3_{it}) - (p_t + D_t) \quad , \tag{12}$$

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

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

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

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

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 4.1 percent, roughly 25% larger 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 5th and 6th rows of Table 1. Holding risk aversion fixed, 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 6), the Sharpe ratio and equity premium fall from 45.9 percent to 38.6 percent and 4.1 percent to 2.1 percent, respectively.

To summarize, retirement, defined here as old agents receiving fixed incomes, 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.

2.2 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^h_{it}. 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 (*e.g.*, 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 (i) we can make the model more realistic along certain dimensions which are important for calibration (*e.g.*, the demographic structure) and (ii) the model will display partial risk-sharing behavior — an undeniable aspect of U.S. data on income and consumption — even with unit root idiosyncratic shocks. With this in mind, we make the following changes to the framework of Section 2.

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 production technology. The main reason for adding production is computational tractability: the resulting price of the risky asset will always be equal to unity. Agents rent capital and labor to a single 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

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

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, w controls the average wage rate, and δ_t is the depreciation rate on capital. The depreciation rate is stochastic:

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

where δ controls the average and s is, approximately, the standard deviation of r_t.[2]

This production process delivers three key ingredients: (i) the model is tractable (solving the analogous endowment economy is substantially more difficult),(ii) the volatility of aggregate consumption growth can be calibrated realistically, and (iii) the average return on human capital - i.e., the wage rate - can be substantially less volatile than the return on equity. As we show below, the implied equilibrium volatility of the return on equity is not too far (accounting for standard errors) from its empirical counterpart. Each ingredient is critical for our question. The first is obvious. The second 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

[2] 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).

impact of idiosyncratic risk. The fourth is instrumental in determining which age cohorts hold equity in equilibrium and, consequently, whether or not idiosyncratic risk is priced.

These ingredients come at a cost. They imply, for example, excessively volatile investment and output, a feature shared by most existing production-based models should they be calibrated to have realistic variability in asset returns. We do not resolve such issues here. We view our model in the same way we view an endowment economy; as an economy with a potentially unrealistic production side which, nevertheless, yields informative restrictions on consumption and asset returns.

Endowments.

The endowments (2)–(7) are of a special form required to support an autarkic outcome. Since this is no longer required, and because of the incorporation 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 nontradeable endowment now takes the form of labor efficiency units, not units of the single good.[3] At time t the ith working agent of age h is endowed with n_{it}^h units of labor which they supply inelastically. Retirees, agents for whom h exceeds a retirement age $\hat{H}$, 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 $z_{it}^0 = 0$. For computational reasons, 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: $y_{it}^h = w_t n_{it}^h$.[4]

[3]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.

[4]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.

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), 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.

2.2.1 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 of agents across ages, idiosyncratic shocks and wealth. The existence of aggregate shocks implies that, necessarily, μ must 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,

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

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. The third equation rules out shortselling (which turns out to be innocuous) and the the fourth restricts terminal wealth to be non-zero. Note that, beyond terminal wealth, we do not impose borrowing constraints.

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

$$\begin{aligned} V_h(\mu, Z, z_h, a_h) \ = \ \max_{k'_{h+1}, b'_{h+1}} \Big\{ u(c_h) + \\ \beta E \left[V'_{h+1}(G(\mu, Z, Z'), Z', z'_{h+1}, k'_{h+1} r(G(\mu, Z, Z'), Z') + b'_{h+1}) \right] \Big\} \ , \end{aligned} \tag{21}$$

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 Storesletten, Telmer, and Yaron (2001a).

2.3 Quantitative Properties

Our model now has three main motives for trade: the life-cycle distribution of idiosyncratic shocks, the life cycle distribution of the ratio of human to total wealth, and the possibility that $\rho < 1$. In order to focus on life-cycle issues and maintain a tangible link with the Constantinides-Duffie benchmark, we concentrate on the case of $\rho = 1$.

We calibrate the above economy according to the criteria outlined in Appendix A. The most important features are as follows.

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 $\rho < 1$ estimates. This results in $\sigma_E = 0.0871$ and $\sigma_C = 0.1457$.

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

3. The standard deviation of aggregate consumption growth is set to match 3.3 percent. As equations (8) and (9) emphasize, realistic properties for aggregate consumption are essential here, just as they are in representative agent models. The cost, in our case, is excessively volatile output and investment, something which is commonplace in models with production. Full details are provided in Storesletten, Telmer, and Yaron (2001a)

4. The standard deviation of the return on capital affects the aggregate volatility of consumption growth. For the reasons stated above we equate the volatility of consumption growth across the various economies by altering the volatility of the return to capital. The latter critically depends on the stochastic depreciation process (18) – so in essence we look for a depreciation process that will guarantee a 3.3% volatility for consumption and that in turn pins down the volatility of the return to capital. We therefore have little to say about why the return on the equity market is as variable as it is. What we can say, however, is that the main consequence — the return on financial capital being substantially more volatile than the return on human capital — has stark implications for life-cycle portfolio choice and, therefore, for how idiosyncratic shocks interact with asset pricing.

5. Young agents are born with zero assets and retired agents receive zero labor income. This serves as the primary motive for trade. It also results in a realistic life-cycle distribution of human to financial capital — younger agents hold most of the former whereas older agents hold most of the latter — which, as we'll see, plays an important role in portfolio choice.

6. Retired agents comprise roughly 20 percent of the population.

Table 2 reports the Sharpe ratio, the risk free rate, and the first two moments of the risky return in this economy. Each row describes a different economy: in the case of complete markets, there are no idiosyncratic shocks to earnings and only aggregate shocks are operative; in the case of 'no ccv' there are idiosyncratic shocks but they are homoskedastic with respect to aggregate shocks. Rather than search for a risk aversion that would match the equity premium or the Sharpe ratio we investigate these economies for two alternative levels of risk aversion (3 and 8).[5] This table demonstrates several features of the effects of trade. First, the life cycle model, in spite of the effects of retirement as analyzed earlier, can still generate a substantial equity premium. In fact for a risk aversion of 8, the implied risk premium and Sharpe ratio are about 2.5% and 33% respectively – approximately the same magnitude as for the retirement economy in Table 1. Second, adding counter-cyclical variation to idiosyncratic shocks contributes about .5% to the equity premium and about 7% to the Sharpe ratio (an increase of about 25%). In summary, trade in conjunction with life-cycle features provide results that to a large extent are consistent with the no-trade Constantinides-Duffie model (with and without retirement).

