NBER WORKING PAPER SERIES

A STRUCTURAL RETIREMENT MODEL

Alan L. Gustman

Thomas L. Steinmeier

Working Paper No. 1237

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 1983

This project is partially funded under contract number J-9-P-1-0077 from Pension and Welfare Benefit Programs, U.S. Department of Labor. Roger Gordon, Richard Ippolito, Olivia Mitchell, Robert Moffitt, David Stapleton, Steven Venti and participants in a seminar at NBER all provided helpful comments. Any remaining errors are, of course, our own responsibility. The research reported here is part of the NBER's research program in Labor Studies and project in Government Budget. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research or the Department of Labor.

NBER Working Paper #1237
November 1983

A Structural Retirement Model

ABSTRACT

The model analyzed here constrains most work on the main job to be full time. Partial retirement requires a job change and a wage reduction. Estimates of utility function parameters and their distributions incorporate information on age of leaving the main job and of full retirement. These estimates determine the slope at different ages and the convexity of within period indifference curves between compensation and leisure. Even though age specific dummy variables are not used, the model closely tracks retirement behavior. Policy analysis based on earlier models with simpler structures is shown to be misleading.

Alan L. Gustman
Department of Economics
Dartmouth College
Hanover, NH 03755
603-646-2641

Thomas L. Steinmeier
Department of Economics
Texas Tech University
Lubbock, TX 79409
806-742-2203

I. Introduction

Analysis of retirement behavior has progressed rapidly from the estimation of reduced form equations which were linked only loosely, if at all, to any theoretical framework, to the structural analysis presented in the path-breaking study by Roger Gordon and Alan Blinder (1980). Much of this progress reflects the increasing feasibility of estimating structural models directly using maximum likelihood techniques. However, the usefulness of this approach depends, as always in econometrics, on the accurary of the specification of the underlying model. It is our contention that there remain serious problems in the specification of these models which must be remedied before we have an estimated structural model accurately depicting the retirement decision. The purpose of this paper is to estimate a more appropriately specified model and to assess how the improved specification affects the estimates.

Previous work has had problems in two broad dimensions of specification which will be addressed in this paper.[1] First, available structural retirement models misrepresent the choices facing older workers. Empirical evidence from the micro data sets indicates that most older workers are unable to reduce their work effort below full-time without leaving the job they held in their prime working years. Further, workers who do partially retire and reduce their work effort below full-time, either in the same job or in a different one, generally must accept a reduced hourly compensation rate in order to do so.[2]

---

[1] For a list of additional questions about available studies, see Henry Aaron (1982). To the extent that the questions Aaron raises remain unanswered, especially those questions pertaining to savings, consumption and bequest behavior by older individuals, a considerable amount of fundamental work remains to be done.

[2] For a detailed discussion of these points, see Alan Gustman and Thomas Steinmeier (1982 and 1983a).

The second broad dimension of specification to be addressed in this paper involves the presumed manner by which individuals reach retirement decisions given the opportunities available to them. Clearly, the choice of retirement date is one of the more fundamental aspects of the life cycle decision regarding lifetime labor supply. However, some of the previous work employs utility function specifications which are not appropriate for a life cycle analysis.[1] Other work recognizes the life cycle nature of the retirement decision but makes drastic simplifications in the model in order to facilitate solution and estimation.[2] None of the previous structural retirement studies is based on a complete solution to a life cycle model for hours of work at various ages.

As would be expected, misspecifying the opportunity set and the manner in which individuals reach retirement decisions leads to biased parameter estimates. But in addition, the use of an overly simplified specification has led previous investigators to overlook important information pertaining to retirement--information which helps to explain the choices individuals make. In particular, whether or not an individual partially retires at all, and if so for how long, provides important clues about the nature of his preferences

---

[1]Included in this group is the work of Anthony Zabalza, Christopher Pissarides and M. Barton (1980), who do, however, use a model which treats partial retirement in a more realistic manner than many of the other papers.

[2]For example, Gordon and Blinder consider a three-period model in which the three periods are "past," "present," and "future." Assuming full-time work in the past, they derive two formulae for the reservation wage in the present, depending on whether or not the individual works in the future. Whether or not the individual works in the future, however, depends on the very parameters they are trying to estimate. Hence, Gordon and Blinder face the chicken-and-egg problem: the proper reservation wage equation to use for a particular individual depends on the parameter values, but in order to estimate the parameter values, they must already know which equation applies to each individual. In their empirical estimation, they sidestep this problem by using a reservation wage equation which is a compromise of the two derived reservation wage equations, without trying to decide which one is right for a particular person.

and the way they are changing over time. But the specifications employed in most previous studies are unable to use these clues. Gordon and Blinder, in their innovative work, use a model in which everyone would partially retire at some time, since in their model partial retirement entails no penalty in terms of a lower hourly compensation.[1] Olivia Mitchell and Gary Fields (1983) assume that all work takes place on a full-time basis. No partial retirement is allowed in their model. Gary Burtless and Robert Moffitt (1983) allow either partial or full retirement in their model, but not both. They assume that an individual who does partially retire works on a part-time basis at a constant number of hours for the rest of his life.[2]

To meet the objective of this paper, we construct and estimate a structural, life cycle retirement model which remedies these deficiencies. The model specification assumes that individuals can partially retire and reduce their work effort below full-time, but that in doing so, most will incur a reduction in their hourly compensation rate. This specification allows the estimation procedure to use information on the timing of transitions among full-time work, partial retirement, and full retirement to gauge the nature of preferences regarding consumption and leisure and the way that these preferences may be changing as individuals become older. Retirement decisions in the model are based on a complete solution to the life cycle labor supply

---

[1]In this context, it is of interest to note a comment contained in a recent paper by Alan Blinder (1982, p. 54 fn.): "In the 'if-I-had-it-to-do-over-again' department, one thing I regret is assuming ... that the individual can work as few hours as he pleases. It is a convenient assumption, since it enabled us to translate the utility function into a reservation wage function and then simply compare the reservation wage to the market wage. However, I have grave doubts about its realism.

[2]In contrast to the assumption adopted by Burtless and Moffitt, Alan Gustman and Thomas Steinmeier (1983b) report that the average duration of partial retirement is about three years.

problem, with optimal control techniques being utilized to find this solution.

The remainder of the paper may be outlined as follows. In Section II the life cycle model is developed along the lines suggested above, and a method for solving this model is sketched. Section III introduces the stochastic specification of the model and derives the likelihood function used in the estimation. Section IV discusses how the retirement sequence and the budget constraint for various individuals are calculated from the information in the Retirement History Survey, which is the principal data set used in this study. Parameter estimates are presented in Section V for a sample of older white males who are not self-employed, and these estimates are compared to analogous estimates obtained when the model is constrained as it has been by previous investigators. In Section VI the model is used to track retirement probabilities at different ages. This model does not, as do a number of other, incorporate dummy variables for ages 62 and 65 in the utility function, and consequently whether or not it produces the observed peaks of retirement activity at those ages is one test of whether it generates realistic retirement behavior. In this section, the model is also applied to investigate the sensitivity of retirement behavior to long-term wage growth and to the incentives provided by Social Security benefits, pension benefits, and mandatory retirement. The results are compared to those obtained from a model which ignores the possibility of partial retirement work at a reduced compensation rate. The final section discusses further policy applications and presents conclusions.

II. A Life Cycle Model of Retirement Behavior.

The model employed in this paper is a dynamic life cycle model which reflects the incentive effects of pension plans and Social Security, and which incorporates the fact that hours can be reduced below full-time, but only at the cost of a reduced compensation rate. The natural framework for analyzing models of this type is the theory of optimal control.[1] In this framework, the individual is presumed to be trying to maximize a lifetime utility function of the form

$$U = \int_0^T u[C(t), L(t), t]\, dt$$

where C(t) is consumption at time t, L(t) is leisure at time t, and T is the relevant time horizon over which the maximization is done. Leisure is presumed to be measured in units such that it is constrained to range in value between 0 and 1. The utility function is maximized with respect to consumption, leisure, and work effort subject to the lifetime budget constraint

$$A_0 + \int_0^T e^{-rt}\, [Y(t) - C(t)]\, dt = 0 \tag{1}$$

where A_o is initial assets, Y(t) is compensation at time t (including any increments to the present value of Social Security and pension benefits), and r is the real interest rate. Y(t) and L(t) take on values of $Y_F(t)$ and 0, respectively, if the individual chooses full-time work; otherwise, they take on values of $Y_p[H_p(t)]$ and $1-H_p(t)$, respectively, where $H_p(t)$ is the amount of labor supplied to the partial retirement job. $Y_p[\cdot]$ is the function relating partial retirement work effort to compensation and reflects the effects of Social Security and pension rules as well as the actual wage rate on compensation.

The calculations may be reduced, along the lines suggested by Thomas

---

[1]Similar results could be obtained using a discrete time framework with fairly short time intervals, much as similar solutions are obtained from analogous systems of differential and difference equations. In the empirical work, the optimal control model is implemented using one-year time periods.

MaCurdy (1981), to the problem of choosing C(t), L(t), Y(t), and full-time work vs. partial retirement work vs. full retirement at each moment in time so as to maximize the quantity

$$u[C(t), L(t), t] + ke^{-rt}[Y(t) - C(t)]$$

where the value of the parameter k is constant across time, where L(t) is subject to the restrictions indicated above, and where the expression $e^{-rt}[Y(t) - C(t)]$ is recognized as the discounted value of savings at time t.[1] In this problem, the individual generates utility during the current period in three ways: directly through consumption, directly through leisure, and indirectly through savings which are then converted to direct utility in other periods. The individual chooses consumption and work effort, and thereby savings, so as to maximize the utility, both direct and indirect, generated during the period. In this formulation, the parameter k has a natural interpretation as the marginal utility of discounted savings either generated or used during the period. Its value depends upon the wage offers and the shape of the utility function throughout the entire life cycle, and it is thus the vehicle through which decisions and opportunities in other periods affect the choices in the current period.

To determine the appropriate value of k to use in the maximization, note that the savings generated at each moment in time, as determined by the above maximization problem, must just satisfy the lifetime budget constraint [equation (1)] over an individual's lifetime. For a value of k that is too large, the individual will work so much and consume so little over the life cycle that the lifetime budget constraint will show a surplus, which cannot be

---

[1] For detailed discussion of the solution of the control problem, see Gustman and Steinmeier (1983c, pp. 4-12).

optimal. If the value of k is too low, the solution of the optimization problem at each moment in time will create a deficit in the lifetime budget constraint and hence will not be feasible. Only for the proper value of k will the results of the period-by-period optimization problem just satisfy the lifetime budget constraint. Unfortunately, except under special circumstances it is not possible to write down a closed-form analytic expression for k, and hence for the functions C(t) and L(t) which solve the optimal control problem. For a specific utility function, though, the values can be readily approximated using numerical techniques on a computer.

The above problem may be restated as choosing L(t), Y(t), and full-time work vs. partial retirement work vs. full retirement at each moment in time so as to maximize the expression

$$u^*[Y(t), L(t), t, k] = \sup_{C(t)} u[C(t), L(t), t] + ke^{-rt}[Y(t) - C(t)] \ .$$

This allows the problem to be analyzed using standard indifference curves. Figure 1 illustrates a set of indifference curves between income and leisure that is associated with u*.[1] The curvature of these curves is related to the elasticity of substitution between consumption and leisure in the original utility function u. The segmented line ABC represents income and leisure opportunities available to the individual if he partially retires, with the kink at B reflecting a feature of the opportunity set such as effect of an earnings test associated with Social Security which, even after recomputation, is not actuarially fair. The point D represents income available if the individual

[1]Note that these indifference curves reflect more than just preferences, since their shape depends through the parameter k on the opportunities available over the rest of the life cycle. Also, it may be shown with a little work that the slope of these indifference curves depends on L(t) but not on Y(t), with the implication that all the indifference curves at a particular time t are vertical displacements of one another.

Figure 1

Choice of Income and Leisure at Time t.

