AW
7/22/08
1:00 pm

How Does Modeling of Retirement Decisions at the Family Level Affect Estimates of the Impact of Social Security Policies on Retirement?

Alan L. Gustman

Dartmouth College and NBER

and

Thomas Steinmeier

Texas Tech University

July, 2008

The research reported herein was conducted pursuant to a grant from the U.S. Social Security Administration (SSA) to the Michigan Retirement Research Center, Grant No. UM08-03, with a subcontract to the National Bureau of Economic Research. The basic model of retirement at the family level on which this analysis is based was modified as part of our research under NIA grant 5R01AG 22956 (Gustman). The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, MRRC, NIA or the NBER.

How Does Modeling of Retirement Decisions at the Family Level Affect Estimates of the Impact of Social Security Policies on Retirement?

Alan L. Gustman and Thomas L. Steinmeier

Abstract:

This paper applies structural models of retirement and saving of two earner couples to explore the effects on retirement of two actuarially neutral policies, which we know from previous work can have a substantial effect on retirement if heterogeneity in time preference rates is allowed. The main question being investigated here is whether using a model that explicitly incorporates the retirement interactions of two working spouses yields different results from using a much simpler model that treats the retirement decisions of the second spouse as exogenous. The findings indicate that unless the question of interest is specifically related to joint retirement issues, the effects of the two actuarially neutral policies being investigated are roughly equal whichever model is estimated. The implication is that estimates of the effects in two earner households are not much compromised if simpler models that treat the second spouse's retirement as exogenous are used.

A second question explored in the paper is whether two earner and one earner households can be combined in the analysis. The effects of policy changes are clearly different for one earner and two earner households, but there is some evidence that the principal difference is due to the differing budget sets of the two groups. Though the estimated preference parameters are significantly different, the critical parameters governing responses to policy changes are similar. As a result, it seems plausible that unless the question being investigated involves looking at these two groups separately, the overall impact of the policy changes may be adequately assessed by combining the two groups.

A third question involves the magnitude of the effects for these two specific policy changes. Increasing the Social Security early entitlement age from 62 to 64 would reduce the level of retirement for husbands from two earner households by 4.4-4.6 percentage points at age 62, and by 5.1-5.7 percentage points for wives. In contrast, this policy change would induce husbands from one earner households to reduce the level of retirement by 10.2 percentage points at age 62. In a system of personal accounts, offering Social Security benefits as a lump sum instead of as an annuity would increase the level of retirement for husbands from two earner households by 7.1-8.1 percentage points at age 62 and by 8.9 percentage points for husbands in one earner households, and by 2.8-3.2 percentage points for wives in two earner households.

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

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

JEL numbers H55, J26, J14, J16, J32, E21, D31, D91, I3.

How Does Modeling of Retirement Decisions at the Family Level Affect Estimates of the Impact of Social Security Policies on Retirement?

Alan L. Gustman and Thomas L. Steinmeier

I. Introduction

When there is heterogeneity in time preference, policies that on average are actuarially neutral may nevertheless affect retirement behavior (Gustman and Steinmeier, forthcoming). Although for the average member of the population there are no kinks in the budget constraint that would encourage retirement at a particular age, actuarially neutral policies, such as increasing the Social Security early entitlement age from 62 to 64, will encourage those with high rates of time preference to delay retirements from age 62 to age 64 (Gustman and Steinmeier, 2005). For similar reasons, we find that heterogeneity in time preference accounts for the most prominent feature of the retirement hazard, the spike in retirements at age 62. This is despite the design of the Social Security benefit structure, which is roughly (but not exactly) actuarially neutral around age 62.

At the same time, studies of the retirement behavior of two earner couples suggest that their decisions may be strongly interdependent.[1] Younger wives accelerate their retirement to exit work closer to the time their older husbands retire, while their husbands delay their retirement to increase coordination.

These findings raise two questions. First, will the retirement responses of members of two earner households be less sensitive than in one earner households, or even insensitive to actuarially neutral policies? Second, are estimates of the effects of policies on the retirement of

[1] A few examples among many studies include Hurd (1990), Gustman and Steinmeier (2000, 2004b), Maestas (2001), Michaud (2003) and Blau and Gilleskie (2006).

members of two earner households sensitive to the complexity with which decision making in two earner households is modeled?

To answer these questions, this paper examines the effects of two policy changes on retirement and saving in two earner households. First it considers the effects of increasing the early Social Security entitlement age from 62 to 64. Second, starting with a regime of personal accounts, it considers the effects of changing from a system in which payouts take the form of mandatory annuities, to a system in which benefits are paid out as a lump sum.

Section II contains a further discussion of retirement decision making in two earner households, focusing on the question of how the interactions of spouse decision making may affect retirement responses to Social Security policies. The basic model of retirement decision making by two earner couples is presented in Section III. Section IV presents the estimates of that model and compares key parameters with those found for a sample of one earner couples. Section V explores the channels through which the retirement decisions of two earner couples interact and simulates the effects of changes in Social Security policies as we successively remove various features of the model governing the strength of the interactions of retirement decisions of two earner couples. This allows us to determine the extent to which estimates of responsiveness to actuarially neutral retirement policies are sensitive to interdependence in the family model. Section VI compares the effects on retirement in one and two earner households from adopting two actuarially fair policies. Section VII draws conclusions.

II. How may retirement responses to Social Security policies be affected by the interactions of decisions in two earner households?

There are a number of avenues through which retirement decisions of a spouse can affect retirement responses to incentives from retirement programs, including actuarially neutral policies. Consider a case where the husband's evaluation of retirement depends on his wife's retirement status. Assume, as it appears for many in the data (Gustman and Steinmeier, 2000), that the husband does not want to retire before his wife. Suppose the husband is two years older, and suppose further that the wife retires at age 62. In this case, the husband, who might otherwise want to retire at 62, delays his own retirement until age 64 in order not to retire before his wife. Now consider a policy change, such as increasing the early entitlement age to age 64, which is designed to encourage individuals to delay retirement from age 62 until age 64. If the wife were not involved, the husband would have retired at 62, and this policy change would clearly provide him with a substantial incentive to delay retirement. With the wife in the picture, however, he has already delayed retirement until age 64, and it is unlikely that the policy would have any further effect on the husband. In this example, the consideration of joint retirement mutes the response to incentives to policy changes designed to delay retirement.

Differences in preferences between one and two earner households may also lead to a difference in their responsiveness to retirement policies. The interactions of retirement decisions in two earner households are absent from one earner households. There also may be fundamental differences in time preference as each spouse consider the current and future welfare of their husband or wife. There may be fundamental differences in attitudes created by a lifetime of both spouses working outside the home, while unequally sharing responsibilities toward taking care of the house. At one extreme, when the husband stops working in a one earner household, the wife may continue to take care of the house. At another extreme, when the husband stops working in a two earner household, the wife may put the husband to work -- I worked two jobs my whole life, now it's your turn. This may mean that husbands in two earner households are less willing to retire. Other families may blend these two extremes. There are other possibilities. Two earner couples may have postponed time together over their lifetimes and may enjoy joint leisure more than one earner couples. For these or other reasons, there may be differences in time preference rates between the two sets of households which in turn would affect the responsiveness to the various policies.