To understand these results it is instructive to first understand the complete markets economy as a useful benchmark. Figure 1 provides the wealth distribution in

[5]Interestingly, Cogley (2002) uses a risk aversion of 8 in calibrating the equity premium based on time varying cross-sectional moments of consumption growth from his findings from the CEX coupled with a plausible levels of measurement errors.

both the complete markets as well as our benchmark ccv-economy. A noticeable feature of this figure is the concentration of wealth among the middle aged agents. In addition, note that the overall wealth (the area under the figure) is equated across the two economies. Finally, note that in the economy with idiosyncratic shocks the young save relatively more than in the complete markets economy. This is simply a manifestation of the precautionary savings motive in the incomplete markets economy (there is a long literature discussing this in partial equilibrium; Aiyagari (1994) is a classic reference within an infinite horizon general equilibrium model, while Storesletten, Telmer, and Yaron (2004a) analyze this motive extensively within a life-cycle model).

The key to understanding the equity premium that arise within our model's life cycle economy lies in the portfolio choices. The portfolio behavior is depicted in Figures 2 and 3. The first of these two figures depicts the level of stock and bond holdings while the second depicts the portfolio shares – that is the figure reports the fraction of financial wealth invested in bonds and stocks for the average individual of each age. The portfolio choices are displayed for both the complete markets and the ccv incomplete market economy.

In the complete markets economy all agents share a standard deviation of consumption growth of 3.3% that is equal to that of aggregate consumption growth. Retirees have large investment positions in bonds and workers (age 30 and above) have relatively large positions in stocks. Early in life, due to CRRA preferences and zero initial wealth, working agents start by having long position in bonds. A few years later these agents move into long positions in stocks while shorting bonds. Over time agents decrease the portfolio shares invested in equity (see Figure 3) – a feature due to the effect put forth in Bodie, Merton, and Samuelson (1992). The idea being that agents decrease their allocations to risky assets in order to maintain a constant riskiness level of their overall portfolio. This follows since in retirement the retirees no longer face labor earnings which are assumed to be substantially less risky than equity (see an extensive discussion of this issue in Jagannathan and Kocherlakota (1996)). The outcome of all of these effects is that even in the complete markets economy we observe a hump-shaped age profile of portfolio holdings.

Next we turn to the portfolio holdings in the ccv economy. Relative to the complete markets economy we see a clear shift in portfolio holdings. Workers reduce their exposure to equity holdings because of the counter-cyclical risk dimension and consequently increase their relative bond holdings. Retirees on the other hand hold relatively more stocks. Note, that this follows exactly the intuition described in our earlier analytical example. That is the introduction of counter cyclical variation leads retirees to bear more aggregate risk. The key feature of our economy with trade is that quantitatively the equity premium is retained in spite of the opportunities to shift aggregate risk to the retirees. The reason is that absent Bodie, Merton, and Samuelson (1992) motive, the retirees could have handled all of the "shifted" aggregate risk and no large premium would be required. A retiree holding all of their wealth as equity would face substantially more aggregate risk than would a worker holding all of their *financial* wealth as equity, because the worker also owns some human wealth. The result, as we see in Figure 3, is that retirees choose to diversify and allocate some of their wealth toward bonds. Thus, the presence of risky return and the lack of earnings shocks inhibit retirees from absorbing *all* of the aggregate risk that the young workers would like to shield away from themselves. To counteract this, the wealthy middle aged agents stand to fill this gap but demand a premium for doing so.

For our story to hold it is important that the age portfolio profiles do not change substantially as we look at other parts of the wealth distribution, as all the figures thus far were for the *average* investor. It turns out the portfolio choice patterns underlying the model's equity premium are consistent across the wealth distribution. The portfolio choice profiles underlying our analysis utilized the average investor. Figure 7 shows that the hump shape profile is also a robust feature for agents in the top 10% of the wealth distribution. Figure 3 shows the hump share profile is also present in the complete markets economy. This due to a large extent because agents in our economy start out with zero wealth. In the incomplete markets case wealth does affect the age profile of portfolio investment; it is clear that the share invested in risky equity is substantially larger relative to the average wealth as depicted in Figures 3. Figures 8 and 9 provide the portfolio share age profile as a function of the aggregate state for both the complete market and ccv economy respectively. The upshot is that these portfolio profiles do not change in a significant manner across recessions and expansions, although as expected there is a slight increase in equity investment during expansions.

To summarize, the intuition for the equity premium in the economy we present is quite simple: Young workers dislike the risky asset because of the countercyclical nature of the idiosyncratic risk they face. Retirees dislike it because it has a highly variable return and they no longer receive labor income. What Figure 3 shows is that middle-aged workers represent a bridge between the two; they dislike aggregate risk for the same reasons, but in each case to a lesser degree. They hold part of their wealth as financial wealth and therefore care less than the young do about idiosyncratic shocks. They face the same variability in equity returns as the old, but their labor income mitigates their overall exposure to aggregate shocks. The end-result is that the middle-aged hold levered equity by issuing bonds to the young and the old. The resulting hump-shaped pattern in equity ownership 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).[6]

[6]Brown (1990) shows that non-tradeable labor income can generate hump-shaped portfolio rules in age. Amerkis and Zeldes (2000) discuss a similar phenomenon. Our computational solution also features hump-shaped decision rules (with age) for the share of financial wealth held as bonds. The

This outcome has a natural interpretation in terms of the intergenerational redistribution of aggregate risk. With the introduction of counter-cyclical variation risk, young workers face an additional source of aggregate risk, manifested in the conditional variance of their idiosyncratic risk process.[7] Thus, the first order effect of efficiently allocating counter-cyclical risk is in shifting aggregate risk towards the retirees. However, the return on capital is more variable than the aggregate component of the return on labor. This implies that retirees would like in general to reduce their exposure to the risky asset. Thus, an efficient allocation will stop short of transferring aggregate risk completely to the retirees, resulting instead in the hump-shape portfolio profile pattern we see in Figure 3.

2.4 Intergenerational Risk Sharing

An 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 better characterizes the world. This seems important for the question, 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, our economy displays a type of contingent, self-insurance behavior.

