

The Social Security Early Entitlement Age In A Structural Model of Retirement and Wealth

by

Alan L. Gustman
Dartmouth College and NBER
and
Thomas L. Steinmeier
Texas Tech University

June, 2002

Analyses of the effects of the social security early entitlement age were supported by a grant from the U.S. Social Security Administration (SSA) to the Michigan Retirement Research Center, with a subcontract to the National Bureau of Economic Research. Estimation of the model on which these simulations are based was supported by a grant from the National Institute on Aging (1R01AG13913-01A1) to NBER. The opinions and conclusions are solely those of the authors and should not be construed as representing the opinions or policy of SSA, NIA or any other agency of the Federal Government, the Michigan Retirement Research Center, or the National Bureau of Economic Research. Alan L. Gustman is Loren Berry Professor of Economics at Dartmouth College, Department of Economics, Hanover, N.H. 03755 (alan.l.gustman@dartmouth.edu). Thomas L. Steinmeier is Professor of Economics, Texas Tech University, Department of Economics, Lubbock, Texas 79409 (Thomas.Steinmeier@TTU.edu).

Abstract

This paper specifies and estimates a structural life cycle model of retirement and wealth and applies that model both to understand the role of the social security early entitlement age in creating a peak in retirements at age 62, and to simulate the effects of postponing the Social Security early entitlement age from 62 to 64.

The model includes a set of budget equations and a utility function. Data are from the first five waves of the Health and Retirement Study and are confined to married men. The budget equations fully incorporate the complex incentives from social security (relying mainly on respondents' earnings records), wage offers for full and partial retirement work, the incentives created by pensions (measured from employer provided plan descriptions), as well as the influence on retirement and saving of health status, family structure, and constraints from the firm side, such as layoffs and inability to reduce hours on the main job. Parameters of the utility function reflect the influences of time and leisure preference and vary among individuals. Estimation is based on the general method of moments.

Our estimates suggest that leisure and time preference are widely distributed among the population, with a bimodal distribution of time preference. Discount rates are either very low or very high. Those with high discount rates find the actuarial adjustments in social security benefits, which use a 3 percent real interest rate, to be inadequate. Once they reach age 62, the benefit accrual profile declines with age. This is the major explanation for the spike in retirement activity at 62. Liquidity constraints from inability to borrow on social security and pension benefits add to this effect.

Simulations with the model suggest that raising the social security early entitlement age from age 62 to 64 will shift about three fifths of the bunching of retirements at age 62 to age 64. The bunching amounts to about 8 percent of the population, so raising the social security early age of entitlement will have a substantial effect on the social security system and its finances.

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

Thirty years ago, the most common retirement age was 65.¹ Economically, this age made a lot of sense. This was the social security normal retirement age, and working beyond this age resulted effectively in giving up a substantial part of social security benefits that would otherwise have been paid. While additional work did result in an increase in later benefits, the increase was small in comparison to the benefits given up. In essence, potential workers felt a substantial reduction in the net reward for continued work, and many of them chose retirement instead. This was reinforced by the provisions of many private pension plans, which similarly penalized workers for work beyond 65, and by mandatory retirement provisions.

Today, the most common retirement age is 62, with 65 as only a secondary peak.² It is relatively easy to see why the peak at age 65 has diminished. Additional work beyond 65 now results in a much more substantial increases in social security benefits, reducing the incentives to retire (Gustman and Steinmeier, 1985; 1991). Age discrimination laws have eliminated formal age 65 mandatory retirement provisions in most jobs and severely curtailed the ability of defined benefit pensions to penalize workers for working beyond 65, and in any case defined benefit pensions are covering a continually diminishing share of the labor force. The reduction in retirement at age 65 thus appears to be a rational response to changes in incentives.

The new peak at age 62 is a different story. The obvious thing that happens at age 62 is that workers become eligible for early retirement benefits under social security, albeit at a reduced level. Working beyond 62 still results in giving up current benefits, but unlike the situation thirty years age at age 65, current work beyond 62 increases future benefits in a way

¹Burtless and Moffitt (1984); Gustman and Steinmeier (1984); Packard and Reno (1988).

²Burtless (1999); Coile, Diamond, Gruber and Jousten (1999); Quinn (1999); CBO (1999); Gustman and Steinmeier (2000a).

that is approximately actuarially fair. There may be some incentives from pensions to retire at 62, but these incentives have been eroding over time for the same reasons that the pension incentives to retire at age 65 have been reduced. In short, it is much harder to come up with a convincing story why workers are currently retiring in such numbers at age 62 than it was to explain the peak in retirement at age 65 thirty years ago.

The difficulty in explaining retirement at age 62 has serious and unfortunate implications for analyzing social security policy. It is well known that social security will face a funding crisis in the next several decades, and there is much debate as to how best to avoid the crisis. One step was taken in 1983 by gradually increasing the social security normal retirement age to 67 in the coming two decades, with concomitant reductions in benefits for those who take benefits at the early retirement age of 62. Some proposals include increasing the normal retirement age further, and some include increases in the early retirement age, both to encourage individuals to work longer and to avoid having benefits that are reduced too much. The only problem is that no one is quite sure what will happen to retirement if the early retirement age is increased. Although it is clear that social security is responsible for the spike in retirements at age 62, it is not clear what characteristic of social security is causing retirements to cluster at age 62 in the first place. There has not been any experience with changing the early retirement age since it was introduced in the early 1960's, so regression analysis using the early retirement age as a right-hand variable is not likely to be useful. The most promising approach is to use a structural retirement model, but the challenge is to come up with a structural model which replicates the retirement peak at age 62 without forcing the result by including discontinuities in the utility function at age 62.

To overcome this, we are going to focus on a feature of the utility function which is

usually assumed in structural models to be constant across the population, namely, the time preference rate. Allowing time preference rates to be heterogeneous across individuals simultaneously solves several problems in making structural retirement models more realistic. First, it solves the problem of creating incentives for individuals to retire at age 62. Individuals with high time preference rates who are considering retiring at age 62 will see the loss of current benefits clearly, but they will devalue the increase in future benefits even though those increases are actuarially fair. Secondly, individuals with high time preference rates are also likely not to have saved very much and may not be able to finance consumption if they retire before 62. Thirdly, heterogeneous time preference rates can explain the huge differentials in assets even among households with similar lifetime incomes, differences which are very difficult to explain otherwise. Fourth, they explain common complaints against the earnings test, in which individuals seem to see the lost benefits as a tax without considering the offsetting future benefit increases. If the individuals have a high time preference rate, they will in fact be ignoring the future benefit increases. Finally, higher time preference rates helps to explain why so many individuals collect benefits at the earliest opportunity, even when there might be some actuarial advantage to delaying for a year or two.

The next section looks at some of the previous work on this topic, while Section III examines some of the retirement patterns that a model must accommodate. The next two sections describe the model and the data that are used to estimate it. Section VI describes the estimation procedure and presents the basic estimates. The next section presents basic simulations with the model and simulations with alternative policies, particularly regarding the early retirement age. The final section contains some final thoughts.

II. Previous Literature

The literatures on saving and retirement converge in a model that would explain the spike in retirements at age 62 by heterogeneity in time preference.³ Consider first the literature on saving. Studies of wealth accumulation suggest there is strong evidence of heterogeneity in time preference. For example, using HRS data, Venti and Wise (1999) find a very wide dispersion of wealth within lifetime earnings classes, attributing these differences to heterogeneity in preferences for saving rather than to luck in securing high returns. Samwick (1998), using data from the Survey of Consumer Finances, finds that time preference rates are dichotomous, either clustering around 7 percent and below (there is also a peak in frequency at a negative rate, which Samwick interprets as reflecting a bequest motive), or 17 to 20 percent and above.

