

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Risk Aspects of Investment-Based Social Security Reform

Volume Author/Editor: John Y. Campbell and Martin Feldstein, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-09255-0

Volume URL: <http://www.nber.org/books/camp01-1>

Publication Date: January 2001

Chapter Title: Investing Retirement Wealth: A Life-Cycle Model

Chapter Author: John Y. Campbell, João F. Cocco, Francisco J. Gomes, Pascal J. Maenhout

Chapter URL: <http://www.nber.org/chapters/c10600>

Chapter pages in book: (p. 439 - 482)

Investing Retirement Wealth A Life-Cycle Model

John Y. Campbell, João F. Cocco, Francisco J. Gomes,
and Pascal J. Maenhout

During the past few decades, American households have begun to display increasing financial sophistication and awareness of rates of return on alternative investments. At the same time, the implicit rate of return on contributions to the social security system has declined as the system has matured, and this rate of return is projected to decline further in the twenty-first century in response to unfavorable demographic trends (Geanakoplos, Mitchell, and Zeldes 1999). Not surprisingly, politicians and the public have become interested in the possibility of moving to a privatized system in which retirement contributions earn market-based rates of return.

Unfortunately, it is not straightforward to compare alternative retirement systems. Three important issues affect the comparison and invalidate simple rate-of-return calculations. First, the return on the current system is low in part because of the overhang of unfunded liabilities. Past generations have received a gift that must be paid off before the economy can enjoy the steady-state benefits of any new system. Second, capital income taxation puts a wedge between pretax and after-tax rates of return. Welfare calculations should take account of the tax revenue generated by capi-

John Y. Campbell is the Otto Eckstein Professor of Applied Economics at Harvard University and a research associate of the National Bureau of Economic Research. João F. Cocco is assistant professor of finance at London Business School. Francisco J. Gomes and Pascal J. Maenhout are graduate students in economics at Harvard University.

Campbell gratefully acknowledges the financial support of the National Science Foundation, Cocco of the Banco de Portugal, Cocco and Gomes of Fundação para a Ciência e Tecnologia, Portugal, and Maenhout of the Fund for Scientific Research Flanders. The authors are grateful for the helpful comments of David Cutler, David Weil, Amir Yaron, and other conference participants.

tal accumulation (in some systems, this tax revenue is forgiven, and private retirement accounts earn higher pretax rates of return). Third, returns on alternative financial assets can differ if these assets have different risk characteristics. A valid comparison of rates of return must adjust for risk.

This paper focuses on the last issue, the evaluation of alternative investments with different risk characteristics. From the point of view of households, the current social security system represents a defined-benefit pension plan in which income realizations through life are tied to annuity payments in retirement. This is similar to a system in which households are forced to accumulate a riskless asset in a retirement-savings account.

Some commentators have recently argued that households would be better off if their retirement savings were invested in risky assets such as equities that have a higher average return. This could be achieved within the present system if the social security trust funds were invested in equities; alternatively, within a privatized system, household retirement accounts could include equity investments.

If a household can borrow to invest in equities, however, then the accumulation of riskless assets within a social security account need not restrict the household's overall portfolio. The household can undo riskless social security accumulation by borrowing outside the retirement account; the household's overall portfolio can be made just as risky as if the retirement account itself were invested in equities. Thus, any claimed benefits for a change in the social security system must depend on the presence of portfolio constraints that prevent this sort of asset transformation.

Two different sorts of constraints are potentially relevant. First, a household may be unable to borrow at the riskless interest rate to finance equity investments (Constantinides, Donaldson, and Mehra 1998). Second, a household may face fixed costs of equity market participation; if these fixed costs exceed the benefit of participation, the household may hold no equities (Abel 1998). These constraints may affect different households differently. The first constraint is particularly likely to bind on a household whose unconstrained optimal equity position is particularly large, while the second constraint is particularly likely to bind on a poor household with little total wealth. These different sorts of households may be differentially affected by a social security reform that alters portfolio constraints.

Clear understanding of these issues requires a well-developed normative theory of optimal portfolio choice over the life cycle. Until very recently, however, theoretical work on this subject lagged far behind the familiar theory of single-period optimal portfolio choice. Samuelson (1969) and Merton (1969, 1971) showed that there are conditions under which long-lived investors choose the same portfolios as single-period investors so that the investment horizon is irrelevant; unfortunately, these conditions are highly restrictive in that they include power utility, returns on safe and

risky investments that are independently and identically distributed (i.i.d.) over time, and, most disturbing of all, the absence of labor income.

In the last few years, economists have returned to this topic and have begun to explore long-run portfolio choice when these restrictions are relaxed. Ross (1999) shows how horizon effects can emerge from more general models of preferences. Brennan and Xia (1998), Campbell and Viceira (in press), and Wachter (1998b) consider changes over time in the riskless real interest rate, while Balduzzi and Lynch (1999), Barberis (2000), Brandt (1999), Campbell and Viceira (1999), Kim and Omberg (1996), Samuelson (1991), and Wachter (1998a) consider changes over time in the equity premium, and Brennan, Schwartz, and Lagnado (1997) and Liu (1998) allow a more complex pattern of variation in both the real interest rate and the equity premium. The effect of labor income on portfolio choice has been explored by Bertaut and Haliassos (1997), Bodie, Merton, and Samuelson (1991), Cocco, Gomes, and Maenhout (CGM) (1998), Gakidis (1997), Heaton and Lucas (1997), Storesletten, Telmer, and Yaron (1998), and Viceira (in press), among others.

In this paper, we concentrate on the effect of labor income. The theoretical literature on this subject can be loosely summarized as follows. A household with labor income has an implicit holding of a nontradable asset, human capital, that represents a claim to the stream of future labor income. This nontradable asset can “crowd out” explicit asset holdings. If labor income is literally riskless, then riskless asset holdings are strongly crowded out, and the household will tilt its portfolio strongly toward risky assets (Bodie, Merton, and Samuelson 1991). If the household is constrained from borrowing to finance risky investments, the solution may be a corner at which the portfolio is 100 percent risky assets. If labor income is risky but uncorrelated with risky financial assets, then riskless asset holdings are still crowded out, but less strongly; the portfolio tilt toward risky assets is reduced (Viceira, in press). If labor income is positively correlated with risky financial assets, then risky assets can actually be crowded out, tilting the portfolio toward safe financial assets.

Under the assumption that income shocks are uncorrelated or only weakly correlated with stock returns, these results suggest that households that expect high future labor income—discounted at some appropriate rate and measured relative to financial wealth—should have the strongest desire to hold stocks. In a life-cycle model with a realistic age profile of income, the discounted value of expected future income increases relative to financial wealth in the very early part of adulthood but peaks fairly early and then declines as workers approach retirement. This suggests that fairly young (but not the very youngest) households are the most likely to be affected by borrowing constraints that limit their equity positions. While empirical evidence on household portfolio allocation is fragmentary, a few recent empirical papers have found that, over the life cycle,

household portfolios have hump-shaped equity positions and U-shaped positions in safe assets, consistent with the message of the theoretical literature (Bertaut and Haliassos 1997; Heaton and Lucas 2000; Poterba and Samwick 1997).

A complicating factor is that many households, particularly younger and poorer ones, appear to hold no equities at all. This is inconsistent with simple frictionless models of optimal portfolio choice but may be explained if there is a fixed cost of participating in equity markets. Such a fixed cost would deter young households from buying equities, but, later in the life cycle, these households might find it worthwhile to begin participating if their wealth levels are high enough to justify paying the cost.

In this paper, we explore the quantitative importance of these effects by solving a calibrated life-cycle model of consumption and portfolio choice with labor income uncertainty. The model is set in partial equilibrium and takes as given the stochastic properties of income and asset returns. We closely follow CGM (1998) but augment their model to allow us to explore alternative retirement-savings systems and fixed costs of equity market participation.

We also ask whether heterogeneity across households is likely to have a large effect on optimal investment patterns. This issue is important for the debate over social security privatization. A privatized system can allow greater individual choice over the investment of retirement wealth, but opponents argue that some individuals may fail to invest optimally and that privatization may increase administrative costs. Whatever the merits of these arguments, it is important to understand the potential gains from individual choice in the absence of administrative costs and optimization failures. To explore this issue, we compare the labor income risk of individuals working in different sectors of the economy and study the sensitivity of optimal choices to differences in the rate of time preference and the coefficient of relative risk aversion.

The organization of the paper is as follows. Section 11.1 lays out our life-cycle model and calibrates the parameters. Section 11.2 presents benchmark results in graphic form. Section 11.3 explores heterogeneity across households. Section 11.4 conducts a welfare analysis, and section 11.5 concludes.

11.1 A Life-Cycle Model of Portfolio Choice

11.1.1 Model Specification

Time Parameters and Preferences

We let t denote adult age. The investor is adult for a maximum of T periods, of which he works the first K . For simplicity, K is assumed to be exogenous and deterministic. We allow for uncertain life span in the man-

ner of Hubbard, Skinner, and Zeldes (1994). Let p_t denote the probability that the investor is alive at date $t + 1$, conditional on being alive at date t . Then investor i 's preferences are described by the time-separable power utility function

$$(1) \quad E_i \sum_{t=1}^T \delta^{t-1} \left(\prod_{j=0}^{t-1} p_j \right) \frac{C_{it}^{1-\gamma}}{1-\gamma},$$

where C_{it} is the level of date t consumption, $\gamma > 0$ is the coefficient of relative risk aversion, and $\delta < 1$ is the discount factor. We assume that the individual derives no utility from leaving a bequest.

The Labor Income Process

Investor i 's age t labor income, Y_{it} , is exogenously given by

$$(2) \quad \log(Y_{it}) = f(t, Z_{it}) + v_{it} + \varepsilon_{it} \quad \text{for } t \leq K,$$

where $f(t, Z_{it})$ is a deterministic function of age and other individual characteristics Z_{it} , ε_{it} is an idiosyncratic temporary shock distributed as $N(0, \sigma_\varepsilon^2)$, and v_{it} is given by

$$(3) \quad v_{it} = v_{i,t-1} + u_{it},$$

where u_{it} is distributed as $N(0, \sigma_u^2)$ and is uncorrelated with ε_{it} . Thus, log income is the sum of a deterministic component that can be calibrated to capture the hump shape of earnings over the life cycle and two random components, one permanent and one transitory. We assume that the temporary shock ε_{it} is uncorrelated across households, but we decompose the permanent shock u_{it} into an aggregate component ξ_t and an idiosyncratic component ω_{it} , uncorrelated across households:

$$(4) \quad u_{it} = \xi_t + \omega_{it}.$$

This decomposition implies that the random component of aggregate labor income follows a random walk, an assumption made by Fama and Schwert (1977) and Jagannathan and Wang (1996). While macroeconomists such as Campbell (1996), Campbell and Mankiw (1989), and Pischke (1995) have found empirical evidence for short-term persistence in aggregate quarterly labor income growth, the simplification to a random walk should have little effect on optimal consumption and portfolio choice over the life cycle.

Financial Assets

We assume that there are two assets in which the agent can invest: a riskless asset with gross real return \bar{R}_f , which we call *Treasury bills*, and a

risky asset with gross real return R_t , which we call *stocks*. The excess return on the risky asset, $R_{t+1} - \bar{R}_f$, is given by

$$(5) \quad R_{t+1} - \bar{R}_f = \mu + \eta_{t+1},$$

where η_{t+1} , the period $t + 1$ innovation to excess returns, is assumed to be i.i.d. over time and distributed as $N(0, \sigma_\eta^2)$. We allow innovations to excess returns to be correlated with innovations to the aggregate component of permanent labor income, and we write the correlation coefficient as $\rho_{\xi\eta}$. We also allow for fixed costs of equity market participation: to have access to the stock market, the investor must pay a one-time monetary fixed cost equal to F .

Retirement and Liquid Wealth

We model a system of mandatory saving for retirement in the following simple way. During working life, the individual must save a fraction, θ , of current labor income as retirement wealth. Under this assumption, disposable labor income, Y_{it}^d , is given by

$$(6) \quad Y_{it}^d = (1 - \theta)Y_{it} \quad \text{for } t \leq K.$$

The amount θY_{it} is added to retirement wealth, denoted by W_{it}^R . During working life, retirement wealth is illiquid; the individual cannot consume it or borrow against it. At age K , retirement wealth is rolled into a riskless annuity so that the individual receives in each of the retirement years the annuity value corresponding to W_{iK}^R . This assumption of riskless annuitization affects the portfolio choices of older investors. An interesting extension of our work would be to allow investors to choose between riskless and variable annuities.

We consider several alternative systems governing the investment of retirement wealth during working life. In the first system, the individual is forced to hold retirement wealth in riskless assets. This implies that $W_{it}^R = B_{it}^R$, where B_{it}^R is the dollar amount of retirement wealth that investor i has in riskless assets. In alternative systems, retirement wealth is partially or fully invested in risky assets, but the allocation remains constant over time and is not controlled by the investor. We interpret this as the Social Security Administration managing the individual's retirement account. For this reason, the fixed cost of investing in stocks applies only to investments outside the retirement account.