The risk-sharing behavior of our model can be measured by the relative volatility of consumption to that of income. In the 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 roughly 35 percent (see Deaton and Paxson (1994) or Storesletten, Telmer, and Yaron (2004a)), so our model exhibits too little risk sharing. Figure 4 reports the cross-sectional variance in the *growth rate* of consumption, which is more directly related to the essence of risk sharing: the equalization of marginal rates of substitution. In this case, we see a larger difference between our model and the autarkic outcome. Autarky implies that, for workers, the graph is flat at 0.17,

cross-sectional average in Figure 2 inherits this shape because financial capital's share of total wealth is, on average, an increasing function of age.

[7] In Storesletten, Telmer, and Yaron (2001b) we examine the welfare consequences of this form of aggregate risk more explicitly. We find that the welfare costs of business cycles can be quite large, should the elimination of business cycles also imply the elimination of heteroskedasticity in idiosyncratic shocks.

as shown. 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.

As mentioned earlier, our model's asset pricing implications crucially depend on aggregate risk being shifted across generations. This should manifest itself into evidence of intergenerational risk sharing. To that account Figure 5 plots the *cohort* specific standard deviation of consumption growth. That is we first aggregate consumption over all members of cohort j and then compute the standard deviation of this consumption growth with respect to all cohorts of a specific age. The graph clearly shows a pronounced difference between the age-profile of cohort standard deviation for the ccv-economy and the homoskedastic economy. While the latter profile is pretty much flat across age, the aggregate risk measured by the cohort's standard deviation increases with age for the economy with counter cyclical variation risk.

In an attempt to formalize the degree of cohort risk sharing Figure 6 plots the exposure of each cohort's consumption growth to aggregate consumption fluctuations. Specifically, we let $x_{iht} \equiv= \log(c_{iht}/c_{i,h-1,t-1})$ denote the consumption growth of person i of age h at time t. We define aggregate consumption growth as $X_t \equiv \frac{1}{I}\frac{1}{H}\sum_i \sum_h$. Defining δ^c_{iht} to be the idiosyncratic consumption component of agents belonging this cohort, the cohort's consumption growth, x_h, can be written as

$$x_{ht} = x_h + a_h X_t + \epsilon_{ht} + \delta^c_{iht}$$

where by construction $\sum_i \delta^c_{iht} = 0$. The choice of a_h sets $corr(\epsilon_{ht}, X_t)$ to zero, so that ϵ_{ht} can be interpreted as the cohort specific component of consumption growth which is unrelated to X_t. Thus, the constant a_h is interpreted as the share of aggregate risk in consumption. The figure displays the fact that for the complete markets and the homoskedastic economies, age does not play a key role and cohorts are equated in terms of bearing aggregate risk. On the other hand for the economy with counter cyclical variation risk, aggregate risk bearing is monotonically increasing with age. This is consistent with the idea that the young are shifting risk to older generation. The older generations in return demand a premium for shielding the young from aggregate risk.

3 Conclusions

Our main question is whether idiosyncratic labor-market risk matters for the pricing of aggregate risk. An inescapable aspect of idiosyncratic risk is that it necessarily has a life-cycle component: the young face more than the old. This arises both directly — workers face shocks but retirees don't — and indirectly, in terms of the inevitable life-cycle pattern in the ratio of human wealth to total wealth. These life-cycle effects are of first-order importance for the question. They imply that a substantial fraction of the population don't care very much about the very shocks which drive the model, and therefore present a challenge to the asset-pricing story. Nevertheless, our main conclusion is that idiosyncratic risk matters and the equity premium can be sizeable even with life-cycle effects present.

What drives our model is an interaction between idiosyncratic and aggregate risk which goes beyond the countercyclical-volatility effect emphasized by Constantinides and Duffie (1996), Mankiw (1986) and many subsequent papers. The converse of the fact that younger agents face the most idiosyncratic risk is that older agents face the most aggregate risk. Idiosyncratic risk is difficult to transfer across generations. Aggregate risk is not. Our framework suggests that how the aggregate risk is shared, and how this interacts with the nontradeable distribution of idiosyncratic risk, is important for asset pricing. If, for instance, an equilibrium features equity ownership increasing with age, then the effect of idiosyncratic risk will be diminished relative to the Constantinides and Duffie (1996) model. If equity ownership decreases with age, the opposite will hold. Our calibration generates an intermediate case: equity ownership increases until the late working years and then declines into retirement. The asset pricing effects of idiosyncratic shocks and the countercyclical-volatility effect remain important and our model generates significant risk premia that are consistent with those generated within the infinite-horizon Constantinides and Duffie (1996) model.

More specifically, our model is driven by life-cycle variation in the ratio of human wealth to financial wealth and the fact that idiosyncratic risk 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 countercyclical-volatility 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 effects imply that 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. The risk premium which supports this allocation reflects both the countercyclical-volatility risk emphasized by Constantinides and Duffie (1996), and a "concentration of aggregate risk" upon the middle-aged, alluded to by Mankiw (1986).

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 the countercyclical volatility 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.

Altonji, J. G., F. Hayashi, and L. J. Kotlikoff, (1991), Risk sharing, altruism, and the factor structure of consumption, NBER working paper number 3834.

Altug, S. and R. A. Miller, (1990), Households choices in equilibrium, *Econometrica* 58, 543–70.

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.

Attanasio, O. and S. J. Davis, (1996), Relative wage movements and the distribution of consumption, *Journal of Political Economy* 104, 1227–1262.

Attanasio, O. and G. Weber, (1992), Consumption growth and excess sensitivity to income: evidence from U.S. micro data, Unpublished manuscript, Stanford University.

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

Benzoni, L., P. Colin-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.

Cochrane, J. H., (1991), A simple test of consumption insurance, *Journal of Political Economy* 99, 957–76.

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.

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.

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.

Jagannathan, R. and N. Kocherlakota, (1996), Why should older people invest less in stocks than younger people?, *Federal Reserve Bank of Minneapolis, Quarterly Review* 20, 11–23.

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

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.

Mace, B. J., (1991), Full insurance in the presence of aggregate uncertainty, *Journal of Political Economy* 99, 928–56.

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.

Olovsson, C., (2004), Social security and the equity premium puzzle, Working paper, IIES.

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

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

Sarkissian, S., (2001), The cross-country consumption dispersion and the world business cycle, forthcoming, *Review of Financial Studies.*

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

Storesletten, K., C. I. Telmer, and A. Yaron, (2001a), Asset pricing with idiosyncratic risk and overlapping generations, Working paper, Carnegie Mellon University.

