Asset Pricing with Idiosyncratic Risk and Overlapping Generations[1]

Kjetil Storesletten,[2] Christopher I. Telmer,[3] and Amir Yaron[4]

February 2006
(first draft: June 1996)

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-income risk. However, in spite of this channel, our model can still account for about 75% of the average equity premium and Sharpe ratio observed on the U.S. stock market. The main reason that idiosyncratic risk matters for asset pricing is because it inhibits the intergenerational sharing of aggregate risk.

[1]A predecessor of this paper circulated as "Persistent Idiosyncratic Shocks and Incomplete Markets." In addition to numerous participants at seminars and conferences, we thank Dave Backus, Geert Bekaert, David Bowman, Martin Browning, John Cochrane, Wouter den Haan, Mick Devereux, Ron Gallant, René Garcia, Rick Green, John Heaton, Burton Hollifield, Mark Huggett, Narayana Kocherlakota, Per Krusell, Deborah Lucas, Pierre Mella-Barral, Bob Miller, Michael Parkin, Monika Piazzesi, Victor Ríos-Rull, Amlan Roy, Tony Smith, Paul Willen, Steve Zeldes and Stan Zin for helpful comments and suggestions. We have benefitted from the support of NSF grant SES-9987602 and the Rodney White Center at Wharton. Valuable research assistance was provided by Xiaorong Dong.

[2]University of Oslo and CEPR

[3]Tepper School of Business, Carnegie Mellon University

[4]Wharton School, University of Pennsylvania and NBER

1 Introduction

The essence of Mehra and Prescott's (1985) equity premium puzzle is that investing in equity looks like a free lunch; an individual who consumes aggregate consumption places far greater value on the stock market than does the stock market itself. A large literature has asked if aggregation lies at the heart of the puzzle. The idea is that individuals face idiosyncratic risks, 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 subject to substantial risks and is larger in magnitude than financial wealth. 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 average 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.

Much of the literature — exemplified by Constantinides and Duffie (1996) and Mankiw (1986) — has overcome this challenge by assuming that aggregate shocks and the *volatility* of idiosyncratic shocks are negatively related. We label this effect *countercyclical cross-sectional variance*, or CCV. We begin by examining the quantitative merits of CCV. We calibrate a version of the Constantinides-Duffie model based on estimates from the Panel Study on Income Dynamics (PSID), taken from Storesletten, Telmer, and Yaron (2004b). We find that by incorporating idiosyncratic risk the risk aversion coefficient required to account for the U.S. Sharpe ratio is 7.8. Absent idiosyncratic risk a value of 13.6 is necessary. The Constantinides-Duffie model, therefore, is successful at resolving the equity premium puzzle in the (narrow) sense that, given estimates from labor market data, very high risk aversion is not required.

Where do we go from here? The Constantinides-Duffie model abstracts from trade and imperfect risk-sharing. Both are important features of the world. We incorporate them and ask how the model's asset-pricing implications change. Trade occurs for many reasons. We emphasize life-cycle reasons. We think that they are of first-order importance for the question. Our reasons are as follows.

'Human wealth' is the value of an individual's future labor income. The main question — that of our paper, Constantinides and Duffie (1996), and most related work — boils down to whether idiosyncratic shocks to human wealth affect the valuation of financial wealth. The distribution of human wealth necessarily has a life-cycle dimension: the young have more than the old. The same is true, therefore, of idiosyncratic risk. Our premise is that this life-cycle distribution of idiosyncratic risk is important for the equity premium puzzle. Consider, for example, retired people. They comprise roughly 20 percent of the adult population, they participate in equity

markets at an even higher rate, yet they face little if any labor market risk. If, as the Constantinides-Duffie model suggests, the solution to the equity premium puzzle is labor-market risk which makes equity more risky, then why don't retirees hold lots of equity, thus resurrecting the puzzle?

Our answer is that they do, but that it's not a complete disaster for the idiosyncratic risk story. We use an OLG model with non-tradable labor-income risk. With a risk-aversion coefficient of 8, our model's Sharpe ratio is 33%, which is halfway between the complete-market value of 25% and the Constantinides-Duffie value of 41% (which matches U.S. data by construction). What drives this is CCV. We show that, in a life-cycle context, CCV motivates the young to hold relatively little equity, thus concentrating aggregate risk on older agents and increasing the market risk premium. 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 told not to.

The economic force through which CCV operates in our model is a generalization of Bodie, Merton, and Samuelson (1992) (BMS). In their model, life-cycle wages are like a bond. This induces agents to hold a lower fraction of their financial wealth 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, like the retirees, hold relatively little stock. That is, portfolio rules are hump-shaped.[1] 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. 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 instanace, Benzoni, Collin-Dufresne, and Goldstein (2004)).

Our results can also be interpreted in the context of intergenerational risk sharing. In our model, the main reason that idiosyncratic risk matters for asset pricing is that it inhibits the intergenerational sharing of *aggregate* risk. We demonstrate this by showing that, in the absence of CCV, aggregate risk sharing resembles that of the analogous complete-market allocation; it is uniform across generations. CCV changes this. It induces a transfer of aggregate risk from the young to the old, who demand a risk premium to bear it.

There is much related work which examines the quantitative implications of the

[1]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)).

Constantinides and Duffie (1996) model.[2] To understand our contribution, it is important to understand their 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, asset 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 (2001) 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 focuses on these restrictions.[3] 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.

The remainder of our paper 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 3. Section 5 concludes.

[2]Our paper is also related to a large body of work on asset pricing with heterogeneous agent models. Most closely related is Krusell and Smith (1997), which serves as an important benchmark, both in terms of computational methods and results. Our paper is primarily distinguished by its emphasis on the life cycle. In addition, they require extreme borrowing constraints (essentially zero) in order to generate risk premiums, whereas borrowing constraints play no role in our study. Other related papers include 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]Krusell and Smith (1997), emphasize the Constantinides-Duffie model's unrealistic implications for financial wealth distribution. We focus, instead, on its implications for the distribution of financial wealth relative to human wealth. Given the inherent nature of the question — How do nontradeable shocks to human wealth affect the pricing of financial wealth? — this seems of primary importance.

2 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,

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

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,

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

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

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

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 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 $\gamma = 8.6$.[4] To facilitate a direct comparison with the numerical results (shown

[4]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).

