Life Cycle Portfolio Choice and Asset Market Effects of the Baby Boom

Robin Brooks*

September 2000

Abstract

This paper explores life cycle portfolio choice in a Diamond – style neoclassical growth model with overlapping generations, in which agents make a portfolio decision over risky equity and safe bonds. It generates portfolio behavior whereby optimizing agents shift from stocks to bonds as they age. They behave this way because of implicit holdings of a non-traded asset, human capital, which they run down over the life cycle. The paper finds that this portfolio shift is robust to a range of specifications. It focuses especially on the effects of pay-as-you-go social security, which effectively flattens the life cycle profile of human capital by transferring wage income from workers to retirees. The paper finds that the portfolio shift from stock to bonds is unchanged for realistic levels of the payroll tax. The paper uses the model to simulate the general equilibrium effects of a baby boom and baby bust on asset returns. It finds evidence of significant asset market effects, with the expected return on retirement savings of boomer cohorts up to 20 percent below returns to earlier generations. Though social security could in principle offset these asset market effects – effectively providing insurance against the risk of being born into a large cohort – the paper finds that it fails to do so for realistic levels of the payroll tax.

JEL Classification Numbers: E27, G11, G12, H55

Keywords: equity premium, population aging, pension reform

---

* International Monetary Fund, Research Department, 700 19th Street, N.W., Washington, D.C. 20431; Tel: (202) 623-6236; Fax: (202) 623-6334. Email: rbrooks2@imf.org. Special thanks are due to Christopher Sims. I am grateful to Andrew Abel, Orazio Attanasio, Hamid Faruqee, Peter Isard, Ivailo Izvorski, Michel Juillard, Narayana Kocherlakota, Eswar Prasad and Ed Prescott as well as seminar participants at the 2000 Midwestern Macro Conference, the 2000 Society of Economic Dynamics meetings and the 2000 World Congress of the Econometrics Society for helpful comments and suggestions. Any errors are mine. An earlier version of this paper was published as IMF Working Paper WP/00/18 in February 2000 and can be downloaded from http://www.imf.org.

I. Introduction

This paper explores life cycle portfolio choice in a Diamond (1965) – style neoclassical growth model with overlapping generations, in which agents make a portfolio decision over risky equity and safe bonds. The model features two sources of aggregate uncertainty, a technology shock to production and random population growth, and generates portfolio behavior whereby optimizing agents shift from stocks to bonds as they age. This portfolio shift is driven by implicit holdings of a non-traded asset, human capital, which is run down over the life cycle. Agents see this asset, which can be thought of as generating labor income, as a substitute for the riskless asset even though labor income is risky. This is because the correlation between labor income and equity returns is close to zero so that, like the bond, human capital provides an income stream that mitigates the effects of a low stock return on consumption. Agents therefore shift from stocks to bonds to compensate for the decline in human capital over the life cycle, in effect maintaining a constant share of total wealth in the risky asset. The paper goes beyond the existing literature by using a general equilibrium model to explore the determinants of life cycle portfolio choice. It shows that a broad range of specifications generate a correlation between labor income and equity returns that is near zero – a result that matches estimates from panel data and confirms the simulation results of Jagannathan and Kocherlakota (1996), Cocco, Gomes and Maenhout (1999) and Heaton and Lucas (2000) who use partial equilibrium models without production. The paper also explores whether the implications of the model for portfolio behavior are robust to the inclusion of pay-as-you-go social security, which effectively flattens the life cycle profile of human capital by transferring wage income from workers to retirees. It finds that portfolio behavior is basically unchanged for realistic levels of the payroll tax.

The paper goes on to simulate the general equilibrium effects of a baby boom and baby bust on asset returns. It draws inspiration from several academic studies and articles in the financial press that speculate about what will happen to financial markets when the Baby Boomers retire. Broadly speaking there are two opposing views in this literature. One holds that the Baby Boomers will be selling their assets to a smaller generation of young investors when they retire. At this point excess supply of financial assets will lead to a decline – perhaps a crash – in the real value of financial assets. The opposed view maintains that forward-looking financial markets are pricing assets to incorporate the aging of the Baby Boomers so that a market meltdown is unlikely. The simulated baby boom and baby bust generates significant asset market effects, with the expected return on retirement savings of baby boomer cohorts up to 20 percent below returns to earlier generations. Though a defined-benefit social security system could in principle offset these asset market effects – effectively providing insurance against the risk of being born into a large cohort – the paper finds that it fails to do so for realistic levels of the payroll tax.

Since the model has only aggregate sources of uncertainty, the focus of the paper is on the asset pricing implications of heterogeneity across rather than within generations. Using the Sharpe ratio of the equity premium (the average of the excess return scaled by its standard deviation) to control for the riskiness of the model environment, the paper finds that this measure grows in excess of 15 percent as technology shocks become more persistent over time. This result compares favorably with the Mehra and Prescott (1985) benchmark of 37 percent and illustrates that the model economy represents a significant departure from complete markets. Agents are unable to diversify away a significant amount of generation-specific risk because they can trade with only a limited number of other generations in the overlapping generations setting. The model thus underscores the asset pricing implications of aggregate risk at low frequencies.

The paper is organized as follows. Section II introduces the overlapping generations model, while Section III discusses its parameterization. Section IV solves the model numerically and characterizes life cycle portfolio behavior. Section V simulates a baby boom and baby bust and Section VI concludes.

II. The Model Economy

This section describes the Diamond – style overlapping generations model in which agents live for four periods. In childhood agents depend on their parent for consumption. In young working-age agents supply labor inelastically. Out of after-tax wage income they consume for themselves and their children and make a portfolio decision over risky capital and safe bonds, the two assets in the model economy. In old working-age agents again supply labor inelastically. They earn returns on financial assets in addition to the after-tax wage and consume only for themselves, their children having entered young working-age and left the household. They also decide on what mix of financial assets to hold going into retirement. In retirement agents no longer supply labor and consume down their savings, there being no bequests. They may also receive a retirement benefit, which is financed out of taxes levied on the workforce.

The model has four periods because of two considerations. First, the childhood and retirement periods allow the model to capture the general equilibrium effects of changes in youth and old-age dependency. Second, the two working-age periods allow for a riskfree one-period bond in zero net supply, so that agents make a portfolio allocation decision in addition to the standard consumption-saving choice. If the risk premium on capital has a low correlation with the after-tax wage, young workers will short the bond and invest in equity because this diversifies some of the risk to next period income. Old workers will be willing lenders because this provides a safe income stream in retirement, when they no longer work and consume out of savings. The paper uses the general equilibrium setting of the model to explore what factors determine portfolio allocation over the life cycle. It shows that for a broad range of parameterizations agents shift the allocation of their financial wealth from stocks to bonds as they approach retirement – portfolio behavior for which there is mounting empirical evidence.[1]

Agents optimally change their holdings of stocks and bonds over the life cycle because they implicitly hold a non-traded asset, human capital, which they run down and convert into financial wealth as they age. A simple example, in which labor income is riskfree, will help to illustrate this point. Assume there are two assets, a riskfree bond and risky capital, while total wealth consists of financial wealth plus human capital, the present value of future labor income. With perfect markets for borrowing and lending agents shift from holding stocks when young to bonds when old, even selling short the riskless asset in youth when the value of certain future labor income is large. The goal of this shift is to maintain a constant share of total wealth in risky capital over the life cycle. Cocco, Gomes and Maenhout (1999) and Heaton and Lucas (1997) extend this intuition to risky labor income. They show that the optimal share of financial wealth in equities is decreasing with age, provided that the correlation between shocks to labor income and stock market returns is low, because agents effectively view human capital as a substitute for the riskfree asset even though it is risky. This paper extends their analysis – they take production as exogenous – by showing that a calibrated general equilibrium model generates a low correlation of labor income with returns on risky capital, so that agents shift from stocks to bonds as they age and run down their human capital.

Although heterogeneity within generations – in the form of uninsurable idiosyncratic income shocks – has been shown to be important in resolving the empirical difficulties associated with representative consumer models, the focus in this paper is on heterogeneity across generations.[2] The model therefore has two sources of aggregate uncertainty: a technology shock to production and stochastic population growth. These two sources of uncertainty are quite different in their implications for asset pricing. Agents observe a technology shock in the period it is realized. In contrast demographic shocks can be thought of as having one period advance warning – children do not become active decision makers until young working-age. The childhood period therefore plays an important role in modeling demographic uncertainty. Once a cohort is born and its size known, the model assumes that all demographic uncertainty associated with it is resolved because progress through the age distribution is perfectly predictable. The model therefore focuses on demographic uncertainty associated with aggregate fertility trends, while ignoring that associated with length of life.[3]

The model abstracts from several important issues. First, it ignores market imperfections such as transactions costs or borrowing constraints, although these may well explain the limited degree of stock market participation by households.[4] Second, the model represents a closed economy and is therefore mute on investors' ability to diversify demographic risk by investing in regions with different population characteristics. However, as the paper demonstrates below, most countries with developed asset markets have very similar demographic characteristics, so that the closed economy can be thought of as representing the developed world as a whole.[5] Third, the model abstracts from bequests.[6] In principle such intergenerational transfers should have important implications for portfolio allocation over the life cycle, although empirical studies suggest they are of minor importance to most households.[7] Finally, the model ignores the effects of an important asset, housing, on holdings of financial assets over the life cycle. Obviously it is of interest to extend the model to address these issues – in order to explore the sensitivity of its asset market implications – but this is beyond the scope of the paper.

The evolution of the age distribution over time is depicted in Table 1. The age distribution in period t consist of N_{t-1} young workers, N_{t-2} old workers, and N_{t-3} retirees. The period t cohort of children is determined by $N_t=(1+n_t)N_{t-1}$ where n_t is the realization of a stationary cohort growth shock. Cohort growth differs from population growth in that it relates the size of the youngest generation to that of its parent, rather than to the rest the population.