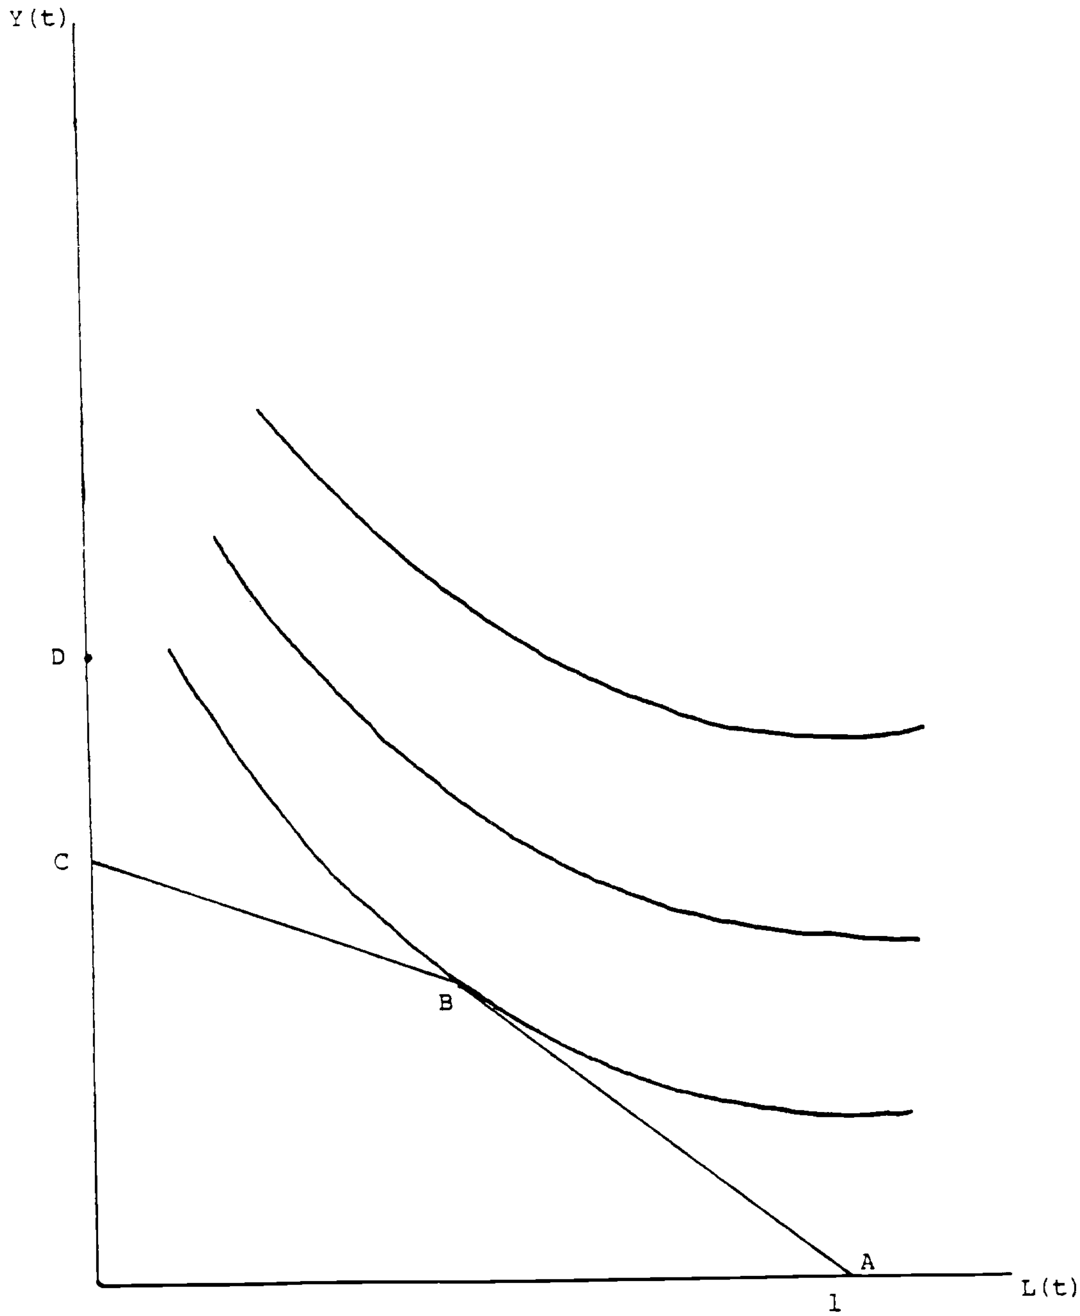

works full-time at the relatively higher compensation rate. In the diagram as illustrated, the individual would be partially retired and working just enough not to have any of his earnings subject to the earnings test.

In the empirical analysis, we use the following CES specification for the utility function:

$$\text{(2)} \qquad u[C(t), L(t), t] = \operatorname{sign}(\delta)\{[C(t)]^{\delta} + e^{\underline{X}_t\underline{\beta}+\epsilon}[L(t)]^{\delta}\}$$

where $\underline{X}_t$ is a vector of explanatory variables which affect the relative weight of leisure in the utility function at time t, $\underline{\beta}$ is the associated vector of parameters which is presumed to be constant across both time and individuals, ε is a time-invariant stochastic term affecting the relative weight of leisure for the individual, and δ (with $\delta \leq 1$) is a time-invariant stochastic term defining the within-period elasticity of substitution between consumption and leisure.[1,2] This specification of the utility function follows the specification used by Gordon and Blinder. In this specification, the within-

[1]A more general formulation of the CES utility function would raise the expression in brackets to the power ν/δ and premultiply the result by the time preference factor $\exp(-\eta t)$. In the context of the model, however, the principal indication of time preference is the rapidity with which consumption declines over time (since $\underline{X}_t$ includes age, the separate effect of time preference cannot be readily identified on the basis of labor supply behavior alone), and the principal indication of ν/δ is the degree to which consumption rises or falls when an individual retires. As is the case with most longitudinal data sets, the Retirement History Survey contains relatively poor consumption data, which in turn implies poor identification of both ν and η. Later sensitivity analysis indicates that the estimated parameter values are not very sensitive to the omission of the time preference term; similar sensitivity analysis for the exponent is precluded because the resulting nonseparability of the utility function very considerably increases the computational complexity of the estimation procedure. The omission of these parameters from the utility function should not seriously impair the usefulness of the model for the analysis of labor supply and retirement behavior, but it does suggest that a great deal of caution should be exercised in using the model to examine issues regarding either consumption or savings behavior.

[2]If the utility function were not premultiplied by sign(δ), then the
(footnote 2 continued on next page)

period elasticity of substitution is calculated as $\sigma = 1/(1 - \delta)$.

If the individual has a δ relatively close to unity, the indifference curves in Figure 1 will be nearly linear, while if δ is negative, the indifference curves will be more sharply curved. Over time, one would expect that the effect of the time-dependent variables in $\underline{X}_t$ would be such that the individual places an increasing weight on leisure as he grows older, which would be reflected in Figure 1 as an increasingly steep set of indifference curves. If the indifference curves are nearly linear, this will at some point in time induce the individual to quit full-time work at a point like D and to retire fully at A. On the other hand, if the indifference curves are sharply curved, the individual is likely to find it to his advantage to spend some time partially retired along the segmented line ABC after he quits full-time work. Over time, as the indifference curves become still steeper, he reduces his partial retirement work effort until finally he fully retires at point A.

III. The Estimation Technique

The estimation procedure uses two basic pieces of information to infer the nature of the distribution of δ. The first is the fraction of individuals who partially retire to the fraction who proceed directly from full-time work to full retirement. If the distribution of δ is bunched around unity, the model predicts that almost no one should ever partially retire, whereas if

---

(footnote 2 continued from previous page)
marginal utility of consumption would be negative for negative values of δ. This would restrict economically meaningful values of δ to lie within the range of 0 to 1 and would preclude elasticities of substitution less than one. To avoid an interpretation of negative utility when the factor sign(δ) is included in the specification, one could include a large positive constant term which is added to u for individuals with a negative value of δ. Such an additive constant in the utility function will have no influence on the optimal time paths of C(t), L(t), or the choice of jobs.

the bulk of the distribution represents low values of δ, the model predicts that almost no one would fully retire directly after full-time work without first spending time in partial retirement. The second piece of information relevant to establishing the nature of the distribution of δ is the general relationship between retirement age and the overall level of wages over the life cycle. If the distribution of δ is clustered around unity, then consumption and leisure are highly substitutable, and the higher the compensation path the later the individual should wish to retire. For lower average values of δ, consumption and leisure are complements, and the retirement age should be negatively related to the overall level of the compensation path. The estimated distribution of δ will reflect both the relative number of individuals who do and do not partially retire and the observed correlation between the retirement age and the overall level of compensation over the life cycle. The estimates may not reflect either piece of information perfectly, but will in general represent a compromise between the two.

The second major issue to be resolved by the estimation procedure is how fast the coefficient of the leisure term in the utility function is increasing over time, which governs how rapidly the indifference curves illustrated in Figure 1 are becoming steeper as the individual grows older. Again, there are two important pieces of evidence. The first is the response of retirement behavior to changes in compensation such as those occurring at ages 62 and 65 for many individuals. If the indifference curves are becoming steeper quite rapidly, then an individual will be sensitive to small changes in compensation rates for only a short period of time, and any peaks in retirement activity due to age-related changes in compensation are likely to be quite small. Just the reverse is true if the indifference curves are becoming steeper only slowly. A second piece of evidence bearing on the speed at which the indifference curves are becoming steeper is the length of time that an individual who does partially retire spends in that state. If the indifference curves become steeper only slowly, then many of the individuals who do partially retire should spend a considerable amount of time partially retired. If the indifference curves are rapidly becoming steeper, then the average length of a spell of partial retirement should be short. The estimation procedure adjusts the coefficient of the time-related variables in the coefficient of the leisure term in the utility function to reflect both of these pieces of evidence, although again the actual estimates will represent a compromise between the two.

Before describing the estimation procedure, it will be helpful first to specify the nature of the stochastic distributions for δ and ε. The stochastic distribution of δ is assumed to come from the following exponential distribution:

$$f(\delta) = \gamma e^{-\gamma(1 - \delta)}$$

where γ is a positive parameter defining the distribution.[1] For the stochastic distribution of ε, which affects the steepness of the indifference curves, the normal distribution was used:

---

[1]The fact that $f(\delta)$ is a single-parameter distribution means that the median, mean, and variance are all uniquely related to one another. Our earlier work used a two-parameter distribution, namely a truncated normal with parameters μ_δ and σ_δ^2. Estimates of μ_δ tended to lie well above unity, so that the part of the distribution below unity consisted entirely of the extreme lower tail of the distribution, and the estimates of μ_δ and σ_δ^2 appeared to be rather unstable. The inability to estimate a two-parameter distribution is less surprising in light of the kind of information available to the estimation procedure in order to identify the distribution, as discussed at the beginning of this section. Substitution of the exponential distribution for the normal distribution eliminated the instability of the parameter estimates, since there was now only a single parameter to be estimated, and it also provided improved values for the likelihood function.

$$\varepsilon|\delta \sim N(\rho\delta,\ \sigma_\varepsilon^2)$$

In this distribution, the parameter σ_ε^2 describes the conditional variance of individual preferences regarding leisure, and the parameter ρ provides a means by which the two stochastic terms in the model may be correlated.

The estimation procedure used in this study is maximum likelihood, which requires the calculation of the probability that each individual would have chosen the retirement behavior which was in fact observed, conditional on the opportunities that were available to him. To clarify how this probability was calculated for each individual, consider a set of values for the parameters of the model. For convenience, let this set of values be denoted by the vector $\underline{\theta}$, which includes δ, σ_ε^2, ρ, and $\underline{\beta}$. For a given set of parameter values, and for the wage opportunities available to the individual, the observed retirement behavior of the individual places limits on the values of the stochastic variables δ and ε. Panel a of Figure 2, which pertains to an individual whose retirement behavior is observed at only a single point in time, illustrates the basic calculation. In this figure, Ω_F defines a region of values for δ and ε for which the individual would chose to continue working full-time. Ω_P and Ω_R define analogous regions for partial retirement and for full retirement.[1] The region Ω_F corresponding to working full-time is an area with (algebraically) low values of ε. In the utility function, a low value of ε implies that the individual derives relatively little utility from leisure, and he is likely to be working full-time in order to gain the income necessary to derive utility from consumption goods. In contrast,

[1]Finding the boundaries for these regions is the difficult part of the calculation. The computations are considerably eased by the assumed separability of the utility function, but they are nevertheless quite messy. For a more detailed description of the calculations, see Appendix A.