Storesletten, K., C. I. Telmer, and A. Yaron, (2001b), The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk, *European Economic Review* 45, 1311–1339.

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.

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 2 and Table 1, and then goes on to describe the calibration of the economies with trade, presented in Section 2.2 and Table 1. 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^2_{t+1}, 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 2.2, 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,

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

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:

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

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,

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

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,

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

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 2.2, 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 $\sigma_E^2 = 0.00758$ and $\sigma_C^2 = 0.02122$. The resulting values for a and b are $a = 0.0181$ and $b = -0.2067$.

Table 1
Asset Pricing Properties – No Trade Economies

	Risk	Riskfree Rate		Equity Premium		Sharpe
	Aversion	Mean	Std Dev	Mean	Std Dev	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)	7.3	1.30	5.87	3.41	10.35	41.2
No Retirement (match EP)	8.0	1.30	6.55	4.11	11.4	45.9
Retirement (SR)	7.3	1.30	5.00	2.35	8.59	34.4
Retirement (EP)	8.0	1.30	5.62	2.87	9.49	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 2.1.2).

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_E^2	σ_C^2	Riskfree Rate Mean	Equity Premium Mean	Equity Premium Std Dev	Sharpe Ratio
Complete markets	3	0.965	3.3	0	0	4.4	0.72	7.5	9.6
no CCV	3	0.953	3.9	0.0114	0.0114	2.0	0.80	7.3	11.0
estimated CCV	3	0.953	3.7	0.0168	0.0059	2.7	0.85	7.2	11.8
Complete markets	8	0.96	3.3	0	0	4.8	2.82	11.0	25.7
no CCV	8	0.809	3.3	0.0114	0.0114	–	2.00	8.8	25.0
estimated CCV	8	0.801	3.3	0.0168	0.0059	2.5	2.43	7.6	32.0
large CCV	8	0.797	3.3	0.0204	0.0023	2.0	2.51	6.6	38.0

'Models with Trade,' are described in Section 2.2. The calibration procedure is discussed in the text. All economies are calibrated so that aggregate consumption volatility is 3.3%. The 'Homoskedastic Economy' is distinguished by the volatility of idiosyncratic shocks *not* varying with aggregate shocks. The 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

