University of Michigan Retirement Research Center

Working Paper

WP 2004-067

Personal Accounts and Family Retirement

Alan L. Gustman and Thomas L. Steinmeier

MRRC

Project #: UM02-06

"Personal Accounts and Family Retirement"

Alan L. Gustman
Dartmouth College and NBER

Thomas L. Steinmeier
Texas Tech University

January 2004

Michigan Retirement Research Center
University of Michigan
P.O. Box 1248
Ann Arbor, MI 48104

Acknowledgements

This work was supported by a grant from the Social Security Administration through the Michigan Retirement Research Center (Grant # 10-P-98358-5). The opinions and conclusions are solely those of the authors and should not be considered as representing the opinions or policy of the Social Security Administration or any agency of the Federal Government.

Personal Accounts and Family Retirement

Alan L. Gustman
Thomas L. Steinmeier

Abstract

This paper constructs a model of retirement and saving by two earner couples. The model includes three dimensions of behavior: the joint determination of retirement and saving; heterogeneity in time preference; and the interdependence of retirement decisions of husbands and wives. Estimation is based on panel data from the Health and Retirement Study covering the period 1992 to 2000.

When husbands postpone their retirement so they can retire together with their typically younger wives, the spike in retirement at age 62 is smeared to later ages. Thus retirements differ between one and two earner families. We find both an asymmetry in which husbands prefer their wife to be retired before they retire, and a clear distaste of many husbands to retiring when their wives are in poor health, while the wives are willing to stay at home with sickly husbands.

We simulate a system of personal Social Security accounts based on a 10.6 percent contribution rate over the lifetime. One version allows individuals to make lump sum withdrawals at retirement instead of annuitizing. This program would increase the retirement rates of husbands at age 62 by about 15 percentage points compared to the current system. Adding a lump sum option, by itself, would increase retirements at 62 by about 6 percentage points.

Authors' Acknowledgements

This paper was supported by a grant from the U.S. Social Security Administration to the Michigan Retirement Research Center (UM02-06), with a subcontract to the National Bureau of Economic Research. The opinions and conclusions are solely those of the authors and should not be construed as representing the opinions or policy of the Social Security Administration or any other agency of the Federal Government, the Michigan Retirement Research Center, or the National Bureau of Economic Research. Alan L. Gustman is Loren Berry Professor of Economics at Dartmouth College, Department of Economics, Hanover, N.H. 03755 (alan.l.gustman@dartmouth.edu). Thomas L. Steinmeier is Professor of Economics, Texas Tech University, Department of Economics, Lubbock, Texas 79409 (Thomas.Steinmeier@TTU.edu).

I. Introduction

Most individuals around the age of retirement are part of a couple. In two-earner couples, the retirement decisions of such individuals are intertwined with the retirement decisions of their spouses, for at least two reasons. First, the rules governing Social Security payments to the couple depend in a fairly complex way on the retirement decisions of both spouses. Secondly, the retirement decision by either spouse may directly influence the attractiveness of retirement to the other spouse. Given these interdependencies, it would seem desirable to allow for family effects when trying to model the retirement effects of potential changes in the Social Security system.

In the context of two-earner couples, there are a couple of important empirical regularities that should be reflected in any plausible model that attempts to explain retirement. The first of these is a tendency of the spouses in a two-earner couple to retire at about the same time. The near simultaneous retirement of both spouses happens considerably more frequently that would occur if the retirement decisions of the spouses was independent.[1] The second empirical regularity is the tendency for individuals, both husbands and wives, to retire disproportionately around the age of 62, which is not coincidentally the age of early retirement under the Social Security system. In order to provide credible predictions for potential changes in Social Security, a model should be able to reflect both of these empirical regularities.

Generating the age 62 spike in retirement without using age shifters has been a difficult challenge for retirement models. Models that assume a perfect capital market and homogeneous preferences, as most do, suggest that today's Social Security benefit formula has little effect on retirement behaviour. The Social Security benefit structure is now roughly actuarially fair, and

---

[1] See, for example, Hurd (1990), Blau (1998), Coile (1999), Gustman and Steinmeier (2000 and forthcoming), Blau and Gilleskie (2001), Johnson and Favreault (2001) and Maestas (2001).

so does not generate any incentives to retire at particular ages, at least when it is assumed that the discount rate guiding each household's behaviour is the same as the real rate used in calculating the actuarial adjustments for deferred Social Security benefits. So models that measure retirement incentives in present value terms gain no traction in trying to explain retirement using the current Social Security benefit formula. Moreover, even though the introduction of Social Security, and benefit enhancements adopted in subsequent years, created windfalls for previous generations, benefits paid to those reaching retirement age today let them just about break even on their payroll tax contributions.[2] So there is little wealth effect from Social Security to affect retirement outcomes.

Yet the retirement data clearly suggest that Social Security is affecting retirements, especially at age 62. We have recently developed a model that jointly explains retirement and saving and allows for heterogeneous time preference. This model attributes the spike at age 62 to individuals with high time preference who find it preferable to accept Social Security benefits as soon as they become available (Gustman and Steinmeier, 2003).[3] Normally, if an individual delays retirement beyond age 62 and thus foregoes current Social Security benefits, later benefits are increased in an amount that is roughly actuarially neutral, assuming the appropriate discount rate is about 3 percent real. Individuals with a time preference rate much higher than 3 percent, however, will find the future benefits to have considerably less value than the current benefits foregone. In this circumstance, delaying retirement effectively means giving up benefits, and this reduces perceived current compensation. The reduction in compensation at age 62, in turn,

[2] Nor does Social Security create strong wealth differences among covered families. Spouses of higher earners are less likely to work, yet they benefit from spouse and survivor benefits. For this and other reasons, despite its progressive benefit formula, Social Security is not highly redistributive among families with different incomes. For evidence on the variation in returns to Social Security by family income, see Gustman and Steinmeier (2001).

[3] Earlier studies analyzing the joint relation between retirement and saving include Diamond and Hausman (1984) and Kahn (1988). Rust and Phalen (1997) and Blau and Gilleskie (2003) analyze retirement and health insurance, assuming that the capital market is inoperative. French (2002) estimates a model jointly explaining retirement and saving, but his model does not explain the spike in retirement at age 62.

induces many of these individuals to retire at that age. A further factor inducing high time preference individuals to retire at age 62 is that many of them will not have saved enough to allow them to retire before 62. Consequently, many such individuals who would have liked to retire before age 62 must wait until Social Security benefits become available before retiring. A main drawback with this model as previously estimated, and indeed of other models focusing on the age distribution of retirement, is that they treat the retirement decision as an individual decision and, in the case of couples, take the decision of the spouse as exogenous.

A related literature models the retirement decisions of two-earner couples. The principal focus of this family retirement literature is on the degree to which husbands and wives retire at about the same time, and not on the age distribution at which retirement occurs. The empirical work has established that husbands and wives are more likely to retire at about the same time than would occur if the retirements were unrelated. Empirical models of retirement decisions in a family context have successfully explained the coordination in the timing of retirement between husbands and wives, largely as a result of preferences rather than as a result of retirement incentives. However, these models do not contain the elements necessary to explain jointly the tendency of husbands and wives to elevate retirement at age 62.

This paper integrates two lines of research in retirement modeling. It combines our recent model developed to jointly explain retirement and saving outcomes for men, which also explains the spikes in retirement activity at age 62, with a model of family retirement decision making, which explains the coordination of retirement outcomes by two-earner couples. The model employs the backward induction method used in dynamic programming to derive the best retirement dates for both spouses. Estimation is based on biannual panel data from the Health and Retirement Study covering the period 1992 to 2000. Once estimated, the model is applied to examine the long run effects on retirement of adopting personal Social Security accounts and of annuitizing those accounts.

We find that husbands, especially those who value time with their wives in retirement, frequently postpone their retirement so they can retire together with their typically younger wives. This effect reduces the influence of the factors that generates the spike in retirement at age 62, instead smearing somewhat this peak to later ages. In addition, having a career wife means that the household has a larger percentage of income that does not depend on the husband's earnings, and vice versa. Accordingly, the effect on marginal utility from a given dollar amount of incentives is less. This further mutes the economic incentives to retire at a specific date such as at age 62.

The paper explores the long run effects on retirement of the higher returns that would be realized in personal Social Security accounts, and the effects on retirement of providing these benefits in the form of annuities vs. a lump sum. According to estimates available from current retirement models, introducing personal accounts as a substitute for the current Social Security system with its actuarially fair formula should have a small effect on retirement outcomes, reflecting the income effect from higher returns. In contrast, we find that in the long run a complete substitution of personal individual accounts for the current benefit formula would have major effects on retirement. Similarly, allowing lump sum payouts of returns to individual accounts instead of mandatory annuitization would also have major effects on retirement.

Section II presents the model. The sample from the HRS survey is derived in Section III. Retirement outcomes are described in Section IV. Section V provides details of the estimation procedure, while the model estimates are presented in Section VI. Retirement effects of adopting individual accounts are simulated in Section VII. Section VIII concludes the paper.