Figure 2. Regions of the Stochastic Variables Corresponding to Particular Retirement Sequences.

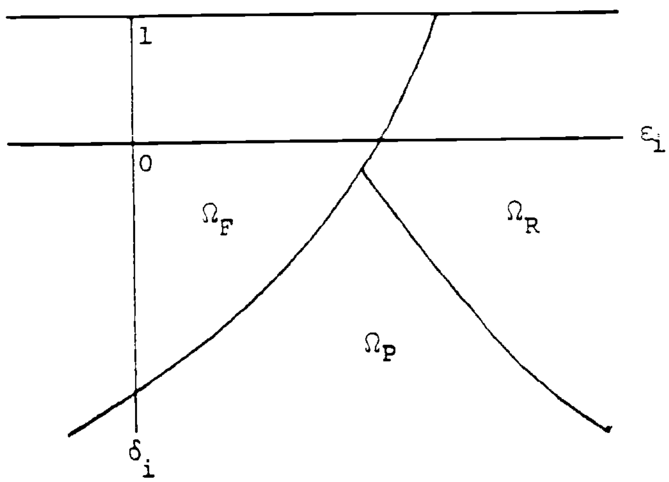

(a) Single Year

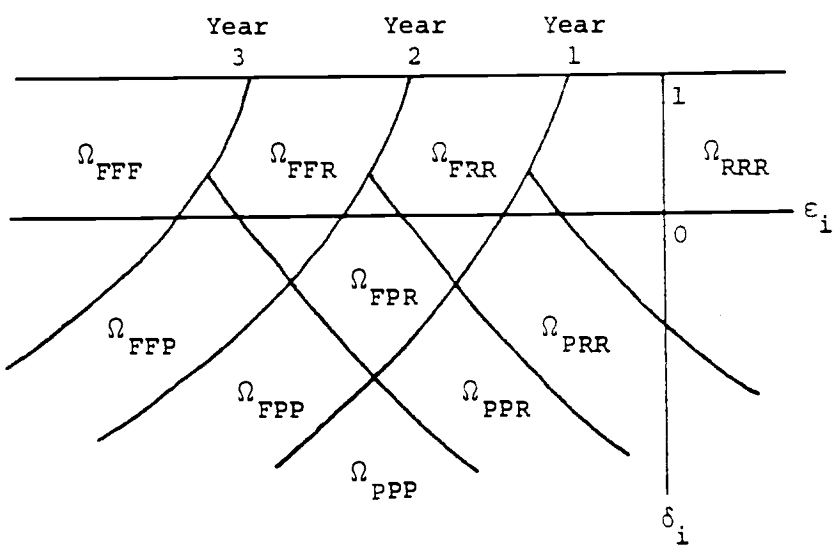

(b) Multiple Years

in the region Ω_R corresponding to full retirement, the values of ε are algebraically high, indicating that the individual places great value on leisure and is likely to be retired. The precise location of the boundary depends on the individual's current wages relative to his opportunities over the course of his life cycle. It may be located in the zone of either positive or negative ε's. For the region Ω_p corresponding to partial retirement, the most noticeable characteristic is that none of the region extends above a certain critical value of δ. In Figure 1, a high value of δ is associated with indifference curves with relatively little curvature, and under such circumstances the individual is likely to be found either working full-time at point D or fully retired at point A. For lower values of δ, which are associated with more sharply curved indifference curves, the individual will be partially retired for some values of ε. The lower the value of δ, in general, the wider will be the range of ε for which the individual will be partially retired at a particular time t.

Similar regions can be defined for an individual who is observed in several different years. Figure 2(b) illustrates a typical case involving observations in three different years. In this diagram, the right-hand inverted Y separates the (δ, ε) plane into three parts according to whether the individual would be fully retired, partially retired, or not retired at all in Year 1. These boundaries are derived in exactly the same fashion as the boundaries were derived for the case with only one observation. Combinations to the right of the inverted Y would lead to the individual being fully retired in Year 1, while combinations below or to the left would lead to the individual being partially retired or working full-time, respectively, in that year.

By Year 2 (presumed to be later), the effect of the time-dependent explanatory variables in $\underline{X}_t$ in the utility function will cause the individual to place a relatively greater weight on leisure as he grows older. In general, a relatively greater weight on leisure in the utility function will serve to shift the inverted Y, which defines the boundaries of the three regions, to the left. The region for which the individual will still be working full-time has shrunk, and the region for which the individual will be fully retired has expanded. The same process repeats itself again in Year 3, with the boundaries among full retirement, partial retirement, and full-time work being defined by the left-hand inverted Y in the diagram. The three inverted Y's defined in this manner divide the (δ, ε) plane into regions corresponding to various retirement sequences. For instance, the region Ω_{FPR} defines a region of values for δ and ε for which the individual would be working full-time in Year 1, partially retired in Year 2, and fully retired in Year 3.

To calculate the value of the probability that individual i would have chosen the observed retirement sequence S_i, first find the region $\Omega_{S_i}(\underline{\beta})$ which defines the combinations of δ and ε which would have caused the individual to have followed the observed sequence S_i.[1] Note that the boundaries of this region depend explicitly on the parameter vector $\underline{\beta}$, since these parameters in part determine how the individual will behave when confronted with a given time path of compensation. The probability that the stochastic variables would have taken on values in the region $\Omega_{S_i}(\underline{\beta})$, and hence would have generated the sequence S_i, is

[1]The retirement sequences contain two important pieces of information: the timing of the individual's departure from full-time work and the timing of his entry into full retirement. Partial retirement is implied if the timing of these two events is not the same. It is the fact that this estimation procedure uses these two pieces of information, whereas previous studies have used at most one, that permits, and indeed requires, the use of two stochastic variables in the specification.

$$Pr_i(S_i;\ \underline{\theta}) = \iint_{\Omega_{S_i}(\underline{\beta})} f(\delta,\ \varepsilon;\ \gamma,\ \sigma_\varepsilon^2,\ \rho)\ d\varepsilon\ d\delta$$

where $f(\delta,\ \varepsilon;\ \gamma,\ \sigma_\varepsilon^2,\ \rho)$ is the joint probability density function of the stochastic variables δ and ε. The value of this density function for specific values of δ and ε depends on the parameters γ, σ_ε^2 and ρ, as noted explicitly in the function.

The likelihood function is calculated as the product of the probabilities in the sample:

$$\ell(\underline{\theta}) = \prod_{i=1}^{N} Pr_i(S_i;\ \underline{\theta})$$

where N is the number of individuals in the sample. Maximum likelihood estimates of the parameters of the model are found by maximizing this likelihood function with respect to the parameters.

IV. Empirical Specification.[1]

The data set used in this study is the Retirement History Survey (RHS), a random longitudinal sample of approximately 11,000 households. The RHS sample consists of households whose heads were between the ages of 58 and 63 in 1969. Detailed questionnaires were administered to these households every two years from 1969 through 1979. This study uses the survey results through 1975, which were the latest results available at the time the study was begun. Due to the complexity of the estimation procedure, the sample actually used in the estimates and simulations in this paper is formed by taking every tenth household from the RHS. The sample is further restricted to white males who were not self-employed when working full-time; the effects of differences in constraints, opportunity sets, and preferences among males vs. females and blacks vs. whites is important and complicated enough to require separate

---

[1]Further details of the empirical specification are contained in Gustman and Steinmeier (1983c).

analysis. Observations are dropped if critical information is missing, although a major effort is made to impute missing information if at all possible.[1] The final sample consists of 494 observations.

In view of the complexity of estimating an optimal control model, a parsimonious choice is made for the set of explanatory variables to be included in the vector $\underline{X}_t$. The explanatory variables used in the empirical analysis include a constant, age (X_1), a dummy variable equal to unity if the individual has previously experienced a long-term health problem (X_2), and vintage (X_3), expressed as the last two digits of the year of birth. The coefficients associated with these variables are β_0, β_1, β_2, and β_3. For positive values of β_1 and β_2 (the coefficients of age and the health dummy variable, respectively), the utility function of an individual places an increasing emphasis on leisure over time, causing the within-period indifference curves to become increasingly steep. It is this increasing weight on leisure which eventually causes retirement in the model.

The retirement sequences used in the empirical analysis are formed on the basis of a question which asked the respondents whether they considered themselves to be completely retired, partially retired, or not retired at all.[2]

---

[1]Missing information of one of two types was the primary reason that observations were dropped; either there was insufficient information to impute partial retirement wages, or there was not at least one observation of full-time wage. At least one actual observation for a full-time wage is required because an important determinant of behavior in this model is the ratio of full-time compensation to partial retirement compensation, and an actual observation on a full-time wage provides a general indication as to whether this ratio is likely to be high or low for the individual. This requirement does not, however, cause observations to be dropped simply because individuals were already partially or fully retired when the survey began. Normally, the required information is provided in answer to questions about the individual's last job.

[2]Distinguishing between full-time work and partial retirement work on the basis of self-response eliminates the need for arbitrary distinctions as to the individual's retirement status based on hours of work per week, weeks of (footnote 2 continued on next page)

This question was asked in every survey year, and since it is a relatively innocuous question at the beginning of a major section of the survey, it was almost always answered whenever a survey was administered. For most individuals, the retirement sequence consists of a set of four observations on retirement status (one observation every two years from 1969 to 1975). In cases where the individual dropped out of the sample either because of death or refusal to answer the questionnaire, the retirement sequence includes fewer than four observations. The estimation procedure described in the previous section, however, can be used for any number of observations on retirement status, and hence it is possible to use whatever observations on retirement status are available for each individual in the data set.[1]

A problem arises with what may be called "reverse" sequences, which occur when an individual reported himself to be not retired at all after previously reporting himself to be partially or fully retired, or when he reported himself to be partially retired after previously reporting himself to be fully retired.[2] The problem is that the model described previously (as well as those

(footnote 2 continued from previous page) work per year, and/or sharp drops in wage rates (as in Burtless and Moffitt (1983)). It also reduces considerably the problem of missing information which would otherwise arise if data on hours, weeks, and/or wage levels are required to distinguish between the two types of work. For a comparison of the retirement statistics obtained using reported as opposed to objective measures of retirement, and an illustration of the sensitivity of parameter estimates to the definition of retirement status in the context of reduced form equations, see Alan Gustman and Thomas Steinmeier (1981 and forthcoming).

[1]In most cases, the individual's response regarding retirement status is taken at face value. However, the sample contains some individuals who said that they were partially retired even though they had not held any job for two years or more. For these individuals, the response is not taken at face value, and for purposes of the estimation procedure the individual is classified as retired.

[2]In the data set we use, less than 1 percent of the sequences in which the individual did not partially retire involve reversals, but about 20 percent of the sequences in which the individual partially retired do. For further discussion of the extent of this problem, see Gustman and Steinmeier (1983b and c).

of Gordon and Blinder and other studies) normally predicts a progression from being not retired at all to being fully retired, possibly with a spell of partial retirement in the transition. Consequently, observations involving reverse flows have a zero probability of occuring in the model and would preclude maximum likelihood estimates from being obtained. These reverse flows may represent the reactions of individuals to unexpected changes in circumstances, or they may be the result of individuals making errors in their maximization calculations. Either possibility could be incorporated into the model, but at the cost of a considerable increase in the complexity of the estimation procedure. Rather than getting involved in a still more complicated estimation procedure, we observe that the reverse parts of the sequences contribute little useful information for bounding the combinations of δ and ε which the individual is likely to have had, and for this reason during the estimation procedure we simply ignore the parts of the observed retirement sequence S_i which cause reverse sequences.[1]

The model requires two time paths for compensation, one for full-time work and one for partial retirement work. Compensation has three components which are considered in this paper: the wage offer, the increment in the total discounted value of any pension benefits accruing on the job, and the increment in the total discounted value of Social Security benefits accruing for the additional work. For the wage offers, separate equations are estimated

[1]For example, consider the sequence FRFR, where F indicates a survey in which the individual considered himself to be not retired at all and R indicates a survey in which the individual considered himself to be fully retired. The basic information in this sequence is that sometime between the first and fourth survey the indifference curves became steep enough that the individual quit full-time work and retired, and the indifference curves were not sharply enough curved that the individual partially retired. This same basic information is preserved if the retirement sequence is treated as FxxR, where x represents an observation not used in the estimation procedure.

for full-time work and for partial retirement work.[1] Wage streams over the life cycle are imputed using these estimated equations in conjunction with the individual's job history as best it can be reconstructed from the information in the survey. The overall height of the wage path is determined from whatever wage observations are available, and the profile of the path is drawn from the effects of the tenure, experience, and health variables in the wage equations, with the tenure variable inferred from the job history.[2] If there is no indication that the individual faced mandatory retirement, it is presumed that he could have continued indefinitely in his last full-time job. If the individual did indicate that he faced mandatory retirement, it is presumed that after the mandatory retirement date, he could have begun a new full-time job in the same occupation, but at a wage rate which reflected a drop in tenure to zero.

