

WP/00/18

IMF Working Paper

What Will Happen to Financial Markets When the Baby Boomers Retire?

Robin Brooks

WP/00/18

INTERNATIONAL MONETARY FUND

Research Department

What Will Happen to Financial Markets When the Baby Boomers Retire?

Prepared by Robin Brooks¹

Authorized for distribution by Peter Isard

February 2000

Abstract

This paper explores whether changes in the age distribution have significant effects on financial markets that are rational and forward-looking. It presents an overlapping generations model in which agents make a portfolio decision over stocks and bonds when saving for retirement. Using the model to simulate a baby boom–baby bust demonstrates that returns to baby boomers will be substantially below returns to earlier generations, even when markets are rational and forward-looking. This result is important because the current debate over how to reform pay-as-you-go pension systems often takes historical returns on financial assets—and on the equity premium—as given.

JEL Classification Numbers: E27, G11, G12, H55

Keywords: equity premium, population aging, pension reform

Author's E-Mail Address: rbrooks2@imf.org

¹ Special thanks are due to Christopher Sims. I am also grateful to Luis Cubeddu, Hamid Faruqee, Wouter Den Haan, Peter Heller, Peter Isard, Ivailo Izvorski, Narayana Kocherlakota, Douglas Laxton, and Eswar Prasad for helpful comments. Any errors are mine.

Contents

Page

I. Introduction	4
II. The Paper vis-à-vis the Literature	7
III. The Model.....	8
IV. Equilibrium	11
V. The State Variables.....	11
VI. The Solution Method.....	12
VII. Accuracy of the Solution Method.....	13
VIII. Model Parameterization	14
IX. Simulation Results	16
X. Simulating a Baby Boom-Baby Bust	23
XI. Conclusion.....	33

Text Tables

1. The Age Distribution over Time	9
2. Steady State Values for the Model without Uncertainty	16
3. Descriptive Statistics of the Consumption-Saving Model.....	17
4. Correlations based on Third-Order Approximation	18
5. Descriptive Statistics of the Consumption-Investment Model	18
6. Correlations based on Third-Order Approximation using (4)	20
7. Descriptive Statistics of the Consumption-Investment Model with Social Security	20
8. Correlations based on Third-Order Approximation	21
9. Descriptive Statistics of the Consumption-Investment Model using (4')	22
10. Correlations based on Third-Order Approximation	22
11. Correlations based on Third-Order Approximation	25

Figures

1. Population Growth	23
2. Returns on Capital and the Riskless Asset	24
3. Wage Income	24

4. Return Differential between Stocks and Bonds	25
5. Per Capita Consumption of Young Workers	25
6. Per Capita Equity Holding of Young Workers	26
7. Per Capita Bond Holding of Young Workers	26
8. Per Capita Equity Holding of Old Workers	26
9. Per Capita Bond Holding of Old Workers	27
10. Ratio of Old to Young Workers	27
11. Per Capita Consumption of Retirees	28
12. Discounted Lifetime Utility	28
13. Wage Income	30
14. Workers per Retiree	30
15. Per Capita Retirement Benefit	30
16. Returns on Capital and the Riskless Asset	31
17. Wage Income	31
18. Return Differential between Stocks and Bonds	32
19. Per Capita Consumption of Retirees	33
20. Discounted Lifetime Utility	33
References	35

I. INTRODUCTION

This paper explores the effects of changes in the age distribution on returns to financial assets. The aging of the baby boomers and speculation over possible effects on financial markets has raised the profile of this issue, both in the financial press and in academic circles.² Broadly speaking, there are two opposing views. The first holds that retiring baby boomers will be selling their assets to a smaller generation of young investors. This will drive asset prices down, leaving baby boomers with a smaller nest-egg than anticipated. The second maintains that forward-looking financial markets are pricing assets to incorporate the aging of the boomer generation. As a result, there will not be a market meltdown when the baby boomers retire. This paper bridges the gap between these opposing arguments by asking the following question: can demographic change, which is slow-moving and predictable, have a significant impact on financial markets that are rational and forward-looking? It presents a model in which rational, forward-looking agents of different ages trade in financial assets, and uses this framework to simulate a baby boom-baby bust of the kind observed in many developed countries over the post-war period. The main finding of the paper is that changes in the age distribution have significant effects on asset returns, even when investors are rational and forward-looking, and that these effects have important implications for the welfare of baby boomers and surrounding cohorts.

The model used is a stationary overlapping generations model with two assets: shares of ownership in risky capital and a riskless one-period bond that is in zero net supply. The representative agent lives for four periods: childhood, young working-age, old working-age, and retirement. In childhood, the agent relies on her parent for consumption and is not a decision maker. In both working-age periods she supplies labor inelastically and earns a wage. In retirement she consumes down her savings, there being no bequests. The model features only aggregate uncertainty, a technology shock to production and random population growth, and is solved numerically using the parameterized expectations approach. Although agents' degree of risk aversion is constant over time, they invest as if increasingly risk averse with age: young workers short the riskless asset in order to hold equity, while old workers hold mostly the riskless asset. This portfolio behavior stems from the risk and life cycle characteristics of a nontradable asset, human capital, which agents implicitly hold. Young workers anticipate receiving wage income in old working-age, so that next period consumption does not depend on savings alone. In addition, since the return on capital is positively but imperfectly correlated with wage income, equity is an attractive investment because it will diversify the effects of an adverse technology shock. In contrast, old workers' investment decision reflects the fact that next period consumption is out of savings alone. As a result, they largely eliminate consumption risk by investing mostly in the riskless asset.

Using the model to simulate a baby boom-baby bust yields the following effects. First, there is an aggregate saving effect on asset returns as changes in the age distribution affect aggregate saving and therefore the real interest rate. During the boom this effect will push up returns on capital and the riskless asset, a result of higher aggregate consumption

² For examples of articles in the financial press, see Passell (1996) and Colvin (1997).

because of relatively large cohorts of children. During the bust aggregate saving is relatively high, pushing returns on both assets down. Second, the return differential between stocks and bonds changes over the demographic shift. This effect derives from the fact that agents shift from stocks to bonds as they age. At the turning point of the boom-bust, when a large boomer cohort of old workers trades with a smaller cohort of young investors, this investment behavior generates excess demand for the riskless asset. As a result, the bond return falls sharply relative to the return on capital, and the return differential rises. Third, wage income moves inversely with the size of the labor force, even though capital accumulation in the model is endogenous.

The quantitative effects of the boom-bust can be summarized as follows. During the baby boom the expected one-period returns on capital and the riskfree asset rise above their steady states by 3.2 and 7.8 percent respectively. On an annualized basis this means that the expected return on capital rises above its steady state by ten basis points, while the riskfree rate is 23 basis points higher. In contrast, during the baby bust the expected one-period return on capital falls below its steady state by up to 3.8 percent, while the riskfree rate falls by up to 10.6 percent relative to its steady state. This translates respectively into up to ten and 24 basis points on an annual basis. The greater sensitivity of the riskfree rate to the boom-bust is significant for older investors who want to minimize consumption risk in retirement. These magnitudes are also substantial relative to the impact of other fundamentals on asset returns. And while they are small relative to the recent run-up in stock indices, this simulation exercise holds non-demographic fundamentals constant over the boom-bust and ignores the possibility of a speculative bubble.

These effects go against baby boomers, especially those in the tail-end of the baby boom. But are baby boomers worse off? The simulated boom-bust consists of two boom followed by two bust periods, so that the first boomer cohort has relatively more children than the second. This reduces consumption per head of the first boomer cohort in parenthood below the steady state. If utility of young workers is additively separable in their consumption and that of their children, this effect dominates adverse asset market effects. This means that the lifetime utility of the first boomer cohort is below steady state, while that of the second is above because positive consumption effects as a parent more than offset the effects of lower returns on retirement savings. In other words, adverse asset market effects are second-order. The fact that raising children is costly is more important. This result is reversed if young workers' utility is defined over household consumption. This specification reduces the effective weight of consumption when young, so that adverse asset market effects dominate earlier positive consumption effects. Since the asset market implications of the model are qualitatively unchanged across specifications, and since it is not obvious how to model parents' utility, the reversal of the welfare result is interesting. It also points to a deficiency of the model, since humans derive utility not only from consumption, but also from having children. As such, the welfare implications of the model should be viewed as incomplete.