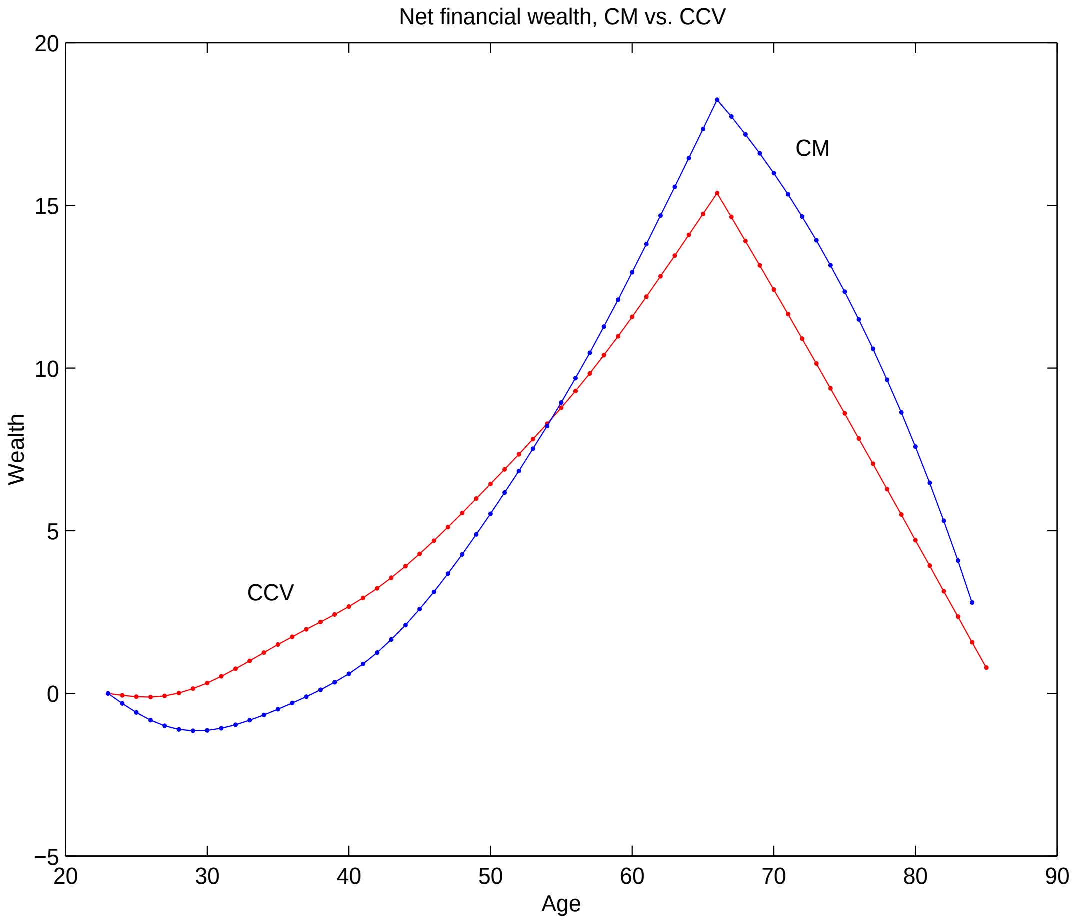

Figure 2
Quantity of Bonds and Stocks, by Age

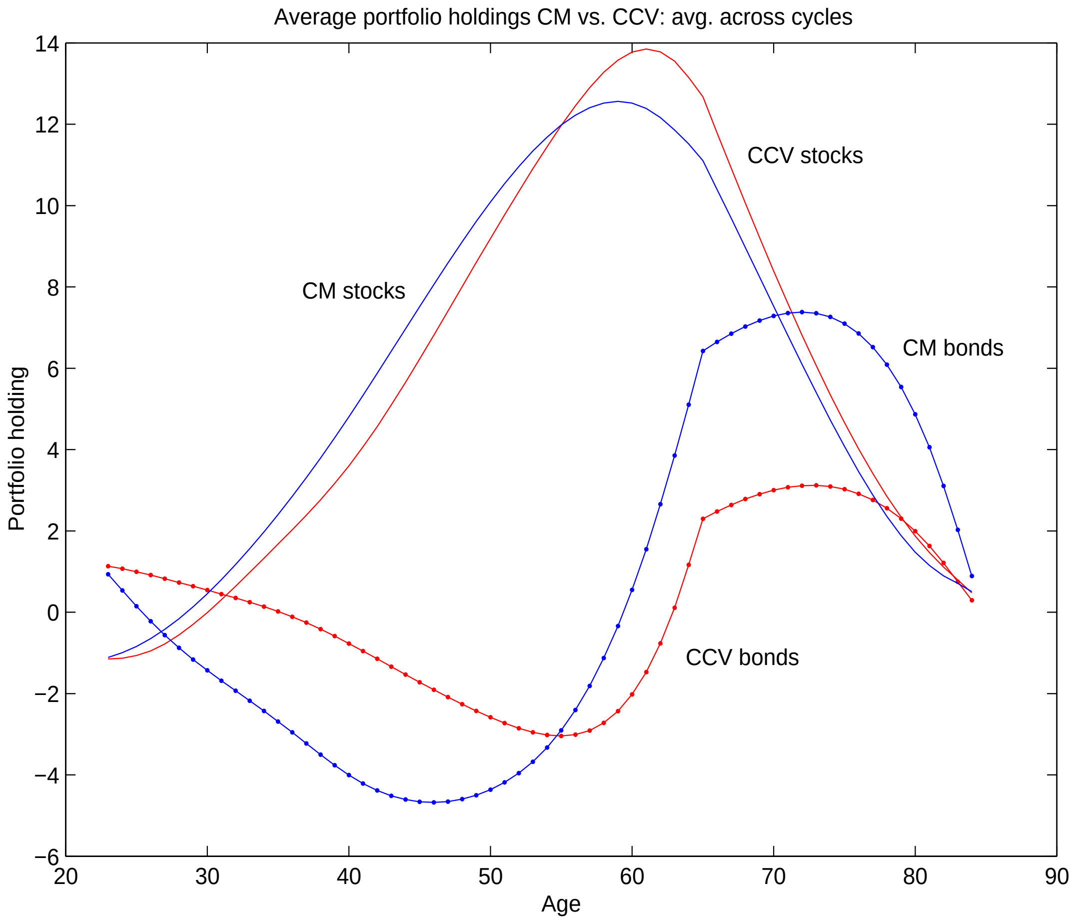

Figure 3
Bond and Stock Portfolio Shares, by Age

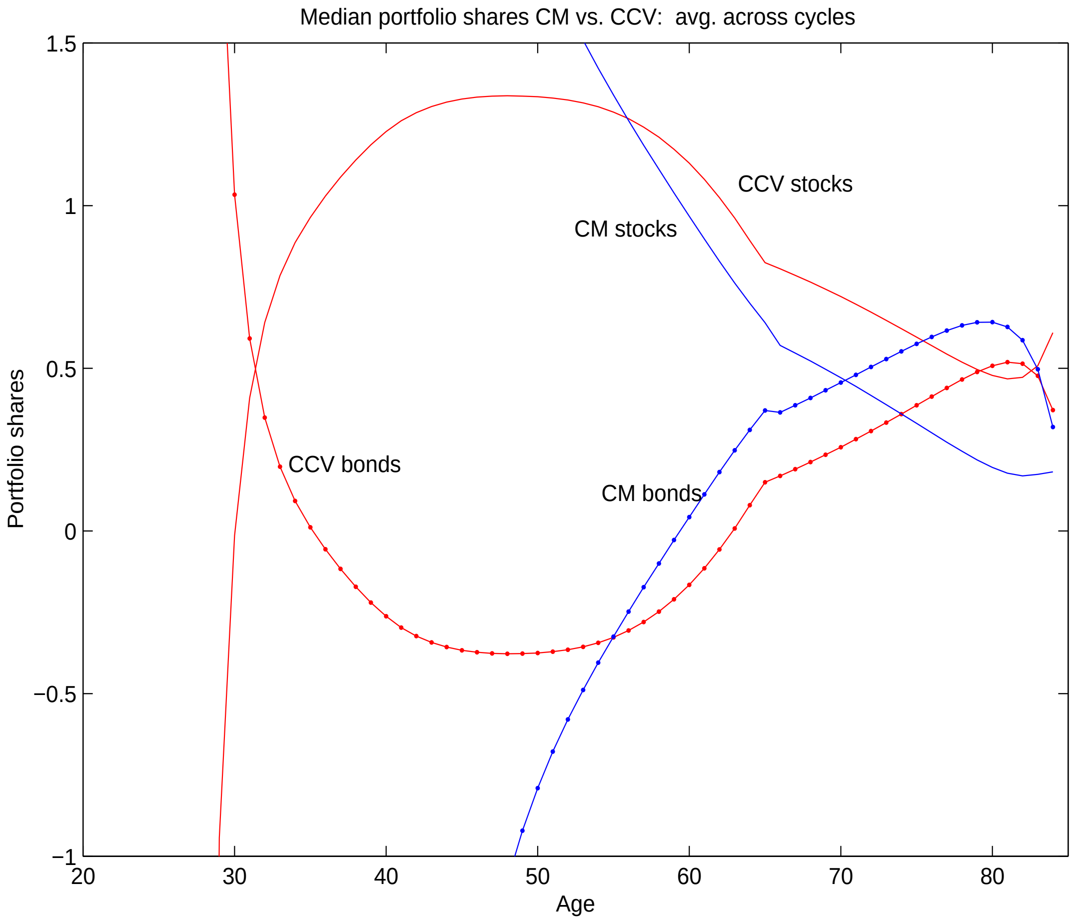

Figure 4
Variance of Log Consumption

