NBER WORKING PAPER SERIES

INTEGRATING RETIREMENT MODELS

Alan L. Gustman
Thomas Steinmeier

Working Paper 15607
http://www.nber.org/papers/w15607

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2009

This paper was supported by a grant from the National Institute on Aging, 5 R01 AG024337. All opinions and views expressed herein are those of the authors and do not necessarily reflect the views or opinions of either the National Bureau of Economic Research or of the National Institute on Aging.

Integrating Retirement Models
Alan L. Gustman and Thomas Steinmeier
NBER Working Paper No. 15607
December 2009
JEL No. C61,D31,D91,E21,H55,I3,J14,J16,J26,J32

ABSTRACT

This paper advances the specification and estimation of models of retirement and saving in two earner families. The complications introduced by the interaction of retirement decisions by husbands and wives have led researchers to adopt a number of simplifications to increase the feasibility of estimating family retirement models. Our model relaxes these restrictions. It includes the extended choice set created when each spouse makes an independent retirement decision. It also includes the full range of complexity found in dynamic-stochastic models of retirement decision making, so far analyzed only in the context of single earner households. Retirement outcomes include full retirement, partial retirement and full-time work. Reverse flows from states of lesser to greater work are also included. The preference structure incorporates heterogeneity in time preference, varying taste parameters for full-time and part-time work, and the possibility of changes in preferences after retirement. The opportunity set reflects the full range of nonlinearities created by pensions and Social Security. Financial returns are stochastic. Exogenous shocks such as layoffs are also included. Estimation is based on data from the Health and Retirement Study.

The solution method is based on backward induction. We show that this method is superior to a method based on a Nash equilibrium, providing plausible behavioral predictions when Nash equilibrium criteria fall silent.

In contrast to some recent studies, the findings suggest the flow of wives into the labor force in the last few decades has probably reduced the amount of husbands' work. The model also provides plausible responses to various policies. For example, we find that any effort to promote opportunities for partial retirement as a means to increase overall work is likely to be unsuccessful as any induced decline in full retirements is offset by a decrease in full-time work.

Alan L. Gustman
Department of Economics
Dartmouth College
Hanover, NH 03755-3514
and NBER
ALAN.L.GUSTMAN@DARTMOUTH.EDU

Thomas Steinmeier
Department of Economics
Texas Tech University
Lubbock, TX 79409
thomas.steinmeier@ttu.edu

I. Introduction

As the baby boomers begin to swell the ranks of the retired and increase demands on our retirement programs, retirement behavior and retirement saving are becoming an even more central concern of policy makers. Researchers and policy makers must fully understand retirement behavior if they are to determine how rules and regulations governing the Social Security system and private pensions influence retirement and saving outcomes, program participation and costs, and the welfare of retirees.

Researchers continue to make progress toward a comprehensive econometric model of retirement and saving. Our aim in this paper is to make contributions on two levels. First, we contribute to the analysis of retirement and saving at the family level by introducing a level and complexity in the decision making of each spouse not found in previous work. Second, we bring together into a model of family retirement details regarding the choice set, preferences and constraints previously found only in models of retirement in one earner households, and typically not found together in a single model, even of one earner households.

Among the features in our family retirement model: 1. Outcomes include full-time work, partial retirement and full retirement. 2. People are forward looking in their decision making. 3. Saving and retirement are jointly determined. 4. The analysis is structural. It specifies separately the preferences and constraints guiding each individual's behavior. 5. Saving and retirement are modeled on the assumption that one cannot borrow on future income or Social Security. Thus liquidity constraints are incorporated into the analysis. 6. There is heterogeneity in time preference. This means that the response rates to future rewards from wages or from postponing Social Security or pension claiming differ among members of the population. It also means, consistent with the distribution of wealth, that some will be well prepared for retirement and others poorly prepared. 7. There are two utility functions, two decision makers, and one budget constraint. Each spouse makes his or her own decisions. However, the decisions are made with awareness of the other spouse's decision making, preferences and opportunity set, implications for the income and wealth of the household, for joint consumption, and for the welfare of the spouse. 8. In this specification, the solution to the model is not as obvious as in the case with a single (unitary) utility function. Household decision making is solved through backward induction. We show that this set up often provides different results from a solution based on a Nash equilibrium. That is, the Nash equilibrium, which assumes that each spouse takes the decision of the other spouse as given, is not always equivalent to methods in which the choice of each spouse takes into account the possible responses of the other spouse. For example, when decisions arrived at independently by each spouse, including awareness of the responses of their spouse, would leave both spouses worse off, the husband and wife are assumed to avoid this outcome. 9. The interactions of the decisions of each spouse are fully integrated into the estimation procedure. 10. The analysis is dynamic, following retirement outcomes from full-time to part-time work and or retirement, and following saving over the life cycle. 11. The analysis is stochastic. People may reduce their work effort over time, then subsequently increase it. Some people may return to the labor force after retiring fully, or increase their hours of work after partially retiring, as circumstances change, with some events foreseen and others not. Others may change their decisions as they realize they have made an error, or as decisions of husbands and wives interact in ways that have not previously been analyzed. 12. Incentives from pensions and Social Security are included on an individual basis for each individual in the sample. The analysis includes the effects of the very sharp, nonlinear incentives from defined benefit pensions that still account for two thirds of the pension wealth of those on the cusp of retirement. To do this, the estimation makes use of linked data in calculating the opportunity set, including detailed pension plan descriptions obtained from employers and Social Security records. It does not use typical or representative pension incentives. 13. There are minimum hours constraints at work. Many jobs typically held in the prime working years and paying higher wages require full-time work or none at all. Partial retirement is not allowed on the main job unless specified. 14. The wage offer depends on tenure and on hours of work. Typically when one leaves a full-time job, the wage offer drops substantially, as it often does for part-time work. 15. Layoffs are incorporated in the analysis.

The paper contributes to the analysis of decision making at the family level. Alternative rankings of husbands and wives among the many potential outcomes are modeled. This involves a richer array of outcomes than has been considered in earlier models of family decision making. Each spouse decides on full, partial or nonretirement, the path to be taken in future years, the implications for the behavior of the spouse and responses to changes in that behavior.

Econometric estimation of this model is based on the Health and Retirement Study. The resulting model is sufficiently detailed to facilitate policy analysis, allowing researchers to change the specification of the equations representing the incentives created by Social Security, pensions, labor market opportunities, including the availability of partial retirement or spouse employment opportunities, and other constraints on individual behavior. With this model it is possible to judge the determinants of retirement and saving behavior, and to estimate the likely effects on both retirement and saving of those at different parts of the income and wealth distribution.

We have conducted simulations to highlight important interdependencies that are lost in models that ignore the interaction of decisions reached by husbands and wives. For example, given a certain juxtaposition of preferences, higher wealth at the beginning of the period, all other things being equal, may cause the work effort of one of the spouses to increase, even if the work effort of the other spouse remains the same. This is contrary to the usual implication from a unified model where a pure wealth effect should cause a decline in labor supply. Interactions between the decisions of husbands and wives also may induce a path where labor supply is first reduced, then increased.

In a longer paper of the same title written for the National Institute on Aging, we have included simulations of the effects of a number of policy initiatives, including increasing the early entitlement age to Social Security, the effects of the trend from defined contribution to defined benefit plans, and a loss in assets of 25 percent, simulating the effects of the recent recession. To limit the length of this paper, the empirical analysis focuses on a detailed description of decision making at the household level. Simulations are then presented to highlight the unique features of the model.

Section II highlights the many gaps in the retirement literature that have motivated the present study. Section III presents the stochastic, dynamic model of retirement where previously omitted dimensions of retirement and saving are included in an econometric model of decision making by the two earner family. The decision process in the two earner family is explored in Section IV. Data and estimation are discussed in Section V. Section VI presents a number of counter intuitive simulations designed to highlight the importance of specifying the full set of interactions in the two earner family. Section VII concludes. Two appendices have been posted with this working paper on the NBER web site. Appendix 1 describes the construction of the variables used in this study. Appendix 2 explores a number of issues encountered in estimating time preference from information on asset levels and retirement.

II. Gaps in The Retirement Literature

Among the various dimensions of behavior that a retirement model should explain, a first requirement is that it explain the major flows among the various retirement states, including reverse flows.[1] In the process, it must also explain such major features of the retirement hazard as the spikes in retirement at ages 62 and 65. Second, the model should explain the joint distribution of retirement by husbands and wives. Retirement data clearly show that couples coordinate their retirement decisions.[2] A third requirement is that model outcomes are consistent with the wealth accumulated for retirement by each household.

The majority of empirical retirement studies are reduced form or quasi reduced form. The dependent variable typically is some measure of retirement or labor force activity. Independent variables include a measure of incentives affecting labor-leisure choice, including some type of pension and/or Social Security delta, measuring the change in the present value of benefits if retirement is postponed. Typically these models assume that individuals are forward looking (Coile and Gruber, 2000). Consequently, decision makers consider not only the

---

[1] Gustman and Steinmeier (1984) and Maestas (2010, forthcoming) provide detailed descriptive data documenting the various flows among retirement states.

[2] Couples coordinate their retirement despite age differences and differences in incentives to retire (Hurd, 1990; Gustman and Steinmeier, 2000b, 2004; Maestas, 2001; and Michaud, 2003). A general overview of bargaining within the family is provided by Lundberg (1999). Bourguignon, Browning and Chiappori (2006) discuss decision making at the family level, including the inconsistencies between the data and parameters estimated for a unified model in which a single utility function is used to describe the consumption and labor supply of the family. For a nonparametric analysis of the determinants of household consumption that focuses on the problems of a unitary model in analyzing household consumption, but restricts the sample of couples to those that are fully employed and thus does not include an analysis of the determinants of labor supply, see Cherchye, De Rock and Vermeulen (2009).

immediate change in benefits from postponing retirement (or benefit claiming), but also the option value of being able to continue at work.

One important limitation of reduced form models should be noted. They are unable to explain a key feature of the retirement hazard, the spike in retirements at age 62. It is not that the authors are unaware of the importance of the spike in retirement. Reduced form models often include age specific dummy variables on the right hand side of a retirement equation. The coefficient on a variable indicating the individual is age 62 is typically highly significant and very large. This coefficient estimate indicates the height of the retirement spike and age 62, but it does not explain the spike. The problem is that as soon as Social Security benefits become available at age 62, those with high time preference claim their benefits. Because the Social Security benefit structure is actuarially fair, the changes in the present value of benefits with retirement typically included in reduced form retirement equations are zero. Nor can a reduced form equation with an age dummy at 62 be used to analyze the effects of changing policies that have generated the spike. More generally, reduced form models confound the effects of unmeasured preferences with the economic incentives created by retirement programs, limiting the ability to predict the likely effects of policy changes on retirement outcomes.

A related approach uses difference-in-difference analysis to analyze the effects of changes in government programs. Krueger and Pischke’s (1992) analysis of the effects of the notch in Social Security on retirement is a well known early example. In cases, where there are no unmeasured differences between the control group and the experimentals that affect retirement, the difference-in-difference analysis will successfully indicate the effects of the change being studied on retirement. However, this approach is not useful for analyzing the likely effects of new and untried policies, such as raising the Social Security early entitlement age.

In view of inability of reduced form equations to explain the key features of retirement outcomes and the related limitations of reduced form difference-in-difference analyses, researchers have continued working to improve structural retirement models. These models incorporate the separate influences of preferences and elements of budget constraint, including incentives shaped by Social Security and related retirement policies.[3]

Although the specifications of structural models are much more complex than reduced form equations, most structural models rely on simplifying assumptions to facilitate estimation. Unfortunately, in many cases these simplifications are not innocuous. As we will see, oversimplification frequently limits the ability of the model to explain key characteristics of retirement behavior. Even worse, oversimplification sometimes confounds various causes and/or creates biased estimates.

Many structural retirement models ignore partial retirement and/or the complexity of retirement dynamics.[4] Some models include partial retirement but ignore reverse flows (Gustman and Steinmeier, 1986a), while other include retirement dynamics, but ignore partial retirement. For example, while Blau and Gilleskie (2006) pay a great deal of attention to the dynamic structure of the dependent labor market status measure, Cutler, Liebman and Smyth

[3] The strength of reduced form studies is that they do not impose a rigid structure, so they are not subject to the kind of specification error that can affect the interpretation of key parameters in a structural model. Moreover, there is a fear that identification of structural models is based on functional form. As will be seen when we discuss our estimated model, however, estimation is based on many very different moments, while a standard specification is used for the utility function, greatly loosening the relation between identification and the form of the utility function. Moreover, identification is to a large part determined by nonlinearities in the budget constraint caused by retirement programs. These nonlinearities are estimated from original pension documents obtained from employers, as well as from Social Security rules.

[4] An important early study by Berkovec and Stern (1991) included reverse flows in their analysis. That study ignored the role of pensions, as described below.