The second thread is the retirement literature.⁴ Most relevant are structural retirement models. These models typically explain retirement activity under the assumption of perfectly operating capital markets. They do a good job in explaining the peak in retirements in years that pension accrual rates peak and then decline, and have been applied to analyze retirements in particular firms (Fields and Mitchell, 1984; Stock and Wise, 1987), retirements in the general population (Gustman and Steinmeier, 1986), and retirements across countries (Gruber and Wise, 1999). Some have tried to use reduced form models with forward looking measures of benefit accrual (Coile and Gruber, 2000). Neither the reduced form nor the structural models have satisfactorily explained the peak in retirements at age 62. For example, Coile and Gruber include age dummies, and the coefficient on the dummy for age 62 is strong and significant. The problem is that all of these models assume perfectly operating capital markets, so that actuarially

³For a review of these literatures as well as a discussion of inconsistencies among them, see Gustman and Juster (1996).

⁴For a survey, see Lumsdaine and Mitchell (1999).

fair adjustments in social security benefits do not generate an incentive to retire at age 62. Similarly, the related literature on social security incentives does not explain the very early claiming of benefits.⁵ Despite what appears to be an actuarial incentive to postpone claiming, social security benefits are claimed as soon as they become available.

Diamond and Hausman (1984) and Kahn (1988) are two important contributions to the literature that do consider the joint relation between retirement and saving. Rust and Phalen (1997) and French (2002) present structural models of retirement and saving that consider the effects of imperfectly operating capital markets. Rust and Phalen go to the extreme of assuming that capital markets and insurance markets are both completely inoperative. Their main focus is on retiree health insurance, however. They assume that those without retiree health insurance could not purchase insurance in the market. This is contrary to fact. Many in their sample were not observed to purchase health insurance simply because they continued to work during the sample period, while others were free to purchase health insurance but chose not to. Their methodology leads to an overestimate of the effects of market imperfections on retirement. French (2002), using PSID data and imputing representative pension profiles gleaned from our earlier work, but without benefit of restricted pension or social security data, finds no effect of raising the social security early entitlement age on simulated retirements. This finding is in sharp contrast to the one we report below.

III. Retirement Statistics

In order to construct a sensible model, it is helpful to have a good idea of the pattern of

⁵See the discussion in Gustman and Steinmeier (2002a) of the various inconsistencies in reduced form models; e.g., benefits are discounted by a constant interest rate, but benefit claiming is assumed to take place as soon as possible.

retirement statistics that the model is attempting to replicate. Figure 1a presents data from the Health and Retirement Study (HRS) on the percentage of individuals who have left full-time work and those who have retired completely. Figure 1b looks at the net retirements for full-time workers, calculated as the difference in the retirement rate at one age and the value at the previous age. These should more properly be called “pseudo-retirements”, since they are based not on individual retirements but on the differences of percentage retired for adjacent ages. Figure 1b includes comparable data from the Retirement History Survey, which was taken during the late 1960's and 1970's (Gustman and Steinmeier, 2001).⁶ This figure clearly shows the shift of the retirement peak from age 65 to age 62.

In these figures, individuals working over 30 hours per week and more than 1560 hours per year are counted as full-time. Individuals working more than 100 hours per year but less than 25 hours per week or 1250 hours per year are counted as part time, and individuals not doing any work at all are counted as fully retired. Individuals who fall between full time and part time or between part time and retired are classified on the basis of self reports. The observations are made on the dates of the surveys; since both surveys are every-other-year surveys, figures for adjacent years generally refer to two separate groups of respondents.

In addition to the spikes in retirement, the other thing to note in Figure 1a is the pattern of part-time work. At ages before 60, most individuals who are working at all are working full-time and only a small fraction are working part-time. By the late 60's however, the number working part-time approaches and even exceeds the number working full-time. One of the aspects of retirement that we would like the model to capture is the increasing preference for part-time over

⁶ Since the youngest respondents in the RHS were 58 in the first wave of the survey, pseudo-retirements could not be constructed prior to age 59.

full-time work for those individuals still in the labor force at later ages.

Figure 2 presents the retirement percentages for individuals who report that they are in fair or poor health at the time of the survey. The number of observations at age 60 in this figure are roughly a sixth of the number of observations at the corresponding age in Figure 1, and the behavior of the percentages is correspondingly a bit more erratic. However, there are two things which are readily apparent in comparing Figure 2 to Figure 1. First, and as no surprise, the percentage of workers in poor health who are retired is uniformly higher than the corresponding percentages of healthy workers. Second, and somewhat surprisingly, the percentage of part-time workers among those in poor health is much lower, both as a percentage of the population and as a percentage of those working, than is true for healthy workers. To the extent that poor health reduces work effort, it evidently prompts workers to drop out of the labor force completely and retire rather than to reduce their work effort to part-time. Again, this is a pattern of behavior that we will want the model to capture if possible.

IV. The Model

The model is a rather standard life cycle model consisting of a utility function and a budget constraint. The utility function is given by

$$U = \int_0^T e^{-\rho t} s_t \left[\frac{1}{\alpha} C_t^\alpha + h_t L_t \right] dt$$

In this equation, T is the maximum lifespan and s_t is the probability that the individual will survive until time t . C_t is consumption in real terms at time t , and L_t is the utility value of leisure (retirement) at time t . L_t takes on a value of 0 for full-time work, 1 for complete retirement, and a value L_p for part-time work. If leisure is assumed to have diminishing returns, and if part-time work is half-time, L_p should lie between 0.5 and 1. h_t is a term which measures

the relative value of leisure relative to consumption, and is assumed to be increasing in value as the individual ages and finds work more difficult.

This utility function contains several elements which may be considered to be heterogeneous between different individuals. Different individuals almost certainly have different time preference rates as reflected by the parameter ρ in the discount factor $e^{-\rho t}$. They may also have different preferences for leisure as reflected in h_t . Finally, individuals may be heterogeneous with regard to L_p , which measures the relative attractiveness of part-time vs. full-time work.

Since, for married individuals, a substantial part of social security benefits accrue to the spouse as either spouse benefits or survivor benefits, it is critical to include the value of the spouse benefits in the model. To do this, we model the death of the two spouses as stochastic and include three consumption terms in the utility function: one for the instance where both partners are alive at time t , one where only the husband is alive, and one where only the wife is alive. The associated survival probabilities refer to the three states and implicitly contain the transition probabilities that one of the spouses will pass away and leave the other spouse remaining. In recognition that consumption is more valuable while both spouses are alive, the consumption function is adjusted so that the marginal utility of C_t for a couple is equal to the marginal utility of approximately $0.7 C_t$ for a remaining spouse. The respondent is assumed to give equal weight to his own and his spouses consumption, although if the time preference rate is high the value of survivors benefits will not be very high in any case.

The individual is assumed to start out with no assets. The asset accumulation over time is given by

$$A_t = (1+r)A_{t-1} + W_t(1-L_t) + E_t + B_t - C_t$$

where A_t is the level of real assets at time t , r is the real interest rate, W_t is the real wage rate, E_t is the earnings of the spouse, and B_t is the level of social security and/or pension benefits at time t . If the individual is working, the wage rate may depend on whether the work is full-time or part-time. The level of benefits at time t depends on the previous decisions of the individual as to when to leave full-time employment and when to retire fully. Note that this model does not calculate the value of accruals to social security and pensions directly, but the value of the accruals is implicit in the model because work during one period will affect the value of social security and pension benefits in later periods. The implicit value of these accruals, of course, depends strongly on the time preference rate.

In the model, A_t is restricted to be non-negative for each period. This assumption, along with heterogeneous time preferences, is an important part of what drives this model. Otherwise, individuals with high time preferences could simply borrow against their future labor income and social security and pension benefits and consume most of their lifetime income early on. The requirement for nonnegative assets prevents this from happening, and instead such individuals are likely to consume essentially all of their income each period, acquiring very little in the way of savings.