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 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.

2.2 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 3.

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

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

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

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

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

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

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

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

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

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 4.2 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, 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.

3 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 (*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:

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

where δ controls the average and s is, approximately, the standard deviation of r_t.[5]

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. 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 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.

[5]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).

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.[6] 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: $y_{it}^h = w_t n_{it}^h$.[7]

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.

3.1 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, $\mathcal{H} = \{1, 2, \ldots, H\}$, $\mathcal{Z}$

[6] 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.

[7] 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.

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 Appendix B.

3.2 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 $\rho = 1$. The above economy is calibrated 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 $\sigma_E = 0.0768$ 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 roughly 7%, 3 percentage points less than the U.S. sample value.

4. 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.

5. 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 for this class of models. 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. We think of our 'benchmark economy' as that with a risk aversion coefficient of 8, which is close to the value which delivers the sample Sharpe ratio in the Constantinides-Duffie model (see Table 1). We also report moments for a more standard value of 3.

We see several important features of how trade and idiosyncratic risk affect asset pricing. First, our life cycle model generates a price of risk of approximately the same magnitude as the autarkic retirement economy in Table 1. With risk aversion of 8, the implied Sharpe ratio is about 33%.[8] Second, adding counter-cyclical variation to idiosyncratic shocks increases the Sharpe ratio by about 6 percentage points (an increase of about one fourth).

We take several lessons from Table 2. First, while the Constantinides-Duffie model can account for the U.S. Sharpe ratio of 41% with risk aversion of 8, life-cycle effects mitigate our model's ability to do the same. This is due to the existence of retirees who do not face labor market risk. Nevertheless, our model still delivers a Sharpe ratio of 33%, which compares to the complete-market value of 26% and the Mehra-Prescott model's value (with risk aversion of 8) of 25%. Second, the existence of retirees also dampens the effect of CCV. CCV delivers an additional 16 percentage points to the Sharpe ratio in the Constantinides-Duffie model, but only 8 points in our model.[9] Finally, and surprisingly, the implications of our model are essentially the same as the no-trade model, with retirement, from Table 1! Trade, *per se*, is not central for the asset-pricing implications of idiosyncratic risk. By including retirees into the highly-tractable Constantinides-Duffie framework, we get implications which are quite similar to a much richer model, most notably one which motivates trade with a realistic life-cycle distribution of human to financial capital.

4 Explaining the economic forces at work

4.1 The distribution of aggregate risk

To understand our results, it is instructive to begin by considering the allocations of consumption. In what follows, we discuss Figures 1-9. All graphs refer to the case

[8] We focus on the Sharpe ratio because the volatility of the risky return varies across our different economies. As described above, this is because the shock variability, from equation (18), is chosen so that aggregate consumption variability remains constant across economies. In Table 2, for instance, we see that higher CCV is associated with lower volatility in equity returns. This is because, *ceteris paribus*, higher CCV results in higher aggregate consumption variability which, then, necessitates us to reduce the variability of the depreciation shocks.

[9] We have ignored frictions such as portfolio constraints and trading costs. It is quite likely that an important aspect of the overall impact of CCV involves its interaction with such frictions. See, for example, Alvarez and Jermann (2001), Gomes and Michaelides (2004) and Lustig (2004).

with risk aversion of 8. The qualitative features of all these graphs are identical for the low risk aversion economies.

The Constantinides-Duffie economies in Section 2 showed that the implication of incorporating retirement is that that aggregate risk gets concentrated on agents with less labor-income risk. For example, in Section 2.2, the standard deviation of the aggregate part of log consumption growth for the retirees is larger than the workers by a factor of $-b(1+\gamma)/2 = 1.75$ with $\gamma = 8$. A more direct way of measuring exposure to aggregate (equity) risk is to measure how closely individual consumption growth covaries with the excess return. We label this covariance as "aggregate risk bearing". In the Constantinides-Duffie retirement economy aggregate risk bearing is, again, 1.75 times larger for retirees than for workers.

The exact same mechanism is at work in our retirement economy with trade. In Figure 1 we plot the aggregate risk measure across age for three different economies. Consider first the CCV economy. The retirees' bear relatively much aggregate risk. Their consumption growth covaries more with the excess return than does aggregate consumption (0.003 versus 0.0024). The young dislike aggregate risk (due to CCV) and, by trading in stocks and bonds, effectively transfer it to the old. That is, just as in the the Constantinides-Duffie economy, aggregate risk bearing is increasing with age.

This is not the case in the complete-market economy. Marginal rates of substitution are equated across age and, therefore, aggregate risk bearing is constant across age because (due to numerical approximations, the graph is not completely straight). The latter is also true in the "No CCV" economy. The risk-bearing line is flat because, in the absence of CCV, the young evaluate aggregate risk in the same manner as the old, thereby negating any concentration of aggregate risk. The result of this, shown in Table 2, is that the Sharpe ratio under complete markets is roughly the same as under no-CCV. Again, this echos the Constantinides-Duffie model in that asset risk premia are unaffected by idiosyncratic risk that is independent of aggregate risk.

The overall point of Figure 1 is that CCV inhibits intergenerational risk sharing. To emphasize this, consider the following ratio, plotted in Figure 2: the covariance between individual consumption growth and aggregate consumption growth, divided by the variance of aggregate consumption growth. The ratio is roughly unity in the complete-market and no-CCV economies, but is increasing until retirement in the CCV economy (aggregate consumption growth is highly correlated with the excess return in our economy).

But wait a minute! The middle-aged workers (in the CCV economy) are also exposed to labor income risk. Why do they take on more equity risk than the young workers? The reason is life-cycle savings. Aging implies financial wealth accumulation which reduces overall exposure to labor income risk. Figure 3 plots the average age profile of financial wealth in the benchmark CCV economy and the complete-market

economy. As agents get closer to retirement, they accumulate life-cycle wealth because their human capital (*i.e.*, the expected discounted value of future labor earnings) becomes smaller. In addition, we see that, because average earnings peaks at age 50, the young borrow against future earnings and have (on average) negative financial wealth until age 28 in the CCV economy and age 38 in the complete-market economy. This will be instrumental for the next topic, portfolio choice.[10]

4.2 Portfolio behavior

Now that we understand the distribution of aggregate risk, we can explain the portfolio allocations. Figures 4 graphs the levels and Figure 5 graphs the (age-specific average) share of stocks as a fraction of financial wealth. The portfolio choices are displayed for both the complete-market and the CCV incomplete-market economies.