(2006, p. 19) note they make a trade-off by eliminating partial retirement as an option in the model.

Another major problem is that until recently, structural models of retirement did not explicitly consider the joint determination of retirement and wealth. Saving was something that went on in the background and parameter estimates did not take account of the level of wealth accumulated by the household. These shortcomings were addressed by Rust and Phelan (1997), who estimated a structural model that jointly considered retirement and saving. But they greatly oversimplified the budget constraint by assuming that capital and insurance markets were inoperative. French (2005) estimates a model with joint saving and retirement, a line of research that continues in French and Jones (2004). Van der Klaauw and Wolpin (2006) also make a very useful contribution in their joint analysis of retirement and wealth.

However, these and other structural models continue to oversimplify the relation between retirement and saving. For example, French (2005) assumes the same time preference parameter for all individuals. As Cutler, Liebman and Smyth (2006) note, “French’s approach does not allow for heterogeneity in preferences or heterogeneity in the data generating process for the state variables. Moreover, French uses non-separable preferences that constrain the income and substitution effect of an increase in wage levels to cancel.”

Returning to the complex retirement variable, as Gustman and Steinmeier (2005) show, a model must be able to explain the spikes in retirement observed in the data at ages 62 and 65, as well as the very wide distributions of wealth at each level of lifetime earnings.[5] However, other simplifications found in recent contributions to the structural retirement literature also

[5] Gustman and Steinmeier (1999) and Venti and Wise (1999) demonstrate the extent to which asset accumulation varies among those who have the same lifetime earnings.

importantly affect the ability of the model to reproduce retirement and wealth outcomes. For example, Van der Klaauw and Wolpin (2006) allow for heterogeneity in the consumption parameter, but not in the time preference rate. Heterogeneity in the consumption parameter is not sufficient to simultaneously generate a reasonable wealth distribution and a spike in retirement at age 62. Unless there is some arbitrary parameter in the utility function at age 62, we are essentially back to the same problem that plagues the reduced form option value model.

Some investigators who recognize the important effects of simplification resort to estimating for limited subsamples to avoid having to include certain complexities in their models. For example, a researcher whose model cannot address nonlinearities, such as those created by discontinuities commonly found in defined benefit pension formulas, may limit the sample to those who do not have a defined benefit pension. Rust and his colleagues (in Benitez-Silva et al., 1999 and Benitez-Silva et al., 2004) limit the sample to exclude those who are covered by a pension. Van der Klaauw and Wolpin (2006) eliminate from their sample those who have a DB plan on their current job, as well as those who had a DC plan at any time. French (2005) uses summary measures of pension incentives rather than the individual accrual profiles relevant to the particular observation. Bound et al. (2006) assume, counterfactually, that all defined contribution assets and nonpension wealth are paid out as an annuity. In addition, they do not allow saving in their model. As a result, all payments are assumed to be immediately consumed and there is too rigid a linkage between the timing of income, including the assumed timing of income from assets, and the timing of expenses. Resulting parameter estimates are highly unlikely to be representative of the population values, especially since a large majority of couples approaching retirement age who are covered by a pension continue to be covered by a defined benefit pension plan (Gustman, Steinmeier and Tabatabai, 2010a).

Gustman and Steinmeier (2005) models the joint determination of retirement and wealth, while Gustman, Steinmeier and Tabatabai (2010b) allows reversals of flows from states of greater to lesser retirement. These studies are able to reproduce spikes in the retirement hazard of appropriate size at ages 62 and 65. However, these papers take the retirement decision of the spouse to be exogenous.

Some retirement studies include the effects of joint retirement decision making of husbands and wives, but most do not. Most studies include families, both those with single earners and those including two earners, but they either ignore spouse earnings, or if included treat spouse earnings as exogenously determined. For example, Bound et al. (2006) confine their estimation to single individuals, and thus do not deal with interactions between spouses at the level of the family.

Turning now to studies that do focus on joint retirement decision making, some models such as Coile (1999) are clearly reduced form. They will be of only limited help in policy analysis, and the parameters estimated cannot be directly related to the deeper parameters from the utility function and separated from the elements of the budget constraint.

There are a few structural studies of the determination of retirement within the household. To simplify analysis to make room for considerations arising at the family level, the studies currently available do not include many of the advances found in structural retirement models of the behavior of individuals. For example, Gustman and Steinmeier (2000) focuses on the joint determination of retirement, but the analysis does not include partial retirement, reverse retirement flows, or the joint determination of saving. In the absence of consideration of the joint determination of household wealth together with the retirement decisions of each spouse, the models lose their ability to generate the key spike in retirement at age 62. Blau and Gilleskie (2006) pay a great deal of attention to the dynamic structure of the dependent labor market status measure, but again do not include saving or related behavior. Other studies of retirement suffering from these or related limitations include Blau (1997, 1998) and Michaud (2003). Gustman and Steinmeier (2004), a model that jointly addresses retirement and saving decisions in a family context, nevertheless ignores partial retirement.

Once independent decision making is allowed within the context of a family model, it is useful to consider the various ways in which the retirement decisions of spouses may be related. Gustman and Steinmeier (2000) estimate the separate effects of correlation in budget constraints, correlation in preferences, and dependence of preferences of one spouse on the retirement status of the other.

A simple model for analyzing the retirement dates of the two individuals assumes a household utility function which is a weighted average of the utility functions of the two spouses. Blundell et al. (2001) uses a static, working life framework. There is no saving, and the emphasis seems to be on the issue of whether consumption is private to the two spouses or public between them. The analysis considers the relative weight in the decision process given to each spouse. In this setting, the problem reverts to maximizing a single function with respect to the budget constraint. The weights can be variable, as in Maestas (2001), where the weights depend on the responses of the two spouses to questions relating to which spouse has more influence on the financial decisions of the household.[6]

---

[6] Blundell et al. (2001) use weighted utility functions, but in this model (and unlike Maestas) the weights have an unobserved effects component. Since they eventually linearize virtually everything, this does not have serious consequences, but it would introduce some large difficulties in a more structural environment.

The weighting model can reproduce joint retirement in approximately the right amounts. However, the weighting model does not reproduce the spike in retirements at age 62.[7]

Consider another simplification in the basic specifications adopted in these models. In modeling retirement within the family, Maestas (2001) attempts to distinguish the effects of correlated preferences from bargaining power within the family and introduces information from the HRS on the relative influence of each spouse on decision making within the family. However, she assumes that husbands respond to leisure complementarity while the wife does not, a relationship that should be a product of her analysis.

The models of family retirement also simplify their treatment of pensions. Even though employer provided pension data are very useful in identifying the sources of incentives independently affecting the behavior of husbands and wives, most family retirement models do not include the separate incentives of pension plans held by each spouse. Thus this important source of identification is not available to studies that eliminate those with a pension from the sample, or who use files with imputed pension wealth that by design cannot catch the nonlinearities creating the spike in DB accrual profiles at early retirement age (e.g., Cutler, Liebman and Smythe, 2006). Taking an analogous approach, Maestas (2001) throws out observations with DB plans so she can

[7] Note that it is in general never optimal for the standpoint of one spouse to have the remaining spouse retire earlier, since this does not increase the utility of the first spouse's retirement and only reduces the income available for consumption. Since weighting means that the retirement of each spouse is in effect a compromise between the desires of the two spouses, it tends to eliminate the spikes in retirement that would otherwise be apparent. In our attempts to estimate the model using weighted functions, the retirement spikes of the two spouses were small fractions of 1 percentage point. Thus while we expect interdependence of spouse decision making to reduce the age 62 spike somewhat, and to spread retirements to other years, using a weighted function reduces the age 62 spike much too much, virtually eliminating it.

ignore the pension spikes, losing an important potential source of model identification between husbands and wives in the process. Michaud (2003) also ignores the spikes in accruals. Turning once again to Maestas (2001), she also assumes there is a perfect capital market and that one can freely borrow from Social Security and pension wealth. Once the assumption of perfect capital markets is relaxed, as discussed previously, it is important to allow for heterogeneity in time preference, something none of these models allows. Gustman and Steinmeier (2000, 2004), precursors to the current paper, provide an exception.

All of these efforts by previous researchers have contributed to our understanding of joint retirement behavior within the household, and suggest a number of basic conclusions. They provide an important foundation for ongoing analyses. But the extent of omissions and simplifications creates doubt about the reliability of their findings. To address the remaining questions, we turn to a more comprehensive specification of the family retirement model, and attempt to estimate it using all of the information available in the HRS for describing retirement and saving behavior.

III. The Dynamic, Stochastic Model

The utility functions of the two spouses are fairly standard functions of consumption (a public good within the household) and labor supply over the lifetime. For the husband, the lifetime utility function is give by:

$$U_h = \sum_{t=0}^{T} \left[e^{-\rho t} \sum_{m=1}^{3} s_{m,t} \left(\tfrac{1}{\alpha} C_{m,t}^{\alpha} + e^{X_t^h \beta_h + \varepsilon_h} V_L^h (L_t^h, L_t^w) \right) \right].$$

In this function, $C_{m,t}$ is consumption at time t in survival state m, where m is an indicator whether both spouses are still living at time t, only the husband is living at time t, or only the wife is living. $s_{m,t}$ is the probability that the household will be in state m at time t. L_t^h and L_t^w are the leisure amounts of the husband and wife, respectively, at time t, and V_L^h is the value of leisure to the husband. Note that the value of leisure to the husband may depend on the amount of leisure of the wife in the same period. The exponential form preceding V_L^h is a multiplicative factor for the value of leisure. It consists of a standard linear form $X\beta$ plus an individual effect ε which reflects the strength of the husband's preferences for retirement over work. The elements of X contain a constant, age, and health status. As age increases, work gradually becomes more onerous and retirement more desirable. When the utility of retirement exceeds the utility of consumption from the income earned from work, retirement occurs.

Leisure can take on three values associated with full-time work, partial retirement, and full retirement. The value of retirement for the husband can be modified if the wife is also retired. For the case where the wife is working full-time, the basic value of leisure V_L^h for the husband is normalized to zero if he is working full-time, unity if he is fully retired, and V_P^h if he is partially retired. Note that if leisure is a normal good, which we assume that it is, the value of V_P^h should fall between one-half and one if partial retirement is equated with approximately half-time work. The closer the value is to one, the greater the value of partial retirement is relative to full retirement, and the more frequently and longer should be the spells of partial retirement.

The model allows for the value of retirement to be increased if the husband prefers to spend time with the wife, and the wife is also retired. In the case without partial retirement, this can be accomplished simply by adding a term to the $X\beta$ linear form that has whether the spouse is retired as an additional variable. With partial retirement, however, the picture is a little more complex. The general idea is that the wife's retirement adds to the utility of the husband's retirement only up to the point of the husband's retirement. If the husband is partially retired, it doesn't matter whether the wife is fully retired or partially retired, since the additional leisure of the wife if she is fully retired doesn't add anything to the husband's leisure during the time that the husband is at work. If the husband is fully retired, however, it does make a difference whether the wife is partially retired or fully retired, since only the part of his leisure that he shares with the wife is augmented. If the augmentation factor is g, the values of the function V_L^h as a function of its two arguments can be written as follows:

		Husband's Retirement Status		
		Full-Time Work	Partial Retirement	Full Retirement
Wife's	Full-Time Work	0	V_P^h	1
Retirement	Partial Retirement	0	$g\,V_P^h$	$g\,V_P^h + (1 - V_P^h)$
Status	Full Retirement	0	$g\,V_P^h$	g

If the husband is working full time, he doesn't have any leisure, and the value of the leisure is zero regardless of the retirement status of the wife. If the husband is partially retired, the value of leisure is V_P^h if the wife is working full-time, and this gets multiplied by a factor of g if the wife is working at least part time. If the husband is fully retired, the value of leisure is unity, and this again gets multiplied by a factor of g if the wife is also fully retired. If the wife is partially retired, then we can divide the husband's leisure into two parts. The partial retirement leisure has a value of V_P^h, and this gets multiplied by g because the wife is also partially retired. The remaining leisure has a value of $1 - V_P^h$, but this leisure does not get multiplied because the wife is not there for this leisure. The total value of the husband's leisure is the sum of these two parts.