For married individuals, the budget constraint also includes the earnings, social security benefits, and pension benefits for the spouse. In this paper, the spouse's retirement decision is taken as exogenous, so the income of the spouse is taken as essentially exogenous. We intend to generalize the model in later work to include the possibility that the retirement decisions of the couple are joint, but that is a substantial undertaking and will probably preclude looking at partial retirement. The current model does consider the effects of the respondent's work decisions on

the social security benefits which will be available to the spouse, however.⁷

V. Data

The dataset used in this study is the Health and Retirement Study (HRS).⁸ This is a nationally representative sample of households which contains at least one individual born between 1931 and 1941. The study was started in 1992 and conducts interviews with the respondents every two years. The last year for which data are available is 2000. Since the ages of these individuals was 51 to 61 in 1992, many of the respondents have retired over the sample period, making it a nearly ideal data set for studying retirement. In addition, and unlike the previous Retirement History Survey, the HRS interviewed the spouses with the same questionnaire, even if the spouses were not in the specified birth cohorts. This gives invaluable information about the financial resources that the spouses are contributing to the household.

An important advantage of the HRS is that in addition to the household interviews, the survey attempted to collect information from the social security records of the respondents and from the pension plans of the employers, if applicable. Social security records could be obtained only with the explicit permission of the respondents, and ultimately they were obtained for about three-quarters of the sample. These are valuable not only because they provide an earnings history which is difficult for the respondents to recall accurately, but they also allow potential social security benefits to be calculated much more accurately. The survey also attempted to collect the summary plan descriptions from the employers of the respondents. These plans were

⁷For models of family retirement that assume a perfectly operating capital market, see Gustman and Steinmeier (2000b, 2002b).

⁸The Health and Retirement Study is supported principally by a grant from the National Institute on Aging to the Institute for Social Research at the University of Michigan. Additional support is provided by the Social Security Administration and other federal agencies.

coded into a program that enables one to calculate pension benefits, given the earnings and retirement ages of the respondents. The plans were collected and successfully coded for about two-thirds of the respondents who had a pension in their last observed job, and for about one-third of the respondents who were covered in a previous job. Together, the social security and pension records allow a much more accurate computation of the benefits which account for half of the typical retiree's post-retirement income.

Table 1 provides an accounting of the derivation of the final sample. The focus of this study is on married males for two reasons. First, it appears that it would be inappropriate to include both males and females in the same analysis, given the differences other researchers have found in their retirement patterns. Secondly, among the males, most of the unmarried males have previously been married, but because of the way the HRS is constructed there is little information available on the former spouses. Thus, the study starts off with 4,928 potential male respondents who were married as of the first interview in 1992.

It is evident from the table that there are the two main causes of deletion from the final sample. The first is that the respondent has not been committed to the labor force, which in this study is defined as having worked full-time in at least half the years after age 40. Respondents not satisfying this criterion are occasional workers at best, and for them the very concept of retirement becomes somewhat hazy. The other major cause of deletion is that the respondent reported that there was a pension in the last full-time job, but the survey failed to obtain a usable pension from the firm. Since the incentives to retire vary so widely in pensions, the lack of a usable pension means that a potentially very important piece of the retirement decision is unobserved. Note that missing pensions in previous jobs are not causes for deletion, since those jobs are not involved in the retirement decision. Rather, they are imputed.

There are several other causes of deletion which are not completely trivial but not as important as the previous two. In some cases the spouse was not interviewed, and in other the spouse was interviewed but failed to provide enough information to ascertain whether she was a steady worker or an occasional worker. Some individuals did not provide the hours information necessary to calculate the retirement status at any age. Others had business assets that were greater than the sum of financial wealth plus real estate wealth. Since assets are used to indicate time preference rates, and since individuals holding large business assets are usually holding them for reasons other than time preference, we felt it best to exclude these individuals from the estimates. Other minor causes of exclusion are an inability to determine if substantial jobs were covered by social security and an inability to calculate earnings because the individual never reported earnings in the household survey and either did not give permission for obtaining the social security record or had several years of earnings over the limit in the record.

The dependent variables in the estimation are the retirement statuses as of the survey dates, as discussed in the last section. Observations after an individual has applied for social security disability benefits are not counted. Essentially, we are assuming that disability severe enough to obtain disability benefits effectively ends the individual's ability to make the decision whether or not to work. These observations are treated as right censored, much as though the respondent had died or dropped out of the survey. In Section VII we argue that this treatment of DI applicants has only a small effect on our estimates of the reaction of retirements to changes in the age of early social security benefit entitlement.

The main explanatory variables are the streams of potential earnings and benefits that would accrue to the individuals if they retire at alternative ages. For respondents who allowed access to the social security records, earnings prior to 1991 are taken from those records,

augmented by the self-reported data for periods of non covered work or years in which earnings were at the earnings limit. Jobs are separated into full-time and part-time jobs based on the information in the respondent interviews. Earnings after 1991 and, for those who had already retired, potential earnings for years after retirement, are projected from the experience and tenure coefficients taken from a wage regressions reported in the Appendix. If the individuals indicate that they were fired or laid off from the last full-time job, the projections assume a new job beginning with zero tenure. For those without social security records, the earnings sequences are projected on the basis of the job histories in the 1992 respondent interviews, with periods between the reported jobs supplemented by information in the 1996 survey on the total amount of work experience. If respondents report that they can reduce effort in a full-time job to half-time, the hourly wage rate in partial retirement is assumed to be the same as for full-time work, unless observed part-time wages indicate otherwise.

For individuals who reported that they had pensions, the pension benefits for alternative retirement dates are calculated by the pension provider program supplied by the HRS. For defined contribution plans, these benefits are assumed to be available as a lump sum, while for defined benefit plans they are available as an annuity. Benefits are imputed for jobs previous to the last full-time job if the plan information is missing, but if the plan is missing for the last full-time job, the individual is omitted from the sample on the basis that the observation is missing very important incentives affecting the retirement date.

Social security benefits are also calculated for respondents and their spouses. These benefits include those available to a couple and those available to either spouse on the death of the other spouse. The benefits are calculated based on the expected retirement date of the spouse, but for alternative retirement dates for the respondent. Thus, they include the work

incentives for the respondent that operate through the respondent's own benefits and the spouse and survivor benefits for the spouse. As mentioned before, the model includes uncertainty to the extent that either spouse has a random chance of mortality, and the available benefits depend on this mortality experience.

The other explanatory variables in the analysis are age, cohort, and health status. Health status is a binary variable that takes on a value of one if the individual has previously reported poor health in at least two consecutive surveys or is in the first year of a stretch of poor health that will ultimately last at least two surveys. This criterion is chosen because respondents frequently report poor health in one survey but not in subsequent surveys, but if the report persists for two surveys it is frequently more or less permanent. In this measure, respondents are classified in poor health if in response to the question "Would you say your health is excellent, very good, good, fair, or poor?" they answer "fair" or "poor."

VI. Estimates

The first step in obtaining estimates is to specify the stochastic elements to the model. These concern primarily the components of the heterogeneity terms having to do with leisure. We follow this discussion with a discussion of the estimation technique and finish with the results of the estimation.

Stochastic Elements

The term h_t in the model, which is the coefficient of the leisure term, is taken to be of the exponential form $e^{\beta X_t + \varepsilon}$. In this form, β is taken to be a coefficient vector and X_t is a vector of explanatory variables including a constant, age, health status, and cohort. As the age variable increases over time, leisure becomes more attractive for reasons of gradually

deteriorating physical capacities and similar factors which eventually result in retirement. It is important to note that this term explicitly does not contain any binary age variables or splines in age which might encourage retirement at a particular age. The variable ϵ in this expression is a term randomly drawn from a normal distribution, but which is constant over time for a particular individual. The idea here is that tastes for leisure (or conversely, tastes for work) vary over the population, with some individuals who have a high taste for leisure activities and others who do not mind work, or even enjoy it.