Table 1: The Age Distribution over Time

Period	Children	Young Workers	Old Workers	Retirees
t	N_t	N_{t-1}	N_{t-2}	N_{t-3}
$t+1$	N_{t+1}	N_t	N_{t-1}	N_{t-2}
$t+2$	N_{t+2}	N_{t+1}	N_t	N_{t-1}
⇓	⇓	⇓	⇓	⇓

The representative young worker in period t maximizes expected lifetime utility (1) subject to the budget constraints in young working-age (2), old working-age (3) and retirement (4).

$$V_t = \lambda(1+n_t)\frac{(c_t^0)^{1-\theta}}{1-\theta} + \frac{(c_t^1)^{1-\theta}}{1-\theta} + \beta E_t\left[\frac{(c_{t+1}^2)^{1-\theta}}{1-\theta}\right] + \beta^2 E_t\left[\frac{(c_{t+2}^3)^{1-\theta}}{1-\theta}\right] \quad (1)$$

$$(1+n_t)c_t^0 + c_t^1 + s_{et}^1 + s_{bt}^1 = w_t(1-\eta_t) \quad (2)$$

$$c_{t+1}^2 + s_{et+1}^2 + s_{bt+1}^2 = w_{t+1}(1-\eta_{t+1}) + (1+r_{et+1})s_{et}^1 + (1+r_{ft})s_{bt}^1 \quad (3)$$

$$c_{t+2}^3 = (1+r_{et+2})s_{et+1}^1 + (1+r_{ft+1})s_{bt+1}^1 + bw_{t+2} \quad (4)$$

Young workers care about their children's consumption c_t^0 with a discount factor λ, but do not have an altruistic bequest motive. β reflects their subjective rate of time preference while θ is their coefficient of relative risk aversion. Young workers determine their own consumption c_t^1, that of their children $(1+n_t)c_t^0$ and decide how many shares s_{et}^1 in risky capital to hold, investing s_{bt}^1 in the safe bond. When they make this portfolio decision the period $t+1$ bond return is known. It is therefore denoted r_{ft} because it is in the period t information set. In contrast the return on capital is not observed until period $t+1$ because it depends on the period $t+1$ technology shock. It is denoted r_{et+1} because it is in the period $t+1$ information set. Young workers' disposable income is $w_t(1-\eta_t)$ where η_t is the payroll tax rate. In old working-age they receive a return on their portfolio of financial assets, selected in young working-age, in addition to the after-tax wage. In retirement they may receive a retirement benefit, which is determined by an exogenous replacement rate b. In the absence of bequests the budget constraint for retirees amounts to a decision rule, so that every period there are just two generations of active decision makers: young and old workers.

Output is generated by a constant-returns-to-scale neoclassical production function.[8] Factor markets are efficient so that capital and labor are rewarded their marginal products:

$$r_{et} = \alpha K_{t-1}^{\alpha-1}(A_t L_t)^{1-\alpha} - \delta \quad (5)$$

$$w_t = (1-\alpha)K_{t-1}^{\alpha} A_t^{1-\alpha} L_t^{-\alpha} \tag{6}$$

δ is the depreciation rate and α determines the share of output rewarded to capital. K_{t-1} is generated by investment decisions of young and old workers in period $t-1$, so that although it enters period t production it is dated as of period $t-1$. A_t is the realization of a stationary technology shock. The labor force L_t consists of young and old workers so that $L_t = N_{t-1} + N_{t-2}$. In equilibrium the capital stock used in period $t+1$ production is determined by share holdings of young and old workers chosen in period t:

$$K_t = N_{t-1} s_{et}^1 + N_{t-2} s_{et}^2 \tag{7}$$

The equilibrium condition for the riskless asset (8) illustrates that the bond is a vehicle for intergenerational lending and borrowing. As a result the bond return loads on the relative size of the two working-age generations, while the return on capital depends on the labor force, the sum of these cohorts.

$$0 = N_{t-1} s_{bt}^1 + N_{t-2} s_{bt}^2 \tag{8}$$

This difference in how the two assets load on the age distribution provides an intuition for why the bond return is more sensitive to population dynamics than the return on capital. To a degree this greater sensitivity is a matter of construction. Increasing the number of periods would raise the number of overlapping generations, so that changes in the ratio of borrowers to lenders would be less pronounced.[9] The zero net supply restriction also brushes over the fact that in reality the supply of bonds is endogenous. In principle an infinitely-lived agent such as government could manage the supply of bonds to offset any effects on the riskfree rate – effectively using government debt to insure agents against the asset market effects of demographic shocks. But this is exactly the point the zero net supply restriction helps to make, by focusing on the asset market effects of population dynamics in the absence of active government.

The model abstracts from government activity with the exception of an unfunded pay-as-you-go social security system.[10] Given a replacement rate b, government collects payroll taxes on wages w_t at a rate η_t from all workers:

$$\eta_t = \frac{bN_{t-3}}{N_{t-1} + N_{t-2}} \tag{9}$$

$N_{t-3}/(N_{t-1} + N_{t-2})$ is the retiree to worker ratio so that the payroll tax rate moves with this ratio to balance social security over a demographic shift. This version of social security provides retirees with a defined benefit in retirement, granting them access to wage income at the replacement rate, and is relatively more costly for small generations. Effectively this version of social security provides insurance against the risk of being born into a large cohort, reflecting Bohn (1999) who uses a Diamond-style overlapping generations model to argue that smaller generations are better off than larger ones because they supply labor when the capital-labor ratio is high and save when it is low.[11] A key question for this paper is then whether the asset market implications of the model are sensitive to the inclusion of a social security system with realistic replacement rates.

Imposing these equilibrium conditions and aggregating across period t budget constraints yields the social resource constraint. At the aggregate level the model economy corresponds to a simple one-sector business cycle model.

$$C_t + K_t - (1-\delta)K_{t-1} = K_{t-1}^{\alpha}(A_t L_t)^{1-\alpha} \tag{10}$$

Both sources of aggregate uncertainty are log-normally distributed such that

$$\ln A_t = \phi \ln A_{t-1} + \varepsilon_t \tag{11}$$

$$\ln N_t = \rho \ln N_{t-1} + v_t \tag{12}$$

where $\varepsilon_t \sim N(0,\sigma_\varepsilon)$ and $v_t \sim N(0,\sigma_v)$. For $\phi > 0$ the model economy has 6 state variables. Two of these are endogenous and represent the distribution of wealth across working-age cohorts. Three describe cohort size (N_t, N_{t-1}, and N_{t-2}) and another the level of technology (A_t).

III. Model Calibration

The parameterization of the model reflects the fact that each period represents 20 years. The subjective discount rate β is set at 0.6, which corresponds to a pure rate of time preference φ of about 2.5 percent per year where $\beta=1/(1+\varphi)^{20}$. The discount factor applied to the utility of children λ is set in the range of 0.6 – 1, where 0.6 is equivalent to an annual discount factor of 0.975. The coefficient of relative risk aversion θ is set in a range of 0.5 – 2 to explore the sensitivity of the results with respect to different values for the elasticity of intertemporal substitution. The share of output rewarded to capital α is set in a range 0.3 – 0.4, while depreciation is assumed to occur at a rate of 5 percent per year so that $\delta = 1\text{-}0.95^{20} = 0.65$. When the model includes the pay-as-you-go social security system, the replacement rate b is set at 30 percent so that the steady state payroll tax rate η amounts to 15 percent. Special attention is given to the role of the two sources of uncertainty because – given a period length of 20 years – there is little data to estimate the relevant parameters with any accuracy. Based on pooled data for North America and the European Union countries for the postwar period $\sigma_\varepsilon = 0.1$ and $\sigma_v = 0.01$, a reflection of the fact that more risk is associated with changes in non-demographic fundamentals. ϕ and ρ are each set in ranges of 0 – 0.99 so that – though the model is stationary – A_t and N_t range from being *iid* to approaching a random walk. A key focus of the paper will be how agent behavior changes with persistence of these shocks.[12]

Computational constraints dictate that the model be a fairly stylized version of reality – it omits a key asset held over the life cycle (housing) and abstracts from important market imperfections (borrowing constraints). As a result it is unlikely that generated portfolio behavior will match the data exactly. A more realistic goal for the paper is to use the general equilibrium setting of the model to identify key determinants of portfolio behavior over the life cycle – with a special focus on the role of human capital in the portfolio decision over traded financial assets.

IV. Portfolio Choice over the Life Cycle

The model is solved numerically using the parameterized expectations approach (PEA) of Den Haan and Marcet (1990). The essence of the PEA is that the conditional expectations of agents represent functions $g{:}R_+^{\,m} \rightarrow R_+$ of the state variables. This insight is used to substitute each of the conditional expectations in the first-order conditions of the model with functions $\Pi(\Theta_t,\psi)$ where Θ_t represents the state variables in period t. The functional form Π and the vector of parameters ψ are then chosen to make $\Pi(\Theta_t,\psi)$ as close as possible to g. The paper uses Monte Carlo simulations to implement the PEA, first drawing sequences of length T for the technology shock and the age distribution. The first iteration then picks starting values for the endogenous states as well as ψ and solves out the model for the T periods. Next it fits the conditional expectations by minimizing with respect to ψ the mean squared error between the actual realization in period $t+1$ and the expectation at t of that realization, represented by a polynomial in Θ_t. The following iterations repeat this process until a fixed point in ψ is reached. The accuracy of the solution is tested using a statistic developed by Den Haan and Marcet (1994), which checks for orthogonality between the Euler equation residuals and a vector of variables in the period t information set – using new draws for the technology shock and the age distribution series. The appendix discusses in greater detail how the model is solved using the PEA.

Table 2 characterizes the PEA solution of the model, describing agent behavior and asset returns for a baseline specification where λ=0.6, θ=1.2, α=0.3, ϕ=0, ρ=0.99 and b=0 (for ρ=0.99 cohort growth n_t is serially uncorrelated). It presents the means (μ) and standard deviations (σ) of key variables over the simulation period (T=1000), drawing on the decision rules implicit in the fixed point estimate for ψ to solve out the model for the A_t and N_t series used in the Monte Carlo simulations. It does this for first, second and third-order polynomial approximations to expectations – using the same draws for A_t and N_t each time – so that the focus is on how agent behavior changes as the formation of expectations becomes non-linear in the states.