For individuals who indicated that they were eligible for pensions on full-time jobs, the pension component of compensation is imputed on the basis

---

[1]All compensation figures are deflated to 1967 levels using the index of hourly earnings taken from the Economic Report of the President, 1981, Table B-36. The estimated wage equations are thus net of the effects both of general price increases and of productivity increases. The issue of selectivity bias in the estimation of these wage equations is addressed in Gustman and Steinmeier (1983c).

[2]The use of tenure and experience as a basis for drawing the profiles results in relatively smooth profiles over time, which in turn precludes individuals in the model from switching retirement states because of highly fluctuating wage offers. For the partial retirement wage equation, tenure and experience variables are excluded in order to simplify the control solution to the life cycle model, but this exclusion does not result in any serious bias. Indeed, it might, whether convenient or not, have been necessary. In partial retirement wage equations for jobs started after 55, these variables are not significant at conventional levels. But the real problem is that for partial retirement wage equations for jobs started before age 55, the coefficients are significant, but they imply, for those with over 40 years of experience and ten years of tenure, an implausibly high rate of decline in the wage offer (over 7% per year) for each additional year of experience.

of information regarding the normal retirement age, provisions for early retirement, the early retirement age if applicable, and the amount of expected or actual pension benefits. No information is available in the RHS on the type of pension plan or on the actuarial reduction rate which was applied to benefits begun before the normal retirement age. In order to calculate the present value of benefits, we assume that each individual faced a defined benefit plan wherein the annual benefits are calculated by multiplying the years of service times some dollar amount or some proportion of final average salary, and adjusted actuarially if the individual retires before the normal retirement age:[1]

$$P = a\ b\ T$$

where P is the dollar value of annual benefits, a is an actuarial adjustment which lowers benefits if the individual begins collecting them before normal retirement age, b is a value for a given individual which is assumed to grow in proportion to an index of nominal wages until the individual quits the job and begins collecting benefits, and T is the number of years of service accrued when the individual begins to collect benefits. In a final average salary plan, b is interpreted as a measure of the final average salary times an appropriate fraction, whereas in a plan stipulating benefits as a specified dollar amount times years of service, b is interpreted as this specified amount. In either case, b would be expected to grow roughly proportionately

[1]According to a study for the Urban Institute, which is based on the BLS Level of Benefits Survey, more than 80 percent of the surveyed participants belong to a plan which either ties their benefits to final pay or in which they receive a dollar amount for each year of service (Sara Hatch et al., 1981, p. 25). This study also reports data on actuarial reduction rates which we use in the pension benefit calculations. These reduction rates vary by broad industry and occupation.

to the index of nominal wages until the individual starts collecting benefits. The value of the pension component of compensation is calculated by computing the present discounted value of pension benefits if the individual does not work another year in the job, computing the expected benefits if he does work another year, and taking the difference. The net effect of the various influences on pensions is to cause this increment in pension value to decline noticeably around the age of early retirement and to decline sharply again around the age of normal retirement, a result which is consistent with the recent work of Edward Lazear (1982), Laurence Kotlikoff and David Wise (1983), and others. There are likely to be considerable errors in using the procedure adopted here, which is essentially to use the reported data to calculate the pension component as best we can, but the procedure does capture the sharp drops at the early and normal retirement ages which are the main feature of pension accruals.

The calculation of the Social Security component of compensation closely parallels the calculation of the pension component. Thus, the Social Security component is derived by calculating the present value of expected Social Security benefits if the individual continues to work for another year and the present value of the benefits if the individual does not work another year, and then taking the difference. The calculations consider the effects of Social Security which operate through the individual's own benefits and, if married, through the spouse's benefits and her potential widow's benefits, and they incorporate two factors regarding Social Security benefits which have received recent attention in the literature. First, the effects of automatic benefit recomputation due to increased average monthly earnings, as stressed by Alan Blinder, Roger Gordon, and Donald Wise (1980), are included.[1] Second,

[1]The calculation of average monthly earnings uses the information in the Social Security earnings records in the RHS until the individual reached age (footnote 1 continued on next page)

the early retirement penalty for retirement before 65 reflects the fact that cost-of-living adjustments were based in this period on the entire primary insurance amount, a practice which, as noted by Richard Burkhauser and John Turner (1982), effectively causes the early retirement penalty amount not to be indexed. The time pattern of the Social Security component of compensation is similar to the pension component, namely, for most individuals there is a sharp drop at the early retirement age and a further sharp drop at the normal retirement age. The first drop appears to be due to the fact that the penalty for retirement before age 65 is less than actuarially fair once the effects discussed by Burkhauser and Turner are considered. The second drop reflects the fact that after age 65 the actuarial adjustments were very small during the sample period.

After the compensation paths are calculated in the manner just described, a problem arises because sixteen (out of 494) individuals in the sample were working full-time even though annual compensation in full-time work was less than the potential annual compensation would have been in partial retirement work, a situation which, not surprisingly, would not occur in the context of the model and which therefore would preclude maximum likelihood estimates from being obtained. This is very likely caused either by problems in the wage imputation procedure, particularly for partial retirement wages, or by misclassification of the individual as full-time when in fact he was partially retired. As with the reverse sequences problem discussed before, it would be possible to accommodate such observations in the context of a considerably

---

(footnote 1 continued from previous page) 55. After that the wages imputed in the procedure previously discussed are used. This is done to make sure that the calculations fairly closely reproduce the effect of additional earnings on average monthly earnings while at the same time allowing for the fact that after age 55, the Social Security earnings records increasingly fail to reflect potential earnings for those who partially or fully retire.

more complicated model, but in view of the fact that such observations seem to present very little information on either the slope or the curvature of the underlying indifference curves of the individuals, for the estimation at hand the best course of action appears to be simply to drop these observations.

V. Parameter Estimates.

The values of $\underset{\sim}{X}_t$, the compensation paths, and the retirement sequences as derived in the last section are inserted into the estimation procedure described in Section III in order to obtain values for the parameters of interest. In the course of this estimation, the relevant time horizon for the life cycle calculations is assumed to be from age 25 to age 85, with the individual assumed to be working full-time before age 55 and to be completely retired after age 72. Since the period of observation covers at most the age range from 58 to 69, these assumptions never contradict any observed retirement sequence, and they reduce the computation considerably. Further, since the utility function is separable, the only effect of a violation of either assumption would be relatively weak effects operating through the marginal utility of lifetime income.

One final difficulty remains to be resolved before estimates can be obtained. This pertains to the assumption in many previous studies, including that of Gordon and Blinder, that $L(t) = 0$ at full-time work. As can be seen in equation (2), this is inconsistent with values of δ less than zero and hence with an elasticity of substitution between consumption and leisure which is less than unity. There is no obviously "correct" choice for the value of $L(t)$ corresponding to full-time work, and the value chosen will affect the interpretation of the indifference map (i.e., a given value of $L(t)$ will represent a larger or smaller percentage of full-time work effort) and the interpretation of the parameter estimates. However, the particular value chosen should not greatly affect simulation results as long as the same definition is used in both the estimation and the simulation, except that a value close to zero is probably a poor choice. For the estimated results presented here, L(t) takes on a value of 0.46 for full-time work.[1]

Table 1 presents estimates for several variants of the model discussed in previous sections. The first column indicates the parameter estimates and the estimated standard errors for the basic model. All of the parameters except for β_3, the coefficient of vintage in the utility function, are significant at standard levels. The central parameters in the model (γ, σ_ε, and β_1) appear to be determined fairly precisely.

Of this group of parameters, perhaps the most important is β_1, which describes the effect of age in the utility function. The estimated parameter value of 0.21 implies that the coefficient of the leisure term in the utility function is growing by about 23 percent per year.[2] In turn, this implies that the indifference curves in Figure 1 are becoming steeper at the rate of

---

[1]The average workweek for those in full-time jobs was 45.31 hours, which is 54 percent of the presumed available time of 84 hours per week. It might appear that an alternative approach would be to treat the value of L(t) corresponding to full-time work as a parameter to be estimated. However, the identification of this parameter would be extremely weak; indeed, in the estimates which treat the number of hours in partial retirement work as fixed, the only identification for this parameter would come from the assumed functional forms of $f(\delta)$ and $f(\varepsilon|\delta)$. Under these circumstances, it appears preferable to choose a reasonable value for the L(t) corresponding to full-time work and to note that the interpretation of the other parameters depends to some degree on the value chosen.

[2]Since β_1 is part of an exponential expression, the percentage figure is calculated as $100(e^{0.21} - 1)$. Similar comments apply to the other percentage figures cited in this paragraph.

Table 1

Parameter Estimates.[a]

		Two Job Model with Variable Hours in Partial Retirement Work[b]	Single Job Variable Hours Model	Two Job Model with Fixed Hours in Partial Retirement Work[b]	Single Job Fixed Hours Model
β_0:	constant	-5.02 (0.10)	-25.82 (0.55)	-7.01 (0.11)	-24.33 (0.31)
β_1:	coefficient of age	0.21 (0.01)	0.36 (0.05)	0.24 (0.02)	0.39 (0.07)
β_2:	coefficient of dummy for long-term health problem	0.84 (0.10)	1.41 (0.24)	1.05 (0.13)	1.22 (0.22)
β_3:	coefficient of vintage	0.06 (0.03)	0.08 (0.05)	0.07 (0.04)	0.10 (0.04)
γ:	parameter for distribution of δ in two job models	0.61 (0.04)		0.42 (0.03	
δ:	simple parameter in single job models		-0.15 (0.07)		0.15 (0.04)
σ_ε:	standard deviation of ε	0.99 (0.02)	1.72 (0.18)	1.25 (0.02)	1.35 (0.19)
ρ:	correlation parameter in two job models	-8.78 (0.22)		-7.09 (0.13)	
σ:	calculated median elasticity of substitution between consumption and leisure	0.87	0.87	0.60	1.18
number of observations		478	561	478	490
value of log-likelihood		-1084.18	-653.75	-1111.85	-603.85

---

[a] Standard errors are in parentheses.

[b] There was a large off-diagonal element in the information matrix involving σ_ε and ρ which made the information matrix nearly singular. This created some amount of fluctuation in the estimated standard errors due to machine rounding.

23 percent per year.[1] The estimated value of β_2, which is the coefficient of the health variable, indicates that the onset of a long-term health problem causes the indifference curves to increase in slope by 132 percent, which is equivalent to the effect caused by an increase of almost four years in age. The estimated value of 0.99 for the parameter σ_ε indicates that there is a very large amount of individual variation in the preference for leisure. A single standard deviation in the error term reflecting individual preferences for leisure is associated with a 169 percent difference in the slope of the indifference curves measured at a particular age. It is this large variation in individual preferences, combined with the relatively smaller yearly effects of age on the slope of the indifference curves, which causes individuals to retire over a range of ages beginning in the mid-50's and extending into the 70's. For the parameter γ, which governs the distribution of the within-period elasticity of substitution between consumption and leisure, the estimated value of 0.61 implies a median elasticity of substitution of 0.87, which is slightly on the inelastic side of unity.

In order to investigate what might happen if partial retirement jobs are not distinguished from full-time jobs in the estimation, Column 2 of the table presents estimates for a model which considers only a single job for any individual, and in which only the states of working and retirement are distinguished.[2]

---

[1] The 23 percent per year figure should be interpreted as a linear approximation to the effect of age on the utility of leisure during the observed age range; it would probably be inappropriate to extrapolate this figure back to younger ages.