Program rules also affect the opportunity set, creating a number of differences in retirement incentives between single earner and two earner households. There may be differences in the rate at which additional earnings increase the PIA. For instance, if single earners have a more continuous work history, then additional years of earnings are less likely to replace years of zero earnings in the PIA calculations. In addition, because single earner households do not have earnings income from the spouse, the reduction in income from one or another spouse retiring will be more severe in single earner relative to dual earner households. This may account for the greater responsiveness of single earner households to the increase in the early entitlement age. Rules governing spouse and survivor benefits create differences in retirement incentives between one and two earner households. Actuarial adjustments to Social Security are probably better for single than for dual earner husbands. Again due to the presence of spouse and survivor benefits, the overall returns to Social Security are probably better for single than for dual earner households. This also means, when considering actuarially neutral policies within a system of personal accounts, that a switch to personal accounts would increase wealth less for single than for dual earner households.

Our analysis will test the prediction that the joint pull of the retirement decisions among two earner couples, the dilution of the relative importance of any particular dollar incentive facing one of two earners, and the diminished importance of the liquidity constraints in some circumstances, are likely to leave two earner couples less responsive, or unresponsive, to retirement incentives from actuarially neutral or other policies.

III. A Model of Retirement Decision Making

A structural approach that includes both the decisions to retire and to save is required to address any questions about incentive effects of actuarially neutral policies. Although benefit adjustments that are actuarially neutral on average typically offer an increase in real benefits of six to eight percent for each year benefit receipt is postponed, for those with high time preference rates, this increase in future benefits is too low to compensate for postponing benefit receipt. In Gustman and Steinmeier (2005) we estimate that forty percent of the population has a time preference rate greater than ten percent, and for most of them, the time preference rate is considerably above ten percent. Accordingly, studies that ignore heterogeneity in time preference miss the implication that many seniors will decide that an actuarially fair increase in benefits is too small to compensate them adequately for postponing their retirement, so that actuarially neutral policies will speed their retirement.

In contrast to a structural approach, typical reduced form retirement equations do not allow for heterogeneity in time preference. As a result, retirement studies that rely on a simple regression of retirement on a pension or Social Security delta would fall silent when it came to addressing the retirement effects of actuarially neutral policies. Specifically, in the case of actuarially neutral policies, the pension or Social Security deltas measuring the effects of additional work on the present value of expected benefits are all zero. Consequently, a simple retirement equation would predict a retirement effect of zero from adopting an actuarially neutral policy. For similar reasons, a simple reduced form retirement equation cannot explain the spike in retirement at age 62, the Social Security early entitlement age.

A structural approach also allows us to specify and estimate the importance of each avenue through which retirement decisions interact within the household. We then can use these estimates to determine how each avenue of interaction affects the impact of particular policies on retirement.

Our approach to modeling retirement decision making at the family level differs from earlier models of family retirement in three ways. First, it jointly models the decision not only to retire by each spouse, but also to save. It therefore allows for the influence of heterogeneity of time preference in the decisions to save and to retire. Second, the model builds in an explicit measure of each spouse's preference for spending more time with their spouse during retirement. Third, it is solved using the backward induction method of dynamic programming, while allowing the spouses to choose alternatives that will increase their joint welfare.

We use a slightly modified version of a model of retirement decision making in two earner households from Gustman and Steinmeier (2004a).[2] The model has two utility functions, two decision makers, and one budget constraint. Utility of each spouse is a function of joint consumption and own labor supply over the lifetime. For the husband, the lifetime utility function is given by

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

where C is consumption and L is leisure. m is an indicator of 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 takes on a value of 0 if the

---

[2] There are three major modifications to our earlier model. First, the empirical work reported here interpolates labor market outcomes between biannual waves of the Health and Retirement Study. This leads to sharper spikes in retirement activity that in our previous working paper, which included only biannual observations. Second, additional moments are added to the maximization in this paper to clarify the nature of interdependence in the retirement decisions of the two spouses. Third, the model in this paper does not include indicators of vintage, which although included have not played a significant role in any of our previous studies.

husband is working (or has died) and a value of 1 if he is retired. The exponential expression preceding L is thus the utility value of retirement in period t. It consists of a standard linear form $X\beta$ and an individual effect ε which reflects the strength of the husband's preferences for retirement over work. The elements in X contain a constant, health status, age, and whether the spouse is retired. As age increases, work gradually becomes more onerous and retirement more desirable. When the utility value of retirement exceeds the utility of consumption from the income earned from work, retirement occurs. In this model, to keep the calculations more tractable, retirement is taken as irreversible.

The utility function of the wife is symmetric:

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

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:

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

Assets, which are assumed to be contrained to be non-negative, grow at the real interest rate r. The second term on the right side is the husband's earnings, and the fourth term is the wife's earnings. Potential earnings are assumed to be non-stochastic in this model. 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 experiences.

The interactions of the decisions of husbands and wives are described in the Appendix. Data are from the Health and Retirement Study (HRS), and include restricted Social Security and employer provided pension data. The first wave of the HRS was in 1992. Interviews are conducted every two years. The initial wave of the survey began with 12,654 respondents from households with at least one member who was 51 to 61 in 1992 (born from 1931 to 1941). Retirement is estimated based on behavior through 2002. In our estimates retirement is defined to be exit from full-time work, which is specified here to be working at least 30 hours per week and 1560 hours per year, or at least 25 hours per week and 1250 hours per year and self reported to be working full-time.

Estimation is based on the method of simulated moments. The following moments are used in the estimation:

Number of Moments	Description
	Husbands:
13	Percentage retired at each age 54-66
5	Percentage retired at 55, 58, 60, 62, 65, given own health problem
5	Percentage retired at 55, 58, 60, 62, 65, given wife's health problem
5	Percentage retired at 55, 58, 60, 62, 65, given low lifetime family potential earnings
5	Percentage retired at 55, 58, 60, 62, 65, given high lifetime family potential earnings
	Wives:
13	Percentage retired at each age 54-66
5	Percentage retired at 55, 58, 60, 62, 65, given own health problem
5	Percentage retired at 55, 58, 60, 62, 65, given husband's health problem
5	Percentage retired at 55, 58, 60, 62, 65, given low lifetime family potential earnings
5	Percentage retired at 55, 58, 60, 62, 65, given high lifetime family potential earnings
	Percentage of couples with both spouses retired as of the survey dates, given that:
6	Both spouses enjoy retirement more if the spouse is retired
6	The husband enjoys retirement more if the wife is retired
6	The wife enjoys retirement more if the husband is retired
6	Regardless of whether or not the spouses enjoy each other's retirement
	Percentage of couples:
1	With the wife retired in 1992, husband working in 1998
1	With the wife retired in 1994, husband working in 2000
1	With the wife retired in 1996, husband working in 2002
1	With the husband retired in 1992, wife working in 1998

1	With the husband retired in 1994, wife working in 2000
1	With the husband retired in 1996, wife working in 2002

There are 96 moments total. Among the moments for the husbands and wives, the percentages retiring at each age help identify the standard deviations of the epsilon terms, and the spike at age 62 helps identify the coefficient of age. The moments related to health problems help identify the corresponding coefficients of health, and the moments related to potential earnings help to identify the consumption parameter. The moments related to joint retirement help identify the corresponding joint retirement parameters, and the last six moments help identify the correlation of the epsilons. The higher the correlation, the less frequent should be cases where couples retire many years apart, and vice versa. The tendency to retire many years apart is what these moments capture.