We start with the complete-markets economy. As discussed above, the distinguishing feature of the complete markets allocation is that consumption growth rates across agents are equated with the rate of aggregate consumption growth. However, the portfolio policies which support this allocation differ because agents of different ages have different amounts of human capital and financial wealth. Consider first the retirees. Their policies are to 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 (10.9% versus 3.3%). Retirees have zero labor income, so, in order to replicate aggregate consumption risk, they hold lots of their wealth in bonds.

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.

[10]Due to the precautionary motive, the young save relatively more in the economy with idiosyncratic shocks than in the complete markets economy. Note, however, that overall wealth (the area under the graph) is equated across the two economies. This is a feature of our calibration approach. See e.g. Storesletten, Telmer, and Yaron (2004a) and Gourinchas and Parker (2002) for more extensive analysis of the precautionary savings motive within a life-cycle models.

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 are now in a position to describe behavior in the 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-market 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.

4.3 Sensitivity analysis

- 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.

- The age portfolio profiles in Figure 5 do not change much for other parts of the wealth distribution, or change much over the business cycle. As all the figures thus far were for the *average* investor, Figure 6 documents portfolio holdings of the top 10% of the wealth distribution. The portfolio holdings in this figure show that the portfolio choice patterns underlying the model's equity premium

[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).

are relatively constant across the wealth distribution. Figures 7 and 8 provide the age profile of portfolio shares as a function of the aggregate state for both the complete market and the CCV economy respectively. The upshot is that these portfolio profiles do not change in any significant manner across recessions and expansions, although as expected there is a slight increase in equity investment during expansions.

4.4 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).

Self-insurance behavior in 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 at least 35 percent (see Deaton and Paxson (1994), Storesletten, Telmer, and Yaron (2004a), or Heathcote, Storesletten, and Violante (2005)), so our model exhibits too little risk sharing. Figure 9 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. As shown, autarky implies that, for workers, the graph is flat at 0.17. 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 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.

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.

den Haan, W. and A. Marcet, (1994), Accuracy in simulations, *Review of Economic Studies* 61, 3–17.

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.

Gourinchas, P. O. and J. Parker, (2002), Consumption over the life cycle, *Econometrica* 70 (1), 47–89.

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.

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

Krusell, P. and A. A. Smith, (1998), Income and wealth heterogeneity in the macroeconomy, *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.

Lustig, H., (2004), The market price of aggregate risk and the wealth distribution, Working Paper, Stanford University.

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, 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, (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 3 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),

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

and

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

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:

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

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 3, 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 3, 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 $b = -0.1652$.

Calibration of Models with Trade

The models in Section 3 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 $Z \in \{0.9725, 1.0275\}$. 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 $\delta = 0.164$. The average wage rate, w, is set equal to $(1-\theta)\,E(K)^{\theta}N^{-\theta}$. The parameter s is chosen so that the standard deviation of the risky return, r_t, is 10 percent. The secular growth rate, by which all quantities are normalized, is 1.5 percent.

In Section 3.2, the persistent component of hours worked, z_{it}^{h}, follows a unit root process with innovations governed by a four-state Markov chain, two states corresponding to an expansion and the other two a contraction. The conditional variances, σ_E^2 and σ_C^2, are set to 0.0095 and 0.0467 respectively, which are taken from Storesletten, Telmer, and Yaron (2004b), Table 1, and then scaled down so that the unconditional variance matches that of the $\rho = 0.92$ process (a value of 0.0281). In the case of mean-reverting shocks, we implement their directly, by imposing the following restrictions on the approximated 4-state discrete Markov process: (i) the conditional variance in recessions and booms should be 0.181 and 0.037, (ii) the autocorrelation should be 0.92, (iii) the transition matrix Γ should be symmetric. It is infeasible to match the autocorrelation exactly, so the Markov chain we use has an autocorrelation of only 0.890. The elements of the process z are $z \in \{-2.385, 0.646\} \equiv \{z_{rl}, z_{rh}\}$ in recessions and $z \in \{-0.904, 0.467\} \equiv \{z_{bl}, z_{bh}\}$ in booms. The transition matrix is

	(z'_{rl})	(z'_{rh})	(z'_{bl})	(z'_{bh})
(z'_{rl})	.784	.016	.196	.004
(z'_{rh})	.016	.784	.004	.196
(z'_{bl})	.196	.004	.784	.016
(z'_{bh})	.004	.196	.016	.784

The parameters κ_h are chosen so as to match the PSID mean age profile in earnings.

The demographic structure is calibrated to correspond to several simple properties of the U.S. work force. 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.

Risk aversion is set to match the Sharpe ratio (see Table 1), and the discount factor, β, is chosen so that the average capital to output ratio generates aggregate consumption variability of 3.7 percent. This results in $\beta = 0.69$ and an associated (average) capital to output ratio of 1.8.

The following Table illustrates the aggregate properties of our economy. The sample size for the U.S. data is chosen to be the same as that used by Mehra and Prescott (1985)

Table B1: Aggregate Moments

Panel A: Population Moments of Growth Rates, Theoretical Economy

	Std Dev	Autocorrelation	Correlation with Output
Output	0.930	-0.307	1.000
Investment	1.214	-0.309	0.998
Consumption	0.038	-0.117	0.935

Panel B: Sample Moments of Growth Rates, U.S. Economy, 1929-1982

	Std Dev	Autocorrelation	Correlation with Output
Output	0.062	0.561	1.000
Investment	0.358	0.225	0.143
Consumption	0.036	0.353	0.471

U.S. sample moments are based on annual NIPA data, 1929-1982. Theoretical moments are computed as sample averages of a long simulated time series.

As is discussed in the text, the production side of our economy is unrealistic. Aggregate consumption variability, however, matches the data, as does the variability of the risky asset return (discussed above). Our model does not resolve the well-known problems with production-based asset pricing models. It is best viewed in the same way one views any endowment economy: a model with realistic properties of aggregate consumption which can be supported by some, potentially unrealistic, production technology. Alternatively, one can view our economy as featuring a linear technology — commonplace in the finance literature — where we are explicit about the implications of from the production side of the model — not very commonplace in the finance literature.