[2] With a full set of parameters, the calculation procedure for the maximum likelihood of the single job model does not converge. More precisely, it approaches a maximum asymptotically as the parameter values approach infinity. The problem appears to be that the stochastic structure with two stochastic variables is more complicated than can be estimated when only the states of working and retirement are distinguished. When the elasticity of substitution (footnote 2 continued on next page)

From the table, the most important difference between the two models is that the single job model yields a substantially higher estimated value of β_1 and therefore a substantially more rapid estimated rate at which an individual's indifference curves become steeper as he grows older.[1] This suggests that retirement is more influenced by the rapid steepening of the indifference curves and will be less influenced by the incentives provided by the budget constraint in the single job model than in the two job model, a conjecture that will be the subject of further investigation in the next section. The greater value of β_1 could also be expected to be associated with a narrower dispersion of retirement ages, but its effect in this regard is at least partly offset by the higher value of σ_ε, which indicates a more heterogenous set of preferences regarding leisure.

With regard to the elasticity of substitution, the estimates from the single job model are slightly on the inelastic side of unity and are in fact not very far from the median estimate from the two job model. Both estimates are very much lower than the results of Gordon and Blinder, who found an elasticity of about 10. This result is particularly puzzling in that the

---

(footnote 2 continued from previous page) is treated as a parameter rather than as a stochastic distribution, as is usually done when similar models are estimated by others, the maximum likelihood procedure does converge, as reported in the table. With the elasticity of substitution treated as a parameter, however, ρ and β_0 are not separately identified, and only β_0 is estimated.

[1]Since none of the models in Table 1 are nested in the statistical sense, the standard specification tests using the log-likelihood values are not applicable for determining which of these models, if any, is superior. To illustrate the difficulties, consider an individual with a retirement sequence FPR. In the two job model, he would be assigned the probability of small diamond-shaped area such as Ω_{FPR} in Figure 2. In the single job model, he would be assigned the probability for a range of values of ε, since in that model δ is not considered to be stochastic. The inappropriateness of comparing the log-likelihood values of the two models arises because the probabilities of individual observations in the two models are not comparable. The same kind of problem arises in interpreting an average log-likelihood value, since it is averaging probabilities of areas such as Ω_{FFF} with probabilities of areas such as Ω_{FPR} in Figure 2

single job model appears to be very close to the kind of model that Gordon and Blinder in fact estimated. One possible source of difference is the previously noted caveat that the interpretation of this elasticity is affected by the choice of L(t) corresponding to full-time work, but even so the Gordon-Blinder elasticity still appears large in comparison to the results of this study. In order to investigate this discrepency further, recall that Gordon and Blinder obtained their elasticity estimate from a reservation wage equation of the form

$$\ln W_R = (1 - \delta) \ln Y + \ldots$$

where W_R is the current reservation wage, Y is a measure of full lifetime income, and the coefficient of Y is the inverse of the elasticity. With the Gordon-Blinder estimates, a 10 percent increase in the entire lifetime compensation stream (and therefore Y) increases the current reservation wage by only 1 percent, which should induce later retirement because current wage offers are also higher by 10 percent. Thus an elasticity of 10 carries with it the implication that a long-term secular increase in real compensation levels should cause the average retirement age of successive cohorts to be sharply increasing, which is at variance with postwar experience. It is possible to speculate that their high estimate of the elasticity may be due to an error-in-variables problem with Y, both because of the difficulties in measuring Y from the RHS data and because the definition of full income used in their estimates does not correspond exactly with the theoretical construct due to the problem noted in footnote 2 of page 2 of this paper. It is difficult to judge whether this can account entirely for the discrepency between our results and those of Gordon and Blinder, however.

Column 3 of Table 1 presents estimates for a two job model in which hours in partial retirement work is fixed at a level below full-time effort.[1] An assumption of fixed hours in partial retirement work may be preferable to an assumption of variable hours in a model which, in order to keep the complexity of the model within reasonable limits, abstracts from possible fixed costs of employment and lower limits on work effort even in the partial retirement job. Column 4 presents estimates for a single job model with the fixed hours assumption. In this model the individual is constrained to work full-time if he works at all, and empirically the individual is treated as retired unless he was working full-time. In comparing the fixed hours models with each other and with the corresponding variable hours models of Columns 1 and 2, it is apparent that the fixed hours assumption does not greatly affect the general conclusions that were drawn from the estimates of the variable hours models. In particular, it is still the case that the single job model yields much higher estimates for β_1 than does the two job model, and the estimates from both fixed hours models are fairly close to the corresponding estimates from the variable hours models. Under the fixed hours assumptions, the elasticity estimate from the single job model is somewhat higher than the estimate from the two job model, but it is still only a small amount above unity and remains very much smaller than the Gordon-Blinder estimate for this parameter. Estimates of the remaining parameters vary by as much as 15 to 20 percent from the corresponding estimates from the variable hours models, which is a noticeable but not an overwhelmingly large difference. In general, the estimates do not appear to be extremely sensitive to whether fixed hours or variable hours are assumed.

[1]The fixed amount of work effort in partial retirement jobs is assumed to be 65 percent of full-time work effort, reflecting the fact that the average workweek for those who were partially retired was 29.61 hours, or 65 percent of the average workweek for those in full-time jobs.

The model is tested for sensitivity to several other assumptions, but it appears to be fairly robust with respect to all of the assumptions investigated.[1] Two of these tests are important enough to warrant specific mention. First, the parameter estimates of the model move hardly at all when the federal income tax is applied to the wage earnings for purposes of constructing the compensation streams. This result implies that any omitted consideration which has a primary effect of shifting the entire wage streams up or down (e.g., Social Security contributions) is unlikely to have more than a relatively minor impact on the estimated parameter values. Secondly, the introduction of a 10 percent time preference, via a term of the form $e^{-0.1t}$ which premultiplies the utility function, has only a very minor impact on the estimated results. In these results, the estimated value of β_1 increases almost exactly enough to offset completely the effect of time preference in the coefficient of the leisure term in the utility function. This result suggests that the omission of the time preference factor does not matter very much as long as the model is used to investigate issues pertaining to labor supply and retirement behavior.

VI. Implications of the Estimated Model.

The properties of the model may be illustrated by using it to simulate retirement behavior. The simulations are accomplished as follows. For each individual in the sample, the value of the vector $\underline{X}_t$ and the compensation paths $Y_F(t)$ and $Y_P[H_P(t)]$ are calculated according to the procedure described above. The individual will retire, according to the life cycle model, at some age R, the value of which depends on the stochastic variables δ and ε.

---

[1]Complete results for these sensitivity tests may be found in Gustman and Steinmeier (1983c).

Let $\omega_i(R)$ be the set of all combinations of δ and ε which imply that this individual would retire at age R. Then the probability that an individual with the characteristics indicated by $\underline{X}_t$ and facing the compensation path $Y_F(t)$ and $Y_p[H_p(t)]$ will retire at age R is found by integrating the probability density function for δ and ε over the region $\omega_i(R)$:

$$f_i(R) = \iint_{\omega_i(R)} f(\delta, \varepsilon)\, d\varepsilon\, d\delta$$

For the entire sample, the simulated percentage of individuals who retire at a given age R is found by taking the average value of $f_i(R)$ over the sample:

$$f(R) = \frac{1}{N} \sum_{i=1}^{N} f_i(R)$$

The same procedure can be used to calculate the distribution of any other statistic of interest concerning the simulated retirement decisions of the sample.

Figure 3 presents a set of results derived from simulations of the two job model with fixed partial retirement hours.[1] The first panel of this figure indicates the simulated percentages of individuals who first retire from full-time work at the indicated ages, and the second panel indicates comparable percentages of individuals who fully retire at the ages. The main feature of these results appears to be the substantial peaks in retirement activity at ages 63 and 65. The two panels also indicate the analogous statistics derived from direct observation of white male-headed households in

[1]The tables of Appendix B present a more detailed set of the results of the simulations on which Figures 3 and 4 are based. The fixed hours models were used to minimize the potential problems resulting from the fact that the model did not incorporate any possible effects caused by fixed costs of employment or by some minimum hours constraints on jobs held while partially retired.

Figure 3

Simulations with the Two Job, Fixed Hours Model.

