

Financial Risk, Retirement, Saving and Investment

Alan L. Gustman and Thomas L. Steinmeier



Project #: UM06-12

“Financial Risk, Retirement, Saving and Investment”

Alan L. Gustman
Dartmouth College and NBER

Thomas L. Steinmeier
Texas Tech University

September 2006

Michigan Retirement Research Center
University of Michigan
P.O. Box 1248
Ann Arbor, MI 48104
<http://www.mrrc.isr.umich.edu/>
(734) 615-0422

Acknowledgements

This work was supported by a grant from the Social Security Administration through the Michigan Retirement Research Center (Grant # 10-P-98358-5). The findings and conclusions expressed are solely those of the author and do not represent the views of the Social Security Administration, any agency of the Federal government, or the Michigan Retirement Research Center.

Regents of the University of Michigan

David A. Brandon, Ann Arbor; Laurence B. Deitch, Bingham Farms; Olivia P. Maynard, Goodrich; Rebecca McGowan, Ann Arbor; Andrea Fischer Newman, Ann Arbor; Andrew C. Richner, Grosse Pointe Park; S. Martin Taylor, Grosse Pointe Farms; Katherine E. White, Ann Arbor; Mary Sue Coleman, ex officio

Financial Risk, Retirement, Saving and Investment

Alan L. Gustman and Thomas L. Steinmeier

Abstract

This paper considers the prospects for adding choice of portfolio composition to a life cycle model of retirement and saving, while preserving the ability of the model to continue to explain the course of saving and retirement. If eventually successful, such a modification might be used to improve understanding of retirement and saving behavior both under the current Social Security system, and under variations involving personal accounts. In particular we consider the implications of separating parameters that now reflect both risk aversion and time preference.

We explore a number of barriers to developing a specification that is consistent with observed saving, retirement and investment choices. In our previous model with exponential consumption, individuals would hold portfolios exclusively in stocks, contrary to observation. Changing the exponent of consumption can reduce stock holdings below 100%, but at the cost of implausible retirement behavior. Introducing a separate parameter for risk aversion can restore plausible retirement behavior, but the pattern of stock holdings is too high, especially at younger ages, for plausible values of the risk aversion parameter.

At the moment, no easy solution is at hand to this fundamental problem now being engaged by financial economists. This suggests that models of retirement and saving may, for the immediate future, be forced to constrain portfolio composition to correspond with levels observed in the data, postponing the inclusion of portfolio mix as a choice variable until further progress is made in modeling that behavior. This does not necessarily reduce the efficiency of life cycle models of retirement and saving. Rather it recognizes that portfolio choice may be influenced by behavior that is not fully consistent with that posed by a life cycle model. Individuals may, for example, be accepting recommendations from planners or firms that they would not otherwise follow if they fully understood how to balance risk and return in portfolio choice in the same way they balance risk and return in their saving and retirement decisions. If these behavioral considerations govern their portfolio choice, while retirement and saving are determined by life cycle considerations, a model that correctly constrains portfolio composition may in fact generate parameter estimates that accurately reflect the forces governing retirement and saving behavior.

Financial Risk, Retirement, Saving and Investment

Alan L. Gustman and Thomas L. Steinmeier

In recent papers we have attempted to develop structural models that jointly explain retirement and saving, allowing the individual to choose among full time or part time labor market activities, and to choose whether to remain with their long term employer or leave that employer (e.g., see Gustman and Steinmeier, 2005a). We found that this approach is very useful for explaining such key features of retirement as the peak in retirements at age 62, a phenomenon that could not be explained by models that posit retirement is determined by the actuarial fairness or unfairness of the Social Security benefit structure. A model that jointly explains the effects of retirement and saving is also useful for understanding the likely effects of personal accounts and such features as the availability of lump sum payouts instead of mandatory annuitization (Gustman and Steinmeier, 2005b, 2005c). We also have modified that model to introduce uncertainty about returns to assets both before and after retirement and changes in leisure preferences after retirement (Gustman and Steinmeier, 2002, 2006), allowing us to explain reverse flows in retirement from states of lesser to greater work.

Although recent versions of our model treat returns to retirement assets as stochastic, the specification developed to date does not allow savers to choose the composition of their asset portfolios, whether held outside a retirement saving vehicle, or in 401k or other DC plans, IRAs or in other retirement saving accounts. Rather, although asset levels are endogenously determined as the individual chooses optimal levels of consumption, labor market activity and therefore saving, asset composition is set exogenously to conform to typical values found in empirical data.

This paper considers the prospects for adding choice of portfolio composition to a life cycle model of retirement and saving, while preserving the ability of the model to continue to explain the course of saving and retirement. If eventually successful, such a modification might be used to improve understanding of retirement and saving behavior both under the current Social Security system, and under variations involving personal accounts. In particular we consider the implications of estimating separate parameters to represent risk aversion and time preference, whereas one parameter now reflects both risk aversion and time preference.

We explore a number of barriers to developing a specification that is consistent with observed saving, retirement and investment choices. In our previous model with exponential consumption, individuals would hold portfolios exclusively in stocks, contrary to observation. Changing the exponent of consumption can reduce stock holdings below 100%, but at the cost of implausible retirement behavior. Introducing a separate parameter for risk aversion can restore plausible retirement behavior, but the pattern of stock holdings is too high, especially at younger ages, for plausible values of the risk aversion parameter.

The discussion is divided into five sections. Section I discusses risk and return tradeoffs to various portfolios. Section II integrates the trade off between risk and return into a standard utility function specification. The model is expanded in Section III to include portfolio investment in a life cycle model that also includes retirement and saving. Section IV derives the actual risk vs. return tradeoff and discusses the implications of the estimated tradeoff for portfolio choices and observed retirement and saving behavior, focusing on the challenge of having a given set of parameters explain the multiple outcomes of concern in a model of retirement, saving and investment choice. Since the model is dynamic, it is not only necessary

for outcome levels to be plausible, but for the paths of outcome to conform to observed data. Section V concludes.

I. The Risk-Return Tradeoff in Historical Data

In this section we examine five data series from the Ibbotson annual databooks. These series are the total returns from large stocks, corporate bonds, long-term government debt, intermediate-term government debt, and short-term government debt. All of these series are from 1926 through 2003. The large stock is roughly equivalent to the modern day S&P 500 index. Long-term government debt is at a maturity level of roughly 20 years. Intermediate-term government debt is at a maturity level of roughly 5 years, and short-term government debt is at a maturity level of roughly one month.