IV. Model Estimates

Table 1 reports the coefficients estimated for the base model of retirement and saving in two earner households. The overall fit of the model, at least with respect to the moments used in the estimation, is good. The q-statistic, which is distributed as a chi-squared distribution if the model is correct, has a value of 94.26 with 80 degrees of freedom. This compares with a critical value of 101.88 at a 5% significance level and 96.58 at a 10% significance level, indicating that there is little evidence against the model at any usual degree of significance.

The coefficient on own age is a key determinant of the sensitivity of retirement to policy changes. If the coefficient on age is very high, the disutility of work is rising quickly with age, so respondents will not be sensitive to marginal changes in incentives from policy changes. If the coefficient is low, then there is time for policy changes to influence retirement. Actuarially neutral policies of the type considered in this paper would not change retirement incentives for those with a time preference rate that is equal to the rate of return, but for those whose time preference rate is greater than the rate of return, even actuarially neutral policies can create strong incentives. In Table 1, the coefficient on own age is .197 for men, and .170 for women, with both coefficients statistically significant and not very different from each other.

Among the remaining variables, the results suggest that husbands strongly prefer retirement if their wives are also retired. This appears to be true whether or not the man enjoys spending time with his wife. The size of this effect is equivalent to roughly the effect of being 4.9 years older among those men who like spending time with their wife and 4.2 years for men who do not necessarily like spending time with their wife. In contrast, the wife's retirement decision is not much directly influenced by her husband's retirement status, whether she likes to spend time with her husband or not. Having poor own health increases retirement by the equivalent of 5.5 additional years of age for men, and 4.9 years of additional age for women. If their spouse is in poor health, men prolong their time in the labor force, while there is no effect on the wife's retirement from having a husband in poor health. For both husbands and wives, the standard deviation of leisure is highly significant. The correlation between the unmeasured preference for leisure by husbands and wives is also significant.

To consider the implications of the model, for each observation 10,000 draws of ε_h and ε_w are obtained from a bivariate normal distribution with standard deviations σ_h and σ_w and correlation ρ_ε. For each draw, the retirement ages of the two spouses are calculated using the backward induction method and the results are tabulated for all the draws and all the individuals in the sample.

Table 2 shows the simulated joint retirement probabilities at each year of age for husbands and wives, as well as the marginal probabilities. Whatever the retirement age of their spouse, the retirement of both husbands and wives peaks at age 62. Although the estimates confirm the importance of joint retirement, the pull of interdependent retirement decisions does not obscure the retirement spike at age 62 for either spouse. The simulations suggest only a very small spike at age 65 for men, and no spike at 65 for women from two earner households. This is probably because the model abstracts away from Medicare.

The tendency of couples to retire together is also captured by the model. This aspect of the simulations is reported in Table 3. The spike at zero years difference – they retire at the same time -- includes 17.5 percent of couples. Given the differences in their ages, if they do not retire in the same year, husbands are somewhat more likely to retire before their wives.

V. Sensitivity to Modeling Sources of Interdependence in Two Earner Families

Consider next the basic model of retirement in two earner families, and eight variants of that model. The purpose of these variants is to determine the sensitivity of responses to roughly actuarially neutral policy changes to the sources of joint decision making within the household. Essentially we change the extent and sources of interdependence and observe the effects on the level of retirement for husbands and wives. The models proceed from the base version, described in Section III and estimated in Section IV, which includes all sources of interdependence. Next we consider models that eliminate the sources of interdependence one by one, leading ultimately to models that eliminate all the parameters of the model that generate interdependence. The specifications of these models is as follows:

Model 1 is the basic model described in Section III and estimated in section IV.

Model 2 uses the same parameter estimates as model 1. In this model, each of the 10,000 simulations per individual are done twice. The first time the simulation is computed identically to model 1; that is, the retirement ages are computed according to the joint model. In the second simulation, the wife's retirement age is held constant at the age calculated in the first simulation.

The same values of the epsilons are used as in the first simulation, and the spouse retirement variables are still included. The husband's retirement age is then recalculated assuming that the wife's retirement age is exogenous at the age calculated in the first simulation. The reported results of this model use the second set of simulations.

The idea of this model is to hold the wife's retirement as fixed and to eliminate interdependence in the retirement decision itself, as described in the Appendix, as opposed to interdependence arising from correlated epsilons and the presence of spouse retirement variables. Note that interdependence in the retirement decision itself only occurs if, in a particular year, the decisions of both spouses depend on the decisions of the other spouse. The majority of retirement decisions, however, are cases where one spouse retires before the other, where there is no contemporaneous interdependence of the decisions.

Model 3 is similar to model 2, except that after the wife's retirement date is established from the joint retirement model, the husband is given a new draw of epsilon. In essence, the difference is that in model 2, the two epsilons are still correlated, but in model 3 the epsilon of the husband is independent of the wife's retirement. This model essentially removes correlated preferences as a source of joint retirement.

Model 4 takes the wife's retirement to be exogenous at the observed age (or the expected age if she has not retired by the last survey), and the husband's retirement model is estimated and simulated. Although the wife's retirement is exogenous, the husband's spouse retirement variables are retained.

Model 5 is analogous to model 4, except here the spouse retirement variables are omitted. This model removes the last source of correlated retirement in the preference sets. The resulting model is essentially the model that would be estimated for husbands, paying no attention to joint retirement and taking the wife's income as exogenous.

Models 6-9 are analogous to Models 2-5 except that these models look at the wife's retirement and take the husbands retirement as exogenous.

Tables 4 and 5 consider how the effects of policies vary as we move models where joint retirement is a central feature to models in which the retirement of one spouse is taken as exogenous in the estimation of retirement of the other spouse. Table 4 deals with changes which might occur if the Social Security early entitlement age were to increase from age 62 to age 64. This change would be roughly actuarially neutral, since the benefit adjustment rate is roughly actuarially neutral between those ages. For each model, the distribution of retirement ages for the husband and/or wife is calculated from that model. Next, the opportunity set for each model is changed so that Social Security can only be collected at age 64 or later, and another distribution of retirement ages is calculated for that model. For each model and for each retirement age, the percentage of individuals retired if the early entitlement age is 62 is subtracted from the percentage of individuals retired if the early entitlement age is 64, and those results are reported in the table. For instance, in model 1, the simulations show that at age 62, 55.5 percent of husbands would be retired if the early entitlement age is 62, but only 50.9 percent would be retired if the early entitlement age were to become 64. The difference is a decline of 4.6 percent, as reported in the first column of Table 4.