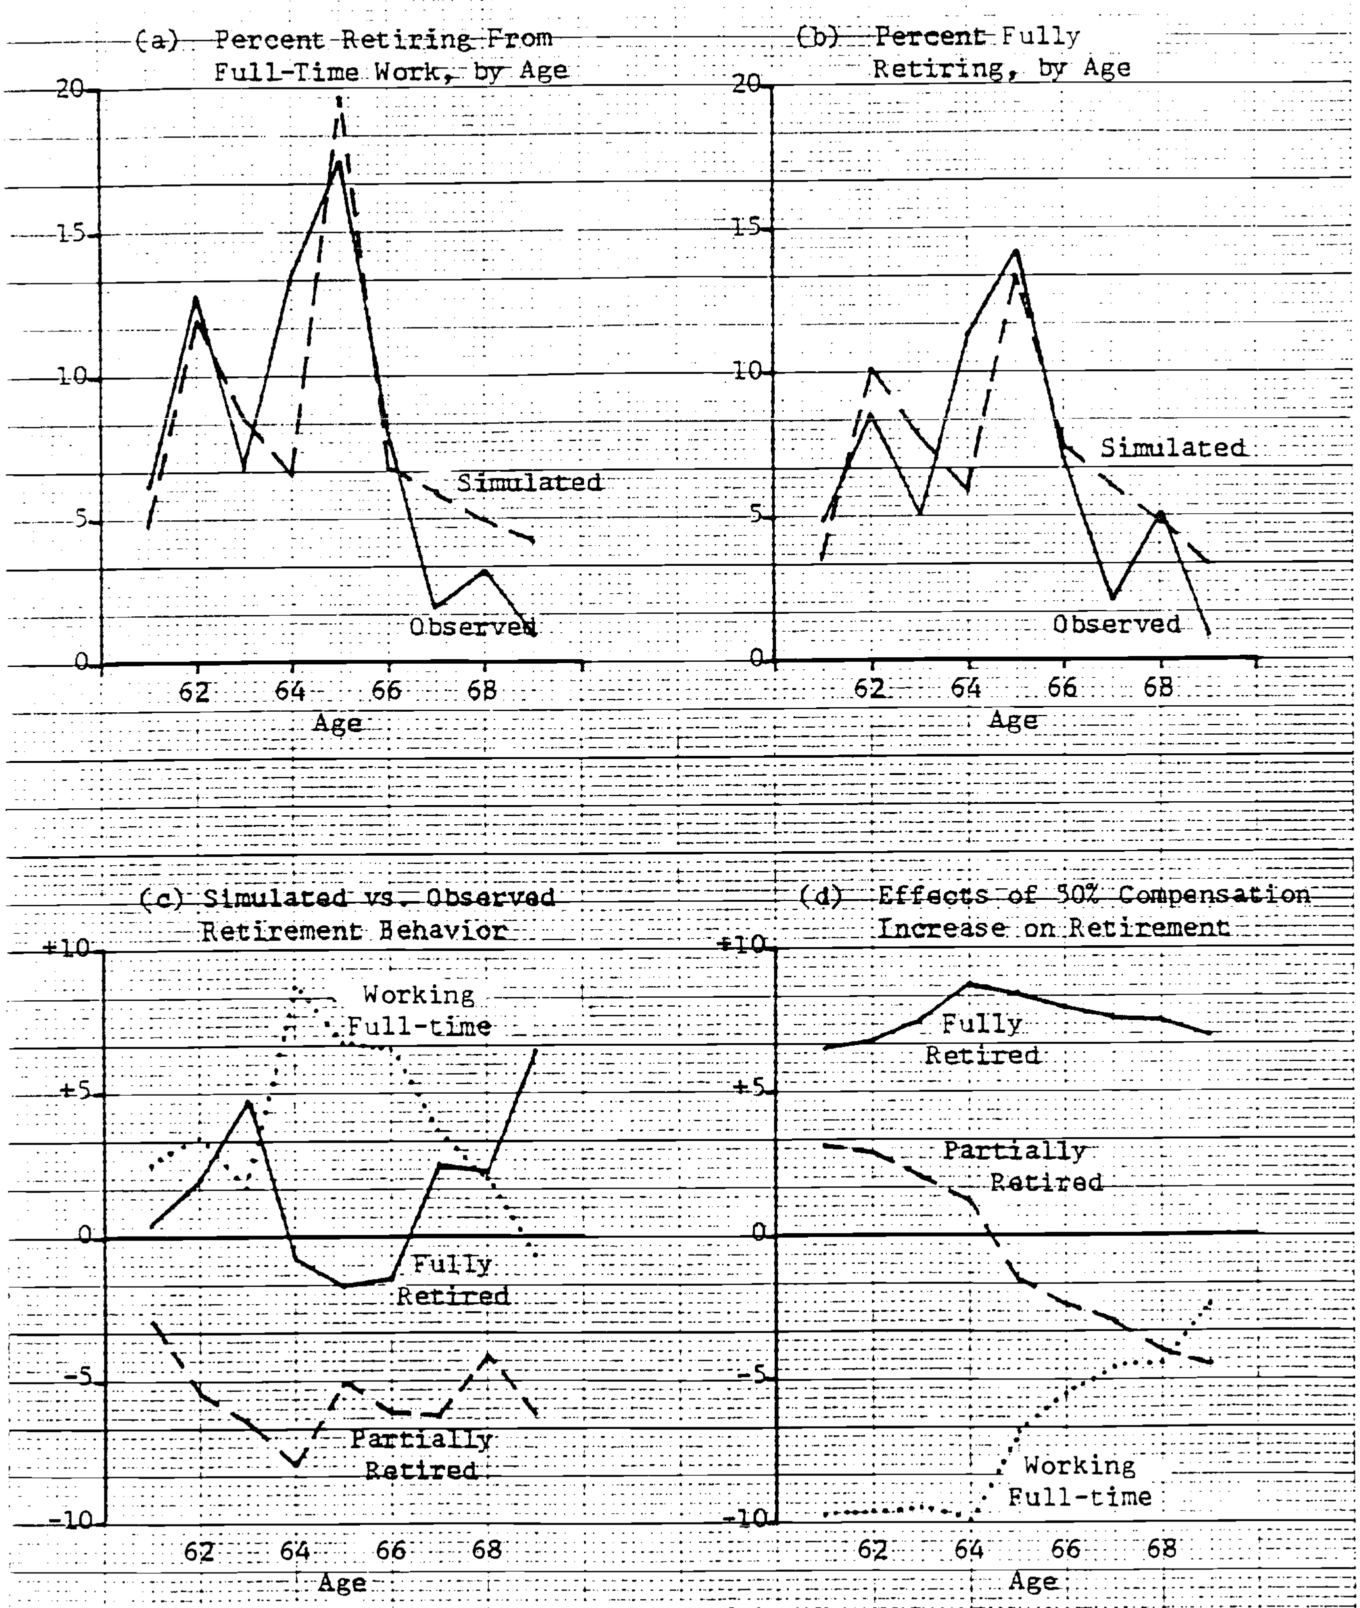

) Squares to the Inch

the RHS.[1] It is encouraging that the simulated results fairly closely parallel the observed distributions, and in particular that the two peaks at age 62 and 65 are at the same age and of approximately the same magnitude in both the simulated and observed distributions. It should be emphasized that the peaks in the simulated results are not the result of dummy age variables in the utility function which produce sudden increases in the desire for leisure, as has been the case in many previous studies.[2] Rather, the peaks in the simulated results are the result of the effects of pension programs and the Social Security program in changing the effective compensation of many individuals at these ages.

The third panel of the figure compares the simulated and actual percentages of individuals who were working full-time, partially retired, and fully retired by age. For instance, at age 65, the percentage of individuals who were fully retired was 1.7 percentage points lower in the simulation than was actually observed at this age. This result is indicated by the fact that the solid line in this panel is slightly below the axis at this age. The simulation appears to track the percentage of individuals who were fully retired to within 2 or 3 percentage points at most ages. However, it overstates the percentage who were working full time, particularly in the age

---

[1] These figures were calculated by taking the differences in the percentages of individuals who were in particular retirement states in adjacent age categories. For example, 41.1% of the sample reported themselves to be fully retired at age 64, and 55.3% reported themselves to be fully retired at age 65. The difference of 14.2% is inferred to have fully retired at age 65. The figures for the fraction of individuals who are in the various retirement states at different ages are given in Table V-1 in Gustman and Steinmeier (1981).

[2] For instance, the reservation wage equation of Gordon and Blinder includes a dummy variable for a finite jump at age 65, a dummy variable for age 62 or greater which is interacted with a Social Security wealth variable, and a dummy variable for age 65 or greater which is interacted with a pension dummy variable. All of these tend to make the reservation wage change suddenly at ages 62 and 65.

range 64-66, and it understates the percentage who were partially retired. The understatement of partial retirement in the simulation does not appear to arise because a low percentage of individuals partially retire in the simulation; approximately one third of the individuals in the simulation partially retire at some time, which is roughly the same fraction which is found in the observed behavior. Rather, it appears to be caused by a coefficient on the age variable in the utility function which implies that individuals who do partially retire spend too short a time in that state. In this regard, it may be recalled that the coefficient of the age variable plays two roles in the model. First, it governs the height of the peaks of retirement activity at ages 62 and 65; the greater this coefficient, the less sensitive retirement activity will be to the economic incentives in the compensation streams, and the lower will be the resulting peaks of retirement activity. Secondly, the coefficient governs the length of time that any individual who does partially retire spends in that state. A greater value of the coefficient will in general result in shorter periods of partial retirement. The estimated coefficient seems to have done a reasonable job in reflecting the peaks of retirement activity, but it appears that one of the trade-offs that the estimation procedure had to accept in order to do this was to predict periods of partial retirement that were somewhat too short.

The final panel of Figure 3 investigates the implications of a 50 percent increase in the entire compensation streams for both full-time and partial retirement work. This might correspond, for example, to the effects of increases in labor productivity over a period of two to four decades. The results indicate that such an increase in compensation would reduce the percentage of individuals working full-time by about 10 percentage points up through age 64 and by a smaller amount thereafter. The percentage of individuals who are fully retired increases by a comparable amount, and the effects on the percentage of partially retired individuals is mixed. In terms of the average retirement age, such an increase in real wages would cause the average age to decline by a little less than one year. This result is perfectly plausible in light of recent experience and suggests that a substantial part of the gradual decline in the retirement age in recent decades might be interpreted as an effect of the increasing levels of real income made possible by advances in productivity.

Figure 4 addresses an issue discussed in the last section, namely: what difference does it make if we use the estimates from one of the single job models rather than the estimates from the two job model? The single job model used in these simulations is the fixed hours model, in the estimation of which individuals were treated as working if they were working full-time and as retired otherwise. The first two panels of Figure 4 present information from the single job model comparable to the information in the first three panels of Figure 3 regarding the two job model. From the first panel of Figure 4, it is apparent that this model also produces peaks in retirement activity of approximately the correct magnitude at ages 62 and 65. In comparing the first panels of Figures 3 and 4, however, the overall distribution of retirement ages appears to be somewhat more concentrated for the single job model, with more of the distribution between the ages of 62 and 65 and less of the distribution before 62 or after 67. This impression is confirmed by the second panel of Figure 4, which indicates that the single job model tends to underpredict retirement by about 8 percentage points at age 62 and to overpredict it by about 5 percentage points at age 68.[1]

[1]This result, which may at first appear curious, arises because the estimation procedure is particularly sensitive to individuals who are observed (footnote 1 continued on next page)

Figure 4

Properties of the Single Job, Fixed Hours Model.