The result that the return differential between stocks and bonds shoots up at the turning point of the boom-bust is arguably the most interesting result of the paper. As noted above, it derives from two characteristics of the model: the shift from stocks to bonds over the life cycle, and the limited number of agents trading at any one point in time. A critic of

this result might say: if baby boomers know they will retire in roughly 10 years, and that they may face adverse asset markets at that point, why not lock in wealth ahead of time? But this is exactly what boomer cohorts do in the model. In old working-age they invest in the riskless asset, in order to reduce consumption risk down the road. It is precisely this behavior, and the demographic imbalance, which drives down the bond relative to the stock return.

A number of papers have recently noted that a risk associated with switching to individual retirement accounts is that investors may be subject to adverse movements in asset prices that persist over time.³ This paper makes this argument explicit, by presenting a framework in which investors are subject to cohort-specific risk that is linked to demographic change. Furthermore, the model corresponds to a world in which pay-as-you-go pension systems have been replaced with individual retirement accounts, with no government regulation over agents' portfolio decision. Thus the key result of the paper, from a policy point of view, is that the historical distribution of asset returns may be inappropriate for computing the gains to investors from switching to individual accounts, since the baby boomers are only now approaching retirement. Indeed the model suggests that baby boomers will earn returns that are substantially below returns to previous generations.⁴ This argument should not be understood as an endorsement of pay-as-you-go systems, since a defined-benefit pension system can be fully funded. Indeed, the paper augments the model with a simple pay-as-you-go pension scheme and shows that such a system does not eliminate cohort-specific risk that comes from demographic change. Instead the paper underlines the role of government as an infinitely-lived agent, one that can insure agents against cohort-specific risk by adjusting government borrowing over time to stabilize the riskfree rate. Of course such a move would reflect a political consensus, since it involves transfers of wealth across generations.

There are a number of objections to the approach in this paper. It presents a closed economy model, which ignores investors' ability to insure against cohort-specific risk by holding an internationally diversified portfolio. Though this is an important consideration, most countries with significant asset markets have experienced post-war demographic shifts similar to the US. In effect, the model should be thought of as representing the developed world as a whole. Of course, there are regions with very different age distributions, such as Africa or parts of Asia. However, it is unlikely that assets in these markets have risk-return characteristics attractive to baby boomers preparing for retirement.⁵

³ See, for example, Heller (1998) and Hemming (1998).

⁴ For papers that use the historical distribution of asset returns to compute gains from switching to individual retirement accounts see, for example, MaCurdy and Shoven (1992) and Feldstein and Rangelova (1998).

⁵ For a discussion of home country bias in portfolio allocation, see French and Poterba (1991).

The model also ignores bequests. Bergantino (1998) finds that, using data from the Survey of Consumer Finances, intergenerational transfers are of minor importance to households. 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 per spouse. This evidence suggests that intergenerational transfers are of minor importance to most households, especially relative to wage income, in determining the life-cycle path of asset holdings.

A further shortcoming of the model is that it does not replicate the equity premium observed in the data. In large part this is because what is called equity in the model is not a levered asset, in the sense that there is no corporate debt. Adjusting for this, the model supports a Sharpe ratio of roughly ten percent, the same order of magnitude as Storesletten, Telmer, and Yaron (1997) whose overlapping generations model has agent-specific, persistent income shocks. In contrast, Constantinides, Donaldson, and Mehra (1998) generate an equity premium in an overlapping generations model in which young investors are borrowing constrained, which underscores that this model features no market imperfections that might generate an equity premium. The implicit assumption is that the behavior of relative returns over a demographic shift can be adequately characterized in the absence of such features.⁶

The next section presents a brief overview of related papers. Section three presents the model, with subsequent sections devoted to equilibrium conditions and the solution method. Section eight discusses calibration of the model, while section nine characterizes the solution. Section ten uses the model solution to simulate the effects of a baby boom-baby bust on asset returns. Section 11 concludes.

II. THE PAPER VIS-À-VIS THE LITERATURE

The rationale for agents with constant risk aversion to substitute from equity to bonds as they age has been previously explored by Jaganathan and Kocherlakota (1996).⁷ They make the point that investors have fewer working years ahead of them as they age. Assuming that most investors' labor incomes are poorly correlated with stock returns, they demonstrate that it is rational for agents to shift the composition of their financial wealth from stocks to less risky assets as they grow older. This paper demonstrates that this behavior obtains even

⁶ For a review of recent papers on the equity premium, see Kocherlakota (1996).

⁷ A more recent discussion of the role of labor income as a non-traded asset and its impact on investment behavior over the life cycle can be found in Campbell, Cocco, Gomes, and Maenhout (1999).

when labor income and stock returns are positively correlated, a result of the interaction of the life-cycle features of the model with the risk-return characteristics of its assets.⁸

A number of empirical papers have recently explored the link between changes in the age distribution and financial markets. Bergantino (1998) presents evidence linking the level of real stock prices in the US to time series for aggregate demand of financial assets, which are derived from Survey of Consumer Finances data on portfolio composition by age. Brooks (1998) finds that real stock and bond prices across developed countries are positively related to the share of the population that is middle-aged, using the cross-section dimension of the data to control for unobserved fundamentals. In contrast, Poterba (1998) fails to find a significant relationship between returns on a range of assets and different measures of the age distribution. However, his focus on asset returns rather than prices could be interpreted as effectively measuring the high frequency correlation between demographic change and returns, when intuition would suggest that the relationship should be strongest at low frequencies. In essence this is the problem that bedevils the empirical analysis to date, since the effective number of observations is so small that statistical tests have limited power. Perhaps this is the strongest argument for the simulation exercise that follows.

III. THE MODEL

The representative agent lives for four periods: childhood, young working-age, old working-age, and retirement. In childhood the agent makes no decisions of her own, with consumption, c_t^0 , determined by the parent, the next older cohort. In young working-age the agent inelastically supplies one unit of labor and earns a wage w_t . Out of wage income, she consumes c_t^1 for herself, and assigns $(1+n_t)c_t^0$ to her offspring, where n_t is the period t population growth rate. In addition, she may hold shares of ownership in risky capital, s_{et}^1 , and invest s_{bt}^1 in a riskless, one-period bond, which is in zero net supply. The budget constraint of a period t young worker, born in period $t-1$, is therefore:

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

The agent born in period $t-1$ reaches old working-age in period $t+1$. She again supplies one unit of labor inelastically, earning w_{t+1} , and receives income from stock and bond holdings chosen in the previous period. Out of total income she consumes only for herself, since her children, having entered young working-age, are now self-sufficient. She consumes c_{t+1}^2 , invests s_{et+1}^2 in risky capital, and s_{bt+1}^2 in the safe asset. The budget constraint of a period $t+1$ old worker is:

⁸ The approach in this paper contrasts with Bakshi and Chen (1994) who hypothesize that agents become more risk averse with age. They find that the average age of the US population is positively correlated with future excess returns on stocks over treasury bills.

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

Here r_{et+1} is the return on equity held from period t into period $t+1$, while r_{ft} is the return on the riskless asset, held from period t into period $t+1$. The agent born in period $t-1$ reaches retirement in period $t+2$. In retirement she receives no wage income, and since there are no bequests, she consumes down her savings. The budget constraint of a period $t+2$ retiree is thus:

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

Preferences are described by an additively separable utility function. The expected lifetime utility of a period t young worker is given by:

$$V_t = (1 + n_t) \frac{(c_t^0)^{1-\theta}}{1-\theta} + \frac{(c_t^1)^{1-\theta}}{1-\theta} + \beta E_t \left[\frac{(c_{t+1}^2)^{1-\theta}}{1-\theta} \right] + \beta^2 E_t \left[\frac{(c_{t+2}^3)^{1-\theta}}{1-\theta} \right] \quad (4)$$

The subjective discount factor is given by β , where $0 < \beta < 1$, and θ is the coefficient of *constant* relative risk aversion, so that the representative agent does not become more risk averse with age. This rather conventional specification of preferences is adopted to focus attention on the interaction of the life-cycle dimension of the model with the representative agent's investment decision. 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 according to $N_t = (1 + n_t)N_{t-1}$, where n_t is the realization of a stationary population growth shock. Table 1 describes the evolution of the age distribution over time:

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}
↓	↓	↓	↓	↓

Output is generated according to a Cobb-Douglas production technology:

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

where K_{t-1} is generated in period $t-1$ by the investment decisions of young and old workers. A_t is the realization of a stationary, labor-augmenting technology shock. L_t consists of young and old workers, so that $L_t = N_{t-1} + N_{t-2}$. The two factors of production are rewarded their marginal products.

