

NBER WORKING PAPER SERIES

RETIREMENT AND THE
STOCK MARKET BUBBLE

Alan L. Gustman
Thomas L. Steinmeier

Working Paper 9404
<http://www.nber.org/papers/w9404>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2002

The research reported herein was performed pursuant to a grant from the National Institute on Aging (1R01AG13913-01A1). The opinions and conclusions are solely those of the authors and should not be construed as presenting the opinions or policy of NIA or the National Bureau of Economic Research.

© 2002 by Alan L. Gustman and Thomas L. Steinmeier. All rights reserved. Short sections of text not to exceed two paragraphs, may be quoted without explicit permission provided that full credit including, © notice, is given to the source.

Retirement and the Stock Market Bubble
Alan L. Gustman and Thomas L. Steinmeier
NBER Working Paper No. 9404
December 2002
JEL No. J26, J14, J32, E21, D31, D91, I3, C61, H55

ABSTRACT

This paper specifies and estimates a structural dynamic stochastic model of the way individuals make retirement and saving choices in an uncertain world, and applies that model to analyze the effects of the stock market bubble on retirement behavior. The model includes individual variation both in retirement preferences and in time preferences. Estimates are based on information covering the period 1992 through 2000 from the Health and Retirement Study (HRS), a panel survey of retirement age respondents and their spouses.

The extraordinary returns in the stock market in the late 1990's, which more than doubled stock prices and unexpectedly increased the value of a mixed portfolio by nearly 60 percent, increased retirement for the HRS sample of workers by over 3 percentage points by the turn of the century and would have decreased the average retirement age by about a quarter of a year if it had not been interrupted. The subsequent decline in the market, which very nearly wiped out the gains that had been made during the preceding surge, effectively neutralized the effect of the preceding stock market gains on retirement. The effects of the bubble were to increase retirement as long as the bubble continued, but any continuing effects of the bubble after its end will probably be minimal.

Alan L. Gustman
Department of Economics
Dartmouth College
Hanover, N.H. 03755
and NBER
alan.l.gustman@dartmouth.edu

Thomas L. Steinmeier
Department of Economics
Texas Tech University
Lubbock, TX 79409
Thomas.Steinmeier@TTU.edu

I. Introduction

The stock market has had a truly roller-coaster ride over the last decade or so. During the decade of the 1980's, the New York Stock Exchange index rose an average of 12 percent per year, falling to 5 percent per year in the first half of the 1990's. In the five year period starting in 1995, the market rose by almost 21 percent per year compounded. Over the course of this five year period, a stock portfolio would have been worth about two and a half times as much at the end of the five years as at the beginning. And then the bubble burst. During the period from the beginning of 2000 to the end of September 2002, the market has fallen by 13 percent per year. From the peak in August 2002, the cumulative decline has been 34 percent, so that a little more than a third of stock market wealth has evaporated in two years.

These events have had enormous implications for the segment of the population on the cusp of retirement. In accordance with usual life cycle behavior, many workers had amassed substantial savings to finance their retirement years. Much of this saving was invested in the stock market, either directly or through 401k or other defined contribution pension plans. In the late 1990's these workers watched with interest as their portfolios grew almost beyond their wildest expectations. Since the annual returns of the stock market exhibit an almost zero autocorrelation, it seems unlikely that the returns of the late 1990's were anticipated in any major way. All of a sudden, these workers had the wherewithal to do things that were previously beyond their grasp, and in particular they had the means to retire somewhat earlier than they might have anticipated, without sacrificing expected post-retirement consumption.

The bust in the last couple of years has just the opposite effect. Retirement-age individuals who were still working watched their nest eggs shrink, and many have undoubtedly

concluded that it would be prudent to continue the income from working until either the market recovers or they have had an opportunity to rebuild the nest egg to a level sufficient to ensure a comfortable retirement. Individuals who had recently retired faced an even more unappetizing menu. They could reduce consumption to match their reduced resources or risk running out of money well before their expected lifetimes. Alternatively, they could go back to work, often at wages which were substantially lower than the wages they were used to bringing home in their last years of working.

In this paper we investigate the effects of the recent stock market bubble on retirement behavior. To do so, we construct and estimate a structural econometric model that explains retirement and saving in an uncertain world. The model allows for heterogeneous time preferences, with some individuals willing to save and others being more impatient. The heterogeneous time preferences enable the model to accomplish the difficult task of explaining the joint peaks in retirement activity at ages 62 and 65. The model includes a stochastic return on asset balances, including balances in defined contribution pension plans. A stochastic return on assets is essential for any kind of consideration of the stock market bubble. The model also allows for other realistic features of the life cycle model, including heterogeneous retirement preferences and the possibility of returning to work (albeit at a reduced wage) if assets decline precipitously in value or if individuals find that retirement is not as appetizing as they expect. All of these features are indispensable elements to analyze the effects of the stock market bubble on savings and retirement. In the end, we use this model to investigate the effects on retirement of both the sharp rise in the stock market in the late 1990's and the subsequent collapse in the last couple of years.

In the next section, we discuss some of the previous work both on the issue of wealth influences on retirement and on the nature of stochastic structural modeling. The following section lays out the structural dynamic stochastic model we use to characterize the retirement and asset accumulation decisions. Section IV talks about the data used in the analysis, and the next section discusses the procedure by which the unknown parameters of the model are estimated. In Section VI we present the results of the estimation and examine the characteristics of the base simulation, including the ability of the simulation to capture many of the basic patterns observed in retirement. The following section turns to simulations of alternative scenarios, including scenarios where the pattern of stock returns were more typical. A comparison of these alternatives to the base simulation indicates the degree to which the extraordinary behavior of the market in recent years has altered retirement. A final section summarizes conclusions to be reached from the study.