II. The 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 given by

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

where C is consumption and L is leisure. m is an indicator 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 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:

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

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.

With two utility functions, two decision makers, and one budget constraint, the solution of the model is not as obvious as is the case with a single utility function. 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	Husband's Decision: Retire
Wife's	Work	A	B
Decision	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 $U_D^w > U_B^w$, 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 $U^w_A > U^w_C$ and $U^w_D > U^w_B$ 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 $U^h_D > U^h_A$, the husband will find it advantageous to retire, and the wife will follow. If $U^w_D > U^w_A$, the wife will retire and the husband will follow. If $U^h_A > U^h_D$ and $U^w_A > U^w_D$, 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 $U^h_A > U^h_B$ and $U^h_D > U^h_C$, 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.

Another alternative, and simpler, mechanism for determining the retirement dates of the two individuals is to assume a household utility function which is a weighted average of the utility functions of the two spouses. 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. As with the current model, the weighting model can reproduce joint retirement in approximately the right amounts. However, another important consideration is reproduction of the retirement spikes, especially the spike at age 62. In this respect the weighting model falls far short. 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 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.

III. Data

This paper uses data from the Health and Retirement Study (HRS), a survey of roughly 7600 households with individuals born from 1931 to 1941. If one spouse was in the eligible birth cohorts but the other was not, both spouses were interviewed anyway. 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 though the first five waves of the survey. The HRS also has two important supplements, which are available on a restricted basis. First, the Social Security earnings records through 1991 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 variables necessary to estimate the model are the age and retirement status as of the various survey dates, the actual and potential earnings, pension benefits, the level of wealth as of 1992, health status, and whether or not the respondent enjoyed the prospect of spending time with the spouse in retirement. Work is taken as full-time work, which is defined in this model as working more 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. Respondents are considered to be in poor health after they report “fair” or “poor” health in two successive surveys. Alternative responses are “excellent,” “very good,” and “good” in these questions. A respondent is considered to enjoy time with the spouse in retirement if, in 1992, he or she rated “having more time with spouse” as a “very important” benefit of retirement. About half the respondents chose this response over the alternative responses of “moderately important,”, “somewhat important,” or “not important at all.”

Actual earnings for full-time work are taken from the Social Security earnings record, if available. Earnings over the Social Security limit and in jobs identified as not being Social Security covered are imputed from information in the respondent interview. For respondents who did not give consent to have their earnings records included in the survey, earnings are imputed using the job history part of the survey, which covers the current or last job, the longest job, and other pension-covered jobs. Information about the total number of full-time work years is also used in these imputations. Whether or not the Social Security earnings record is available, for work beyond the last observed date of full-time work, earnings are projected using a wage equation using tenure and experience. If the last job ended involuntarily, the tenure variable is reset to zero in making these projections.

Given the earnings profile, the Social Security formulas are used to calculate Social Security benefits for retirement at various alternative retirement dates. Using the same earnings profile in conjunction with the starting date of the job, pension benefits for alternative retirement dates are calculated using the HRS pension calculator in conjunction with the information in the pension documents supplied by the employers. It should be pointed out that for confidentiality reasons, the HRS collected only the pension documents from the employer and did not try to obtain any information about the respondent, such as starting dates or amounts in pension accounts, from the employer. With this information about the earnings profile, Social Security benefits, and pension benefits, it is possible to calculate a complete household income profile for any combination of retirement dates of the two spouses.

Wealth in 1992 is taken as the sum of financial assets (including IRA's), real estate holdings, and business assets. Excluded are the value of the residence and vehicles, on the grounds that many people may not use these assets to finance consumption in retirement.

Excluded also is the stock value of Social Security and pensions. In this model, Social Security and pensions enter as a flow of benefits in the retirement years. Converting these future benefits into a stock value is ambiguous, given that time preference rates are heterogeneous in the model.

The HRS interviewed almost 4800 couples in 1992; the remaining households were primarily widowed or divorced single households. Not all of the couples could be included in the sample, however. Table 1 gives the various reasons why couples are excluded from the sample and the number of couples lost for each reason. There are three numerically important reasons for exclusion, plus several minor ones. The first major reason for exclusion is that at least one of the spouses was married to a different partner at age 35. In general it is impossible to know the circumstances of the previous marriage and how it contributed to the current situation the observed couple, so these couples are eliminated. The second reason is that one of the partners, usually the wife, does not have a career from which retirement is a meaningful concept. For present purposes, a spouse is considered not to have a career if there is no work after age 49, or if the spouse worked full-time for fewer than half the years between age 40 and the last year of full-time work. For younger spouses who had not reached 49 by the end of the survey, they were considered to have careers if they were working in the last observed survey and had been working full-time 5 of the previous 10 years. This definition of career workers would include, for instance, a wife who dropped out of the labor force during her children's schooling and then returned to work for a substantial period of time. Only 38 percent of the original 4767 couples come from two-earner families who had not been divorced since age 35. The final important reason for exclusion is that the survey did not include the employer report for a pension in the last full-time job. Pensions frequently have strong incentives to retire at specific dates, but without the employer reported pension documents it is very difficult to place those incentives at the correct dates. If those individuals are included, their retirement may be responding to economic incentives that we do not see, leading to biases in the measure strength of economic incentives on retirement. This is not such a problem for pensions in jobs before the last job; missing pensions for those jobs are imputed and the individuals are retained in the sample. Missing pension information reduces the remaining sample by almost a third.

IV. Interdependence and the Retirement Spike at Age 62

To begin the descriptive analysis, it is useful to compare retirement outcomes in two-earner as compared to one-earner families. Table 2 gives the percent of individuals retired by age for both husbands and wives in the two-earner sample of households. So as to minimize the dating problems, these figures are measured as of the survey dates, when the ages and retirement status can be determined fairly accurately from the information in the survey. To illustrate, of the 211 husbands who were age 60 on any of the five survey dates, 31.8 percent of them were retired as of that age. By differencing these numbers between two adjacent ages, we can get the implied number who retired at that given age. For husbands, these pseudo-retirement percentages for ages 60-64 are displayed in the left bars of Figure 1. Notice that for husbands the retirement percentage at age 62 is only a couple of percentage points greater than the retirement percentage at age 63.

This spike in retirements at age 62 is smaller than has been found in the studies mentioned above. The retirement pattern of husbands whose wives are not career workers, given by the middle set of bars of Table 2, helps to solve the puzzle.[4] Comparing the spikes in two-earner and one-earner households, the retirement spike at age 62 is much more noticeable among

[4] The sample of husbands from one-earner households includes those who would have satisfied our selection criterion except that the wives were not career workers. This group amounts to 842 couples, just a little larger than the sample with wives who are career workers.

the one-earner couples. The difference between retirement rates at age 62 and at adjacent ages is even larger when the two groups are combined, as illustrated in the right hand bars of Table 2.

That husbands with career wives do not have quite as strong a tendency to retire at age 62 as do husbands with non-career wives is to be expected. Previous research, cited earlier, has told us that joint retirement will induce husbands and wives to retire at the same time. With an older husband, coordination of retirement means that the husband is postponing retirement until his younger wife leaves the labor market. Thus some husbands who would have left at 62 delay their retirement to coordinate with a younger spouse. There also is another reason a reduced spike in retirement at age 62 might be expected for husbands and wives whose partners are career workers. The fact that the spouse is a career worker probably implies that he or she is generating a substantial amount of family income. This availability of income means that the marginal utility of earnings for either spouse is less than would be the case if only one spouse is earning the bulk of income. The lower marginal utility of earnings in turn suggests that the economic incentives provided by Social Security, pensions, and the like will have less effect on the retirement of either particular spouse. Thus, the presence of a spouse with career earnings would be expected to reduce the ability of economic incentives to induce a retirement spike at age 62.

As further evidence on this issue, we reconsidered a previous model of married males in which the earnings of the spouse was taken as exogenous (Gustman and Steinmeier, 2002). In that model, the simulated retirement spike, defined as the percentage retiring at age 62 minus the average of the percentages retiring at ages 61 and 63, was 8.1 percentage points. When we split that sample into husbands who had career wives vs. husbands with non-career wives, and estimated the model for each group separately, the spike for the sample with non-career wives is

8.7 percent and the spike for the sample with career wives is only 5.4 percent. Hence, the results from that model are also consistent with the magnitude of the spike being smaller for husbands with career wives than for husbands with non-career wives.

Despite the fact that the raw data on the retirement spike at age 62 for husbands with a working spouse roughly accords with our priors, there is a noticeable instability in the retirement percentages, including some highly questionable patterns. In Table 2 for example, between ages 60 and 61 the percentage of husbands who are retired goes up by only 0.1 percentage points, and for wives, the percentage retired is actually lower at age 61 than at age 60.