$$r_{ot} = \alpha K_{t-1}^{\alpha-1} (A_t L_t)^{1-\alpha} - \delta \quad (6)$$

$$w_t = (1 - \alpha) K_{t-1}^{\alpha} A_t^{1-\alpha} L_t^{-\alpha} \quad (7)$$

where δ is the depreciation rate. 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 :

$$K_t = N_{t-1} s_{ot}^1 + N_{t-2} s_{ot}^2 \quad (8)$$

The equilibrium condition that the riskless, one period bond is in zero net supply implies that bond holdings of young and old workers sum to zero. In other words, either young workers borrow from their parents, or they lend to them.

$$0 = N_{t-1} s_{bt}^1 + N_{t-2} s_{bt}^2 \quad (9)$$

Imposing these equilibrium conditions and aggregating across budget constraints of young workers, old workers, and retirees in period t yields the social resource constraint:

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

Output in period t is divided into aggregate consumption and gross investment.⁹

⁹ The model has two sources of aggregate uncertainty: the technology shock to production, and the population growth rate. An investor deciding on how much equity to hold going forward into period $t+1$ does not know the realization of the period $t+1$ technology shock that determines r_{ot+1} . In contrast the return on the riskless asset in period $t+1$ is known in period t . Hence it is denoted r_{ft} , to emphasize that it is in agents' period t information set. Similarly, though K_t enters production in period $t+1$, it is determined by investment decisions made in period t and is therefore also in agents' period t information set. The stochastic population growth rate represents aggregate uncertainty over cohort size. An unexpectedly large cohort of children has repercussions for agents' consumption-investment decision, as they react to a larger consumption requirement in young working-age and the larger labor force in subsequent periods.

IV. EQUILIBRIUM

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

$$c_t^0 = c_t^1 \quad (11)$$

$$(c_t^1)^{-\theta} = \beta E_t \left[(c_{t+1}^2)^{-\theta} (1 + r_{et+1}) \right] \quad (12)$$

$$(c_t^1)^{-\theta} = \beta (1 + r_{ft}) E_t \left[(c_{t+1}^2)^{-\theta} \right] \quad (13)$$

$$(1 + n_t) c_t^0 + c_t^1 + s_{et}^1 + s_{bt}^1 = w_t \quad (14)$$

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

$$(c_t^2)^{-\theta} = \beta E_t \left[(c_{t+1}^3)^{-\theta} (1 + r_{et+1}) \right] \quad (15)$$

$$(c_t^2)^{-\theta} = \beta (1 + r_{ft}) E_t \left[(c_{t+1}^3)^{-\theta} \right] \quad (16)$$

$$c_t^2 + s_{et}^2 + s_{bt}^2 = w_t + (1 + r_{et}) s_{et-1}^1 + (1 + r_{ft-1}) s_{bt-1}^1 \quad (17)$$

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:

$$c_t^3 = (1 + r_{et}) s_{et-1}^2 + (1 + r_{ft-1}) s_{bt-1}^2 \quad (18)$$

(11) through (18) 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 conditions (8) and (9), which are the equilibrium conditions for the stock and bond markets.

V. THE STATE VARIABLES

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

$$w_t^1 = w_t \quad (19)$$

Total wealth of old workers consists of wage income, in addition to stock and bond holdings and returns.

$$w_t^2 = w_t + (1 + r_{et})s_{et-1}^1 + (1 + r_{ft-1})s_{bt-1}^1 \quad (20)$$

w_t^1 and w_t^2 describe the distribution of wealth across the two decision making cohorts in period t 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. Thus the set of period t state variables that describes the model looking ahead to period $t+1$ is:

$$\Theta_t = [w_t^1, w_t^2, N_t, N_{t-1}, N_{t-2}] \quad (21)$$

VI. THE SOLUTION METHOD

The model is solved using the parameterized expectations approach (PEA). This approach is described in detail in Den Haan and Marcet (1990), who solve a one-good, stochastic growth model. More recently, Izvorski (1997) uses the PEA to solve a model with heterogeneous agents and incomplete markets. The essence of the PEA is the observation that the conditional expectations on the right-hand sides of (12), (13), (15), and (16) each represent a function $g: R_+^5 \rightarrow R_+$ of the state variables in Θ_t . This insight is used to substitute each conditional expectation with a function $\Pi(\Theta_t, \psi)$, where Π (the functional form) and ψ (the vector of parameters) will be chosen to make $\Pi(\Theta_t, \psi)$ as close as possible to g . Using the PEA (12), (13), (15), and (16) can be rewritten as:

$$(c_t^1)^{-\theta} = \beta \Psi(\Theta_t, \tau) \quad (22)$$

$$(c_t^1)^{-\theta} (s_{et}^1)^2 = \beta (1 + r_{et}) \Omega(\Theta_t, \gamma) \quad (23)$$

$$(c_t^2)^{-\theta} = \beta \Lambda(\Theta_t, \xi) \quad (24)$$

$$(c_t^2)^{-\theta} (s_{et}^2)^2 = \beta (1 + r_{ft}) \Gamma(\Theta_t, \omega) \quad (25)$$

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 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 (both sequences are drawn only once), and solves out the model for T periods. The PEA then turns to fitting the conditional expectations in (12), (13), (15), and (16), finding 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 .¹⁰

The particular version of the PEA implemented here proceeds to fit the conditional expectations in a step by step approach. For illustration, turning to the conditional expectation in (12), a non-linear least squares estimation for τ is performed:

$$\min_{\tau} \frac{1}{T-1} \sum_{t=1}^{T-1} \left[(c_{t+1}^2)^{-\theta} (1+r_{et+1}) - \Psi(\Theta_t, \tau) \right]^2 \quad (26)$$

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 .¹¹

VII. ACCURACY OF THE SOLUTION METHOD

The basic intuition underlying the PEA is to approximate conditional expectations of period $t+1$ using the information set available to agents in period t . For a successful approximation, the PEA prediction error should therefore be orthogonal to agents' information set at t . It is this intuition that 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.

$$\begin{bmatrix} (c_{t+1}^2)^{-\theta} (1+r_{et+1}) - \Psi(\Theta_t, \tau^*) \\ (c_{t+1}^2)^{-\theta} (s_{et}^1)^2 - \Omega(\Theta_t, \gamma^*) \\ (c_{t+1}^3)^{-\theta} (1+r_{et+1}) - \Lambda(\Theta_t, \xi^*) \\ (c_{t+1}^3)^{-\theta} (s_{et}^2)^2 - \Gamma(\Theta_t, \omega^*) \end{bmatrix} = \varepsilon_{t+1} \quad (26)$$

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

¹⁰ Note that the conditional expectations in (13) and (16) have not been parameterized in the traditional manner. Both equations have been multiplied through 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.

¹¹ The implementation of the PEA follows Den Haan and Marcet (1990). Rather than performing a computationally expensive non-linear least squares estimation to find τ_e , it takes a first-order approximation of $\Psi(\Theta_t, \tau_n)$ around τ_n . Rearranging terms τ_e is then the coefficient vector in an OLS regression.

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

$$\left[\frac{\sum_{t=1}^{T-1} (\varepsilon_{t+1} \otimes v_t)}{T-1} \right] \quad (27)$$

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 .¹²

VIII. MODEL PARAMETERIZATION

The paper presents simulations results for four versions of the model. The first drops the riskfree asset from the analysis, focusing on the consumption-saving decision under uncertainty. The second adds the riskfree asset to the analysis, focusing on the full consumption-investment decision. Comparing these two specifications, the main question is: how does adding the riskfree asset change agent behavior and model characteristics. The third model augments the second specification with a simple pay-as-you-go pension system. This changes the representative agent's budget constraints in young working-age, old working-age, and retirement to:

$$(1+n_t)c_t^0 + c_t^1 + s_{et}^1 + s_{bt}^1 = (1-\pi)w_t \quad (28)$$

$$c_{t+1}^2 + s_{et+1}^2 + s_{bt+1}^2 = (1-\pi)w_{t+1} + (1+r_{et+1})s_{et}^1 + (1+r_{ft})s_{bt}^1 \quad (29)$$

$$c_{t+2}^3 = (1+r_{et+2})s_{et+1}^2 + (1+r_{ft+1})s_{bt+1}^2 + \left(\frac{N_{t+1} + N_t}{N_{t-1}} \right) \pi w_{t+2} \quad (30)$$

¹² The accuracy 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.