A cursory glance at Table 4 reveals that the effects of this change do not depend all that much on the model which is used to estimate them. For the husbands, the effect of the policy change at age 62 range from a decline of 4.1 percent to a decline of 4.6 percent. At the two extremes are estimates of a decline of 4.6 percent if all sources of interdependence are included and a decline of 4.4 percent if the husband's model is estimated with the wife's income taken as exogenous. The same story is true for the wives, with one glaring exception. The figures for model 8 are considerably lower than the figures estimated for any of the other models. This is the model where the husband's retirement is taken as exogenous but the model is reestimated with the spouse retirement variables retained. From model 1, the husband's joint retirement variables are substantial, but the wife's joint retirement variables are small. When the husband's retirement is taken as exogenous, however, the wife's joint retirement variables are estimated to be large, in large part to pick up on the observed joint retirement. The q statistic suggests that this model does not fit the data very well as a result (q = 67.5 vs. a critical value of 52.2 with 37 degrees of freedom). Omitting this model as somewhat suspicious, among the remaining models the estimated effects at age 62 range from a decline of 4.2 percent to a decline of 4.7 percent. At the two extremes for the wives are estimates of a decline of 4.5 percent if all sources of interdependence are included and a decline of 4.7 percent if the husband's model is estimated with the wife's income taken as exogenous. Overall, this table suggests rather strongly that, at least for this policy change, the estimated effects are not much different whether the full joint retirement model is used or the simple one spouse model, with the other spouse taken as exogenous, is used.

Table 5 considers another hypothetical policy change which might be relevant were Social Security to be converted to personal accounts. The base case in this exercise is an environment is which personal accounts have been instituted. The personal accounts receive contributions equivalent to the full Social Security contribution rate (12.4%) and individuals can receive benefits at age 62 in the form of a real annuity. The alternative case allows amounts over the amount necessary to finance a poverty-level annuity to be taken out as a lump sum at age 62. Almost by definition, these two alternatives are actuarially neutral.

The entries in Table 5 are the percentage of individuals retired with the lump sum option minus the percentage retired when benefits must be taken as an annuity. In the base case, for instance, the percentage of husbands retired at age 62 would be 7.1 percentage points higher if the lump sum option were available. By and large, these are individuals with higher time preference rates who take large lump sums and retire as soon as the lump sum becomes available. The main point of Table 5 is that, with the exception of the somewhat questionable model 8, the effects of lump sum availability do not vary too much depending on the model, and in particular whether the model attempts a full description of joint retirement or whether the model takes the behavior of one spouse as exogenous. Among the husbands, the effect of lump sums is 7.1 percent in the full joint retirement model and 8.1 percent in the model where the wives' retirement is taken as exogenous. Among the wives, the effect is 4.2 percent with joint retirement and 3.8 percent with the husbands' retirement taken as exogenous.

VI. Effects of Actuarially Fair Policies on Retirement in One and Two Earner Households

In Gustman and Steinmeier [2005], we estimated a model similar to model 5 for a sample of married men in the HRS, and we found somewhat larger effects for increasing the Social Security early entitlement age to 64 than those reported here in Table 4. The results in the previous section suggest that the differences are not due to modeling joint retirement, since the results of model 5 are reasonably close to those from model 1. The most obvious remaining possibility is that the earlier results were estimated for a sample of married men, whether they came from two earner or one earner households, and the results in the previous section pertain only to the sample of married men whose wives were also career workers.

To investigate whether this explanation may account for the differences between the current results and the earlier results, we estimated model 5 for the sample of married men in the HRS whose wives were not career workers. The households of these married men will be referred to as "one earner households," as opposed to the "two earner households" used in the previous section. Table 6 contains two sets of estimates for model 5, one for husbands in two earner households and the other for husbands in one earner households. Examining these two sets of estimates, it appears that all of the coefficients in one set are within no more than a standard deviation of the values in the other set, with the exception of the constant term. Indeed, a formal test of the hypothesis that the two sets of coefficients are equal is roundly rejected for the full set of coefficients (χ^2 = 54.01vs. a critical value of 12.59 for 6 df), but if the test is restricted to the five coefficients other than the constant, the hypothesis of equality is readily accepted (5.92 vs. a critical value of 11.07 for 5 df).

The top half of Table 7 reports on the effects of the two policies for the husbands of one earner as well as two earner families. The most noticeable thing about this table is the much greater effect of the change in the early entitlement age for the one earner husbands. Whereas the husbands in two earner families reduce retirement by 4.4 and 2.8 percentage points at ages 62 and 63 in response to an increase in the early entitlement age to 64, the corresponding percentage point reductions for husbands in one earner families are 10.2 and 9.3, respectively. The husbands in one earner families thus appear to have a much greater retirement response to the change in the early entitlement age. The same thing is true, but to a much lesser extent, for introducing the lump sum option into the personal account scenario.

There are several potential causes for the differential responses between the husbands in one earner vs. two earner households. One possibility is that, despite the fact that most of the parameters in the estimates for the two groups are similar, the actual differences are responsible for the differential responses. Another is that the two groups have different distributions of time preferences, which are fixed effects in the model not reflected in the parameter estimates. Still another possibility is that differences in the opportunity sets create the differential responses.

To shed light on the first possibility, we simulate the responses of the two earner husbands using the estimated parameters of the one earner households, and vice versa. A problem immediately arises in that when the two earner opportunity sets are used with the one earner parameters, the simulated retirement rates are much too high. For instance, using the one earner estimated parameters, the estimated percentage retired at age 62 is 54.8 percent for one earner husbands but rises to a whopping 71.7 percent when these same parameters are applied to the two earner husbands. The problem appears to be that the earnings of the wife for two earner couples acts like a wealth effect for the husband, lowering the marginal utility of income and causing retirement to occur at an earlier age. In the two earner estimated parameters, the constant term is significantly more negative, which has the effect of reducing the desirability of retirement and thus more or less offsetting the wealth effect. Thus, when the simulations are done for two earner husbands using their own estimated parameters, the retirement rate at age 62 is 53.2 percent.

Our primary interest, however, lies with the parameters other than the constant term. In particular, we would like to know whether the higher age coefficient in the two earner estimates leads to a lower response rate. To focus on this issue, when we simulate one group of husbands using the other group's parameters, we adjust the constant term to maintain the same retirement rate at age 62 as for the base case. The bottom part of Table 7 gives the results. The first and third columns indicate the simulated effects of the two policy changes on two earner husbands, using the coefficients from the one earner estimates but adjusting the constant so that the retirement rate for the base case is still 53.2 percent. The second and fourth columns indicate the effects on one earner husbands, using the coefficients from the two earner estimates but again adjusting the constant so that the retirement rate for the base case is 54.8 percent. A comparison of the columns in the top part of the table with the corresponding columns in the bottom part of the table reveals little difference in the effects, regardless of whether the two earner estimates of the one earner estimates are used. This is consistent with the fact that the two sets of estimates are not significantly different except for the constant.

Since the estimated parameters do not appear to be responsible for the differential effects of the two groups, we next examine the possibility that the differences are somehow due to different distributions of the time preference rates. Table 8 summarizes the time preference distributions for the two groups. The distributions indicate that the time preferences of the one earner households tend to be a bit more extreme on both ends, with the one earner households being 7.1 percent more likely to have a time preference rate below 0.05 than two earner households, and 6.6 percent more likely to have a time preference rate above 0.5.

It is difficult to simulate the effects of inserting one group's time preference rates in the other group's opportunity sets, since time preference is an individual fixed effect and the mapping from one group to the other group is not clear. Nevertheless, it would appear that the fact that the one earners have slightly more high time preference individuals is insufficient to account for the differences in retirement responses, particularly for the increase in the early entitlement age. Among two earner households, the fact that 40% of households have time preference rates in excess on 0.5 leads to a retirement spike at age 62 that is roughly 11%, which in turn reduces retirement at age 62 by approximately 4.4 percent if the early entitlement age is delayed to age 64. Relative to the two earner households, the fraction of households with a time preference over 0.5 percent is about 18 percent higher in one earner households. It seems unlikely that this would be associated with a more than doubling of the retirement effects of (10.2 percent vs. 4.4 percent at age 62) of increasing the early entitlement age to 64.