Table 2: The Baseline Specification – A_t is *iid*

	First-Order		Second-Order		Third-Order	
	μ	σ	μ	σ	μ	σ
c^1	0.1833	0.0141	0.1840	0.0156	0.1836	0.0153
c^2	0.2306	0.0269	0.2308	0.0287	0.2306	0.0264
c^3	0.2871	0.0203	0.2873	0.0260	0.2870	0.0185
s_e^1	0.0922	0.0070	0.1484	0.0191	0.1412	0.0180
s_e^2	0.0589	0.0154	0.0045	0.0055	0.0106	0.0053
s_b^1	-0.0737	0.0064	-0.1300	0.0197	-0.1228	0.0177
s_b^2	0.0737	0.0061	0.1300	0.0196	0.1228	0.0177
r_e	1.1863	0.1846	1.1766	0.2017	1.1849	0.1992
w	0.3216	0.0252	0.3226	0.0263	0.3219	0.0259
r_f	1.1626	0.2042	1.1514	0.2030	1.1682	0.1568
r_e - r_f	0.0237	0.2369	0.0253	0.1931	0.0167	0.1435
K	0.1535	0.0154	0.1554	0.0196	0.1542	0.0188
L	2.0326	0.0892	2.0326	0.0892	2.0326	0.0892
A	0.9990	0.1038	0.9990	0.1038	0.9990	0.1038
Sharpe	0.1001		0.1309		0.1162	
Score	39.7%		7.8%		5.9%	

λ=0.6, β=0.6, θ=1.2, α=0.3, δ=0.65, ϕ=0, σ_ε=0.1, ρ=0.99, σ_v=0.01

Table 2 illustrates that agents shift their financial wealth from stocks to bonds as they age. Young workers hold a portfolio that is short in the riskless asset and long in risky capital, effectively borrowing to invest in equity. Old workers hold very little of the risky asset, investing almost all their retirement savings in the safe bond. This pattern between age and portfolio allocation becomes more pronounced as the order of the polynomial approximation increases. The *Score* statistic indicates that the accuracy of the solution jumps going from the first to the second-order polynomial approximation, while there is little additional gain from adding third-order terms. The non-linearity in the formation of expectations seems to be largely captured by the second-order approximation.

Table 3: The Baseline Specification – A_t is *iid*

	c^1	c^2	c^3	s^1_e	s^2_e	s^1_b	s^2_b	r_e	w	r_f	ep	K	K_{+1}	L	A
c^1	.	0.94	0.25	0.93	0.05	-0.9	0.89	0.30	0.99	-0.4	0.71	0.37	0.92	-0.0	0.92
c^2	.	.	-0.1	0.93	0.04	-0.9	0.94	0.46	0.90	-0.4	0.88	0.17	0.92	-0.0	0.91
c^3	.	.	.	0.04	0.10	0.10	-0.1	-0.5	0.35	-0.2	-0.4	0.69	0.06	-0.0	0.03
s^1_e	.	.	.	.	-0.3	-1.0	0.99	0.38	0.90	-0.4	0.78	0.25	0.91	0.01	0.88
s^2_e	.	.	.	.	.	0.30	-0.3	-0.0	0.05	-0.0	-0.0	0.03	-0.0	-0.1	0.02
s^1_b	.	.	.	.	.	.	-1.0	-0.4	-0.8	0.36	-0.8	-0.2	-0.9	-0.0	-0.9
s^2_b	.	.	.	.	.	.	.	0.44	0.85	-0.4	0.84	0.17	0.90	0.03	0.87
r_e	.	.	.	.	.	.	.	.	0.24	0.54	0.47	-0.7	0.37	0.01	0.65
w	.	.	.	.	.	.	.	.	.	-0.4	0.63	0.42	0.89	-0.0	0.90
r_f	.	.	.	.	.	.	.	.	.	.	-0.5	-0.7	-0.4	-0.0	-0.0
ep	.	.	.	.	.	.	.	.	.	.	.	0.01	0.77	0.02	0.71
K	.	.	.	.	.	.	.	.	.	.	.	.	0.36	0.32	0.01
K_{+1}	.	.	.	.	.	.	.	.	.	.	.	.	.	0.33	0.87
L	.	.	.	.	.	.	.	.	.	.	.	.	.	.	-0.0
A	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.

λ=0.6, β=0.6, θ=1.2, α=0.3, δ=0.65, ϕ=0, σ_ε=0.1, ρ=0.99, σ_v=0.01

What explains the link between age and portfolio allocation? Table 3 presents a matrix of correlations for the second-order approximation. It shows that the correlation of the equity premium (ep) with wage income is 0.63. By shorting the bond and going long in risky capital young workers diversify some of the uncertainty associated with next period income. In contrast old workers have no more wage income going forward. Their goal is to minimize consumption risk in retirement when – in the absence of social security – they depend on savings alone for consumption. As a result they invest almost all their retirement savings in the safe asset. Underlying this portfolio behavior are implicit holdings of human capital, which agents run down over the life cycle and view as a substitute for the riskless asset, even though labor income is risky and non-traded. This feature of the model parallels Cocco, Gomes and Maenhout (1999) who find that agents behave in this way provided that the correlation between wage income and the return on equity is low and positive. The model generates a correlation between the return on capital and wage income of 0.26, which intuition would suggest depends especially on the technology shock being *iid*. Take the case of a positive technology shock that occurs in period *t*. Old workers respond to this temporary shock by smoothing consumption over their remaining lifetime. They raise bond holdings going into retirement – the correlation between s^2_{bt} and A_t is 0.87. In contrast their holdings of risky capital are unresponsive because they want to minimize consumption risk in retirement and the technology shock is uncorrelated over time. Their greater demand for the riskless asset is matched by greater supply – young workers being eager to diversify income risk in old working-age by holding more equity. This raises the supply of capital going into period $t+1$ – the correlation between the period $t+1$ capital stock (K_{+1}) and A_t is 0.87 – which moves factor returns in opposite directions in period $t+1$. Because this second-period effect of technology shocks is almost large enough to offset the positive first-period effect on the correlation of factor returns, the unconditional correlation – an average over first and second-period effects – is low and positive.[13]

Table 4: Highly Persistent Technology Shocks – $\ln(A_t) = 0.8\ln(A_{t-1}) + \varepsilon_t$

	First-Order		Second-Order		Third-Order	
	μ	σ	μ	σ	μ	σ
c^1	0.1860	0.0254	0.1862	0.0315		
c^2	0.2377	0.0671	0.2332	0.0424		
c^3	0.2885	0.0576	0.2907	0.0516		
s_e^1	0.1237	0.0173	0.1402	0.0308		
s_e^2	0.0335	0.0189	0.0144	0.0089		
s_b^1	-0.1037	0.0151	-0.1218	0.0294		
s_b^2	0.1037	0.0151	0.1218	0.0294		
r_e	1.1451	0.2658	1.1729	0.1485		
w	0.3276	0.0448	0.3262	0.0545		
r_f	1.0195	0.4772	1.1431	0.1619		
r_e - r_f	0.1257	0.6405	0.0298	0.1977		
K	0.1598	0.0191	0.1569	0.0306		
L	2.0326	0.0892	2.0326	0.0892		
A	1.0090	0.1709	1.0090	0.1709		
Sharpe	0.1962		0.1505			
Score	31.4%		6.2%			

λ=0.6, β=0.6, θ=1.2, α=0.3, δ=0.65, ϕ=0.8, σ_ε=0.1, ρ=0.99, σ_v=0.01

This intuition suggests that making the technology shock more persistent will weaken the second period effect – there being a lesser need for agents to smooth consumption – so that the unconditional correlation of the return on capital and wage income should rise. Table 4 explores this reasoning, presenting simulation results for a specification that corresponds in all respects to the baseline except for the technology shock, which follows a more persistent process ($\phi = 0.8$). It describes model behavior for the first and second-order approximations – the third-order simulation having not yet solved. The pattern between age and portfolio allocation is robust to greater persistence in the technology shock – agents continue to shift their financial wealth from stocks to bonds as they age. If anything Table 5 shows that the intuition behind the portfolio shift grows stronger because the correlation between the return on capital and wage income is now zero. The non-traded asset, human capital, is now more of a substitute for the safe bond than before. Another way of looking at this is through the correlation of the equity premium with wage income, which is down to 0.37. Young workers can diversify next period income risk even more effectively by shorting the bond and investing in equity.

Table 5: Highly Persistent Technology Shocks – $\ln(A_t) = 0.8\ln(A_{t-1}) + \varepsilon_t$

	c^1	c^2	c^3	s^1_e	s^2_e	s^1_b	s^2_b	r_e	w	r_f	ep	K	K_{+1}	L	A
c^1	.	0.96	0.82	0.91	0.13	-0.9	0.88	0.03	1.00	-0.5	0.43	0.84	0.95	-0.1	0.98
c^2	.	.	0.64	0.88	0.17	-0.9	0.89	0.22	0.94	-0.6	0.64	0.70	0.93	-0.1	0.97
c^3	.	.	.	0.68	0.14	-0.6	0.59	-0.4	0.86	-0.3	-0.1	0.92	0.72	-0.1	0.75
s^1_e	.	.	.	.	-0.3	-1.0	0.99	0.09	0.90	-0.4	0.42	0.75	0.95	0.03	0.91
s^2_e	.	.	.	.	.	0.28	-0.3	0.02	0.14	-0.2	0.18	0.04	-0.0	-0.4	0.13
s^1_b	.	.	.	.	.	.	-1.0	-0.2	-0.9	0.46	-0.5	-0.7	-0.9	-0.0	-0.9
s^2_b	.	.	.	.	.	.	.	0.18	0.87	-0.5	0.51	0.68	0.94	0.04	0.89
r_e	.	.	.	.	.	.	.	.	-0.0	0.19	0.59	-0.5	0.08	-0.1	0.20
w	.	.	.	.	.	.	.	.	.	-0.5	0.37	0.86	0.94	-0.1	0.98
r_f	.	.	.	.	.	.	.	.	.	.	-0.7	-0.5	-0.5	-0.0	-0.4
ep	.	.	.	.	.	.	.	.	.	.	.	0.1	0.5	-0.0	0.49
K	.	.	.	.	.	.	.	.	.	.	.	.	0.82	0.15	0.74
K_{+1}	.	.	.	.	.	.	.	.	.	.	.	.	.	0.14	0.94
L	.	.	.	.	.	.	.	.	.	.	.	.	.	.	-0.1
A	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.