where π is the payroll tax levied on young and old workers and the last term on the right-hand-side in (30) is the per capita retirement benefit. It is worth noting that the retirement benefit is not a riskfree source of income, since it loads on the technology shock via the wage. In effect this extension gives retirees access to wage income, raising the question whether their investment behavior changes as a result. It is also of interest whether this form of pension system protects investors from cohort-specific demographic risk. The fourth specification is based on a modification of the utility function, which in the baseline framework is additively separable in c_t^0 and c_t^1 . In effect, this says that young workers care about their consumption and that of their children in per capita terms, rather than in terms of household consumption. Since it is not immediately obvious whether agents' utility is separable in their consumption and that of their children, an alternative specification is explored, one in which young workers care about household consumption.

$$V_t = \frac{((2+n_t)c_t^1)^{1-\theta}}{1-\theta} + \beta E_t \left[\frac{(c_{t+1}^2)^{1-\theta}}{1-\theta} \right] + \beta^2 E_t \left[\frac{(c_{t+2}^3)^{1-\theta}}{1-\theta} \right] \quad (4')$$

Since (4') leaves open how to allocate household consumption between young workers and their children, it is assumed for simplicity that $c_t^0 = c_t^1$. With preferences represented by (4'), (11) drops away, while (12) and (13) are replaced by, respectively

$$((2+n_t)c_t^1)^{-\theta} = \beta E_t \left[(c_{t+1}^2)^{-\theta} (1+r_{et+1}) \right] \quad (31)$$

$$((2+n_t)c_t^1)^{-\theta} = \beta(1+r_{\beta}) E_t \left[(c_{t+1}^2)^{-\theta} \right] \quad (32)$$

Starting values for w_t^1 and w_t^2 , and for τ , γ , ξ , and ω , are chosen based on a version of the model under certainty, in which the representative agent has perfect foresight, and makes only a consumption-saving decision (the portfolio problem drops away).¹³ In that framework, given the preference and production parameters, steady state values of the choice variables and factor returns are given in Table 2.¹⁴

¹³ Preference and production parameters are chosen to reflect the fact that each period corresponds to roughly twenty years. β is set to 0.6, while θ is set to one. Period utility therefore takes the form $u(c) = \ln(c)$. On the production side, the share of output that goes to capital, α , is set to 0.3, while the rate of depreciation of capital, δ , is taken to be 0.4. π is set at 0.2 in the specification with social security. In the stochastic simulations the stationary technology shock follows $\ln A_t = \phi \ln A_{t-1} + \varepsilon_t$ where $\phi = 0$, $\varepsilon_t \sim N(0, \sigma_\varepsilon)$ and $\sigma_\varepsilon = 0.1$. The period t cohort of children is generated according to $\ln N_t = \rho \ln N_{t-1} + v_t$, with $v_t \sim N(0, \sigma_v)$, $\rho = 0.99$, and $\sigma_v = 0.01$. Period t population growth is then backed out according to $N_t = (1+n_t)N_{t-1}$.

¹⁴ In the perfect foresight case, the model has only one state variable: the capital-labor ratio. Imposing the equilibrium condition that net saving equal net investment, a non-linear equation solver, written by Christopher Sims for Matlab, is used to solve for the steady state capital-labor ratio.

Table 2: Steady State Values for the Model without Uncertainty

	$c_0 = c_1$	c_2	c_3	s_e^1	s_e^2	K	r_e	w^1	w^2
(4)	0.1531	0.2243	0.3287	0.0153	0.1346	0.1496	1.4420	0.3216	0.3589
(4SS)	0.1048	0.1950	0.3628	0.0161	0.0806	0.0967	2.1009	0.2257	0.2756
(4')	0.1539	0.3312	0.3564	0.0795	0.1987	0.2783	0.7932	0.3874	0.5300

With agents' preferences described by (4), young workers consume a large fraction of their wage income, saving under five percent. Having to raise children means they postpone saving largely until old age. While the wage is the only source of income for young workers, it makes up just under 90 percent of total income for old workers, with the remainder generated by savings from young working-age. Retirement income is generated purely out of savings, with per capita retirement consumption well above that of young and old workers. Adding pay-as-you-go social security to the model reduces the need to save for retirement, so that investment in capital by young and old workers falls. As a result, the steady state capital stock declines, pushing the return on capital up and the wage down (the endogenous state variables w^1 and w^2 reflect after-tax income). The steady state retirement benefit amounts to 0.1129, corresponding to a 40 percent replacement rate, with the steady state return on payroll contributions zero as the model is stationary. Switching to (4') to model preferences generates substantial capital deepening. This is because the savings rate of young workers increases substantially, to just above 20 percent, when they care about household rather than per capita consumption, due to the concavity of period utility. This capital deepening makes the representative agent better off, with lifetime income substantially above that when preferences are represented by (4).

IX. SIMULATION RESULTS

Table 3 characterizes the solution to the model based on (4) with only the consumption-saving problem, presenting the descriptive statistics of the model.¹⁵ These are generated using the coefficients on Θ_t at convergence, and simulating the model for the sequence of technology shocks and demographics used for the PEA.¹⁶ The descriptive statistics in Table 3 for the first-order approximation to conditional expectations use an approximation of the following form

¹⁵ In this specification the two first-order conditions relating to bond holdings fall away, leaving only the conditional expectations in (12) and (15) to be parameterized. The set of period t state variables remains Θ_t .

¹⁶ The values for τ , γ , ξ , and ω are taken to have converged when $(\tau_n - \tau_e)/\tau_n$, $(\gamma_n - \gamma_e)/\gamma_n$, $(\xi_n - \xi_e)/\xi_n$, and $(\omega_n - \omega_e)/\omega_n$ are each less than 0.00001.

$$\Psi(\Theta_t, \tau) = \exp(\tau_0 + \tau_1 w_t^1 + \tau_2 w_t^2 + \tau_3 N_t + \tau_4 N_{t-1} + \tau_5 N_{t-2}) \quad (33)$$

in the case of (12), for example. In other words, the algorithm is estimating six coefficients per Euler equation. Table 3 also reports simulation results based on a second-order approximation, which is based on a reduced tensor basis that omits terms of the third order and higher. In that case, the algorithm estimates 21 coefficients per parameterized expectation. Finally, table 3 reports simulation results based on a third-order approximation, again based on a reduced tensor basis that omits all terms of the fourth order and higher.¹⁷

Table 3: Descriptive Statistics of the Consumption-Saving Model

	First-Order		Second-Order		Third-Order	
	μ	σ	μ	σ	μ	σ
c^1	0.1527	0.0116	0.1527	0.0114	0.1527	0.0114
c^2	0.2237	0.0177	0.2237	0.0173	0.2237	0.0173
c^3	0.3273	0.0214	0.3274	0.0214	0.3274	0.0214
s_e^1	0.0153	0.0020	0.0153	0.0020	0.0153	0.0020
s_e^2	0.1342	0.0101	0.1342	0.0104	0.1342	0.0104
r_e	1.4472	0.1721	1.4470	0.1716	1.4471	0.1716
w	0.3207	0.0246	0.3207	0.0247	0.3207	0.0246
w^2	0.3579	0.0278	0.3580	0.0278	0.3580	0.0277
K	0.1519	0.0136	0.1520	0.0140	0.1521	0.0141
L	2.0326	0.0892	2.0326	0.0892	2.0326	0.0892
<i>Score</i>	6.9%		5.6%		5.4%	

The results in Table 3 do not differ perceptibly from the steady state values of the perfect foresight framework. As in that setting, young workers consume a large fraction of their wage income, saving just under five percent on average. As a result, saving for retirement is postponed until old working-age, when agents save just under forty percent of total income. Second, the linear approximation to conditional expectations performs well, with model characteristics barely changing when moving to higher order approximations. The score statistic, the percentage of G statistics for 500 simulations that are above the critical value of the 95th percentile of a χ_{10}^2 distribution, indicates that the first order approximation is successful in fitting expectations, and that gains from higher order approximations are small. The formation of expectations appears to be largely linear in Θ_t . Table 4 reports the correlations matrix associated with the third-order approximation results reported in table 3:

¹⁷ To generate the PEA solution to the model the sequences for technology and the age distribution are drawn for length $T=1000$. The same draws are used across specifications for easier comparison. The results remain qualitatively unchanged for sequences of much greater length.

Table 4: Correlations based on Third-Order Approximation

	c^1	c^2	c^3	r_e	w	w^2
c^1	.	0.9943	0.9684	0.5235	0.9977	0.9943
c^2	.	.	0.9827	0.4687	0.9951	0.9999
c^3	.	.	.	0.3433	0.9688	0.9827
r_e	0.5351	0.4686
w	0.9951
w^2

Table 4 shows that the correlation between wage income and returns on capital is positive as expected, given the Cobb-Douglas production technology. Yet the degree of correlation is far from perfect. Young workers would like to hold more equity because the low degree of correlation offers a measure of diversification against technology shocks. However their ability to accumulate capital is limited, so that the correlation between wage income and total income to old workers is almost perfect.