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

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

   where $A(Z, Z')$ is an $(m+1) \times (m+1)$ matrix (conditional on Z and Z'), and the entries of μ'_m on ξ are zeros (the first m rows of the $m+1$ column). The aggregate shock Z can take on two values, $Z \in \{\underline{Z}, \bar{Z}\}$, so each element in the 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):

$$\begin{aligned} \hat{V}_h(\xi, \mu_m, Z, z, \epsilon, a) \quad = \quad & \max_{b'_{h+1}, k'_{h+1}} \{u(c_h) + \\ & \beta E \left[\hat{V}'_{h+1}(\hat{G}(\mu, Z, Z'), Z', z', \epsilon', k'_{h+1} \cdot R(\hat{G}(\mu, Z, Z'), Z') + b'_{h+1}) \right] \} \end{aligned} \tag{27}$$

[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.

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 500 agents in each age cohort). Using the decision rules obtained in (27), simulate a long sequence of the economy (5100 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 $\int b'_h(\mu_m, \xi^*, Z, z, \epsilon, a)d\mu = 0$.

5. Update $\hat{G}$ by running a linear regression of (ξ', μ'_m) on μ_m 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 C1).

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 C1 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 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.

[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.

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 11 points for aggregate capital, 11 points for the conditional expected equity premium, and 50 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 (and imposing the borrowing constraints and Kuhn-Tucker conditions):

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

   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.

Accuracy of approximation

The solutions on the grid points are exact by construction. To evaluate whether the interpolation between grid points gives rise to systematic Euler equation pricing errors we follow den Haan and Marcet (1994) and use simulation to construct the following moment conditions:

$$g(c, Z, z, R, q) \equiv \frac{1}{T}\sum_{t=1}^{T}\sum_{h=1}^{H}\frac{1}{N_h^*}\sum_{i=1}^{N_h^*}\left[\left(\beta(\frac{c_{i,t+1}^{h+1}}{c_{i,t}^{h}})^{-\alpha}\left[R_{t+1}\ \ \frac{1}{q_t}\right]-1\right)\otimes z_{i,t}^{h}\right] \tag{29}$$

where T is the number of periods in the simulation, N_h^* is the number of *unconstrained* agents within age cohort h, and the instruments are $z_{i,t}^h = \{1, a_{i,t}^h, R_t, R_{f,t-1}\}$. We

calculate the p-value corresponding to the χ^2_4 statistic based on moment conditions $g(\cdot)$ and its covariance matrix. As in den Haan and Marcet (1994) we repeat this for $S = 300$ independent simulations and a sample size of 10000. The percent of p-values that were above (below) the upper (lower) 5% critical value were 5.4% and (4.1%) respectively. By this formal metric the Euler equation errors are quite accurate.

Table C1
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

$X_j \equiv E[(w_{it} - E_t w_{it})^j]$, w_{it} denotes the wealth of agent i at time t, $X_j^* \equiv E[(w_{it}^* - w_{it}^*)^j]$, w_{it}^* denotes the wealth of *working* agent i at time t, $B_j \equiv E[(b_{it} - b_{it})^j]$, 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 $\xi' = a(Z, Z') + X(Z, Z') * B(Z, Z')$

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 labelled 'match SR' and 'match EP,' risk aversion is chosen to match the U.S. Sharpe ratio and the mean equity premium, respectively. Rows labelled 'Retirement' hold risk aversion at these levels and then incorporate retirement, defined as old agents not receiving any idiosyncratic shocks (Section 2.2).

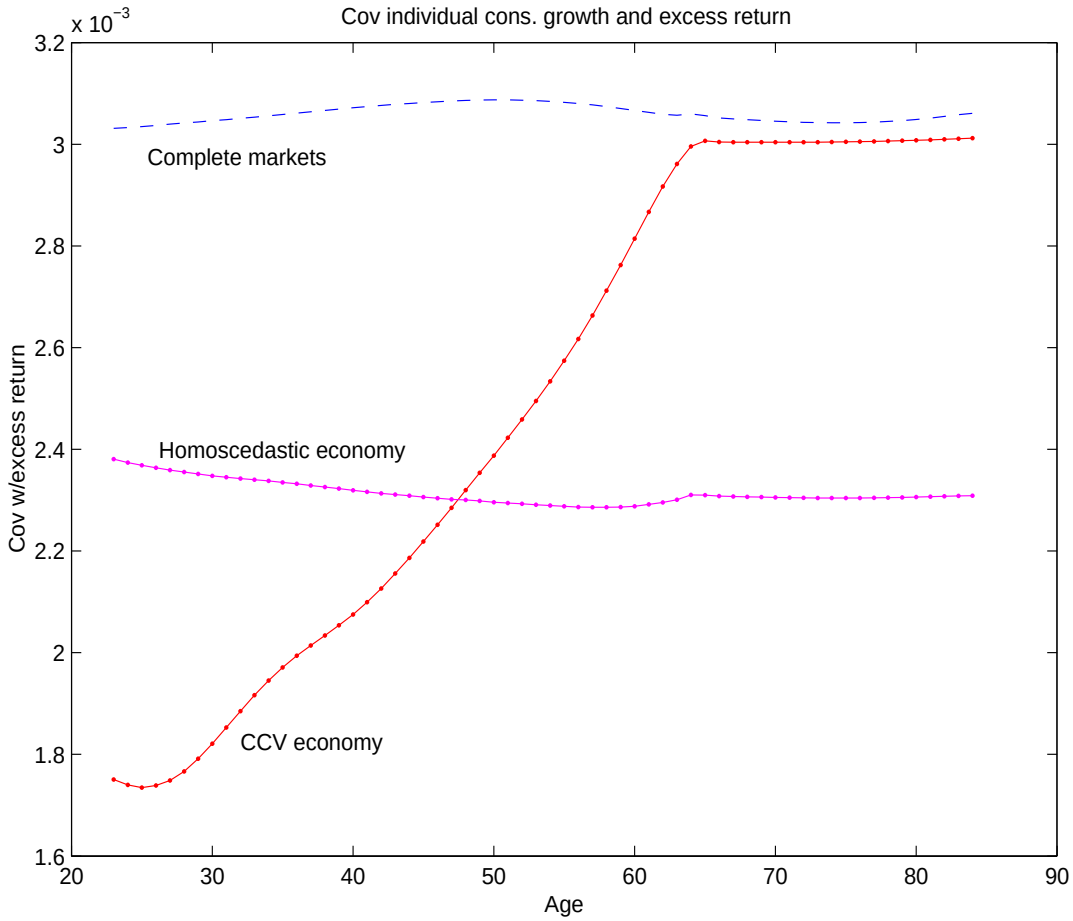
U.S. sample moments are computed using non-overlapping annual returns, (end of) January-over-January, 1956-1996. Estimates of means and standard deviations are qualitatively similar using annual data beginning from 1927, or a monthly series of overlapping annual returns. Equity data correspond to the 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	0.63	7.0	9.0
no CCV	3	0.947	3.3	0.0114	0.0114	2.6	0.80	7.0	11.5
estimated CCV	3	0.948	3.3	0.0168	0.0059	2.3	0.86	6.7	12.8
Complete markets	8	0.96	3.3	0	0	4.7	2.48	9.3	26.8
no CCV	8	0.809	3.3	0.0114	0.0114	2.6	19.8	7.6	26.1
estimated CCV	8	0.801	3.3	0.0168	0.0059	1.3	2.32	7.1	32.6
large CCV	8	0.797	3.3	0.0204	0.0023	1.6	2.51	6.6	38.0