There are a couple of factors which may contribute to these results. First, the number of observations at any particular age is only a little over 200. For an underlying probability of around 0.5, the 95 percent confidence interval for this number of observations is almost 14 percentage points. An additional factor arises from the nature of the survey. The HRS is a biannual survey, so that in general the same individuals were interviewed at ages 56, 58, 60, 62, and so on, while other individuals were interviewed at ages 57, 59, 61, and so on. Hence, comparing the retirement rates between any two adjacent years is risky because it does not compare the same individuals. In a large sample, this might not be such a problem, but the difficulty is compounded by the uncertainties created by the relatively small sample sizes at any particular age.

One way around the latter complication is to look at the comparisons two years apart, when we ought to be looking at more or less the same individuals. Consider first the respondents who were interviewed at even numbered ages. Approximately 16 percent of these husbands retired at ages 61 and 62, computed as the difference between the 47.4 percent who were retired at age 62 and the 31.8 percent who were retired at age 60. From the prior years, it would appear that the number retiring at age 61 is probably no fewer than 3 percent or so, leaving the percent who retire at age 62 on the order of 13 percent. But now consider the husbands who were interviewed at odd numbered ages. For those husbands, approximately 29 percent retired at ages 62 and 63, measured as the difference between the 60.6 percent retired at age 62 and the 31.9 percent retired at age 61. If 13 percent retire at age 62, this means that, even adjusting for the biannual design of the HRS, the retirement rate at age 63 is almost as high, if not higher, than at age 62. Thus the sharp peak in retirement at age 62 apparent for married men as a whole is somewhat smeared for husbands with career wives.

For women the pseudo retirements in Table 2 suggest a clear spike at age 62. However, even here the biannual data tell a different story about the spike at 62, one that is similar to the outcomes for their husbands, though in slightly less stark terms. For wives interviewed at even numbered ages, retirement at ages 61 and 62 combined appears to be about 21 percentage points. Between ages 56 and 60, it appears that the average retirement rate is on the order of 5 percentage points per year, leaving about 16 percentage points at age 62. For wives interviewed at odd numbered ages, retirement is about 28 percentage points at ages 62 and 63 combined, leaving about 12 percentage points retiring at age 63. Again, almost as many wives appear to retire at age 63 as at age 62. For wives, however, the most noticeable fact in the retirement data is that they retire about two years earlier than their husbands, on average.

With this picture of the retirement outcomes for two-earner families in mind, we now turn to the estimation of the retirement model to determine how well it can do in reproducing these outcomes. These estimates will also suggest some reasons for the retirement hazards described above.

V. Estimation Procedure

For the estimation procedure, we must first specify a bit more completely the individual effects of the model. The individual effects consist of a time preference parameter ρ and two retirement preference effects ε_h and ε_w. The time preference parameter is taken to be a fixed effect and is in essence treated as a value to be estimated separately for every individual in the population. It would be desirable to allow the time preferences of the two spouses to vary separately, but most assets are jointly owned and any division of these assets would be arbitrary. The retirement preference effects can be separately identified, however, because we observe when each spouse retires. In this model, the retirement preference effects are treated as random effects coming from a jointly normal distribution with standard deviations σ_h and σ_w and with correlation ρ_ε.

The estimation procedure is the generalized method of simulated moments (GMM). In this procedure, a group of moments is gathered into a column vector **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 parameters. In general, these moments come from an asymptotically normal distribution with a mean value of zero. The estimation procedure seeks the parameter values which minimize $q = \mathbf{m}'\mathbf{W}^{-1}\mathbf{m}$, where $\mathbf{W} = \sum_{i=1}^{n} \mathbf{m}_i\mathbf{m}_i'$. The $\mathbf{m_i}$ vectors are the moments of the individual observations, and the **W** matrix is essentially the observed variance-covariance matrix of the moments. Variances of the estimates are calculated from $\text{var}(\mathbf{\Theta}) = [\mathbf{G}'\mathbf{W}^{-1}\mathbf{G}]^{-1}$, where $\mathbf{\Theta}$ is the vector of parameters and **G** is the derivative of the moments with respect to the parameters. If the model is correctly specified, m is distributed around zero, and q should have a χ^2

distribution with $\ell - k$ degrees of freedom, where ℓ is the number of moments used 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 retirement dates if the spouses have already retired, or the expected retirement dates if one or both have not retired. 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 stream, and with it an associated path of wealth, is computed using the usual backward induction method of dynamic programming models. 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 parameters ε_h and ε_w in the calculations of ρ, since ε_h and ε_w only affect the leisure terms in the utility functions. This approach does require a slight approximation, however, since the actual or expected retirement dates refer to situations where both couples 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 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 not substantially affect that purpose.

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. With these three individual effects, the optimal retirement dates R_h and R_w can be calculated from the methods described in Section III. Another draw is made from the distribution of ε_h and ε_w, and another corresponding pair of retirement dates R_h and R_w is calculated. The process continues through a large number of simulations (10,000 per observation), and the resulting distribution of R_h and R_w is tabulated. The moments used in the estimation are calculated by comparing observed tabulations to the simulated distributions.

The moments used in the estimation are chosen to provide identifying information on the parameters. The moments include, for both husbands and wives, the percentage retired at five ages (55, 58, 60, 62, and 65) and the corresponding percentage retired among the following subgroups: respondents in poor health, respondents whose spouses are in poor health, respondents born before 1934, respondents born after 1938, respondents in the highest third of potential family income, and respondents in the lowest third of potential family income. There are five more moments for the percentage of couples jointly retired in each of the five surveys, plus five more each for joint retirement among couples whose husbands enjoy time with their wives and whose wives enjoy time with their husbands. Finally, there are two moments if the husband or wife reported a retirement between the last survey before turning 62 and the first survey on or after age 62, plus three more moments each for both the husband and wife corresponding to retirement at ages 58 and 65, and to retirement corresponding to the age of early retirement benefits if the respondent has a pension.

VI. Estimates

Table 3 gives the estimated parameters for the model, along with the asymptotic t-statistics. There are several notable things to say about these estimates. First, consider the variables related to whether the spouse is retired or not. To estimate this coefficient for the husband, a binary variable representing whether the wife is retired or not is entered into the $X\beta$ vector for the husband. A positive value of the coefficient indicates that if the wife is retired, the value of retirement is increased for the husband. Two separate coefficients are estimated for this effect, depending on whether or not the husband said that one of the advantages of retirement is that he could spend more time with his wife. Two symmetric coefficients are also estimated for the wife. For the husband, there is an almost significant effect for the wife being retired if the husband says that time with the wife is important. The magnitude of this effect is equal to about three years of age. If the husband does not feel quite so strongly about time with the wife, the effect of her retirement on the value of his retirement is about half as great. The story for the wife is different, however. Whether or not she says that time with the husband is important in retirement, the effects of his retirement on the value of her retirement are very small. This result, that the husband likes to have the wife around when he retires but not vice versa, reconfirms the results of previous work using simpler models (Gustman and Steinmeier, 2000 and forthcoming; Coile, 1999).

Both spouses, not surprisingly, strongly increase their preferences for retirement if their health becomes poor. More surprising are the results of including the health of the spouse. For wives, having a husband in poor health increases the value of her own retirement, and the effect is about two-thirds of the effect if her own health is poor. For husbands, the effect of having a wife in poor health is to reduce strongly the value of his own retirement. In other words, wives appear willing to care for their sick husbands, even to the point of advancing their own retirement, but husbands will significantly delay retirement in order to avoid having to care for their sick wives.[5] Part of this effect may be an income effect, i.e., the income needs increase when the wife becomes sick, but this explanation does not appear to work when it is the husband who becomes sick.

Unfortunately, by themselves the coefficients in Table 3 do not do a very good job of conveying whether the model can reproduce the tendency of couples to retire together. For this, we turn to simulations with the model. The simulations are almost identical to the simulations which are used to generate the moments in the estimation. For each couple in the sample, the value of time preference is calculated in the same manner as discussed in the estimation. 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 described earlier, and the results are tabulated for all the draws and all the individuals in the sample. These are the results discussed below.

One of the purposes of the model is to capture the tendency of couples to retire together, and the success of this endeavor is indicated in Figure 2. Since there are five waves available in the HRS, each spouse who retired during the survey retired between two waves. For instance, the wife might have retired between the 1994 and 1996 surveys. For all couples for whom both

[5] These findings are in contrast to Johnson and Favreault (2001), who find that both men and women are less likely to leave the labor force if their spouses appear to have left the labor force because of health problems.

spouses retired during the survey period (1992-2000), we can calculate the difference between the pair of waves that the husband retired and the pair of waves that the wife retired. For instance, if the husband retired between 1996 and 1998 and the wife retired between 1992 and 1994, the spouses would have retired two periods apart. The husband would have retired in the third two-year period, and the wife would have retired in the first two-year period. If both retired between 1994 and 1996, they would have retired 0 periods apart. The number of periods separating the retirements is depicted on the horizontal axis of the figure. Positive numbers indicate that the wife retired first, while negative numbers indicate that the husband retired first.