λ=0.6, β=0.6, θ=1.2, α=0.3, δ=0.65, ϕ=0.8, σ_ε=0.1, ρ=0.99, σ_v=0.01

Table 5 sheds some light on this counterintuitive result. A positive technology shock now implies a sustained increase in the expected return on capital. Therefore old workers whose main concern is consumption risk in retirement become more responsive to the state of technology in their portfolio decision. Their holdings of risky capital are now positively correlated with technology – the correlation of s^2_{et} with A_t is 0.13 whereas it was basically zero for the *iid* technology shock. A similar rationale applies to equity holdings of young workers. The correlation between s^1_{et} and A_t now lies at 0.91, whereas it was 0.88 for the *iid* shock. In effect the greater degree of persistence in technology makes equity holdings by both working-age cohorts more responsive to the observed state of technology. As a result the supply of capital increases more than for the *iid* technology shock – the correlation between the period $t+1$ capital stock (K_{+1}) and A_t rises to 0.94. This causes the second-period effect of technology shocks to grow in magnitude so that it fully offset the first-period effect – the resulting greater fluctuation in the capital-labor ratio drives the unconditional correlation of the return on capital and wage income to zero.

With the greater degree of persistence in the technology shock the economy becomes a riskier environment. Intuitively this is because a greater degree of persistence raises the probability of being born at a time when factor returns remain low for several periods. As a result agents demand a greater premium to hold risky capital. To quantify this effect the paper follows Storesletten, Telmer and Yaron (1998) who use the risk-adjusted equity premium, the Sharpe ratio, to control for risk. This is important because the assets in the model differ significantly from actual stocks and bonds. What the paper calls equity is really a composite asset that aggregates over different claims on productive capital such as stocks and corporate debt. As a result it is not leveraged and therefore not risky enough. In contrast the standard deviation of the riskfree rate is of the same order of magnitude as that of the return on capital so that – even though the bond return is know when agents make portfolio decisions – it is quite a risky asset. Because of these factors the premium agents demand to hold risky capital is small in absolute terms, around ten basis points on an annualized basis. However the risk-adjusted equity premium or the Sharpe ratio lies around 11 percent for *iid* technology shocks, rising to 15 percent when $\phi = 0.8$. Though this number is about half the Mehra and Prescott (1985) benchmark of 37 percent, the model clearly fails to explain the level of the equity premium or, put differently, why stocks are so volatile relative to returns on physical investment. This may be because investment is completely reversible so that the price of capital is always unity. Adjustment costs would relax this restriction and allow a model with a risky asset closer to actual stocks.[14] That being said this paper compares favorably with Storesletten, Telmer and Yaron (1998) whose model relies on idiosyncratic risk to generate a Sharpe ratio of 16 percent. The fact that both papers have similar risk-adjusted equity premia underlines the asset pricing implications of aggregate uncertainty at low frequencies – a key focus of this paper.

These simulation results are qualitatively unchanged for different parameterizations. As the risk aversion parameter falls – and the elasticity of intertemporal substitution rises – the correlation between the return on capital and wage income rises along with the correlation between the equity premium and wage income. When $\theta = 0.8$ and the technology shock follows an *iid* process (all other parameters are as above) the former rises to 0.46 while the latter reaches 0.76. Agents nonetheless see labor income as a close substitute for the riskfree asset and continue to shift their financial wealth from stocks to bonds as they age. The risk-adjusted equity premium or the Sharpe ratio falls along with the risk aversion parameter, to about eight percent for $\theta = 0.8$ and *iid* technology.[15] The paper also explores – but does not report – the sensitivity of the model with respect to λ, ρ and α. As λ rises towards one the steady state capital-labor ratio of the model falls because youth dependency effects become more important. For changes in the value of ρ in the range 0 – 0.99 there is no perceptible impact on portfolio behavior or the Sharpe ratio – the precise nature of demographic risk appears secondary for asset pricing. The main features of the model also appear robust to different values for α in the range 0.3 – 0.4.

In principle pay-as-you-go social security has important implications for life cycle portfolio behavior. The defined-benefit nature of the system means that agents effectively earn wage income in retirement – at an exogenous replacement rate with the transfer financed out of payroll taxes on working-age cohorts. The pension program therefore flattens the age profile of the non-traded asset, human capital, by smoothing the flow of wage income over the life cycle. This could make the shift from stocks to bonds less pronounced. The paper explores whether this effect is significant for a realistically parameterized pension system. From a broader perspective the risk-sharing properties of social security are interesting in the context of demographic risk. A defined-benefit system tends to benefit large generations who pay lower social security taxes, while the asset market effects of demographic change go in the opposite direction – large cohorts supply labor when the capital-labor ratio is low and save for retirement when it is high. Bohn (1999) shows that the asset market effects dominate when $\alpha > \eta/(1-\eta)$ in a model with risky capital as the only asset. The paper explores whether this result carries over to a more general setting in which agents make a portfolio decision over risky capital and a safe bond. Table 6 returns to the baseline specification with *iid* technology. It presents simulation results for first and second-order approximations – the third-order simulation having not yet converged. The replacement rate for the retirement benefit is assumed to be 30 percent of wage income, which implies a payroll tax rate of 15 percent when the age distribution is in equilibrium.

Table 6: The Baseline Specification with Social Security – b = 0.3

	First-Order		Second-Order		Third-Order	
	μ	σ	μ	σ	μ	σ
c^1	0.1360	0.0105	0.1364	0.0116		
c^2	0.2081	0.0203	0.2088	0.0246		
c^3	0.3165	0.0210	0.3163	0.0225		
s_e^1	0.0490	0.0044	0.0995	0.0121		
s_e^2	0.0524	0.0085	0.0028	0.0026		
s_b^1	-0.0314	0.0035	-0.0820	0.0123		
s_b^2	0.0314	0.0034	0.0820	0.0122		
r_e	1.7770	0.2431	1.7686	0.2666		
w	0.2854	0.0224	0.2860	0.0232		
r_f	1.7558	0.2122	1.7417	0.2205		
r_e - r_f	0.0213	0.2766	0.0269	0.2144		
η	0.1500	0.0017	0.1500	0.0017		
K	0.1031	0.0103	0.1040	0.0131		
L	2.0326	0.0892	2.0326	0.0892		
A	0.9990	0.1038	0.9990	0.1038		
Sharpe	0.0769		0.1252			
Score	37.9%		8.3%			

λ=0.6, β=0.6, θ=1.2, α=0.3, δ=0.65, b=0.3, ϕ=0, σ_ε=0.1, ρ=0.99, σ_v=0.01

Table 6 shows that a realistic parameterization of social security fails to smooth the life cycle profile of human capital sufficiently for there to be a significant effect on portfolio behavior. Table 7 shows that the correlation between the return on capital and wage income is basically unchanged from the baseline model – agents continue to see labor income as a close substitute for the riskless asset on balance – while the correlation between the equity premium and wage income rises somewhat to 0.71. Along with portfolio behavior the asset pricing implications of the model with social security are little changed. The Sharpe ratio remains around ten percent so that social security does little to change the risk-adjusted premium agents demand to hold capital. The retirement system does however crowd out retirement savings so that the average capital stock falls by about one-third. This effect raises the expected return on both stocks and bonds, leaving the risk premium unaffected.

Table 7: The Baseline Specification with Social Security – b = 0.3

	c^1	c^2	c^3	s^1_e	s^2_e	s^1_b	s^2_b	r_e	w	r_f	ep	K	K_{+1}	L	A
c^1	.	0.95	0.59	0.95	0.19	-0.9	0.90	0.32	0.99	-0.4	0.76	0.36	0.92	-0.0	0.92
c^2	.	.	0.34	0.95	0.20	-0.9	0.93	0.50	0.93	-0.3	0.89	0.16	0.92	-0.0	0.95
c^3	.	.	.	0.39	0.28	-0.2	0.25	-0.4	0.66	-0.5	0.02	0.78	0.41	-0.0	0.32
s^1_e	.	.	.	.	-0.1	-1.0	0.99	0.44	0.92	-0.3	0.83	0.20	0.91	-0.0	0.92
s^2_e	.	.	.	.	.	0.16	-0.2	-0.1	0.22	-0.2	0.09	0.23	0.14	0.01	0.13
s^1_b	.	.	.	.	.	.	-1.0	-0.5	-0.9	0.20	-0.9	-0.1	-0.9	0.01	-0.9
s^2_b	.	.	.	.	.	.	.	0.52	0.85	-0.2	0.86	0.09	0.89	-0.0	0.91
r_e	.	.	.	.	.	.	.	.	0.25	0.63	0.60	-0.7	0.40	0.00	0.65
w	.	.	.	.	.	.	.	.	.	-0.4	0.71	0.42	0.90	-0.0	0.90
r_f	.	.	.	.	.	.	.	.	.	.	-0.3	-0.8	-0.3	-0.0	-0.0
ep	.	.	.	.	.	.	.	.	.	.	.	-0.0	0.80	0.02	0.82
K	.	.	.	.	.	.	.	.	.	.	.	.	0.35	0.33	0.01
K_{+1}	.	.	.	.	.	.	.	.	.	.	.	.	.	0.34	0.89
L	.	.	.	.	.	.	.	.	.	.	.	.	.	.	-0.0
A	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.