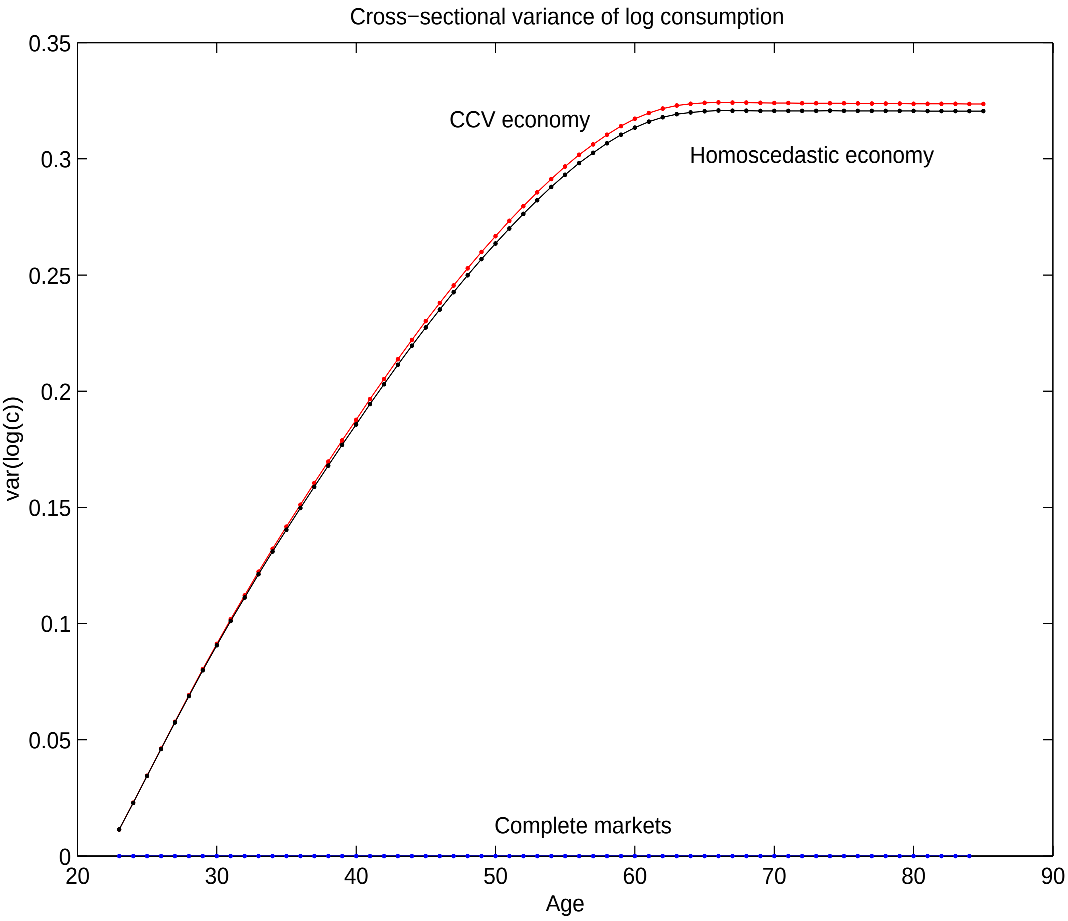

Figure 5
Standard Deviation of Consumption Growth by Age

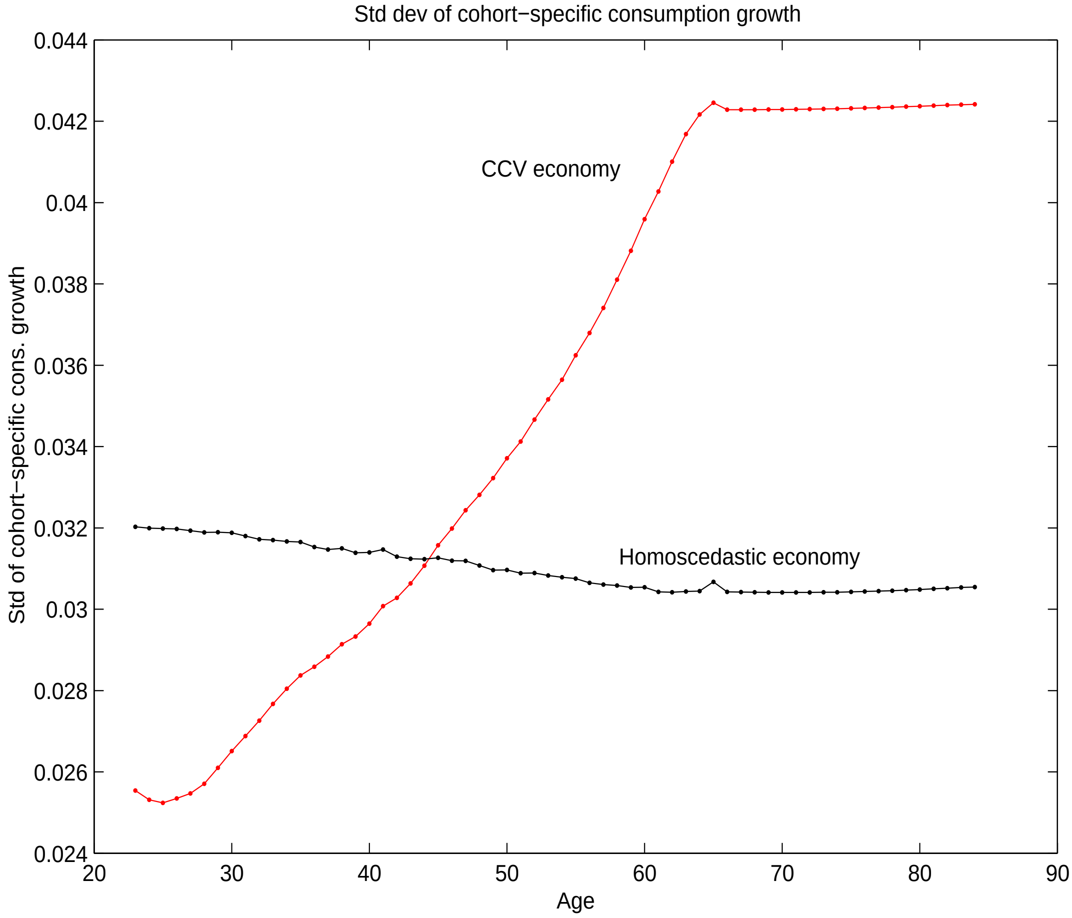

Figure 6
Aggregate Risk Bearing by Age

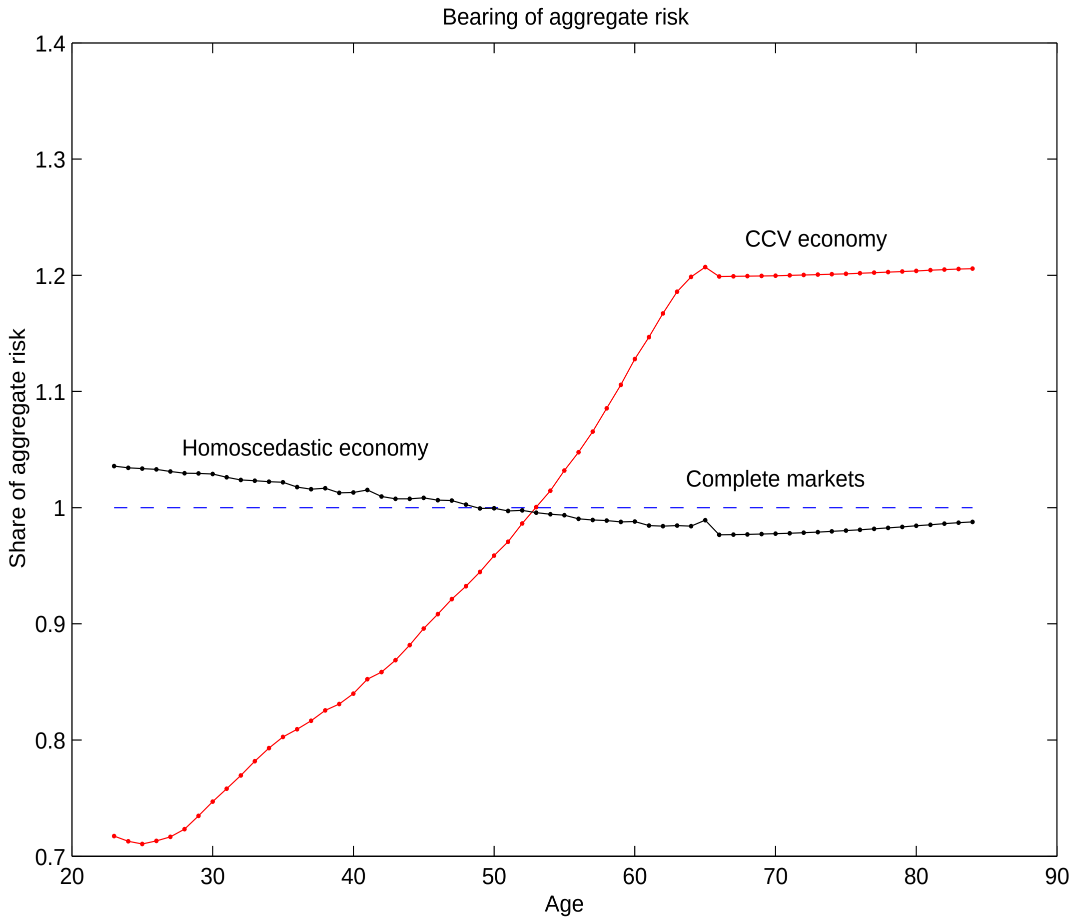

Figure 7