II. Previous Work

Analyses of retirement and of saving behavior both rely on a life cycle framework. Diamond and Hausman (1984) pointed to the joint determination of wealth and retirement. Yet until recently there were no structural econometric studies that estimated retirement and saving as jointly determined outcomes. Studies of saving typically oversimplified the labor market behavior that was assumed to be occurring at the time of retirement, while studies of retirement assumed a standard life cycle pattern of asset accumulation with oversimplified distributions of retirement preferences.¹

¹See Gustman and Juster (1996) for a joint survey of these literatures.

Stock and Wise (1987) and Berkovec and Stern (1991) introduced a stochastic environment. Rust's seminal work (Rust, 1989, 1990) provided an econometric framework for a joint analysis of retirement and saving in a stochastic framework. Rust and Phelan (1993) estimated structural models that jointly considered retirement and saving under the assumption of inoperative capital and insurance markets. French (2002) estimated a model with joint saving and retirement using stylized information on pension incentives.

Samwick (1998) in studying the distribution of wealth has pointed to the importance of a heterogeneous distribution of time preference that has an important mass at both high and low levels of time preference. To explain the spikes in retirement observed in the data at ages 62 and 65, as well as the very wide distributions of wealth at each level of lifetime earnings, Gustman and Steinmeier (2002a) have specified and estimated a structural model of retirement and wealth that allows for these wide differences in time preference.

But to explain the likely responses of retirement and wealth accumulation to the stock market bubble, it is also necessary to incorporate a stochastic environment. Thus in the present paper, we specify and estimate a structural dynamic stochastic model of retirement and wealth that also allows for the wide heterogeneity in saving behavior that is so apparent in the data.

III. A Structural Model of Retirement with Stochastic Asset Returns

The model takes a standard life cycle approach in which individuals maximize expected utility subject to an asset evolution constraint. For each time period t_i , the individual chooses consumption C and leisure L to maximize expected utility:

$$EU_{t_i} = E_t \left[\sum_{t=t_i}^T \left\{ e^{-\rho t} \sum_{m=1}^3 s_{m,t} \left(\frac{1}{\alpha} C_{m,t}^\alpha + h_t L_{m,t} \right) \right\} \right]$$

L takes on a value of 1 if the individual is retired and a value of 0 if he or she is working, and as such it is essentially a retirement indicator. Since L is a binomial variable, it can be entered linearly with no loss of generality.² h_t indicates the strength of the individual's preference for retirement, which may vary from one person to the next. ρ is the time preference rate, which also may vary from one person to the next.

In this paper, we consider individuals who are part of couples. Given the complexity of the calculations, we take the income of the spouse to be exogenous. The individual does take into account the consumption possibilities of the spouse after he or she is deceased. The index m takes on three values indicating whether both members of the couple survive until time t , only the individual survives, or only the spouse survives. $s_{m,t}$ is the probability that the household will have the composition described by m at time t . T corresponds to the maximum age beyond which the household's survival probabilities are too small to matter.

The asset constraint is given by

$$A_t = (1 + r_t) A_{t-1} + W_t (1 - L_{m,t}) + E_{m,t} + B_{m,t} - C_{m,t}$$

A_t is the level of assets at time t , and r_t is the stochastic return on those assets at time t .

²It would be desirable eventually to allow L to take additional values, which would allow work effort to be separated into full-time work, part-time work, and retirement. But introducing additional work effort states complicates the computational problem, and in this paper we consider only two states.

Assets are assumed to start out at 0 at the beginning of the working life and are not permitted to be negative. W_t is the wage rate at time t , which will depend on whether the individual has stayed on his or her primary job or has previously retired and is going back to work. The term $E_{m,t}$ is the income accruing to the spouse, including earnings and pensions. The spouse is assumed to have a retirement date unaffected by the individual's choices, and the term is taken to be zero in states where the spouse is no longer alive.

$B_{m,t}$ is the amount of the individual's pension and the household's social security benefits, both of which will be affected by the individual's retirement decisions. For defined benefit pensions, the benefit amount is determined by the retirement date and continues until death. For defined contribution pensions, the contributions are put into an account and allowed to accrue subject to the same stochastic return as is applied to assets. The account is assumed to be made available to the individual when the individual retires from the primary job. Household social security benefits are calculated according to the social security rules, depending on previous retirement decisions and the composition of the surviving household. Since most individuals claim benefits as soon as eligible (Coile et al., 2002; Gustman and Steinmeier, 2002b), we do not try to model the acceptance decision here and instead assume that the individuals will claim the benefits as soon as they can.

Individuals are assumed to be heterogeneous with regard to both their time preference rate and their retirement preferences. With regard to time preference, we assume that ρ has a different value for different individuals and essentially treat it as a fixed effect in the estimation. Retirement preferences are reflected in the coefficient to the leisure term in the utility function and are characterized by

$$h_t = e^{\beta X_t + \epsilon_t}$$

The linear form βX_t has three terms: a constant, age, and health status. The coefficient of age is taken to be positive, so that retirement gradually becomes more desirable as the individual ages and suffers from the cumulative effects of nature's aging. Note that there is no term which makes the individual suddenly more desirous of retirement at 62, 65, or any other particular age; this means that any spikes in retirement in the simulations are the result of idiosyncracies in the opportunity set, and not the result of preferences. Given the computational costs of modeling health as endogenous, health status is taken as exogenous in this study. This is probably acceptable in a study of the effects of the stock market on retirement, though if the model were to be used for other purposes, such as studying disability, it would be more important to make health endogenous.

The ϵ_t term in h_t reflects the individual's relative preference for leisure. An individual starts out with a value of ϵ drawn from a distribution with mean 0 and standard deviation σ_ϵ and keeps this value until he or she retires from the primary job. Upon retirement, the individual may find that retirement is more or less fulfilling than anticipated, or perhaps the individual may find that they value consumption relatively more than they had thought. In any case, experience allows the value of ϵ to change after retirement, and the model reflects this by allowing the value of ϵ_t to vary after retirement, with values in successive years correlated with a correlation parameter ρ_ϵ . If the individual finds that retirement is substantially less fulfilling than anticipated, a return to work, albeit at a reduced wage as compared to the primary job, may be the optimal decision.