λ=0.6, β=0.6, θ=1.2, α=0.3, δ=0.65, b=0.3, ϕ=0, σ_ε=0.1, ρ=0.99, σ_ν=0.01

V. Simulating the Effects of a Baby Boom and Baby Bust on Asset Returns

Following World War II many developed countries experienced a surge in birth rates that has come to be know as the Baby Boom. The Baby Boom lasted for roughly two decades and was followed by the Baby Bust, a decline in birth rates that lasted another two decades and is commonly linked to the introduction of “the pill” in the early 1970s. Academic interest in this fertility shift appears to have tracked the progress of the Baby Boomers through the age distribution. Over a decade ago Mankiw and Weil (1989) explored the link between a rising number of households in their “homebuying years” and increasing real housing prices during the late 1970s and early 1980s. Subsequently Fair and Dominguez (1991) examined the repercussions of population change on various macroeconomic relationships, with a special focus on consumption and savings behavior. More recently the entry of the Baby Boomers into “peak saving years” has shifted attention towards the asset market effects of population dynamics. Bakshi and Chen (1994), Bergantino (1998), Brooks (1998), Erb, Harvey and Viskanta (1997), Poterba (1998) and Yoo (1994) have each investigated the empirical link between demographic change and real returns on financial assets – finding on balance only weak evidence to support a relationship. Meanwhile the financial press has speculated about what will happen to financial markets when the Baby Boomers retire.[16] One view holds that the Baby Boomers will have to sell their assets to a smaller generation of young investors when they retire. At this point excess supply of financial assets will lead to a decline – perhaps a crash – in the real value of financial assets. An opposed view maintains that forward-looking financial markets are pricing assets to incorporate the aging of the Baby Boomers so that a market meltdown is unlikely. This section goes looking for the market meltdown, using the model to simulate the general equilibrium effects of a baby boom and bust on asset returns.

Figures 1 and 2 plot historical and projected youth dependency ratios for a number of countries.[17] They place the Baby Boom and Baby Bust in the context of a long run trend towards lower fertility, temporarily exacerbated by the Great Depression and World War II as people around the world postponed having children. The Baby Boom is thus part recovery from extremely low birth rates, while the Baby Bust is the return to a longer-term trend of fertility decline, accelerated by the introduction of "the pill." The data clearly illustrate that aggregate fertility is endogenous over the long run – driven in part by economic and political fundamentals – a feature the model ignores because it abstracts from the decision to have children.

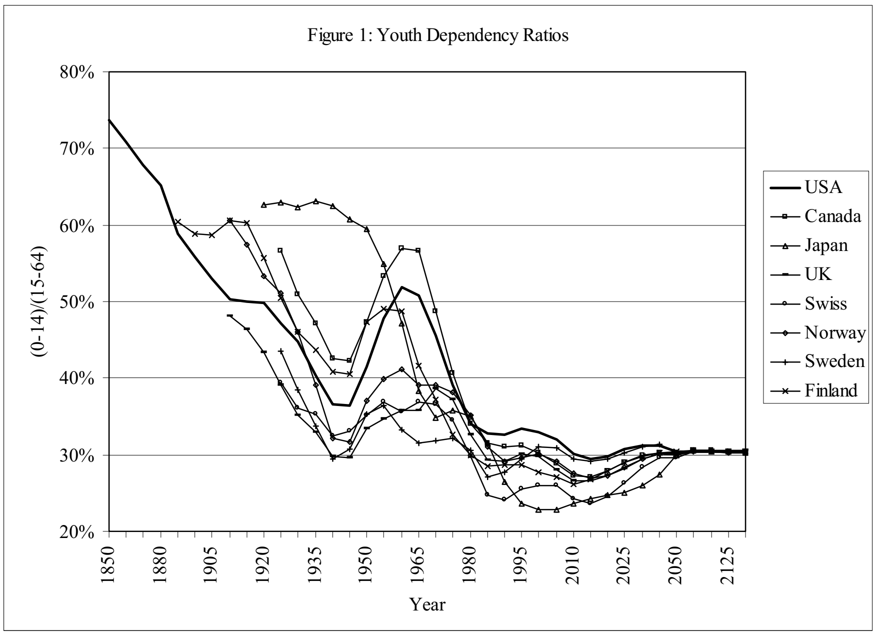

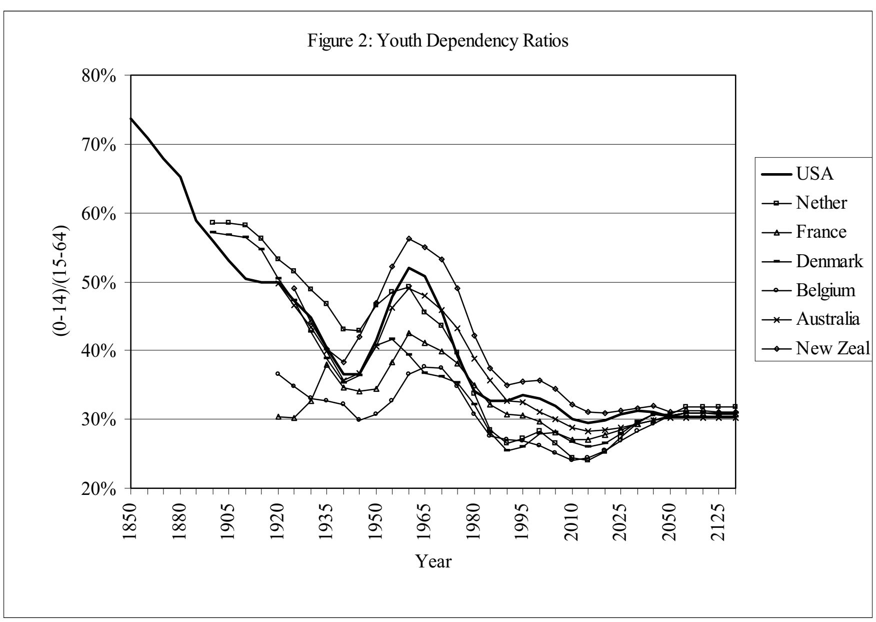

Figures 1 and 2 show that the Baby Boom and Baby Bust is a common feature across the developed world – with one important exception. Japan remained in a high fertility equilibrium until after World War II, unlike most other countries where birth rates were declining well before. It then experienced a precipitous decline in fertility, such that the youth dependency ratio declined from 60 percent in 1950 – the highest ratio among these countries – to 23 percent in 2000 – the lowest. The remarkable scale and speed of this transition has made the Japanese age distribution the odd-man-out among industrialized countries. Figure 3 plots annualized cohort growth over the simulated demographic shift, which should be thought of as a stylized version of the Baby Boom and Baby Bust in North America and the European Union countries.[18]

Figure 3: Annualized Cohort Growth

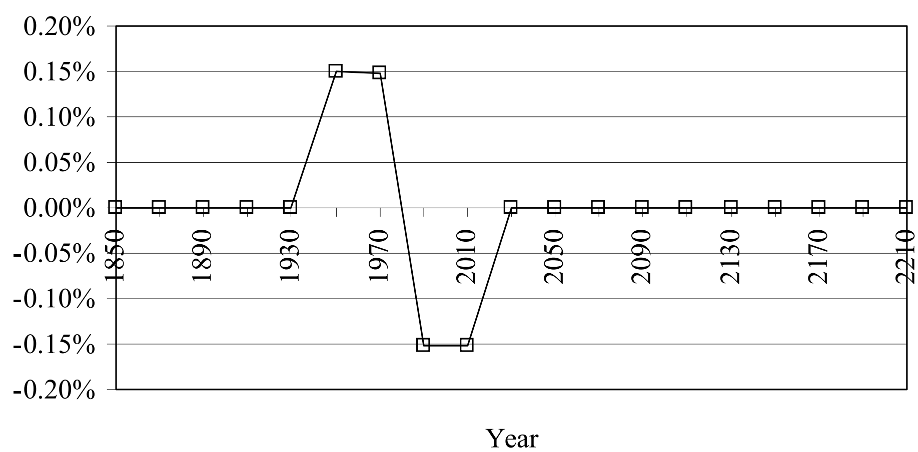

The simulated baby boom and baby bust is four periods long. The boom lasts for two of these periods, beginning in 1950 when cohort growth jumps to three percent (0.15 percent annualized). Cohort growth in the second boom period is also three percent. The bust begins in 1990 when cohort growth falls to –3 percent and remains at that rate into 2010. The age distribution begins its return to equilibrium in 2030 when cohort growth returns to its steady state level of zero. Key features of this demographic shift are that it is large and symmetric. Cohort growth of three percent per period corresponds to about three standard deviations of the cohort growth shock in the model. The shock is therefore large relative to expectations, a reflection of the size of the Baby Boom and Baby Bust relative to the earlier fertility decline, which was more gradual. Because the transition is symmetric it focuses on transitional dynamics rather than capital deepening effects associated with a long run decline in fertility.

Figure 4 plots the expected return on capital (re) and the riskfree rate (rf) over the simulated baby boom and baby bust for a specification with $\theta = 1$, $\lambda = 1$, $\alpha = 0.3$, $\delta = 0.4$ and $b = 0$. The technology shock is assumed constant at its expected value $E(A_t) = 1$, which corresponds to holding non-demographic fundamentals constant over the baby boom and baby bust. The focus is therefore on the effects of demographic change exclusively. Figure 4 shows that the baby boom and baby bust have two distinct effects on asset returns. The first raises returns on both assets during the boom, when aggregate saving is relatively low because of rising youth dependency. It later reduces stock and bond returns when aggregate saving becomes large as boomer cohorts save for retirement and youth dependency declines with the baby bust. This effect reflects changes in the real interest rate as aggregate saving changes over the baby boom and baby bust. The second effect moves the bond return more than the expected return on capital, a reflection of the greater sensitivity of the riskfree asset to changes in the age distribution. Figure 5 shows that this effect produces large swings in the equity premium over the baby boom and baby bust.