The annual rate of inflation is subtracted out from each series. The real returns and associated standard deviations of the series are calculated as in the following table:

	Total Returns				
	Large Stocks	Corp Bonds	LT Govt Debt	IT Govt Debt	ST Govt Debt
Means	9.19	3.23	2.81	2.53	0.78
Standard Deviations	20.54	9.82	10.52	6.94	4.03

The most risky asset, large stocks, had a mean real return of 9.19 percent, with a standard deviation of 20.54 percent. The safest asset, short-term government debt, had a much lower real return of 0.78 percent, with a standard deviation of 4.03 percent. Note that although long-term government debt is dominated in a risk-return sense by corporate bonds, the long-term government debt may still be desirable if it has desirable correlations with the other series of available assets, so we retain it for the time being.

The correlations of these series is given in the following table:

Correlations of Returns

	Large Stocks	Corp Bonds	LT Govt Debt	IT Govt Debt	ST Govt Debt
Large Stocks	1.0000	0.2582	0.2006	0.1414	0.1110
Corp Bonds		1.0000	0.9506	0.9449	0.5777
LT Govt Debt			1.0000	0.9400	0.5706
IT Govt Debt				1.0000	0.7210
ST Govt Debt					1.0000

These figures indicate that the returns of corporate bonds, long-term government debt, and intermediate-term government debt are all highly correlated. Stock returns are not very correlated with the returns from any form of debt, while short-term returns are somewhat correlated with the returns of longer-term debt.

With these returns and variances, it is possible to calculate portfolios that would minimize the standard deviations of returns for any given level of return. These portfolios and the associated mean return and its standard deviation are given in the following table:

Risk-Minimizing Portfolios

Return	Large Stocks	Corp Bonds	LT Govt Debt	IT Govt Debt	ST Govt Debt	Standard Deviation
0.80	0.2	0.0	0.0	0.2	99.6	4.03
1.20	4.9	0.0	0.0	0.5	94.6	4.07
1.60	8.5	0.0	0.0	6.0	85.5	4.32
2.00	11.2	0.0	0.0	15.9	72.9	4.69
2.40	13.9	0.0	0.0	25.8	60.3	5.15
2.80	16.7	0.0	0.0	35.2	48.1	5.67
3.20	19.4	0.0	0.0	45.1	35.5	6.24
3.60	22.1	0.0	0.0	54.9	23.0	6.85
4.00	24.9	0.0	0.0	64.3	10.8	7.49
4.40	28.1	0.0	0.0	71.8	0.1	8.14
4.80	34.1	0.0	0.0	65.8	0.1	8.89
5.20	40.1	0.0	0.0	59.9	0.0	9.74
5.60	46.1	0.0	0.0	53.9	0.0	10.66
6.00	52.1	0.1	0.0	47.8	0.0	11.65
6.40	58.1	0.1	0.0	41.8	0.0	12.68
6.80	64.0	1.1	0.0	34.9	0.0	13.74
7.20	69.2	8.8	0.0	22.0	0.0	14.83
7.60	74.1	19.3	0.0	6.6	0.0	15.94

8.00	80.1	19.4	0.0	0.5	0.0	17.05
8.40	86.8	12.8	0.0	0.4	0.0	18.20
8.80	93.5	6.2	0.0	0.3	0.0	19.37

For example, to get a return of 4.0%, a portfolio of 24.9 percent large stocks, 64.3 percent intermediate-term government debt, and 10.8 percent short-term government debt minimizes the standard deviations of the overall returns and yields a standard deviation of 7.49 percentage points.

This analysis bears out that corporate bonds dominate long-term government debt, so we will drop long-term government debt from further analysis. At low interest rates, short-term debt predominates, while intermediate-term debt predominates at intermediate returns and stocks at high returns. The percentage of stocks in the optimal portfolio steadily increases throughout the range of returns, and it is already a major component of the portfolio even at intermediate returns. This is undoubtedly due to its lack of correlation with the returns of the other kinds of assets available in the portfolio.

Corporate bonds are absent from the optimal portfolio at most rates of return. They do enter at relatively high rates of return, but even here they are not a major component, never reaching more than 20 percent of the optimal portfolio. This leads to the suspicion that eliminating them from consideration might not affect the risk of the optimal portfolio very much. To test this, the optimal portfolio analysis is repeated, but this time forcing corporate bonds to have zero weight in the portfolio. The results are in the following table:

Risk-Minimizing Portfolios Excluding Corporate Bonds				
Return	Large Stocks	IT Govt Debt	ST Govt Debt	Standard Deviation
0.80	0.2	0.2	99.6	4.03
1.20	4.9	0.5	94.6	4.07
1.60	8.5	6.0	85.5	4.32

2.00	11.2	15.9	72.9	4.69
2.40	13.9	25.8	60.3	5.15
2.80	16.7	35.2	48.1	5.67
3.20	19.4	45.1	35.5	6.24
3.60	22.1	54.9	23.0	6.85
4.00	24.9	64.3	10.8	7.49
4.40	28.1	71.8	0.1	8.14
4.80	34.1	65.8	0.1	8.89
5.20	40.1	59.9	0.0	9.74
5.60	46.1	53.9	0.0	10.66
6.00	52.2	47.4	0.4	11.66
6.40	58.2	41.5	0.3	12.69
6.80	64.2	35.5	0.3	13.76
7.20	70.2	29.5	0.3	14.85
7.60	76.2	23.5	0.3	15.97
8.00	82.2	17.5	0.3	17.10
8.40	88.2	11.6	0.2	18.25
8.80	94.2	5.6	0.2	19.41

Without corporate bonds, at returns above 7 percent, some of the weight which in the previous table that had been assigned to corporate bonds is transferred to stocks, but most of the weight is transferred to intermediate-term government debt. The most important feature of this table is that the risk-minimizing standard deviation hardly increases at all if corporate bonds are omitted from the mix. At a rate of return of 8.0 percent, corporate bonds are 19.4 percent of the optimal portfolio, and the standard deviation of the return for the portfolio is 17.05 percentage points. If corporate bonds are excluded, the standard deviation of the return for the portfolio rises only trivially, to 17.10 percentage points. This suggests that while evaluating the risk-return tradeoffs, corporate bonds can probably be omitted from the analysis, and the analysis can concentrate on stocks, intermediate-term debt, and short-term debt.

An additional piece of evidence which tends to support this analysis comes from the distributions of government debt by maturity. The analysis above suggests that intermediate-term debt turns out to be the most important component of the government debt, followed by short-term debt, and with long-term debt playing only a minor role. According to the Treasury (www.publicdebt.treas.gov/opd/opddload.htm), the public debt held by the public as of July 31,

2006 included 0.9 trillion dollars of bills with a maturing of less than one year, 2.4 trillion dollars of notes maturing in less than 10 years, with an average of around 5 years, and 0.5 trillion dollars in bonds, which mature in 30 years or less and with a probable average of around 15 years. Since the average maturities of the bills, notes, and bonds correspond roughly to the short-term, intermediate-term, and long-term government debt used by Ibbotson, the distribution of the treasury debt corresponds in a rough way to what would probably be demanded in optimal portfolios.