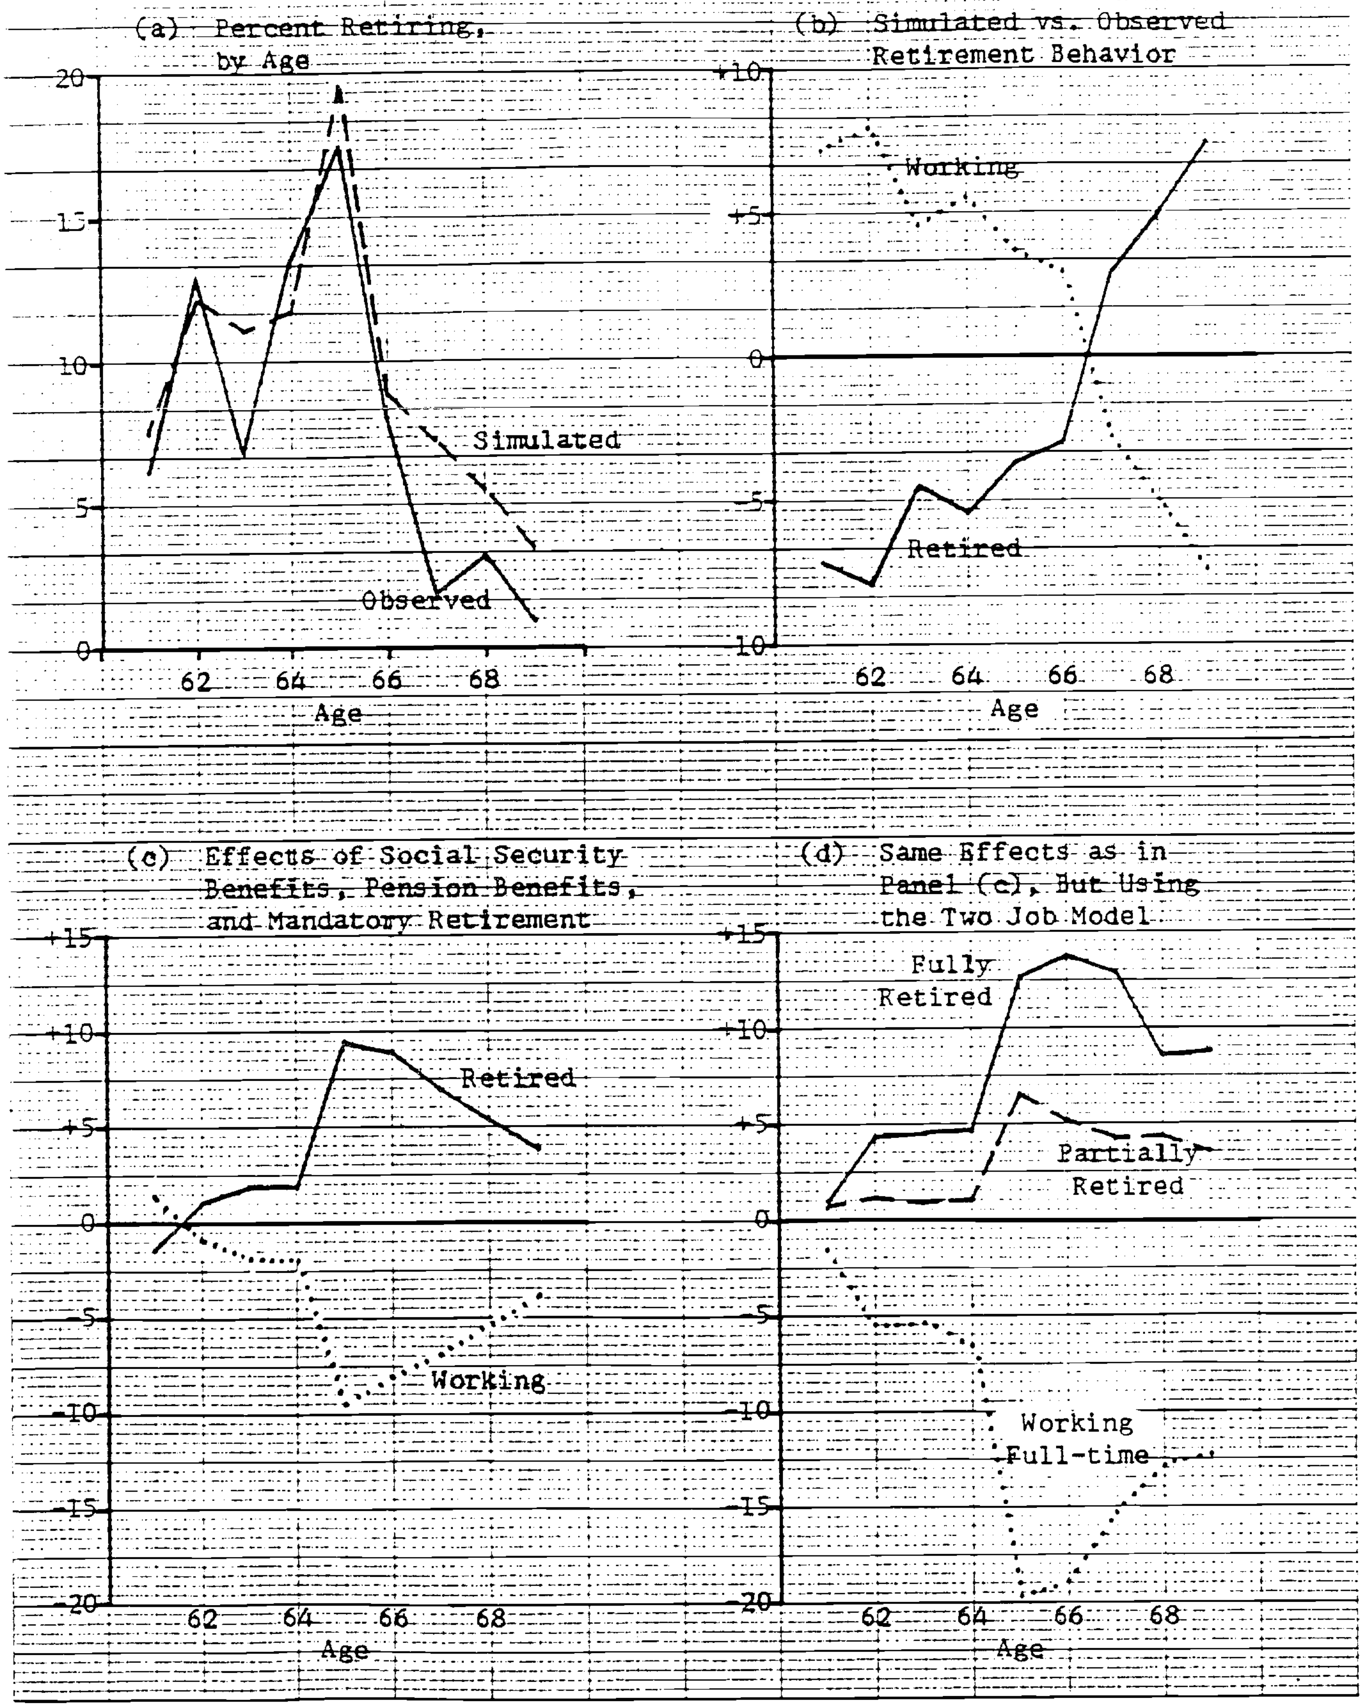

The last two panels in Figure 4 report on the combined simulated effects of Social Security, pensions, and mandatory retirement using both the single job model and using the two job model. The magnitude of these effects is measured by performing a simulation in which the compensation streams include no pension or Social Security benefits, and no mandatory retirement, and comparing the results to a simulation in which all of these effects are present in the compensation streams.[1] It was hypothesized that the higher value of β_1 estimated for the single job model should make retirement behavior less sensitive to the incentive effects provided by Social Security, pensions, and mandatory retirement, and this impression is confirmed by the simulations. For example, in the estimated two job model the combined effects of Social Security and pension benefits and mandatory retirement is to cause the percentage of individuals working full-time at age 66 to fall by 18.9 percentage points. For the estimated single job model, the comparable reduction in the percentage of individuals working at age 66 is 8.1 percentage points.[2] It

---

(footnote 1 continued from page 37) to retire at the same time that their compensation exhibits a sharp decline. (In terms of Figure 2, a sharp decline in compensation causes the boundary between working and retirement to shift more to the left than usual, making possible a large gain in the likelihood function if the estimation procedure assigns a high probability density to this area.) Since the single job, fixed-hours model treats retirement as the departure from full-time work, the association between retirement and the compensation declines at ages 62 and 65 is greater than with the other models, and as a result the estimates imply a more concentrated distribution of retirement in the 62-65 age range than occurs with the other models.

[1]For estimates of the separate effects of pensions, Social Security and mandatory retirement, see Gustman and Steinmeier, 1983c.

[2]Since the sensitivity of retirement behavior to changes in the budget constraint is related to the magnitude of β_1, the single job, variable hours model would exhibit the same kind of reduced sensitivity as the single job, fixed hour model. However, since the estimated value of the standard deviation of ε is greater in the single job variable hours model, this model would imply a less concentrated distribution of retirement ages than does the single job fixed hours model, indicating that a low sensitivity of retirement behavior to changes in the budget constraint is not necessarily associated with a small dispersion of retirement ages.

appears highly likely that the insensitivity of retirement behavior in the estimated single job model relative to the two job model would extend to most potential policy changes which would alter the incentive structure of the individual's lifetime compensation streams. Accordingly, not only is the two job model consistent with evidence cited at the outset as to the importance of minimum hours constraints in full time work and the availability of partial retirement, and thus the correct model to use, but we have now shown that applying the wrong model may have important consequences for predicted retirement behavior and for policy analysis.

VII. Conclusions.

This work has emphasized the significance of a more realistically specified life cycle model in the empirical analysis of retirement behavior than those used in previous studies. The estimated model behaves in a sensible manner, particularly in being able to approximate fairly closely the peaks in retirement activity at ages 62 and 65. Moreover, the model can be used to predict how a policy change will affect the number of people seeking full-time work and partial retirement work, which is the first step in understanding the labor market substitutions which can result from changes in retirement policies. The fact that the model can handle policies which have potentially very complicated effects on an individual's lifetime budget constraint underscores the potential usefulness of this kind of a model in analyzing the effect of alternative policies on retirement behavior.

Earlier work has established the importance of minimum hours constraints in the higher-paying full-time jobs that individuals typically hold in their primary working years, and the current work reinforces earlier suspicions that to ignore these constraints may cause serious biases in terms of understanding the way retirement behavior may respond to various policy changes. The manner in which individuals choose in their later years between full-time work at one level of compensation and partial retirement work at a relatively lower level of compensation yields potentially useful information regarding the nature of the individual's preferences and the rate at which these preferences are changing with age, but this information has been largely ignored in most of the estimation work done to date. Taking these considerations into account affects the estimates of important parameters regarding preferences, and these differences in parameter estimates, in turn, increase the estimated sensitivity of retirement behavior to many potential changes in policy. Indeed, retirement behavior appears to be approximately twice as sensitive to policy changes in the correctly specified model.

REFERENCES

Aaron, Henry J., Economic Effects of Social Security, Washington, D.C: The Brookings Institution, 1982.

Blinder, Alan S., "Private Pensions and Public Pensions: Theory and Fact," Working Paper Number 902, National Bureau of Economic Research, June, 1982.

Blinder Alan S., Gordon, Roger and Wise, Donald, "Reconsidering the Work Disincentive Effects of Social Security," National Tax Journal, December, 1980, 33, 431-442.

Burkhauser, Richard V. and Turner, John, "Can Twenty Five Million Americans Be Wrong? - A response to Blinder, Gordon and Wise," National Tax Journal, December, 1981, 34, 467-472.

Burtless, Gary and Moffitt, Robert A., "The Effect of Social Security on Labor Supply of the Aged: The joint choice of Retirement Date and Post Retirement Hours of Work," mimeo, February, 1983.

Gordon, Rogert H. and Blinder, Alan S., "Market Wages, Reservation Wages and Retirement," Journal of Public Economics, 1980, 14, 277-308.

Gustman, Alan L. and Steinmeier, Thomas L., "Minimum Hours Constraints and Retirement Behavior," in Contemporary Policy Issues, a Supplement to Economic Inquiry, April 1983(a), 77-91.

Gustman, Alan L. and Steinmeier, Thomas L., "Partial Retirement and the Analysis of Retirement Behavior," Industrial and Labor Relations Review, (forthcoming).

Gustman, Alan L. and Steinmeier, Thomas L., "Partial Retirement and Wage Profiles for Older Workers," Working Paper No. 1000, National Bureau of Economic Research, October, 1982.

Gustman, Alan L. and Steinmeier, Thomas L., "Retirement Flows," Working Paper No. 1069, National Bureau of Economic Research, January, 1983b.

Gustman, Alan L., and Steinmeier, Thomas L., "Some Theoretical and Empirical Aspects of the Analysis of Retirement Behavior," report submitted to Department of Labor, February, 1981.

Gustman, Alan L., and Steinmeier, Thomas L., "Structural Retirement Models," Final Report To the U.S. Department of Labor, April, 1983c.

Hatch, Sara P. et al., "Financial Retirement Incentives in Private Pension Plans," Final Report by Urban Institute to the Department of Labor, September, 1981.

Kotlikoff, Lawrence and Wise, David, "Labor Compensation and The Structure of Private Pension Plans: Evidence for Contractual vs. Spot Labor Markets" mimeo, 1983.

Lazear, Edward, "Pensions As Severence Pay," Working Paper 944, National Bureau of Economic Research, July, 1982.

MaCurdy, Thomas E., "An Empirical Model of Labor Supply in a Life Cycle Setting," Journal of Political Economy, December, 1981, 89, 1059-1085.

Mitchell, Olivia S. and Fields, Gary S., "The Economics of Retirement Behavior," Working Paper No. 1128, National Bureau of Economic Research May, 1983.

Sammartino, Frank, "A Model of the Retirement Decision," mimeo, February, 1978.

Zabalza, Anthony, Pissarides, Christopher, and Barton, M., "Social Security, and the Choice Between Full-Time Work, Part-Time Work and Retirement," Journal of Public Economics, October, 1980, 14, 245-276.

Appendix A

Computation Techniques

Two assumptions of the model considerably ease the computational burden of evaluating the likelihood function. They are, first, the assumption that $f_\varepsilon(\varepsilon_i|\delta_i)$ is distributed normally, and secondly, that u is separable with respect to consumption and leisure. This brief appendix will sketch how these assumptions do in fact simplify the calculations.[1]

Consider as an illustration the problem of evaluating the probability that individual i is fully retired in year M. This corresponds to the region Ω_R in Figure 2(a). Using the fact that $f_\varepsilon(\varepsilon_i|\delta_i)$ is distributed normally, this probability can be written as

$$\begin{aligned} Pr_i(R;\ \underline{\theta}) &= \int [\int_{\Omega_{R|\delta_i}} f_\varepsilon(\varepsilon_i|\delta_i)d\varepsilon_i]d\delta_i \\ &= \int \{1 - F[(\varepsilon_{R|\delta_i} - \rho\delta_i)/\sigma_\varepsilon]\}d\delta_i \end{aligned}$$

where $\Omega_{R|\delta_i}$ is the range of values of ε in Ω_R for a given value of δ_i, $\varepsilon_{R|\delta_i}$ is the minimum of this range, and $F[\cdot]$ is the cumulative normal distribution function. Since there are standard routines to evaluate the cumulative normal using function approximations, this simplification permits the (numerical) integration to be performed over one dimension rather than two, provided that a way can be found for evaluation $\varepsilon_{R|\delta_i}$.

The separability of the utility function allows this critical value of ε to be found without iterative techniques. The first step in the procedure

---

[1]The description in this appendix assumes that L(t)=0 for full-time work. If L(t)>0 at full-time work, the appropriate modifications would not significantly increase the complexity of the calculations.

for doing so is to ascertain whether, for a given δ_i, the boundary of Ω_R is with Ω_F or with Ω_P. The answer to this question can be illustrated with the aid of Figure 1. Suppose that for this value of δ_i, ε takes on just the right value so that the indifference curve through A is tangent to the line segment AB of the budget line at time t. If this same indifference curve passes below D in the diagram, then it will be impossible to find a value of ε such that the individual partially retires at time t. This implies that for this value of ε, the boundary of Ω_R is with Ω_F. If the indifference curve passes above D, then for some values of ε the individual will partially retire at time t, and in this case the boundary of Ω_R is with Ω_P.