The main feature of the observed data is that about 38 percent of those couples whose retirements are both observed retired between the same pair of surveys. The effect falls off very rapidly. Only about 15-18 percent retire in adjacent two-year periods, and even fewer retire two or more periods apart. In part, this result stems from the limited length of the sample. There are four two year periods during which both spouses might have retired, but only three pairs of adjacent two-year periods during which the spouses could have retired one period apart. Even considering this, the percentage of couples retiring during the same two-year period is high.

The darker bars indicate the percentage of the sample that is simulated to retire during the same two-year periods between 1992 and 2000. The model does appear to capture the phenomenon of joint retirement fairly well. About 40 percent of the sample for whom both spouses are simulated to retire in this period do so within the same two-year period, while about 14 to 22 percent are simulated to retire in adjacent two-year periods. The model even appears to capture the slight tendency of the husbands to retire first, which should not be too surprising given that the husbands are slightly older. For the full sample (not just those couples both of whose spouses retire in the 1992-2000 time frame), the simulations indicate that about 14 to 15 percent will retire in the same year, and for another 9 percent of the couples, the spouses will retire within one year of each other.

In Figure 3, the top half of the figure shows the percentage of husbands interviewed at even numbered ages retiring in the two year periods between surveys, and the bottom half shows the percentages for husbands interviewed at odd numbered ages. Figure 4 shows the corresponding information for the wives. The white bars indicate the observed percentages retiring at each two-year interval, calculated as differences between the appropriate percentages in Table 1. The dark bars indicate the percentages as calculated from simulations with the model using the parameter values estimated in Table 3. It is evident that the estimated percentages smooth out some of the instabilities present in the observed numbers, although the average retirement percentages appear to be of about the right magnitude.

Figure 5 presents the simulated retirement percentages by single year of age. In this diagram, the white bars are the husbands, and the bars with shading are the wives. There is a clear but somewhat muted peak in retirement at age 62. For the husbands, the percentage retiring at age 62 is about 3 percentage points higher than the average of the percentages at the adjacent ages of 61 and 63. For the wives, the percentage at 62 is about 4.3 percentage points higher than the average of the adjacent ages. In general, the wives have more retirements at ages prior to age 60 and fewer retirements after age 62 than do the husbands.

In summary, the model with joint retirement yields several interesting results. The most striking results in the coefficients are their suggestion of a clear distaste of many husbands to staying at home when their wives are in poor health, while the wives are willing to stay at home with sickly husbands. In terms of joint retirement, the model is approximately able to replicate the amount of joint retirement that is evident in the sample. The model is able to generate a peak in retirement activity at age 62. As in the data, the peak is somewhat more smeared than the results found in previous work which included one-earner families.

VII. Effects of Personal Accounts and Annuitization on Retirement in the Long Run

For couples, the current arrangement for determining family benefits is a complicated calculation involving own benefits, spouse benefits, survivor benefits, reduction factors, earnings tests, and several other assorted factors. In previous studies we have examined the effects on retirement of many of these features in the current law (Gustman and Steinmeier, 1985, 2000a). We also have examined proposals for changing Social Security rules regarding own, spouse and survivor benefits (Gustman and Steinmeier, 2002, forthcoming). With regard to the introduction of personal accounts, we have analyzed the effects of adopting personal Social Security accounts on a voluntary basis (Gustman and Steinmeier, 1998), as well as the recent proposal by President Bush's Commission to Strengthen Social Security (Gustman and Steinmeier, 2003).

In this section, we estimate the retirement effects if the couples in the HRS had been subjected over their full working lives to a system of personal accounts instead of the current Social Security system. We then ask what the retirement effects are of mandating that benefits be taken in the form of annuities, against the alternative of immediately being able to claim account balances as a lump sum once the individual reaches the early entitlement age. The simulation of the effects of annuities does require that the respondent leave enough money in the account to provide a poverty level income throughout retirement.

To simulate the long run effects of individual accounts, we ignore the very thorny transition issues from the current system to individual accounts. Thus as the alternative to the current system, we assume that the entire contributions of the OASI system were placed in personal accounts that grow at a compound rate equal to the geometric mean of a portfolio composed of half stocks and half government bonds over the last 75 years. The difference in returns highlights the limit of the potential effects of adopting individual accounts in the long run. In the initial sets of simulations, these accounts are annuitized at the same rate into an annuity that pays two-thirds to either survivor.

The accounts of both spouses are treated symmetrically, avoiding the situation where the Social Security taxes on the lower earning spouse has essentially no effect on family benefits. The disability part of Social Security, as well as the Medicare tax, are not included in either the taxation or the benefits of this system of personal accounts.

Figure 6 shows the effects on retirement at various ages for several variants of personal accounts. The figure shows only those effects at ages 60 through 64, where most of the action is. The bars at the left side of the figure show the retirement patterns of the base simulation, that is, the simulation assuming that benefits are paid according to the current Social Security formula. The next set of bars considers what would happen if the contributions, both employee and employer, that the HRS individuals had paid into the OASI system had instead been put into a system of personal accounts, with the entire balance annuitized at retirement with a joint and two-thirds annuity. The most notable feature of this counterfactual simulation is that retirement at age 62 would be about three percentage points higher.

The main factor leading to this result is that the level of benefits from the personal accounts are roughly double the benefits accrued under the current system. While both the current system and the system of personal accounts are roughly actuarially neutral between ages 62 and 65, individuals with relatively high time preference rates will perceive a larger loss from continuing at work the higher the benefits. Such individuals will give little weight to the fact that later benefits will be higher, so they will perceive their net earnings (after giving up the current benefit) will be lower the higher that benefit is. Another way to view this is to regard the foregoing of current benefits as a tax, which in fact is what the earnings test is frequently considered.

The next group of bars indicates what might happen if the contribution rate to personal accounts in earlier years had been at the same 10.6 percent that has been in effect in recent years. Since the contribution rate made by HRS respondents was considerably lower in their earlier years of participation, the effect of a uniform 10.6 rate is to raise considerably the amounts in the personal accounts. In some sense, this simulation suggests what will happen in the future as future generations are subject to the higher contribution rate for a longer period of time. As might be expected, the larger amounts in the personal accounts, which translate to higher benefits, accentuate the effects found in the previous simulation. The age 62 peak rises by another 2 percentage points, and is now almost 5 percentage points higher than in the base simulation.

Some personal account proposals, in the spirit of giving individuals more control over their assets, permit part of the accounts to be taken as lump sums. In general, since part of the purpose of Social Security is to avoid poverty in old age, such proposals permit lump sums only to the extent that the remaining annuities are above an amount related to the poverty level. The last two groups of bars in Figure 6 pertain to simulations, one for historical contribution rates and the other for uniform 10.6 percent contribution rates, in which part of the personal accounts can be taken as a lump sum at retirement. The limitation is that the remaining annuity for *each* spouse must be at least as great as the poverty level for a single individual over 65, which in 1992 was $6,929. If the individual's annuity were to be less than this amount, no lump sums would be allowed. In working couples, the minimal benefit amount before both spouses could take lump sums would be $13,858.

The ability to take lump sums at retirement is extremely attractive for individuals with high time preference rates. As seen in Figure 6, the ability to take lump sums increases retirement at age 62 by roughly 5 percentage points relative to the simulations with personal accounts but no lump sums. This suggests that including lump sums in the system can have a very large effect, at least if the ability to take the lump sums depends on retirement. A system which includes a uniform 10.6 percent contribution rate and includes the ability to take lump sums at retirement would produce about a 10 percentage point increase in retirement at age 62 relative to the current system.[6]

Figure 6 looks at the percentage of workers who retire at any given age. Another perspective on the effects of the personal accounts is to ask how they change the percentage of those who are already retired at particular ages. Prior to doing this, however, Figure 7 indicates the percentage retired by age in the base simulation. About half of the husbands are retired by age 62, while half of the wives are retired at age 60.

Figure 8 shows how the personal accounts affect the retirement percentages of the husbands, relative to the base simulation. Whatever the nature of the personal accounts, all of the simulations show that the effects on retirement begin before the eligibility for benefits begins. The number of retirees at age 60 is 2 to 3.5 percentage points higher, and the number at age 61 is 2 to 4.5 percentage points higher. Evidently the prospect of higher benefits under the personal accounts enables individuals with limited savings to retire before even reaching age 62,

[6] All of the simulations assume that the maximum amount of earnings subject to Social Security taxes remains at historical levels. If the current maximums were to be projected backward, the effects of the counterfactual personal accounts would be even larger.

since there is less need for savings to supplement Social Security benefits after age 62.

At age 62, however, the effects of personal accounts on retirement explode. Even the program with the lowest effect, which uses historical contributions and does not allow any lump sums, increases the percentage retired by about 5.5 percentage points. A program which uses uniform contribution rate of 10.6 percent over the lifetime and allows for lump sum withdrawals increases the percentage retired at age 62 by a whopping 15 percentage points over the base simulation with current benefits. The move to a uniform contribution rate by itself appears to increase age 62 retirees by around 3 percentage points, and the addition of lump sums adds another almost 6 percentage points to retirement at that age. Nor do the effects end there. The effects of the personal accounts on the percentage retired diminish only slowly after age 62; by age 65 the percentage retired is still 4 to 8 percentage points higher, depending on the exact nature of the personal accounts.