Portfolio Shares for Top 10 Earnings Percentile, by Age

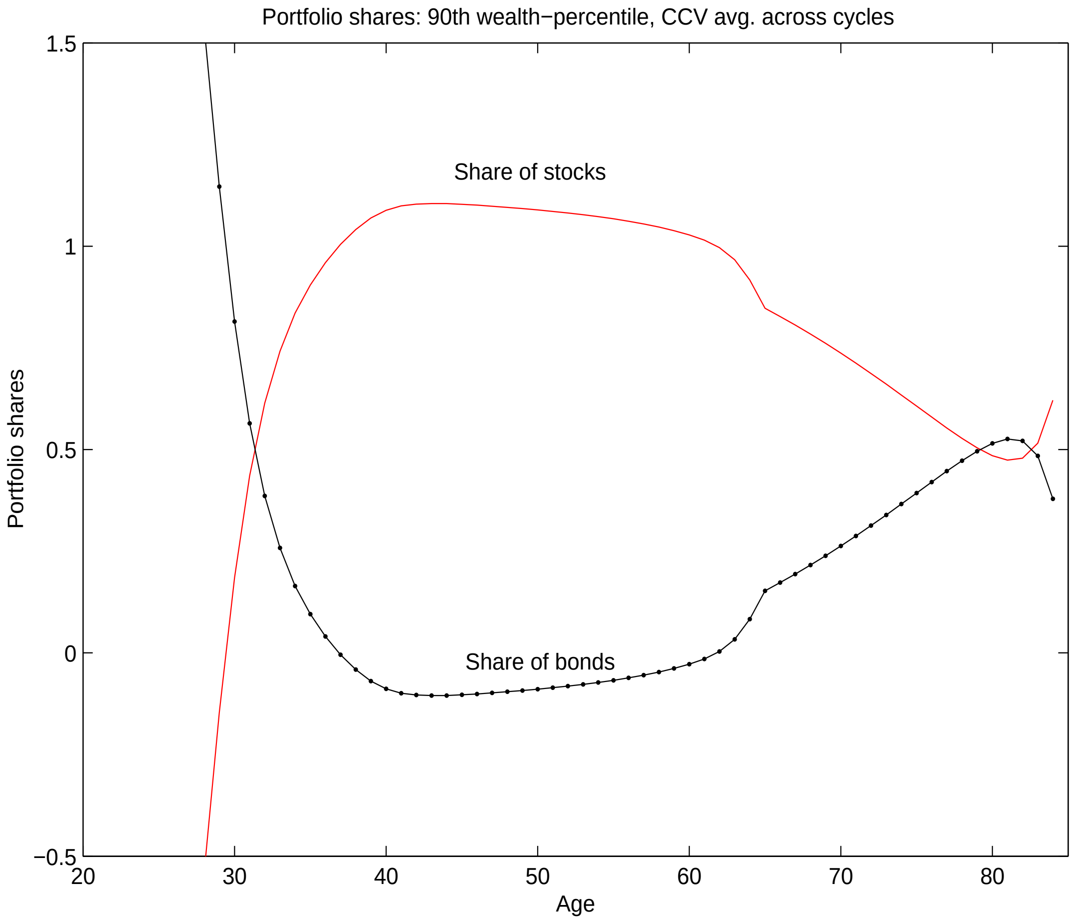

Figure 8
Portfolio Shares conditional on Business Cycle, by Age

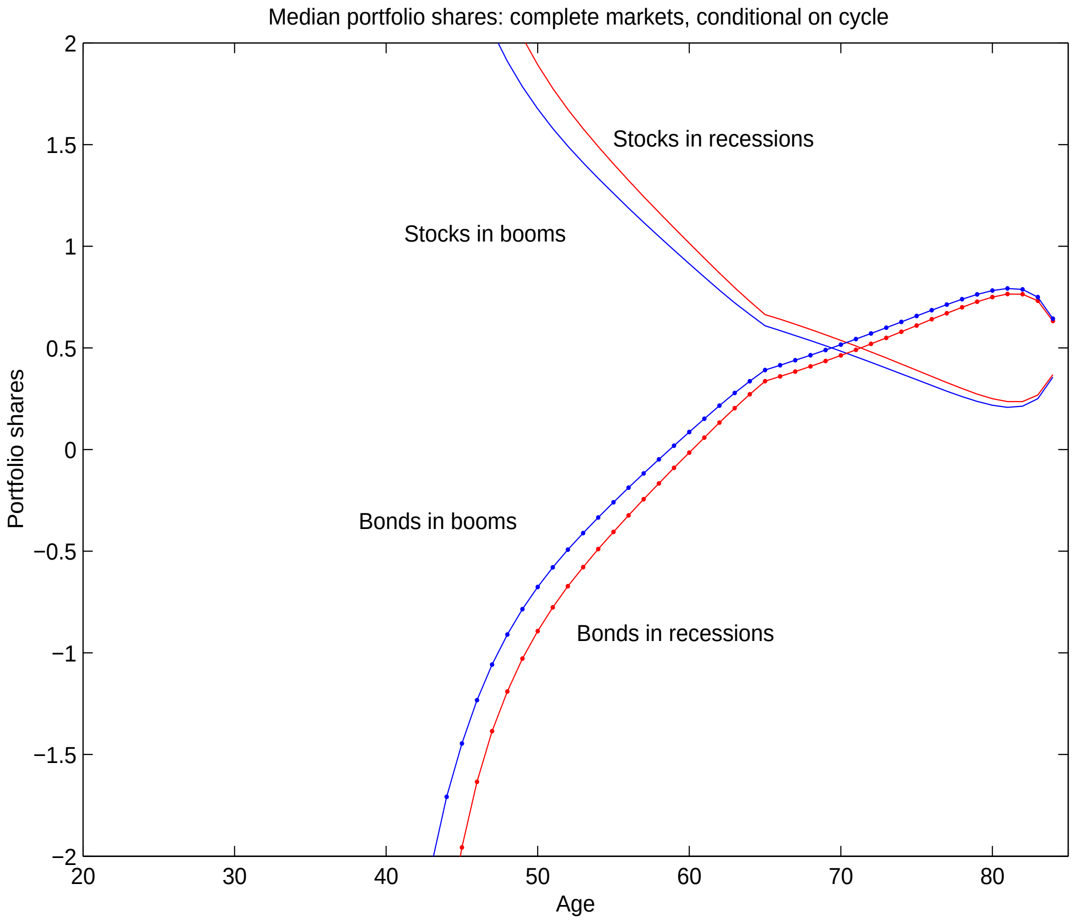

The solid line conditions on aggregate expansions. The dashed