By elimination, the remaining possibility for the differential retirement effects of changing the early entitlement age is the differences in the opportunity sets. This possibility actually seems fairly plausible. According to Table 3, about half of husbands in two earner households retire before their wives. This means that if they were to retire a year or two later, they would still have the earnings from the wife available. For households with high rates of time preference, who in general have little savings, the availability of the wife's earnings may make it possible for the husband to retire at 62 even if the early entitlement age were to be increased to age 64. For a one earner household, a husband from a high time preference household who would like to retire at age 62 is likely to have little savings and no earnings from the wife to fall back on. If the early entitlement age were to be raised to age 64, this husband would be unable to retire at age 62 because doing so would leave the family with no income at all. For this reason, it seems reasonable that changes in the early entitlement age would have a substantially greater impact on men from one earner households than men from two earner households.

Note that such an argument does not come into play for the case of annuitized vs. lump sum payments from personal accounts. The husband from a one earner household would still have income available at age 62 whether or not the lump sums are available, and retirement at age 62 would be feasible in either case. Thus, the requirement to annuitize benefits does not have the same incentive to delay retirement for one earner husbands that the postponement of the early entitlement age to 64 does. This argument is consistent with the results that indicate that the effects of increasing the early entitlement age are much larger for one earner than for two earner households, but the effects of allowing lump sums in personal accounts are only slightly larger for one earner than for two earner households.

These results indicate that if one earner and two earner households were combined into a single group and a model such as model 5 were estimated for the combined group, simulations from such a model would probably underestimate retirement for the one earner households and overestimate it for the two earner households. This is because, with the other parameters roughly the same between the two groups, the estimated value for the constant term is likely to be between the two values found in Table 6. The intermediate value would (algebraically) greater than the value estimated for the two earner households and (algebraically) less than the value estimated for one earner households. In turn, this would result in simulations in which the amount of retirement is too high for the two earner households and too low for one earner households. The implication is that if the purpose is to derive accurate results for the levels of retirement for one earner vs. two earner households, it is probably wise to use separate estimates for each group, or, at a minimum, to allow the constant term to be different between the two groups.

If the purpose of the analysis is to analyze the overall effects of potential policy changes, though, a single equation for the two groups combined may not be so bad. This is because most of the incentives to change retirement depend on parameters in the model other than the constant, and particularly they usually depend on the value of the coefficient on age. Changes in the constant serve to move the entire distribution of retirement up or down, but they do not in general change the relative incentives to retire at various ages. But the differences in the other parameters of the model between one earner and two earner households are not significant at conventional levels, and it appears in the simulations that these differences do not much affect the effects of the policies considered in this paper.

VII. Conclusions

There are three general conclusions that one can take away from these results. First, the analysis confirms that actuarially neutral policies may have substantial effects on retirement behavior. Both policies simulated here, increasing the early entitlement age from 62 to 64, and offering benefits from personal accounts as a lump sum instead of an annuity, would have a substantial impact on retirements. The effects are somewhat greater for husbands in one earner households than in two earner households. This is particularly true for an increase in the early entitlement age, which would reduce retirement at age 62 by 10.2 percentage points for husbands in one earner families but only 4.4-4.6 percentage points (depending on the exact model used) in two earner households. The effect on the retirement of wives in two earner households is slightly greater (5.1-5.7 percentage point reduction at age 62) than that of the husbands.

The policy of allowing lump sums from personal accounts, rather than requiring the personal accounts to be annuitized, raises the retirement levels of husbands at age 62 by 7.1-8.1 percentage points for husbands in two earner households. The effect on the retirement level of husbands in one earner households is only slightly greater, at 8.9 percent. For wives in two earner households, the increase in retirement at age 62 would be 2.8-3.2 percentage points.

The second general conclusion is that among the two earner households, it does not make a great deal of difference to the simulated effects of these policies whether the estimates are from a model which takes detailed account of the interactions between the retirement decisions of the two spouses or from a model which takes the retirement and earnings behavior of the second spouse as essentially exogenous. There are instances where it would be important to model and estimate the interactions of the two spouses, particularly if the object of the analysis is to understand better the joint retirement decision. But for the purposes of understanding the effects of the types of policies being considered here, the need to model the interaction in a detailed way is not evident. This is good news for many analysts, particularly those doing structural modeling. Modeling the decisions of a two earner couple essentially doubles the number of state variables required compared to a model behavior of a single individual. Thus there can be an enormous saving in complexity by simplifying the specification of certain models to be used in policy analysis, allowing interactions at the family level to be ignored.

The third general conclusion relates to the process of using the same model to estimate the effects of policy actions on men from one earner vs. two earner households. It is clear that the two groups have different responses to the policy changes examined in this paper, and particularly to an increase in the early entitlement age. Such a difference may be due to different preferences or to different opportunity sets. The analysis of the last section suggests that the differential effect to the early entitlement age increase may be primarily due to differences in the opportunity sets of the husbands in one earner vs. two earner households, but a word of caution is probably in order here. Although most of the coefficients are not significantly different between the estimates of one earner vs. two earner households, the constant is clearly (algebraically) greater in the estimates for the one earner households. Given that the constant influences the overall level of retirement, any analysis which looks at overall retirement between these two groups should almost certainly conduct separate analyses for the two groups. Even if the focus of the analysis is on the overall effect of policy changes, it would probably be wise to conduct separate analyses of the two groups until more evidence accumulates to suggest that a single analysis of the two groups combined does not distort the results too much.

References

Blau, David M. and Donna B. Gilleskie. 2006. "Health Insurance and Retirement of Married Couples". *Journal of Applied Econometrics*, in Press, available on line, November, http://www3.interscience.wiley.com/cgi-bin/fulltext/113453619/PDFSTART

Gustman, Alan L. and Thomas L. Steinmeier. 2000. "Retirement in a Family Context: A Structural Model for Husbands and Wives". *Journal of Labor Economics* 18(3): (July), pp. 503-545.

________________. 2004a. "Personal Accounts and Family Retirement". NBER Working Paper 10305.

________________. 2004b. "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.

________________. Forthcoming. "Projecting Behavioral Responses to the Next Generation of Retirement Policies". *Research in Labor Economics*.

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.

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

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

Appendix 1: Solution to the Model through Backward Induction[3]

We model the decision making using the same backward induction method that is used to solve dynamic control models. To do this, we note that for any pair of retirement dates (R_h, R_w), the household income streams are completely specified. With these income streams, the optimal consumption path can be calculated. The optimal consumption path and the two retirement dates are sufficient to calculate the lifetime utilities of both partners, contingent on the two retirement dates. These can be denoted as $U_h(R_h, R_w)$ and $U_w(R_h, R_w)$.

Let T be the last year in which the couple can plausibly work. This means that the couple must be retired in T+1. Starting in year T, each spouse considers whether it would be better to work or retire in that year. There are four possible outcomes, which for expositional purposes are denoted in the following table:

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

Both the husband and wife have utility measures of these four outcomes. These are denoted by U_i^j, where i denotes the decision (A, B, C, or D), and j denotes the spouse (h for husband, w for wife). For example, U_C^h would be the husband's utility if he were to work but the wife were to retire. These calculations assume that both spouses work in previous periods; the possibility of prior retirement is considered in subsequent steps.