For the dynamic model, all that is really required from this analysis is the relationship between the expected returns and the standard deviations of those returns. For any given level of assets going into the next period, the model can look at the menu of expected returns and the associated standard deviations and pick the one which has maximum value. The optimal mix of stocks and intermediate and short term debt is the one corresponding to the expected returns and standard deviations. The methodology for doing this is outlined in the next section.

II. Risk and Return in the Utility Function

Suppose that the utility function, as a function of assets, can be locally approximated by a quadratic function around some initial level of assets A_0 . Then the marginal utility can be locally approximated by a linear function. Suppose that the marginal utility of asset level A_0 is given by a , and the slope of the marginal utility function is given by $-b$, where b is a positive parameter. Then the marginal utility as a function of assets can be written as

$$\begin{aligned}
 (1) \quad \text{MU} &= a - b(A - A_0) \\
 &= (a + b A_0) - b A
 \end{aligned}$$

Integrating with respect to A gives, up to a constant, the utility function

$$(2) \quad U = (a + b A_0) A - (1/2) b A^2$$

Suppose that the return next period to \$1 in assets today is normally distributed with a mean of μ and a standard deviation of σ . Then, if we start out with A_0 assets today, next period those assets A_{t+1} will have a normal distribution with a mean of μA_0 and a standard deviation of σA_0 . The expected utility of these assets next period is given by

$$(3) \quad \begin{aligned} E(U_{t+1}) &= E \{ (a + b A_0) A_{t+1} - (1/2) b (A_{t+1})^2 \} \\ &= (a + b A_0) E(A_{t+1}) - (1/2) b E[(A_{t+1})^2] \\ &= (a + b A_0) \mu A_0 - (1/2) b [(\mu A_0)^2 + (\sigma A_0)^2] \\ &= (a + b A_0) \mu A_0 - (1/2) b (\mu^2 + \sigma^2) (A_0)^2 \end{aligned}$$

using the standard relation for the expected values of moments of a normal distribution.

Now suppose that there is a relationship between the mean of the distribution and the standard deviation. For the moment, let that relationship be described by a linear relationship:

$$(4) \quad \sigma = c + d \mu$$

The object of the exercise is to choose that value of μ (and the associated portfolio and standard deviation σ) which maximizes the expected value in equation (3). Substituting from equation (4) into equation (3) yields

$$(5) \quad \begin{aligned} E(U_{t+1}) &= (a + b A_0) \mu A_0 - (1/2) b [\mu^2 + (c + d \mu)^2] (A_0)^2 \\ &= (a + b A_0) \mu A_0 - (1/2) b [(1 + d^2) \mu^2 + 2 c d \mu + c^2] (A_0)^2 \end{aligned}$$

Taking the derivative of equation (5) with respect to μ and setting it equal to zero yields:

$$d E(U_{t+1}) / d \mu = (a + b A_0) A_0 - b [(1 + d^2) \mu + c d] (A_0)^2 = 0$$

Solving for μ yields:

$$(6) \quad \mu = \frac{1}{(1 + d^2)} \left[\frac{1}{A_0} \frac{a}{b} + (1 - cd) \right]$$

as the optimal expected rate of return, and the optimal mix of assets is the mix associated with the minimum standard deviation for this rate of return.

There are several things to note about this result:

1. If the parameter b , which measures the slope of the marginal utility function, is lower, then the optimal value of μ will be higher, and the optimal portfolio will be more weighted towards risky stocks and less toward less risky government debt. This is to be expected, since a lower value of b reflects a less risk averse utility function.
2. The optimal return depends on the ratio of the marginal utility at A_0 (a) to the slope of the marginal utility function (b), but not to a and b separately. Another way of putting this is that anything which shifts the marginal utility function proportionately, and hence changes the marginal utility by the same percent as the slope changes, will not affect the optimal return μ or the optimal portfolio allocation.
3. If the utility function is CRRA in assets, the optimal portfolio should be approximately the same regardless of the level of assets. To see this, consider the CRRA function

$$U = \frac{1}{1 - \gamma} A^{1 - \gamma}$$

$$MU = a = A^{-\gamma}$$

$$d MU / d A = -b = -\gamma A^{-\gamma - 1}$$

Taking the ratio and evaluating at A_0 yields:

$$a / b = A_0 / \gamma .$$

Rearranging yields

$$(1 / A_0) (a / b) = 1 / \gamma$$

Substituting into equation (6) gives

$$(7) \quad \mu = \frac{1}{(1+d^2)} \left[\frac{1}{\gamma} + (1 - cd) \right]$$

This expression depends on the parameters of the risk-return tradeoff and the coefficient of relative risk aversion, but not on the level of assets. Hence, if the utility of asset function is given by a constant relative risk aversion function, the optimal return and the optimal mix of assets will be approximately constant.

4. Before retirement, one can always at least partially offset a realization of better or worse than expected returns by retiring a little sooner or later. Put another way, the ability to adjust the retirement date provides some insurance against the risk of asset returns. This insurance should flatten to some degree the marginal utility curve. After final retirement, of course, the ability to absorb bad realizations of returns by changing the retirement date vanishes, so that the marginal utility curve is steeper. All this is to say that for a given level of assets, the marginal utility curve should be flatter before retirement than after retirement. In terms of the model, the parameter b (which is the negative of the slope of the marginal utility function) will be lower before retirement than after. Since there is an inverse relation between b and μ through equation (6), this implies that the optimal value of μ will be higher before retirement than after. This in turn means that the optimal portfolio before retirement will be riskier, with a higher percentage of stocks, before than after retirement.

III. Modifications to the Empirical Model

The analysis of the previous section has serious implications for the empirical model we have been using to date. In previous estimates, the exponent of the consumption term has always been near zero, so that the value of γ in the previous section would be near unity. Looking at

the tables in the first section, an increase of 7 percentage points in the returns results in an increase in the standard deviation of about 15 percentage points, so that the relationship between the mean and the standard deviation is approximately $\sigma = 0.04 + 2(\mu - 1) = -1.96 + 2\mu$. (Recall that μ is the total amount received for an investment of \$1 and hence includes the principal.) This means, in the notation of the previous section, that $c = -1.96$ and $d = 2$.