The individual carries several state variables from one period to the next; these are variables which are consequences of past decisions and random events which have a bearing on the current decision. Three state variables are applicable in all periods. These are the level of assets A_t , the level of preferences ϵ_t , and whether or not the individual is still in the primary job. If the individual is still in the primary job and that job had a defined contribution pension, there is another state variable relating to the size of the defined contribution balance. After the individual has left the primary job, additional state variables related to the value of pension benefits and social security benefits are introduced.

Solution to the model.

As with any stochastic dynamic model, the solution of the model proceeds backwards. In the final year T , for each point in a grid of points in the primary job-retirement preference-asset-pension-social security space, the individual calculates the utility value of consumption if he or she were to work and the utility value of consumption plus leisure if the individual were to retire. The primary job variable refers to whether the individual has the option of working in the primary job this period; if not, the decision to work refers to a post-retirement job. In the terminal period, where there is no value to having assets left over, the individual will simply consume all available assets and income. The individual chooses to work or not depending on which choice yields the highest total utility, and the marginal utility of assets is the marginal utility associated with that consumption choice.

The solution moves on to the next to the last period. The individual must again decide whether to work or not and how much to consume, except that this time the consumption decision is not necessarily to consume all available resources. Rather, for each potential

consumption decision, there will be a certain amount of assets left over. Using the marginal utility values for the final period, and allowing for the probabilities of rates of returns on these assets, the probabilities of changes in the retirement preferences, and the time preference rate, the expected marginal utility of these assets can be calculated. The individual would choose that level of consumption which just matches the marginal utility of a dollar of current consumption with the expected marginal utility of leaving the dollar in assets to be consumed in the future. If necessary, future asset values between the grid points are interpolated. The individual does this calculation both for the case of working and not working for each cell in the primary job-retirement preference-asset-pension-social security space, and chooses whichever case yields the highest level of utility, including expected future utility from the assets to be carried forward.

This procedure is carried on to period $t-2$, $t-3$, and so on to period 1, which is presumed to correspond to the age just after the individual finishes school. In effect, the model calculates out a complete contingent decision tree for work and asset accumulation. To simulate the individual, we start with no assets and a random draw of retirement preferences ϵ . The model tells what the work and consumption choices will be, which in turn tells what the assets to be carried into the next period will be. The realized rate of return on those assets determines the level of assets at the beginning of the next period. If the individual retires, a random draw will determine the change in retirement preferences. If the individual chooses to work in the initial period and has a defined contribution pension, the realized rate of return will determine the increase in the value of the balance in the pension. For the new level of assets and retirement preferences, the individual decides again whether to retire or not and how much to consume. If the individual worked in the primary job in the previous period and retires, the amount of

potential pension and social security benefits (which may not be collectable yet) are determined. The calculations for the remaining periods are done similarly.

In the estimations and simulations, the model uses a grid with 40 points for assets, 10 for defined contribution balances, 17 for retirement preferences, and 4 each for pensions and social security values. The number of calculations each period is related to the product of the number of points for each of these variables, multiplied again by the two categories for the primary job variable.

IV. Data and the Sample

The primary data used in this analysis is the Health and Retirement Study (HRS), a panel survey of retirement age respondents and their spouses. The survey started in 1992 with about 7,600 households with at least one individual who was 51 to 61 years old at the time. The households have been surveyed at two year intervals since then; we use information through the 2000 survey. The survey attempts to gather comparable information about both spouses in a couple, and it has matched social security and pension records for many respondents. For the estimation and simulation of the model, we extract information on the retirement behavior of the respondents, on the amount of wealth they have accumulated, and on their lifetime resources and incomes.

In the model, respondents are considered to be working if they are essentially working full-time. The criteria for work is more than 30 hours of work per week or more than 1560 hours per year. The 1560 is 40 hours per week times 39 weeks per year and thus should be sufficient to cover occupations like teachers who would be classified as working even if they are interviewed

over the summer. In addition, if an individual is working at least 25 hours a week and 1,250 hours per year, and self-reports themselves to be not retired at all, we classify them as working. We recognize that this classification omits part-time work and partial retirement, but these are omitted at this stage in order to keep an already complicated model tractable. In order to reduce classification error, all work measurements used in the estimation are taken as of the survey dates. We do not try to interpolate the actual work effort to dates between the surveys or before the first survey.

Wealth is measured in the 1992 survey and is used only to get an estimate of the respondent's time preference, which as indicated before is regarded as a fixed effect in the model. Roughly speaking, time preference is taken to be the value for which the observed wealth is consistent with the wealth that would be generated by the model. In this regard, wealth is considered to be wealth that could be used to finance retirement, and includes financial wealth, real estate, and business assets. Given the likelihood of measurement error in the wealth of individual respondents, we do not try to use information about the measured increases or decreases in wealth between surveys in the course of estimating the model.

Lifetime resources include a number of components. For about three-quarters of the sample, earnings are reflected in attached social security records, adjusted by self reported information for cases over the earnings maximum or for uncovered jobs. The remaining individuals have earnings histories constructed from self reports about current and last jobs, other jobs lasting longer than 5 years, and information about the total number of years worked. Projecting the potential earnings for returning to work after an individual retires is done by using

the experience and tenure coefficients from estimated wage equations.³ Pension information is available for about two-thirds of individuals with pensions. For defined benefit pensions, the annuity is calculated based on the summary plan description provided by the firm. For defined contribution pensions, the contribution amounts are calculated year by year, and the balances are allowed to grow within the model at the stochastic rate of return. Social security benefit amounts for various potential retirement dates are calculated according to the social security rules, assuming the respondents collect benefits when eligible to do so.

The earnings and pension income of the spouse are calculated and, for the purposes of this model, are taken as exogenous in order to keep the model computationally tractable. Also taken as exogenous are inheritances and other types of income. Income and capital gains related to assets, of course, are treated as part of the stochastic rate of return to assets.