line conditions on aggregate contractions.

Figure 9
Portfolio Shares conditional on Business Cycle, by Age

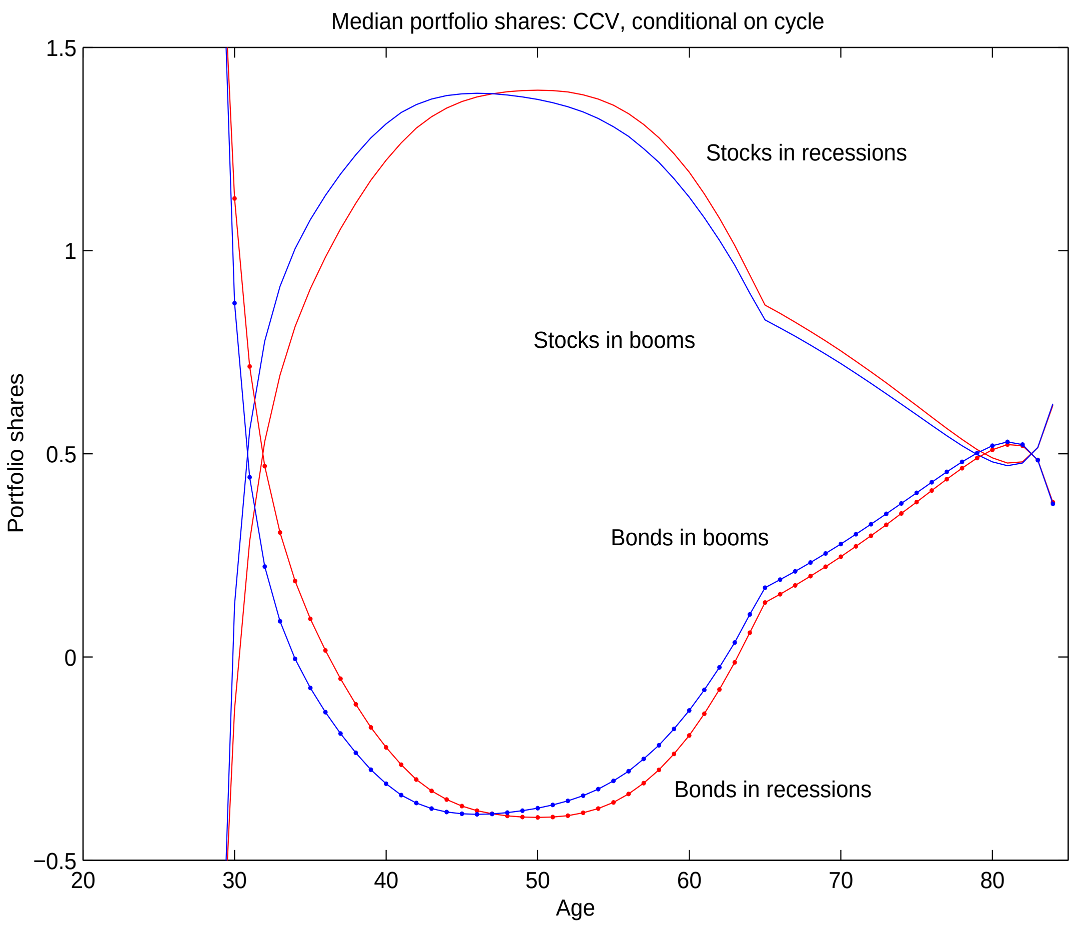

The solid line conditions on aggregate expansions. The dashed line conditions on aggregate contractions.