The other heterogeneity term affecting leisure is L_p , the utility value of part-time work as opposed to complete retirement. If L_p has a value close to 0.5, then the individual regards the disutility of an hour of work as about the same regardless of whether the work is full-time or part-time. If it has a value closer to unity, then the individual does not mind part-time work too much but really dislikes having to work full-time. Again, this term is assumed to come from a random draw from a distribution for the population. Initial work used a truncated normal distribution to allow for the possibility that the distribution has a peak between 0.5 and unity and falls off to either side, but the implied peak of the distribution was always considerable less than 0.5, and the range between 0.5 and unity was always on one of the tails. For this reason, we changed to an truncated exponential distribution function $k e^{\delta L_p}$, where k is the factor required for the distribution function to integrate to unity between 0.5 and 1. If δ is positive, values of L_p toward unity will be more common, while if δ is negative, values near 0.5 will be more common.

The model in this form, however, produces part-time work which is too common at younger ages and too rare at later ages. As discussed in the section on descriptive retirement statistics, the ratio of part-time to full-time work increases at older ages. To accommodate this,

we add an age term to the distribution: $\delta = \delta_0 + \delta_a \text{Age}$. If δ_a is positive, the distribution of L_p shifts upward over time, so that part-time work becomes relatively more attractive to full-time work, and hence more common. In this formulation, individuals do not get a completely new draw of L_p at every age. Rather, they get a draw from L_p at a particular age, and they are assumed to maintain the same relative position in the distribution of L_p 's even as that distribution shifts over time. A similar problem occurs with respect to health. Recall that the descriptive results indicated that health greatly reduced the amount of partial retirement. When we tried to introduce health into the δ term in the same way we introduced age, the coefficient tended toward minus infinity, indicating that when a health problem causes individuals to retire, it usually causes complete retirement rather than partial retirement. To reflect these results, the value of L_p is set to 0.5 for individuals with health problems.

The treatment of ϵ and L_p in this model is technically known as random effects, wherein the values are randomly drawn from assumed distributions, and obtaining the parameters of these distributions are part of the estimation. On the other hand, time preference, which is the third element of heterogeneity in the model, is treated as a fixed effect. Recall that the symptom of high time preference rates is low assets, and vice versa. Hence the model uses the asset levels in the 1992 survey to infer the time preference rate. That is, time preference is calculated as the value which would produce the observed amount of assets in 1992. Assets for this purposes are taken to be financial, real estate and business assets. There are two circumstances where this calculation is slightly modified. First, some individuals have more assets than they would be calculated to have even at a zero time preference rate. The calculations do allow for inheritances, but such individuals may have had better investment returns than the calculations allow, or they may be higher for other reasons. In any case, they almost undoubtedly do signal low time

preference rates, and are assigned a value of zero for the time preference rate. A second group has zero assets, even allowing for defined contribution lump sums which would be available if they retired. Because they are at a corner solution with regard to assets, the time preference rates for such individuals cannot be completely determined. But because only individuals with very high time preference rates would have zero assets on the verge of retirement, such individuals are assigned an arbitrarily high time preference rate which essentially implies that they will consume their income every period.

The parameters to be estimated are as follows: α , the consumption parameter, the four elements of β , the linear term describing the value of leisure, δ_0 and δ_a , which describe the relative attractiveness of partial retirement, and σ_e , which gives the variation of leisure preferences. It is helpful to consider the kind of information which would help identify these parameters. β_0 determines the average retirement age, and σ_e governs the general width of the distribution of retirement ages. β_h , the coefficient of health, determines how much earlier individuals in poor health retire, and β_c does the same thing with respect to older vs younger cohorts. β_a , the coefficient of the age variable, determines how sensitive retirement ages are to economic incentives. A high value for β_a means that the value of leisure is rising so fast that the individual will retire at about the same time regardless of economic incentives, while a low value signifies the opposite. Higher values of δ_0 will produce more partial retirement, while higher values of δ_a will bring about a more rapid growth in partial retirement with age. Finally, α essentially determines the retirement rates of high income vs. low income respondents. A value of α close to one means that the marginal utility per dollar of income is nearly constant. The utility value of earnings for high income individuals is great relative to the utility value of leisure, and such individuals should retire relatively late. If α is strongly negative, then the

marginal utility of consumption declines quickly, and high income individuals who can achieve a given level of consumption with relatively little work should retire early.

The Estimation Procedure

In previous work we have used maximum likelihood as an estimation procedure. However, it has become apparent that maximum likelihood estimates are quite sensitive to a relatively small number of cases where the model would predict that the probability of an observed result is very small or zero. To illustrate, consider an individual with 20 years of tenure in a job whose pension increases dramatically with 20 years of tenure. If the observed retirement is at 21 years, the log likelihood is reduced moderately compared to a retirement at 20 years. However, the log likelihood would be dramatically reduced if the individual were found to retire after 19 years on the job, since such behavior makes sense only if the individual gives economic factors a very low weight in the retirement decision. However, in many cases it is difficult to distinguish these situations, especially since the 1992 HRS did not ask the month of hire.

The usual result of very low likelihood observations is to drive up the coefficient of the age variable in β , which in turn means that retirements occur more as a result of preferences and less as a result of any response to incentives. We have become uneasy with the disproportionate effect of these individuals on the estimates, especially given the asymmetric effects of these individuals on the likelihood function and uncertainties regarding the construction of the budget sets and the health variables. Therefore, we have turned to an approach which gives less weight to these observations.

The generalized method of moments (GMM) estimator provides one such approach. GMM seeks parameter values which minimize

$$q = m' w^{-1} m,$$

where m is a vector of the difference between the observed and theoretical moments and w is a weighting matrix.⁹ The theoretical moments are the moments, such as the percentage of individuals retired, calculated from the model, and these are compared to the corresponding actual moments, such as the observed percentage of individuals retired. Note that the theoretical moments depend on the parameter values. The weighting matrix w is given for large samples by

$$w = \sum_{i=1}^n m_i m_i'$$

where i runs over the observations and m_i is the column vector of moment differences for observation i . The matrix w is $L \times L$, where L is the number of moments.

If the model is correct, then q is a random variable distributed like a χ^2 distribution with $L - k$ degrees of freedom, where k is the number of parameters being estimated. If the number of moments is equal to the number of parameters, then there enough parameters so that each theoretical moment can be adjusted to match the observed moment exactly, and q ought to be zero. With more moments, it is no longer possible to guarantee that the theoretical moments can match the observed moments exactly, but if the model is correct, the theoretical moments should not differ from the actual moments by a large amount. The value q essentially measures how large the amount has to be before the model is called into question. In GMM, the variances of the parameters is calculated by

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

where θ is the vector of parameters and G is the matrix of the derivatives of the moments with respect to the parameters.

⁹ The notation is taken from Greene (2000) but is similar to the notation in many other econometrics texts.

For the present estimation, we use 46 moments. 13 of the moments are the percentages working full-time for each age between 54 and 66. The remaining moments are taken at ages 55, 58, 60, 62 and 65, since these have been common ages in retirement plans. For each of these ages, moments are calculated for the percentage fully retired, the percent of those with a health problem who are working full-time, the percent of those with a health problem who are fully retired, the percent of those born before 1934 who are working full-time, the percent of those born after 1938 who are working full-time, the percent of those with lifetime incomes below \$1,250,000 who are working full-time, and the percent of those with lifetime incomes above \$1,900,000 who are working full-time. The income figures and vintages are chosen to divide the sample roughly into thirds. Two moments are lost because those born before 1934 could not have been 55 in the survey time frame, and those born after 1938 could not have reached 65 within this time frame.