Investors also have liquid wealth outside their retirement accounts. We denote liquid wealth of investor i at date t by W_{it}^L and liquid holdings of bills and stocks by B_{it}^L and S_{it}^L , respectively. We assume that the investor faces the following borrowing and short-sales constraints:

$$(7) \quad B_{it}^L \geq 0,$$

$$(8) \quad S_{it}^L \geq 0.$$

The borrowing constraint (7) ensures that the investor’s allocation to bills in both the liquid and the retirement accounts is nonnegative at all dates. It prevents the investor from borrowing against future labor income or retirement wealth. The short-sales constraint (8) ensures that the investor’s allocation to equities is nonnegative at all dates.

The Household’s Optimization Problem

In each period of a household’s working life ($t \leq K$), the timing of events is as follows. The investor starts the period with liquid wealth W_{it}^L and retirement wealth W_{it}^R . Then labor income Y_{it} is realized. Following Deaton (1991), we denote *cash on hand* in period t by

$$(9) \quad X_{it} = W_{it}^L + (1 - \theta)Y_{it}.$$

The investor must decide how much to consume, C_{it} , whether to pay the fixed cost of entering the stock market (if he has not done so before), and how to allocate the remaining cash on hand between stocks and bills. We denote the proportion of liquid wealth invested in stocks by α_{it}^L . The proportion of retirement wealth invested in stocks, α_{it}^R , is given exogenously by the retirement system and does not vary over time, so $\alpha_{it}^R = \alpha_i^R$ for all t . We consider different values for α_i^R .

Next-period liquid wealth and retirement wealth are then given by

$$(10) \quad W_{i,t+1}^L = R_{p,i,t+1}^L [W_{it}^L + (1 - \theta)Y_{it} - C_{it} - (f_{it} - f_{i,t-1})F],$$

$$(11) \quad W_{i,t+1}^R = R_{p,i,t+1}^R [W_{it}^R + \theta_t Y_{it}],$$

where f_{it} is a binary variable that equals zero until the investor pays the fixed cost of entering the stock market and equals one thereafter, and $R_{p,i,t+1}^j$ is the return on the portfolio held from period t to period $t + 1$:

$$(12) \quad R_{p,i,t+1}^j \equiv \alpha_{it}^j R_{t+1} + (1 - \alpha_{it}^j) \bar{R}_j, \quad j = L, R.$$

Here, α_{it}^L is freely chosen when $f_{it} = 1$ and equals zero when $f_{it} = 0$.

After retirement ($t > K$), the problem takes the same form except that retirement wealth no longer accumulates. Instead, it is annuitized and provides riskless income $A(W_{iK}^R)$. After-tax labor income $(1 - \theta)Y_{it}$ in (9) and (10) is replaced by $A(W_{iK}^R)$.

The problem that the investor faces is to maximize (1) subject to the working-life and retirement versions of (2)–(12) plus the constraints that consumption must be nonnegative at all dates. The control variables of the problem are $\{C_{it}, \alpha_{it}^L, f_{it}\}_{t=1}^T$. The state variables are $\{t, X_{it}, W_{it}^L, v_{it}, f_{i,t-1}\}_{t=1}^T$. The problem is to solve for the policy rules as a function of the state variables, that is, $C_{it}(X_{it}, W_{it}^R, v_{it}, f_{i,t-1})$, $\alpha_{it}^L(X_{it}, W_{it}^R, v_{it}, f_{i,t-1})$, and $f_{it}(X_{it}, W_{it}^R, v_{it}, f_{i,t-1})$.

Numerical Solution

This problem cannot be solved analytically. We derive the policy functions numerically by discretizing the state space and the variables over which the choices are made and by using Gaussian quadrature to approximate the distributions of the innovations to the labor income process and risky asset returns (Tauchen and Hussey 1991). The problem is then solved by using backward induction. In period T , the investor consumes all his wealth, and the value function coincides with the instantaneous utility. In every period t prior to T , and for each admissible combination of the state variables, we compute the value associated with each level of consumption, decision to pay the fixed cost of entering the stock market, and share of liquid wealth invested in stocks. This value is equal to current utility plus the expected discounted continuation value. To compute this continuation value for points that do not lie on the grid, we use cubic spline interpolation. The combinations of the choice variables ruled out by the constraints of the problem are given a very large (negative) utility such that they are never optimal. We optimize over the different choices using grid search.

When the fixed cost of equity market participation F is equal to zero, we simplify the solution by exploiting the scale independence of the maximization problem and rewriting all variables as ratios to the permanent component of labor income.

11.1.2 Calibration

Time Parameters and Preferences

Adult age starts at age twenty for households without a college degree and at age twenty-two for households with a college degree. The age of retirement is set to sixty-five for all households. The investor dies with probability one at age one hundred. Prior to this age, we use the mortality tables of the National Center for Health Statistics to parameterize the conditional survival probabilities, p_j for $j = 1, \dots, T$. We set the discount factor δ to 0.96 and the coefficient of relative risk aversion γ to 5. In variations of the benchmark case, we also consider investors who are extremely impatient with $\delta = 0.80$, comparatively risk tolerant with $\gamma = 2$, and extremely risk averse with $\gamma = 10$.

The Labor Income Process

To estimate the labor income process, we follow CGM (1998). Here, we briefly describe the data and estimation method.

We use the family questionnaire of the Panel Study on Income Dynamics (PSID) to estimate equations (2) and (3), which give labor income as a function of age and other characteristics. Families that are part of the

Survey of Economic Opportunities subsample are dropped to obtain a random sample. Only households with a male head are used as the age profile of income may differ across male- and female-headed households and relatively few observations are available for female-headed households. Retirees, nonrespondents, students, homemakers, and household heads younger than twenty (twenty-two for college-educated households) or older than sixty-five are also eliminated from the sample.

Like CGM (1998) and Storesletten, Telmer, and Yaron (1998, 1999), we take a broad definition of labor income so as implicitly to allow for insurance mechanisms—other than asset accumulation—that households use to protect themselves against pure labor income risk. Such insurance mechanisms include welfare programs that effectively set a lower bound on the support of nonasset income, endogenous variation in the labor supply of both male and female household members, financial help from relatives and friends, and so on. Thus, we define *labor income* as total reported labor income plus unemployment compensation, workers' compensation, social security, supplemental social security, other welfare, child support, and total transfers (mainly help from relatives), all this for both head of household and, if present, his spouse. Observations that still report zero for this broad income category are dropped. Labor income defined this way is deflated using the consumer price index, with 1992 as the base year. The sample starts in 1970, so a household appears at most twenty-four times in the sample. Households with fewer observations are retained in the panel.

The estimation controls for family-specific fixed effects. To control for education, the sample is split into three groups: households without high school education, a second group with high school education but without a college degree, and, finally, college graduates. This sample split is intended to accommodate the well-established finding that age profiles differ in shape across education groups (Attanasio 1995; Hubbard, Skinner, and Zeldes 1994). For each education group, the function $f(t, Z_{it})$ is assumed to be additively separable in t and Z_{it} . The vector Z_{it} of personal characteristics, other than age and the fixed household effect, includes marital status and household composition. Household composition equals the additional number of family members in the household besides the head and, if present, his spouse.

Ideally, one should also control for occupation. Using PSID data, this is problematic because, from the 1975 wave onward, the majority of the unemployed report no occupation and are categorized together with people who are not in the labor force. But modeling unemployment as a switch in occupation is inappropriate as the possibility of unemployment through layoff is one of the main sources of labor income risk. We explore this in greater detail in section 11.3 below.

To obtain age profiles suitable for the simulation model of life-cycle

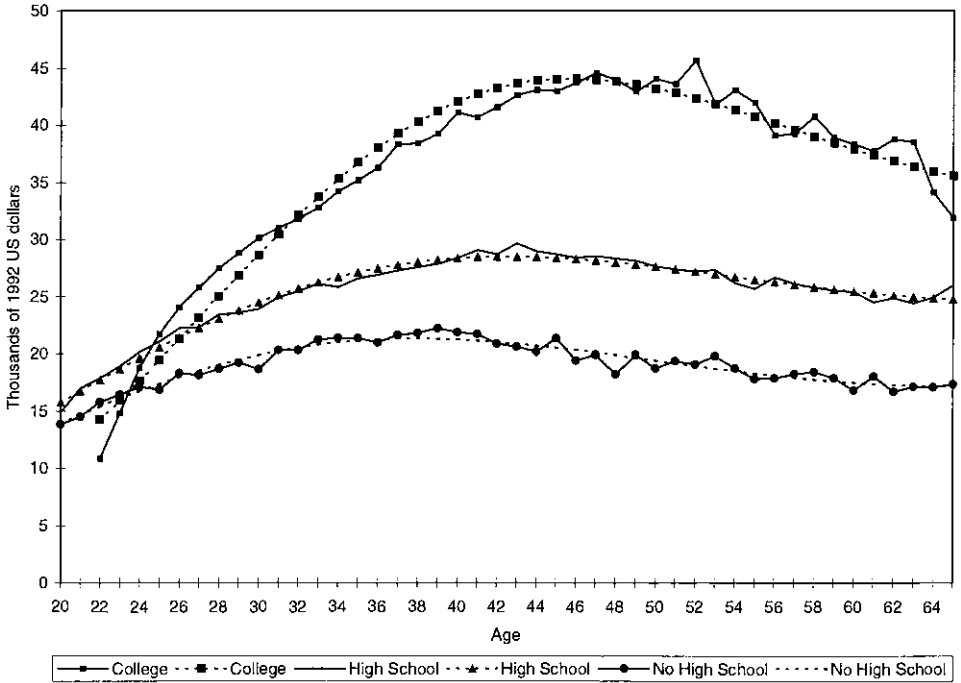


Fig. 11.1 Labor income profiles (age dummies and fitted polynomials)

Note: The solid line gives the underlying age dummies, the dotted line a third-order polynomial fitted to these age dummies.

portfolio choice, we fit a third-order polynomial to the age dummies estimated from the PSID. The resulting income profiles are similar to those used in Attanasio (1995), Carroll and Summers (1991), and Gourinchas and Parker (1996). They are plotted in figure 11.1, along with the underlying age dummies for each of the three education groups.

To estimate the variances of permanent and temporary shocks to labor income, we follow Carroll and Samwick (1997). Defining Y_{it}^* as

$$(13) \quad \log(Y_{it}^*) \equiv \log(Y_{it}) - \hat{f}(t, Z_{it}),$$

then

$$(14) \quad \text{var} [\log(Y_{i,t+d}^*) - \log(Y_{it}^*)] = d\sigma_u^2 + 2\sigma_\varepsilon^2.$$

We estimate σ_u^2 and σ_ε^2 by running an OLS regression of $\text{var}[\log(Y_{i,t+d}^*) - \log(Y_{it}^*)]$ on d and a constant term. We find that groups with less education tend to have more variable transitory income shocks but less variable permanent shocks than groups with more education. Table 11.1 reports these variances.

Table 11.1 Baseline Parameters

Description	Parameter Value
Retirement age (K)	65
Discount factor (δ)	.96
Risk aversion (γ)	5
Variance of transitory shocks (σ_{ε}^2):	
No high school	.1056
High school	.0738
College	.0584
Variance of permanent shocks (σ_{η}^2):	
No high school	.0105
High school	.0106
College	.0169
Sensitivity to stock returns (β):	
No high school	.0956
High school	.0627
College	.0733
Correlation with stock returns ($\rho_{\varepsilon\eta}$):	
No high school	.3280
High school	.3709
College	.5155
Riskless rate ($\bar{R}_f - 1$)	.02
Mean excess return on stocks (μ)	.04
Standard stock return (σ_{η})	.157
Fixed cost (F)	0 or 10,000
Social security tax rate (θ)	.10

We use a similar procedure to estimate the correlation between labor income shocks and stock returns, $\rho_{\varepsilon\eta}$. The change in $\log(Y_{it}^*)$ can be written as

$$(15) \quad \Delta \log(Y_{it}^*) = \xi_t + \omega_{it} + \varepsilon_{it} - \varepsilon_{i,t-1}.$$

Averaging across individuals gives

$$(16) \quad \overline{\Delta \log(Y_{it}^*)} = \xi_t.$$

The correlation coefficient is then easily computed from the OLS regression of $\overline{\Delta \log(Y_{it}^*)}$ on demeaned excess returns:

$$(17) \quad \overline{\Delta \log(Y_{it}^*)} = \beta(R_{t+1} - \bar{R}_f - \mu) + \psi_t.$$

As an empirical measure for the excess return on our stylized risky asset, we use CRSP (Center for Research in Securities Prices) data on the New York Stock Exchange value-weighted stock return relative to the Treasury-

bill rate. For all education groups, the regression coefficients are strikingly low and insignificant. To allow for potential lags in the realization of labor income, we repeat the exercise with the excess stock return lagged one year. The relation becomes much stronger: the regression coefficient now varies from 0.06 to 0.10 and the correlation coefficient from 0.32 to 0.52, as reported in table 11.1. Interestingly, the correlation of labor income with the stock market is larger and more significant for households with higher education.¹

In our portfolio-choice model, allowing for lags in the relation between innovations in stock returns and permanent shocks to labor income unfortunately requires an additional state variable. We therefore assume that the correlation is contemporaneous. The model requires the variances of both ξ_t , the aggregate permanent labor income shock that is correlated with stock market risk, and ω_{it} , the idiosyncratic permanent shock to labor income. The first variance is obtained immediately as the variance of $\Delta \log(Y_{it}^*)$. Subtracting this variance from the total variance of u_{it} gives then the variance of ω_{it} .