Figure 4: Annualized Equity and Bond Returns

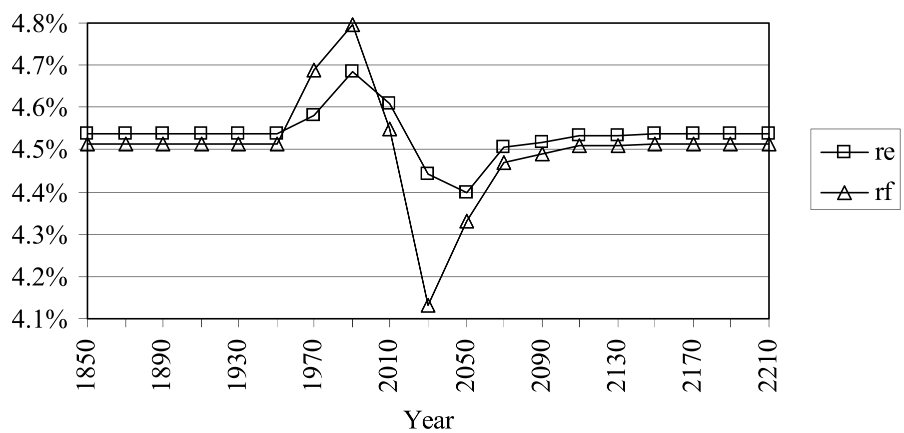

Figure 5: Equity Premium

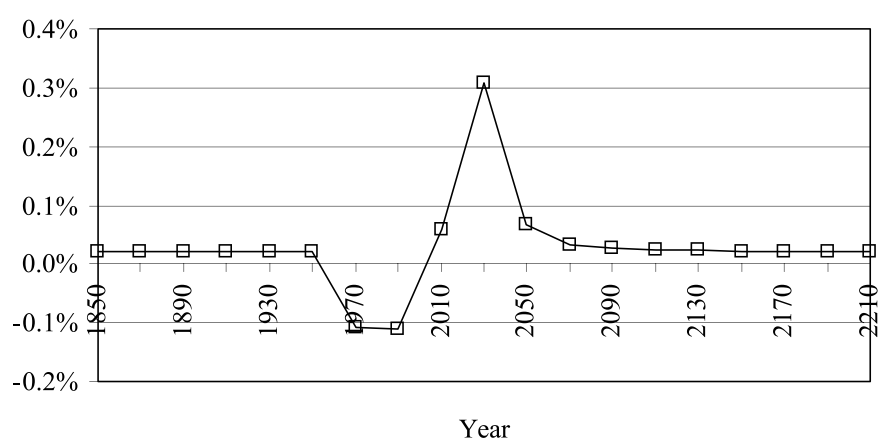

The effect on the equity premium is driven by the interaction of the baby boom and baby bust with portfolio decisions and is discussed period by period. Though the baby boom begins in 1950 the expected return on capital and the riskfree rate remain in steady state in that period.

This is because the labor supply and capital stock remain in equilibrium, the latter having been determined by portfolio decisions in the previous period that were unaffected by the baby boom. In effect the baby boom impacts asset markets with a one period lag, a key feature of the overlapping generations model. The riskfree rate in 1970 rises because the increase in youth dependency in 1950 raises the consumption requirement of young workers while their wage income remains unchanged. This raises their borrowing requirement or the incipient supply of the riskfree asset. With demand for the riskfree asset basically unchanged – old workers in 1950 demand the steady state amount of bonds going into retirement – the riskfree rate must rise to clear the bond market. This effect pushes the riskfree rate above the expected return on capital, which rises because the capital-labor ratio in 1970 falls as the first boomer cohort enters the labor force. In the 1990 period the riskfree rate and the expected return on capital rise still further, with the equity premium roughly unchanged.

Figure 6: Ratio of Old to Young Workers

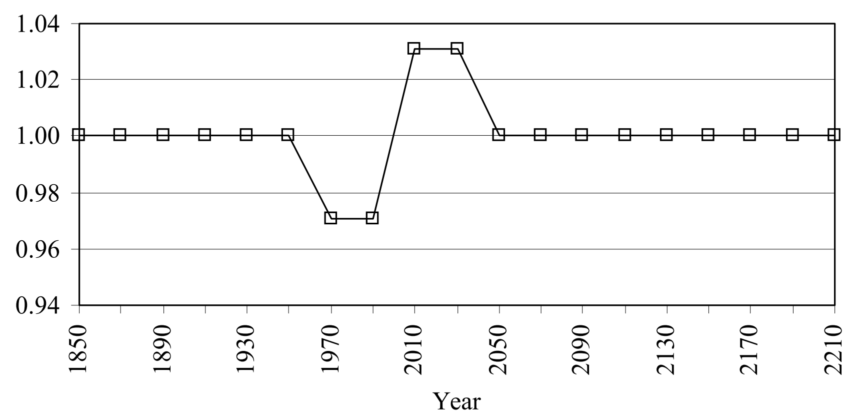

The further rise in the riskfree rate is linked to a demographic imbalance on the bond market, illustrated by the ratio of old to young workers in Figure 6. In 1970 this ratio falls below steady state as boomers enter the labor force. This exacerbates the incipient excess supply of bonds that stems from the ongoing baby boom, which raises the consumption needs of young households. The expected return on capital rises in 1990 because the capital-labor ratio falls still further. In 1990 the boom turns to bust. Young workers are now raising relatively fewer children, which reduces their borrowing requirement. This decline in the incipient supply of the riskless asset is partly offset by a decline in demand, a consequence of declining wage income. On balance the riskfree rate in 2010 falls to just above equilibrium and the equity premium turns positive again. The expected return on capital in 2010 falls because the first cohort of the baby bust has entered the workforce, raising the capital-labor ratio. The riskfree rate falls to a record low in 2030 because the last of the boom cohorts (old workers) trades on the bond market with the first of the baby bust generation (young workers). This imbalance results in excess demand for the bond, which is exacerbated by the fact that young workers have a small borrowing requirement – they are after all raising the second baby bust cohort. Though the expected return on capital in 2030 falls further as the capital-labor ratio continues to rise, the equity premium peaks at around 30 basis points. Beyond 2030 the expected return on capital and the riskfree rate return to their steady state levels as the remnants of the baby boom and baby bust work their way though the age distribution.[19]

The simulation generates significant asset market effects over the boom and bust. They hit one generation in particular, the younger of the two boomer cohorts, which wants to place its retirement savings in the riskless asset but must trade with the first of the baby bust cohorts. The resulting excess demand for bonds pushes the 2030 riskfree rate 15 percent below its equilibrium level and 20 percent below its peak in 1990. This effect is significant – certainly for retirement consumption of tail-end baby boomers, which is around ten percent below that of their parents and grand-parents – and perfectly consistent with rational, forward-looking financial markets. It highlights that markets in overlapping generations models are fundamentally incomplete, in that non-overlapping generations cannot trade. The interaction of this feature of the model with the life cycle portfolio shift from stocks to bonds is what dooms tail-end baby boomers. Of course the magnitude of this effect depends on a number of assumptions. First, the extent of market incompleteness is declining in the number of periods of the model. As the number of overlapping generations increases, swings in the number of borrowers versus lenders will become less extreme so that changes in the riskfree rate will be less pronounced. Put differently, as the number of working-age cohorts in the model rises, more generations are able to trade on financial markets, which improves the extent to which they can diversify generation-specific risk. Second, an infinitely-lived agent such as government can overcome the incompleteness of the model, by facilitating trade between non-overlapping generations. Indeed, as Bohn (1999) points out, a defined-benefit social security system is *ex ante* efficient in that it acts as an insurance mechanism against the misfortune of being born into a large cohort. And while in principle social security could offset the asset market effects of demographic shifts, this paper finds that it fails to do so for realistic levels of the payroll tax.[20]

This result bears on the current debate over Social Security reform because it highlights the role of Social Security as a mechanism for intergenerational risk-sharing. Moreover the above simulation shows that – in a world where Social Security has been replaced with individual retirement accounts – demographic shocks have significant asset market effects that hit some generations harder than others. Past returns and the historical equity premium are therefore a poor guide to returns investors can expect over the coming demographic transition.

VI. Conclusion

This paper explores life cycle portfolio choice in a Diamond– style neoclassical growth model with overlapping generations, in which agents make a portfolio decision over risky equity and safe bonds. It generates portfolio behavior whereby optimizing agents shift from stocks to bonds as they age and motivates this portfolio shift in terms of implicit holdings of a non-traded asset, human capital, which is run down over the life cycle. The paper goes beyond the existing literature in using a general equilibrium model to explore the determinants of life cycle portfolio choice and finds that the shift from stocks to bonds is robust across a wide range of specifications. It focuses especially on the robustness of the results to the inclusion of pay-as-you-go social security, which effectively flattens the life cycle profile of human capital by transferring wage income from workers to retirees. The paper finds that portfolio behavior is unchanged for realistic levels of the payroll tax.

The paper uses the model to simulate the general equilibrium effects of a baby boom and baby bust on asset returns. It finds evidence of significant asset market effects, with the expected return on retirement savings of boomer cohorts up to 20 percent below returns to earlier generations. Though a defined-benefit social security system could in principle offset these asset market effects – effectively providing insurance against the risk of being born into a large cohort – the paper finds that it fails to do so for realistic levels of the payroll tax.

This result bears on the current debate over Social Security reform because it highlights the role of Social Security as a mechanism for intergenerational risk-sharing. Moreover the paper shows that – in a world where Social Security has been replaced with individual retirement accounts – demographic shocks have significant asset market effects that hit some cohorts harder than others. Past returns and the historical equity premium are therefore a poor guide to returns investors can expect over the coming demographic transition.