Figure 9 reflects the effects of the same personal account proposals on the wives. In general, increases in retirement percentages are not quite as large as for the husbands. Part of the explanation for this is that at any age, substantially more of the wives are already retired, so there is less scope for any measure to increase retirement even more. Another part of the explanation may be that the personal accounts have caused the husbands to retire earlier, which reduces family income. This reduction in income may in turn somewhat offset of the incentives toward increased retirement. In any case, even though the retirement effects are smaller for the wives than for the husbands, they are still substantial.

Overall, these are huge increases in retirement. They are particularly discouraging if part of the purpose of introducing personal accounts is to reduce the incentives to retire early that are inherent in the current system. The magnitude of the personal account benefits, coupled with the fact that they can be received only after retirement, provides a strong incentive to retire when the benefits are available, especially for households with relatively high time preference rates. The incentives are considerably compounded if the possibility exists of receiving substantial lump sums upon retirement.

One way to mitigate these retirement effects is to reduce the contribution rates so that the benefits are roughly comparable to benefits under the existing system. Indeed, a simulation which cuts the contributions (and hence the benefits) in half, and uses historical contribution rates with no lump sum withdrawals allowed, yields retirement probabilities less than one percentage point different from the base simulation. The drawback of doing this, of course, is the possibility that a period of below average returns in the stock market would leave beneficiaries with considerably less than they have now and would increase the probability of poverty among the aged. A secondary consideration is that so doing would merely reduce the disincentives to continued work to the present levels; it would not eliminate them.

Another way to mitigate the retirement effects would be to make the receipt of benefits, both the annuity and the lump sum, independent of employment. This would mean that individuals could collect benefits at age 62 whether or not they continued working. The drawback to this approach is that people would be able to spend part of the funds that are intended for retirement while they are still working, and this would lower the amount of benefits paid during retirement. This is essentially the consideration that prevented Congress from abolishing the earnings test prior to 65 when it abolished the earnings test for recipients over 65. A secondary drawback of this strategy is that family income would drop more precipitously upon retirement if benefits are started before retirement.

The main problem is that eliminating disincentives to further work is fundamentally inconsistent with paying benefits that depend on current earnings from work. This is true under the current system, and it would be true under personal accounts as well. For those with relatively high time preference rates, any delay in receiving benefits is costly in terms of utility, even if future benefits are increased in an actuarially fair manner.

VIII. Conclusions

Retirement behavior differs between one-earner and two-earner couples sufficiently that these two groups should be analyzed separately, especially in the context of studies of behavioral responses to current or potential policies. To model retirement outcomes in two-earner families, one must include three dimensions of behavior: the joint determination of retirement and saving; heterogeneity in time preference; and the interdependence of retirement decisions of husbands and wives.

Among the specific results, the model is approximately able to replicate the amount of joint retirement that is evident in the sample. We continue to find an asymmetry in which husbands prefer their wives to be retired before they retire. Moreover, there is a clear distaste of many husbands to retiring when their wives are in poor health, while the wives are willing to stay at home with sickly husbands. The model is able to generate a peak in retirement activity at age 62. As in the data, the peak is more smeared for husbands in two-earner families than was true for previous analyses of husbands in general, taking the wives earnings as exogenous.

Using a model that incorporates the three key dimensions of behavior, we analyzed the effects of introducing a system of personal accounts instead of the current Social Security system. If returns to individual accounts matched the returns observed in historical data, retirements by two-earner couples would take place much earlier than observed under the current

Social Security system. The primary reason they would retire earlier is that for workers with high discount rates, the reward to continued work is significantly reduced once the individual becomes eligible to receive the higher benefits generated by individual accounts. For similar reasons, allowing lump sum settlements on personal accounts would substantially accelerate retirement rates.

Altogether, our findings suggest that a system which includes a uniform 10.6 percent contribution rate over the lifetime and allows the individual to make lump sum withdrawals at retirement instead of annuitizing benefits (above poverty level) would produce about a 15 percentage point increase in retirement for husbands at age 62 relative to the current system. By itself, allowing individuals to opt out of annuitization and take a lump sum instead, increases retirements at 62 by about 6 percentage points. These effects are also observed at higher ages.

References

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

Blau, David M. and Donna B. Gilleskie. 2001. "Health Insurance and the Retirement of Married Couples". Mimeo.

________________. 2003. "The Role of Retiree Health Insurance in the Employment Behavior of Older Men". NBER Working Paper 10100.

Coile, Courtney. 1999. "Social Security, Pensions and the Retirement Decisions of Couples". Mimeo.

Diamond, Peter A. and Jerry A. Hausman. 1984. "Individual Retirement and Savings Behavior". *Journal of Public Economics* 23: 81-114.

French, Eric. 2002. "The Effects of Health, Wealth , and Wages on Labor Supply and Retirement Behavior". Xeroxed. Federal Reserve Bank of Chicago.

Greene, William H. 2000. *Econometric Analysis*, Fourth Edition. Upper Saddle River: Prentice Hall.

Gustman, Alan L. and Thomas L. Steinmeier. 1985."The 1983 Social Security Reforms and Labor Supply Adjustments of Older Individuals in the Long Run". *Journal of Labor Economics* 3(2): 237-253.

________________. 1991. "Changing Social Security Rules For Work After 65: Proposed Policies And Their Effects", *Industrial And Labor Relations Review* 44(4):733-745.

________________. 1998. "”Privatizing Social Security : First Round Effects of a Generic, Voluntary, Privatized U.S. Social Security System". In Martin Feldstein, editor. *Privatizing Social Security*. Chicago: University of Chicago Press: 313-357.

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

________________. 2001. "How Effective Is Redistribution Under the Social Security Benefit Formula". *Journal of Public Economics*. Volume 82(1), October, pp. 1-28.

________________. 2002. "The Social Security Early Retirement Age In A Structural Model of Retirement and Wealth". NBER Working Paper 9183. September.

________________. 2003. "Retirement Effects of Proposals by The President's Commission to Strengthen Social Security". NBER Working Paper 1030. October.

________________. Forthcoming. "Social Security, Pensions and Retirement Behavior Within the Family". *Journal of Applied Econometrics*. NBER Working Paper 8772.

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

Johnson, Richard W. and Melissa M. Favreault. 2001. "Retiring Together or Working Alone: The Impact of Spousal Employment and Disability on

Retirement Decisions". Center for Retirement Research at Boston College. March.

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

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

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

Table 1
Derivation of Sample from HRS Data Set

	Observations Lost	Observations Remaining
Married couples in 1992, both interviewed		4767
At least one spouse divorced since age 35	746	4021
No career job		
Husband	461	3560
Wife	1765	1795
Ages for retirement are inconsistent		
Husband	161	1634
Wife	81	1553
Ambiguity about whether jobs are social security covered		
Husband	25	1528
Wife	10	1518
FT years unavailable in wave 3 or social security record		
Husband	0	1518
Wife	107	1411
No FT earnings in social security record or self report		
Husband	11	1400
Wife	13	1387
No self reported earnings, and social security earnings over limit		
Husband	10	1377
Wife	9	1368
Relatively large business assets	112	1256
No Pension Provider record in last job		
Husband	302	954
Wife	187	767
Final sample (couples)		767

Table 2
Observed Retirement Rates

Age	Percent Retired: Husbands	Wives
54	13.9	18.7
55	13.0	22.5
56	15.6	28.2
57	15.5	28.6
58	24.5	47.7
59	27.8	39.1
60	31.8	47.2
61	31.9	46.6
62	47.4	68.2
63	60.6	74.4
64	67.5	81.4
65	75.7	84.9
66	77.0	93.1
67	82.1	100.0

Table 3
Parameter Estimates

	Coefficients	t-statistics
α: Coefficient of Consumption	-0.19	1.25
Husband's Parameters		
β_o: Constant	-10.032	36.56
β_a: Age	0.397	2.26
Wife Retired for Those Who		
β_{se}: Enjoy Time with Wife	2.84	1.82
β_s: Do Not Enjoy Time with Wife	1.43	1.23
β_h: In Poor Health	4.26	3.40
β_{sh}: Wife in Poor Health	-4.76	2.34
β_v: Vintage (Birth Cohort)	-0.06	0.45
σ_ε: Standard Deviation of ε	5.43	8.07
Wife's Parameters		
β_o: Constant	-10.078	47.78
β_a: Age	0.172	3.88
Husband Retired for Those Who		
β_{se}: Enjoy Time with Husband	0.51	0.30
β_s: Do Not Enjoy Time with Husband	0.12	0.07
β_h: In Poor Health	2.60	2.26
β_{sh}: Husband in Poor Health	1.78	1.73
β_v: Vintage (Birth Cohort)	0.16	1.72
σ_ε: Standard Deviation of ε	3.47	5.36
ρ_ε: Correlation of ε_h and ε_w	0.23	0.71
q-statistic	88.42	
Number of Observations	767	