Other Parameters

The riskless real interest rate is assumed to be constant at 2 percent. We set the equity premium μ equal to 4 percent. This is well below the long-run historical average but represents a reasonable compromise between that average and lower forward-looking estimates based on the observation that stock prices have tended to increase in recent years relative to corporate earnings (Blanchard 1993; Campbell and Shiller 1998). We set the standard deviation of innovations to the risky asset σ_η to 0.157. Recall that the classic formula for the risky-asset portfolio share, under power utility with i.i.d. returns and no labor income, is $\mu/\gamma\sigma_\eta^2$. With these parameters, the implied risky asset share would be about one-third at the benchmark risk aversion of 5; we find higher optimal shares in our model only because of the presence of labor income. We set the fixed cost of equity market participation to zero in the benchmark case, but we go on to consider a \$10,000 fixed cost.

The proportion of labor income θ that is added to retirement wealth is equal to 10 percent of current labor income when retirement wealth accumulates at the riskless rate. This value implies an average replacement ratio at age sixty-five of 60 percent. When retirement wealth is also invested in stocks, we either fix θ at the same 10 percent value or adjust it so as to maintain the replacement ratio at 60 percent. We will show that the value of θ has a very important effect on our results. Table 11.1 summarizes the parameters used in the baseline case.

1. We also examined the relation of labor income shocks with lagged returns on long-term government bonds. These results are reported in sec. 11.3 below.

11.2 Benchmark Results: A Graphic Summary

The first comparison that we consider is between a system with riskless retirement accumulation ($\alpha^R = 0$) and one in which at each age half of retirement wealth is invested in stocks ($\alpha^R = .5$). In the latter system, we reduce the social security tax rate from 10 to 6 percent so that, on average, the replacement ratio is the same in both systems and equal to 0.6. At retirement, the account is annuitized at the riskless interest rate so that, on average, and given survival probabilities, the system has zero balance.

To study the behavior of the variables in the model, we calculate cross-sectional averages across ten thousand households receiving different draws of income and asset returns and plot them against age. Figure 11.2 plots labor income net of social security contributions, consumption, liquid wealth, and retirement wealth for households with a high school degree (the life-cycle patterns for other education groups are similar). Figure 11.2*A* illustrates the system in which retirement wealth is fully invested in the riskless asset, and figure 11.2*B* illustrates the system in which retirement wealth is partially invested in stocks.

In both systems, the average consumer is borrowing constrained early in life. Consumption tracks net income very closely, and little savings accumulate outside the retirement account until after age forty. These limited savings early in life are driven by the precautionary savings motive; thus, like Gourinchas and Parker (1996), we find that younger consumers are buffer-stock savers rather than life-cycle savers in the classic sense. Consumption rises with income early in life because of borrowing constraints and falls later as increased mortality drives up the effective rate of time preference; thus, consumption profiles are hump shaped over life, as found in the literature on life-cycle consumption behavior.²

Investment of some retirement wealth in stocks has an income effect. Because the average return on stocks is higher than the average return on bills, and because younger consumers have neither the desire nor the liquid wealth to offset a shift of retirement wealth into stocks, the shift increases average lifetime resources. Since we reduce the social security tax rate to keep the average replacement ratio constant across systems, the investment of retirement wealth in stocks frees up resources in the working years. These additional resources are consumed early in life since, at this stage, households are borrowing constrained.

Of course, the investment of retirement wealth in stocks has a cost: it imposes additional risk on households. In midlife, households react by increasing their precautionary saving, accumulating more liquid wealth, and consuming less relative to income. After retirement, the additional

2. We could generate a more pronounced hump shape in consumption if we added age-specific preference shocks to the model.

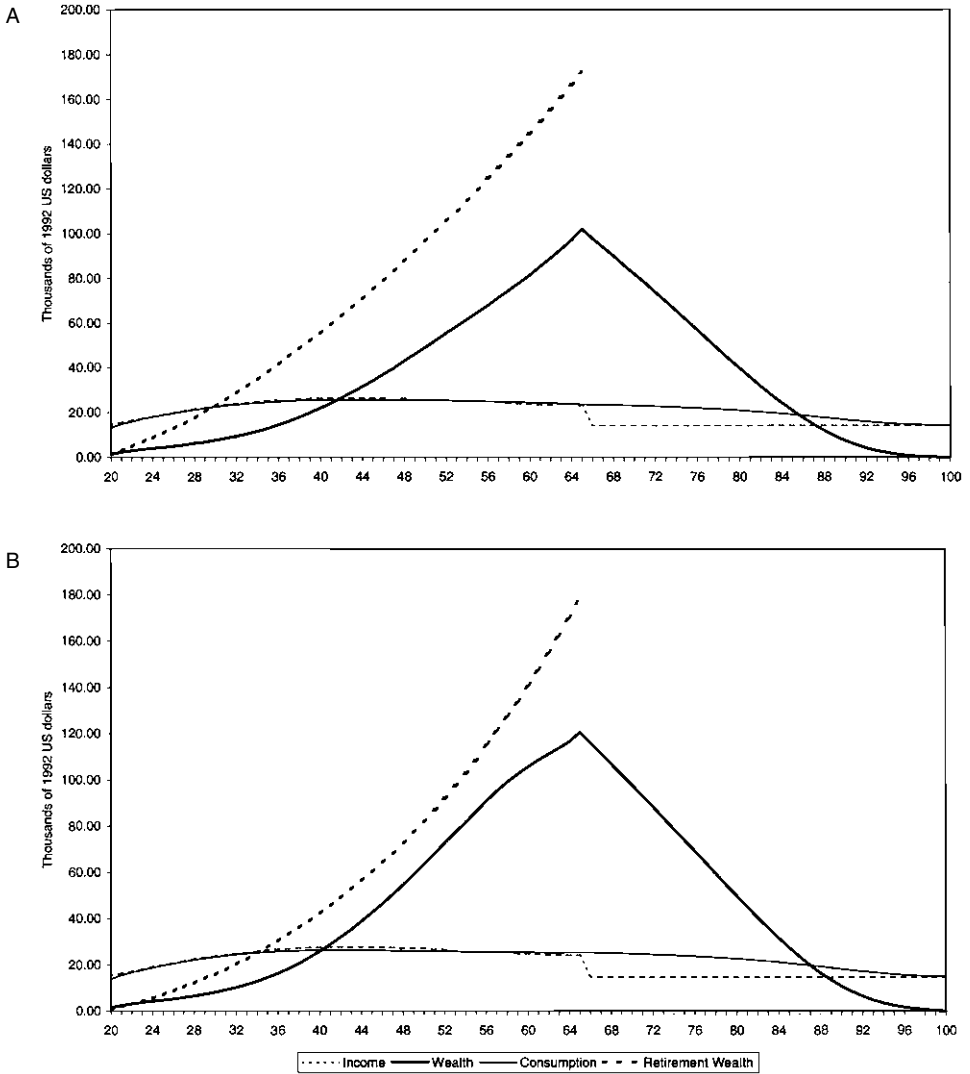


Fig. 11.2 Consumption, income, wealth, and annuity: A, Retirement wealth fully invested in the safe asset; B, Retirement wealth invested 50/50 in risky/safe asset

wealth is run down since, at this stage, retirement wealth is rolled into a riskless annuity. These patterns show up in the paths of consumption relative to income in figure 11.2.

Figure 11.3 plots liquid wealth and liquid holdings of equities and bills over the life cycle. In each retirement system, the borrowing constraint binds for young households; they would like to take more equity risk but

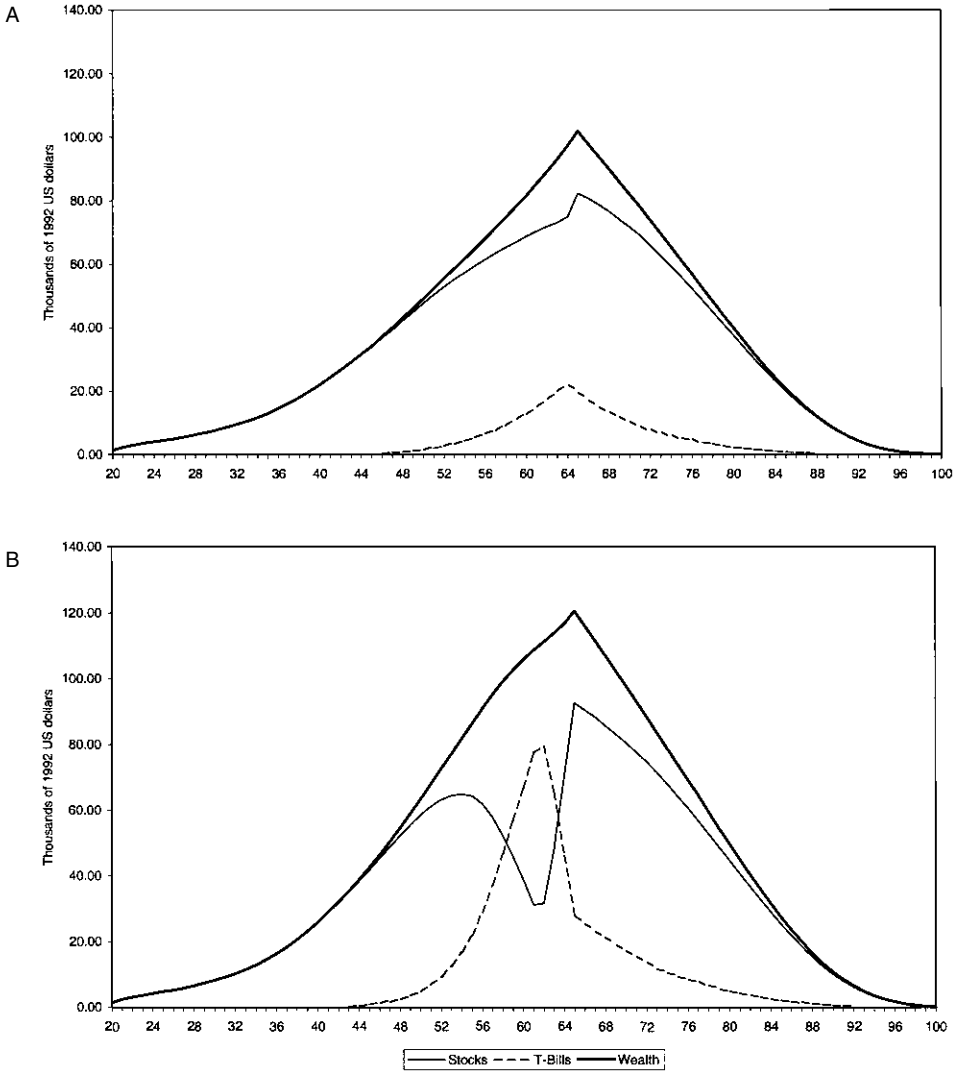


Fig. 11.3 Wealth, stocks, and Treasury bills: A, Retirement wealth fully invested in the safe asset; B, Retirement wealth invested 50/50 in risky/safe asset

are unable to do so. For approximately the first twenty years of life, they hold 100 percent of their portfolios in the form of equity. Households in midlife hold bills, but these holdings decrease again after retirement.

Figure 11.4 plots the portfolio share of stocks in liquid wealth. The crucial variables for portfolio composition are liquid wealth, retirement wealth, and future labor income. In the model, although future labor

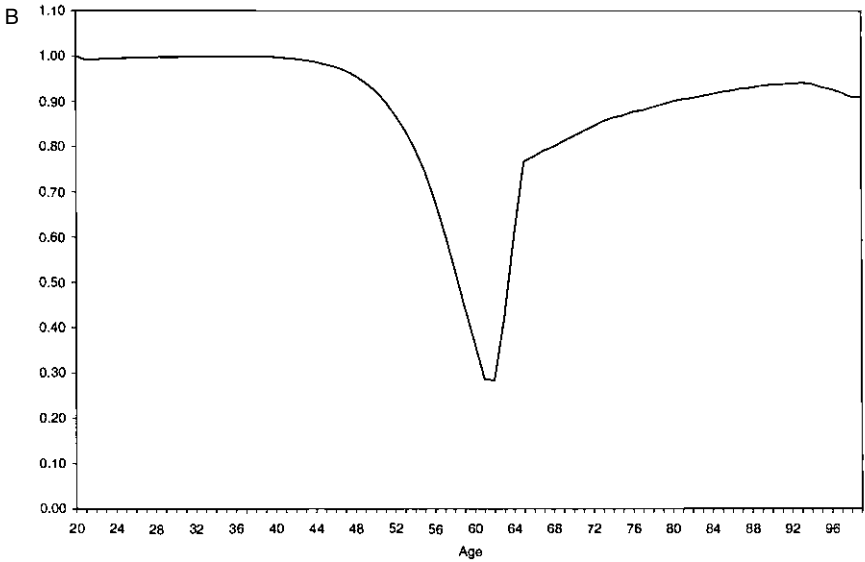
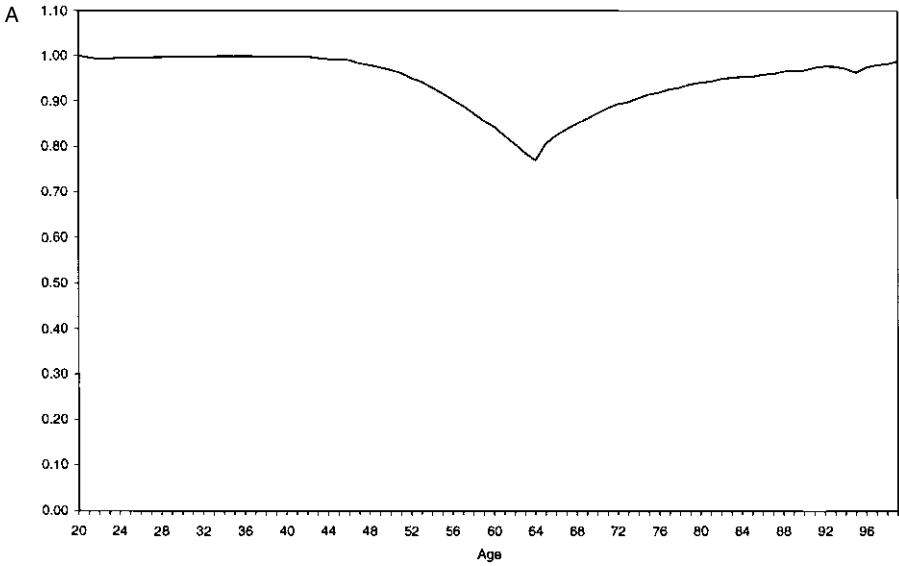