Appendix: Solving the Model using the Parameterized Expectations Approach

Maximizing expected utility period t young workers choose c_t^0, c_t^1, s_{et}^1, and s_{bt}^1 such that

$$c_t^0 = \lambda^{1/\theta} c_t^1 \tag{A1}$$

$$\left(c_t^1\right)^{-\theta} = \beta E_t \left[\left(c_{t+1}^2\right)^{-\theta} \left(1 + r_{et+1}\right)\right] \tag{A2}$$

$$\left(c_t^1\right)^{-\theta} = \beta \left(1 + r_{ft}\right) E_t \left[\left(c_{t+1}^2\right)^{-\theta}\right] \tag{A3}$$

$$\left(1 + n_t\right) c_t^0 + c_t^1 + s_{et}^1 + s_{bt}^1 = w_t \left(1 - \eta_t\right) \tag{A4}$$

are satisfied, taking factor returns and the return on the riskless asset as given. Period t old workers choose c_t^2, s_{et}^2, and s_{bt}^2 such that

$$\left(c_t^2\right)^{-\theta} = \beta E_t \left[\left(c_{t+1}^3\right)^{-\theta} \left(1 + r_{et+1}\right)\right] \tag{A5}$$

$$\left(c_t^2\right)^{-\theta} = \beta \left(1 + r_{ft}\right) E_t \left[\left(c_{t+1}^3\right)^{-\theta}\right] \tag{A6}$$

$$c_t^2 + s_{et}^2 + s_{bt}^2 = w_t \left(1 - \eta_t\right) + \left(1 + r_{et}\right) s_{et-1}^1 + \left(1 + r_{ft-1}\right) s_{bt-1}^1 \tag{A7}$$

are satisfied, again taking factor returns and the return on the riskless asset as given. Consumption of the period t retiree cohort is given by:

$$c_t^3 = (1+r_{et})s_{et-1}^2 + (1+r_{ft-1})s_{bt-1}^2 + bw_t \tag{A8}$$

(A1) through (A8) represent a system of eight equations that characterize individual consumption $(c_t^0, c_t^1, c_t^2, c_t^3)$ and investment behavior $(s_{et}^1, s_{bt}^1, s_{et}^2, s_{bt}^2)$ for given wage and return distributions. In equilibrium, consumption and investment decision rules that maximize expected utility at the individual level must be consistent with the equilibrium conditions for the stock and bond markets.

The model has only two active decision makers: young and old workers. Both make their consumption-investment decision based on total wealth, which for young workers is simply after-tax wage income.

$$w_t^1 = w_t(1-\eta_t) \tag{A9}$$

Total wealth of old workers consists of the after-tax wage, in addition to stock and bond holdings plus returns.

$$w_t^2 = w_t(1-\eta_t) + (1+r_{et})s_{et-1}^1 + (1+r_{ft-1})s_{bt-1}^1 \tag{A10}$$

w_t^1 and w_t^2 are the distribution of wealth across working-age cohorts and are endogenous state variables. In addition the age distribution with the exception of the retiree cohort, which will not live to see the next period, represents an exogenous state variable. Assuming that the technology shock is *iid*, the set of period *t* state variables is then:

$$\Theta_t = [w_t^1, w_t^2, N_t, N_{t-1}, N_{t-2}] \tag{A11}$$

Using the parameterized expectations approach to replace the conditional expectations in (A2), (A3), (A5), and (A6) with polynomials in the state variables yields:

$$(c_t^1)^{-\theta} = \beta\Psi(\Theta_t, \tau) \tag{A12}$$

$$(c_t^1)^{-\theta}(s_{et}^1)^2 = \beta(1+r_{ft})\Omega(\Theta_t, \gamma) \tag{A13}$$

$$(c_t^2)^{-\theta} = \beta\Lambda(\Theta_t, \xi) \tag{A14}$$

$$(c_t^2)^{-\theta}(s_{et}^2)^2 = \beta(1+r_{ft})\Gamma(\Theta_t, \omega) \tag{A15}$$

It is worth noting that the expectations in (A13) and (A16) have not been parameterized in the traditional manner. Both equations have been multiplied by functions of the respective equity holdings. This modification is based on Izvorski (1997) and addresses an indeterminacy in the system of Euler equations and aggregate equilibrium conditions that arises in models that solve for equilibrium holdings of two or more assets.

Given two sequences of length *T* for the technology shock and the age distribution, assuming starting values for w_t^1 and w_t^2, and given values for τ, γ, ξ, and ω, it is then possible to solve out the model for *T* periods. The PEA begins with exactly this step. It draws two sequences of length *T* for the technology shock and the age distribution and solves out the model for the *T* periods. The PEA then turns to fitting the conditional expectations in (A2), (A3), (A5), and (A6), solving for the coefficients in the polynomial approximations that minimize the mean squared error between the actual realization in $t+1$ and the expectation at t of that realization based on Θ_t. The particular version of the PEA implemented here proceeds in a step by step approach. Fitting the conditional expectation in (A12), a non-linear least squares estimation for τ is performed:

$$\min_{\tau} \frac{1}{T-1} \sum_{t=1}^{T-1} \left[\left(c_{t+1}^2 \right)^{-\theta} \left(1 + r_{et+1} \right) - \Psi\left(\Theta_t, \tau \right) \right]^2 \tag{A16}$$

At the *n*'th iteration a new value τ_{n+1} is generated according to $\tau_{n+1} = \lambda\tau_n + (1-\lambda)\tau_e$ where τ_e is the estimate for τ from the non-linear least squares estimation. Given τ_{n+1}, γ_n, ξ_n, and ω_n, the system is solved out again for *T* periods and the algorithm proceeds to fit the other conditional expectations in turn. This procedure is repeated until the algorithm reaches a fixed point in τ, γ, ξ, and ω.

Rather than performing a computationally expensive non-linear least squares estimation to find τ_e, the paper takes a first-order approximation of $\Psi(\Theta_t, \tau_n)$ around τ_n following Den Haan and Marcet (1990). Rearranging terms τ_e is then the coefficient vector in an OLS regression. The basic intuition behind the parameterized expectations approach is to approximate the conditional expectations of period $t+1$ using information available to agents at t. For a successful approximation the prediction error should therefore be orthogonal to agents' information set at t. This intuition lies behind an accuracy test developed by Den Haan and Marcet (1994). The accuracy test checks for orthogonality between the Euler equation residuals and a vector v_t of variables in agents' period t information set.

$$\begin{bmatrix} \left(c_{t+1}^2 \right)^{-\theta} \left(1 + r_{et+1} \right) - \Psi\left(\Theta_t, \tau^* \right) \\ \left(c_{t+1}^2 \right)^{-\theta} \left(s_{et}^1 \right)^2 - \Omega\left(\Theta_t, \gamma^* \right) \\ \left(c_{t+1}^3 \right)^{-\theta} \left(1 + r_{et+1} \right) - \Lambda\left(\Theta_t, \xi^* \right) \\ \left(c_{t+1}^3 \right)^{-\theta} \left(s_{et}^2 \right)^2 - \Gamma\left(\Theta_t, \omega^* \right) \end{bmatrix} = \varepsilon_{t+1} \tag{A17}$$

where τ^*, γ^*, ξ^*, and ω^* are the PEA parameter estimates at convergence. For any $m \times 1$ vector v_t in agents' period t information set, the statistic

$$G = (T-1)\left[\frac{\sum_{t=1}^{T-1}(\varepsilon_{t+1}\otimes v_t)}{T-1}\right]'\left[\frac{\sum_{t=1}^{T-1}(\varepsilon_{t+1}\otimes v_t)(\varepsilon_{t+1}\otimes v_t)'}{T-1}\right]^{-1}\times$$

$$\left[\frac{\sum_{t=1}^{T-1}(\varepsilon_{t+1}\otimes v_t)}{T-1}\right] \tag{A18}$$

has an asymptotic χ^2 distribution with degrees of freedom given by $4\times m$. The vector of state variables Θ_t is chosen for v_t. This test is implemented in the following manner. Given τ^*, γ^*, ξ^*, and ω^* at convergence, the model is simulated N times, each time for different draws of the technology shock and the age distribution. For these N simulations, the frequency with which the G statistic is greater than the critical value of the 95th percentile of a χ^2_{20} is reported. If the percentage of G statistics above the critical value of the 95th percentile is substantially greater than five percent, this is evidence against accuracy of the solution.

References

Abel, A. "Will Bequests Attenuate the Predicted Meltdown in Stock Prices When Baby Boomers Retire?" Unpublished Manuscript, March 2000

Abel, A. "The Effects of a Baby Boom on Stock Prices and Capital Accumulation in the Presence of Social Security." Unpublished manuscript, October 1999

Altonji, J. Hayashi, F. and Kotlikoff, L. "Risk Sharing between and within Families." Econometrica 64 (1996): 261 – 294

Ameriks, J. and Zeldes, S. "How do Household Portfolio Shares Vary with Age?" Unpublished manuscript, February 2000

Attanasio, O. and Violante, G. "The Demographic Transition in Closed and Open Economies: A Tale of Two Regions." Unpublished manuscript, February 2000

Bakshi, G. and Chen, Z. "Baby Boom, Population Aging, and Capital Markets." Journal of Business, Vol. 67, No.2, 1994: 165 - 202

Bergantino, S. "Life Cycle Investment Behavior, Demographics, and Asset Prices." Dissertation, M.I.T., September 1998

Bohn, H. “Social Security and Demographic Uncertainty: The Risk Sharing Properties of Alternative Policies.” NBER Working Paper 7030, March 1999

Brooks, R. “Population Aging and Global Capital Flows in a Parallel Universe.” Unpublished manuscript, September 2000

Brooks, R. “What Will Happen to Financial Markets When the Baby Boomers Retire?” IMF Working Paper, WP/00/18, February 2000

Brooks, R. “Asset Market and Savings Effects of Demographic Transitions.” Dissertation, Yale University, December 1998

Cocco, J., Gomes, F. and Maenhout, P. “Consumption and Portfolio Choice over the Life Cycle.” Unpublished manuscript, 1999

Colvin, G. “How to Beat the Boomer Rush.” Fortune, August 18, 1997, 59 - 63

Constantinides, G., Donaldson, J. and Mehra, R. “Junior Can’t Borrow: A New Perspective on the Equity Premium Puzzle.” Mimeo, May 1998

Constantinides, G. and Duffie, D. “Asset Pricing with Heterogeneous Consumers.” Journal of Political Economy 104 (1996): 219 - 240

Den Haan, W. and Marcet, A. “Solving the Stochastic Growth Model by Parameterizing Expectations.” Journal of Business & Economic Statistics 8(1990): 31 – 34

Den Haan, W. and Marcet, A. “Accuracy in Simulations.” Review of Economic Studies 61 (1994): 3 – 18

Erb, C., Harvey, C. and Viskanta, T. “Demographics and International Investments.” Financial Analysts Journal (July/August 1997): 14 - 28