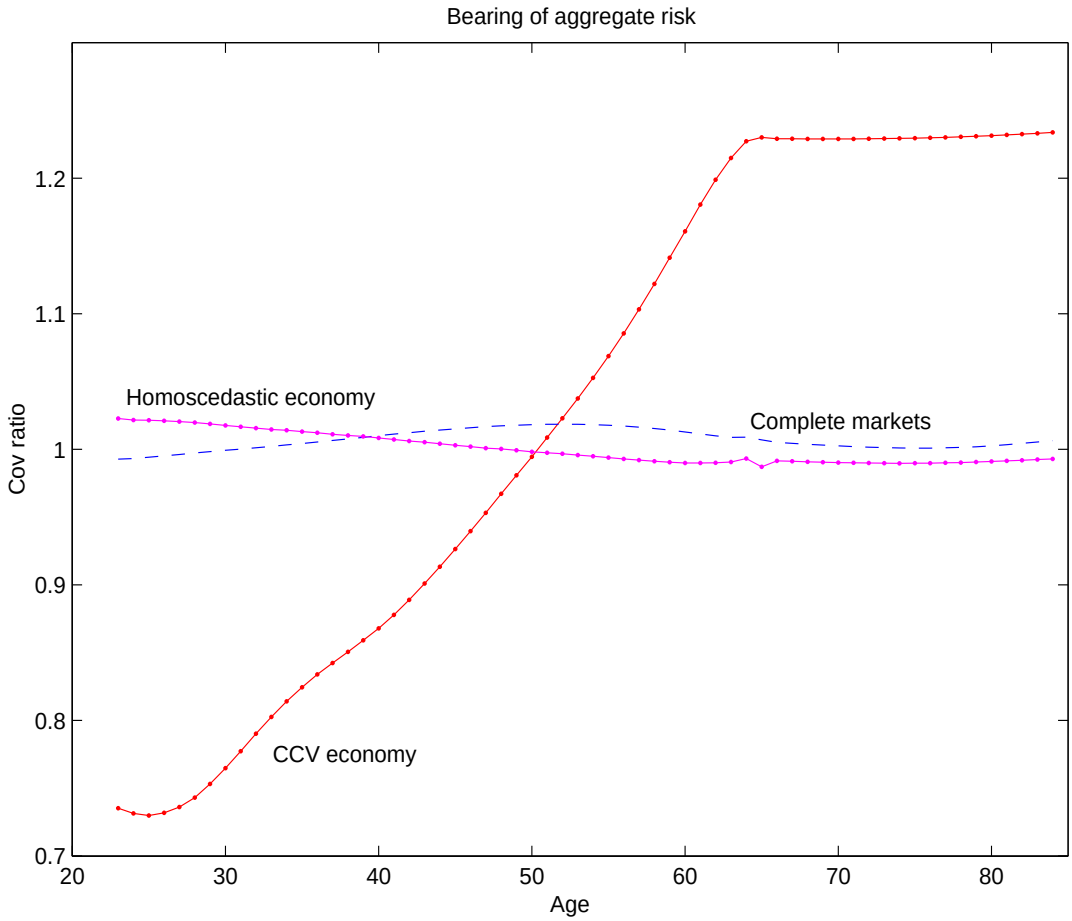
'Models with Trade,' are described in Section 3. The calibration procedure is discussed in the text. All economies are calibrated so that aggregate consumption volatility is 3.3%. The 'Homoskedastic Economy' is distinguished by the volatility of idiosyncratic shocks *not* varying with aggregate shocks. 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
Aggregate Risk Bearing by Age