Figure 1
Observed Retirements at Ages 60-64 by Group

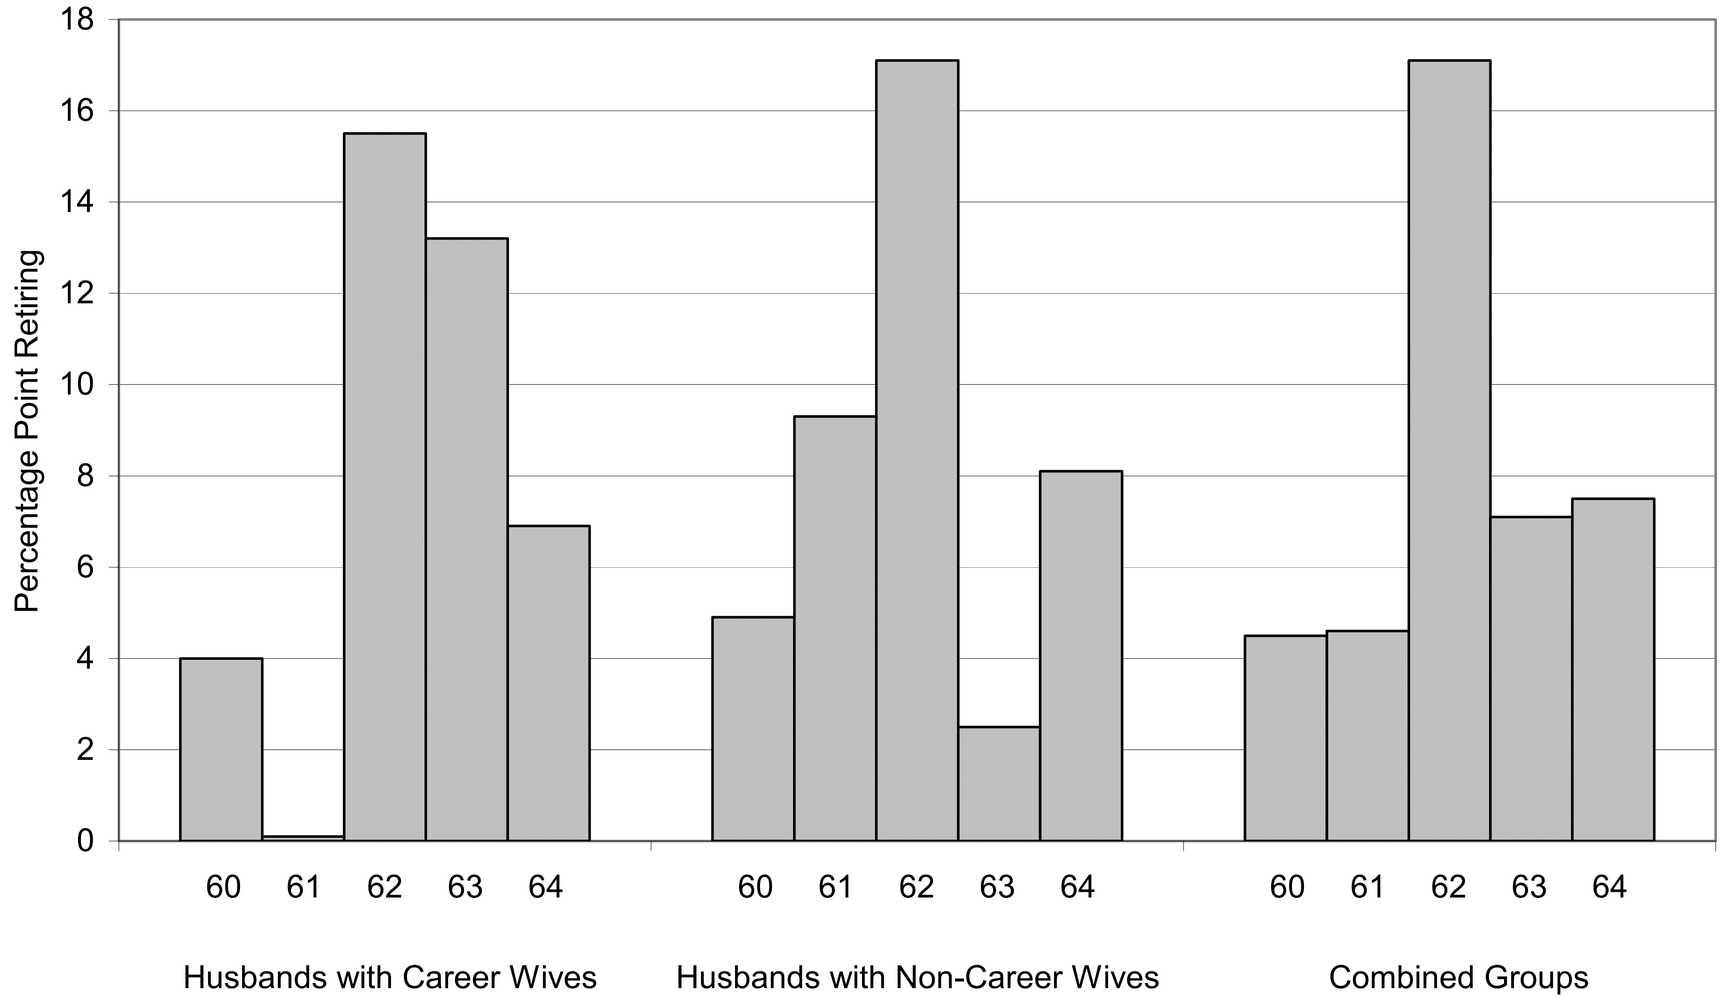

Figure 2
Differences in Retirement Waves

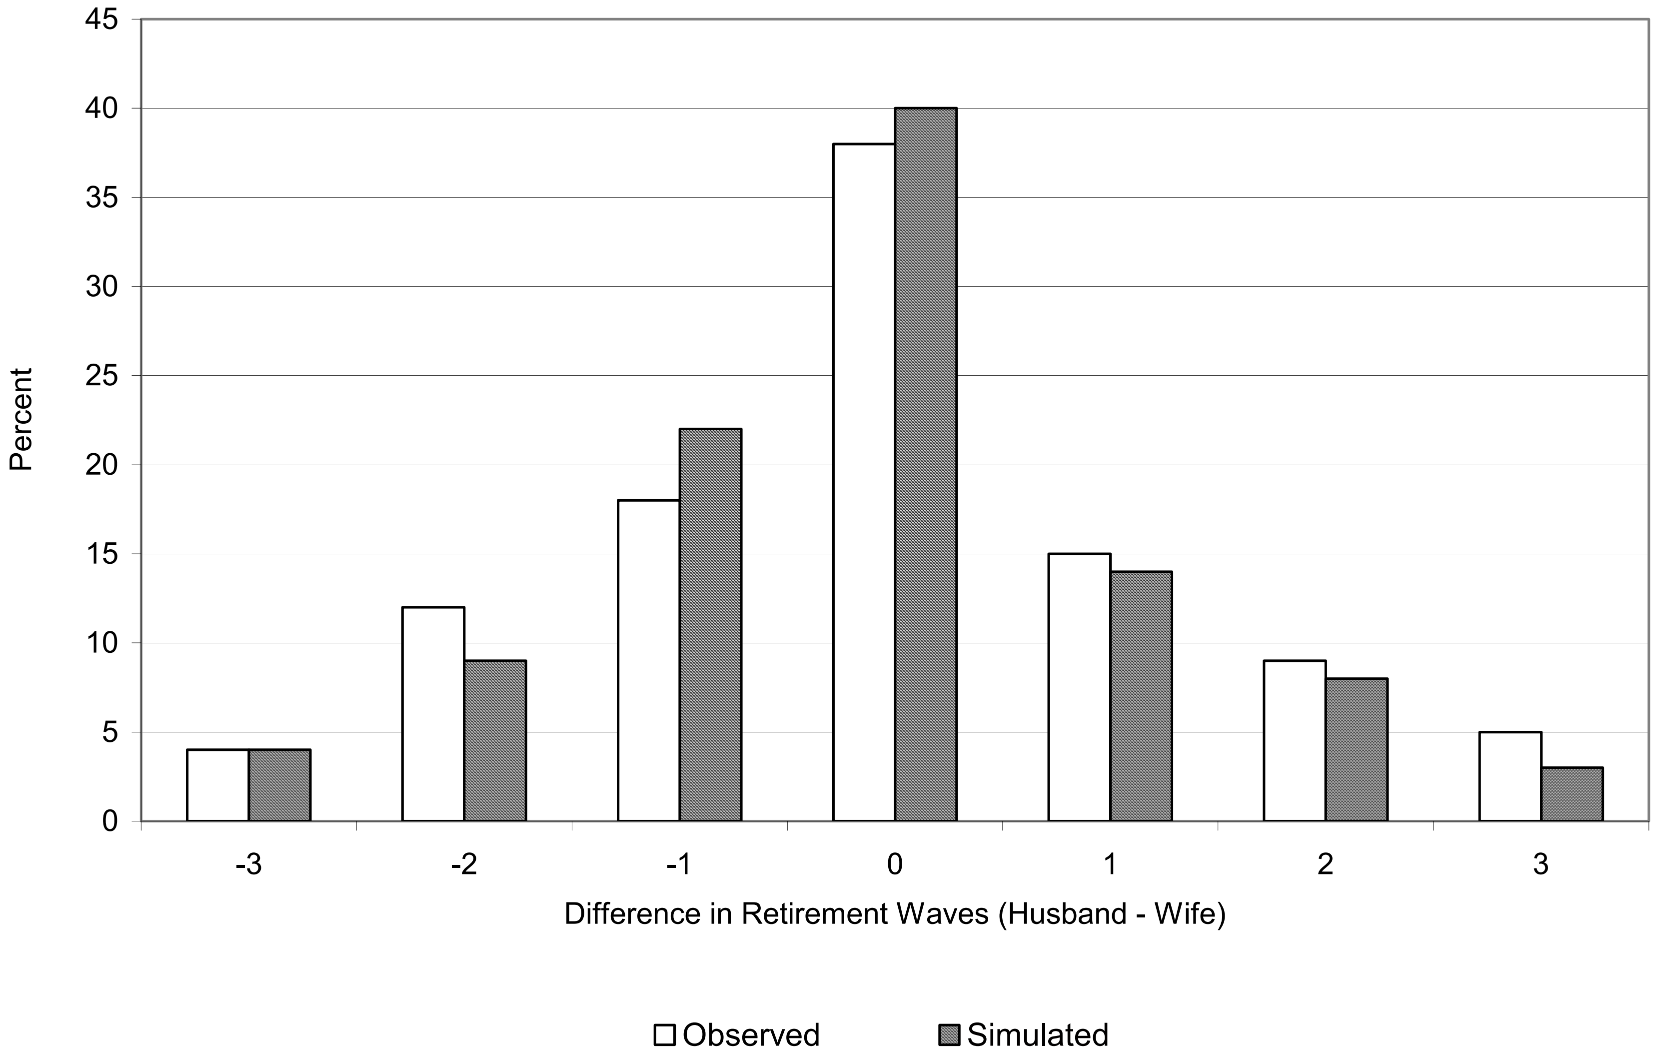