One possibility is that one of the spouses, say the husband, finds it advantageous to retire regardless of the wife's decision. This would occur if $U_B^h > U_A^h$ and $U_D^h > U_C^h$. In this case, the wife will compare the two combinations where the husband retires, which are B and D. If

[3] Reproduced from Gustman and Steinmeier (2004a).

$U_D^w > U_B^w$, the wife will also retire and the combination D will be chosen. If the inequality goes the other way, the wife will work and the combination B will be chosen. A symmetric process occurs if the wife finds it advantageous to retire regardless of the husband's decision.

A similar situation occurs if one of the spouses, say the wife, finds it advantageous to work regardless of the husband's decision. This would occur if $U_A^w > U_C^w$ and $U_B^w > U_D^w$. Knowing that the wife will work in any case, the husband then compares combinations A and B. If $U_A^h > U_B^h$, the husband will choose to work and combination A will be chosen. If the inequality runs the other way, he will retire and combination B will be chosen. As before, the situation is symmetric if the husband finds it advantageous to retire regardless of the wife's decision.

If the decision of each of the spouses depends on the other spouse's decision, things become a little more complicated. Suppose that both spouses find it advantageous to do what the other spouse is doing, be it work or retirement. This occurs if $U_A^h > U_B^h$ and $U_D^h > U_C^h$ for the husband and $U_A^w > U_C^w$ and $U_D^w > U_B^w$ for the wife. This might occur if each spouse valued retirement much more if the other spouse was also retired. In this situation, if either spouse retires, the other spouse wants to retire as well. The husband compares utility at combinations A and D. If $U_D^h > U_A^h$, the husband will find it advantageous to retire, and the wife will follow. If $U_D^w > U_A^w$, the wife will retire and the husband will follow. If $U_A^h > U_D^h$ and $U_A^w > U_D^w$, both spouses will find it advantageous not to retire. Combination A will be chosen, and both spouses will continue to work.

Another possibility is that one spouse wants to do what the other one does, but the other spouse wants to do the opposite. For instance, the husband may have $U_A^h > U_B^h$ and $U_D^h > U_C^h$, while the wife may have $U^w_C > U^w_A$ and $U^w_B > U^w_D$. The wife's preferences may arise if she is on the borderline between work and retirement, and if she values retirement roughly the same regardless of whether the husband is retired. In this case, the wife might want to retire if the husband is working, but if the husband retires, there will be less income and a higher marginal utility of income, and this may cause her to want to work. A key to analyzing this possibility is to note that for the wife, $U^w_A > U^w_B$. If the wife is working, then $L_w = 0$, and the work/retirement decision of the husband cannot affect the leisure term in the wife's utility function. However, having the husband work increases income and consumption, to the benefit of the wife. Thus, if the wife is working, she is unambiguously better off if the husband also works.

This argument enables us to establish a strict preference ordering of the four combinations for the wife: C > A > B > D. What the wife would really like is for her to retire but the husband to work. But she knows that if she retires, the husband will too, leading to the least desirable of the four combinations from her viewpoint. So she will not want to retire. The husband knows that if he retires, the wife will certainly not. But if she does not retire, his best alternative is to avoid retirement as well. Thus, both spouses will continue working. The same reasoning will apply if it is the wife who wants to do the same thing as the husband but the husband wants to do the opposite of the spouse.

The final possibility is that both spouses want to do the opposite of the other. In this case, the ordering of the four combinations can be established for both spouses. For the husband, it is B > A > C > D, and for the wife it is C > A > B > D. The first choice for each spouse is to retire while the other spouse works. But if both spouses retire, they will arrive at combination D, which is the least desirable result from both spouses' point of view. Rather than assume that both rush to retire before the other, we assume that the spouses cooperate enough to avoid a result which is mutually undesirable. Since combination A is the second best alternative for both spouses, we assume that in this case the two spouses agree to both work rather than to race to see who can retire first.

The above discussion deals with every possible ordering of preferences among the four work/retirement combinations for both the husband and wife (with the exceptions of those orderings which are inconsistent with the model) in year T. The next step is to analyze what happens in the year T-1. The two spouses face much the same decision in the previous year that they do in year T. The primary difference is how the utility values associated with the four work/retirement combinations are calculated.

At time T-1, the husband's utility associated with combination D (both spouses retire) is $U_D^h = U_h(T-1, T-1)$, and the wife's utility for combination D is calculated analogously. The utility of both spouses associated with combination A is simply the utility that is associated with the optimum combination if both spouses retire after T-1. These utility values have already been calculated in the previous step. For combination B, the husband is retired but the wife is working in period T-1. This means that the wife can retire at time T or T+1. The utility of the wife in this case is given by $U_B^w = \max_{R_w > T-1} [U_w(T-1, R_w)]$. That is, the wife picks whether to retire at T or T+1, and the resulting utility is the wife's utility for combination B at time T-1. The husband's utility for combination B is the husband's utility if he retires at time T-1 and the wife retires at her optimal choice between T and T+1. The situation for combination C is completely symmetric with combination B.

The choice of the two spouses in year T-1 is determined in exactly the same way as in year T, using the utility values for combinations A, B, C, and D just described. This procedure is repeated for each preceding year in turn until the point where both spouses are presumed to certainly be working. The earliest period when either spouse finds it optimal to retire is that spouse's retirement date. If only one spouse finds it optimal to retire that year, the remaining spouse determines his or her retirement date conditional on the first spouse's retirement date.

This mechanism differs in some important respects from the mechanism we have used in previous papers investigating joint retirement. In those papers, the spouses had separate utility functions but chose their retirement dates using a Nash equilibrium concept. For example, suppose that the spouses were born in the same year and that the Nash equilibrium was that the wife retires at age 63 and the husband retires at age 65. This means that if the husband retires at age 65, the wife has higher utility retiring at age 63 than at other ages, for instance at age 62 or age 64. Similarly, if the wife retires at age 63, the husband has higher utility retiring at age 65 than at other ages, for instance at age 64 or age 66. In other words, in the Nash equilibrium concept the retirement decisions of both spouses are optimal for those spouses, given what the other spouse is doing.

What is missing in the Nash equilibrium concept is the following. Suppose that the wife in the previous example were to consider retiring at age 62 rather than at the Nash equilibrium age of 63. If that were her decision, then the husband's optimal decision might be to retire at age 66 rather than age 65. The wife's decision in the Nash equilibrium would be to compare $U_w(R_h = 65, R_w = 63)$ with $U_w(R_h = 65, R_w = 62)$, which holds the husband's retirement date constant, whereas the correct comparison for the wife's decision as to whether to retire at 63 or 62 should be between $U_w(R_h = 65, R_w = 63)$ and $U_w(R_h = 66, R_w = 62)$. That is, the wife should allow for the fact that if she retires at 62 rather than 63, she knows that the husband will prefer to retire at age 66 rather than 65. The mechanism used in this paper correctly allows for this possibility. When the wife is trying to determine whether to retire at age 62, the process assumes that if she does retire at age 62, the husband will subsequently retire at age 66, while if she does not retire at age 62, she will retire at age 63, and the husband will retire at age 65.