Table 5 reports the descriptive statistics for the full consumption-investment problem. The first, second, and third-order approximations to the conditional expectations correspond to those discussed above.

Table 5: Descriptive Statistics of the Consumption-Investment Model

	First-Order		Second-Order		Third-Order	
	μ	σ	μ	σ	μ	σ
c^1	0.1529	0.0117	0.1532	0.0129	0.1531	0.0127
c^2	0.2248	0.0274	0.2250	0.0280	0.2248	0.0268
c^3	0.3271	0.0242	0.3271	0.0261	0.3270	0.0218
s_e^1	0.0949	0.0063	0.1449	0.0170	0.1413	0.0167
s_e^2	0.0556	0.0137	0.0065	0.0047	0.0096	0.0044
s_b^1	-0.0795	0.0052	-0.1296	0.0179	-0.1260	0.0168
s_b^2	0.0795	0.0048	0.1296	0.0178	0.1260	0.0168
r_e	1.4408	0.1820	1.4375	0.1948	1.4410	0.1938
r_f	1.4198	0.2010	1.4186	0.1727	1.4268	0.1481
w	0.3213	0.0251	0.3217	0.0258	0.3214	0.0256
w^2	0.3598	0.0391	0.3611	0.0454	0.3604	0.0431
$r_e - r_f$	0.0210	0.2459	0.0190	0.1722	0.0142	0.1417
K	0.1529	0.0150	0.1538	0.0182	0.1533	0.0177
L	2.0326	0.0892	2.0326	0.0892	2.0326	0.0892
<i>Sharpe</i>	0.0854		0.1103		0.1002	
<i>Score</i>	39.7%		7.8%		5.9%	

Comparing tables 3 and 5 yields a number of insights. First, introducing the riskless asset changes the pattern of asset accumulation in a fundamental way. Without the riskless asset, young workers hold almost no equity because their consumption needs are large while

raising children. Once the riskless asset is introduced, young workers short the bond in order to invest in equity. Equity holdings of young workers rise to above forty percent of their wage income (for second- and third-order approximations to conditional expectations). Old workers lend to young workers because this gives them a safe income stream in retirement. By deciding to hold almost no equity, they largely eliminate consumption risk in retirement.

Second, with the introduction of the riskless asset, agents act as if they are increasingly risk averse with age, although the coefficient of relative risk aversion is constant over the life-cycle. Young workers short the riskless asset, investing in risky equity, while old workers hold primarily the riskfree bond. This behavior comes from the interaction of the life-cycle features of the model with the risk-return characteristics of stocks and bonds. Young workers have an additional working period ahead of them, so that consumption in the next period does not depend entirely on savings. This makes investing in high yielding, yet risky equity attractive, especially since labor and interest income are imperfectly correlated. But since consumption requirements for young workers are high, they must borrow to hold shares of ownership in capital, generating supply of the riskless asset. In contrast, old workers' retirement consumption depends entirely on savings. Faced with the choice of investing in risky equity or safe bonds, they hold bonds almost exclusively, generating demand for the riskless asset. Minimizing consumption risk in retirement appears to dominate old workers' portfolio decision.

The annualized return on capital in the model is about 4.56%. Compared to the consumption-saving framework, the introduction of the riskless asset produces slight capital deepening, as old workers do not divest themselves entirely of shares of ownership in capital. The annualized return on the riskfree asset amounts to 4.53%, so that the model fails to generate an equity premium of the order observed historically in the US and other developed markets. One reason for this is that the shares of ownership in capital do not correspond to equity as it is commonly understood, since the shares in the model are not a levered asset, i.e. there are no corporate bonds. The return on capital going to shareholders in the model should be thought of as subsuming equity, corporate bonds, and other forms of lending and ownership. Though the implications for asset returns are not of the order observed in the data, the Sharpe ratio supported by the model is around ten percent, an order of magnitude comparable to Storesletten, Telmer, and Yaron (1997). This result is encouraging since the model features only aggregate uncertainty, while the model of Storesletten, Telmer, and Yaron includes persistent, agent-specific shocks. The Sharpe ratio benchmark associated with Mehra and Prescott (1985) is 37 percent.

With agents making a consumption-investment decision, the formation of expectations becomes non-linear in Θ_t . While the score statistic improves only marginally when going from the first to a second-order approximation in the absence of the riskless asset, it improves by a factor of five when going from the first to a second-order approximation, and differences in portfolio allocation between young and old workers become more pronounced. Looking to the third-order approximation, it appears that the key non-linearity in the formation of expectations is captured by second-order terms. Table 6 presents the correlations matrix associated with table 5:

Table 6: Correlations based on Third-Order Approximation using (4)

	c^1	c^2	c^3	r_f	r_e	w	b	$r_e - r_f$
c^1	.	0.9740	0.4186	-0.3532	0.3693	0.9964	0.9759	0.8742
c^2	.	.	0.2291	-0.2225	0.5191	0.9640	0.9994	0.9426
c^3	.	.	.	-0.7965	-0.6084	0.4665	0.2312	0.0001
r_f	0.6867	-0.3829	-0.2237	-0.1057
r_e	0.3297	0.5182	0.6503
w	0.9651	0.8512
w^2	0.9427
$r_e - r_f$

Table 6 illustrates the benefits to young workers of being able to borrow and hold equity. The correlation of consumption with wage income in old working-age falls relative to the model without the riskless asset. And while the correlation of retirement consumption with the return on capital is 0.3433 in the absence of the riskless asset, it is -0.6084 in the consumption-investment framework. Essentially, agents have eliminated their exposure to technology shocks by holding mostly the riskless asset. Table 6 also illustrates that the riskfree rate and the return on equity tend to move together, because effects related to changes in aggregate saving dominate effects that drive returns on both assets in opposite directions.

Table 7 reports the descriptive statistics of the consumption-investment model with pay-as-you-go social security. The first, second, and third-order approximations to the conditional expectations are as above.

Table 7: Descriptive Statistics of the Consumption-Investment Model with Social Security

	First-Order		Second-Order		Third-Order	
	μ	σ	μ	σ	μ	σ
c^1	0.1045	0.0081	0.1046	0.0088	0.1046	0.0088
c^2	0.1948	0.0192	0.1953	0.0230	0.1952	0.0235
c^3	0.3613	0.0239	0.3612	0.0239	0.3611	0.0241
s_e^1	0.0462	0.0040	0.0918	0.0105	0.0929	0.0118
s_e^2	0.0504	0.0072	0.0054	0.0022	0.0042	0.0030
s_b^1	-0.0301	0.0031	-0.0758	0.0104	-0.0769	0.0116
s_b^2	0.0301	0.0030	0.0758	0.0104	0.0769	0.0116
r_e	2.1097	0.2448	2.1060	0.2656	2.1080	0.2701
r_f	2.0905	0.2084	2.0850	0.1943	2.0824	0.2425
w	0.2813	0.0219	0.2816	0.0225	0.2815	0.0226
w^2	0.2753	0.0259	0.2765	0.0335	0.2764	0.0342
b	0.1125	0.0089	0.1127	0.0091	0.1126	0.0091
$r_e - r_f$	0.0192	0.2876	0.0210	0.1956	0.0256	0.2412
K	0.0982	0.0094	0.0987	0.0116	0.0986	0.0118
L	2.0326	0.0892	2.0326	0.0892	2.0326	0.0892
<i>Sharpe Score</i>	0.0668		0.1074		0.1061	
	29.8%		6.7%		5.5%	

Portfolio behavior of the representative agent does not change qualitatively with the introduction of the pay-as-you-go pension system. Investors still shift the composition of financial wealth from stocks to bonds as they grow older, possibly because the retirement benefit is not riskfree. If this is the case it underlines how important consumption risk in retirement is for old workers. Of course this finding is conditional on the nature of the retirement benefit. If the benefit were riskfree, there might be substantial repercussions for portfolio behavior. Table 8 presents the correlations matrix for the model with social security:

Table 8: Correlations based on Third-Order Approximation

	c^1	c^2	c^3	r_f	R_e	w	b	$r_e - r_f$
c^1	.	0.9648	0.6852	-0.3444	0.3683	0.9960	0.9858	0.8423
c^2	.	.	0.4925	-0.1883	0.5479	0.9478	0.9420	0.9311
c^3	.	.	.	-0.7278	-0.3399	0.7333	0.7434	0.2616
r_f	0.6786	-0.3852	-0.3640	-0.0720
r_e	0.3160	0.3252	0.8032
w	0.9899	0.8119
b	0.8032
$r_e - r_f$


The per capita retirement benefit is negatively correlated with the return on the riskless asset, implying that in principle it might help insure agents against cohort specific risk in retirement. However the return on the retirement benefit loads on two factors, the technology shock and the number of workers per retiree. As a result, it is possible that baby boomers who have few children may in fact receive a low per capita benefit, because the workforce is relatively small when they retire. Adding social security leaves the risk-return characteristics of the model broadly unchanged, the smaller capital stock aside. The Sharpe ratio supported by the model remains around ten percent. As above, the non-linearity in expectations appears to be captured by the second-order approximation, as the score statistic improves only marginally when adding third-order terms.