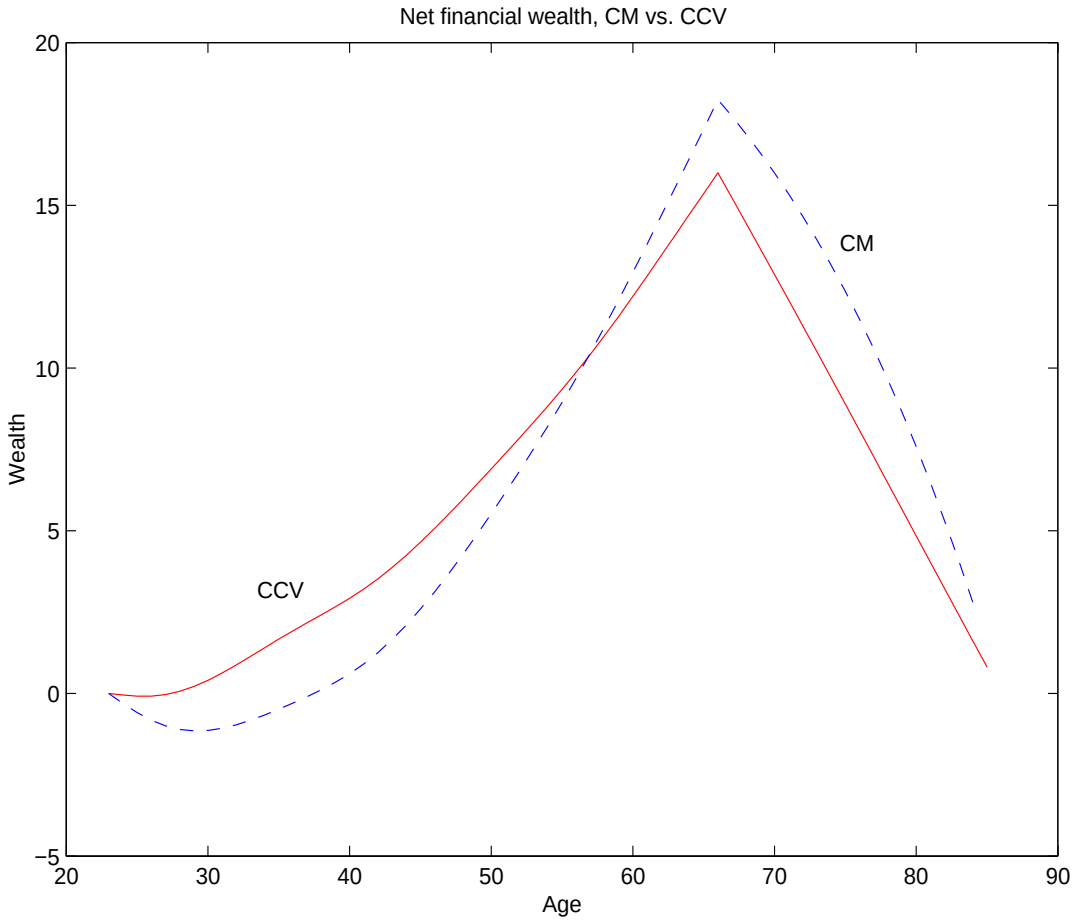
The graph plots the measure of "aggregate risk bearing" which is described in Section 4.1 of the text: the covariance between individual consumption growth and the excess return on equity.

Figure 2
Aggregate Risk Bearing Ratio



The graph shows the age-profile of the aggregate risk bearing ratio described in the text. The numerator is the covariance between individual consumption growth and aggregate consumption growth. The denominator is the variance of aggregate consumption growth.

Figure 3
Financial Wealth by Age



The graph shows the age profile of average financial wealth for two different economies: complete markets (dashed line), and the benchmark CCV economy (solid line).