Fig. 11.4 Liquid portfolio share invested in stocks: *A*, Retirement wealth fully invested in the safe asset; *B*, Retirement wealth invested 50/50 in risky/safe asset

income is risky, it can be thought of as implicit holdings of a riskless asset. Innovations to labor income are positively correlated with innovations to stock returns, but this correlation is not sufficiently large for future labor income to resemble more closely stocks than bills. Since early in life the implicit holdings of the riskless asset in the form of future labor income are large, the investor wishes to invest what little liquid wealth he has fully into stocks. From age forty on, liquid wealth increases relative to future labor income and retirement wealth so that implicit holdings of the riskless asset become less important. This induces a shift in the composition of liquid wealth toward bills. After retirement, liquid wealth is run down more rapidly than the implicit annuitized holdings of the riskless asset. As this happens, the implicit holdings of the riskless asset become relatively more important, inducing a shift in portfolio composition back toward stocks.

These considerations explain the life-cycle patterns in both figure 11.4A and figure 11.4B. There are, however, important differences in magnitudes between the two figures. In figure 11.4B, the midlife decrease in the share invested in stocks is much more dramatic. Since the investor already holds risky assets in the retirement account, he wishes to hold a safer liquid portfolio.

Another way to understand these results is to compare the patterns of current utility of consumption over the life cycle. Figure 11.5A shows the ratio of average current utility for households in the 50/50 system to the average current utility of households in the 100/0 system. We see that most of the gain from investing retirement wealth in stocks occurs early in life. The source of this gain is the higher levels of consumption that a lower social security contribution allows. Return risk in the retirement account allows some households to end up poorer so that, after age fifty-five, current utility is on average higher in the 100/0 system.

Return risk also increases the dispersion of utility. The standard deviations of current utility across households with different income and return realizations are higher in the 50/50 system than in the 100/0 system. Figure 11.5B reports the ratio of these standard deviations in the two systems. These *ex post* differences raise important practical issues for designers of retirement systems because they may create an incentive to bail out cohorts negatively affected by lower stock-return realizations.

Of course, a proper welfare analysis requires the discounting of current utility over the life cycle. We defer such an analysis to section 11.4 below.

One limitation of the results reported so far is that they counterfactually predict 100 percent stock market participation among younger investors. However, we can modify this prediction, with little effect on other aspects of the model, by adding a fixed cost of stock market participation. Figure 11.6 reports results with a \$10,000 fixed cost. The fraction of households that have paid the fixed cost and the average share of assets invested in

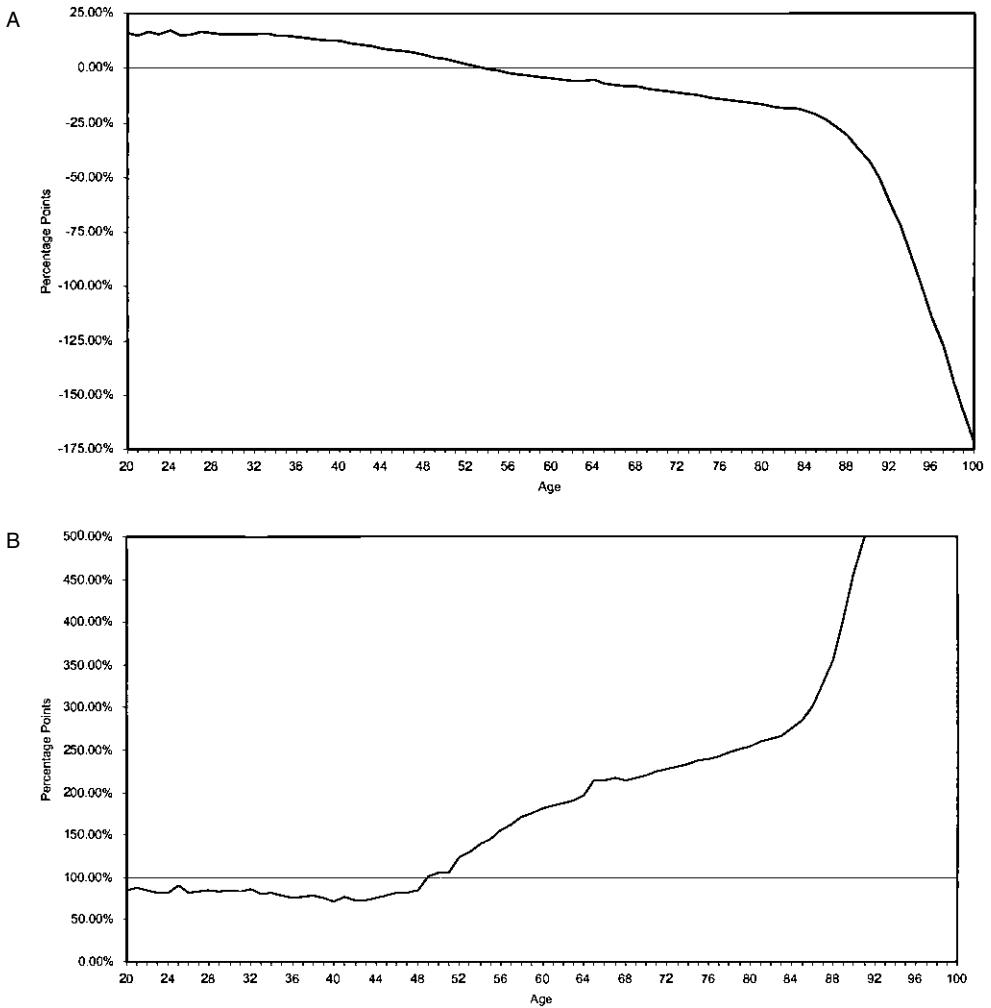


Fig. 11.5 A, Utility gain; B, Utility dispersion (ratio of standard deviations)

stocks are plotted for each retirement system. Early in life, the two series move almost perfectly together, showing that young investors are either entirely in or entirely out of the market; later in life, all investors have paid the fixed cost, and the model behaves much as it did in the absence of the cost.

11.3 Heterogeneity

In the previous section, we studied a representative household with a high school education but no college degree. Results are similar for repre-

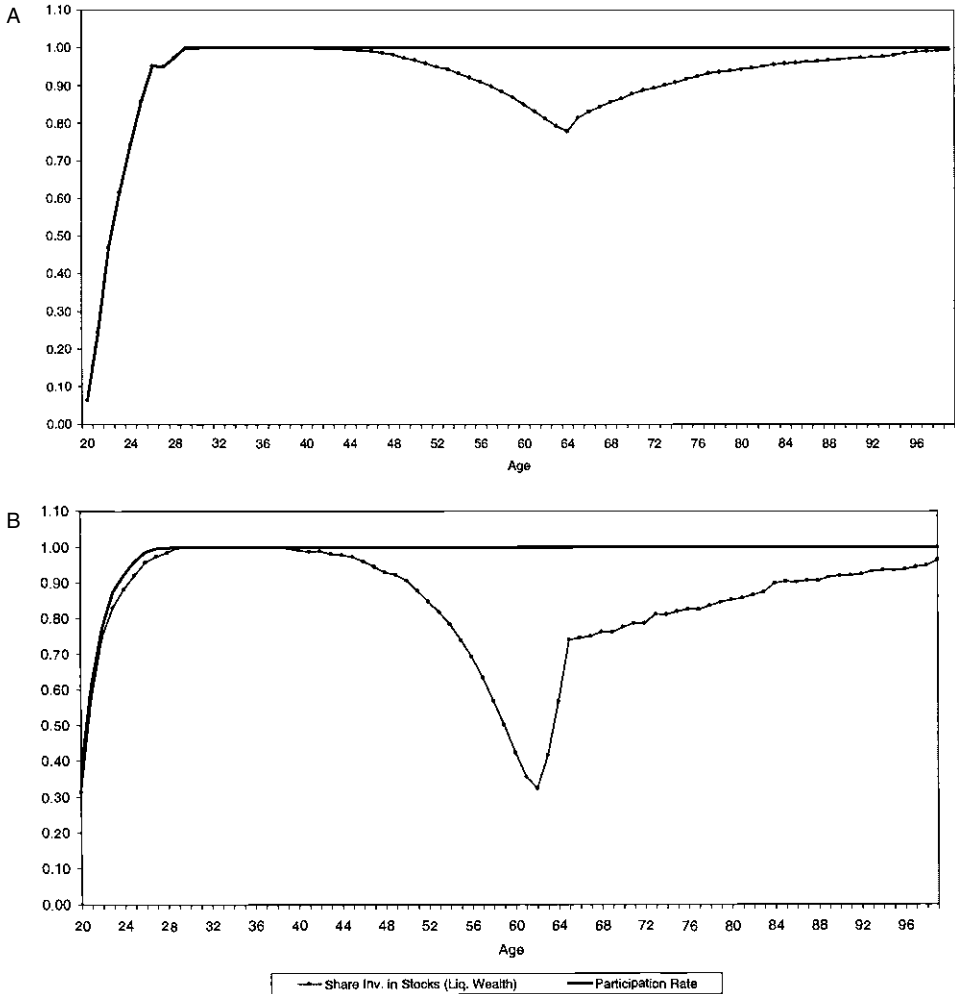


Fig. 11.6 Liquid portfolio share invested in stocks and participation rate: *A*, Retirement wealth fully invested in the safe asset and fixed cost of entering in the stock market; *B*, Retirement wealth invested 50/50 in risky/safe asset and fixed cost of entering the stock market

sentative households in the other two education groups. However, households may differ along other dimensions. For example, labor income processes may differ for households that work in different industries and for self-employed households. Also, some households may be more impatient or risk averse than others, as found by Barsky et al. (1997). These differences across households may have important effects on optimal investment strategies. In this section, we consider this issue.

11.3.1 Measuring Heterogeneity in Labor Income

We first consider variation in the stochastic structure of the labor income process across industries and then study differences between self-employed and non-self-employed households.

We use the two-digit SIC classification to split households into twelve different industries. Starting in 1972, the PSID reports both the current industry of the household head if currently working and the last industry if currently unemployed. This is the information that we use. Three caveats apply, however. First, we ignore the industry of the spouse. This might be problematic because the spouse's labor income is added to the head's labor income, yet it might have quite different risk characteristics. Second, on average, 16 percent of our respondents switch industries each year. Business-cycle considerations (like the anticipation of a recession) might force people out of cyclic sectors and into less volatile industries. As we do not model the switching decision, our estimates of the sensitivities of labor income shocks to financial market risks for different industries might be biased. However, we did not find any significant effects when we regressed the number of industry switchers onto innovations in business-cycle indicators. Third, there is a timing issue because the labor income reported in the PSID is for the previous calendar year while the industry concerns the current job.

Table 11.2 reports the number of household-year observations for each of the thirty-six different education-industry cells. There is tremendous variation across industries, with particularly small numbers in mining, per-

Table 11.2 Cell Sizes

Industry	No			Total
	High School	High School	College	
Agriculture	765	1,165	381	2,311
Mining	119	360	98	577
Manufacturing	3,116	7,319	2,430	12,865
Construction	1,414	2,938	494	4,846
Transportation, communications	828	3,172	701	4,701
Trade	1,072	4,661	1,445	7,178
Finance, real estate, insurance	109	995	965	2,069
Business services	439	1,437	548	2,424
Personal services	159	358	108	625
Recreation	39	236	110	385
Professional services	356	1,738	3,883	5,977
Public administration	339	2,357	1,097	3,793
Self-employed	1,115	3,562	2,242	6,919
Non-self-employed	7,640	23,174	10,018	40,832
Total	8,755	26,736	12,260	47,751

sonal services, and recreation. These industries are omitted from tables 11.3–11.5 below. As a further cutoff, we drop cells in which any PSID wave contains fewer than twenty observations; these cells are left blank in tables 11.3–11.5. We do, however, include observations in these small cells when estimating column and row totals, that is, when reporting the results for a given industry across all education levels or for a given education group across all industries.