Table 9 reports the descriptive statistics of the consumption-investment model with young workers deriving utility from household consumption. The first, second, and third-order approximations to conditional expectations are as above.

Across specifications a number of stylized facts emerge. With the introduction of the riskless asset, agents short the riskless asset when young in order to invest in risky equity. Old workers are willing lenders because they want to minimize consumption risk in retirement. In short, the investment horizon matters in the model, with agents making portfolio decisions as if they are increasingly risk averse with age. This feature of the model is robust to changes in preferences or adding a pay-as-you-go pension scheme. In addition, the risk-return characteristics of the model are stable across specifications.

X. SIMULATING A BABY BOOM-BABY BUST


This section simulates the effects of a baby boom-baby bust on factor returns, the riskfree rate, investor behavior, and welfare. The simulations are based on the PEA solution, using third-order approximations to parameterize conditional expectations. Rather than allowing the technology shock to vary randomly, as is the case when the model is solved, the technology shock is now held fixed at one (its mean value), essentially holding all non-demographic fundamentals constant. The baby boom-baby bust lasts four periods, beginning in period t when population growth (or more precisely, growth relative to the parent cohort) jumps to two percent, up from zero in the steady state. In period $t+1$ population growth is two percent again, switching to -2 percent in periods $t+2$ and $t+3$. From period $t+4$ population growth returns to zero. Figure 1 plots population growth over time.¹⁸



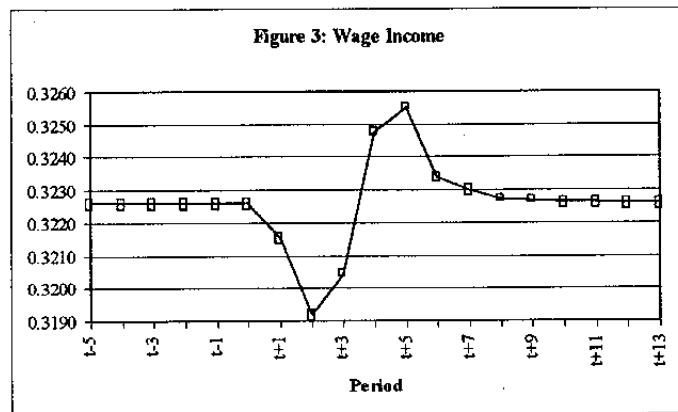
Beginning with the baseline model summarized in Table 5, Figure 2 presents returns on capital and the riskless asset over the demographic shift (these are period returns over roughly 20 years). Returns on capital and the riskless asset move broadly together, rising during the early stages of the boom when relatively large cohorts of children push up

¹⁸ The simulated boom-bust is similar to the post-war demographic transition in many developed countries (see Brooks (1998)). The results below are qualitatively unchanged for other forms of demographic shift, a simple baby boom for example.

aggregate consumption, reducing aggregate saving. Once boomer cohorts enter the workforce, aggregate saving rises, an effect that is compounded by the fact that the boom turns to bust and consumption of children becomes relatively less important.




Thus Figure 2 illustrates that demographic shifts lead to changes in aggregate saving over time, causing the real interest rate to vary. As the real interest rate changes, it pushes returns on stocks and bonds in the same direction. The movement in returns on the risky asset are mirrored in wage income, which moves inversely with the size of the labor force. This reflects changes in the capital-labor ratio, which falls during the boom stage of the transition and rises during the bust, even though capital formation is endogenous to the model.¹⁹ Figure 3 presents wage income over the demographic transition, illustrating that boomer cohorts receive lower wage income than in equilibrium.

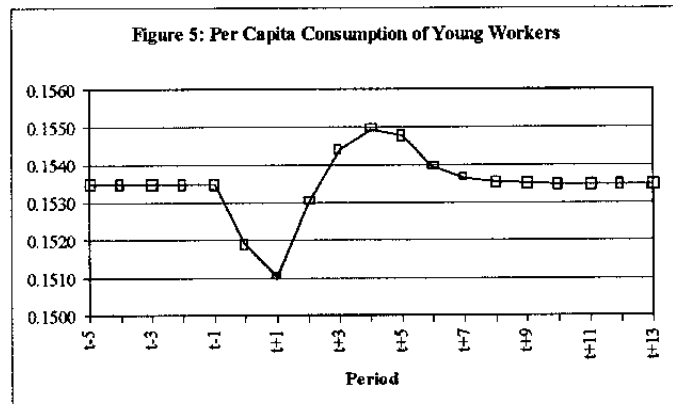


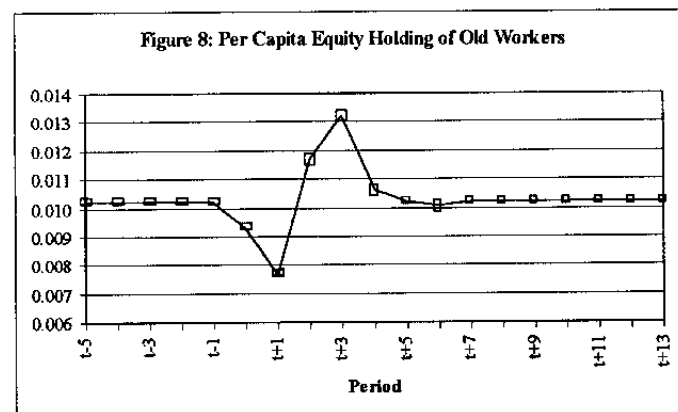

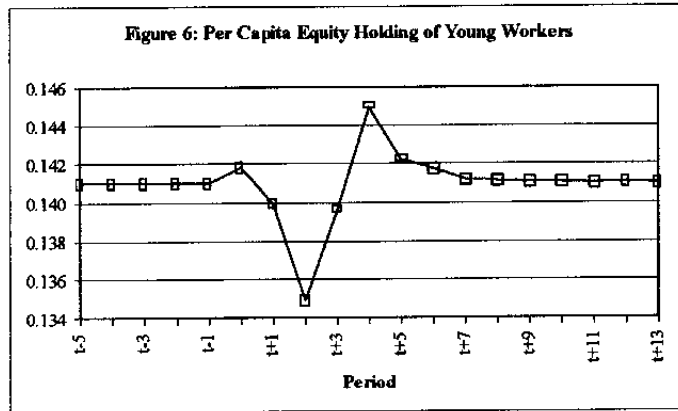
¹⁹ Of course this result depends on the parameter choice for the elasticity of intertemporal substitution.

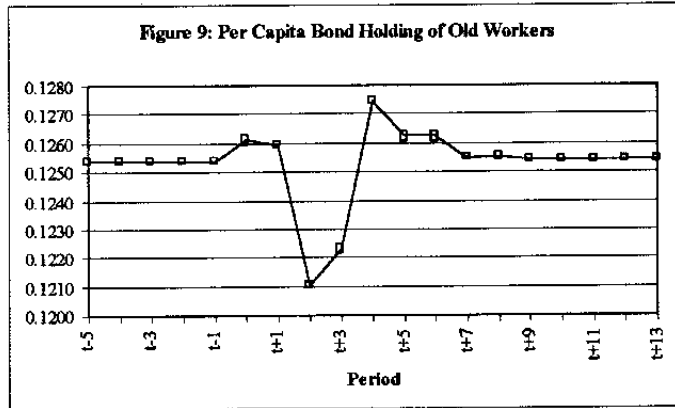
Figure 4 plots the return differential over the demographic shift. It falls sharply in period $t+1$, then rises and peaks in period $t+4$, before returning to its stochastic steady state level towards period $t+7$. This path for the return differential stems from the interaction of agent behavior and the demographic shock. The boom begins in period t , when a large cohort of children is born. Factor returns and the riskfree rate in period t are unaffected, determined by decisions made in period $t-1$. As a result per capita consumption of young workers falls as they feed a relatively large cohort of children, as seen in Figure 5.