Figure 4
Quantity of Bonds and Stocks, by Age

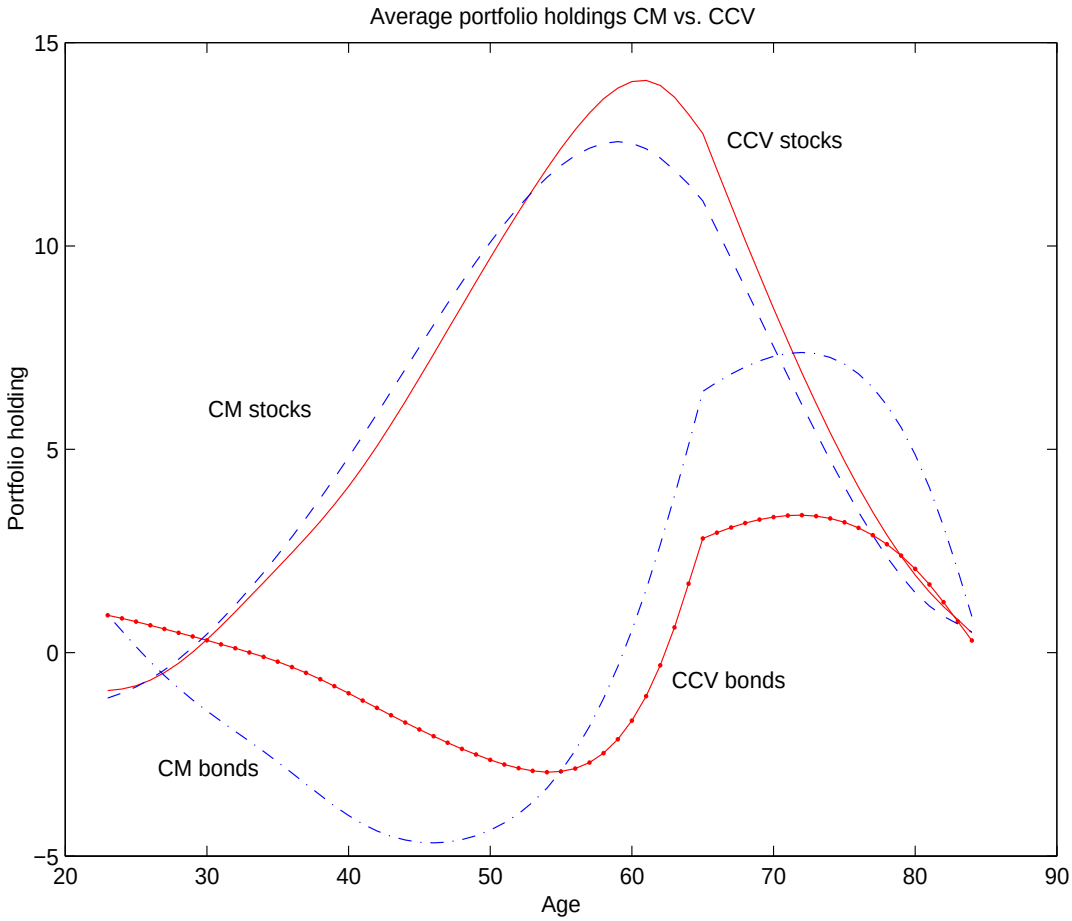


Figure 5
Bond and Stock Portfolio Shares, by Age

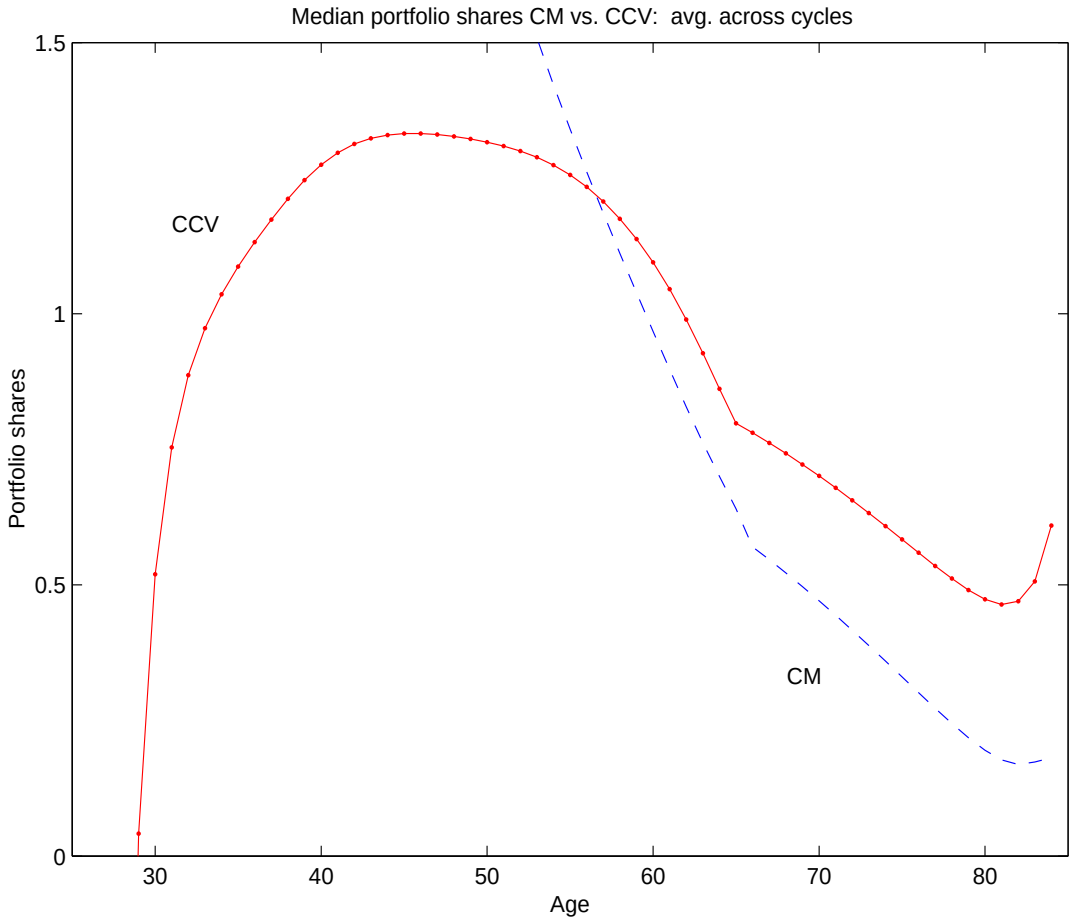


Figure 6
Portfolio Shares for Top 10 Earnings Percentile, by Age

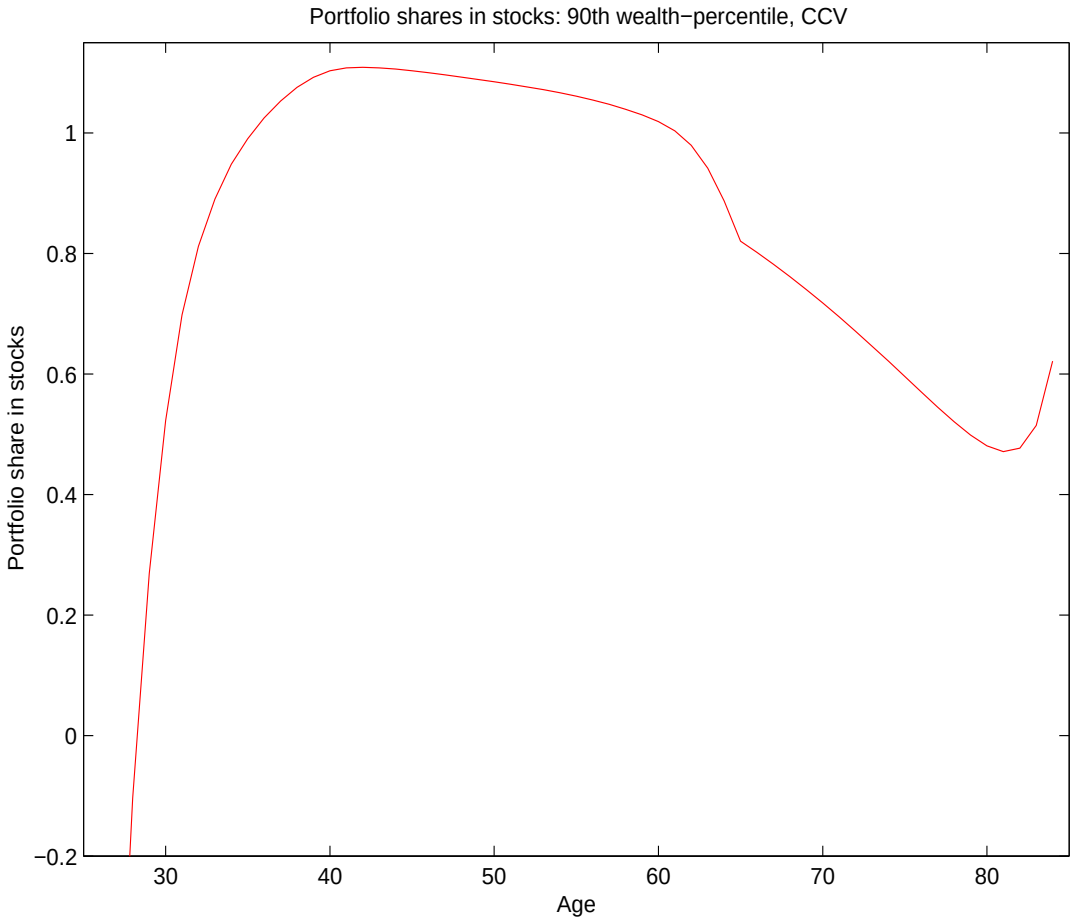