This mechanism has another advantage over the Nash concept which will become more apparent in further research as attention shifts to models with stochastic elements. In a stochastic model, the retirement dates depend on the outcomes of the stochastic process or processes. When one spouse wants to retire, the optimal date of the other spouse quite possibly depends on the outcomes of stochastic processes not yet determined. Thus, the concept of a Nash equilibrium becomes problematic, since the outcome matrix as a function of the two retirement dates becomes stochastic. The usual method to solve these stochastic dynamic programming problems is backward induction, beginning with the terminal period. Since the decision process for retirement described above also uses backward induction, it can easily be integrated into the solution of a stochastic dynamic programming problem.

Table 1: Parameter Estimates for Model of Retirement and Saving in Two Earner Families

	Coefficients	t-statistics
α: Coefficient of Consumption	-0.58	6.43
Husband's Parameters		
β_o: Constant	-10.64	91.80
β_a: Age	0.197	3.55
Wife Retired for Those Who		
β_{se}: Enjoy Time with Wife[a]	4.92	2.46
β_s: Do Not Enjoy Time with Wife[a]	4.21	2.12
β_h: In Poor Health[a]	5.54	4.51
β_{sh}: Wife in Poor Health[a]	-3.06	1.74
σ_ε: Standard Deviation of ε[a]	6.27	10.11
Wife's Parameters		
β_o: Constant	-10.37	42.39
β_a: Age	0.170	4.62
Husband Retired for Those Who		
β_{se}: Enjoy Time with Husband[a]	0.51	0.22
β_s: Do Not Enjoy Time with Husband[a]	0.26	0.13
β_h: In Poor Health[a]	4.92	3.91
β_{sh}: Husband in Poor Health[a]	0.06	0.05
σ_ε: Standard Deviation of ε[a]	5.07	9.41
ρ_ε: Correlation of ε_h and ε_w	0.54	2.11
q-statistic	94.263	
Number of Observations	851	

[a] In order to facilitate estimation and interpretation, these coefficients are relative to the corresponding age term. Thus, the actual value of β_s for the husband is $0.829 = 0.197 \times 4.21$.

Table 2: Simulation of Joint and Marginal Distributions of Retirement Ages
(Base Case with Current ER Age, Current Earnings Test, No Personal Accounts)
(Percent)

Retirement Age of Wife	Retirement Age of Husband: <54	54	55	56	57	58	59	60	61	62	63	64	65	66	67	68	>68	Row Total	Cumulative Total
	Joint Probabilities																	Marginal	
<54	1.00	.41	.81	.85	1.02	.98	1.09	1.37	1.40	2.27	1.08	1.03	1.00	.55	.41	.32	.92	16.5	16.5
54	.15	.07	.12	.16	.22	.24	.20	.27	.26	.42	.20	.21	.21	.14	.08	.06	.13	3.1	19.6
55	.29	.08	.28	.31	.25	.34	.34	.41	.38	.42	.27	.28	.27	.17	.12	.09	.21	4.5	24.1
56	.21	.10	.14	.21	.32	.34	.36	.45	.35	.47	.27	.33	.29	.20	.14	.10	.24	4.6	28.7
57	.23	.09	.17	.16	.27	.36	.38	.50	.42	.47	.31	.34	.37	.23	.15	.12	.25	4.8	33.5
58	.35	.11	.16	.16	.25	.35	.44	.51	.54	.60	.35	.38	.39	.25	.17	.12	.30	5.4	38.9
59	.33	.12	.15	.14	.21	.28	.41	.54	.52	.65	.40	.44	.45	.32	.21	.15	.34	5.7	44.6
60	.44	.14	.19	.20	.22	.36	.41	.55	.67	.85	.58	.60	.63	.42	.28	.21	.47	7.2	51.8
61	.50	.12	.16	.17	.20	.32	.44	.47	.54	.90	.54	.63	.66	.45	.31	.24	.58	7.2	59.0
62	.69	.22	.32	.27	.32	.45	.56	.68	.71	1.60	.75	.80	.80	.61	.47	.33	.93	10.6	69.6
63	.29	.08	.10	.11	.14	.20	.23	.29	.34	.57	.38	.48	.57	.47	.33	.24	.53	5.3	74.9
64	.19	.07	.08	.10	.13	.18	.20	.25	.28	.48	.31	.40	.52	.42	.37	.26	.55	4.9	79.8
65	.21	.07	.08	.08	.12	.18	.19	.28	.25	.47	.29	.33	.43	.38	.34	.28	.57	4.5	84.3
66	.17	.06	.06	.07	.09	.14	.15	.18	.19	.37	.20	.25	.29	.29	.25	.22	.52	3.5	87.8
67	.14	.04	.05	.05	.07	.11	.13	.15	.15	.31	.17	.18	.22	.20	.19	.18	.46	2.8	90.6
68	.09	.03	.04	.04	.06	.09	.09	.13	.13	.27	.14	.15	.17	.16	.14	.13	.41	2.3	92.9
>68	.22	.09	.09	.09	.14	.23	.26	.37	.38	.76	.42	.46	.54	.48	.41	.38	1.80	7.1	100.0
	Marginal Probabilities																		
Column Total	5.5	1.9	3.0	3.2	4.0	5.1	5.9	7.5	7.5	11.9	6.6	7.3	7.9	5.7	4.4	3.4	9.2		
Cumulative Total	5.5	7.4	10.4	13.6	17.6	22.7	28.6	36.1	43.6	55.5	62.1	69.4	77.3	83.0	87.4	90.8	100		

Table 3: Simulated Distribution of Couples by Differences in Years When Husband and Wife Retire

	Years of Difference in Retirement	Percent of Couples
	>10	8.2
	9	2.0
	8	2.3
	7	2.7
Wife	6	2.8
Retires	5	3.4
First	4	3.4
	3	3.2
	2	3.1
	1	2.3
	0	17.5
	1	5.2
	2	4.5
	3	4.1
Husband	4	3.8
Retires	5	3.6
First	6	3.2
	7	3.1
	8	2.9
	9	2.6
	>10	15.8

Table 4: Differences In Retirement Outcomes When Early Retirement Age Is 64 Minus When It Is 62, By Model