The utility function of the wife is symmetric:

$$U_w = \sum_{t=0}^{T} \left[e^{-\rho t} \sum_{m=1}^{3} s_{m,t} \left(\tfrac{1}{\alpha} C_{m,t}^{\alpha} + e^{X_t^w \beta_w + \varepsilon_w} V_L^w (L_t^h, L_t^w) \right) \right]$$

Where the superscripts and subscripts w refer to the wife's utility and leisure. The budget constraint for the family is given by the asset evolution equation:

$$A_{t+1} = (1 + r_t) A_t + (1 - L_t^h) W_t^h + B_t^h + (1 - L_t^w) W_t^w + B_t^w + I_t - C_{m,t}$$

All quantities in the equation are measured in real terms. Assets, which are constrained to be non-negative, grow at the real interest rate r_t. The second term on the right side is the husband's earnings, and the fourth term is the wife's earnings. The third and fifth terms are the husband's and wife's pension and Social Security benefits, respectively. Although not indicated by the notation, these benefits depend on the past work and retirement decisions. In the case of Social Security, these can even depend on the past work and retirement decisions of the spouse. The term I_t is any inheritances that the household may receive, and the last term is household consumption. Note that consumption is dependent on the survival state of the household, and that the budget constraint must hold regardless of the mortality experience.

The stochastic structure of the model is as follows. The real interest rate r_t is stochastic and is assumed to come from the actual distribution of asset returns, as documented by Ibbotson. The time periods in this model are annual, and we assume that asset returns are uncorrelated over time, which is approximately the case. The time preference parameter ρ is assumed to be heterogeneous over the population of households, and we treat the value of this parameter as being a fixed effect whose value is estimated for every household in the population.[8] The

[8] It would be preferable to allow time preference to have separate values for each of the two spouses. However, allowing for separate time preference parameters for each spouse within the household would introduce enough additional complexity to preclude our being able to estimate the model.

mortality outcomes reflected in the underlying the survival state probabilities $s_{m,t}$ are also considered stochastic.

The two epsilon terms in the linear forms multiplying leisure, ε_h and ε_w, are considered to come from normal distributions with a mean of zero and standard deviations of σ_h and σ_w, respectively, and with a correlation parameter ρ_ε. These terms allow husbands and wives to have different preferences for leisure vs. work, and the correlation allows for the possibility that one reason that the spouses in households may retire at about the same time is that the two spouses share the same attitudes toward leisure vs. work. For estimation purposes, the initial values of these two epsilon terms are treated as random effects. After the individual first leaves full-time work, the value of leisure may suffer an unanticipated change. For instance, it may fall if the individual finds that leisure is not as pleasurable as anticipated. This may provide one reason why individuals return to work after retirement. To provide for this possibility, the model allows the value of the epsilons to change over time after the initial retirement, with correlation parameters over time of ρ_h and ρ_w for the husband and wife, respectively. We considered the possibility that the spouses might anticipate that the values of leisure might change after retirement, but this led to problems that are discussed in Appendix 2. As a result, the model assumes that these changes are completely unanticipated.

The final stochastic elements, which are stochastic from the estimator's point of view but not from the individual's point of view, are the values of partial retirement, V_P^h and V_P^w. These are assumed to come from an exponential distribution $f(V_P^h) = c e^{\gamma_t^h V_P^h}$ defined over the interval 0.5 to 1, where c is a factor necessary to make the integral of the distribution over the allowable range equal to 1, as must be the case for all distributions. The higher γ_t^h, the greater the probability that the husband's value of V_P^h will be close to 1, which would make it more likely that he will go through a period of partial retirement. Conditional on employment, the probability of partial retirement increases with age. Accordingly, we allow the value of γ_t^h to increase with age according to the equation $\gamma_t^h = \gamma_o^h + \gamma_1^h a_t^h$, where a_t^h is the husband's age at time t. We characterize the husband's relative preference for partial retirement as not changing. Consequently, the model assumes that the husband's relative position within this distribution does not change even as the entire distribution changes to higher values of V_P^h. The values V_P^h for the wife are treated in a symmetric way.

The sequence of events in the model is as follows. In every period, the couple starts with a level of assets. If the husband was still in his career job in the previous period and if that job had a defined contribution pension, he may also have a given level of defined contribution pension assets. If he was not in the career job in the previous period but had a defined benefit pension in that job, he may also have a pension benefit amount, which may or may not be currently collectable. The same situation with regard to pension amounts also applies to the wife. Further, the household may be eligible to collect a given level of Social Security benefits, either currently or sometime in the future. Given their current situation, the spouses make their decisions as to what their retirement status (working full-time, partial retirement, or full retirement) will be during the current period, and how much they will save or spend down from their accumulated assets. At the end of the current period, a random draw is taken from the distribution of asset returns, and assets plus any defined contribution amounts are assumed to grow at that rate of asset returns. One of the spouses may die, causing a transition between the survival states. Depending on the retirement status decisions made in the current period, the amount of future defined benefit pensions and Social Security amounts may also be increased.

The model assumes that before age 50, the spouses are working at whatever their earnings were before age 50. After age 70, the spouses are assumed to be retired. These assumptions are made in order to reduce computational burden somewhat. Between ages 50 and 70, the spouses make their decisions about working full-time, partially retiring, or fully retiring. With two utility functions, two decision makers, and one budget constraint, the solution of the model is not as obvious as in the case with a single utility function. The mechanism by which they make these decisions is relatively complex and is described in detail in the next section. Note that if one of the spouses is below 50 years of age or above 70, and the other is within the 50 to 70 age range, only the spouse within the age range will need to make a work/retirement decision. At all ages, however, decisions are made as to how much to consume, given income in the current period and the assets owned when the couple entered the period.

The number of state variables depends on what period of life the spouses are in, and whether they are still in their career jobs. For all ages, one of the state variables is the amount of accumulated assets. Below age 50, the primary additional state variable is the amount accumulated in defined contribution pension plans, if the career job had these kinds of pensions. Up to age 70, each spouse has a value of ε, which relates to the preference for work vs. retirement. After age 70, ε is irrelevant, since past that age individuals are presumed to be retired. Between ages 50 and 70, the most important state variable is whether each spouse is still in his or her career job. If so, the associated state variable is the amount in the defined contribution plan. If the spouse has left the career job, there are two state variables associated with pensions and a third associated with Social Security. One pension variable relates to the level of pension benefits, and the other is a binary variable related to whether the individual can currently receive the benefits or must wait until a later date, as would occur if the individual had left the career job before becoming eligible for early retirement benefits. The Social Security variable is one of a pair of such variables (one for the husband and one for the wife) that define the level of joint Social Security benefits that the couple is eligible for. The pension and Social Security variables are initially set at the point the individual leaves the career job. For the pension variable, this value is adjusted downward over time, since few pensions are fully indexed for inflation. For the Social Security variable, the value for future periods may be increased if the individual returns to full-time work or engages in partial retirement work, an increase that is due either to a reduction in the early reduction amount or an increase in the delayed retirement amount. However, the amount of the Social Security value that may be collected in the current period may be reduced if the individual works either full-time or part-time and the earnings from that work exceed the earnings test limit of Social Security. Past age 70, the pension amounts (except for the inflation erosion) and the Social Security amounts are fixed at their levels at age 70.

In the estimations and simulations, all continuous variables (assets, the epsilons, pension and Social Security amounts) are broken down into a vector of discrete amounts. If required, any amount between these discrete amounts is interpolated.

IV. The Decision Process

This section describes the process by which, in each period, the husband decides how much he will work and the wife decides how much she will work. Each spouse is cognizant of the possibility that their own decision may influence the decision of the spouse, and they take account of these ramifications in making their own decision. Since these decisions are being made in a backwards induction framework, it is assumed that the spouses have already calculated how their current decisions will affect their future decisions.

A. Joint Decisions with a Work/Retirement Choice.

To begin this analysis, it will help to look first at a framework where each spouse is deciding either to work full time or retire. This leaves the analysis of partial retirement until later in the section. In such a framework, the joint decision result has four possibilities: both spouses work; both spouses retire; the husband works and the wife retires; and the husband retires and the wife works. These possibilities can be represented by the following grid:

		Husband's Decision	
		Work	Retire
Wife's	Work	A	B
Decision	Retire	C	D

Thus, decision B is for the husband to retire and the wife to work, and similarly with the other alternatives.

Both the husband and the wife have utility values associated with each of these alternatives. Let U_B^h be the husband's utility if the wife works and the husband retires. This utility will include the utility of decisions in future periods, given the income and consumption in the current period and any restrictions that the current decision to work places on future decisions. For instance, the decision to be retired in the current period may mean that in any future periods the wage rate for full-time work will be lower than if the husband had continued in his career job. Similarly, U_C^w would denote the wife's utility if she retires and the husband works full time.

If $U_A^h < U_B^h$, it means that the husband prefers to retire, given that the wife is working full-time. If $U_C^h > U_D^h$, it means that the husband prefers working full-time, give that the wife is retired. This situation can be represented by drawing horizontal arrows between the corresponding decision points, as follows:

$$\begin{array}{ccc} A & \rightarrow & B \\ C & \leftarrow & D \end{array}$$

where the arrows point to the preferred decision between the two choices. If $U_A^w > U_C^w$, it means that the wife prefers working full time to retirement, given that the husband is also working full-time. If $U_B^w < U_D^w$, it means that the wife prefers retirement to working, given than the husband is retired. The situation for the wife can be represented by vertical arrows between the corresponding decision points:

$$\begin{array}{ccc} A & & B \\ \uparrow & & \downarrow \\ C & & D \end{array}$$

where again the arrows point to the preferred decision between the two choices. We can represent the combined preferences for both the husband and the wife by including both sets of arrows, as follows:

$$\begin{array}{ccc} A & \rightarrow & B \\ \uparrow & & \downarrow \\ C & \leftarrow & D \end{array}$$

Since it takes four arrows to denote the relative preferences of both the husband and wife, and since any of these arrows can point in two directions, there are sixteen combinations of arrows that represent any combination of preferences of the two spouses. These sixteen combinations may be represented as follows:

1. $\begin{array}{ccc} A & \rightarrow & B \\ \downarrow & & \downarrow \\ C & \rightarrow & D \end{array}$	2. $\begin{array}{ccc} A & \rightarrow & B \\ \downarrow & & \downarrow \\ C & \leftarrow & D \end{array}$	3. $\begin{array}{ccc} A & \leftarrow & B \\ \downarrow & & \downarrow \\ C & \rightarrow & D \end{array}$	4. $\begin{array}{ccc} A & \leftarrow & B \\ \downarrow & & \downarrow \\ C & \leftarrow & D \end{array}$
5. $\begin{array}{ccc} A & \rightarrow & B \\ \downarrow & & \uparrow \\ C & \rightarrow & D \end{array}$	6. $\begin{array}{ccc} A & \rightarrow & B \\ \downarrow & & \uparrow \\ C & \leftarrow & D \end{array}$	7. $\begin{array}{ccc} A & \leftarrow & B \\ \downarrow & & \uparrow \\ C & \rightarrow & D \end{array}$	8. $\begin{array}{ccc} A & \leftarrow & B \\ \downarrow & & \uparrow \\ C & \leftarrow & D \end{array}$
9. $\begin{array}{ccc} A & \rightarrow & B \\ \uparrow & & \downarrow \\ C & \rightarrow & D \end{array}$	10. $\begin{array}{ccc} A & \rightarrow & B \\ \uparrow & & \downarrow \\ C & \leftarrow & D \end{array}$	11. $\begin{array}{ccc} A & \leftarrow & B \\ \uparrow & & \downarrow \\ C & \rightarrow & D \end{array}$	12. $\begin{array}{ccc} A & \leftarrow & B \\ \uparrow & & \downarrow \\ C & \leftarrow & D \end{array}$
13. $\begin{array}{ccc} A & \rightarrow & B \\ \uparrow & & \uparrow \\ C & \rightarrow & D \end{array}$	14. $\begin{array}{ccc} A & \rightarrow & B \\ \uparrow & & \uparrow \\ C & \leftarrow & D \end{array}$	15. $\begin{array}{ccc} A & \leftarrow & B \\ \uparrow & & \uparrow \\ C & \rightarrow & D \end{array}$	16. $\begin{array}{ccc} A & \leftarrow & B \\ \uparrow & & \uparrow \\ C & \leftarrow & D \end{array}$

We will analyze each of these combinations in turn. Since the combinations below the main diagonal of this chart (combinations 5, 9, 10, 13, 14, and 15) are mirror reflections of combinations above the main diagonal, we need only to analyze the combinations on or above the main diagonal, and the results can be applied to the mirror images.

1. This is a case where the husband wants to retire, given the wife's decision, and the wife wants to do likewise. But the arrows do not take account of the fact that each spouse's decision may influence the other spouse's decision. The husband has two choices: to retire regardless of the wife's decision or to wait to see what the wife will do and then decide what he wants to do as a result. If he retires, he knows that his wife will also retire, and the result will be combination D. If he waits to make his decision, the decision shifts to the wife. She knows that if she retires, he will too, and the result will be combination D. If neither one takes the decision to retire, they will stay at combination A. Thus, the final choice for both spouses is between combination A and combination D. If either spouse prefers combination D to combination A, that spouse will retire, and the other spouse will follow, and the final result will be combination D. But if both spouses prefer combination A to combination D, neither spouse

will have the incentive to be the first to retire, and the final result will be combination A, which both spouses prefer.