To calculate the theoretical moments for a given set of parameters, we simulate the model for each individual 10,000 times and take the percentage of the simulations which result in the given labor force status at the stipulated age. The moments are thus simulated moments, but due to the large number of simulations the simulated moments should not differ appreciably from the theoretical moments.

Estimates

The estimated parameter values are presented in Table 2. All of the estimated parameters have the expected signs, and with the exception of the cohort coefficient, which is numerically very small, all of the coefficients are significantly different from zero at conventional levels. The q value for the model is 46.92, which is considerably below the 95% limit of approximately 53.38 for a χ^2 distribution with 38 degrees of freedom. This indicates that there is nothing in

the model which is strongly rejected by the evidence in the sample, at least using these 46 moments.

Perhaps the most important parameter in the model is the coefficient of age in the β vector. As mentioned before, this is the coefficient that determines how sensitive retirement in the model is to economic incentives. The parameter value here of 0.076 means that leisure is becoming more valuable in utility terms by almost 8 percent per year. This is considerably less than the values of the parameter we have obtained for similar models using maximum likelihood, which have ranged from 0.24 up. With the parameter of that magnitude, simulations show no more than a minor increase in retirement at age 62 at best. As we shall see in the next section, the parameter obtained with GMM produces a spike of approximately the correct size.

Among the other parameters of interest, poor health has about the same effect on retirement as being four years older, while there appear to be almost no cohort effects at all. The standard deviation of leisure preference has a value not quite six times the value of age coefficient, suggesting that roughly two-thirds of retirements should occur within a 12 year range around the median retirement age. This appears to match approximately the statistics presented in Figure 1. Please note that the parameters for health, cohort, and the standard deviation of ϵ are measured relative to the age coefficient. This means, for example that the actual health coefficient in the model is 3.92 times the age coefficient of 0.076, or 0.298.

Since the time preference rates are estimated as fixed effects, another interesting result to come out of the estimates is the distribution of time preference rates. This distribution is illustrated by Figure 3. The distribution is illuminating. About 40 percent of these respondents have time preference rates of less than 5 percent per year, while another 21 percent have rates between 5 percent and 10 percent. While this means that over half of the sample has a time

preference rate that is not excessively high, considerably less than half the sample appears to have a time preference rate that can be said to approximate the risk-free real interest rate, which by most accounts is considerably less than 5 percent. At the other end of the spectrum, 27 percent of the respondents appear to have essentially an earn-it-and-spend-it mentality and have not saved any assets at all. These are the individuals who are probably most vulnerable to changes in social security or any other change that affect income security in retirement.

VII. Analysis of Changes in Early Retirement Age

We analyze the effects of an increase in the early age of eligibility for social security benefits by comparing simulations of the model using the current early retirement age with simulations of a model in which the early retirement age has been increased. To do the simulations, we use the same sample that we use in the estimation, including the implied values of the time preference parameter for each respondent. 1000 random combinations of the leisure preference parameter ϵ and the partial retirement leisure L_p are drawn, and the model is solved to yield the optimal age of retirement from full-time work and the optimal age of complete retirement. These are tabulated to yield retirement distributions for the entire sample.

Base simulations

Figure 4 gives the results of the simulations using the current early retirement age. We shall call these simulations the base simulations. Figure 4a gives the percent of respondents who are simulated to be retired from full-time work and to be completely retired by age. These may be compared to the comparable graph in Figure 1a. Figure 4b gives the percent retiring from full-time work for each age, and compares it to the observed percent retiring from full-time work.

The most important feature of Figure 4 is that it does a credible job of reproducing the

retirement spike at age 62, which is crucial for any attempt to model the effects of changing the early retirement age. It does not do this perfectly, since the simulated percent retiring at age 62 is a couple of percentage points below the corresponding number in the actual data. It is a little closer when we consider the difference between the percentage retiring at age 62 and the percentage retiring at the adjacent ages. For the actual data, the spike is the difference between the 15.2 percent and the average of 5.3 and 7.4 percent, yielding a spike of 8.9 percentage points above the adjacent years. For the simulated data, the corresponding spike is 8.1 percentage points. Thus, the simulated spike and the spike in the actual data are within a percentage point or so of each other.

The simulation also picks up the spike in retirement from full-time work at age 65 fairly well. It underpredicts complete retirement at age 62 and overpredicts it at age 65. The simulated percentages retired from full-time work and those completely retired are within a few percentage points of the actual numbers. The main difference is that the simulated percentages of retirement seem to increase a little more rapidly between ages 55 and age 62 than is in fact the case. However, the simulations do replicate the pattern that part-time work is an increasing fraction of total work at older ages.

In these results, it is crucial to remember that the simulated spike does not arise in any way because the preferences contain some kind of feature which causes the individual to want to retire preferentially at age 62. Indeed, being able to generate a spike without relying on such discontinuous preferences is one of the primary objectives of the model. If the preferences were to exhibit a tendency to retire at 62, then it would become very difficult to judge whether, and to what extent, such preferences would shift to another age if the early retirement age were to be moved. Instead, the spike in this model comes entirely from the interaction of continuous

preferences with the budget set facing the individuals in the sample. There is no need to make arbitrary assumptions about how preferences would change if the early retirement age increases; the preferences are assumed to remain constant.

Altering the Early Retirement Age

Figure 5 presents results of simulations in which the early age of entitlement to social security benefits is increased from age 62 to age 64. This age is chosen because it would restore the three year differential between the early and normal retirement ages that was in effect before the 1983 Social Security Amendments, but obviously other ages could be simulated.

The most obvious result of this figure is the percentage of individuals retiring from full-time work at age 62 is reduced by almost 5 percentage points. The size of the spike relative to retirement in adjacent years is 3.0 percent (8.0 percent minus the average of 4.8 and 5.2 percent). Thus, the magnitude of the spike decreases from 8.1 percent to 3.0 percent, a reduction of over three-fifths. It is not difficult to see where the individuals who formerly retired at 62 went; in the simulations with the early entitlement age increased to 64, there is now a substantial spike in retirement from full-time work at 64.

The clear implication of these results is that if the early entitlement age is increased, a substantial fraction of those who currently retire at 62 will move their retirements to the new early entitlement age. Some of these individuals are liquidity constrained and simply cannot afford to retire before the social security benefits become available. But the retirement spike at age 62 also occurs among individuals with pensions, which could provide income between retirement and age 62, and individuals with working spouses, whose earnings could also provide income before age 62. This suggests that another explanation is paramount. Individuals who have high time preference rates simply do not see the value of the future increases in benefits if

they were to work another year. For these individuals, working after becoming eligible for social security benefits simply means that they are giving up a year's worth of benefits. When they become eligible for benefits, their net compensation drops from their earnings to their earnings less their potential benefits, and this may be a substantial drop. Their response to such a drop in the model, and probably in reality, is to retire when they become eligible for benefits, whenever that may be.

What of the remaining individuals who still retire at age 62 even when the early retirement age moves to 64? The remaining part of the spike is probably due to the normal retirement provisions in many pension plans which specify a normal retirement age of 62 in the plan. In many plans, there are reductions for receiving benefits before the normal retirement date, but there are rarely adjustments for retiring later. Thus, workers with normal retirement ages of 62 in their plans will see their benefits growing much more slowly after age 62 than before, and this may contribute to retirement. Age 60 is another common normal retirement age in the plans, which accounts for the spike in the simulations at that age. These simulations assume that the plans will not change, but in reality two things could happen. If the companies picked age 62 initially because that is the social security early retirement age, they conceivably could move their normal retirement age to 64, which would reinforce the effects of the social security change. On the other hand, if the companies chose 62 because they wanted to provide strong incentives to leave at that age, they could alter their plans with even stronger incentives to retire at age 62, and thus offset to some degree the social security change. Knowledge about which of these would in fact happen awaits a better understanding of the factors driving the way firms determine the specifics of their pension plans.