	Husbands					Wives				
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 1	Model 6	Model 7	Model 8	Model 9
Age	Base	Spouse Retirement Exogenous	Epsilons Uncorrel-ated	Parameters Reestim-ated	Spouse Retirement Variables Omitted	Base	Spouse Retirement Exogenous	Epsilons Uncorrel-ated	Parameters Reestim-ated	Spouse Retirement Variables Omitted
50	0.0	0.0	0.0	0.0	0.0	0.0	-0.1	-0.1	0.0	-0.1
51	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-0.1
52	0.0	-0.1	-0.1	-0.1	0.0	-0.1	-0.2	-0.1	-0.2	-0.3
53	0.0	-0.1	0.0	-0.1	-0.1	-0.2	-0.2	-0.2	-0.2	-0.3
54	0.0	0.0	-0.1	0.0	-0.1	-0.2	-0.2	-0.3	-0.2	-0.3
55	-0.1	-0.1	-0.1	-0.2	-0.1	-0.3	-0.2	-0.2	-0.2	-0.3
56	-0.2	-0.1	-0.1	-0.2	-0.1	-0.4	-0.4	-0.4	-0.3	-0.5
57	-0.2	-0.2	-0.2	-0.2	-0.3	-0.5	-0.5	-0.5	-0.6	-0.8
58	-0.1	-0.3	-0.3	-0.3	-0.3	-0.7	-0.8	-0.7	-0.8	-1.2
59	-0.2	-0.4	-0.4	-0.3	-0.4	-0.9	-1.0	-0.9	-0.7	-1.3
60	-0.4	-0.6	-0.5	-0.6	-0.7	-1.0	-1.2	-1.0	-0.6	-1.3
61	-0.9	-1.2	-1.2	-1.1	-1.0	-1.7	-1.8	-1.7	-0.9	-2.0
62	-4.6	-4.6	-4.3	-4.1	-4.4	-5.5	-5.6	-5.1	-2.1	-5.7
63	-3.2	-3.0	-2.8	-2.8	-2.8	-4.5	-4.6	-4.2	-1.5	-4.7
64	0.5	0.4	0.5	0.5	0.9	0.5	0.5	0.7	0.2	0.6
65	0.2	0.1	0.1	0.1	0.2	0.8	0.8	0.9	0.3	0.9
66	0.1	0.1	0.0	0.1	0.2	0.8	0.8	1.0	0.3	1.0
67	0.1	0.1	0.1	0.3	0.4	0.8	0.8	0.9	0.2	1.1
68	0.2	0.2	0.1	0.2	0.5	0.7	0.7	0.8	0.2	0.9
69	0.1	0.2	0.1	0.3	0.5	0.5	0.5	0.7	0.2	0.8

Table 5: Differences In Retirement Outcomes Under Personal Accounts, When Benefits Are Paid As A Lump Sum Minus When They Are Paid As A Mandatory Annuity, By Model, Current Rate Contributions

	Husbands					Wives				
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 1	Model 6	Model 7	Model 8	Model 9
Age	Base	Spouse Retirement Exogenous	Epsilons Uncorrel-ated	Parameters Reestim-ated	Spouse Retirement Variables Omitted	Base	Spouse Retirement Exogenous	Epsilons Uncorrel-ated	Parameters Reestim-ated	Spouse Retirement Variables Omitted
50	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0
51	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.0	0.0	0.0
52	0.0	0.0	0.0	0.0	0.0	0.2	0.2	0.2	0.0	0.0
53	0.0	0.0	0.1	0.0	0.1	0.5	0.5	0.4	0.1	0.2
54	0.1	0.1	0.1	0.1	0.1	0.6	0.7	0.5	0.2	0.3
55	0.1	0.1	0.1	0.1	0.1	0.8	0.9	0.7	0.3	0.6
56	0.1	0.1	0.1	0.2	0.3	1.2	1.2	1.0	0.4	1.0
57	0.2	0.2	0.2	0.3	0.2	2.1	2.2	1.9	0.6	1.7
58	0.3	0.3	0.3	0.4	0.4	3.1	3.1	2.9	0.8	2.4
59	0.5	0.5	0.5	0.5	0.5	4.0	4.0	3.9	0.9	3.1
60	0.7	0.7	0.7	0.7	0.8	4.0	4.1	4.0	0.7	3.3
61	1.2	1.0	1.1	1.0	1.3	4.2	4.2	4.2	0.9	3.8
62	7.1	7.2	7.2	7.7	8.1	3.2	3.1	3.1	2.0	2.8
63	5.7	5.8	5.8	6.5	6.5	2.4	2.4	2.4	1.6	2.0
64	4.2	4.3	4.3	5.2	4.7	1.7	1.8	1.7	1.3	1.5
65	3.0	3.1	3.1	4.1	3.4	1.2	1.1	1.1	1.1	1.0
66	2.2	2.2	2.1	3.0	2.4	0.7	0.6	0.7	0.9	0.7
67	1.4	1.5	1.4	2.2	1.6	0.4	0.3	0.3	0.7	0.3
68	1.1	1.1	1.0	1.6	1.3	0.1	0.0	0.1	0.5	0.1
69	0.8	0.8	0.7	1.3	0.9	-0.1	-0.1	0.0	0.4	0.0

Table 6: Parameter Estimates for Models of Retirement by Husbands in Two Earner and One Earner Families, Assuming Exogenous Labor Force Behavior of Wife

	Two Earner Families	One Earner Families
α: Coefficient of Consumption	-0.64 (3.55)	-0.47 (3.39)
β_o: Constant	-10.411 (99.28)	-9.85 (147.16)
β_a: Age	0.139 (3.61)	0.084 (3.02)
β_h: In Poor Health	5.46 (4.28)	4.40 (2.09)
β_{sh}: Wife in Poor Health	-0.40 (0.28)	0.47 (0.30)
q-statistic	32.801	29.926
Number of Observations	851	866

The figures in parentheses are t-statistics.

Table 7: Differences in Percentage of Husbands Retired Due to Policy Differences in Two Earner and One Earner Households, Assuming Exogenous Labor Force Behavior of Wife

	ER = 64 vs ER = 62		Lump Sum vs. Annuity	
Age	Two Earners	One Earner	Two Earners	One Earner
53	-0.1	-0.1	0.1	0.1
54	-0.1	-0.1	0.1	0.1
55	-0.1	-0.1	0.1	0.1
56	-0.1	-0.2	0.3	0.2
57	-0.3	-0.2	0.2	0.2
58	-0.3	-0.3	0.4	0.3
59	-0.4	-0.6	0.5	0.4
60	-0.7	-1.0	0.8	0.7
61	-1.0	-1.6	1.3	1.4
62	-4.4	-10.2	8.1	8.9
63	-2.8	-9.3	6.5	6.6
64	0.9	1.1	4.7	4.7
65	0.2	0.3	3.4	3.4
66	0.2	0.5	2.4	2.5
67	0.4	0.6	1.6	2.0
68	0.5	0.7	1.3	1.6
69	0.5	0.6	0.9	1.3
	With One Earner Parameters	With Two Earner Parameters	With One Earner Parameters	With Two Earner Parameters
53	-0.1	0.0	0.1	0.0
54	-0.1	0.0	0.1	0.0
55	-0.1	-0.1	0.2	0.1
56	-0.2	-0.1	0.3	0.1
57	-0.3	-0.2	0.3	0.1
58	-0.4	-0.3	0.5	0.2
59	-0.5	-0.6	0.7	0.3
60	-0.6	-1.1	1.0	0.6
61	-0.6	-2.3	1.4	1.4
62	-4.5	-9.9	7.8	9.4
63	-3.2	-8.6	6.4	6.5
64	1.5	0.5	4.9	4.3
65	0.3	0.3	3.8	2.8
66	0.3	0.4	2.9	2.0
67	0.4	0.5	2.3	1.4
68	0.6	0.5	2.0	1.0
69	0.6	0.4	1.6	0.8

Table 8: Distributions of Time Preferences for Two Earner and One Earner Households

Time Preference Range	Two Earner Households	Single Earner Households
0.00 – 0.05	39.6%	46.7%
0.05 – 0.10	20.2	11.1
0.10 – 0.25	8.0	3.8
0.25 – 0.50	1.9	1.6
> 0.50	30.2	36.8