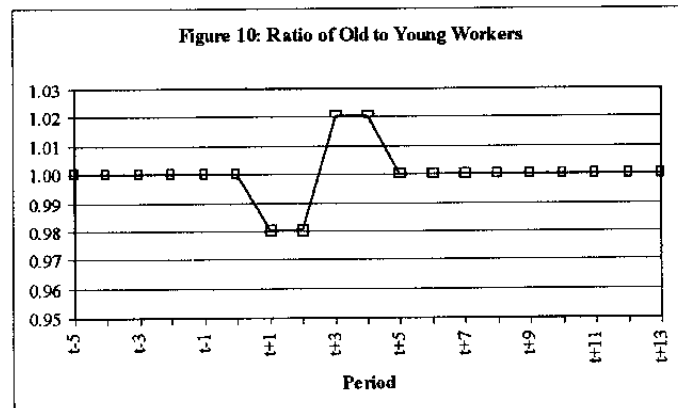
Portfolio behavior of young workers in period t reflects the fact that their household consumption requirement is large. In addition the expected return on capital in period $t+1$ is above its steady state level, as this is when the first boomer cohort enters the workforce. As a result young workers want to borrow more. Although old workers want to reduce their exposure to the technology shock in retirement, young workers' increased desire to borrow results in excess supply of the riskless asset, driving its return up relative to the return on capital. The equity premium in period $t+1$ turns negative.









Period $t+1$ brings investor behavior very similar to that in period t . Young workers again have a large cohort of children to feed, which reduces per capita consumption. But there is an important difference, which is illustrated in Figure 10.



Whereas in period t the ratio of old to young workers is in equilibrium, the cohort of young workers in period $t+1$ (the first boomer cohort) is larger than the cohort of old workers. This demographic imbalance exacerbates the effect of the increased desire to borrow by young workers, pushing the riskfree rate up further in period $t+2$ relative to the return on capital. The return differential thus falls even further. Investor behavior changes dramatically in period $t+2$, the first period of the baby bust. The cohort of young workers (the second boomer cohort) has a small number of children (the first cohort of the baby bust). As a result, per capita consumption of young workers shoots up in period $t+2$, as illustrated by Figure 5. The repercussions for portfolio behavior of young workers are significant. They scale back their investment in equity, partly because they are already well off in utility terms, but also because an increase in aggregate saving is having a downward effect on the return of capital. As a result, their borrowing in the riskfree asset falls significantly relative to the steady state, forcing old workers in period $t+2$ to hold more equity than they would like. This raises the return differential in period $t+3$, with excess demand for the riskfree asset pushing

down its return relative to the return on capital. Period $t+3$ brings the turning point of the demographic shift, as can be seen from Figure 10. It is in this period that the first baby bust cohort of young workers faces off against the last of the boomer cohorts in old working-age. Although equity and bond holdings of young workers in period $t+3$ are close to steady state levels, the demographic imbalance means that, in per capita terms, old workers are forced to hold more equity than they would like, since aggregate supply of the riskless asset is below demand. As a result the equity premium peaks in period $t+4$.

Figure 11 plots the repercussions of these asset market effects on consumption of retirees. Returns on both assets are just above their steady state levels in period $t+3$, but not enough so to offset other factors (low wage income and many children) that hamper the ability of the first boomer cohort to accumulate wealth. As a result, retirement consumption of the first boomer cohort is below the steady state. Adverse asset market effects have a strong impact on retirement consumption of the second boomer cohort. In period $t+4$ the decline in the riskfree rate pushes their retirement consumption substantially below that of their parents. The situation is reversed in utility terms, however. Figure 12 plots discounted lifetime utility, evaluated at young working-age, meaning that the discounted lifetime utility of the first boomer cohort, in young working-age in period $t+1$, is displayed in period $t+1$.




Although per capita consumption in retirement is greater for the first boomer cohort than for the second, it is worse off in utility terms. This is because the first boomer cohort is forced to scale back consumption per head in young working-age, because of its many children. In contrast, the second boomer cohort has relatively few children, so that its consumption per head in young working-age is near steady state, even though its wage income is significantly below steady state. This difference has important repercussions for lifetime utility, because (4) gives a greater weight to per capita consumption in young working-age than in later periods. As a result, even though the second boomer cohort has lower per capita retirement consumption than its parents, it is better off in utility terms because consumption per head in young working-age is higher. In other words, the asset market effects of the boom-bust are of second-order importance to agents' welfare. The fact that raising children is costly is far more important.²⁰

How big are the asset market effects of the boom-bust? The annualized expected return on capital in the steady state is 4.54 percent (see Figure 2). It peaks in period $t+2$ at 4.64 percent, and bottoms out at 4.45 percent in period $t+5$. The annualized steady state return on the riskfree asset is 4.52 percent. It peaks in period $t+2$ at 4.75 percent, and falls to 4.31 percent, its lowest level, in period $t+4$. These effects are significant, especially for older investors who want to minimize consumption risk in retirement. The cumulative return on the riskfree asset in period $t+4$, the return to the second boomer cohort, is seven percent below steady state. In contrast, the riskfree rate in period $t+2$, the return to parents of the first boomer cohort, is 7.8 percent above steady state.²¹ These effects are also significant relative to other sources of uncertainty. The largest change in the return on capital, the decline going from period $t+3$ to period $t+4$, amounts to 30 percent of the standard deviation of the period return (see Table 5). Meanwhile the largest change in the riskfree rate, the decline going from period $t+3$ to period $t+4$, amounts to 76 percent of the standard deviation of the period return.

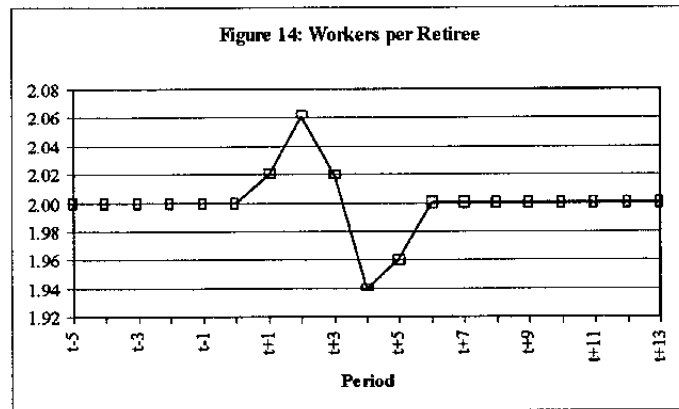
Adding the pay-as-you-go pension scheme to the model leaves its key features unchanged (see Table 7). In particular, agents continue to shift financial wealth from stocks to bonds as they age, even though they receive an income stream in retirement that is independent of savings. Figure 13 plots the wage over the boom-bust, for the model with social security.

²⁰ This result highlights an omission of the model, which does not model the decision to have children. If agents derive happiness from having children, this welfare result might well be reversed.

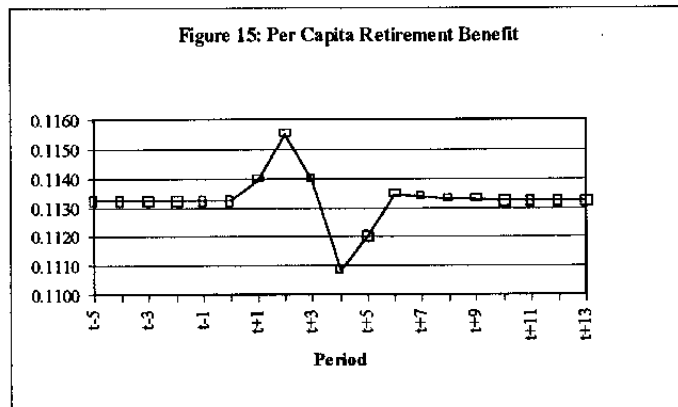
²¹ It is worth noting that the effect on stock returns is small compared to the run-up in stock indices over the past 20 years. The average real return on the Ibbotson Associates large stock index from 1979 – 1998 is 13.43 percent, relative to 3.53 percent for the period 1959 – 1978. In comparison, the period $t+2$ return on capital is only 3.8 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 ignores changes in other fundamentals (the technology shock is fixed at its mean value of one), or the possibility of a speculative bubble.



The wage is substantially above its steady state level in period $t+4$, in contrast to the riskfree rate. As such the pay-as-you-go pension system could insure agents against cohort specific asset market effects. However it fails to do so because the per capita retirement benefit depends on the number of workers per retiree, a measure that falls to its lowest level in period $t+4$, as seen in Figure 14.