2. In this case the husband prefers to retire if the wife works, but to work if the wife retires. The wife wants to retire, given either decision of the husband. The key to understanding this case is to compare the husband's utility of combinations A and C. These are two combinations where the husband is working, but the wife is either working or retired. Since the husband is working and the utility value of leisure is zero, whether the wife works or not does not affect the value of the leisure term in the husband's utility function. However, if she works, there will be additional income to the family, and this additional income will increase the consumption component of the husband's utility function. Hence, the husband's utility will be higher if the wife works, and he will prefer combination A to combination C. Since in this case the husband prefers combination C to combination D and combination B to combination A, this gives a strict ordering of the husband's preferences among the four combinations: B > A > C > D, where combination B is the most preferred and combination D is the least preferred.

The husband might prefer combination B to combination A, but he knows that if he retires his wife will certainly follow suit, leaving him with combination D, his least preferred result. To avoid this, the husband will continue working. The wife knows that the husband will continue working in any case, so her choice is between combination A and combination C. According to her preferences, she chooses combination C. Thus, combination C is the final result in this case.

3. Here, the wife's preference is to retire, given a retirement decision by the husband. For the husband, his preference is to do the same thing as his wife does, given her decision. This case is similar to case 1. The husband knows that if he retires first, his wife will follow, leading to

combination D. The wife knows that if she retires first, her husband will follow, again leading to combination D. If neither retires, they will remain at combination A. The final decision depends on the relative preferences between combinations A and D by the two spouses. If either spouse prefers combination D to combination A, that spouse will retire, and the other spouse will follow suit. If both spouses prefer combination A to combination D, neither spouse will want to be the first to retire, and combination A will be the result.

4. This is a case where the husband wants to work, given either decision of his wife, and the wife wants to retire, given either decision of her husband. The husband knows that if he retires, his wife will too, and combination D will be the result. However, by the same reasoning as in case 2, the husband will prefer combination A to combination C, and given the preferences in this case he will prefer combination C to combination D. By transitivity, this means that both combinations A and C are preferred to combination D, so the husband will not want to be the first to retire. Given that the husband is continuing to work, his wife's choice comes down to combination A or combination C, and given her preferences, she will choose combination C. Thus, the final result in this case is combination C.

6. In this case, both partners want to do the opposite of their spouse's choice, given their spouse's choice. By the same reasoning as was applied to case 2, the four outcomes can be ranked by the husband as B > A > C > D, with combination B as the most preferred and combination D as the least preferred. Symmetric reasoning applied to the wife's preferences yields a ranking by the wife of the four outcomes as C > A > B > D, with combination C as the most preferred and combination D as the least preferred. The ideal situation for the husband would be for him to retire first. In this case, the best choice for the wife would be for her to continue working, leading to choice B. This is the best outcome from the husband's perspective,

but only the third choice outcome for the wife. For the wife, her ideal situation is for her to retire first. If she does this, her husband will continue to work, and the outcome will be combination C. This is her most preferred outcome, but not so good for the husband.

The incentives for this case are for each spouse to be the first to retire. But if they both retire, the result will be combination D, which is the worst result for both spouses. We assume in this case that the spouses cooperate enough that they want to avoid the worst result, and that neither one retires. This would lead to combination A, which although not the most preferred outcome for either spouse, is the second choice outcome for both spouses. This avoids a destructive race to see who can be the first to retire, which risks producing the least preferred outcome for both spouses.

7. This is a case where the husband wants to do the same thing as the wife, given the wife's decision, but she wants to do the opposite of the husband, given his decision. The key to analyzing this case is to look as the wife's ordering of the four outcomes. By the same reasoning as used in case 6 (which is symmetric to the reasoning used in case 2), her ordering of the four outcomes is C > A > B > D, where C is her most preferred outcome and D is her least preferred outcome. She knows that if she retires, her husband will follow suit, and the result will be combination D, which is her least preferred outcome. Thus, she will not want to retire. This leaves the husband with the choice between combinations A and B, and by his preferences in this case he will choose not to retire. This leaves the final result in this case as combination A, where both spouses continue working.

8. Here, the husband wants to continue working, given either choice by the wife, while the wife prefers to work if the husband retires and retire if the husband works. For the wife, by the same reasoning used in case 5, the ordering of her preferences is C > A > B > D. If the wife

retires first, the husband will continue working, and the resulting combination C will give her the most preferred outcome. If the husband were to retire first, the wife would continue to work, and the resulting outcome would be combination B. The final result depends on the husband's evaluation of combination B vs. combination C, an ordering that is not a priori determined by the preferences in this case. If the husband prefers combination C, he may simply let the wife retire first. If the husband prefers combination B to combination C, the situation reverts to much the same situation as occurred in case 6. Both spouses will have an incentive to retire first, but if both succeed in retiring, the result will be combination D. This combination is the least preferred result for the wife. For the husband, the ordering of the outcomes becomes A > B > C > D, so that combination D is his least preferred outcome as well. Thus, both spouses would have an incentive to avoid the consequences of both of them trying to retire first, and both retiring. We assume that in this case, the spouses cooperate enough that neither of them tries to be the first to retire, and the resulting outcome of combination A is more preferred by both spouses than is combination D, which would be the result if both of them tried to retire first and both succeeded. The final outcome in this case is combination C if the husband prefers combination C to combination B, and combination A otherwise.

11. In this case, both spouses prefer to do the same thing as the other spouse, given the other spouse's decision. The husband knows that if he retires his wife will as well, and the resulting outcome will be combination D. Symmetrically, the wife knows that if she retires, her husband will too, with combination D as the result. If neither retires, the result will be combination A. The final outcome depends on the assessment of both spouses of combination A relative to combination D. If either spouse finds that combination D is preferable to combination A, that spouse will retire and the other spouse will follow, with combination D as the final result.

If both spouses find that combination A is preferable to combination D, neither will want to be the first to retire, and the final outcome will be combination A.

12. The preferences in this case indicate that the husband wants to work, given either decision by the wife, while the wife wants to do the same thing as the husband, given the husband's decision. This is perhaps the most complicated of the cases. The husband knows that if he retires, the wife will also retire, and the resulting outcome will be combination D. If the husband prefers combination A to combination D, he will want to continue working. Since the wife will want to continue workings if the husband continues working, the resulting outcome will be combination A. If, on the other hand, the husband prefers combination D to combination A, he will have an incentive to retire. If the wife prefers combination D to combination C, she will allow him to retire, and she will follow, resulting in combination D being the outcome. If, however, she prefers combination C to combination D, she would want to be the first to retire, since if she retires first, the husband will continue working. This outcome would also be preferable to the husband over combination D, so he will refrain from retiring and allow the wife to retire first. The final outcome will be combination A if the husband prefers it over combination D. Otherwise, the outcome will be combination D if the wife prefers it over combination C, and if not the outcome will be combination C.

16. This case reflects a situation where both spouses want to continue working, given either decision of the other spouse. The husband knows that if he retires first, the wife will continue to work. This would result in combination B, which the husband finds less preferable than combination A. Similarly, the wife knows that if she retires first, the husband will continue to work. For her, the resulting combination C is less preferable than combination A. Thus, both

spouses will have an incentive not to retire, and the final result is that both of them will work, yielding combination A as the result.

These results may be summed up in the following table:

	Husband's Preferences			
Wife's Preferences	A $\rightarrow$ B C $\rightarrow$ D	A $\rightarrow$ B C $\leftarrow$ D	A $\leftarrow$ B C $\rightarrow$ D	A $\leftarrow$ B C $\leftarrow$ D
A B $\downarrow$ $\downarrow$ C D	A if $U^h_A > U^h_D$ and $U^w_A > U^w_D$, otherwise D	C	A if $U^h_A > U^h_D$ and $U^w_A > U^w_D$, otherwise D	C
A B $\downarrow$ $\uparrow$ C D	B	A	A	C if $U^h_C > U^h_B$, otherwise A
A B $\uparrow$ $\downarrow$ C D	A if $U^h_A > U^h_D$ and $U^w_A > U^w_D$, otherwise D	A	A if $U^h_A > U^h_D$ and $U^w_A > U^w_D$, otherwise D	A if $U^h_A > U^h_D$, otherwise D if $U^w_D > U^w_C$, otherwise C
A B $\uparrow$ $\uparrow$ C D	B	B if $U^w_B > U^w_C$, otherwise A	A if $U^w_A > U^w_D$, otherwise D if $U^h_D > U^h_B$, otherwise B	A

In the table, the columns represent the husband's preferences. For instance, the second column corresponds to the case where $U^h_A < U^h_B$ (the husband prefers to retire, given that the wife is working) and $U^h_C > U^h_D$ (the husband prefers to work, given that the wife is retired). The rows refer to the wife's preferences. In several cases, the preferences between working and retiring, given the retirement state of the spouse, are not sufficient by themselves to determine the decision, in which case other relationships between the utilities of one or both spouses are required to ascertain which outcome will be chosen.

B. Decisions with Partial Retirement.

The addition of partial retirement introduces a third choice for both the husband and the wife, but the general method of analysis remains the same. To analyze the choices with partial retirement, we denote the nine possible outcomes as follows:

		Husband's Retirement State		
		Working	Partially Retired	Fully Retired
Wife's	Working	A	B	C
Retirement	Partially Retired	D	E	F
State	Fully Retired	G	H	I

Again, each of these states has a utility value associated with it for both the husband and the wife. Each spouse knows that the other spouse may react to their decision, and they make their own decisions with regard to the other spouse's reactions.

The general approach is similar to the case previously described with only the two choices of working or retirement. At the beginning of each period each spouse calculates the utility of retiring, given the other spouse's reaction. Each spouse then ascertains whether it would be better to go ahead and retire, given the other spouse's reaction, or continue working and allowing the other spouse to make the decision either to continue working or retire. If both spouses conclude that they are better off working, or if only one spouse concludes that it is in his or her interest to retire first, then those decisions determine the final outcome. If each spouse concludes that it is in his or her interest to be the first to retire, then there is a possible conflict in the decisions that is in need of resolution.

To be more specific, consider the case of the husband deciding whether to be the first to retire. If he initially retires fully, then the wife must decide, given the husband's decision to retire fully, whether to continue working, to partially retire, or to retire fully herself. Call the utility value of the resulting combination U^h_{FR}, which indicates the final utility if the husband makes the initial decision to fully retire. An alternative is for the husband to initially retire partially. In that case, the wife must again decide whether to continue working, to retire partially, or to retire fully. If she decides to retire partially, that corresponds to combination E on the diagram above. But the repercussions may not stop there. In response to her decision to retire partially, the husband may decide to retire fully as a result, and it is possible that this may further induce the wife to retire fully. The analysis from combination E may be done in the same fashion as the two-by-two analysis done previously, considering only the combinations E, F, H, and I. From E, any of these combinations are possible as each spouse decides whether it is in his or her best interest to retire fully, given that both spouses are at least partially retired. An alternative for the wife is that she may decide to retire fully, given that her husband has partially retired. In this case, she takes into account that her husband may decide to retire fully if she does. Given the possible reactions of the husband, then, the wife compares the utility of the final outcomes and decides whether to partially retire or fully retire, given that her husband partially retires. For the husband, the utility of the final outcome, given his initial decision to partially retire, may be denoted as U^h_{PR}. For the husband, the decision as to whether he is better off by initially retiring partially or fully is determined by comparing U^h_{FR} and U^h_{PR}. H would make the choice associated with the greater utility. That utility may be denoted as U^h_R, which may be interpreted as the utility that the husband would derive if he were the first to make the decision to retire in some form. The wife makes a symmetric decision as to what her utility would be if she were the first to make the decision to retire in some form. Call this utility U^w_R.

Each spouse must then make the decision whether it would be advantageous to retire first. If the husband retires first, his utility would be U_R^h. If he does not retire first, then the wife would choose whether to retire or to keep working. That is, she would choose between her retirement choice, with utility U_R^w, or the outcome where both spouses were working full-time, which for her has a utility value of U_W^w. Let the higher utility value be denoted as U_{wC}^w, where the subscript wC stands for "wife's choice." Correspondingly, let the husband's utility for this outcome be U_{wC}^h. This is basically the utility for the husband if he initially just keeps on working. The husband then compares U_R^h to U_{wC}^h to decide whether he wants to make the initial move towards retirement. If U_R^h is the higher value, he wants to retire initially because if he leaves the initial decision to his wife, the utility value will be lower. If U_{wC}^h is the higher value, on the other hand, he is content to let the wife make the initial decision to retire if she so chooses because that will yield a higher utility to the husband than if he retires first. Needless to say, the wife is going through a similar process to decide whether it is to her advantage to be the first to retire.