Disability Insurance

At this point, a number of retirement experts will object to these results. “The respondents who retire at age 62 in the base model will not delay retirement until age 64 if you increase the normal retirement age,” they will say. “Instead they are likely to enter the disability insurance program and retire at age 62 anyway, and you have not considered disability insurance.” They would be correct. We have not considered disability insurance, at least up to this point. All of the descriptive statistics, all of the estimation, and all of the simulations are for individuals who have not entered the disability insurance program. Entry into the disability insurance program is treated much like death in the analysis: it exogenously drops a respondent out of the sample. The experts’ complaint is that it is not exogenous, and particularly a change like increasing the early retirement age could lead to large enough changes in the take-up of disability insurance to affect the results.

All of this is true, but it is possible to bound the effect that the failure to consider disability insurance could have on the results. To do this, we first consider the number of men who enter the disability rolls at each age, as given by the Annual Statistical Supplements to the Social Security Bulletin for 2000 and 2001, Table 6A4:

Age	1999	2000
55	16,500	13,300
56	13,400	14,100
57	14,800	15,800
58	13,100	12,300
59	15,700	13,800
60	14,300	14,600
61	15,500	15,400
62	13,500	13,900
63	7,300	9,000
64	7,200	6,900

It is important to note that these entries are individuals, not thousands or millions of

individuals. In both years, the number of new disability beneficiaries appears to be about 15,000 individuals at each age between 55 and 61. The trend of these numbers by age up to and including 61 appears to be relatively flat; there certainly does not appear to be any evidence that the number of awards is increasing precipitously as the social security early retirement age approaches. After the attainment of the early retirement age, there appears to be a slight tail-off of awards, on the order of perhaps 1,500, at age 62, and a further tail-off of 7,000 to 8,000 at ages 63 and 64.

If there were no Social Security early retirement benefits, it would appear at worst that the number of disability applications would continue to be about 15,000 per year. Thus, the upper bound of the increase in disability beneficiaries if there were no early retirement benefits appears to be about 1,500 at age 62 and 7,500 at age 63. But what is the relevant population? In the sample we are looking at career workers, which we have defined to be those who work in half the years between age 40 and retirement. Half is also the number relevant fraction to being disability insured (20 quarters of coverage in the last 10 years), so we can look at the number who are disability insured in the years prior to retirement. If we look at the age range 50-54, we find that 83.75% of the males were disability insured in the years before a career worker would have retired (2000 Annual Statistical Supplement, Tables 4C2 and 4C5).¹⁰ Applying this to the 60-64 populations, it appears that there are about 4.63 million career males who either are disability insured or would be disability insured if they were still working. This translates to about 925,000 disability insured males per year of age.

Doing the divisions, this implies that the number of disabled workers might have

¹⁰ Interestingly, in Table 1 the percent of married males who are considered to be career workers is 83.36 percent (3,947 out of 4,767).

increased by a maximum of 0.16 percent of the population at age 62 and 0.81 percent at age 63 if the early retirement age were raised to age 64. In the results, the retirement spike at age 62 is reduced from about 8 percent of population to about 3 percent of the population at age 62 by the increased early retirement age. A change of 0.16 percent is simply too small to have any significant bearing on these results. Although the total number of workers who begin disability insurance in any given year is sizeable (about 330,000 workers in 2000), most of the disability awards are made to individuals far away from retirement, and the number who enter the rolls near retirement are relatively small.

Time Preference vs. Leisure Preference

In the formulation of the utility function above, the time preference rate applies to the total utility in any period vis-a-vis utility in another period. The trade off between consumption and leisure within a period is not directly affected by the time preference rate. Nevertheless, it seems reasonable to expect that individuals with high time preference rates may also have a higher preference for leisure. This hypothesis is further suggested by the fact that when the retirement distributions in Figure 1 are computed for high vs. low time preference groups, individuals with high measured time preference rates tend to retire earlier than do those with low time preferences.

A correlation between time preference and leisure preference can be accommodated in the model by including the time preference parameter in the βX term, which indicates the strength of leisure preferences. When the time preference term is included, however, its estimate is insignificant and very small numerically. To investigate this result further, we ran the simulations separately for groups with low, medium, and high time preferences. Even without any correlation between time preference and leisure preference, the group with high time

preferences are simulated to retire earlier than the others, and by just about the same magnitude as we find in the observed retirements. Evidently the possibility of future rewards to current work makes current work more attractive for those with low time preference rates, and this leads to later retirement.

VIII. Conclusions

In many of the currently estimated retirement models, it is difficult to generate the observed retirement peak at age 62 without using binary variables for age. The principal difficulty is that the social security early retirement penalty is approximately actuarially fair between ages 62 and the normal retirement age. And if a model does not predict a bunching of retirement at the early retirement age of 62, it is unlikely to have much to say about what will happen if that age is changed.

There are several clues in the data that point to heterogeneous time preference rates as a way around this problem. One clue is the wide variation in asset holding among the population. The variation occurs almost regardless of the level of lifetime income and is too wide to be explained by things like different educational levels or different rates of return on assets. A second clue is the widespread dissatisfaction expressed with the social security earnings test, which reduces benefits with additional earnings above an exempt amount. Many individuals seem to regard the foregone benefits as a tax and do not appear to value at all the fact that later benefits will be increased. A third clue is the high take-up rate of social security benefits when individuals first become eligible for them, even though the actuarial value of benefits is frequently increased by delaying benefits to a later age.

In this paper, we estimate a retirement model which allows for heterogeneous time

preferences. We show that such a model can indeed generate a bunching of retirement at age 62 which is nearly as large as the observed magnitude. Some of this bunching is undoubtedly due to the fact that individuals with high time preferences will not have saved and hence will not have the assets to finance retirement before 62. Perhaps a more important reason is that individuals with high time preferences will heavily discount the future benefit increases that delaying retirement fosters, and hence will see the foregone benefits as a reduction in net compensation.

Simulations indicate that if the early retirement age is increased, perhaps three-fifths of the bunching of retirements at age 62 will shift to the new early retirement age. Since the bunching amounts to about 8 percent of the population, an increased retirement age from three-fifths of them will have a substantial effect on the social security system and its finances that should not be overlooked.

References

- Burtless, Gary. 1999. "An Economic View of Retirement". In Henry J. Aaron, Editor. *Behavioral Dimensions of Retirement Economics*. Washington, D.C.: Brookings, 7-42
- Burtless, Gary and Rober Moffitt. 1984. "The Effect of Social Security Benefits on the Labor Supply of the Aged". In H. Aaron and G. Burtless, editors. *Retirement and Economic Behavior*. Brookings: Washington, D.C..
- Coile, Courtney, Peter Diamond, Jonathan Gruber and Alain Jousten. 1999. "Delays in Claiming Social Security Benefits". NBER Working Paper 7318.
- Coile, Courtney and Jonathan Gruber. 2000. "Social Security and Retirement". NBER Working Paper 7830.
- Congressional Budget Office. 1999. "Raising the Earliest Eligibility Age for Social Security Benefits". Washington, D.C.
- Diamond, Peter A. and Jerry A. Hausman. 1984. "Individual Retirement and Savings Behavior". *Journal of Public Economics* 23: 81-114.
- Fields and Mitchell. 1994. *Retirement, Pensions and Social Security*. MIT Press.
- 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.
- Gruber, Jonathan and David A. Wise. 1999. Editors. *Social Security and Retirement Around the World*. Chicago: University of Chicago Press for NBER.
- 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. 1984. "Modeling the Retirement Process for Policy Evaluation and Research". *Monthly Labor Review* 107(7): 26-33.
- _____. 1985. "The 1983 Social Security Reforms and Labor Supply Adjustments of Older Individuals in the Long Run". *Journal of Labor Economics*, Vol. 3, No. 2 (April, 1985) pp. 237-253.
- _____. 1986. "A Structural Retirement Model." *Econometrica*, Vol. 54, No. 3

(May, 1986) pp. 555-584.