Figure 3
Two Year Retirement Rates

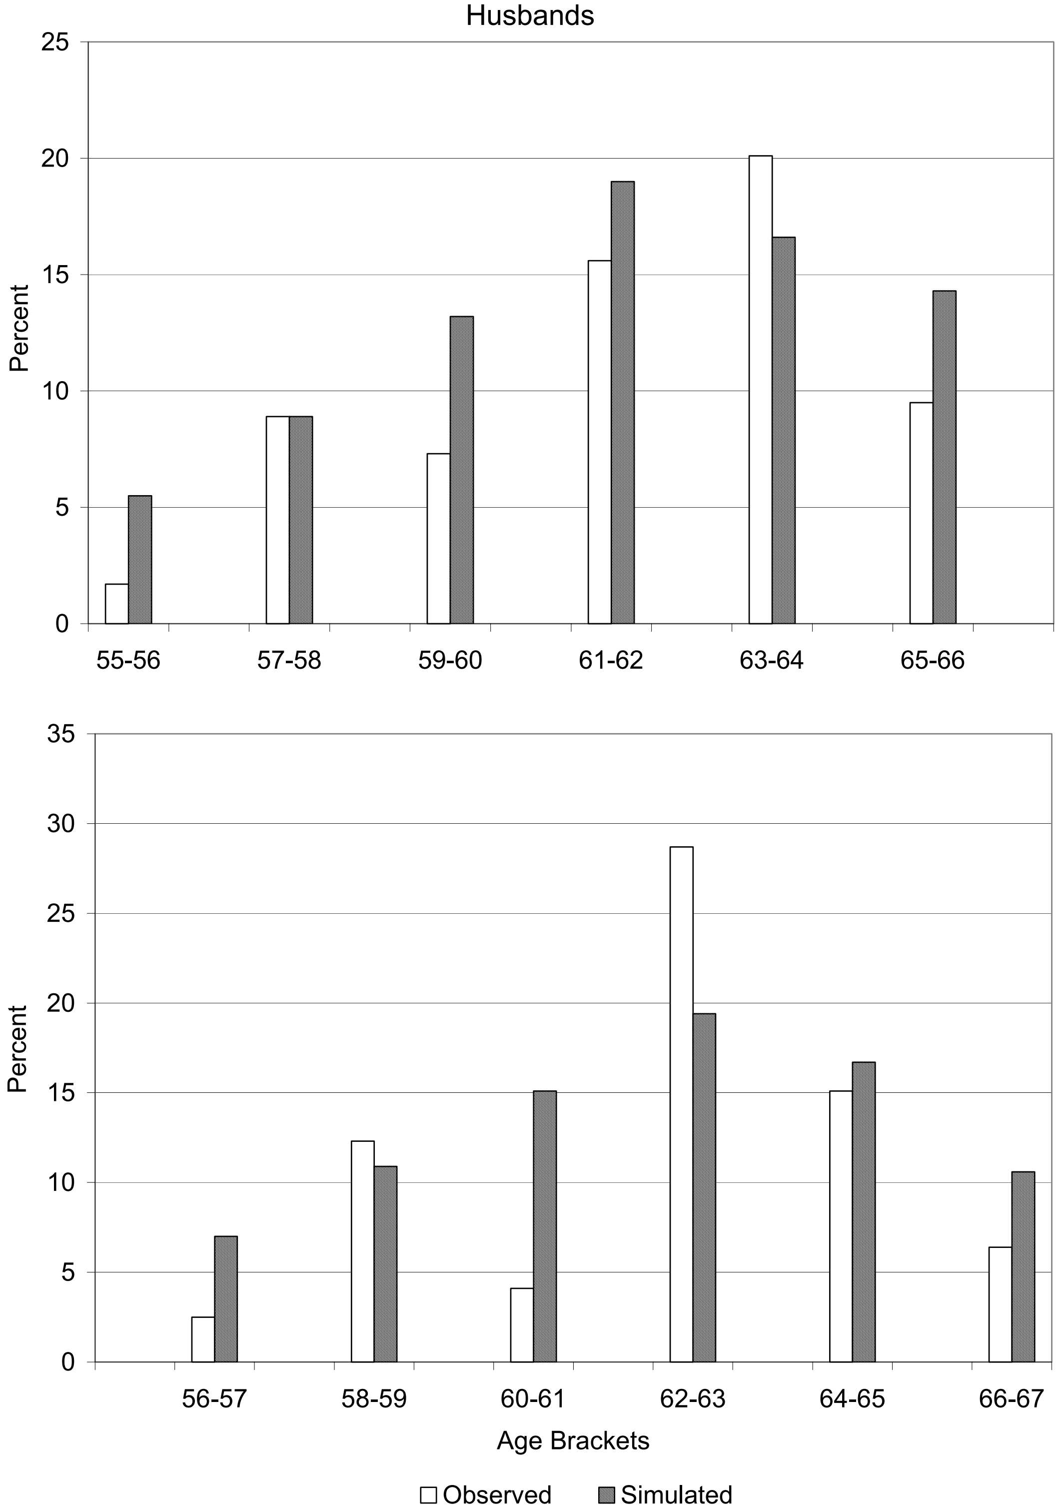

Figure 4
Two Year Retirement Rates

Wives

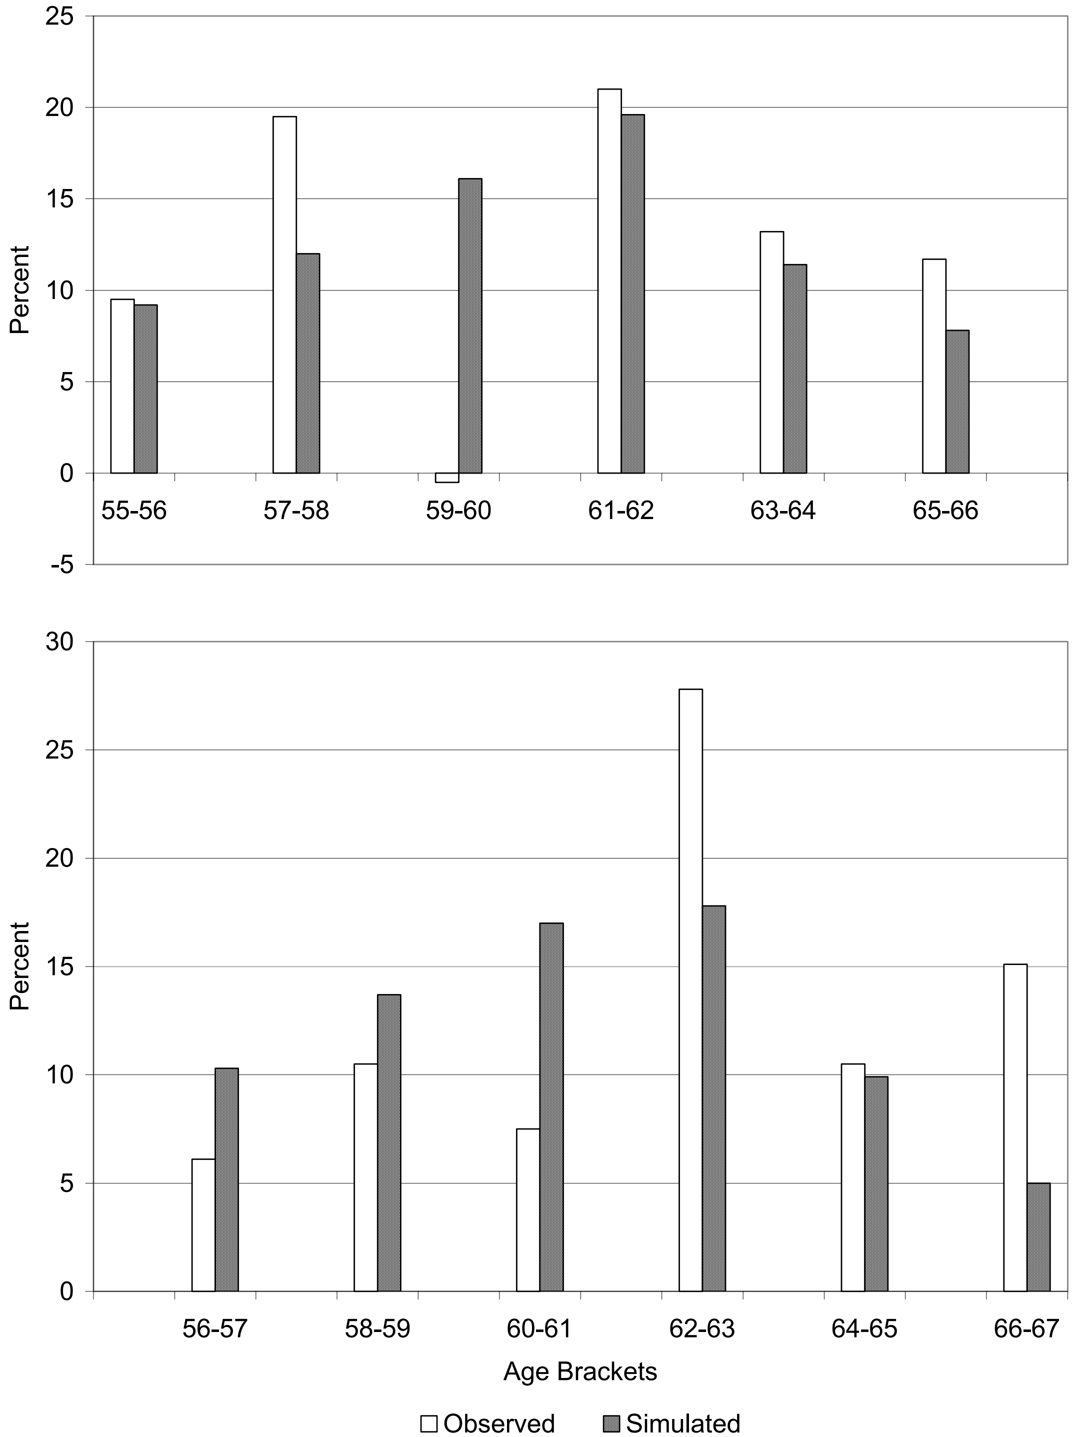