Plugging these two values into equation (7), and using $\gamma = 1$, yields a value of μ equal to about 1.184, which corresponds to an annual return of 18.4%. That is, with $\gamma = 1$, and using the actual risk-return tradeoff, individuals would want to be 100% in stocks. This result is consistent with two papers on optimal portfolio allocation, those of Benzoni, Dufresne, and Goldstein (2005) and Gomes and Michaelides (2005). In order for individuals to be induced to hold financial assets other than stocks, it seems that the coefficient of relative risk aversion parameter has to be approaching a value of 5.

Having a value of γ equal to 5, however, creates other difficulties for the model. In our current model, this value would mean that additional consumption would have sharply lower marginal utility at higher income and asset levels, so that individuals with higher incomes would want to retire much earlier. In fact, with a sufficiently high value of γ , individuals would be approximately target income workers. That is, they would work long enough to achieve a target level of income, and any income beyond that would produce so little utility that they would rather take the leisure. Hence, a high value of γ is inconsistent in our current model with the fact that individuals retire at more or less the same time on average, regardless of their level of income.

These results suggest that if we are going to examine both retirement behavior and asset allocation in the same model, we will need to modify our original model somewhat. The most

straightforward way to do this appears to be to introduce an additional parameter to separate risk aversion from the marginal utility of consumption. To do this, introduce a value function for consumption:

$$V_t^C = \left\{ C_t^{1-\gamma} + \beta E_t \left[(V_{t+1}^C)^{1-\gamma} \right] \right\}^{\frac{1}{1-\gamma}}$$

Through recursive substitution, this can be shown to be equivalent to

$$V_t^C = \left\{ C_t^{1-\gamma} + \sum_{i=1}^D \beta^i E_t \left[C_{t+i}^{1-\gamma} \right] \right\}^{\frac{1}{1-\gamma}}$$

In this form, it can be seen that the value function for consumption is constant returns to scale, that is, if consumption and expected consumption double at every point in time, the value function also doubles.

A value function for leisure can be defined in a corresponding manner:

$$\begin{aligned} V_t^L &= h_t L_t + \beta E_t (V_{t+1}^L) \\ &= h_t L_t + \sum_{i=1}^D \beta^i E_t (h_{t+i} L_{t+i}) \end{aligned}$$

where $h_t = e^{X_t \delta + \varepsilon_t}$. Since L is restricted between zero and one, this function does not need the exponents.

In each period, the individual is presumed to maximize

$$(8) \quad V_t = \frac{1}{1-\alpha} (V_t^C)^{1-\alpha} + V_t^L$$

The exponent α on the consumption term serves much the same effect that the same coefficient did in the previous model. If the value of α is less than one, then the marginal utility of consumption will be declining slowly, and the utility value of the income from a year's worth of work will be increasing with the lifetime wage rate. This should cause high-income individuals

to retire later. The reverse is true if the value of α is above one; in this case, the utility value of the income from a year's worth of work will be declining at higher wage levels, and high-income individuals should retire earlier.

In this revised model, α serves much the same purpose that it has in previous models. The value of γ , on the other hand, serves the purpose of establishing risk aversion. Note that, at time t , and given the initial assets and the current income, the consumption decision boils down to maximizing

$$\frac{1}{1-\gamma} \left\{ C_t^{1-\gamma} + \sum_{i=1}^D \beta^i E_t [C_{t+i}^{1-\gamma}] \right\}$$

This is exactly the same consumption problem that was addressed in the previous section, and the same solution methods can be used to establish the optimal choice of assets. Given the asset allocation decision, the consumption/savings decision boils down to determining the level of savings for which the value of a dollar of current consumption in the above formula just matches the discounted expected value of the future consumption which would result if the dollar were saved. This is exactly the same problem that has been dealt with in previous models. The current income is determined by looking at the total value of the alternative retirement decisions and choosing the decision with the highest value, as has been done in previous models.

In practice, then, this model has two differences from the previous models. First, there is the asset allocation decision, which can be addressed using the techniques described in the previous section. Second, there is the decision regarding the alternative retirement states. In the previous model, this was solved basically by calculating the total expected utility value for each of the potential retirement states and choosing whichever was highest. The same kind of thing can be done with the revised model, except that the value in equation (8) is used for each

alternative. The additional parameter can be addressed by introducing another moment into the estimation process, namely the percentage of assets held in stocks.

IV. Implications for Portfolio Choice

The actual relationship between return and risk, as calculated in the Section II, is not exactly linear, so these results of that section hold only approximately. In the numerical simulations using the risk-return tradeoff, the following procedure is used. First, note that equation (3) was derived without using the linear risk-return tradeoff. At the point where the portfolio decision has to be made for a particular level of assets, the model will have already calculated the marginal utility of assets and its slope. Using equation (3), the expected utility for a particular point of the tradeoff curve is calculated. Using a grid of points similar to those generated for the tables at the end of the last section, we then evaluate which risk-return point yields the highest expected utility. This then would be the optimal expected return, and the associated portfolio mix is the optimal mix for the particular level of assets.

The results of these calculations are described in the table below:

MU Elas- ticity wrt Wealth	Rate of Return	Pct Stocks
0.5	9.19	100.0
1.0	9.19	100.0
1.5	8.93	96.1
2.0	7.37	72.7
2.5	6.43	58.6
3.0	5.81	49.3
3.5	5.36	42.5
4.0	5.03	37.5
4.5	4.77	33.6
5.0	4.56	30.5
5.5	4.39	27.9
6.0	4.12	25.7
6.5	3.84	23.8
7.0	3.60	22.1
7.5	3.40	20.8

The first column is the elasticity of marginal utility with respect to wealth. In a utility function that is crra in wealth, this is equal to the parameter γ . The next two columns are the optimal rate of return and the corresponding percentage of stocks in the portfolio. At low values of γ , the optimal portfolio consists entirely of stocks. At values of γ above about 1.5, the percentage of stocks begins to decline. It reaches 50% at a value of γ around 3, and by a value of γ equal to 5 it is down to around 30%. Note that at even at extremely high values of γ , the percentage of stocks is still nontrivial.

So far, so good. A value of 3 for γ is certainly within the range of estimates in the literature, and it would yield a percentage of stocks in the portfolio of around 50%. However, there is a catch. The value of γ so far refers to the elasticity of marginal utility with respect to wealth, but the γ 's in the model refer to consumption. That is, we and others suppose that the utility of consumption is given as proportional to:

$$\frac{1}{1-\gamma} C^{1-\gamma} ,$$

whereas the γ on page 9 refers to assets.