There are four possible outcomes to this process. Each spouse may conclude that it is not in his or her best interest to initiate retirement, in which case both will continue to work full-time. A second possibility is that the husband may conclude that it in his best interest to partially or fully retire, knowing that the wife may ultimately partially or fully retire as a result, and the wife concludes that it is in her interest to let the husband do so. A third possibility is the symmetric case where the wife concludes that it is in her best interest to initiate retirement, and the husband concludes that it is in his interest to let her do so. In both the second and third cases, there is no conflict between the spouses as to what to do, so the spouse that initially decides on retirement can do so, and the other spouse will respond accordingly. The fourth case is that both spouses may decide that it is to their advantage to make the initial retirement decision before the other. In this last case there is a conflict between the decisions of the two spouses which needs to be resolved.

If both spouses attempt to carry out their retirement desires simultaneously, the final result depends on whether those desires were for partial retirement or full retirement. If both spouses want to partially retire, then they would be a combination E in the diagram above. At that point, one or both of them could decide that they would prefer to be fully retired, according to the decision dynamics relating to the combinations E, F, H, and I. If the decision of one of the spouses is full retirement, then the remaining spouse would have to choose whether, given the full retirement decision of the first spouse, the second spouse wanted to be partially or fully retired. If the decision of both spouses is to be fully retired, then that is the choice if both spouses pursue their desired for retirement. Whatever the case, let the utility of the two spouses if they both pursue retirement be given as U^h_{bR} and U^w_{bR}, where bR signifies that both spouses are pursuing their retirement goals.

This is not the end of the story, however. It may well be that having both spouses working dominates the scenario where both spouses try to retire. This would occur if $U^h_W > U^h_{bR}$ and $U^w_W > U^w_{bR}$. In this case, it is assumed that both spouses know that if they retire, the other spouse will retire as well. For both spouses, it would be better to not retire as long as the other spouse didn't retire, and as a result both spouses would continue working. The economic rationale for this case is that if one spouse retires, he or she loses not only his or her earnings, but the earnings of the other spouse as well. Even if it turns out that having both spouses working dominates having both retired, it may be true that allowing the husband to retire first (outcome hC) or the wife to retire first (outcome wC) dominates, for both spouses, having both spouses working. If that occurs, then one spouse will allow the other spouse to make the retirement decision which delivers a higher utility for both of them. Otherwise, if having both spouses working dominates having both retired, the final outcome will be for both spouses to be working.

C. Unexpected Implications of the Decision Process.

There are a couple of implications of this decision process that might not be expected, but which have entirely logical explanations. The first is that retirement reversals can occur without any changes in preferences (other than the increasing desire for retirement over time) or any changes in effective compensation. Normally, without changes in preferences or compensation, one would expect that once an individual retired, he or she would stay retired, since nothing has changed to bring them out of retirement. But in the present model, that is not always the case.

This effect can be illustrated with simple work/retirement decisions, without the need to introduce additional complications via partial retirement. Suppose that in a given year, the husband's and the wife's utility matrices look as follows:

		Husband's Decision	
		Work	Retire
		Husband's Payoff Matrix	
Wife's	Work	5.032577	5.026383
Decision	Retire	5.023245	5.015767
		Wife's Payoff Matrix	
Wife's	Work	5.042582	4.997644
Decision	Retire	5.072982	5.026156

Given these utility matrices, the husband prefers to work for any given retirement state of the wife, and the wife prefers to be retired for any given retirement state of the husband. This corresponds to case 4 of the analysis in the previous part of this section, and in that analysis we concluded that in this case, the husband will work and the wife will be retired.

Suppose that a year later, the utility matrices look as follows:

		Husband's Decision: Work	Husband's Decision: Retire
Husband's Payoff Matrix			
Wife's Decision	Work	4.924987	4.931819
	Retire	4.915709	4.921649
Wife's Payoff Matrix			
Wife's Decision	Work	4.932621	4.889473
	Retire	4.967829	4.923471

The utility values are generally a little bit lower, largely reflecting the higher mortality rates as individuals age. The wife still prefers retirement regardless of the retirement state of the husband, but due to his being a year older, the husband now prefers to be retired regardless of the retirement state of the wife. This corresponds to case 1 in the previous analysis. For that case, the outcome depended on the comparisons of combination A, where both spouses worked, with combination D, where both spouses are retired, for both spouses. For both spouses, combination A is preferred, so the final result is that both spouses will be working. Since in the previous year, the husband was working but the wife was retired, this means that the wife has switched back to working from retirement even without any change in compensation or retirement preferences.

The essence of this case can be summarized as follows. In the first year the wife is retired, and the husband is working but is close to the border between working and being retired. In the following year the husband crosses that border and would be retired if he were just giving up his own earnings by doing so. Each spouse knows that if they retire rather than work, the other spouse will also retire, because for each of them the extra consumption from their own earnings is insufficient to compensate for the lost leisure. However both spouses know that if they retire and as a consequence the other spouse retires as well, they will lose both incomes. Since in this example the loss of both incomes outweighs the loss of leisure for either spouse, neither spouse sees an advantage in retiring as long as the other spouse does not retire. The wife essentially returns to working as an inducement for the husband to continue working and provide additional income which benefits them both.

A second rather unexpected implication of the decision process is that higher wealth at the beginning of the period, all other things being equal, may cause the work effort of one of the spouses to increase, even if the work effort of the other spouse remains the same. This is contrary to the usual implication that a pure wealth effect should cause a decline in labor supply.

This effect can again be illustrated with a simple work/retirement decision, without the need to introduce additional complications via partial retirement. Suppose that at a given level of initial wealth, the husband's and the wife's utility matrices look as follows:

		Husband's Decision	
		Work	Retire
		Husband's Payoff Matrix	
Wife's	Work	4.288884	4.266913
Decision	Retire	4.266930	4.261674
		Wife's Payoff Matrix	
Wife's	Work	4.312152	4.271539
Decision	Retire	4.315635	4.271357

Using the previous analysis, this corresponds to case 8. The husband prefers to work regardless of the retirement state of the wife. The wife prefers to be retired if the husband works, but if he is retired she prefers to work. In the analysis of that case, the outcome depended on a comparison of combinations B and C for the husband (B is the combination where the husband retires but the wife works, and C is the combination where the husband works and the wife retires). By a slight margin, the husband prefers combination C to combination B, so the final result is that the husband works and the wife retires.

At a slightly higher level of initial wealth, the two utility matrices look as follows:

		Husband's Decision: Work	Husband's Decision: Retire
		Husband's Payoff Matrix	
Wife's Decision	Work	4.289692	4.267817
	Retire	4.267776	4.262585
		Wife's Payoff Matrix	
Wife's Decision	Work	4.312960	4.272443
	Retire	4.316480	4.272268

Note that all of the utility amounts, conditional on work status, have increased due to the increased wealth. The pattern of utilities still fits the criteria for case 8, but in this case the husband prefers combination B to combination C. In this case both spouses would like to be the first to retire, since they know that if they retire first the other spouse will stay on the job. However, if both retire, they will both arrive at combination D, the least preferred result, so they both continue to work. Compared to the case with lower initial wealth, the husband's work effort has remained the same, but the wife's work effort has increased in the case with higher initial wealth.

The essence of this situation may be explained as follows. At the lower level of wealth, the husband finds it advantageous to work even if the wife retires. At a slightly higher level of wealth, however, the husband finds retirement somewhat more appealing, he would prefer to be retired and have the wife working rather than the other way around. In this case, both spouses would like to retire and have the other work, but if they both retire they will be in the worst possible scenario for both of them. To forestall this, both agree to continue working. In essence, the wife has changed her behavior to keep the husband working even though he would otherwise retire at the higher wealth level.

D. Nash Equilibrium vs. Sequential Decisions.

The difference between Nash equilibria and the results of the decision process described above can be readily seen in the bivariate work/retirement scenario, so we will examine it in this scenario rather than introduce the further complications of partial retirement. The simplest case is illustrated in case 1 in part A of this section. In this case, both spouses would find it to their advantage to retire, given either decision of the other spouse. The Nash equilibrium in this case is straightforward: both spouses would retire. And yet, the possibility remains that both spouses might be better off if both of them continued to work. The decision process described in this paper recognizes that both spouses realize that if either of them retire, the other spouse will retire as well, and hence they are both reluctant to retire for fear of inducing the other spouse to retire as well. The essential difference is that the Nash equilibrium assumes that the actions of one spouse have no impact on the decisions of the other spouse, which may be appropriate for a pair of prisoners held incommunicado, but is probably less appropriate for spouses. In contrast, the decision process of this paper assumes that both spouses realize that their own choices may affect the choices of the other spouse.

There may also be a case where there is no Nash equilibrium, yet the decision process described in this paper will deliver an unambiguous and defensible solution. In terms of the classification of cases in Part A of this section, cases 7 and 10 have no Nash equilibria. These are cases where one spouse wants to do the same as the other spouse, but the second spouse wants to do the opposite of the first spouse. In these cases, no Nash equilibrium can make both spouses completely happy with their decision, given the decision of the other spouse. Case 7 was the case where the husband wanted to do the same thing as the wife, but the wife wanted to do the opposite of the husband. Utilizing the sequential decision making process, the wife knows that if she retires, the husband will as well, and that is her least preferred outcome. So she continues working and the husband does too. The wife would really prefer that she retire and the husband stays working, but she knows that this is an impossible outcome, because she knows what the husband will do in response to her retirement. So she settles for her second best outcome, and the decision process has a definite and defensible result in spite of the fact that there are no Nash equilibria in this case.

In any case, as we have shown, the backward induction method does not always yield the exact same result as the Nash equilibrium, even in cases where a Nash equilibrium exists. More generally, the Nash equilibrium, which assumes that each spouse takes the decision of the other spouse as given, is not always equivalent to methods in which the choice of each spouse takes into account the possible responses of the other spouse.

V. Data and Estimation.

Estimation in this study is based on data from the Health and Retirement Study (HRS), a survey of roughly 7600 households with individuals born from 1931 to 1941.[9] The study uses observations from households of married couples for whom both spouses appear to have career jobs from which retirement would be a meaningful concept. Details of the sample restrictions, as well as the construction of the variables used in the estimation, are described in Appendix 1.

[9] The HRS is conducted by the University of Michigan under a grant from the National Institute on Aging (grant number U01AG009740).

Both spouses in HRS households were interviewed separately. This is true even if one spouse was in the eligible birth cohorts but the other was not. The survey started in 1992, at which time these individuals were between 51 and 61 years old, and the same households were re-interviewed every two years thereafter. This study uses the retirement data through the first six waves of the survey, by which time the eligible birth cohorts would have been 61 to 71 years of age. By age 61 half of the individuals would have retired, and virtually all by age 71, so this period includes the overwhelming majority of retirements for the sample individuals.

The HRS also has two important supplements, which are available on a restricted basis. First, Social Security earnings records are attached for about 75 percent of the sample, allowing fairly precise estimates of Social Security earnings and benefits for this part of the sample. Secondly, for respondents who indicated that they had pensions, the survey obtained and coded the summary pension documents from the employers of about two thirds of those in the sample with a pension on their current job. This enables a much more precise determination of the retirement incentives of pensions than is normally obtainable from the respondents themselves.

The estimation method is the generalized method of simulated moments (GMM). In this procedure, a group of moments is gathered into a column vector $\mathbf{m}$. These moments are generally the difference between some observed statistic, such as the percentage retired as of a specific age, and the percentage that is simulated for the sample using specified values of the parameter. In general, these moments come from an asymptotically normal distribution with a mean value of zero. The estimation procedure seeks the parameter value which minimize $q = \mathbf{m'W^{-1}m}$, where $\mathbf{W} = \sum_{i=1}^{n} \mathbf{m_i m_i'}$. The $\mathbf{m_i}$ vectors are the moments of the individual observations, and the $\mathbf{W}$ matrix is essentially the observed variance-covariance matrix of the moments. Variances of the estimates are calculated from $\text{var}(\Theta) = [\mathbf{G'W^{-1}G}]^{-1}$, where $\mathbf{\Theta}$ is the vector of parameters and $\mathbf{G}$ is the derivative of the moments with respect to the parameters. If the model is correctly specified, $\mathbf{m}$ is distributed around zero, and q should have a χ^2 distribution with $\lambda - k$ degrees of freedom, where λ is the number of moments and k is the number of parameters estimated.

To construct the moments for a specific observation using a specific set of parameters, we need first to estimate the value of the time preference parameter ρ for the couple associated with that observation. To do this, we use the observed full-time retirement dates if the spouses have already retired, or the expected retirement dates if one or both spouses have not retired. If the retirement from full-time work was into partial retirement, a second retirement date from partial retirement into full retirement is established in the same way. These retirement dates fix the leisure parts of both utility functions as well as the complete earnings stream for the couple. What remains is to compute the consumption stream conditional on a value of ρ. The consumption model is computed using the usual backward induction method of dynamic programming models. In this model, the rate of return on assets is taken to be stochastic, using the observed distribution of asset returns. The model is then simulated to obtain an associated path of wealth, and during this simulation the realized asset returns are used. The calculated amount of wealth in a particular year, say 1992, can then be compared with the actual wealth observed in that year. The calculated wealth depends on the assumed value of ρ, and this parameter is adjusted up or down until the calculated wealth matches the actual wealth. In these calculations, observed wealth is taken to be the sum of financial, real estate, business assets, and non-pension retirement assets (e.g., IRA's).