We reestimate age profiles of labor income for the shorter PSID sample beginning in 1972 and including industry dummies in the vector Z_{it} of personal characteristics. That is, we allow industry to influence only the level and not the shape of the age profile for a household with a given amount of education. We then estimate the stochastic model of labor income separately for each education-industry cell.³

Table 11.3 reports the total variance of income, and its decomposition into permanent and transitory components, for each different education-industry cell. Agriculture has by far the highest variance of labor income shocks. Other industries subject to significant labor income shocks are construction and business services. The variance decomposition indicates that labor income shocks for construction workers without a high school degree are entirely temporary. At the other extreme, permanent income shocks are especially important for college graduates in public administration. As a general pattern, the relative importance of permanent shocks seems to increase with education attainment. This was already documented for the column totals but seems robust within individual industries.

The bottom of table 11.3 splits the sample in a different way, by distinguishing self-employed from non-self-employed households. (We included both types of households in the industry analysis since there are too few self-employed households to allow an industry decomposition.) Income variability is dramatically larger for self-employed households. Income shocks are entirely temporary for the self-employed without a high school degree but are disproportionately permanent for the self-employed in the two higher education groups.

Table 11.4 considers heterogeneity in the sensitivity of different households' income shocks to lagged stock returns. Table 11.5 repeats the exercise replacing lagged stock returns with lagged returns on long-term government bonds. Unfortunately, the small cell sizes mean that the results are often statistically insignificant for individual industries, but there are many interesting patterns. Stock market risk seems especially relevant for people in manufacturing, construction, and public administration. Interest-rate risk shows up for agriculture, professional services, and finance, real

3. We also report results aggregating across industries and education groups. Note that these differ slightly from the results reported in sec. 11.2 above because, here, our sample starts in 1972 and we include industry dummies in the estimation of age profiles for labor income.

Table 11.3 Variance Decomposition

Industry	No High School		High School		College		All	
	Total	Permanent/ Total	Total	Permanent/ Total	Total	Permanent/ Total	Total	Permanent/ Total
Agriculture3094	.13503166	.0786
Manufacturing	.0579	.0973	.0506	.0912	.0469	.2483	.0529	.1315
Construction	.1097	.0000	.1030	.13941080	.1451
Transportation, communications	.1119	.3058	.0553	.08330672	.1616
Trade	.0997	.1707	.0890	.1254	.0721	.2586	.0895	.1591
Finance, real estate, insurance0930	.1264	.0826	.2382	.0906	.1471
Business services1275	.18201188	.1789
Professional services0702	.1180	.0732	.1953	.0773	.1478
Public administration0470	.1657	.0361	.2594	.0450	.1882
Self-employed	.3695	.0000	.2199	.2192	.1985	.1909	.2302	.1634
Non-self-employed	.0655	.0762	.0535	.0969	.0418	.1768	.0545	.1124
All	.1057	.0683	.0829	.1306	.0744	.2390	.0868	.1413

Note: All positive variances are significantly different from zero at the 0.1 percent level. Two variances estimated to be negative have been set to zero.

Table 11.4 Regression of Permanent Aggregate Shock on Lagged Excess Stock Returns

Industry	No High School			High School			College			All		
	β	ρ	ρ	β	ρ	ρ	β	ρ	ρ	β	ρ	ρ
Agriculture3507	.333317542218
Manufacturing	.1025	.3381	.2786	.07500577	.20570829*3524
Construction	.3197**	.6577	.2423	.07701333**5527
Transportation, communications	-.1778	.2983	.3234	.048802941879
Trade	.1163	.2128	.0245	-.0050	-.0440	.109500750283
Finance, real estate, insurance0693	.0360	-.0093	.022404611281
Business services0265	-.015608041949
Professional services0860	.03200622	.345704792644
Public administration2302	.06342195**	.50941072**4237
Self-employed	.0916	.0964	.1652	.07531869*	.416411193159
Non-self-employed	.1022*	.3970	.4148	.0681*0462	.29980710**4638
All	.0937	.3170	.3785	.0689*0721**	.49430762**4562

*Significant at the 10 percent level.

**Significant at the 5 percent level.

Table 11.5 Regression of Permanent Aggregate Shock on Lagged Excess Bond Returns

Industry	No High School		High School		College		All	
	β	ρ	β	ρ	β	ρ	β	ρ
Agriculture5750**	.42445735**	.5083
Manufacturing	.2020**	.4674	.0750	.2786	.1050	.2627	.1221*	.3637
Construction	.1371	.1977	.1788*	.39611378*	.4005
Transportation, communications	-.1723	.2027	.0632	.29430485	.2177
Trade	.1822	.2324	.0835	.2608	.0012	.0022	.0754	.1980
Finance, real estate, insurance1562	.2074	.1995	.3283	.1936*	.3779
Business services	-.2202	.2638	-.0941	.1600
Professional services0199	.0374	.1319**	.5138	.1174**	.4540
Public administration1539*	.3926	.1230	.2000	.1329*	.3685
Self-employed	.3653	.2693	.1912	.2941	.1391	.2173	.2124**	.4204
Non-self-employed	.1223	.3330	.1021**	.4361	.1290**	.5874	.1152**	.5279
All	.1511*	.3586	.1168**	.4501	.1316**	.6326	.1300**	.5457

*Significant at the 10 percent level.

**Significant at the 5 percent level.

estate, and insurance in addition to the stock market-sensitive sectors. Among college graduates, the self-employed are especially exposed to stock market risk, while interest-rate risk is far more important for the non-self-employed. This finding supports the conclusion of Heaton and Lucas (2000) that privately owned business risk is an especially important substitute for stock market risk in the portfolios of many wealthy households.

11.3.2 Effects of Heterogeneity on Portfolio Choice

In this section, we illustrate the effects of investor heterogeneity on optimal consumption and portfolio choice. First, we consider heterogeneity of preferences, calculating optimal behavior for highly risk-averse investors with $\gamma = 10$ and impatient investors with $\delta = 0.8$. Second, we consider differences in labor income risk of the sort illustrated in section 11.3.1 above. To highlight these differences, we simulate the behavior of households whose income is particularly risky and highly correlated with asset returns: self-employed college graduates.

Table 11.6 shows average consumption and liquid wealth (in thousands of dollars) and the share of liquid wealth invested in stocks for different age groups. This is a more compact way for us to summarize the informa-

Table 11.6 Life-Cycle Profiles

Age	Baseline Case			$\gamma = 10$		$\delta = 0.8$		Self-Employed	
	0/100	50/50 ^a	50/50	0/100	50/50	0/100	50/50	0/100	50/50
<i>Consumption</i>									
20–35	20.22	20.26	21.07	20.13	20.88	20.53	20.61	25.09	26.10
36–50	25.48	26.11	26.15	25.12	24.86	26.50	26.47	38.39	38.28
51–65	24.61	26.06	25.46	24.23	24.54	23.94	23.78	35.23	35.35
66–80	22.43	26.64	24.19	22.65	24.46	15.95	15.73	32.67	34.02
81–100	16.98	25.61	18.27	19.04	20.60	14.27	14.27	27.26	28.70
<i>Wealth</i>									
20–35	5.94	5.77	6.39	8.20	7.92	3.39	2.73	12.84	13.75
36–50	29.34	23.17	35.87	39.28	50.57	7.25	5.64	65.75	81.99
51–65	75.77	40.34	96.26	100.16	126.78	10.23	7.83	173.70	195.02
66–80	77.28	26.06	92.81	105.50	128.71	5.71	4.71	159.76	169.73
81–100	13.60	4.15	18.40	30.85	39.94	0.11	0.11	46.75	52.00
<i>Liquid portfolio share in stocks</i>									
20–35	1.00	0.93	1.00	0.97	0.97	0.99	0.99	0.57	0.53
36–50	0.99	1.00	0.98	0.95	0.72	1.00	1.00	0.91	0.68
51–65	0.88	0.81	0.61	0.61	0.08	1.00	0.90	0.57	0.14
66–80	0.90	0.96	0.85	0.57	0.49	1.00	1.00	0.54	0.51
81–100	0.92	0.58	0.90	0.68	0.57	1.00	1.00	0.61	0.53
<i>Tax rates (%)</i>									
	10.00	10.00	6.00	10.00	6.00	10.00	6.00	10.50	6.75

^aRefers to the scenario where the tax rate is held constant at the same level as in the 0/100 case, implying a higher average replacement ratio.

tion presented graphically in figures 11.2–11.4 above. The first three columns of table 11.6 use the baseline parameters and consider retirement systems with $\alpha^R = 0$, $\alpha^R = 0.5$ and the same 10 percent tax rate as with $\alpha^R = 0$, and $\alpha^R = 0.5$ and a lower 6 percent tax rate that maintains the same average replacement ratio as with $\alpha^R = 0$.

The next two columns of table 11.6 present results for a higher risk-aversion coefficient of 10 rather than 5. Since the tax rates do not depend on γ , they are the same as in the baseline case. To understand the results for higher γ , it is important to remember that, with isoelastic preferences, this parameter measures both risk aversion and prudence (Kimball 1990). Greater prudence increases precautionary savings and explains why highly risk-averse investors consume less, and save more, until age sixty-five. After this age, the precautionary savings motive is reduced since there is no labor income risk and retirement wealth is converted into a riskless annuity. Thus, highly risk-averse investors consume more after retirement.

Table 11.6 also shows that, as one would expect, highly risk-averse investors have a lower portfolio share in stocks. One interesting pattern that is not visible in the table is that, very early in life, these investors' equity portfolio share is increasing with age. This pattern does not show up for investors with $\gamma = 5$ because, in early life, these investors are constrained by their inability to borrow to finance equity investments. The reason for this increasing pattern is explained in CGM (1998). In the presence of an increasing labor income profile, the annuity value of future labor income, equivalent to implicit holdings of the riskless asset, increases with age at first as peak earnings years move closer in time. Investors respond to this increase by shifting liquid wealth toward risky financial assets. Later on, the annuity value of future labor income starts to decrease as peak earnings are realized and retirement approaches; investors respond by shifting out of stocks in middle age.

The next two columns of table 11.6 show optimal consumption, wealth, and portfolio allocation for impatient households with $\delta = 0.8$. These households consume more early in life (roughly up to age fifty) and less later. They accumulate almost no wealth, never holding more than about \$10,000 in liquid assets. What little wealth they do accumulate they hold in stocks; their exposure to the stock market is so small that they are extremely tolerant of equity risk.

The last two columns of table 11.6 report results for self-employed college-educated households. The preference parameters are the same as in the baseline case, but the results are quite different from the first three columns of table 11.6, which apply to households with only a high school education. The tax rate in the 50/50 system is set to 6.75 percent to maintain the average replacement ratio for self-employed college-educated households. Looking at the share of liquid wealth invested in stocks, there are two distinctive features: the share invested in risky financial assets is

much lower, and it exhibits a clear hump shape. The higher variance of labor income shocks and the large positive correlation between the latter and innovations to stock returns crowd out investment in risky financial assets. This effect is particularly strong early in life, when the investor has accumulated little liquid wealth.

11.4 Welfare Analysis

We have presented detailed results describing the effects of alternative retirement systems on consumption, wealth accumulation, and portfolio choice. We now turn to the welfare implications of these systems.

We evaluate each system by discounting current utilities back from the end to the beginning of adult life (twenty, or twenty-two for college-educated households). We renormalize discounted utility into consumption-equivalent units so that a 5 percent increase represents the increase in utility that would be produced by a 5 percent increase in consumption at every date. We calculate the expectation of discounted lifetime utility at age twenty across all realizations of income and risky-asset returns; in order to measure the variability of outcomes, we also calculate the standard deviation of discounted lifetime utility across these realizations. Finally, we assess welfare effects on retired households by repeating the same exercise discounting back only to age sixty-five.

The top panel of table 11.7 reports results for the benchmark case of a household with risk aversion $\gamma = 5$ and a high school education. In the top row, all retirement wealth is invested in the riskless asset; in the second row, half of retirement wealth is invested in stocks, but the tax rate is held constant at 10 percent; in the third row, half of retirement wealth is invested in stocks, and the tax rate is lowered to 6 percent to maintain the average replacement ratio at 60 percent; and, in the fourth row, all retirement wealth is invested in stocks, and the tax rate is lowered to 3.75 percent to maintain the average replacement ratio. The third and fourth columns give the mean and standard deviation of discounted lifetime utility in consumption-equivalent units (thousands of dollars of consumption) at age twenty, the fifth and sixth columns report the same moments at age sixty-five, and the remaining columns report the percentage changes in these moments relative to the benchmark case of riskless retirement wealth.

The table shows that a shift from a riskless retirement system to one that is 50 percent invested in stocks, with a lower payroll-tax rate, increases the welfare of the typical twenty-year-old household by 3.7 percent. Most of the welfare gain is due to the reduction in the tax rate; the household gains only 0.5 percent if the payroll-tax rate is held constant. The standard deviation of lifetime utilities across households with different income and asset-return realizations also increases modestly, by 3.4 percent if the payroll-tax rate is reduced and 2.9 percent if it is held constant.