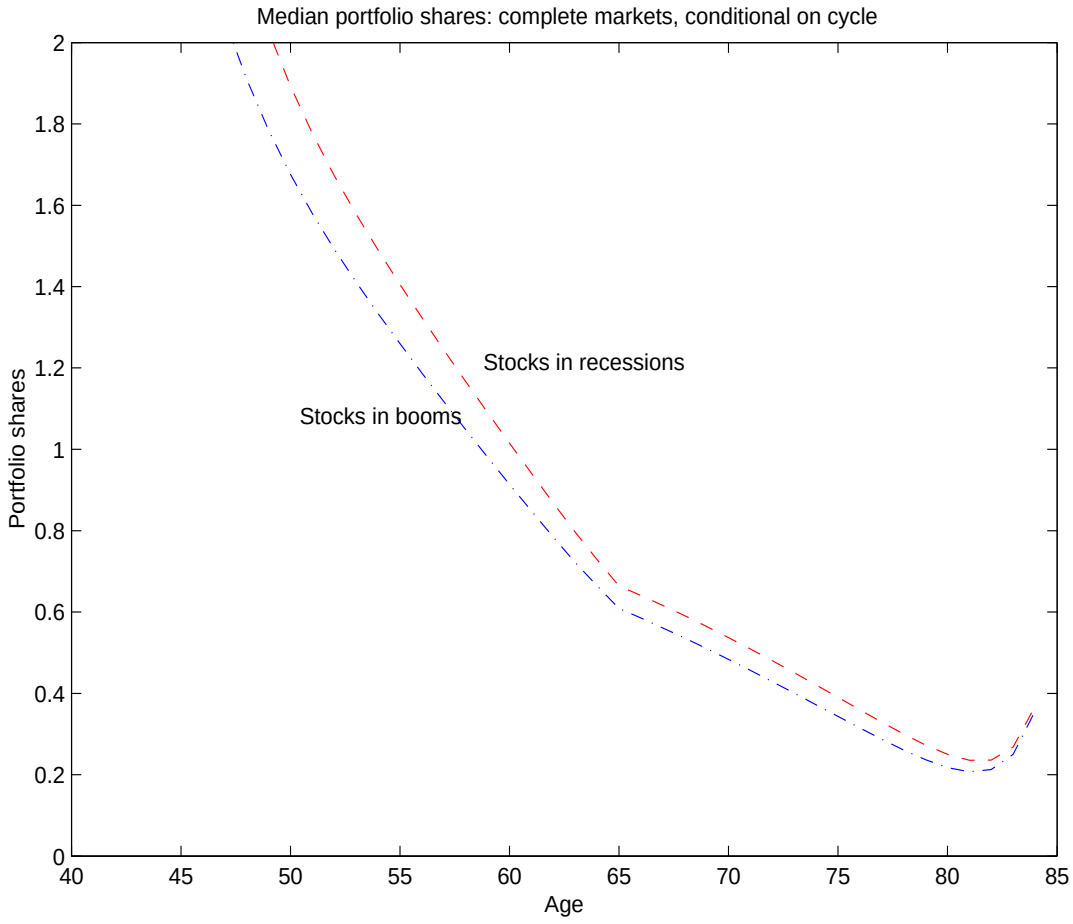
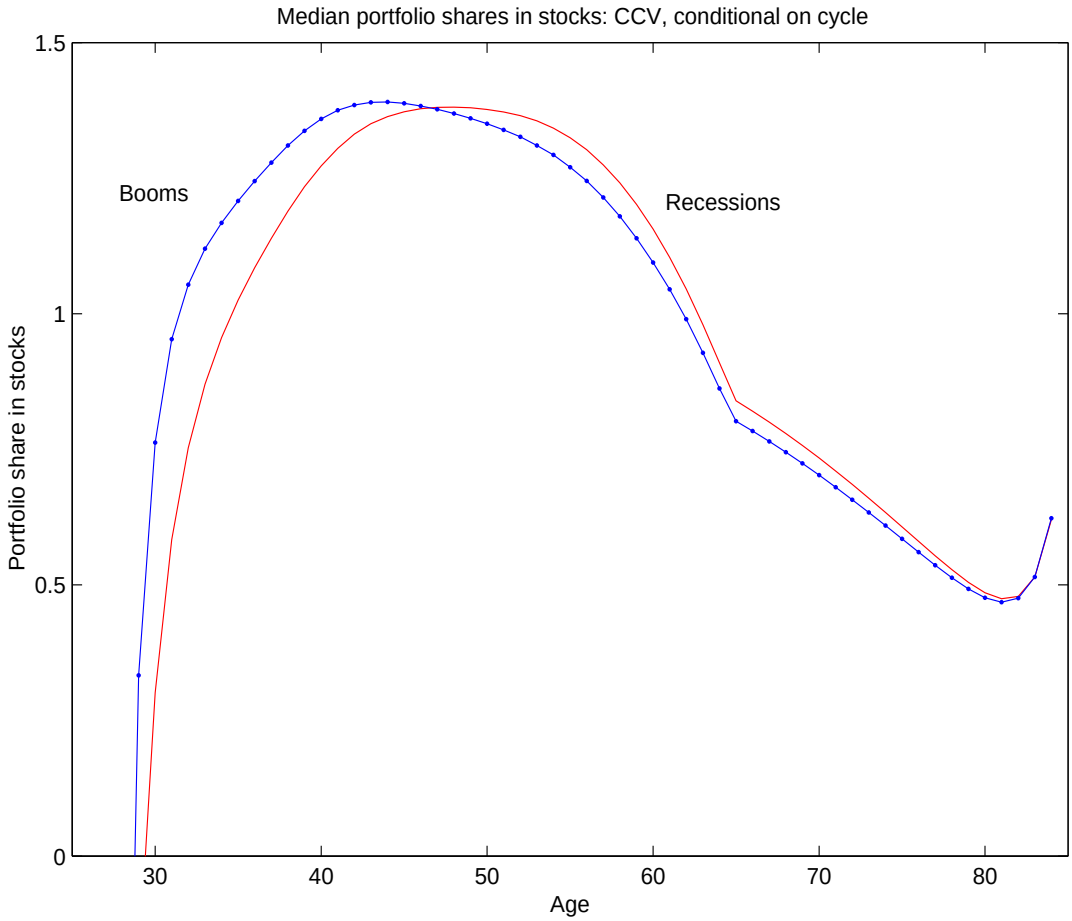


Figure 7
Portfolio Shares conditional on Business Cycle, by Age



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

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
Variance of Log Consumption