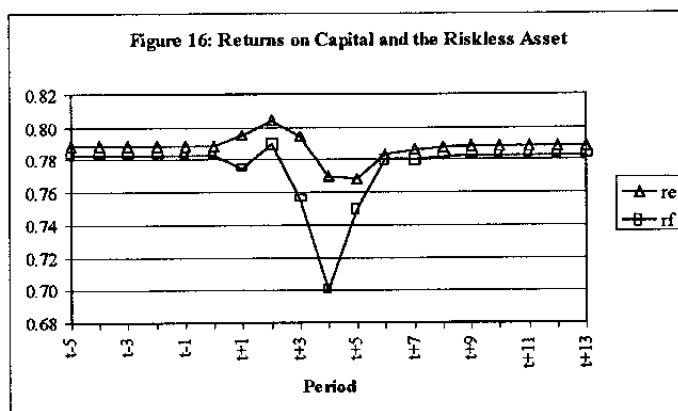


As a result the per capita retirement benefit falls to its lowest level in period $t+4$, since the decline in the number of workers per retiree far outweighs the effect of the higher wage. This can be seen in Figure 15.




This simple pay-as-you-go pension scheme thus fails to insure agents against cohort-specific asset market effects that are tied to demographics, because the same demographic imbalance that pushes down the return on the riskfree asset in period $t+4$ also reduces the per capita retirement benefit. Within this closed economy framework, the model thus highlights what is needed to insure agents against asset market effects of demographic shifts: an infinitely-lived agent who varies the supply of debt in order to smooth the path of the riskfree rate. In essence this would amount to making transfers of wealth across generations. A government acting in this way would reflect a political consensus that it should insure agents against cohort-specific risk.²²

Why insure agents against adverse asset market effects, if these effects are second-order? As noted above, using (4) to represent lifetime utility assigns a greater weight to per capita consumption in young working-age than later periods. Rewriting young workers' utility in terms of household consumption, using (4'), leaves the asset market implications of the model unchanged, while reversing the welfare result. Figure 16 plots returns on the risky and the riskfree assets, Figure 17 plots the wage, and Figure 18 plots the return differential over the boom-bust.



²² A different version of the pay-as-you-go pension system, in which the retirement benefit is exogenous and the payroll tax rate endogenous, might do a better job at insuring agents against cohort-specific risk.


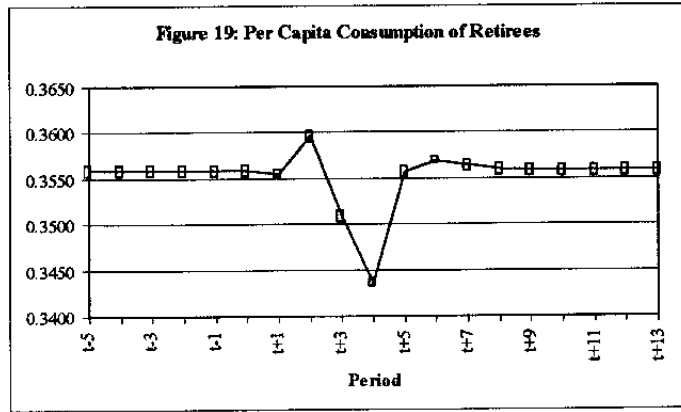


Simulating the boom-bust using (4') to model lifetime utility puts the earlier results in perspective. The key result, the decline in the riskfree rate relative to the return on capital as boom turns to bust, is stable across specifications. In period $t+3$ the last cohort of the baby boom is forced to trade with the first cohort of the baby bust. As above, this disparity in size results in excess demand for the riskless asset, pushing down the riskfree rate to clear the bond market. Figure 19 plots consumption of retirees, illustrating that the decline in the riskfree rate puts boomer cohorts at a distinct disadvantage as they go into retirement. Figure 20 plots discounted lifetime utility over the demographic shift, again from the perspective of young working-age. With young workers' utility defined over household consumption, boomer cohorts are now unambiguously worse off in utility terms than surrounding smaller cohorts. It is important to emphasize that investment behavior and model characteristics are stable whether young workers think in terms of per capita or household consumption. What changes is the importance of consumption in young working-age relative to consumption in later periods. With the reduced weight assigned to consumption in young working-age relative to later periods, asset market effects of demographic shifts become first-order. The reversal of the welfare result raises the interesting question of how to model parents' utility relative to that of their children.²³

The magnitude of the asset market effects is comparable across specifications. The annualized expected return on capital in the steady state, using (4') to model preferences, amounts to 2.95 percent. This return peaks in period $t+2$ at 2.99 percent, falling to a minimum of 2.89 percent in period $t+5$. The annualized return on the riskfree asset is 2.93 percent in the steady state. It peaks at 2.95 percent in period $t+2$, and reaches a low of 2.69 percent in period $t+4$. Again these effects are significant, particularly for older investors who want to minimize consumption risk in retirement. The cumulative return on the riskfree asset in period $t+4$, the return to the second boomer cohort, is 10.6 percent below the steady

²³ The reversal of the welfare result stems from the concavity of period utility, and would not obtain if period utility were linear in consumption, for example. As noted above, the model abstracts from agents' decision to have children. As such the discussion on the welfare effects of the boom-bust should be seen as incomplete.

state return. As above these effects are also significant relative to other sources of uncertainty. The largest change in the return on capital, the decline going from period $t+3$ to period $t+4$, amounts to 20 percent of the standard deviation of the period return (see Table 9). Meanwhile the largest change in the riskfree rate, the decline going from period $t+3$ to period $t+4$, amounts to 65 percent of the standard deviation of the period return.



XI. CONCLUSION

This paper demonstrates that changes in the age distribution have significant effects on financial markets. This result is consistent with rational, forward-looking behavior and derives from the fact that in an overlapping generations model there is only a limited number of cohorts that trade at any point in time, in the absence of an infinitely-lived agent. A government that varies the stock of debt over time to insure agents against cohort-specific risk in asset markets would make up for this market incompleteness.

To summarize, in a setting where investors are rational and forward-looking, changes in the age distribution have an aggregate saving effect on factor returns and the riskfree rate. In addition, since agents shift from stocks to bonds as they age, the return differential increases when boom turns to bust. Since in the model only two cohorts trade at any point in time, agents cannot diversify this risk away, even though it has a sizeable impact on retirement consumption. It is worth noting that a pay-as-you-go pension scheme fails to

insure agents against these cohort-specific asset market effects, since the same demographic effects that move returns against aging boomer cohorts also reduce the per capita retirement benefit. The model thus suggests that, to insure agents against these effects, government should vary the supply of debt over time to smooth the riskfree rate. In essence this would amount to making transfers of wealth across generations, reflecting a political consensus that insurance against the asset market effects of demographic shifts is worthwhile.

Within the context of the model the effects of demographic transitions on asset returns are non-trivial, both in absolute terms and relative to the effects of other fundamentals. And while the effects are small relative to the recent run-up in stock indices, the simulated baby boom-baby bust holds non-demographic fundamentals constant and ignores the possibility of a speculative bubble.

References

- 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
- Brooks, R. "Asset Market and Savings Effects of Demographic Transitions." Dissertation, Yale University, December 1998
- Campbell, J., Cocco, J., Gomes, F. and Maenhout, P. "Investing Retirement Wealth: A Life-Cycle Model." Mimeo, March 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
- 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
- Feldstein, M. and Rangelova, E. "Individual Risk and Intergenerational Risk Sharing in an Investment Based Social Security System." NBER Working Paper 6839, December 1998
- French, K. and Poterba, J. "Investor Diversification and International Equity Markets." *American Economic Review* 81 (1991): 222 - 226
- Heller, P. "Rethinking Public Pension Reform Initiatives." IMF Working Paper, WP/98/61, April 1998
- Hemming, R. "Should Public Pensions be Funded?" IMF Working Paper, WP/98/35, March 1998
- Ibbotson Associates. "Stocks, Bonds, Bills and Inflation: 1999 Yearbook." 1999
- Izvorski, I. "Asset Pricing with Heterogeneous Agents." Mimeo, 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

Kocherlakota, N. "The Equity Premium: It's still a Puzzle." *Journal of Economic Literature*, Vol. XXXIV (March 1996): 42 - 71

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

MaCurdy, T., and Shoven, J. "Stocks, Bonds and Pension Wealth." *Topics in the Economics of Aging*, NBER, 1992

Storesletten, K., Telmer, C. and Yaron, A. "Persistent Idiosyncratic Shocks and Incomplete Markets." Mimeo, June 1997