Table 11.7 Welfare Analysis

	Tax Rate (%)	Welfare Age 20		Welfare Age 65		Gain Age 20 (%)		Gain Age 65 (%)	
		Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Baseline:									
0/100	10.00	18.67	1.75	22.69	3.42				
50/50 ^a	10.00	18.76	1.80	27.30	7.41	0.48	2.86	20.32	116.67
50/50	6.00	19.36	1.81	24.46	6.22	3.70	3.43	7.80	81.87
100/0	3.50	19.59	1.96	24.92	9.77	4.93	12.00	9.83	185.67
$\gamma = 10$:									
0/100	10.00	16.08	2.17	22.53	2.54				
50/50 ^a	10.00	16.05	2.19	27.54	6.62	-0.18	0.95	22.22	160.81
50/50	6.00	16.78	2.26	24.42	5.68	4.35	4.15	8.39	123.62
$\gamma = 2$:									
0/100	10.00	20.71	1.17	22.03	3.27				
50/50 ^a	10.00	21.18	1.30	29.86	8.57	2.31	11.15	35.57	161.92
50/50	6.00	21.50	1.33	23.11	6.25	3.83	13.79	4.91	90.94
		Welfare Age 22				Gain Age 22			
College-educated self-employed:									
0/100	10.50	17.64	2.90	31.20	5.09				
50/50 ^a	10.50	17.63	2.89	37.96	10.13	-0.07	-0.36	21.66	99.11
50/50	6.75	18.31	2.95	32.83	7.79	3.80	1.72	5.22	53.05

^aRefers to the scenario where the tax rate is held constant at the same level as in the 0/100 case, implying a higher average replacement ratio.

Households also prefer the risky retirement system at age sixty-five; their discounted utility is on average 7.8 percent higher if the payroll-tax rate is reduced and 20.3 percent higher if it is held constant (for then they have accumulated greater wealth over their working life and at age sixty-five are able to enjoy the benefits). A further shift in the retirement system to 100 percent equity investment, with a payroll-tax rate that is lower again, produces even larger utility gains but also considerably greater variability of outcomes.

The next panel of table 11.7 considers a highly risk-averse household with $\gamma = 10$. This household is not usually constrained in its equity holdings under the riskless retirement system, and it actually loses slightly if the retirement system is shifted into equities without any reduction in the payroll-tax rate. But, if the payroll-tax rate is reduced, the highly risk-averse household actually gains more than the benchmark household, 4.4 percent rather than 3.7 percent. The reason is that, under power utility, higher risk aversion implies a lower elasticity of intertemporal substitution—that is, a stronger desire to smooth consumption intertemporally—and the higher disposable income associated with a lower tax rate allows the household to smooth consumption more effectively over the life cycle.

The next panel of table 11.7 considers a comparatively risk-tolerant household with $\gamma = 2$. The results here are the mirror image of those for the highly risk-averse household. The risk-tolerant household gains more than the benchmark household if the payroll-tax rate is fixed (for it is particularly anxious to relax constraints on its equity holdings) but gains less if the payroll-tax rate is reduced (since improved opportunities to smooth consumption over the life cycle are less important for this household).

The last panel of table 11.7 reports results for a household headed by a self-employed college graduate. This household has risky labor income that is unusually correlated with the stock market, so its desired stockholdings are smaller than those of the benchmark household; it actually loses slightly from the investment of retirement wealth in equities with a fixed tax rate. On the other hand, this household also has a particularly pronounced hump shape in labor income, so it is particularly anxious to smooth consumption over the life cycle. The household gains 3.8 percent from a risky retirement system with a lower payroll-tax rate, comparable to the results for the benchmark household.

All the results in table 11.7 can be criticized on the grounds that they are derived from a model in which there is no role for a social security system. Since households save and invest their liquid wealth optimally, the mandatory saving and rigid asset allocation of the retirement system can only reduce their welfare. In this setting, any reform that effectively reduces the scale of social security will increase welfare.

As a partial response to this concern, in table 11.8 we report a “paternalistic” welfare analysis in which the government uses different utility parameters than those of the household itself. In the first panel, the household is impatient, with $\delta = 0.8$; left to its own devices, it will do very little saving for retirement, as illustrated in table 11.6 above. The government, however, discounts utility using $\delta = 0.96$, as in our benchmark case. Here, there is a modest welfare gain of 2.0 percent from a shift of retirement wealth toward equities with a fixed tax rate but a small welfare loss of -0.6 percent if the payroll-tax rate is reduced. The reason, of course, is that, from the government’s point of view, people make poor use of their tax cuts, spending them early in life and failing to save enough for retirement.

The second panel of table 11.8 reports results when the household is highly risk averse, with $\gamma = 10$, but the government is risk tolerant, with $\gamma = 2$. In this case, a shift of retirement wealth toward equities forces households to take on more risk, which improves their welfare from the government’s point of view. Welfare rises by 1.7 and 2.4 percent, respectively, in the cases with fixed and reduced payroll-tax rates.

All the results that we have reported so far assume that the riskless real interest rate is fixed at 2 percent and the equity premium at 4 percent. As

Table 11.8 Paternalistic Welfare Analysis

	Tax Rate (%)	Welfare Age 20		Welfare Age 65		Gain Age 20 (%)		Gain Age 65 (%)	
		Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
δ too low:									
0/100	10.00	18.13	1.64	17.26	1.31				
50/50 ^a	10.00	18.50	1.77	25.33	8.05	2.04	7.93	46.76	514.50
50/50	6.00	18.03	1.62	17.03	1.20	-0.55	-1.22	-1.33	-8.40
γ too high:									
0/100	10.00	20.67	1.09	20.26	2.75				
50/50 ^a	10.00	21.02	1.22	30.59	7.34	1.69	12.12	50.96	166.82
50/50	6.00	21.17	1.21	27.39	6.24	2.42	11.01	35.19	126.91

^aRefers to the scenario where the tax rate is held constant at the same level as in the 0/100 case, implying a higher average replacement ratio.

a final exercise, in table 11.9 we consider variations in these parameters. This is a valuable check on the robustness of our basic results; it also enables us to consider what might happen to welfare if the investment of retirement wealth in risky assets reduced the equilibrium equity premium.

The top panel of table 11.9 repeats the first three rows of table 11.7 for the benchmark case. The second and third panels consider two alternative scenarios in which the equity premium is 1 percentage point lower at 3 percent. In the first alternative, the equity premium falls, but the riskless interest rate is unchanged, so the expected equity return falls by 1 percentage point. This is a scenario envisaged by critics of social security investment in equities who worry that such a reform would drive up stock prices and drive down expected stock returns. In the second alternative, the equity premium falls, but the riskless interest rate rises by 1 percentage point, leaving the expected equity return unchanged. This scenario is predicted by general equilibrium models in which the return to risky capital is fixed by technology, such as Cox, Ingersoll, and Ross (1985) or Abel (chap. 5 in this volume).

Within each of the alternative scenarios, the welfare gains produced by risky investment of retirement wealth are similar to but slightly smaller than those in the benchmark case. In rows 6, 8, 11, and 13, table 11.9 compares welfare in the alternative scenarios with risky investment of retirement wealth to welfare in the benchmark case with riskless investment of retirement wealth. This is a crude way to capture the possibility that risky investment of retirement wealth might lower the equity premium. It turns out that the results are critically dependent on the way in which the equity premium falls. If it falls through lower stock returns, as in the first alternative, then welfare gains are reduced from 3.7 to 2.8 percent. If it falls through a higher riskless rate, as in the second alternative, then welfare gains are actually increased to 5.0 percent.

Table 11.9 Welfare Analysis with Alternative Mean Asset Returns

	Tax Rate (%)	Welfare Age 20		Welfare Age 65		Gain Age 20 (%)		Gain Age 65 (%)	
		Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Baseline case ($\mu = 4\%$, $\bar{R}_F = 1.02$):									
0/100	10.00	18.67	1.75	22.69	3.42				
50/50 ^a	10.00	18.76	1.80	27.30	7.41	0.48	2.86	20.32	116.67
50/50	6.00	19.36	1.81	24.46	6.22	3.70	3.43	7.80	81.87
Alternative 1 ($\mu = 3\%$, $\bar{R}_F = 1.02$):									
0/100	10.00	18.57	1.70	22.15	3.53				
50/50 ^a	10.00	18.63	1.78	25.13	6.64	0.33	4.43	13.48	87.89
^b 50/50						-0.22	1.73	10.75	94.05
^b 50/50	7.25	19.20	1.82	25.49	6.74	3.38	6.69	15.11	90.74
						2.82	3.93	12.35	97.00
Alternative 2 ($\mu = 3\%$, $\bar{R}_F = 1.03$):									
0/100	7.00	19.30	1.79	24.20	4.16				
50/50 ^a	7.00	19.32	1.84	26.76	6.88	0.12	2.81	10.60	65.38
^b 50/50						3.50	5.16	17.96	101.16
^b 50/50	4.75	19.61	1.83	25.07	6.11	1.61	2.23	3.60	46.88
						5.03	4.57	10.49	78.65

^aRefers to the scenario where the tax rate is held constant at the same level as in the 0/100 case, implying a higher average replacement ratio.

^bComputes the welfare gain relative to the baseline case, i.e., under the initial assumption for asset returns of $\mu = 4\%$ and $\bar{R}_F = 1.02$.

11.5 Conclusion

Decisions about the quantity and form of retirement saving are among the most important that a typical household takes in the course of a lifetime. Despite the importance of the issue, until very recently financial economists have had little quantitative understanding of the factors that should affect this decision. This gap in our knowledge has made it hard to give sound advice to policy makers considering reforms in retirement systems.

In this paper, we have built a partial equilibrium life-cycle model that can be used to explore the properties of alternative systems. In our benchmark case, we find a welfare gain equivalent to 3.7 percent of consumption from the investment of half of retirement wealth into equities, accompanied by a reduction in the social security tax rate to maintain the same average replacement rate of income in retirement. The main channel through which these gains are realized is that a lower social security tax rate helps households smooth their consumption over the life cycle. The gains from equity investment of retirement wealth are smaller, about 0.5

percent of consumption, when the social security tax rate is held constant at its initial level. Interestingly, in our model, particularly risk-averse households are particularly keen to smooth consumption and thus experience even larger gains from reduced tax rates made possible by equity returns on retirement wealth. This is true despite the fact that risk-averse households are less enthusiastic equity investors.

While these results are encouraging for proponents of equity investment by the social security system, we note two caveats. First, lower social security tax rates reduce welfare by 0.6 percent of consumption in a model in which investors are extremely impatient, with a low time-discount factor of 0.8, but the government judges their welfare using a higher time-discount factor of 0.96. This calculation is a crude way to capture factors that might lead to inadequate private saving and justify a mandatory retirement-saving system. In this model, however, there is a substantial gain of 2.0 percent from social security equity investment with fixed tax rates. These findings suggest that the appropriate adjustment of tax rates will depend on detailed assumptions about the behavior of households but that, under a wide range of assumptions, there are welfare gains to be had by investing some retirement wealth into equities.

Second, a system with partial investment of retirement wealth in the equity market has greater variability of outcomes across cohorts. Particularly negative outcomes for some cohorts might provoke pressure for political bailouts, and the anticipation of bailouts might change the consumption behavior modeled here. This is an important issue for research on social security reform.

Using data from the PSID, we study heterogeneity in labor income processes across households. We find that some households—particularly self-employed college graduates—are exposed to much greater volatility in their labor income than are typical households. The labor income of these households also tends to be more highly correlated with returns on stocks and long-term bonds. This heterogeneity affects optimal investment strategies and may help justify social security reform that includes an element of personal choice.

We also consider the possibility that the investment of retirement wealth in equities might reduce the equity premium. We do not build a general equilibrium model to study this issue, but we do compare results under alternative assumptions about the equity premium. We find that it matters greatly whether the equity premium falls through a decline in the expected return on stocks or through a rise in the riskless interest rate. In the former case, the welfare gains from investing retirement wealth in equities are reduced, while, in the latter case, they are actually greater than in our benchmark model.

In evaluating our results, it is important to be aware of several respects in which our model is oversimplified. First, we consider only self-financing

retirement systems in which there is no net payment from any household to any other. Thus, we ignore the redistributive features of the present social security system, and we have nothing to say about the overhang of liabilities to previous generations implied by the present system. Second, we assume that asset returns are independently and identically distributed. Thus, we ignore the variation in real interest rates and equity premia that is the subject of much recent research. Third, we abstract from the existence of owner-occupied housing. This is an important omission since housing is the main component of wealth for many people. Cocco (1998) takes a first step toward the realistic incorporation of housing into a life-cycle model. Fourth, we assume that labor income shocks have constant variances. Some researchers have argued that the variance of idiosyncratic shocks to labor income is higher when the economy is weak and risky-asset returns are low; this can have important effects on the demand for risky assets, as shown by Mankiw (1986), Constantinides and Duffie (1996), and Storesletten, Telmer, and Yaron (1998). Finally, we assume that labor income and the retirement age are exogenous to the household. Bodie, Merton, and Samuelson (1991) point out that households with flexible labor supply can afford to hold riskier portfolios because they can adjust to negative asset returns both by changing their consumption and by changing their labor supply. An important task for future research is to incorporate these and other realistic complications into the basic life-cycle model of portfolio choice.