To see what difference this makes, we run the model (with fixed retirement) using a value of γ (for consumption) equal to 5, which is toward the high end of the values used in the literature. The wages for the individual are relatively high: \$100,000 at age 25, increasing by one percent per year in real terms. The individual works until age 64 and then retires. Joint Social Security benefits are \$25,000 per year, with \$14,000 for the survivor. To keep things simple, there are no spouse earnings and no pensions. Ex post returns are the expected rate of returns for the portfolio, and the rate of time preference is zero. The simulation yields the following time path, conditional on both spouses surviving:

MU Elas-

Age	Year	ticity wrt Wealth	Pct Stocks	Rate of Return	Wealth plus Returns	Income	Con- sumption	Wealth at End of Year
25	1961				0	100000	100000	0
26	1962	0.00	100.0	7.2	0	101000	101000	0
27	1963	0.00	100.0	7.2	0	102010	102010	0
28	1964	0.00	100.0	7.2	0	103030	103030	0
29	1965	0.00	100.0	7.2	0	104060	104060	0
30	1966	0.00	100.0	7.2	0	105101	104796	305
31	1967	0.00	100.0	7.2	326	106152	104796	1682
32	1968	0.01	100.0	7.2	1803	107214	104794	4223
33	1969	0.01	100.0	7.2	4526	108286	104788	8024
34	1970	0.02	100.0	7.2	8600	109369	104776	13193
35	1971	0.04	100.0	7.2	14141	110462	104757	19846
36	1972	0.06	100.0	7.2	21271	111567	104730	28108
37	1973	0.08	100.0	7.2	30126	112683	104694	38114
38	1974	0.10	100.0	7.2	40851	113809	104648	50012
39	1975	0.13	100.0	7.2	53602	114947	104592	63957
40	1976	0.17	100.0	7.2	68550	116097	104526	80120
41	1977	0.20	100.0	7.2	85873	117258	104450	98681
42	1978	0.25	100.0	7.2	105766	118430	104364	119833
43	1979	0.29	100.0	7.2	128437	119615	104270	143782
44	1980	0.35	100.0	7.2	154105	120811	104170	170747
45	1981	0.40	100.0	7.2	183006	122019	104068	200957
46	1982	0.47	100.0	7.2	215386	123239	103967	234658
47	1983	0.54	100.0	7.2	251507	124472	103863	272115
48	1984	0.61	100.0	7.2	291653	125716	103773	313595
49	1985	0.70	100.0	7.2	336112	126973	103706	359379
50	1986	0.79	100.0	7.2	385183	128243	103650	409776
51	1987	0.89	100.0	7.2	439197	129526	103643	465080
52	1988	1.01	100.0	7.2	498472	130821	103679	525614
53	1989	1.13	100.0	7.2	563353	132129	103788	591694
54	1990	1.27	100.0	7.2	634178	133450	103990	663638
55	1991	1.43	98.7	7.1	710861	134785	104296	741350
56	1992	1.61	89.4	6.7	790760	136133	104617	822276
57	1993	1.80	79.8	6.2	873217	137494	104937	905774
58	1994	2.01	72.5	5.8	958649	138869	105283	992236
59	1995	2.22	65.6	5.5	1046841	140258	105660	1081438
60	1996	2.46	59.6	5.2	1137758	141660	106053	1173366
61	1997	2.70	54.4	5.0	1231495	143077	106497	1268075
62	1998	2.96	49.8	4.7	1328057	144508	106997	1365568
63	1999	3.25	45.6	4.5	1427367	145953	107434	1465886
64	2000	3.54	42.1	4.4	1529713	25000	107895	1446817
65	2001	3.55	42.0	4.3	1509744	25000	108347	1426397
66	2002	3.56	41.9	4.3	1488358	25000	108782	1404576
67	2003	3.57	41.8	4.3	1465512	25000	109200	1381311
68	2004	3.58	41.7	4.3	1441163	25000	109600	1356563
69	2005	3.59	41.6	4.3	1415269	25000	109985	1330284
70	2006	3.60	41.4	4.3	1387780	25000	110199	1302581
71	2007	3.61	41.3	4.3	1358783	25000	110377	1273406
72	2008	3.62	41.1	4.3	1328254	25000	110511	1242743
73	2009	3.63	41.0	4.3	1296173	25000	110597	1210576
74	2010	3.64	40.8	4.3	1262528	25000	110615	1176913
75	2011	3.65	40.7	4.3	1227381	25000	110579	1141802
76	2012	3.65	40.7	4.3	1190726	25000	110484	1105242
77	2013	3.66	40.6	4.3	1152558	25000	110328	1067230

78	2014	3.66	40.5	4.3	1112875	25000	110102	1027773
79	2015	3.67	40.4	4.3	1071682	25000	109800	986882
80	2016	3.68	40.3	4.3	1028992	25000	109409	944583
81	2017	3.68	40.2	4.3	984845	25000	108913	900932
82	2018	3.68	40.2	4.3	939347	25000	108326	856020
83	2019	3.68	40.3	4.3	892523	25000	107640	809883
84	2020	3.68	40.2	4.3	844410	25000	106846	762565
85	2021	3.68	40.2	4.3	795055	25000	105916	714138
86	2022	3.68	40.3	4.3	744593	25000	104862	664731
87	2023	3.67	40.4	4.3	693111	25000	103679	614432
88	2024	3.67	40.4	4.3	640676	25000	102339	563337
89	2025	3.66	40.5	4.3	587430	25000	100829	511601
90	2026	3.65	40.7	4.3	533531	25000	99157	459373
91	2027	3.64	40.8	4.3	479083	25000	97276	406807
92	2028	3.62	41.1	4.3	424329	25000	95219	354110
93	2029	3.61	41.3	4.3	369391	25000	92927	301464
94	2030	3.58	41.7	4.3	314525	25000	90421	249103
95	2031	3.55	42.0	4.3	259931	25000	87680	197251
96	2032	3.51	42.4	4.4	205864	25000	84701	146163
97	2033	3.47	42.9	4.4	152582	25000	81490	96092
98	2034	3.42	43.4	4.4	100339	25000	78049	47289

The first two columns are self-explanatory. The third column gives the elasticity of marginal utility with respect to wealth, which is the quantity $(bA_0)/a$ in the previous discussion. The fourth and fifth columns give the percent stocks and the associated rate of return of the portfolio. The remaining columns give the real wealth from the previous period augmented by the returns, real income, real consumption, and the real wealth at the end of the period, which is the augmented wealth plus income minus consumption.

The consumption and wealth accumulation of this individual look pretty much as one might expect. With a time preference parameter of zero and a high value of γ , the individual really strives to maintain a relatively even level of consumption. As expected, the individual accumulates enormous amounts of wealth, equal to around 10 times income just before retirement.

At first glance, the portfolio behavior of this individual also makes sense. He starts off with a holding of 100% stocks, and in the few years before retirement at age 64 he reduces this to 40% stocks. After that, the percentage in stocks declines a little further, but not much. The

surprising thing about this behavior is the suddenness with which the stock holdings decline. Most investment advice, such as the $110 - \text{age}$ formula, yields a much more gradual reduction of stocks in the portfolio. In contrast, the model simulation gives a 60% reduction in stocks in the span of around 8 years.