Heaton J. and Lucas, D. “Portfolio Choice in the Presence of Background Risk.” The Economic Journal 110 (2000): 1 - 26

Heaton J. and Lucas, D. “Market Frictions, Savings Behavior, and Portfolio Choice.” Macroeconomic Dynamics 1 (1997): 76 - 101

Higgins, M. “Demography, National Savings and International Capital Flows.” International Economic Review, Vol. 39 No. 2 (1998): 343-369

Higgins, M. and Williamson, J. “Asian Demography and Foreign Capital Dependence.” NBER Working Paper 5560 (1996)

Ibbotson Associates. “Stocks, Bonds, Bills and Inflation: 1999 Yearbook.” 1999

Izvorski, I. “Asset Pricing with Heterogeneous Agents.” Unpublished manuscript, September 1997

Jagannathan, R. and Kocherlakota, N. “Why should Older People Invest Less in Stocks than Younger People?” Federal Reserve Bank of Minneapolis Quarterly Review, Summer 1996

Mehra, R. and Prescott, E. “The Equity Premium: A Puzzle.” Journal of Monetary Economics 15(1985): 145 - 161

Passell, P. “The Year is 2010. Do you know where your Bull is?” The Sunday New York Times, March 10, 1996

Poterba, J. “Population Age Structure and Asset Returns: An Empirical Investigation.” Mimeo, 1998

Storesletten, K., Telmer, C. and Yaron, A. “Persistent Idiosyncratic Shocks and Incomplete Markets.” Unpublished manuscript, February 1998

Vissing-Jorgensen, A. “Limited Stock Market Participation and the Equity Premium Puzzle.” Unpublished manuscript, December 1999

Yoo, P. “Age Distributions and Returns of Financial Assets.” Federal Reserve Bank of St. Louis Working Paper 94-002B, February 1994

---

[1] Yoo (1994) uses three separate cross-sections from the Survey of Consumer Finances to analyze age patterns in asset allocation. He finds that within each of the cross-sections, the share of financial wealth in equities increases over the working life and declines after retirement, generating a hump-shaped pattern. His result is robust to the inclusion of measures of human capital and total wealth. Poterba and Samwick (1997) use pooled data from three separate cross-sections from the Survey of Consumer Finances to distinguish the separate impact of age and cohort membership on ownership of financial assets. They find that the age-profiles for equity ownership and portfolio shares for directly held corporate stock increase at younger ages and decrease at later ages (a hump-shaped pattern). Heaton and Lucas (2000) use three separate cross-sections of data from the Survey of Consumer Finances to analyze household asset portfolios. They find that households older than 65 hold a smaller fraction of financial wealth in stocks than do younger households. Ameriks and Zeldes (2000) use pooled data from separate Surveys of Consumer Finances and panel data from TIAA-CREF to show that equity ownership has a hump-shaped pattern with age – though equity shares conditional on ownership are nearly constant across age groups.

[2] Equilibrium models that allow for consumer holdings of bonds consistent with the observed per capita supply of bonds are hard to construct. Constantinides and Duffie (1996) point out that this difficulty may be resolved if consumers' idiosyncratic income shocks are sufficiently persistent and heteroscedastic. Storesletten, Telmer and Yaron (1997) investigate the role of idiosyncratic income shocks in the context of a calibrated overlapping generations model. In their model idiosyncratic shocks fail to explain a large fraction of the equity premium.

[3] Bohn (1999) constructs a Diamond-style overlapping generations model in which only a fraction of every generation survives to the next period. This setting allows him to model uncertainty associated with both aggregate fertility and life expectancy.

[4] Strong risk aversion and the existence of significant holding costs may explain the limited degree of stock market participation. In a recent study using panel data from the PSID Vissing-Jorgensen (1999) finds that the likelihood of participation in the stock market in one period is strongly correlated with participation in the past,

which may be because of significant entry or participation costs. The asset market implications of borrowing constraints are explored by Constantinides, Donaldson and Mehra (1998), in an overlapping generations model that takes production as exogenous.

[5] A number of papers, among them Attanasio and Violante (2000), Brooks (2000) and Juillard et. al. (2000), use overlapping generations models of large open economies to simulate the impact of projected differences in population trends across regions on international capital flows. All papers point to a turning point in regional savings – investment balances between 2010 and 2030 when rapid population aging in the developed world will cause it to become a capital importer. This switch will be financed by developing regions who are projected to become capital exporters. However this literature likely overstates the magnitude of implied capital flows because it uses models with perfect foresight, so that developing country assets are perfect substitutes for those of the developed world.

[6] See Abel (2000) for a Diamond-style overlapping generations model in which agents have a bequest motive. The paper finds that a baby boom generates an increase in stock prices and that stock prices are rationally anticipated to fall when the baby boomers retire, even though – as emphasized by Poterba (1998) – they do not spend all their assets during retirement.

[7] Bergantino (1998), using Survey of Consumer Finances data, finds that fewer than 25 percent of households report having ever received a substantial inheritance, trust, or transfer. Of those that did, the median value in 1995 dollars was about $17,000 per spouse, or about 60 percent of the median annual income. Altonji et. al. (1996) find that private intergenerational risk sharing is highly imperfect empirically. Since the model does not feature lifetime uncertainty, there are no accidental bequests.

[8] The magnitude of factor price movements depends on the elasticity of factor substitution. A Cobb-Douglas technology implicitly imposes a unit elasticity, while a CES production function with an elasticity greater than one would imply smaller factor price changes. Bohn (1999) argues that the Cobb-Douglas specification is appropriate and that it may even understate factor price movements given production function estimates.

[9] The assumption that each period represents about 20 years imposes the unintended restriction that asset returns are perfectly serially correlated over this time frame. Obviously it is of interest to increase the number of periods within the 80-year span of the model.

[10] This approach is still quite general because government transfers matter only through the generational accounts of different cohorts. The social security system could therefore be regarded as subsuming other intergenerational transfers.

[11] Other specifications for social security are possible. Holding the payroll tax fixed and allowing the replacement rate to adjust to changes in the retiree to worker ratio would correspond to a defined-contribution system. Bohn (1999) argues that a fully privatized social security system (where $b = \eta = 0$) has essentially the same risk-sharing properties as a defined-contribution pay-as-you-go system. This is because neither the defined-contribution nor the privatized system impose higher taxes on the young when the retiree to worker ratio rises, whereas the defined-benefit system does. In testing the sensitivity of the asset market implications of the model, a defined-benefit specification is therefore more interesting.

[12] These parameter choices are in line with the recent overlapping generations literature. Higgins (1994) chooses α=0.33, β=0.54, δ=0.72, and θ=0.77 in the context of a 3 period model. Higgins and Williamson (1996) appear to have the same parameterization, with the exception that θ is set at 1. Attanasio and Violante (2000), in a model with mortality risk, set α=0.36, β=1.011 per annum, δ=0.5 per annum, and θ=2. Constantinides, Donaldson and Mehra (1998) set β=0.44 and θ= 2, 4, 6 in a three period model without production.

[13] The correlation between wage income and the return on capital also depends on the risk aversion parameter θ and the depreciation factor (1-δ). Increasing both raises the proportion of shocks being saved, strengthening the second-period effect and generating a lower unconditional correlation. Both effects are small within reasonable parameter ranges though.

[14] For a Diamond-style overlapping generations model with adjustment costs see Abel (1999).

[15] Heaton and Lucas (1997) point out that unless the agent is sufficiently risk averse, all savings will be carried out by stock holdings. They show that for a large enough risk aversion parameter (they consider two values: 5 and 8), agents do some saving by holding bonds, even if non-traded income is uncorrelated with stock returns.

[16] For examples of articles in the financial press, see Passell (1996) and Colvin (1997).

[17] The age distribution data up to 1990 were provided by the statistical agencies in each country. The demographic data from 1995 onward are drawn from the World Bank's "World Population Projections: The 1994-95 Edition." This data splice is necessary because the World Bank data go back only to 1990. In 1990 the age distribution data from both sources matches almost perfectly. An important property of the World Bank projections is convergence in population growth over the very long run. As a result youth dependency across countries converges at 30 percent by 2075. This feature of the projections is better interpreted as an "equilibrium condition" rather than as a forecast.

[18] The simulated Baby Boom and Baby Bust is calibrated to cohort growth in a region consisting of North America (NA) and the European Union (EU) countries. Cohort growth rates are generated by aggregating the 5-year age groups in the original data into 20-year age groups that span 0-19, 20-39, 40-59, 60+. Cohort growth is then the growth rate of the "child" cohort over the "young worker" cohort over a 20 year interval. In the 20 year period centered around 1950 annualized cohort growth in NA amounted to 0.5 percent, 0.6 percent in the EU, and 0.6 percent in the combined region. In the period centered around 1970, the peak of the Baby Boom, cohort growth rose to just under 2 percent per year in NA and just under 1 percent in the EU, for a growth rate of about 1.3 percent in the combined region. In the period around 1990, the peak of the Baby Bust, cohort growth turns negative. It is –0.7 percent per annum in NA, –0.9 percent in the EU, and –0.8 percent in the combined region. In the period around 2010 cohort growth remains negative in the EU (-0.7 percent per year) but rebounds to 0.2 percent in NA, for an average of –0.2 percent in the combined region.

[19] It is worth noting that the effect on stock returns is small compared to the run-up in stock prices over the past 20 years. The average real return on the Ibbotson Associates large stock index from 1979 – 1998 is 13.43 percent, relative to 3.53 percent for the period 1959 – 1978. In comparison the 1990 model return on capital is only 4.8 percent above steady state. Changes in the age distribution are thus an insufficient explanation for the recent surge in stock prices. Of course this simulation exercise holds non-demographic fundamentals constant and ignores the possibility of a speculative bubble.

[20] A lower value for λ tends to dampen the youth dependency effects during the boom phase of the transition, so that the equity premium is less likely to turn negative early on. The magnitude of the decline in the riskfree rate in 2030 is not sensitive to this feature of the model. The magnitude of the asset market effects is relatively invariant to other parameter values.