References

- Abel, Andrew B. 1998. The aggregate effects of including equities in the social security trust fund. Wharton School, University of Pennsylvania. Typescript.
- Attanasio, Orazio. 1995. The intertemporal allocation of consumption: Theory and evidence. *Carnegie-Rochester Conference Series on Public Policy* 42:39–89.
- Balduzzi, Perluigi, and Anthony Lynch. 1999. Transaction costs and predictability: Some utility cost calculations. *Journal of Financial Economics* 52: 47–78.
- Barberis, Nicholas C. 2000. How big are hedging demands? Evidence from long-horizon asset allocation. *Journal of Finance* 55 (February): 225–64.
- Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro. 1997. Preference parameters and behavioral heterogeneity: An experimental approach in the Health and Retirement Study. *Quarterly Journal of Economics* 112:537–79.
- Bertaut, Carol C., and Michael Haliassos. 1997. Precautionary portfolio behavior from a life-cycle perspective. *Journal of Economic Dynamics and Control* 21: 1511–42.
- Blanchard, Olivier J. 1993. Movements in the equity premium. *Brookings Papers on Economic Activity*, no. 2:75–118.

- Bodie, Zvi, Robert C. Merton, and William Samuelson. 1991. Labor supply flexibility and portfolio choice in a life cycle model. *Journal of Economic Dynamics and Control* 16:427–49.
- Brandt, Michael W. 1999. Estimating portfolio and consumption choice: A conditional Euler equations approach. *Journal of Finance* 54 (October): 1609–45.
- Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado. 1997. Strategic asset allocation. *Journal of Economic Dynamics and Control* 21:1377–1403.
- Brennan, Michael J., and Y. Xia. 1998. Resolution of a financial puzzle. Anderson Graduate School of Management, University of California, Los Angeles. Typescript.
- Campbell, John Y. 1996. Understanding risk and return. *Journal of Political Economy* 104:298–345.
- Campbell, John Y., and N. Gregory Mankiw. 1989. Consumption, income, and interest rates: Reinterpreting the time series evidence. In *NBER Macroeconomics Annual*, ed. Olivier Blanchard and Stanley Fischer, 185–216. Cambridge, Mass.: MIT Press.
- Campbell, John Y., and Robert J. Shiller. 1998. Valuation ratios and the long-run stock market outlook. *Journal of Portfolio Management* 24:11–26.
- Campbell, John Y., and Luis M. Viceira. 1999. Consumption and portfolio decisions when expected returns are time varying. *Quarterly Journal of Economics* 114:433–95.
- . In press. Who should buy long-term bonds? *American Economic Review*.
- Carroll, Christopher D., and Andrew A. Samwick. 1997. The nature of precautionary wealth. *Journal of Monetary Economics* 40 (September): 41–71.
- Carroll, Christopher D., and Lawrence H. Summers. 1991. Consumption growth parallels income growth: Some new evidence. In *National savings and economic performance*, ed. B. D. Bernheim and John B. Shoven. Chicago: University of Chicago Press.
- Cocco, João F. 1998. Owner-occupied housing, permanent income, and portfolio choice. Harvard University. Typescript.
- Cocco, João F., Francisco J. Gomes, and Pascal J. Maenhout. 1998. Consumption and portfolio choice over the life-cycle. Harvard University. Typescript.
- Constantinides, George M., John B. Donaldson, and Rajnish Mehra. 1998. Junior can't borrow: A new perspective on the equity premium puzzle. NBER Working Paper no. 6617. Cambridge, Mass.: National Bureau of Economic Research.
- Constantinides, George M., and Darrell Duffie. 1996. Asset pricing with heterogeneous consumers. *Journal of Political Economy* 104:219–40.
- Cox, John, Jonathan E. Ingersoll Jr., and Stephen A. Ross. 1985. A theory of the term structure of interest rates. *Econometrica* 53:385–408.
- Deaton, Angus S. 1991. Savings and liquidity constraints. *Econometrica* 59: 1221–48.
- Fama, Eugene F., and G. William Schwert. 1977. Human capital and capital market equilibrium. *Journal of Financial Economics* 5:115–46.
- Gakidis, Harry. 1997. Earnings uncertainty and life-cycle portfolio choice. Massachusetts Institute of Technology. Typescript.
- Geanakoplos, John, Olivia S. Mitchell, and Stephen P. Zeldes. 1999. Social security money's worth. In *Prospects for social security reform*, ed. Olivia S. Mitchell, Robert J. Myers, and Howard Young. Philadelphia: University of Pennsylvania Press.
- Gourinchas, Pierre-Olivier, and Jonathan Parker. 1996. Consumption over the life cycle. Massachusetts Institute of Technology. Typescript.

- Heaton, John, and Deborah J. Lucas. 1997. Market frictions, saving behavior, and portfolio choice. *Macroeconomic Dynamics* 1:76–101.
- . 2000. Portfolio choice and asset prices: The importance of entrepreneurial risk. *Journal of Finance* 55:1163–98.
- Hubbard, Glenn, Jonathan S. Skinner, and Stephen Zeldes. 1994. The importance of precautionary motives for explaining individual and aggregate saving. *Carnegie-Rochester Conference Series on Public Policy* 40:59–125.
- Jagannathan, Ravi, and Zhenyu Wang. 1996. The conditional CAPM and the cross section of expected returns. *Journal of Finance* 51:3–53.
- Kim, Tong Suk, and Edward Omberg. 1996. Dynamic nonmyopic portfolio behavior. *Review of Financial Studies* 9:141–61.
- Kimball, Miles. 1990. Precautionary saving in the small and in the large. *Econometrica* 58:53–73.
- Liu, Jun. 1998. Portfolio selection in stochastic environments. Stanford University. Typescript.
- Mankiw, N. Gregory. 1986. The equity premium and the concentration of aggregate shocks. *Journal of Financial Economics* 17:211–19.
- Merton, Robert C. 1969. Lifetime portfolio selection under uncertainty: The continuous time case. *Review of Economics and Statistics* 51:247–57.
- . 1971. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory* 3:373–413.
- Pischke, Jörn-Steffen. 1995. Individual income, incomplete information, and aggregate consumption. *Econometrica* 63:805–40.
- Poterba, James M., and Andrew A. Samwick. 1997. Household portfolio allocation over the life cycle. NBER Working Paper no. 6185. Cambridge, Mass.: National Bureau of Economic Research.
- Ross, Stephen A. 1999. Samuelson's fallacy of large numbers revisited. Massachusetts Institute of Technology. Typescript.
- Samuelson, Paul A. 1969. Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics* 51:239–46.
- . 1991. Long-run risk tolerance when equity returns are mean regressing: Pseudoparadoxes and vindication of a "Businessman's Risk." In *Money, macroeconomics, and economic policy: Essays in honor of James Tobin*, ed. William C. Brainard, William D. Nordhaus, and Harold W. Watts. Cambridge, Mass.: MIT Press.
- Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron. 1998. Asset pricing with idiosyncratic risk and overlapping generations. Carnegie-Mellon University. Typescript.
- . 1999. The risk sharing implications of alternative social security arrangements. *Carnegie-Rochester Conference Series on Public Policy* 50:213–59.
- Tauchen, George, and R. Hussey. 1991. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. *Econometrica* 59:371–96.
- Viceira, Luis M. In press. Optimal portfolio choice for long-horizon investors with nontradable income. *Journal of Finance*.
- Wachter, Jessica. 1998a. Portfolio and consumption decisions under mean-reverting returns: An exact solution for complete markets. Harvard University. Typescript.
- . 1998b. Risk aversion and allocation to long-term bonds. Harvard University. Typescript.

Comment Amir Yaron

The goal of “Investing Retirement Wealth: A Life-Cycle Model” by John Y. Campbell, João F. Cocco, Francisco J. Gomes, and Pascal J. Maenhout is to analyze quantitatively the effects that different mandatory *retirement* plans have on savings, portfolio choice, and welfare. Campbell et al.’s analysis concentrates on retirement accounts that are based on a defined-contribution system. They ask what the ramifications are of investing at least part of the retirement account in equity, relative to investing it completely in risk-free bonds. Their analysis is motivated by the recent evidence on the projected decline in implicit rates of return on investments in social security and notable plans to alleviate these problems by privatizing components of social security or by investing part of the social security trust fund in equities (see Advisory Council on Social Security 1996). Both plans rely on the historically higher rates of return that stocks have yielded over Treasury bills—the “equity premium.”

It should be clear that, absent any frictions, changes in the investment profile of a defined-contribution retirement account will be completely neutralized by agents’ adjustments to the portfolio of their private savings. In the light of this, Campbell et al. go to great lengths to specify a rich labor supply process and various borrowing and stock market participation constraints. It is the interaction among these factors that governs the quantitative changes in optimal portfolio allocations and hence any changes in savings and welfare.

My analysis of the paper starts by summarizing Campbell et al.’s main findings. I then proceed to analyze, within a two-period model, the effects that borrowing constraints and fixed costs of entering into the stock market have on portfolio allocations—and point toward the factors that I believe yield the results that the authors obtain. The last part discusses some outstanding issues regarding their analysis.

Campbell et al.’s main results can be summarized as follows: (1) When fixed costs are *not* present, they find that most young agents hold 100 percent stocks. (2) Once fixed costs are introduced, most agents participate in the stock market by the age of thirty. (3) The responses of portfolio allocations are qualitatively similar to the case when there are no fixed costs. They find substantial welfare gains, on the order of 3.7–4.9 percent of annual consumption, when half of retirement savings is invested in equity and taxes are adjusted to equate replacement ratios at retirement. (4) Most of the welfare gains are due to the income effect and influence mainly the young. (5) They find evidence for significant differences in income processes across occupations and industrial sectors—a feature that alters agents’ portfolio choice.

Portfolio Choice

There are three factors that play a key role in Campbell et al.’s framework. The first two are borrowing constraints and fixed costs preventing agents from participating in the stock market. The third factor is the labor income profile that agents face. I find it useful to start by analyzing the implications of the borrowing and fixed-cost constraints within a simplified two-period life-cycle model. To do this, it is instructive to think of an environment where an agent maximizes expected utility over two periods by choosing the first-period rate of savings, κ , out of disposable wealth, $W - T$, and the proportion of savings invested in equity, α^L . The proportion of the retirement account that is invested in equity, α^R , is not under the agent’s control. The retirement account is financed by the tax collected in the first period, T . Formally, agents face the following problem:

$$\begin{aligned} &\max_{\{\kappa, \alpha^L\}} U(c_1) + \beta EU(c_2) \quad \text{subject to:} \\ &c_1 = (1 - \kappa)(W - T), \\ &c_2 = \kappa(W - T)R_p^L + TR_p^R, \\ &R_p^i = \alpha^i R + (1 - \alpha^i)R_f, \quad i = L, R, \\ &0 \leq \alpha^i \leq 1, \quad i = L, R, \\ &0 \leq \kappa \leq 1. \end{aligned}$$

The issues with which Campbell et al. are concerned can now be cast in terms of this framework. They ask in a more realistic life-cycle setup, How do κ and α^L change with age? How do κ and α^L respond to an increase in α^R ?

Campbell et al.’s setup is geared toward getting realistic age profiles for κ and α^L in the base case when the retirement accounts are fully invested in risk-free bonds. Table 11C.1 summarizes the changes in first-period consumption, $\Delta\kappa(W - T)$, and the share invested in equity in the liquid investment account, α^L , in response to an increase in the equity share of the retirement account—namely, $\Delta\alpha^R > 0$.

The first entry in table 11C.1, the northwestern quadrant, when there are no borrowing constraints and no fixed costs is the complete frictionless

Table 11C.1 Response to $\Delta\alpha^R > 0$ ($\Delta T = 0$, $\Delta \text{savings} = -\Delta C = \Delta\kappa[W - T]$)

	No Fixed Costs	Binding Fixed Costs
No borrowing constraints	$\Delta C = 0$ $\Delta\alpha^L < 0$	$\Delta C > 0$ $\Delta\alpha^L = 0$
Binding borrowing constraints	$\Delta C > 0$ $\Delta\alpha^L = 0$	$\Delta C = 0$ $\Delta\alpha^L = 0$

environment mentioned earlier. In this situation, agents undo the government's increased investment in equity in the retirement account. Hence, the amount invested in equity in the private liquid account must decline to ensure that the overall investment in risky equity out of total wealth is unchanged. Consequently, there are no changes in consumption and savings.

The southeastern block in table 11C.1 describes the situation in which agents both are constrained from borrowing and are facing binding fixed costs preventing entry into the stock market. In this case, a one-dollar switch into equity in the retirement account increases second-period income by $E(R) - R_f$ —the “equity premium.” Agents then desire to borrow and invest in stocks even more. However, as long as the stock market participation costs and borrowing constraints are still binding, agents are constrained at zero bond- and stockholdings in the liquid savings account ($B^L = S^L = 0$). As a consequence, agents cannot change their first-period consumption or savings.

The northeastern block of table 11C.1 describes the situation in which agents are not constrained from borrowing but are facing fixed costs that prevent them from entering the stock market. Again, agents benefit from the “equity premium,” namely, for each dollar invested in the retirement account, they now receive an extra $E(R) - R_f$ dollars in the second period. The fact that agents are constrained from investing in the stock market implies that, had this constraint not existed, they would save more by investing part of first-period wealth in equity. Agents do not choose to increase their bondholdings (and hence save more) because they were already at their constrained optimum before the change in the retirement portfolio. Therefore, the best that agents can do is to smooth consumption by increasing consumption in the first period. Hence, savings are reduced.