To find the value of ε such that the indifference curve through A is tangent to AB, first note that with the CES utility function, the maximization problem at time M becomes

$$\text{maximize} \quad \text{sign}(\delta)\{[C(M)]^{\delta} + e^{\underline{X}_i\underline{\beta}+\varepsilon}[L(M)]^{\delta} + k\ e^{-rM}[Y(M) - C(M)]\}$$

where for convenience the i subscripts have been dropped. Since consumption, leisure, and income are additively separable in this expression, and since the definition of Y(M) does not involve consumption, this maximization problem separates into the following two problems:

$$\text{maximize} \quad \text{sign}(\delta)[C(M)]^{\delta} - k\ e^{-rM}\ C(M)$$

$$\text{maximize} \quad \text{sign}(\delta)\ e^{\underline{X}_M\underline{\beta}+\varepsilon}[L(M)]^{\delta} + k\ e^{-rM}\ Y(M)$$

For the indifference curve through A to be tangent to AB, the second maximization problem must have a tangent (i.e., not a corner) solution at L(M) = 1 when $Y(M) = W_p(M)\ [1 - L(M)]$, where $W_p(M)$ is the wage rate for partial retirement work at time M. For this to happen, it must be true that the derivative of the maximand of the second problem with respect to L(M) must be equal to zero when evaluated at L(M) = 1:

$$|\delta|\ e^{\underline{X}_M\underline{\beta}+\varepsilon} - k\ e^{-rM}\ W_p(M) = 0$$

Solving this expression for (e^{ε}/k) yields:[1]

$$e^{\varepsilon}/k = e^{-\underline{X}_M\underline{\beta}}\ e^{-rM}\ W_p(M)/|\delta|$$

Since, given values of $\underline{\beta}$ and δ, the right-hand side of this equation can be calculated, the requirement that the problem should have a tangent solution at L(M) = 1 at time M fixes a value for a ratio involving ε and k. However, it does not allow the two variables to be separately determined.

Fortunately, the question of whether or not the indifference curve tangent to AB at A passes above or below D can be answered without resolving the ratio (e^{ε}/k) into its numerator and denominator separately. Since k must always be positive, the maximand of the second problem can be written equivalently as

$$\text{(A.1)} \qquad \text{maximize} \qquad \operatorname{sign}(\delta)(e^{\varepsilon}/k)\ e^{\underline{X}_M\underline{\beta}}\ [L(M)]^{\delta} + e^{-rM}\ Y(M)$$

Using the value of (e^{ε}/k) just determined, the issue of whether the indifference curve passes above or below D reduces to whether this expression is greater at A or at D. Substituting the appropriate values of L(M) and Y(M) at points A and D into this expression and making the comparison implies that the indifference curve tangent to AB at A passes below D if

$$e^{-rM}\ Y_F(M) > \operatorname{sign}(\delta)(e^{\varepsilon}/k)\ e^{\underline{X}_M\underline{\beta}}$$

---

[1]A similar procedure can be used to find the value of the expression (e^{ε}/k) under the assumption that hours of work in partial retirement are fixed at some level below full time hours.

and passes above it otherwise.

If this indifference curve passes above D, then the boundary of Ω_R is with Ω_P, and the critical value of ε is the one associated with the indifference curve which is just tangent to AB at A. This tangency is interpreted as meaning that the individual is just on the borderline between partial and full retirement at time M. The value of the ratio (e^{ε}/k) associated with this indifference curve has already been found; what remains is to find the value of ε separately. To find this value, first note that at any time t, the second maximization problem may be written as

$$\text{maximize} \quad \text{sign}(\delta)(e^{\varepsilon}/k)\ e^{\underline{X}_t\underline{\beta}}\ [L(t)]^{\delta} + e^{-rt}\ Y(t)$$

After substituting the appropriate value for (e^{ε}/k), this maximization problem determines the labor supply for every time t in the life cycle. From this, it is possible to calculate the present discounted value of total lifetime earnings, which may be denoted as Y_T. Now consider the maximization problem involving consumption at time t:

$$\text{maximize} \quad \text{sign}(\delta)[C(t)]^{\delta} - k\ e^{-rt}\ C(t)$$

From this, consumption for any period t can be calculated as

$$\text{(A.2)} \qquad C(t) = [k\ e^{-rt}/|\delta|]^{\frac{1}{\delta-1}}$$

The present discounted value of the consumption stream must be equal to Y_T in order to satisfy the lifetime budget constraint

$$Y_T = \int_0^T e^{-rt}\ C(t)\ dt$$

Since Y_T has already been determined, this relationship provides a means of calculating k once the expression for C(t) in equation (A.2) is substituted.

With this value of k, and with the value of (e^{ε}/k) previously calculated, it is possible to compute the value of ε for which the individual would be just on the borderline between partial and full retirement at time t, given the particular value of δ. This is the value of ε which marks the boundary between Ω_R and Ω_P for this value of δ.

If the indifference curve tangent to AB at A passes below D, then the boundary of Ω_R is with Ω_F. In this case, the critical value of ε is the one associated with the indifference curve which passes simultaneously through A and D, that is, the value of ε for which the individual is just on the borderline between working full-time and retiring fully. Since this indifference curve passes through both points, for the correct value of ε the maximand in equation (A.1) must have the same value at both points. Substituting the appropriate values of L(M) and Y(M) at these points into the expression and setting the two results equal to each other yields the equation

$$e^{-rM}\ Y_F(M) = \text{sign}(\delta)(e^{\varepsilon}/k)\ e^{\underline{X}_M\underline{\beta}}$$

This equation may be solved for the critical value of (e^{ε}/k) for which an indifference curve passes through both A and D at time M. Using this value of the ratio (e^{ε}/k), the separate value of k may be inferred, as in the last case, by using the lifetime budget constraint. The value of k can then be used to calculate the critical value of ε for which the individual would be just on the borderline between full-time work and full retirement, or alternatively interpreted, the value of ε which marks the boundary between Ω_R and Ω_F.

The two types of boundaries discussed in the last two paragraphs are sufficient to calculate the probability associated with Ω_R. For probabilities involving Ω_F or Ω_P, however, it is necessary to calculate the value of ε for a third type of boundary, namely the boundary between Ω_F and Ω_P. At this boundary point, the individual is just on the borderline between full-time work and partial retirement at time M, and the critical value of ε is the one associated with an indifference curve which is tangent somewhere along the segmented line ABC and which passes through D. Unlike the values of ε for the other kinds of boundary points, the critical value of ε for this boundary cannot be calculated directly but must be approximated by an iterative procedure. Essentially, this iterative procedure begins with some point along ABC. Using the methods discussed earlier, the value of (e^{ε}/k) associated with the indifference curve tangent to ABC at this point is found (at a kink point such as B, a tangent indifference curve is considered to be any indifference curve which is locally above ABC except at B). With this value of (e^{ε}/k), the value of the maximand in question A.2 is examined both at the point and at D to determine whether this indifference curve passes above or below D. If it passes above D, the next point considered along ABC will be to the left of the original point, and otherwise the next point will be to the right of the original point. This process is continued until an indifference curve is found which is tangent to ABC and which passes within a given tolerance distance from D.[1] From the value of the ratio (e^{ε}/k) associated with this indifference curve, the values of ε and k may be separated using the lifetime budget constraint in the manner described before.

The foregoing discussion suggests how to calculate any of the three kinds of boundaries for a given year M, which is sufficient to allow the probabilities of Ω_F, Ω_P, and Ω_R for that year to be computed. For a multiyear sequence,

---

[1]If the segmented line ABC has non-convex bends, as may be associated, for instance, with the exhaustion of Social Security benefits due to the earnings test so that further earnings are not subject to the reduction, then there may be more than one indifference curve which is tangent to the segmented line and which passes through D. In such a case, we are interested in the one which has the lowest ratio (e^{ε}/k), that is, the one with the smallest (in absolute value) slope at point D.

the procedure is very similar, except that the calculation of the limits of Ω_S, where S is the sequence, will in general involve the calculation of boundary points associated with more than one year. For example, consider the area Ω_{PRR} in Figure 2(b). In order to calculate the probability of this region, it is necessary to calculate the boundaries between Ω_F and Ω_P and between Ω_P and Ω_R in Year 1 and the boundaries between Ω_P and Ω_R in Year 2. Let the values of ε associated with these three boundary points be denoted as $\varepsilon_{1,F}$, $\varepsilon_{1,R}$, and $\varepsilon_{2,R}$, respectively. Then the probability of the sequence PRR can be calculated as

$$\Pr(\text{PRR};\ \underline{\theta}) = \int\{F[(\varepsilon_{1,R} - \rho\delta_i)/\sigma_\varepsilon] - F[(\max[\varepsilon_{1,F},\ \varepsilon_{2,R}] - \rho\delta_i)/\sigma_\varepsilon]\}d\delta_i$$

Similar expressions may be readily derived for any other sequence of the kind illustrated in Figure 2(b).

Appendix B

Table B.1
Simulations with the Two Job, Fixed Hours Model.

age	percent retiring, by age: from full-time work	percent retiring, by age: from all work	percent working full-time	percent partially retired	percent fully retired
	A. Base Simulation.				
61	4.8	4.0	78.8	4.6	16.6
62	11.9	10.1	66.9	6.4	26.7
63	8.5	7.8	58.4	7.1	34.5
64	6.6	5.8	51.8	7.9	40.3
65	19.7	13.3	32.1	14.3	53.6
66	6.8	7.3	25.3	13.8	60.9
67	5.8	6.1	19.5	13.5	67.0
68	4.9	4.8	14.6	13.6	71.8
69	4.2	5.0	10.4	12.8	76.8
	B. Observed Results.[a]				
61	6.2	4.8	76.3	7.5	16.2
62	12.8	8.5	63.5	11.8	24.7
63	6.8	5.1	56.7	13.5	29.8
64	13.6	11.3	43.1	15.8	41.1
65	17.4	14.2	25.3	19.4	55.3
66	8.0	7.0	17.7	20.0	62.3
67	1.9	2.1	15.8	19.8	64.4
68	3.2	5.1	12.6	17.8	69.5
69	0.9	0.8	11.7	18.0	70.3
	C. Real Compensation Streams Increased by 50%.				
61	5.7	4.9	69.1	7.8	23.1
62	11.8	10.4	57.3	9.2	33.5
63	8.4	8.5	48.9	9.1	42.0
64	7.0	7.0	41.9	9.1	49.0
65	16.6	13.0	25.3	12.7	62.0
66	5.5	6.8	19.8	11.4	68.8
67	4.9	5.8	14.9	10.5	74.6
68	3.8	4.7	11.1	9.6	79.3
69	3.1	4.4	8.0	8.3	83.7
	D. Pension and Social Security Benefits and Mandatory Retirement Eliminated				
61	4.9	4.1	80.4	4.0	15.6
62	8.0	6.8	72.4	5.2	22.4
63	8.6	7.6	63.8	6.2	30.0
64	6.4	5.7	57.4	6.9	35.7
65	6.0	5.2	51.4	7.7	40.9
66	7.2	6.2	44.2	8.7	47.1
67	9.4	8.8	34.8	9.3	55.9
68	7.3	7.4	27.5	9.2	63.3
69	4.8	4.7	22.7	9.3	68.0

[a] Calculated as the difference in percentages in the indicated state in adjacent age categories. See footnote 1 on p. 35 in the text for further discussion.

Table B.2

Simulations with the Single Job, Fixed Hours Model.

age	percent retiring, by age	percent working	percent retired
A. Base Simulation			
61	6.5	83.5	16.5
62	12.1	71.4	28.6
63	11.1	60.3	39.7
64	11.7	48.6	51.4
65	19.6	29.0	71.0
66	8.8	20.2	79.8
67	7.2	13.0	87.0
68	5.4	7.6	92.4
69	3.5	4.1	95.9
B. Observed Results[a]			
61	6.2	76.3	23.7
62	12.8	63.5	36.5
63	6.8	56.7	43.3
64	13.6	43.1	56.9
65	17.4	25.3	74.7
66	8.0	17.7	82.3
67	1.9	15.8	84.2
68	3.2	12.6	87.4
69	0.9	11.7	88.3
C. Pension and Social Security Benefits and Mandatory Retirement Eliminated.			
61	7.0	82.1	17.9
62	9.5	72.6	27.4
63	10.5	62.1	37.9
64	11.7	50.4	49.6
65	12.0	38.5	61.5
66	10.1	28.3	71.7
67	8.4	19.9	80.1
68	6.9	13.0	87.0
69	5.1	7.9	92.1

[a] Calculated as the difference in percentages in adjacent age categories. For this model, an individual is counted as working if he holds a full-time job, and as retired otherwise.