The prospective distribution of rates of return is based on the historical series calculated by Ibbotson Associates (2002) for various asset classes from 1926 through 2001. To combine these figures into a single rate of return, we look at households in the HRS who had at least some stocks and/or bonds. We rank these households by total financial assets and take the middle 10 percent of the households to assess the distribution of assets among stocks, long-term bonds, and short-term financial instruments. We find that, on average, about 50 percent of these assets are in stocks, 5 percent in bonds, and 45 percent in bank accounts and certificates of deposit. We proxy the returns on these assets by using the Ibbotson returns for large company stocks, long-term government bonds, and treasury bills, respectively. The arithmetic mean of the returns for this

³ The wage equation is reported in Appendix Table 1 in Gustman and Steinmeier (2002a). This equation suggests a wage drop of about 25 percent for an individual with 20 years of tenure who retires and then tries to return to work.

weighted average of securities is 5.3% with a standard deviation of 11%, and the geometric mean is 4.7%.

Table 1 details the derivation of the estimation sample from the pool of potential observations. We focus on married men in this analysis for two reasons. Since other studies have suggested that it is probably inappropriate to assume the same retirement parameters for men and women, we restrict the analysis to men. Secondly, since many of the unmarried respondents are in fact divorced or widowed, we omit these respondents because the HRS has little information of the earnings of the former spouse.

Table 1 notes that most of the deletions from the sample come from two sources. The most important is the group that had pension coverage in the most recent job but for whom pension information was not collected from the firms. Since pensions often provide strong incentives to retire at specific dates, the lack of pension information means that we are missing an important influence on retirement. Consequently, we feel it best to omit these observations. The second major omitted group consists of respondents with spotty employment records. This group includes those who have no full-time work after age 50 and also those who work full-time fewer than half the years between age 40 and the last observed year of full-time work. For such respondents, the very concept of retirement is ambiguous, since for much of the time they are not working anyway. There is a third, much less common, group omitted from the sample. These are respondents whose business assets comprise the majority of their wealth. These individuals probably hold these assets for reasons related to the necessity of carrying on a business and not for reasons of acquiring potential retirement assets, and it seems questionable whether it is appropriate to try to infer time preference from the magnitude of these assets.

The final sample consists of 2,305 respondents for whom we can construct, at least approximately, the details of their earnings and income opportunities, and for whom the model seems appropriate. This is slightly less than half of the number of married men available in the original HRS sample.

V. Estimation of the Model

The parameters to be estimated for the model include the consumption parameter α , the standard deviation σ_ϵ for the retirement preference variable ϵ , the correlation ρ_ϵ of the values of ϵ once the individual leaves the main job, and the coefficients in the linear term βX_t which affects retirement preferences. These coefficients include β_o , the constants, β_a , the coefficient of age, and β_h , the coefficient of health.

Each of these parameters have specific implications for the observed pattern of retirement. The estimation procedure in effect uses these implications to help to identify the values of the parameters. β_o determines the average age of retirement, and σ_ϵ governs the range of retirement ages. β_a is perhaps the most important parameter for our purposes, since it determines the sensitivity of retirement to economic incentives. If β_a is high, the preference for retirement is increasing very rapidly with age, and economic incentives have little room to affect the retirement age. The opposite is true if β_a is relatively low. ρ_ϵ affects the probability of returning to work after retirement; if ρ_ϵ is low, then individuals may experience large shifts in retirement preferences after retiring, and a return to work is more likely. α affects how retirement varies with general lifetime income. If α is near unity, then the marginal utility of income diminishes little at high income levels, and high income workers should retire later in order

to reap the benefits of their higher wages. The opposite is true if α is non-trivially negative. Finally, β_h reflects the degree to which bad health encourages retirement, over and above its impact through earnings.

To estimate these parameters, we use the Generalized Method of Moments (GMM) method. We choose this method rather than the more common maximum likelihood estimator because it is less sensitive to a relatively small number of cases where the observed result is not very likely given the model. A notable example is someone who retires after 19 years with an employer, when by staying 20 years the individual would be eligible for a substantially higher pension. This could occur because the years of service with the employer are not measured accurately, or because the individual faced some event which caused retirement despite the pension incentives, or for a whole host of other reasons. In the context of the model, this would be a very low probability event, and maximum likelihood would give great weight to increasing the probability if possible. One way to raise this probability is to choose estimates which make individuals insensitive to economic incentives. Such estimates, however, have the difficulty that they cannot reproduce, even approximately, the spikes in retirement observed at ages 62 and 65.

The GMM estimator begins with a series of moments, which are nothing more than the difference between some observed characteristic of the sample and the expected value of that characteristic implied by the model. One example is the percentage of the sample that is retired during one of the surveys. If the model is approximately valid, the observed value of the characteristic should be close to the expected value implied by the model, and the difference should be fairly small.

The chosen moments are grouped into a vector \mathbf{m} .⁴ For the present model, the moment vector includes retirement at all ages between ages 54 and 66. For ages 55, 58, 60, 62 and 65, the vector also includes retirement if the observation is in the upper third of lifetime income, retirement if the observation is in the lower third of lifetime income, and retirement if the individual is in poor health. These moments help to establish the values of α and β_h , as discussed in the previous paragraph. The final set of moments in the vector are measures of the frequency with which individuals return to work in one survey, given that they were retired in the previous survey. All in all, we use 32 moments in the vector \mathbf{m} .

Generalized method of moments seeks to minimize the quantity

$$q = \mathbf{m}' \mathbf{w}^{-1} \mathbf{m}$$

where

$$\mathbf{w} = \sum_{i=1}^n \mathbf{m}_i \mathbf{m}_i' .$$

If the model is correct, then each moment is a measure of an observed value around its expected value. In the limit, \mathbf{m} should have a multivariate normal distribution. In the expression for \mathbf{w} , \mathbf{m}_i is the vector of moments of a single observation, so that \mathbf{w} is essentially an estimate of the variance-covariance matrix of \mathbf{m} . q is thus approximately a quadratic form and, if the model is correct, it should have a χ^2 distribution with the degrees of freedom given by the number of moments less the number of parameters being estimated. The variance of the estimates is given by

⁴ The notation for the discussion of GMM is taken from Greene (2000) but is similar to the notation in many other econometric texts.

$$\text{Var}(\boldsymbol{\theta}) = [\mathbf{G} \mathbf{w}^{-1} \mathbf{G}]^{-1}$$

where $\boldsymbol{\theta}$ is the vector of parameters to be estimated and \mathbf{G} is the derivative of the moments with respect to the parameters.