_____. 1991. "Changing Social Security Rules For Workers Over 65: Proposed Policies And Their Effects", *Industrial And Labor Relations Review*.

_____. 1992. "The Stampede Toward Defined Contribution Pension Plans: Fact or Fiction?" *Industrial Relations* 31(2): 361-369.

_____. 1999. "Changing Pensions in Cross-Section and Panel Data: Analysis with Employer-provided Plan Descriptions". In *Proceeding, National Tax Association, 1998*: 371-377.

_____. 2000a. "Employer-provided Pension Data in the NLS Mature Women's Survey and in the Health and Retirement Study". In *Research in Labor Economics* 19: 215-252.

_____. 2000b. "Retirement in Dual Career Families: A Structural Model". *Journal of Labor Economics* 18(3): 503-545.

_____. 2001. "Retirement Outcomes in the Health and Retirement Study". *Social Security Bulletin* 64(2): 3-17.

_____. 2002a. "Retirement and Wealth". NBER Working Paper 8229.

_____. 2002b. "Social Security, Pensions and Retirement Behavior within the Family". NBER Working Paper 8772.

Kahn, James A. 1988. "Social Security, Liquidity, and Early Retirement". *Journal of Public Economics* 35: 97-117.

Lumsdaine, Robin L. and Olivia S. Mitchell. 1999. "New Developments in the Economic Analysis of Retirement". In Orley Ashenfelter and David Card, editors. *Handbook of Labor Economics*. Amsterdam, Holland: 3261-3308.

Packard, Michael D. and Virginia P. Reno. 1988. "A Look At Very Early Retirees". In R. Ricardo Campbell and E.P. Lazear, editors. *Issues in Contemporary Retirement*. Stanford: Hoover Institution: 243-265.

Quinn, Joseph F. 1999. "Retirement Patterns and Bridge Jobs in the 1990s". Employee Benefit Research Institute Brief Number 206.

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 J.P. Smith and R.J. Willis, Editors. *Wealth, Work and Health. Innovations in Measurement in the Social Sciences*. Ann Arbor, University of Michigan Press: 87-120.

Table 1
Reasons for Deletions from 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	-0.26	8.57
Parameters in β			
β_0	Constant	-9.85	304.21
β_a	Coefficient of Age ^a	0.076	5.21
β_h	Coefficient of Health ^d	3.92	4.57
β_c	Coefficient of Cohort ^{b,d} (Year of Birth)	-0.01	0.07
Parameters in δ			
δ_0	Constant	-1.78	3.53
δ_a	Coefficient of Age ^c	0.46	3.10
σ_ϵ	Standard Deviation of ϵ^d	5.79	7.44
	q value		46.92
	Number of observations		2305

Several variables are differenced from their approximate means in the sample in order to facilitate estimation. They are:

^a The actual variable is age - 62.

^b The actual variable is cohort - 1936.

^c The actual variable is age - 65.

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

Appendix Table 1

Hourly Wage Rates for Full-Time Workers
Using Median Regression

	Coefficient	Absolute t-statistic
Constant	-5.936	291.90
Tenure	0.018	14.23
Tenure squared (times 10 ⁻²)	-0.021	-6.65
Experience	0.005	3.40
Experience squared (times 10 ⁻²)	-0.006	1.99
Education:		
Less than High School	-0.130	12.80
Some College	0.139	13.94
College Graduate	0.331	24.11
Advanced Degree	0.431	32.26
Race:		
Black	-0.059	5.51
Hispanic	-0.067	4.61
Married	0.086	9.87
Fair or Poor Health	-0.036	3.56
Self Employed	-0.071	5.56
Industry:		
Manufacturing	0.029	3.15
Government	0.056	3.46
Occupation:		
Management/Professional	0.227	22.91
White Collar	0.033	3.61
Pension coverage	0.244	27.30
Number of Observations	25,416	
Pseudo R ²	0.18	

Dependent variable is log of the hourly wage rate.

Appendix Table 2

Hourly Wage Rates for Part-Time Workers Relative to Full-Time Workers Using Median Regression

	Coefficient	Absolute t-statistic
Constant	0.033	0.33
Experience	0.010	1.56
Experience squared (times 10 ⁻²)	-0.019	1.72
Education:		
Less than High School	-0.175	3.78
Some College	0.081	1.85
College Graduate	0.095	1.64
Advanced Degree	0.437	7.49
Race:		
Black	-0.027	0.52
Hispanic	-0.049	0.68
Married	0.028	0.68
Fair or Poor Health	0.010	0.22
Number of Observations		1766
Pseudo R ²		0.03

The dependent variable is the log of the ratio of the part-time hourly wage to the full-time hourly wage with the effects of experience and tenure removed. The regression excludes part-time observations in the same job as the full-time observations.