The use of actual or expected retirement dates to calculate ρ avoids the necessity of using the values of the random effects (ε_h, ε_w, V_P^h, and V_P^w) in the calculations of ρ, since the random effects affect only the leisure terms in the utility functions. This approach requires a couple of approximations. First, the actual or expected retirement dates refer to situations where both spouses survive and do not tell us what would have happened if one of the spouses had died. In general, the retirement ages would be expected to depend on the survival experience of both spouses, since if one spouse dies the income stream of the surviving spouse will be altered. However, since the pre-retirement mortality rates are relatively low, and since the bulk of lifetime income will have been earned before the observations begin, we make the approximation that the retirement dates of one spouse do not depend on whether the other spouse survives. The second approximation is the approximation discussed in Appendix 2. The uncertainty in retirement dates in the complete model probably induces couples to hold more assets than if the retirement dates were certain, and even then the amount of wealth associated with observed retirement dates and asset returns is not completely fixed. However, the discussion in Appendix 2 suggests that this is probably not a serious problem. The primary purpose of calculating time preference rates is to distinguish households which place a substantial value on future utility from those for whom present utility is paramount, and this approximation should be sufficient for those purposes.

Once the time preference rate ρ is calculated, we make a random draw from the joint distribution of ε_h and ε_w, given the parameters σ_h, σ_w, and ρ_ε of the distribution. We also make random draws of V_P^h and V_P^w, given the parameters γ_0^h, γ_1^h, γ_0^w, and γ_1^w for those distributions. With these preference parameters, the path of retirement states for both spouses can be calculated using the methods discussed above, and the values associated with the various moments are tabulated. Another draw is made from the distribution of ε_h and ε_w. Another set of retirement states for both spouses is calculated, and the values associated with the various moments are updated. The process continues through a large number of simulations (10,000 per observation). The moments used in the estimation are calculated by comparing the observed tabulations to the simulated distributions. Once the model has been estimated, the calculation of the 10,000 simulations does not add appreciably to the time required, so we use this number in order to make the simulated moments close to the theoretical moments. However, in order to make the calculations feasible, new draws of V_P^h and V_P^w are not made for each of the 10,000 simulations per observation, but new draws of these values are made for each individual.

The moments used in the simulation are chosen to provide identifying information on the parameters. The moments used are as follows:

Description of moments	Number of moments
The percentage retired from full-time work at ages 54-66	
Husband	13
Wife	13
The percentage completely retired at ages 55, 58, 60, 62, and 65	
Husband	5
Wife	5
The percentage retired from full-time work at ages 55, 58, 60, 62, and 65 among families in the bottom third of potential earnings	
Husband	5
Wife	5
The percentage retired from full-time work at ages 55, 58, 60, 62, and 65 among families in the upper third of potential earnings	
Husband	5
Wife	5
The percentage retired from full-time work at ages 55, 58, 60, 62 and 65 among those in poor health	
Husband	5
Wife	5
The percentage of reversals where the respondent was working full-time after having been partially or fully retired in the previous interview	
Husband	5
Wife	5
The percentage of couples in each interview where both spouses	

were working full-time	6
The percentage of cases where one partner was working full-time in one of the later interviews and the other was retired in one of the earlier interviews	6
Total moments	88

The first 26 moments help to establish the overall pattern of retirements, as governed by β_o^h, β_1^h, β_o^w, β_1^w, σ_h, and σ_w. The next 10 moments reflect the amount of partial retirement by age, and hence help to establish the values of γ_o^h, γ_1^h, γ_o^w, and γ_1^w, which determine the distribution of the value of partial retirement. The next 20 moments help to establish the value of α. The higher the value of α, the later will be the retirement of high earning individuals relative to the retirement of lower earning individuals. The next 10 moments, which involve those individuals in poor health, help to determine β_2^h and β_2^w. The following 10 moments, which have to do with reversals, help to establish the values of ρ_h and ρ_w, which are the changes in leisure preferences after retirement which could induce a return to work. The next six moments, which relate to joint retirement, help to determine the values of β_3^h and β_3^w, which govern how much the leisure of one spouse is augmented by the presence of the other spouse. The final six moments, which measure how often one spouse retires early and the other spouse retires much later, help to establish the value of ρ_ε, the correlation between ε_h and ε_w. High values of ρ_ε should mean that cases where one spouse retires early and the other much later occur less often than if the value of ρ_ε is relatively low.

The results of the parameter estimates are given in Table 1. Most if not all of the parameter estimates are significant at conventional significance levels.[10] For 70 degrees of

[10] If the observed moments do not match the moments generated by the model very well, the q value will be high and the prob value will be low, indicating that the model is not doing a

freedom, the 5% critical value of the χ^2 distribution is 90.53 and the 1% critical value is 100.43. The estimated value of 65.09 is well below these critical values, which indicates that there is little evidence that the model does not fit the model well, at least in the dimensions measured by the moments previously listed. Another way of looking at this is that the probability value of this q value is 0.64, which suggests that by chance, the q value would be higher than the estimated amount more than half of the time.[11]

Since this model is a combination of two of our previous models, we can see how different these estimates are from those in the previous models. One of the models (Gustman and Steinmeier, 2008) was a model of joint retirement decisions for both spouses within a family, but that model did not include stochastic returns or any partial retirement. The other model was a

---

very good job of explaining the observed data. The degrees of freedom for the χ^2 distribution is equal to the number of moments minus the number of parameters estimations, which for the current estimation gives a value of 70.

[11] Other studies have also simplified the estimation of the backward induction process. A notable approach that would approximate the solution to the model is suggested by Keane and Wolpin (1997). (The description of the Keane-Wolpin method is taken from the Appendix in their article in the June1997 JPE.) Unfortunately, their approximations do not appear to be applicable to our model. They make use of the fact that there is exactly one error term corresponding to each potential choice, which is to say that the values of the choices are specified as $V_m = f_m(S) + \varepsilon_m$, where m is a choice and S is the set of state variables. Our error terms enter the model in a more structural (and we would argue more realistic) way. While it may be possible to use the Monte-Carlo methods for simulating expected values, it is not possible to use their method for computing $E[\max(V_m)]$ from a transformation of more easily computed $E[V_m]$'s. The reason that their $E[V_m]$'s are easier to compute is that the expectation has to be integrated only over the single ε_m, (in fact, if $E[\varepsilon_m] = 0$, $E(V_m] = f_m(S)$, with no need for integration), while $E[\max(V_m)]$ has to be multiply integrated over all the ε_m 's. Thus, once they estimate the relationship between the $E[\max(V_m)]$'s and the $E[V_m]$'s from a small subsample of the possible state combinations, they can calculate the $E[\max(V_m)]$'s for the remaining state combinations without doing any integration at all! The cost is a relatively more restricted error structure tying the errors to the choices in a one-to-one fashion. The problem for our model is that since there is no strict correspondence between the potential choices and the error terms, the computation of the m $E[V_m]$'s is as complicated as the computation of $E[\max(V_m)]$.

model which included stochastic returns and a mechanism for varying leisure preferences after retirement, but took the wife's labor supply and its associated income as exogenous to the husband's retirement decision.

In comparing the current results with the previous family retirement model, the coefficient of the age variable for the husband is considerably lower in the present model as opposed to the previous model. Since this coefficient to some extent reflects the scope of economic circumstances to affect retirement, one would expect that the husband's retirement decision would be more sensitive to economic circumstances in the current model. For the wife, however, the coefficient of the age variable is very similar between the two models.

The coefficient of the spouse retired variable is about the same for the husbands between the two sets of estimates, but the coefficient of this variable is almost twice as high for the wife in the current set of estimates as opposed to the previous results. Nevertheless, the husband's coefficient is almost two and a half times as great as that for the wife, which leaves unaltered the conclusion from previous work that the husband's utility of retirement is much more sensitive to the presence of the wife than the other way around. Another dimension of joint retirement is the correlation of retirement preferences, as measured by the parameter ρ_{ε}. That parameter is considerably higher in the current estimates than the previous estimates, suggesting that correlated preferences may play a greater role. However, the previous estimate of this parameter was not very precise, and the confidence interval for that estimate is very close to including the current estimate. As a result, it is unclear that this difference in estimates is meaningful.

The coefficients of the health variables are slightly higher in the current model than they were in the previous family retirement estimates, both for the husband and for the wife. It is not completely clear why this should be so, but again the confidence intervals for the previous estimates include the values estimated for the current model. The standard deviation of the leisure preference term, which to a large extent governs the overall spread of the retirement distribution, is very close between the two estimates for the husbands, but the value for the wife is somewhat higher in the current model relative to the previous model.

In comparing the current model with the previous stochastic model (Gustman and Steinmeier, 2006), the most notable difference is that the leisure preferences change much less after retirement in the current model than in the previous model. This is reflected in the autocorrelation parameter ρ_h, which measures how autocorrelated the preferences after retirement are to the preferences prior to retirement. To a large degree, this parameter reflects how often individuals return to work after a period of retirement. As indicated earlier in this report, the mechanism by which the retirement decisions are made for couples introduces some degree of returns to work, even if there were no other mechanism for inducing returns to work. This means that the autocorrelation parameter ρ_h does not have to account for all the returns to work, but only a part of them. As a result, it is not terribly surprising that the autocorrelation parameter is higher in the current model than it was in the previous model, which assumed that the spouse's retirement was exogenous.

VI. Simulations.

Having obtained estimates for the parameters of the model, we now turn to simulations with the model. First we will present results from a simulation using the actual budget sets for the individuals in the sample. We refer to this simulation as the "base" simulation. We will then present simulations in which the wife is not in the labor force, to see what effect the wife's work has on the husband's retirement, and simulations in which partial retirement is not permitted.

Table 2 shows the basic retirement outcomes from the base simulation. For both the husbands and the wives, the last three rows indicate the percentages of the sample at each age who were working full-time, partially retired, or fully retired. As expected, the percentage working full-time falls with age and the percentage partially retired increases with age. For husbands, by age 62 less than half of the sample is still working full time, while for the wives the percentage working full time falls to less than half at age 60. Both husbands and wives have substantial numbers partially retired. For the husbands, the percentage partially retired grows steady throughout the age range illustrated until it reaches almost 20 percent at age 67. For wives, partial retirement grows at earlier ages, reaches a peak of around 17 percent at age 62, and declines thereafter.

The first two columns of the table are what may be referred to as “pseudo” retirement. These are merely the differences between the figures at adjacent ages in the last three rows. For instance, 32.1 percent of husbands were fully retired at age 61 and 42.8 percent at age 62. The difference of 10.7 percent is taken as the pseudo-retirements from work at age 62. It is really a net result of individuals completely retiring at that age less the individuals who had been retired at 61 but who returned to work at age 62. In any case, the prominent feature of these numbers is the spike of retirements at age 62, both for the husbands and to a lesser extent for the wives. The magnitude of the spike is approximately correct relative to the raw data for the husbands but is a little short for the wives. The model is a little less successful in capturing the secondary spike at age 65 for the husbands and not very successful at capturing the secondary spike for the wives. In any case, the current model is much more successful in capturing the spikes than was our previous family model, in which the spikes were substantially muted. As mentioned before, this probably has to do with the much lower estimated coefficent for the age variable for the husband, which makes the husband's retirement behavior more sensitive to economic incentives in the current model.

Table 3 decomposes both the full-time work figures and the partial retirement figures in the last three columns of the previous table. The full-time work percentages are decomposed into two parts, one for individuals who are still working in career jobs and who have never previously either fully or partially retired and the other for individuals who previously fully or partially retired and who have returned to full-time work. For both husbands and wives, the percentage of the sample who have returned to full-time work rises gradually up until about age 60, peaks at around 9 percentage points in the early 60's, and tails off thereafter. At its peak at age 61, more than 16 percent (9.4 / 57.2) of full-time husbands are individuals who have returned to full-time work after a period of full or partial retirement. For wives, the figure is even higher at 21 percent (8.8 / 42.5). Although the numbers of full-time workers decline both in the career jobs and in the return jobs after age 61, the percentage of full-time workers who have been previously retired continues to grow until it reaches 28 percent for husbands at age 67 and 30 percent for wives at the same age.