We note a couple of caveats with regard to the estimation procedure. The first has to do with the time preference parameter, which we mentioned before would be treated as a fixed effect in the estimation. As a result, there is no single estimated parameter for time preference. Rather, for each individual, the value of the time preference parameter is such that, given the values of the other parameters, the resulting value of calculated wealth in 1992 is just equal to the actual value of wealth. In general, high values of wealth are associated with low time preference rate parameters and vice versa, although this relationship may be affected if the individual is expecting high pension or social security benefit streams. In some cases, assets are higher than would be expected even at a zero time preference rate; we take this as evidence that the time preference rate is low and assign it a value of zero. At the other extreme, individual with zero assets are assigned a high enough time preference rate to insure that such an individual consumes all income in every period.

To simplify the calculation of time preference and keep the calculations tractable, we ask what time preference would generate the savings corresponding to the respondent's observed level of assets in 1992, given the expected retirement date and its associated income flows. This approximation allows time preference to be calculated prior to the solution to the rest of the model, which greatly simplifies the computations. Even so, the time preference rate must be recalculated each time the estimation considers a new set of parameters.

The second caveat is in regard to the expected values used to calculate the moments. In a

model as complex as the present one, it is very difficult if not impossible to calculate closed form solutions for the expected probability of retirement at a given age or date. Rather, we calculate the expected probabilities based on simulations of the model. For any given set of parameters, we simulate the model 10,000 times for each observation by taking random draws for the unobserved stochastic retirement preference variables and using the observed values of the rate of return. The expected values of the probability of retirement at various ages and under various circumstances are based on the simulated distributions as an approximation to the theoretical distributions. More properly, this technique should be called the Generalized Method of Simulated Moments, since we use simulated moments rather than theoretical moments. Since the number of simulations is large, however, the simulated moments should be very good approximations to the theoretical moments, and the resulting estimates should be almost identical to those using theoretical moments.

VI. Estimation Results

The coefficients estimated by the simulated GMM procedure are reported in Table 2. The estimated coefficients are all significantly different from zero by conventional standards. Of particular importance, the coefficient of the age variable implies that retirement leisure is increasing in value by 5.4 percent per year. This relatively low value means that economic incentives should be able to have considerable influence on retirement. The autocorrelation coefficient for the leisure preference term is significantly less than unity, which means that individuals can experience changes in their perception of retirement after they begin to experience it. If they find retirement less attractive than anticipated, they may well reverse course and go

back to work for a while, at least until the inexorable march of age finally makes retirement appealing again. Poor health increases the value of retirement leisure by approximately the same amount as being over three years older.

The overall fit of the model is measured by the q-statistic. If the model is correct, the departures of the observed moments from their theoretical means should be random, and the q-statistic should be distributed as a χ^2 with degrees of freedom equal to the number of moments less the number of parameters. In the present context, there are 32 moments and 6 parameters, so the degrees of freedom are 26. The statistic of 20.56 is well within the bounds of a χ^2 with 26 degrees of freedom, which has a 5 percent critical value of 38.89. This means that, at least among the moments considered in the estimation, there is no evidence that the model does not fit the data well.

Further evidence that the model indeed reflects the data is found in Figures 1 through 3. Figure 1 graphs the actual percentages of respondents retired by age against the percentages simulated with the model. In the actual percentages, the respondents are observed only at the dates of the surveys, which amounts to at most 5 ages per respondent. Further, since the surveys were taken two years apart, the percentages in adjacent years do not necessarily refer to the same respondents. The simulated percentages refer to all the respondents at every age; essentially, these are the percentages from the same simulations used to construct the theoretical moments used in the estimation procedure. Despite the fact that these percentages do not necessarily refer exactly to the same groups of respondents, the ability of the model to track actual retirement appears to be quite good.

Figure 2 examines a perhaps more challenging criteria for the estimated model, namely,

the ability to track retirement at various ages. The percentages for “actual” retirement in this figure are calculated as the differences in the percent retired at adjacent ages. Consequently, these numbers effectively net out the flow back into work from the flow into retirement. The simulated percentages follow a similar procedure of netting out the flows into and out of retirement. The most visible features of retirement are the two spikes at 62 and 65, with the spike at 62 being most prominent. The ability of the model to reproduce these two spikes at about the correct height is an important strength of the model. It should be emphasized that the model is able to capture these peaks of retirement without having any age dummies or splines of age in the utility function. As noted in the discussion of previous work, other structural models have had difficulty capturing these peaks without the use of age dummies in the utility function.

The feature of the model which differs from other estimated structural models and which allows it to capture these twin peaks is the distribution of time preference rates. For those with high time preference rates, two effects conspire to encourage retirement at 62. First, some respondents with high time preferences may not have pensions and may have non-working spouses. These individuals may need to wait until age 62 to retire simply because they do not have the assets to retire earlier. Among high time preference respondents, however, there is another factor which is more universally at work which encourages retirement at age 62. At age 62 individuals are first eligible to collect social security benefits. Although delays in collecting benefits result in actuarially fair adjustments to future benefits, high time preference respondents discount those future benefits sharply. As a result, they see the lost current benefits more like a tax which sharply reduces the reward to continued work. Their response to this perceived tax provides a sharp inducement to retire at 62.