Figure 1a
Observed Retirement Patterns in the HRS




Figure 1b
Pseudo-Retirements from Full-Time Work

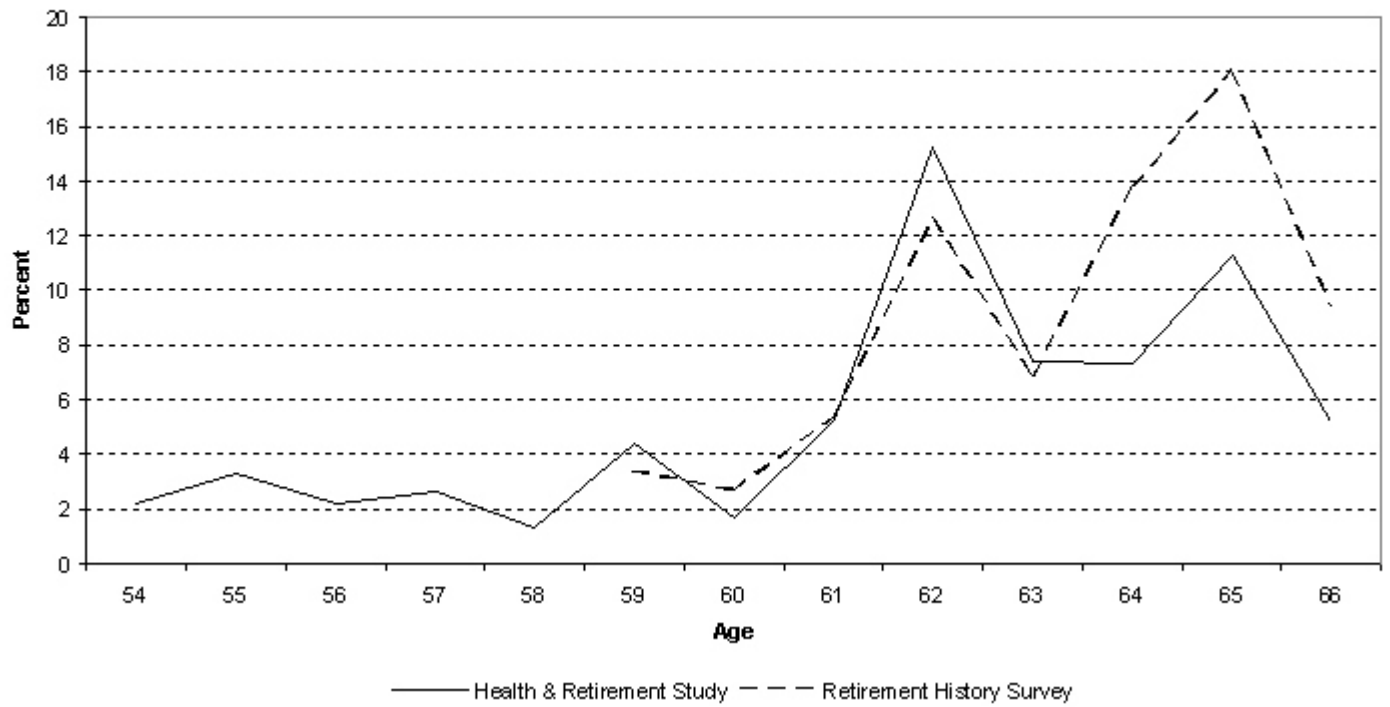


Figure 2
Observed Retirement Patterns for Respondents in Poor Health

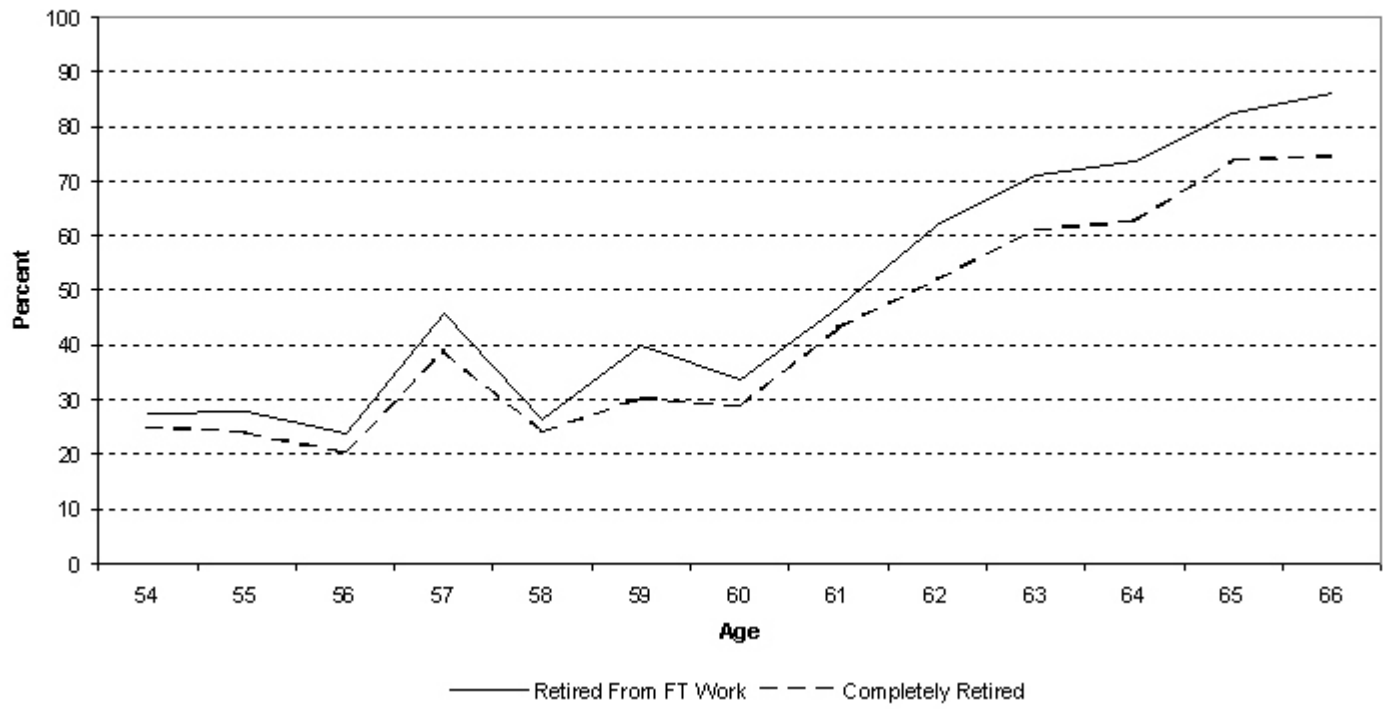


Figure 3
Distribution of Time Preference Rates




Figure 4a
Retirement Patterns Simulated from the Model

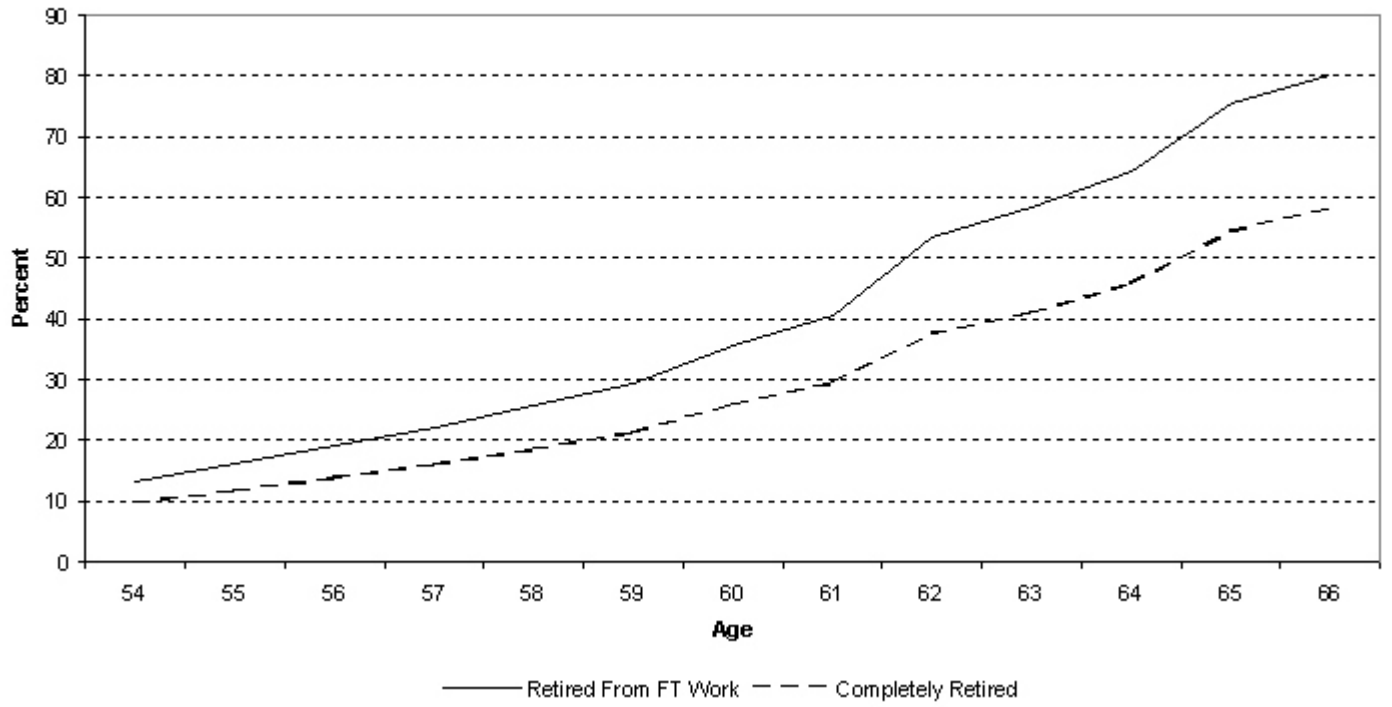


Figure 4b
Retirements from Full-Time Work




Figure 5
Effects of Changing the Early Retirement Age
on Retirements from Full-Time Work