For partial retirement, the decomposition is between those who have partially retired from full-time work without going through a period of full retirement and those who have been fully retired at some previous point. From the table, it is clear that the majority of part-time work is done by individuals who are on a traditional path of moving from full-time work to partial retirement and then to full retirement, but part-time work by individuals who have previously been fully retired is not insignificant. For the husbands, a little over a third of part-time workers in their mid to late 50's have been previously retired. This figure increases during the early 60's until by age 67 around half of the part-time husbands have previously been fully retired. The same trend occurs for the wives, though the percentages of previously retired part-time workers is somewhat lower, particularly in the mid 50's.

Table 3 indicates the percentage of individuals who had returned to full-time work from partial or full retirement and the percentage who had returned to part-time work after full retirement, but it does not indicate the frequency with which husbands and wives go through at least one period of transitioning from a state of greater retirement to a state of less retirement. Table 4 provides this information. For both husbands and wives, around a third returned to full-time work at some point in time after a period of partial or full retirement, with the percentage being slightly higher among wives than among husbands. The percent returning to part-time work after a period of full retirement was around a quarter for both husbands and wives. Since these are not mutually exclusive categories, they cannot simply be added up since some individuals will go through both transitions. Accordingly, the last row of the table suggests that a little over two fifths of husbands go through a period of transition from greater retirement to less retirement, and almost half of the wives do so.

Table 5 addresses one of the issues particular to a model of the retirement behavior of couples, namely, the frequency with which both spouses retire at the same time. The figures in the top part of the table give the frequencies for which the husband retires before the wife, and the figures in the bottom part of the table give the frequencies for which the wife retires before the husband. Despite the fact that on average the husband is a couple of years older than the wife, the husband retires first in only about half the cases, while the wife retires first in around 36 percent of the cases. The remaining 13.5 percent of the cases are instances where both spouses retire in the same year. Note that this is around 10 to 11 percentage points higher than the surrounding figures, which measure the instances where the husband retires a year or two before the wife, or the wife retires a year or two before the husband. This is almost surely a result of the spouse retirement variables in the model, which increase the value of leisure of one spouse if the other spouse is also retired. The correlation in leisure preferences would also lead to an increased tendency for spouses to both retire early or both retire late, but a correlation in leisure preferences should lead to more of a smooth hump in the joint retirement distribution as opposed to a spike, as found in studies of joint retirement outcomes.

Table 6 reports on another distribution of interest, namely, the distribution of time preferences. Not quite half of the households have a time preference rate of less than 5%, while around a third have time preference rates of 50% or higher. The latter group essentially has no financial assets other than forced savings, meaning that they are effectively consuming all of their available income. Only around a sixth of the households fall in the middle, with time preference rates between 5% and 50%, and most of those are in the bracket from 5% to 10%. The clear implication of this result is that models which assume that all households have a uniform time preference rate will almost certainly yield very misleading results for a third of the households, but models which allow two mass points for the time preference distribution, one at a relatively low time preference rate and the other at a relatively high time preference rate, may be a lot closer to the truth. Even here, though, as suggested by Figure 2 in the appendix, there is a very substantial difference between the amount of wealth that would be accumulated by an individual with a 1% time preference rate relative to an individual with a 3% time preference rate, and a model with only two mass points of time preference is unlikely to reflect this difference.

We now turn to two simulations which address a couple of interesting questions. The first has to do with the effect on husband's retirement of having the wife in the labor force. To shed some light on this subject, we examine a simulation in which the wife is assumed to be out of the labor force. The results of this simulation are reported in Table 7. The figures in the table are differences between the simulation holding the wife out of the labor force and the base simulation. As might be expected, excluding the wife from the labor force leads to a substantial increase in full-time work by the husband. From the husband's viewpoint, the income provided by the wife is largely an increase in family resources not associated with his work. Taking away this wealth leads to a wealth effect, which should have the impact of reducing both leisure and consumption. The reduction in leisure is mainly reflected in increased full-time work effort, amounting to over 13 percentage points in the early 60's. Leisure itself, as reflected in the percentage of the sample being fully retired, drops by a comparable amount, again reaching around 13 percent in the early 60's. During the 50's and early 60's, the percentage of the sample who are partially retired is reduced, but after that age partial retirement is increased by the exclusion of women from the labor force. To put it another way, in the early years full-time work is increased substantially at the expense of both partial retirement and full retirement, but in later years full retirement is increased substantially at the expense of both full-time work and partial retirement. In any case, it seems clear that all else being equal, the flow of wives into the labor force in the last few decades has probably reduced the amount of work that the husbands would have done otherwise.

A second simulation looks to see what would be the effect if work were a full time or nothing decision. This is meant to address the issue of the effects of the possibility of part-time work on overall work effort. One of the arguments in favor of expanding the opportunities for part-time work is that increased part-time work would expand overall work effort by allowing older Americans to work into years when health or other issues might make it impossible or difficult to work full-time. Table 8 presents results from a simulation where the option to work part-time is effectively eliminated and the respondents would be forced to choose only between full-time work and full retirement. The figures in this table are the differences between the percentage of individuals in the stated work/retirement category in this simulation and the comparable percentage of individuals in the base simulation. For instance, the first entry in the table, 1.9 percent, means that at age 55, the number of husbands working full-time would be 1.9 percentage points higher if part-time work were eliminated as a possibility. Note that the sum of the two columns for each gender is the percent of individuals partially retired in Table 2. That is, of the 3.4 percent of husbands who were partially retired in the base simulation, if partial retirement were eliminated 1.9 percent of them would switch to full-time work and 1.5 percent of them would switch to being fully retired.

In general, eliminating partial retirement would lead to increases both in the number working full-time, which would increase overall work effort, and in the number fully retired, which would reduce overall work effort. For husbands, the increase in full-time work outweighs the increase in full retirement up until about age 62, which means that up until that age overall work effort would increase if partial retirement were eliminated. After age 62, for the husbands the elimination of partial retirement would cause the increase in full retirement to outweigh the increase in full-time work, which would imply a reduction in overall work effort for this group. For the entire 55 to 67 age range illustrated in the table, the increase in full retirement outweighs the increase in full-time work, but only by a small amount.

Making partial retirement more generally available or more attractive, of course, would have the opposite effect. One might expect that the husbands would exhibit a slight increase in overall work effort as individuals continued to work in partial retirement beyond the date that they would otherwise retire. This effect, however, is largely offset because with the increased availability of partial retirement, individuals would be induced to leave full-time work for partial retirement earlier than they would have left full-time work in the absence of partial retirement. The simulation suggests that this offsetting effect is sufficient to largely negate any effort to promote opportunities for partial retirement as a means to increase overall work. The goal of increasing work effort, in turn, is often promoted as a means of stabilizing Social Security's finances or for other ends.

The pattern for the wives is similar to that for the husbands, although there are some important differences. Before age 62, partial retirement among the wives is more common than among husbands, so the increases in full-time work and in full retirement are both greater than for the husbands. It is still true, however, that before age 62 the increase in full-time work from eliminating partial retirement outweighs the increase in full retirement. After age 62, the reverse is true: the figures are less for the wives than for the husbands, and the increase in the percentage fully retired is larger than the increase in the percentage working full-time. Again, over the entire age range, the increase in the percentage fully retired outweighs the increase in the percentage working full-time, but only by a small amount. Therefore, the implications about the goals of increasing opportunities for partial retirement among the wives are essentially the same as the implications for the husbands outlined in the previous paragraph.

VII. Conclusions.

This paper has integrated many features of retirement models into a single framework. This has made it possible to utilize the full set of labor market information provided by the Health and Retirement Study, including survey responses, pension plan descriptions and Social

Security earnings data provided individually for husbands and wives. The integrated model is much richer than previously specified models of family retirement, allowing each spouse to retire and unretire, transitioning among the states of full-time work, partial retirement and full retirement. It explains in much greater detail the effects of interdependence in the decisions made by each spouse, including clustering of retirements by husbands and wives, while at the same time allowing for forward looking behavior, explaining saving at the family level, incorporating the nonlinear budget constraints from still dominant defined benefit pensions and Social Security, and allowing for the effects of exogenous shocks to market opportunities, health and asset returns.

As we found in our past research, increasing the richness of the model allows us to address phenomena that otherwise cannot be explained by conventional models of retirement. At the family level we are able to isolate the key role of heterogeneity in time preference, allowing the model to explain the wide differences in wealth accumulated by families with similar earnings opportunities. The retirement hazard exhibits the important spike in retirement at age 62 in the face of an actuarially fair Social Security system, captures the extent of partial retirement by each spouse, reproduces the flow from states of greater to lesser retirement, and relates each of the flows for one spouse to the decisions made by the other. We also allow each spouse to have heterogeneous preferences for both full-time and part-time work.

The theoretical discussion increases understanding of the wide variety of situations that families face when approaching the retirement decision. It illustrates how choices focused on one spouse's welfare will take account of the welfare and independent reactions of their mate, incorporating the roles of different preferences and different market opportunities. Allowing for the variety of circumstances facing different families, the theoretical framework incorporates the many different situations facing different couples into a unified framework, and directly into the estimation. We also have shown why, when both spouses realize that their own choices may affect the choices of the other spouse, a solution method based on backward induction is superior to a method based on a Nash equilibrium, including a finding that a solution based on backward induction may provide plausible behavioral predictions when the Nash equilibrium criteria fall silent in attempting to predict the optimal solution.

We have found some surprising implications of the model. For example, retirement reversals can occur without any changes in preferences (other than the increasing desire for retirement over time) or any changes in effective compensation. Also, contrary to the usual implication that a pure wealth effect should cause a decline in labor supply, higher wealth at the beginning of the period, all other things being equal, may cause the work effort of one of the spouses to increase, even if the work effort of the other spouse remains the same.

Simpler specifications cannot simultaneously explain heterogeneity in wealth; liquidity preference and retirement spikes when benefits become available; the high rate of return from states of lesser work to states of greater work despite the assumption of forward looking behavior, and in the absence of changes either in preferences or in market opportunities; the different sequence of retirements by husbands and wives; and why in some cases increased wealth may reduce retirements.

Comparisons of the integrated model of retirement behavior with models that either simplify the retirement options and dynamic, stochastic nature of decision making, or simplify or ignore the importance of interdependence of retirement decision making in a family setting, have shown the value of the estimates obtained with the integrated model. While preserving key findings, such as the conclusion from previous work that the husband’s utility of retirement is much more sensitive to the presence of the wife than the other way around, the integrated model can explain the key features of retirement outcomes that a simple specification of the family model failed to explain, including the ability to simulate with accuracy the spike in retirements at age 62. There also are substantial differences in key coefficients between the integrated more constrained specifications, such as in the coefficient on age, which governs the sensitivity of the retirement response to incentives from policies, as well as changes in preferences after retirement, which are smaller in the integrated model.

We find interesting effects from counter factual experiments. The flow of wives into the labor force in the last few decades has probably reduced the amount of work that the husbands would have done otherwise. In the early years full-time work is increased substantially at the expense of both partial retirement and full retirement, but in later years full retirement is increased substantially at the expense of both full-time work and partial retirement. In another experiment, we find that eliminating partial retirement would lead to increases both in the number working full-time, which would increase overall work effort, and in the number fully retired, which would reduce overall work effort. Altogether, the increase in full retirement outweighs the increase in full-time work, but the difference is small. The simulation suggests that this offsetting effect is sufficient to largely negate any effort to promote opportunities for partial retirement as a means to increase overall work.

Although there is a great deal of effort required to estimate and simulate with a structural model that incorporates many dimensions of behavior, it appears to be worth the effort. One reward is an increased understanding of behavior, providing insight into many dimensions of retirement behavior, and behavior at the family level, that is not otherwise available from more simplified approaches. Another reward is a clearer picture of the likely effects of events, where simplified approaches either fall silent or, worse, may provide misleading guidance for policies. In previous studies we have examined the effects of various policies on retirement. In a longer report to the National Institute on Aging, we also apply the present model to analyze the effects of delaying the Social Security early entitlement age and the effects of substituting defined contribution for defined benefit plans. The implications of these and related changes for saving and retirement simply cannot be understood when behavior must be analyzed with a less structural approach. Since the Health and Retirement Study has been designed to support these complex analyses, there is additional effort required to design and estimate the basic behavioral models incorporating all the relevant dimensions of behavior, but additional effort is not required to collect the requisite data.