Respondents with low time preference rates encounter a similar problem, only at age 65. After age 65, further delays in collecting benefits result in future increases which are less than actuarially fair. As a result, the present value of social security benefits starts to fall, and low time preference workers will perceive this as an offset to the wage. Effectively, the net return to working falls at age 65, and the lower net wage provides an inducement to retire.

Figure 3 illustrates another important feature of the model, namely, its ability to reflect instances where respondents return to work after a period of retirement. The shaded bars indicate the percentage of respondents who were not working in one survey year and who are observed to be working two years later, expressed as a fraction of those who are observed in both years. Such behavior is frequently referred to as a “reversal,” since it reverses the usual life-cycle flow from work to retirement. The figure suggests that between 3 and 4 percent of the sample returns to work in a two year period. This frequency is a substantial proportion of net retirements at most ages, as seen by comparing Figures 2 and 3. In addition, Figure 3 probably understates the true frequency of reversals, since it misses individuals who return to work for only a year or so between the surveys.

The white bars in Figure 3 indicate the simulated probabilities of reversals between the same two year intervals. The model appears to do a good job at simulating the overall level of reversals. There is a bit of an overprediction in the first interval, but for the remaining three intervals the actual and simulated percentage of reversals agree to within 0.1 percentage point. The simulations even do a reasonably good job at prediction the downturn in reversals during the 1996-1998 interval, which might be the result of the more than doubling of the stock market during the preceding three years.

Figure 4 looks at the total and net retirements by age. These results suggest a nontrivial amount of movement back and forth between work and retirement. At many ages the net retirements are only about half of the total, with the other half being offset by returns to work. The major exceptions are the two spike years at age 65 and especially at age 62. At age 62 not only is the total retirement very high, at close to 20 percent, but the percent returning to work is lower than in other years. The result is a very large jump in net retirement.

These results suggest that the model does a credible job at reproducing the basic features of observed retirement patterns. The overall levels of retirement at various ages track the observed retirement levels at those ages reasonably well. The model generates spikes in net retirement activity at ages 62 and 65 that are very close to the observed magnitude. Finally, it generates returns from retirement to work that are also close to the observed magnitude. Thus, the model shows promise of being able to shed light on various retirement issues, including the effects of changing asset returns on retirement.

VII. Effects of the Bubble

In this section use the model to investigate the effects of the tremendous increase in stock prices in the late 1990's and the subsequent collapse in the early 2000's on retirement behavior. The strategy is to simulate the model using the actual stock returns and again using more typical returns, and then to calculate the difference the excess returns make. In addition to doing this exercise for the whole period, it will also be of interest to examine the period of the increase separately from the period of the decline.

Figure 5 gives the results of the simulation using the actual returns from the entire period,

except for the period after 2002 where we use normal returns. Since the sample of individuals is fixed, retirement is increasing as the sample ages over the period. The shaded part of the bars denotes the percent of those who are simulated to be retired in any given year. The white part of the bars indicates the fraction of the sample who were not working at some point in the past but who are now working. The remaining part of the sample is working and has never previously stopped work. Note that as you consider the later years, and therefore an older sample, the fraction of those working who have been retired at some point in the past increases steadily.

The median year of simulated retirement in the sample occurs around 1997. This is important because of the timing of the stock market boom and bust. The boom occurred during the time period when many in the sample were making their retirement decisions, and therefore it seems probable that it had a great potential to affect those decisions. The bust occurred after most of the sample had retired, and consequently the possibility for the bust to affect retirement would seem to have been somewhat less. In later simulations we will examine the comparative effects had the boom or bust occurred at the same point in time in order to standardize for the effects of their occurrence at different points in time.

We first consider the effects of the stock market boom. For this analysis, we will consider the boom to be the years 1995 through 1999. During this five year span the S&P 500 index had annual rates of returns of 37 percent, 23 percent, 33 percent, 28 percent and 21 percent. Since inflation averaged only 2 to 3 percent over this period, almost all of this return was real. Moreover, since stock returns exhibit very little serial correlation over annual periods, it seems safe to assume that these returns were largely unanticipated. Although the typical respondent in the HRS had only about 50 percent of financial assets in stock, it appears that by 2000 anyone

who had started out the period with substantial savings had considerably more assets than they could have anticipated. In many cases the unexpected increase in assets would have exceeded their annual earnings from work. For anyone who was on the verge of retirement, this would have permitted a somewhat earlier retirement without giving up any of the anticipated standard of living. For those who had recently retired and might be considering going back to work, the extra assets reduced whatever incentive there would be in returning to work.

Figure 6 reports on the effect on simulated retirement of those five years of high returns. This figure compares the results of two simulations. In one simulation, normal returns are assumed for the years 2000 and beyond. In the other simulation, normal returns are assumed for the years 1995 and beyond. The only difference is in the years between 1995 and 1999, where one simulation assumes actual returns and the other assumes normal returns.

In this figure, the shaded parts of the bars indicate increase in retirement by individuals who have not been previously retired. The white parts of the bars indicate the increase in retirement caused by individuals staying retired rather than returning to work. The total heights of the bars thus indicate the increase in retirement. The increase in new retirements is somewhat but not overwhelmingly more than the increase in retirement caused by individuals not returning to work, and in this respect the effects approximately mimic the mix of retirements suggested in Figure 4.

The volume of retirements in Figure 6 increases more or less continuously during the period of the high stock market returns. At its peak in 1999, the percent retired in the sample is about 3.3 percentage points higher than it would be without the extra stock returns. The effect continues long after 1999, since the extra returns apply not just to those who retire during this

period but to individuals who would be retiring in the succeeding years. As time goes by, however, the effect of those returns gradually diminishes, both because individuals will have longer to adjust consumption to the higher levels of assets and because more and more of the sample would have retired anyway.