The suddenness of the decline presents two problems for the model. First, according to Gomes and Michaelides (2005), among those who participate in the stock market, about 50-60% of their holdings are in stocks, and this fraction does not vary much by age. Up until a few years before retirement, though, the model predicts that 100% of financial wealth should be in stocks. A second problem is that the model will have difficulty matching the 50% stock holdings observed in the sample in 1992. The portfolio does not reach 50% stock until a couple of years before retirement; six years before retirement the portfolio is 72% stock. Since the sample was on average 56 in 1992, and the median retirement age is 62, this portends trouble for any estimation procedure in matching the observed 50%. The problem is compounded because this individual, with enormous wealth, high income, no pensions, and a low time preference rate, has a combination of characteristics particularly unfavorable to holding stock, as will be explained shortly. Almost any other combination of circumstances and preferences would result in even higher percentages of stock in the portfolio.

The crucial lesson from this exercise is that a high elasticity of marginal utility with respect to consumption (in this case 5) is entirely consistent with a relatively low elasticity of marginal utility with respect to wealth. At age 56, with a wealth-income ratio of over 5, the elasticity of marginal utility with respect to wealth is only 1.70. To understand what is going on in the model, it is necessary to understand this discrepancy.

Consider, for an individual with \$50,000 in consumption and \$5,000 in wealth, the effects of another \$1,000 in wealth. The extra \$1,000 in wealth represents a 20% increase in wealth. But in terms of resources, it increases the amount from \$55,000 to \$56,000, an increase of 1.8%. Over this range, the marginal utility of consumption drops 9% if the elasticity of marginal utility with respect to consumption is 5. The elasticity of marginal utility with respect to assets, however, is $9\% / 20\%$, which is only about 0.45. Essentially, the utility function is relatively flat between \$55,000 and \$56,000, so that the extra \$1,000 in assets does not reduce marginal utility noticeably even though assets are increased by a significant amount.

Ah, but you say, the asset level in this example is unrealistically low. What would happen if the initial assets were \$100,000, rather than \$5,000. Then an additional \$1,000 would represent a 1% increase in assets. This would be associated with an increase in resources from \$150,000 to \$151,000, a 0.67% increase. If the elasticity of marginal utility with respect to consumption is 5, this would reduce marginal utility by 3.33%. The elasticity of marginal utility with respect to wealth would then be $3.33\% / 1\% = 3.33$.

For the individual simulated in the table, the wealth-income ratio reaches 2 at age 47. At this age, however, the elasticity of marginal utility with respect to wealth is 0.54, not 3.33. The solution to this puzzle appears to be that it is not the ratio of wealth to current income which is critical, but the ratio of wealth to current plus prospective non-wealth income. Wealth measures the resources left over from the past, and current and prospective income measure resources available in the future. In the simulation, the elasticity of marginal utility with respect to wealth rises sharply in the three or four years before retirement, which is just the time when the sum of current and prospective income is dwindling most rapidly in percentage terms.

This analysis carries several implications regarding how the elasticity of marginal utility with respect to wealth, and hence the percentage of stocks in the portfolio, will behave in different circumstances. First, consider the effect of a somewhat higher time preference rate. One of the consequences of the higher time preference rate will be the consumption stream will be a bit more tilted toward the early years, with the result that there will be less wealth accumulation. Less wealth means that a given change in wealth will have a larger effect on the percentage change in wealth than on the percentage change in available resources, which in turn will mean that the elasticity of marginal utility with respect to wealth will be lower, with an associated larger percentage of assets invested in stocks. In other words, a higher time preference rate will exacerbate the problem of the model overpredicting the percentage of the portfolio in stocks. To illustrate, the following table does a second simulation on the same individual as before, except that this time the time preference rate is 3 percent per year:

Age	Year	MU Elasticity wrt Wealth	Pct Stocks	Rate of Return	Wealth plus Returns	Income	Consumption	Wealth at End of Year
50	1986	0.46	100.0	7.2	192060	128243	104207	216096
51	1987	0.55	100.0	7.2	231611	129526	103491	257646
52	1988	0.65	100.0	7.2	276145	130821	102797	304169
53	1989	0.76	100.0	7.2	326008	132129	102129	356008
54	1990	0.89	100.0	7.2	381570	133450	101506	413514
55	1991	1.04	100.0	7.2	443204	134785	100960	477029
56	1992	1.20	100.0	7.2	511280	136133	100508	546905
57	1993	1.39	100.0	7.2	586173	137494	100190	623477
58	1994	1.61	89.7	6.7	665106	138869	99912	704062
59	1995	1.84	78.3	6.1	747157	140258	99629	787786
60	1996	2.09	69.4	5.7	832594	141660	99389	874865
61	1997	2.37	61.6	5.3	921293	143077	99165	965205
62	1998	2.67	55.1	5.0	1013344	144508	99003	1058849
63	1999	2.99	49.4	4.7	1108717	145953	98803	1155867
64	2000	3.33	44.6	4.5	1207578	25000	98600	1133978
65	2001	3.33	44.5	4.5	1184694	25000	98391	1111302
66	2002	3.34	44.5	4.5	1160974	25000	98166	1087808
67	2003	3.34	44.4	4.5	1136397	25000	97924	1063473
68	2004	3.35	44.3	4.5	1110943	25000	97666	1038278
69	2005	3.35	44.3	4.5	1084590	25000	97393	1012197
70	2006	3.36	44.2	4.5	1057310	25000	96976	985334

Comparing this to the previous results, it is evident that the elasticities of marginal utility with respect to wealth are uniformly lower, the percentages of stocks in the portfolio are uniformly higher, and the drop in the percentage of stocks around retirement is even sharper. For higher rates of time preference, of course, the effects are even greater.

Secondly, the individual in the initial example is a relatively high earner. Among other things, this means that the ratio of Social Security benefits to peak earnings is relatively low, only about 17% of peak earnings. A lower income individual would have a higher ratio of benefits to earnings. This would have two effects. First, wealth would probably be lower, even as a percentage of earnings, since a greater percentage of future consumption could be financed with benefits. Secondly, in looking at the ratio of wealth to current and future income, the denominator would be higher, leading to a further reduction in the elasticity of marginal utility with respect to wealth. Both of these effects would mean that at lower income levels, with higher Social Security replacement rates, the problem of too high a percentage of the portfolio in stocks will be exacerbated.

To simulate this, we consider an individual with an initial earnings level of half (\$50,000) of that in the first example, and Social Security benefits of \$20,000 while both spouses are alive and \$12,000 after one of them dies. This simulation returns to a zero percent time preference.