Table 1
Parameter Estimates

Parameter Symbol	Description	Estimated Value	Absolute t-Statistic
	Husband's Parameters		
	Parameters for the Value of Leisure		
β_o^h	Constant	-10.37	106.00
β_1^h	Age	0.12	5.43
β_2^h	Own Health	6.00	5.84
β_3^h	Spouse Retired	2.63	2.20
σ_h	Standard Deviation of ε_h	4.75	7.44
ρ_h	Correlation of ε_h After Retirement	0.86	14.01
	Parameters for the Value of Partial Retirement		
γ_o^h	Constant	-0.68	1.47
γ_1^h	Age	0.24	2.58
	Wife's Parameters		
	Parameters for the Value of Leisure		
β_o^w	Constant	-10.35	49.73
β_1^w	Age	0.19	4.16
β_2^w	Own Health	4.25	3.77
β_3^w	Spouse Retired	1.18	0.75
σ_w	Standard Deviation of ε_w	5.04	8.42
ρ_w	Correlation of ε_w After Retirement	0.87	18.12
	Parameters for the Value of Partial Retirement		
γ_o^w	Constant	-6.14	4.02
γ_1^w	Age	0.06	0.32
ρ_ε	Correlation Between ε_h and ε_w	0.67	1.70
α	Consumption Parameter	-0.47	4.80

Number of observations 851
Q Value 65.09

Table 2
Retirement States in Base Simulation, By Age and Gender

	Percent Pseudo-Retiring		Percent		
Age	From FT Work	From All Work	in FT Work	Partially Retired	Fully Retired
			Husbands		
55	3.5	2.7	87.4	3.4	9.2
56	2.5	1.8	84.9	4.1	11.0
57	4.3	3.0	80.6	5.4	14.0
58	4.2	3.4	76.4	6.1	17.5
59	5.7	4.3	70.7	7.5	21.8
60	7.2	5.4	63.5	9.3	27.2
61	6.4	4.9	57.2	10.8	32.1
62	14.7	10.7	42.5	14.7	42.8
63	6.9	5.4	35.6	16.2	48.2
64	6.8	5.7	28.8	17.3	53.9
65	7.9	7.0	20.8	18.3	60.9
66	5.3	5.0	15.5	18.6	65.9
67	3.9	3.5	11.6	19.0	69.4
			Wives		
55	4.6	3.2	76.0	6.9	17.1
56	4.3	3.1	71.7	8.1	20.2
57	5.0	3.5	66.7	9.6	23.7
58	5.6	4.1	61.1	11.1	27.8
59	6.1	4.4	55.1	12.8	32.1
60	6.6	5.4	48.5	14.0	37.5
61	6.0	5.0	42.5	14.9	42.5
62	10.5	8.8	32.1	16.7	51.3
63	6.3	6.9	25.7	16.1	58.2
64	5.1	6.3	20.6	14.9	64.5
65	4.8	6.2	15.8	13.5	70.7
66	3.6	5.4	12.2	11.7	76.1
67	3.1	4.5	9.2	10.2	80.6

Table 3

Returns to Work in Base Simulation, By Age and Gender

Age	Percent in Main Job	Percent to FT Work after Retiring	Percent in PT Work after FT Work	Percent in PT Work after Retiring
		Husbands		
55	82.3	5.1	2.3	1.1
56	78.8	6.1	2.6	1.5
57	73.8	6.8	3.5	1.9
58	68.5	7.9	3.9	2.2
59	62.2	8.5	4.8	2.7
60	54.5	9.0	5.9	3.4
61	47.7	9.4	6.7	4.1
62	35.9	6.6	9.9	4.8
63	28.9	6.7	10.0	6.2
64	22.9	5.9	10.0	7.3
65	16.1	4.8	10.2	8.1
66	11.6	4.0	9.7	8.9
67	8.4	3.2	9.4	9.6
		Wives		
55	69.8	6.1	4.9	2.0
56	64.7	7.0	5.8	2.3
57	59.0	7.7	6.7	2.9
58	52.9	8.3	7.6	3.5
59	46.5	8.6	8.5	4.3
60	39.7	8.8	9.0	5.0
61	33.8	8.8	9.3	5.6
62	25.8	6.3	11.0	5.7
63	20.8	5.0	9.8	6.3
64	16.3	4.3	8.6	6.3
65	11.9	3.9	7.6	5.9
66	8.7	3.5	6.3	5.4
67	6.3	2.8	5.3	4.9

Table 4
Respondent Returning to Work in Base Simulation, By Gender

	Husbands	Wives
Percent returning to full time work after full or partial retirement	29.0	34.4
Percent returning to part time work after full retirement	23.7	25.7
Percent returning to full time work after full or partial retirement or returning to part time work after full retirement	43.2	48.5

Table 5
Distribution of Differences in Retirement
Years in Base Simulation

	Difference in Retirement Dates (Years)	Percent of Households
	10+	19.3
	9	3.1
	8	3.3
	7	3.4
Husband	6	3.5
Retires	5	3.5
First	4	3.7
	3	3.6
	2	3.6
	1	3.5
	0	13.5
	1	2.2
	2	2.8
	3	2.9
Wife	4	3.2
Retires	5	3.5
first	6	3.3
	7	3.4
	8	3.0
	9	2.7
	10+	9.1

Note: The retirement date is the year the individual first retired from full-time work.

Table 6
Distribution of Time Preference
Rates in Base Simulation

Time Preference Rate	Percent of Households
0-5%	44.5
5-10%	12.6
10-25%	8.0
25-50%	1.9
>50%	33.0

Table 7
Simulated Effect on Husbands Retirement Status
of Eliminating Wife's Labor Force Participation

Age	Change in Percent in Full Time Work	Partially Retired	Fully Retired
		Husbands	
55	7.8	-2.4	-5.5
56	9.3	-2.5	-6.7
57	11.1	-3.0	-8.1
58	12.4	-2.8	-9.5
59	12.8	-2.5	-10.3
60	13.1	-2.1	-11.0
61	13.3	-1.5	-11.9
62	9.7	3.5	-13.2
63	8.9	4.0	-12.9
64	7.6	4.5	-12.1
65	6.0	4.3	-10.4
66	5.0	4.6	-9.6
67	4.1	4.8	-8.9

Table 8
Simulated Effect of Eliminating Partial Retirement

Age	Change in Percent in Full Time Work	Fully Retired	Change in Percent in Full Time Work	Fully Retired
	Husbands		Wives	
55	1.9	1.5	3.5	3.4
56	2.4	1.7	4.2	3.9
57	3.4	2.0	5.3	4.3
58	4.2	1.9	6.3	4.8
59	5.0	2.5	7.3	5.5
60	6.1	3.2	8.0	6.0
61	7.0	3.7	8.0	7.0
62	7.8	6.9	6.4	10.2
63	7.9	8.3	6.1	10.0
64	7.3	10.0	5.1	9.8
65	6.9	11.4	4.4	9.1
66	5.9	12.7	3.6	8.1
67	4.9	14.1	2.6	7.6

References

Benitez-Silva, Hugo, Moshe Buchinsky, Hiu Man Chan, John Rust and Sophia Cheidvasser. 1999. "An Empirical Analysis of the Social Security Disability Application, Appeal and Award Process". *Labour Economics* 6: 147-178.

Benitez-Silva, Hugo, Moshe Buchinsky and John Rust. (2004) "How Large is the Bias in Self Reported Disability?" *Journal of Applied Econometrics* , 19(6): 649-670.

Berkovec, James C. and Steven Stern. 1991. "Job Exit Behavior of Older Men". *Econometrica* 59(1): 189-210.

Blau, David. 1997. "Social Security and the Labor Supply of Older Married Couples". *Labour Economics* 4(4): 373-418.

________________. 1998. "Labor Force Dynamics of Older Married Couples". *Journal of Labor Economics* 16(3): 525-629.

Blau, David M. and Donna B. Gilleskie. 2006. "Health Insurance and Retirement of Married Couples". *Journal of Applied Econometrics* 21: 935–953.

Blundell, Ricard, P.A. Chiappori, T. Magnac and C. Meghir (2001). "Collective Labor Supply with Nonparticipation". IFS Working Paper 01/19.

Bound, J., T.R. Stinebrickner and T. Waidmann. 2006. "Health, Economic Resources and the Work Decisions of Older Men". Paper presented at NBER Summer Institute on Aging.

Bourguignon, Francois, Martin Browning and Pierre-Andre Chiappori. 2006. "Efficient Intra-household Allocations and Distribution Factors: Implicatins and Identification". Working paper 2006-02. Centre for Applied Microeconmetrics, Department of Economics, University of Copenhagen.

Cherchye, Laurens, Bram De Rock and Frederic Vermeulen. 2009. "Opening the Black Box of Intrahousehold Decision Making: Theory and Nonparametric Empirical Tests of General Collective Consumption Models". *Journal of Political Economy* 117(6): 1074-1104.

Coile, Courtney C. 2004. "Retirement Incentives and Couples Retirement Decisions". *Topics in Economic Analysis and Policy* 4(1): article 17.

Coile, Courtney and Jonathan Gruber. 2000. "Social Security and Retirement". NBER Working Paper 7830.

Cutler, David M., Jeffrey B. Liebman and Seamus Smyth. 2006. "How Fast Should the Social Security Eligibility Age Rise?" National Bureau of Economic Research, Retirement Research Center Working Paper.

French, Eric. 2005. "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior". *Review of Economic Studies* 72: 395-427.

French, Eric and John Bailey Jones. 2004. “The Effects of Health Insurance and Self-Insurance on Retirement Behavior”. Chicago: Federal Reserve Bank of Chicago, Working Paper 2004-12.

Gustman, Alan L. and Thomas L. Steinmeier. "Modeling the Retirement Process for Policy Evaluation and Research." *Monthly Labor Review* 107(7): 26-33.

________________. 1986a. "A Structural Retirement Model." *Econometrica,* 54(3): 555-584.

________________. 1999. "Effects of Pensions on Savings: Analysis with Data From the Health and Retirement Study". In *Carnegie-Rochester Conference Series on Public Policy* 50: 271-326.

________________. 2000. "Retirement in a Family Context: A Structural Model for Husbands and Wives". *Journal of Labor Economics* 18(3): 503-545.

________________. 2004. “Social Security, Pensions and Retirement Behavior Within the Family”. *Journal of Applied Econometrics* , 19(6): 723-738.

________________. 2005. “The Social Security Early Retirement Age In A Structural Model of Retirement and Wealth.” *Journal of Public Economics* 89(2-3): 441-463.

________________. 2006. “Social Security and Retirement Dynamics”. Michigan Retirement Research Center, Working Paper 2006-121.

________________. 2008. “How Does Modeling of Retirement Decisions at the Family Level Affect Estimates of the Impact of Social Security Policies on Retirement?” University of Michigan, Retirement Research Center, Report on Project UM08-03.

________________. 2009. “How Changes in Social Security Affect Recent Retirement Trends.” *Research on Aging* 31(2): 261–290.

Gustman, Alan L., Thomas L. Steinmeier and Nahid Tabatabai. 2010a. *Pensions in the Health and Retirement Study*. Cambridge: Harvard University Press.

_______________. 2010b. “What the Stock Market Decline Means for the Financial Security and Retirement Choices of the Near-Retirement Population”. *Journal of Economic Perspectives*. Winter.

Hurd, Michael D. 1990. “The Joint Retirement Decision of Husbands and Wives”. In *Issues in the Economics of Aging*, edited by David A. Wise. Chicago: University of Chicago Press: 231-254.

Ibbotson Associates. 2002. *Valuation Edition 2002 Yearbook*. Chicago: Ibbotson Associates.

Keane, Michael P. and Kenneth Wolpin. 1997. “The Career Decisions of Young Men. *Journal of Political Economy* 105(3): 473-522.

Krueger, Alan B., and Pischke, Jorn-Steffen. "The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation." *Journal of Labor Economics* 10 (4) (Oct. 1992): 412 – 437.

Lundberg, Shelly. 1999. "Family Bargaining and Retirement Behavior," in *Behavioral Dimensions of Retirement Economics*, ed. Henry Aaron, Russell Sage/Brookings: 253-272.

Maestas, Nicole. 2001. "Labor, Love & Leisure: Complementarity and the Timing of Retirement by Working Couples". Berkeley: Xerox.

________________. 2010 (forthcoming). "Back to Work: Expectations and Realizations of Work After Retirement". *Journal of Human Resources*.

Michaud, Pierre-Carl. 2003. "Joint Labour Supply Dynamics of Older Couples". IZA Discussion Paper 832.

Rust, John and Christoper Phelan. 1997. "How Social Security and Medicare Affect Retirement". *Econometrica* 65(4):781-831.

Scholz, Seshadri and Khitatrakun. 2006. "Are Americans Saving Optimally for Retirement". *Journal of Political Economy* 114(4): 607-643.

Van der Klaauw, Wilbert and Kenneth I. Wolpin. 2006. "Social Security and the Retirement and Savings Behavior of Low Income Households". *Journal of Econometrics* 145(1-2): 21-42.

Venti, Steven F. And David A. Wise. 1999. "Lifetime Earnings, Saving Choices, and Wealth at Retirement." In James P. Smith and Robert J. Willis, Editors, *Wealth, Work and Health: Innovations in Measurement in the Social Sciences*. Ann Arbor, University of Michigan Press, pp. 87-120.