The effect on the average retirement age can be found by adding the effects at the various ages. When this is done, the conclusion is that the average retirement age for the entire sample is decreased by 0.25 years, or about three months. Considering that a large minority of the sample had no or only limited financial assets, the effect among the part of the sample with non-trivial assets would be considerably larger. Moreover, the magnitude of this effect appears reasonable. According to Venti and Wise (1999, Table 4), the average respondent had about \$79,000 in financial and retirement assets in the first wave of the survey. The 1995-1999 surge in the stock market would cause a portfolio of 50 percent stocks to have risen by about 60 percent more than the usual increase, so the unexpected gain would have been around \$47,000, or somewhat more than a year's worth of average earnings. The simulated decrease of 0.25 years in the average retirement age indicates that respondents do take part of this unexpected gain in the form of increased retirement, but that the major part of it goes toward increased consumption.

Figure 7 considers an analogous exercise to analyze the effects of the sharply lower stock prices in the early years of the new century. The returns to the S&P 500 were negative 9 percent for the year 2000, negative 12 percent for 2001, and a whopping negative 28 percent for the first nine months of 2002, which is the latest available as of the time this study is being done. Considering the small positive inflation in those years, the real returns were slightly worse.

The figure again compares two simulations, one using actual returns until 2002 and

normal returns thereafter and the other using actual returns until 1999 and normal returns thereafter. The negative effects on retirement of the stock declines develop rapidly, reaching a 3.2 percentage point decrease in 2003. This is very nearly the same effect, though in the opposite direction, as the maximum effect of the previous stock market increase, which is somewhat surprising in view of the fact that retirement in this sample is much further along in 2003 than in 1999. However, the diminution of the effect is considerably more rapid in the case of the stock market decline, perhaps as a consequence of the small number of individuals still working after 2005. Interestingly, the fraction of the decline that comes about as a result of the change in initial retirements as opposed to those staying in the retirement state and not returning to work is about the same in the cases both of the stock market boom and its subsequent bust.

Figure 8 considers the boom and the bust together by comparing a simulation which uses the actual returns until 2002 with one which uses normal returns for the entire period after 1995. To get a feel for this simulation, we need to have a comparison of the magnitudes of the extraordinary increases in the first part of the period with the losses in the second part. Over the five year period from the beginning of 1995 to the end of 1999, the cumulative return of the assumed portfolio (50 percent stocks, 5 percent long-term bonds, and 45 percent short term bills) was 57.4 percent over the cumulative normal returns. From the beginning of 2000 until the end of September 2002, the cumulative return of the portfolio was 33 percent below the cumulative normal returns. Combining the two periods, the cumulative return using the actual returns over the entire time exceeded the normal returns by only 6 percent. In other words, individuals lost in the bust essentially all that they had gained in the boom.

Considering this, the results illustrated in Figure 8 should not come as much of a surprise.

Retirement increased rapidly in the late 1990's in response to the stock market surge, but by 2003 retirement would be approximately back to the same level as would have been experienced if the bubble had never happened. After 2003, there is a small negative residual effect on retirement, perhaps due to the fact that at that time, the bust was more recent than the boom.

Figure 9 tries to standardize for the effects of the period during which the boom or bust occurs. For this figure, the base simulation replaces the returns in all years after 1991 with the normal rate of return. A second simulation replaces the returns in the years between 1991 and 1995 with the actual returns from the five years of the boom. The third simulation replaces the years 1993 through 1995 with the returns from the three years of the bust. This approach allows timing of the excess returns to end in the same year, 1995, at the point in time when the frequency of retirement is beginning to peak.

The comparison of the boom simulation with the base simulation is illustrated by the bars above the horizontal axis. The comparison of the bust simulation with the base is shown by the bars below the axis. The effects of the boom build faster and to a higher level, though not by much. Part of the reason perhaps is that the boom takes five years in developing, as opposed to only three for the bust. In any case, the general impression of this figure, which attempts to correct for the different timing of the boom and the bust on this given group of workers, is that the story presented in Figures 6 through 8 are not terribly much affected by the differential timing.

This story is that the extraordinary returns in the stock market in the late 1990's, which more than doubled stock prices and unexpectedly increased the value of a mixed portfolio by nearly 60 percent, increased retirement for the HRS sample of workers by over 3 percentage points by the turn of the century and would have decreased the average retirement age by about a

quarter of a year if it had not been interrupted. The subsequent decline in the market, which very nearly wiped out the gains that had been made during the preceding surge, effectively neutralized the effect of the preceding stock market gains on retirement. The effects of the bubble were to increase retirement as long as the bubble continued, but the continuing effects of the bubble after its end are probably minimal.

VIII. Conclusion

This study has looked at the effects of gyrations in the stock market on retirement. The vehicle of analysis is a structural dynamic stochastic model of the way individuals make retirement and saving choices in an uncertain world. The model includes individual variation both in retirement preferences and in time preferences. Returns to assets are unknown ahead of the fact, although the distribution of these returns is assumed to be known. In each period, the individual chooses whether to work or retire and decides on a level of saving or dissaving to maximize the sum of current and expected future utility, subject to the constraint that assets cannot become negative.

There are three criteria that any model designed to answer this question must satisfy in order to be credible. First, it must be able to generate a reasonably close approximation to the observed pattern of retirement by age, including the large increases in retirement at ages 62 and 65. Secondly, it must include a mechanism by which individuals can return to work that matches the observed returns to work reasonably well. Finally, it must include rates of returns to assets that are unknown beforehand and whose evolution can ultimately affect retirement and savings decisions. We have shown that the model presented and estimated in this analysis satisfies all

three of these criteria.

This model suggests that a string of high return years in the stock market can influence retirement in a significant way. In any particular year, the effect is limited by two factors. One factor is that a large minority of the population nearing retirement age does not have substantial assets either in the form of personal retirement assets or balances in defined contribution pensions. The second factor is that most individuals, even in the retirement age range, are not on the margin of the retirement decision in any particular year, and we would expect that only those who are close to reaching the retirement decision will be affected.

Even considering these two factors, the analysis suggests a stock market boom such as the one experienced in the late 1990's can raise retirement rates of individuals in the affected cohorts by around three percentage points per year in the years immediately following the boom, and that these effects persist for several years thereafter. A stock market collapse of comparable proportions would cause a roughly comparable reduction in retirement. In short, the model suggests there is at least a kernel of truth to the stories in the popular press about the gyrations in the stock market causing people to revise their retirement plans.