Finally, the southwestern block of the table describes the case where there are no fixed costs but the borrowing constraints are binding. In this case, once retirement accounts are invested in equity, agents can potentially reduce their increased second-period income by divesting away from equity. The fact that the borrowing constraint is binding implies that agents still prefer to buy more stocks but are constrained from doing so by the borrowing constraint. In the light of the increased income in the second period, and in order to smooth consumption, agents choose to increase consumption during the first period and continue to allocate 100 percent of their private savings to stocks.

The analysis presented above indicates that an increase in the share invested in equity in the retirement portfolio (i.e., a rise in α^R) should lead to a *reduction* in overall savings and a *reduction* in the share invested in equity in the private liquid accounts. This is true as long as there is some mass of agents who are completely unconstrained (the northwestern quadrant) and some others who are facing binding borrowing constraints but

no fixed costs or binding fixed costs without binding borrowing costs (the southwestern and northeastern quadrants, respectively). Campbell et al. obtain, however, a significant increase in overall savings (see their fig. 11.2 or fig. 11C.1 below) and relatively small changes in the share invested in equity in the liquid savings accounts. The key reason that Campbell et al. find an increase in savings, as opposed to what the simple model presented above suggests, is the stage in the life cycle in which the increased “equity-premium” income is received. In the analysis presented above, the benefits from investing in equity accrue during the second period—the retirement stage. On the other hand, Campbell et al. choose to fix the level of retirement benefits across the different investment experiments. With an increased equity share in the retirement accounts, lower taxes need to be levied in order to reach the level of retirement benefits available before. In the notation used above, this implies a smaller T —that is, Campbell et al. deliver the benefits during the first period. Receiving the benefits in the first period of life allows for a few effects. First, some agents may now be able to enter the stock market, whereas before they were constrained from doing so. Second, because agents now have relatively more income in the first period, they will save more during the first period in order to smooth consumption. This is the main channel that leads to the increase in savings that Campbell et al. find.

Another factor affecting Campbell et al.’s analysis that is absent in the simple two-period model presented above is the fact that labor income is stochastic and has a deterministic hump-shape component. As can be seen in figures 11.2 and 11.3, consumption tracks income very closely in the very early stages of life. The increased income, discussed above, translates almost entirely into an increase in consumption in those highly constrained early stages of life and only later to a rise in savings. The increase in utility shows up in terms of larger consumption during these first years of the life cycle. In summary, welfare, savings, and portfolio choice are quite sensitive to whether the benefits from the “equity premium” are given in the form of increased benefits at retirement or in the form of lower taxes during the working years. Abel (chap. 5 in this volume) makes similar arguments for explaining the different results that he gets concerning capital accumulations when he analyzes a defined-benefit and a defined-contribution system.

It would be interesting to get a quantitative feel for how much the welfare gains Campbell et al. report really depend on the variation in life-cycle labor income processes relative to the alleviation of the borrowing constraint. In figure 11C.1, I plot the combined investment in risky equity both in the liquid and in the retirement account relative to total wealth in both accounts. Namely, I plot $\alpha \equiv (S^L + S^R)/(W^R + W^L)$. Two key features emerge from this figure. First, in both cases, α , the fraction invested in equity as a fraction of total wealth, tends not to move as much over the

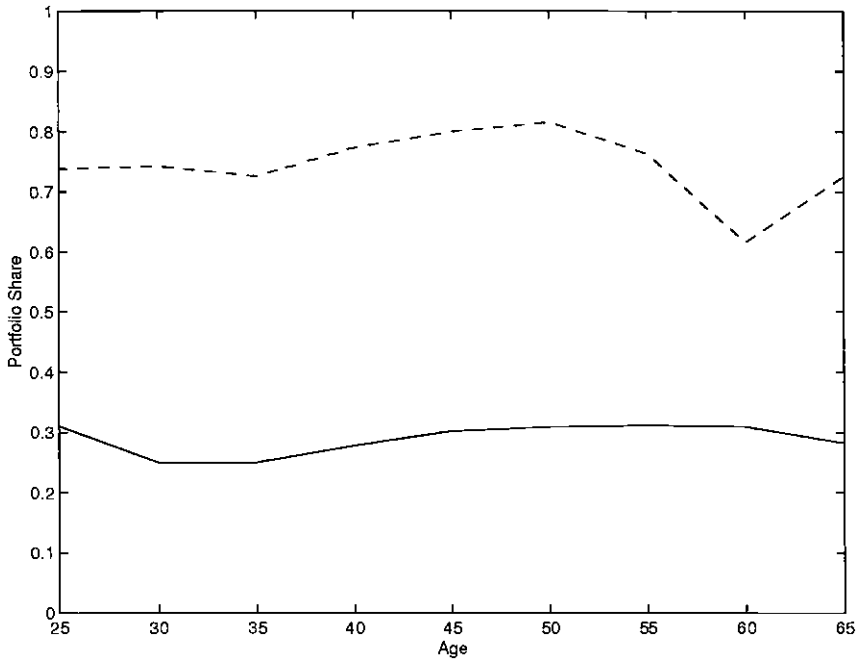


Fig. 11C.1 Portfolio share of stocks out of total wealth: $\alpha = (S^L + S^R)/(W^R + W^L)$

Note: α is based on figs. 11.2 and 11.3 in Campbell et al. and on linearly extrapolating income in between five years. The dotted line shows share out of just financial wealth, and the solid line shows total wealth.

life cycle as α^L (see Campbell et al.'s fig. 11.4). This is of course partly due to the fact that the denominator is now much larger as it includes both retirement and liquid wealth and to the fact that α^L offsets to some degree some of the changes in α^R . The second key observation is that α shifts up in almost parallel fashion once retirement accounts are invested in equity. As mentioned earlier, since agents like to be at 100 percent stocks or very close to it, investing in equity through the retirement account alleviates this constraint and allows agents to get closer to their unconstrained desired portfolio position. It seems that relaxing the borrowing constraint is the dominating factor for the welfare gains that Campbell et al. find while the specific variations in the labor income process are of secondary importance.

Welfare Gains—Other Factors

The welfare gains that Campbell et al. find are large. In the light of this, it is important to consider what factors might alter these welfare gains. Campbell et al. analyze a defined-contribution system. In a defined-benefit system, as in the current social security system, there are important non-

linearities in the benefits formula. First, there is cap for contributions to social security. This implies that wealthy agents are not as negatively affected by having retirement accounts being invested in bonds since their mandatory retirement account is a relatively small portion of their overall retirement plans. Second, social security currently provides a minimum level of benefits. For the very poor, having such a floor level of benefits may be more attractive than the gain from investing their very low level of wealth in equities. Both these factors make me think that, had the authors analyzed a system that is closer to the current one, the welfare gains would be much smaller.

A second important factor that may affect welfare gains, which Campbell et al. acknowledge is absent from their framework, is an account of the transition or debt that would be required in moving from the current system to the one that they in fact analyze. Storesletten, Telmer, and Yaron (1999) show that accounting for the debt that would need to be raised in order to keep up promises to the old is crucial in evaluating welfare gains across alternative social security plans. They show that accounting for the debt burden to the old can cut the welfare gains in half.

Finally, much of the discussion on restructuring social security deals with whether to invest the trust fund in equity or to allow a component of social security to be privatized in a similar fashion to the defined-contribution framework analyzed by Campbell et al. One benefit of the current social security system is its implicit longevity insurance through the annuity provisions that it provides. A crucial aspect of Campbell et al.'s analysis is the fact that, even when the system becomes more privatized, these actuarially fair annuities will be available. There is a long literature documenting how private retirement annuities are far from being fairly priced (see Friedman and Warshawsky 1990). If uninsurable labor income is an important component of overall income, as is the case in Campbell et al.'s analysis, the availability of annuities will crucially affect agents' savings. Storesletten, Telmer, and Yaron (1999) show that, in a general equilibrium framework, the availability of annuities is quantitatively important in comparing alternatives to social security. Hence, a fully privatized defined-contribution system, in which agents need to acquire retirement annuities privately, will not provide as much mortality insurance as the current system as long as private annuity markets are still not completely actuarially fair.

In summary, Campbell et al. provide a rich framework for evaluating the welfare benefits of increasing the equity share of investments in a defined-contribution retirement system. Their analysis provides interesting information about the channels and magnitude by which welfare is increased. A natural next step would be to account for some of the issues mentioned above and reevaluate their results in a general equilibrium framework.

References

- Advisory Council on Social Security. 1996. *Report of the 1994–1996 Advisory Council on Social Security*. Vols. 1 and 2. Washington, D.C.
- Friedman, B., and M. Warshawsky. 1990. The cost of annuities: Implications for saving behavior and bequests. *Quarterly Journal of Economics* 105:135–54.
- Storesletten, K., C. I. Telmer, and A. Yaron. 1999. The risk sharing implications of alternative social security arrangements. *Carnegie-Rochester Conference Series on Public Policy* 50:213–60.

Discussion Summary

James Poterba noted that the young participate very quickly in equity markets in this model. He conjectured that including housing in the analysis would probably eliminate this counterfactual prediction by effectively shortening the horizon of the young. He suggested examining this nontrivial extension of an already quite involved model in another paper. In such an extended model, a tax cut resulting from social security reform will allow young agents to accumulate wealth faster toward a downpayment. This would be an additional source of welfare gains.

Antonio Rangel remarked that, although the model does not assign an explicit role to the government, it could be used to answer an intriguing question, namely, What is the value of relaxing constraints? Rangel envisioned an analysis of the value of being able to buy some of the labor income of the next cohort, both in environments with and in those without constraints, with a subsequent comparison.

Robert Shiller remarked that the authors are probably overestimating the amount of heterogeneity in the PSID data set, for a number of reasons. First, the data are self-reported and therefore likely subject to measurement error. Second, a lot of income changes are essentially job changes, deliberately chosen by agents, rather than exogenous shocks hitting them. For instance, people might move to another city with a higher cost of living and therefore change jobs. Similarly, they might work overtime temporarily because of liquidity constraints due to the purchase of a house or drop out of the labor force for family reasons. Shiller concluded that, while the model does not include these changes in the information sets of agents, in fact they often are important.

Andrew Samwick wondered what could be the reason for the low degree of stock market participation observed among the young as opposed to the high demand for stocks predicted by the model. A first—but minor—reason could be that young households are not the natural tax clientele simply by having low income. Samwick conjectured that a more important

reason might be impatience. Impatient consumers would accumulate only a small buffer stock early in life, as opposed to what is found in the model assuming a relatively low discount rate.

Henning Bohn commented on the idea of making the payroll tax age dependent. Although this seems natural to consider in a model with liquidity constraints early in life, it might be less so in the context of the current system with redistributive components for strategic reasons. What makes the current system viable, according to a vast literature, is that the very young taxpayers making contributions constitute only a small minority. In addition, older taxpayers already have substantial resources sunk into the system by the time they become the median-age voter and have a vested interest in preserving the system. Lowering payroll taxes for the young could erode the support for the current system.

Robert King wanted to follow up on Bohn, but from a different perspective. An interesting result of the paper is that individuals are quite likely to be constrained during the first ten years of the life cycle. Baxter and King (chap. 10 in this volume) did similar simulations in the absence of any asset-allocation choice and obtained the same results, independent of preferences or of the annuity structure. They conducted a welfare analysis very much like the one Campbell et al. report and also obtained considerable welfare gains associated with relaxing those constraints. King concluded that this may well be one of the central insights of this paper and of this project. Practically, it would suggest that the U.S. system should be changed to allow individuals to postpone paying contributions until later in life, rather than to start contributing in the earliest, liquidity-constrained years.

Martin Feldstein remarked in response that anything that lowers the tax rate in equilibrium will have that effect and reduce the burden on the young. *Amir Yaron* concurred but noted that changing the timing of the contributions in an age-dependent way, as suggested by King, would be a more direct approach.

First, with respect to including housing into the analysis, *John Campbell* agreed that this would be an important extension of the model. He noted that one of the authors, João Cocco, has done work on this (see Cocco 1999). The effects of housing on portfolio choice are not straightforward and are more complex than would be anticipated. For one thing, it might actually increase risk taking by the young, and hence the demand for stocks, as owner-occupied housing, providing consumption services, might act as a consumption floor. *Francisco Gomes* also referred in this context to empirical work by Heaton and Lucas (1997) on the relation between mortgages and portfolio composition.

With respect to Samwick's comment that the high demand for stocks by the young might be due to the low discount rate assumed, *Campbell*

remarked that the results are quite robust to changes in this parameter. Previous work considered a rate of time preference of 0.10 and obtained similar results. *Pascal Maenhout* noted that Gakidis (1997) assumes an extreme discount rate of 0.40 in order to match the data in a calibration exercise using the method of simulated moments.

Regarding the discussion of the optimality of age-dependent tax rates, *Gomes* agreed that this was an important prediction of the model. Allowing the young to consume more leads to substantial welfare gains.

Finally, *Campbell* agreed with Shiller on the issue of overstating the amount of heterogeneity but noted that many other papers in a large literature follow a similar procedure and are thus subject to the same criticism.

References

- Cocco, João F. 1999. Owner-occupied housing, permanent income, and portfolio choice. Harvard University. Typescript.
- Gakidis, Harry. 1997. Earnings uncertainty and life-cycle portfolio choice. Massachusetts Institute of Technology. Typescript.
- Heaton, John, and Deborah J. Lucas. 1997. Portfolio choice and asset prices: The importance of entrepreneurial risk. Northwestern University. Typescript.