Age	Year	MU Elasticity wrt Wealth	Pct Stocks	Rate of Return	Wealth plus Returns	Income	Con- sumption	Wealth at End of Year
50	1986	0.25	100.0	7.2	50524	64122	54881	59764
51	1987	0.32	100.0	7.2	64055	64763	54505	74312
52	1988	0.39	100.0	7.2	79648	65410	54129	90930
53	1989	0.48	100.0	7.2	97459	66065	53754	109769
54	1990	0.57	100.0	7.2	117651	66725	53386	130990
55	1991	0.68	100.0	7.2	140395	67392	53031	154757
56	1992	0.80	100.0	7.2	165868	68066	52692	181242
57	1993	0.94	100.0	7.2	194256	68747	52383	210620
58	1994	1.10	100.0	7.2	225742	69435	52121	243055
59	1995	1.28	100.0	7.2	260507	70129	51920	278715

60	1996	1.50	96.3	7.0	298220	70830	51780	317270
61	1997	1.73	83.4	6.4	337480	71538	51652	357367
62	1998	1.99	73.0	5.9	378310	72254	51523	399042
63	1999	2.27	64.4	5.4	420756	72976	51404	442328
64	2000	2.57	57.2	5.1	464837	20000	51244	433594
65	2001	2.56	57.2	5.1	455676	20000	51094	424582
66	2002	2.57	57.2	5.1	446194	20000	50938	415256
67	2003	2.57	57.1	5.1	436382	20000	50774	405607
68	2004	2.57	57.1	5.1	426230	20000	50604	395626
69	2005	2.57	57.0	5.1	415728	20000	50417	385311
70	2006	2.57	57.0	5.1	404894	20000	50164	374730

Again, the percentages of stocks drops from 100% to around 57% in a very short period before retirement.

Third, pensions will have a similar effect. Pensions also raise the replacement rate of post-retirement income, reduce the need to accumulate wealth, lower the elasticity of marginal utility with respect to wealth, and as a result increase the percentage of the portfolio in stocks. The result is pretty much the same as having a higher Social Security replacement rate. This situation can be simulated by using the individual in the first example but adding a \$25,000 pension while the individual is alive:

Age	Year	MU Elasticity wrt Wealth	Pct Stocks	Rate of Return	Wealth plus Returns	Income	Consumption	Wealth at End of Year
50	1986	0.27	100.0	7.2	107403	128243	109268	126378
51	1987	0.34	100.0	7.2	135452	129526	108518	156460
52	1988	0.42	100.0	7.2	167694	130821	107766	190748
53	1989	0.51	100.0	7.2	204444	132129	107021	229552
54	1990	0.61	100.0	7.2	246034	133450	106291	273193
55	1991	0.72	100.0	7.2	292808	134785	105593	322000
56	1992	0.86	100.0	7.2	345120	136133	104948	376305
57	1993	1.01	100.0	7.2	403323	137494	104364	436453
58	1994	1.18	100.0	7.2	467790	138869	103887	502772
59	1995	1.38	100.0	7.2	538871	140258	103575	575553
60	1996	1.61	89.6	6.7	613969	141660	103298	652331
61	1997	1.85	77.9	6.1	692139	143077	103053	732162
62	1998	2.13	68.4	5.6	773427	144508	102832	815102
63	1999	2.42	60.4	5.2	857864	145953	103063	900754
64	2000	2.71	54.2	4.9	945292	50000	103104	892188
65	2001	2.73	53.9	4.9	936176	50000	103185	882991
66	2002	2.75	53.5	4.9	926371	50000	103282	873089
67	2003	2.76	53.2	4.9	915832	50000	103395	862437
68	2004	2.78	52.9	4.9	904513	50000	103522	850991

69	2005	2.80	52.5	4.9	892366	50000	103660	838706
70	2006	2.82	52.2	4.8	879341	50000	103635	825706

Again, the simulation illustrates that considering a pension will induce the percentage of stocks in the portfolio to increase relative to the initial simulation, and that percentage drops from 100% to about half in a shorter period of time just before retirement.

How do these results fit in with the literature? Poterba, Rauh, Venti and Wise (2006) look at nine investment strategies for 401k assets with crra parameters of 0,1, 2, and 4. They consider the accumulation phase only, and measure the outcome when the period of contributions ends. When other wealth is considered, they find that an all-stock portfolio dominates any of the other strategies regardless of the crra parameter. This is close to our result indicating that stocks are dominant until just before retirement. A weakness to their approach is that the amount of non-401k financial assets at the end of the period is fixed, and by implication the asset allocation of the non-401k part of the portfolio is also fixed.

The Gomes and Michaelides (2005) model aims to match two empirical regularities. First, about 50% of individuals participate in the stock market, and this fraction does not vary a great deal by age. Secondly, among those who do participate, about 50-60% of their holdings are in stocks, and this fraction also does not vary much by age. They match the first finding by having a heterogeneous population, half of whom do not save much at all and hence are not in the stock market. This is very similar to our individuals with high rates of time preference. They match the second regularity at ages above 45, but below that the percentage of the portfolio in equities is too high, much like our model. And their results are undoubtedly helped along by the assumption of an equity premium of 4%. The results we derived earlier suggest an equity premium of more like 7% vs. intermediate term bonds and over 8% vs. money market like assets. With the higher equity premium, their model would show 100% equity allocation longer (it is

until around age 35 in their graphs) and a higher allocation after that. They also introduce a one-time fixed cost to be in the stock market for the first time, but the consequences of this assumption are very small: it deters individuals from entering the market for a year or two until they accumulate sufficient assets to make it worthwhile. They additionally have a stochastic wage process in the model, and still they come up with a 100% equity allocation until at least the mid-thirties, at which time savings are twice earnings and are clearly life-cycle, not precautionary, motivated. In short, their model ultimately suffers from much the same problem as does ours, although they appear to have the right explanation for the fraction of individuals not in the market.

Benzoni, Collin-Dufresne and Goldstein (2005) take a somewhat different tack. They postulate a relatively high long-run correlation between earnings and stock market returns, on the order of 0.5. With stocks and earnings so correlated, individuals initially would like to short stocks, but are kept from doing so by borrowing constraints. Their model has no heterogeneity, so the results apply to a typical individual. Despite accumulating savings from the start, the individual remains exclusively in the safe asset until around age 30. Then the fraction of the portfolio in stocks grows gradually until reaching a peak of around 50% at age 60. This model suffers from the opposite problem relative to the previous model in that the fraction in stocks is too low at younger ages. They see the fact that individuals don't hold stocks until around 30 as consistent with younger folks not being big stockholders, but their model doesn't allow for the large number of individuals, especially younger individuals, with almost no assets. The heterogeneity explanation is probably a better explanation for the 50% of individuals who don't participate in the market and who don't have much assets. Heterogeneous individuals could

undoubtedly be introduced into the model, but then the overall percentage who hold stocks and the fraction of assets held as stocks would probably be too low, particularly at younger ages.