More generally, the model suggests that there is a non-trivial reaction of people to unexpected changes in wealth which occur late in their working careers. For instance, a similar response might apply to an unexpectedly large bequest, perhaps because grandma had amassed more wealth than anyone would have guessed. Although a large part of an unexpected change in wealth undoubtedly goes toward lifetime consumption, at least some of it goes toward changes in leisure through changes in retirement.

There are several obvious extensions to the model. Part-time work can be added to the

work effort choices, making it possible to examine flows to and from partial retirement. Health can be modeled as a state variable whose value changes over time, depending on prior health and other factors. Spouse retirement choices can be made endogenous, allowing a consideration of the interplay between the decisions of the two spouses. Wages can also be modeled as a state variable and allowed to evolve over time in a manner which is a priori uncertain. All of these extensions are desirable, but they are all computationally demanding. It appears unlikely that it will soon be feasible to add all of these extensions simultaneously. For the immediate future, we can add extensions selectively, determining which extensions are most important for the particular problem being addressed. Bit by bit, these kinds of structural dynamic stochastic models will become more and more realistic, and they will enable us to address a wider range of problems than is possible at present.

References

- Berkovec, James C. and Steven Stern. 1991. "Job Exit Behavior of Older Men". *Econometrica* 59(1): 189-210.
- Coile, Courtney, Peter Diamond, Jonathan Gruber and Alain Jousten. 2002. "Delays in Claiming Social Security Benefits". *Journal of Public Economics* 84(3): 357-385.
- Diamond, Peter A. and Jerry A. Hausman. 1984. "Individual Retirement and Savings Behavior". *Journal of Public Economics* 23(1-2): 81-114.
- French, Eric. 2002. "The Effects of Health, Wealth , and Wages on Labor Supply and Retirement Behavior". Xeroxed. Federal Reserve Bank of Chicago.
- Greene, William H. 2000. *Econometric Analysis*, Fourth Edition. Upper Saddle River: Prentice Hall.
- Gustman, Alan L. and F. Thomas Juster. 1996. "Income and Wealth of Older American Households: Modeling Issues for Public Policy Analysis". In Eric A. Hanshek and Nancy L. Maritato, editors. *Assessing Knowledge of Retirement Behavior*. Washington, D.C.: National Academy Press, 11-60.
- Gustman, Alan L. and Thomas L. Steinmeier. 2002a. "The Social Security Early Entitlement age in a Structural Model of Retirement and Wealth." National Bureau of Economic Research Working Paper 9183, September.
- _____. 2002b. "Retirement and Wealth". *Social Security Bulletin*, 2001/2002 64(2): 1-26.
- Ibbotson Associates. 2002. *Valuation Edition 2002 Yearbook*. Chicago: Ibbotson Associates.
- Rust, John. 1989. "A Dynamic Programming Model of Retirement Behavior". In David A.

- Wise, Editor, *The Economics of Aging*. Chicago: University of Chicago Press, pp. 359-398.
- _____. 1990. "Behavior of Male Workers at the End of the Life Cycle: An Empirical Analysis of States and Controls. In David A. Wise, Editor, *Issues in the Economics of Aging*. Chicago: University of Chicago Press, pp. 317-379.
- Rust, John and Christopher Phalen. 1997. "How Social Security and Medicare Affect Retirement". *Econometrica* 65(4):781-831.
- Samwick, Andrew. 1998. "Discount Rate Heterogeneity and Social Security Reform". *Journal of Development Economics*. 57: 117-146.
- Stock, James and David A. Wise. 1987. "Pensions and the Option Value of Work and Retirement". *Econometrica* 58(5): 1151-1180.
- Venti, Steven F. And David A. Wise. 1999. "Lifetime Earnings, Saving Choices, and Wealth at Retirement." In James P. Smith and Robert J. Willis, Editors, *Wealth, Work and Health: Innovations in Measurement in the Social Sciences*. Ann Arbor, University of Michigan Press, pp. 87-120.

Table 1
Derivation of the Sample

Reason	Observations Deleted	Observations Left
Males married in 1992		4,928
Spouse not interviewed in 1992	161	4,767
Respondent does not have a career job	793	3,974
Ages for retirement are missing or inconsistent	256	3,718
Ambiguity about whether jobs are ss covered	20	3,698
No FT earnings in ss record or self reported earnings	35	3,663
No self reported earnings, and ss earnings over limit	30	3,633
Relatively large business assets	291	3,342
No Pension Provider record in last job	868	2,474
FT years unavailable for spouse	169	2,305

Table 2
Estimated Results

Symbol	Description	Coefficient Value	t-statistic
α	Consumption parameter	-0.09	2.6
Parameters in β			
β_0	Constant	-9.688	405.2
β_a	Coefficient of Age ^a	0.054	7.8
β_h	Coefficient of Health ^b	3.32	4.7
σ_ϵ	Standard Deviation of ϵ ^b	7.94	10.9
ρ_ϵ	Autocorrelation of ϵ after Retirement	0.59	4.0
q value (26 degrees of freedom)		20.56	
Number of observations		2305	

^a The age variable is differenced from its approximate mean in the sample in order to facilitate estimation. The actual variable is age - 62.

^b These coefficients are expressed relative to the age coefficient, again to facilitate estimation. See text for explanation.

Figure 1
Percent Retired by Age




Figure 2
Net Retirements by Age




Fig 3
Percent Returning to Work




Figure 4
Total and Net Retirements by Age




Figure 5
Percent Retired by Year




Figure 6
Effects of 1995-1999 Stock Market Boom on Retirement




Figure 7
Effects of 2000-2002 Stock Market Bust on Retirement




Figure 8
Effects of 1995-2002 Stock Market Bubble on Retirement




Figure 9
Effects of Boom vs. Bust on Retirement