V. Implications and Conclusions

To summarize the findings:

1. Estimates of our model based only on retirement and saving behavior yield a coefficient of relative risk aversion (γ) of approximately 1. Examination of historical data suggest an equity premium of about 7% vs. intermediate term bonds and over 8% vs. money market like assets. Consistent with earlier findings, individuals would want to be 100% in stocks with a relative risk aversion coefficient of 1 and an equity premium of 7%.
2. In order for individuals to be induced to hold financial assets other than stocks, it seems that the coefficient of relative risk aversion parameter has to be approaching a value of 5. Having a value of γ equal to 5, however, creates other difficulties for the model. In our current model, this value would mean that additional consumption would have sharply lower marginal utility, so that individuals with higher incomes would want to retire much earlier. In fact, with a sufficiently high value of γ , individuals would be approximately target income workers. That is, they would work long enough to achieve a target level of income, and any income beyond that would produce so little utility that they would rather take the leisure. Hence, a high value of γ is inconsistent in our current model with the fact that individuals retire at more or less the same time on average, regardless of their level of income.
3. These results suggest that if we are going to examine both retirement behavior and asset allocation in the same model, we need to modify our original model somewhat. The most straightforward way to do this is to introduce an additional parameter to separate risk aversion from the marginal utility of consumption. Accordingly, we introduce a value function for

consumption. The simple value function for consumption we introduce has constant returns to scale, that is, if consumption and expected consumption double at every point in time, the value function also doubles. After adding a value function for leisure, the new utility function has separate coefficients for risk aversion and for time preference. The additional parameter can be addressed by introducing another moment into the estimation process, namely the percentage of assets held in stocks.

4. We next undertake an exercise in comparative statics. Our analysis explains why a high elasticity of marginal utility with respect to consumption is entirely consistent with a relatively low elasticity of marginal utility with respect to wealth.

5. When we run the model (with fixed retirement) using a value of γ (risk aversion for consumption) toward the high end of the values used in the literature, the model overpredicts the percentage of the portfolio in stocks. In addition, stock holdings decline suddenly as the person approaches retirement age. Most investment advice, such as the 110 – age formula, yield a gradual reduction of stocks in the portfolio. In contrast, the model simulation gives a 60% reduction in stocks in the span of around 8 years. This will make it very difficult for the model to generate portfolios consisting of fifty percent stocks as are found in the data. Moreover, these results are for a rate of time preference of zero; a higher time preference rate exacerbates the problem of overpredicting stock holdings.

6. Social Security raises the replacement rate of post-retirement income, reduces the need to accumulate wealth, lowers the elasticity of marginal utility with respect to wealth, and as a result, for those with low incomes and higher Social Security replacement rates, the problem of too high a percentage of the portfolio in stocks will be exacerbated. Pensions will have a similar effect.

7. We then compare our findings with those in the literature. The bottom line is that current financial models in the literature also have difficulty explaining why most portfolios devote half their assets to stocks, predicting larger shares. Either people are failing to make rational decisions regarding portfolio choice, or there is some aspect of portfolio choice that the models are failing to consider. What that aspect might be, however, is at present not clear. It is possible to modify our model to include portfolio choice while maintaining plausible retirement choices, but the modified model suffers from many of the same problems with regard to overpredicting stock shares.

These findings raise the question: where do we go from here? Potential modifications to dynamic-stochastic models of retirement that would allow portfolio choice to be included also have obvious shortcomings. If we try to estimate a revised model which includes a separate relative risk aversion coefficient γ as an estimated parameter and use the percentage of stocks in the portfolio to pin the value down, we will probably either have an unacceptably high value of q because the model overestimates the fraction of stocks in wealth, or an unacceptably high value of γ , or, more probably, both. Alternatively, assuming a value of γ , we could use the rest of the parameters as previously estimated, and adjust the constant term in β to yield the correct average retirement dates. Then we could explore the behavior of the model under a variety of circumstances. The main advantage of the second approach is that it could illustrate how a variable retirement date can be used as insurance for bad draws of rates of return. However, the insurance will probably mean that the elasticity of marginal utility with respect to wealth will be even lower, which will worsen the problem that the model overestimates the percentage of stocks in the portfolio.

At the moment, no easy solution is at hand to this fundamental problem now being engaged by financial economists. This suggests that models of retirement and saving may, for the immediate future, be forced to constrain portfolio composition to correspond with levels observed in the data, postponing the inclusion of portfolio mix as a choice variable until further progress is made in modeling that behavior. This does not necessarily reduce the efficiency of life cycle models of retirement and saving. Rather it recognizes that portfolio choice may be influenced by behavior that is not fully consistent with that posed by a life cycle model. Individuals may, for example, be accepting recommendations from planners or firms that they would not otherwise follow if they fully understood how to balance risk and return in portfolio choice in the same way they balance risk and return in their saving and retirement decisions. If these behavioral considerations govern their portfolio choice, while retirement and saving are determined by life cycle considerations, a model that correctly constrains portfolio composition may in fact generate parameter estimates that accurately reflect the forces governing retirement and saving behavior.

References

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein. 2005. "Portfolio Choice over the Life-Cycle in the Presence of 'Trickle Down' Labor Income." NBER Working Paper 11247.

Gomes, Francisco J. and Alexander Michaelides. 2005. "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence". Center for Economic Policy Research Discussion Paper No. 4853.

Gustman, Alan L. and Thomas L. Steinmeier. 2002. "Retirement and the Stock Market Bubble." NBER Working Paper 9404.

_____. 2005a. "The Social Security Early Retirement Age In A Structural Model of Retirement and Wealth". *Journal of Public Economics*. Volume 89, Issues 2-3, February 2005, Pages 441-463.

_____. 2005b. "Retirement Effects of Proposals by the President's Commission to Strengthen Social Security". *National Tax Journal*. 58 No. 1 March, 2005: 27-49.

_____. 2005c. "Retirement, Saving, Benefit Claiming and Solvency Under a Partial System of Voluntary Personal Accounts". Final Report, Michigan Retirement Research Center, October.

_____. 2006. "Social Security and Retirement Dynamics". Social Security Administration, Michigan Retirement Research Center. UM05-05.

Poterba, James, Joshua Rauh, Steven Venti and David Wise. 2006. "Lifecycle Asset Allocation Strategies and the Distribution of 401(k) Retirement Wealth". NBER Working Paper 11974.