Figure 5
Simulated Retirement Percentages

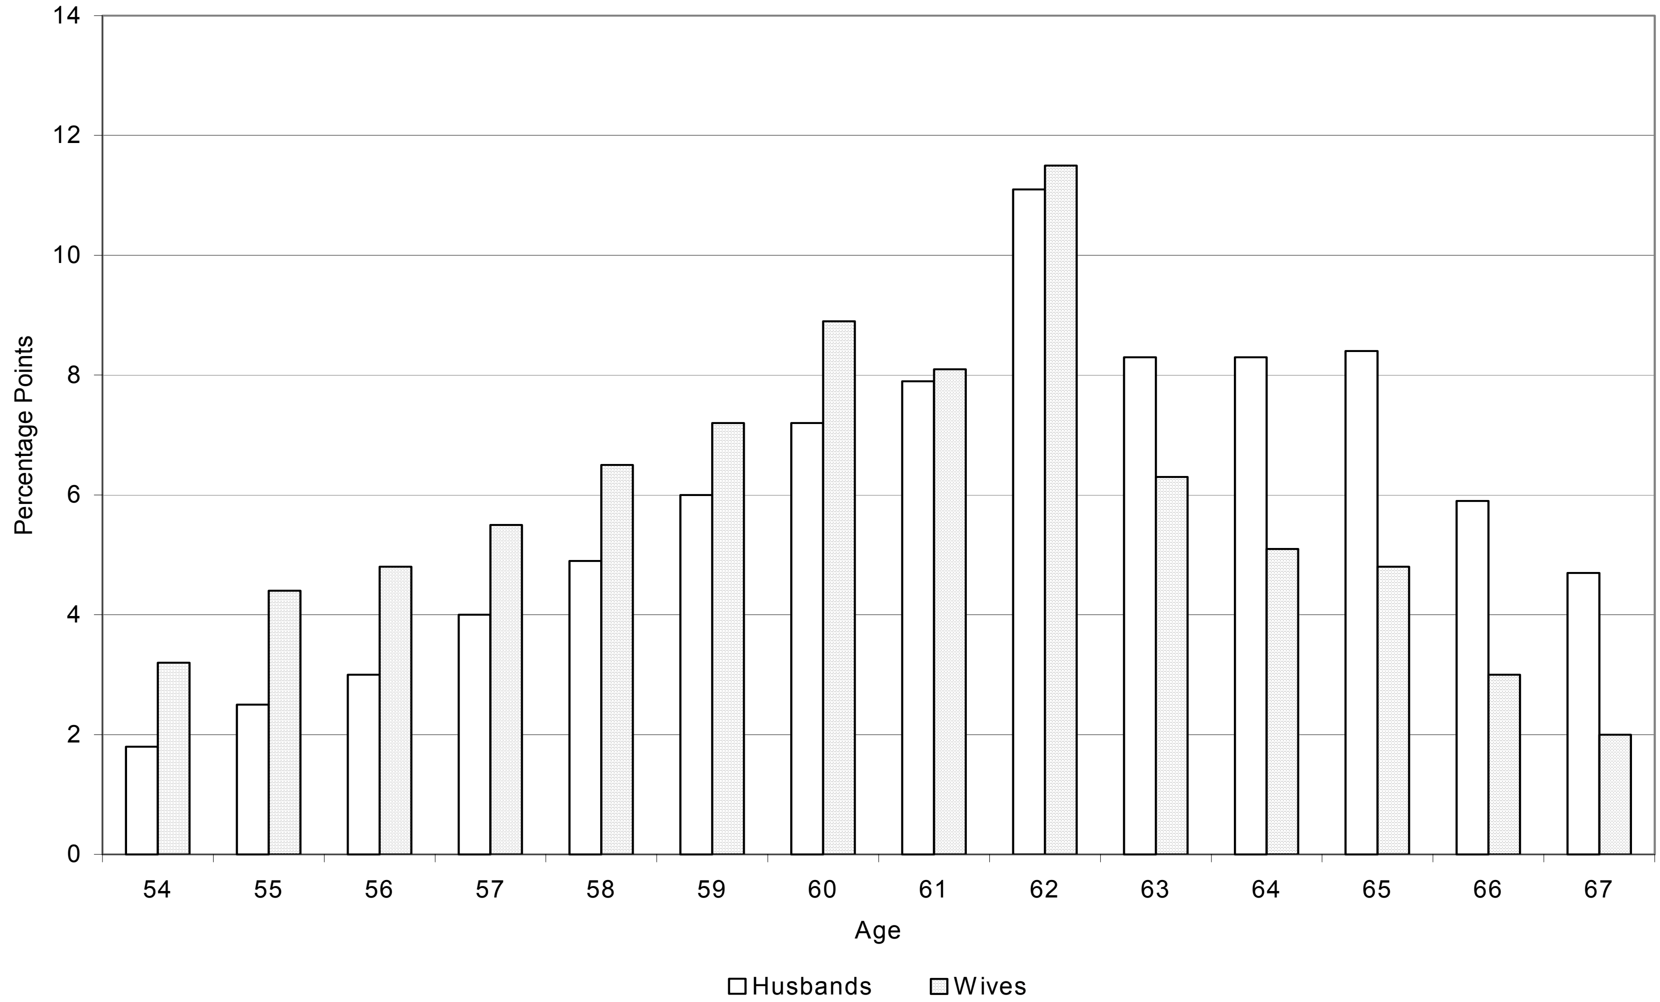

Figure 6
Retirement Percentages Under Alternative Privatization Plans, Husbands

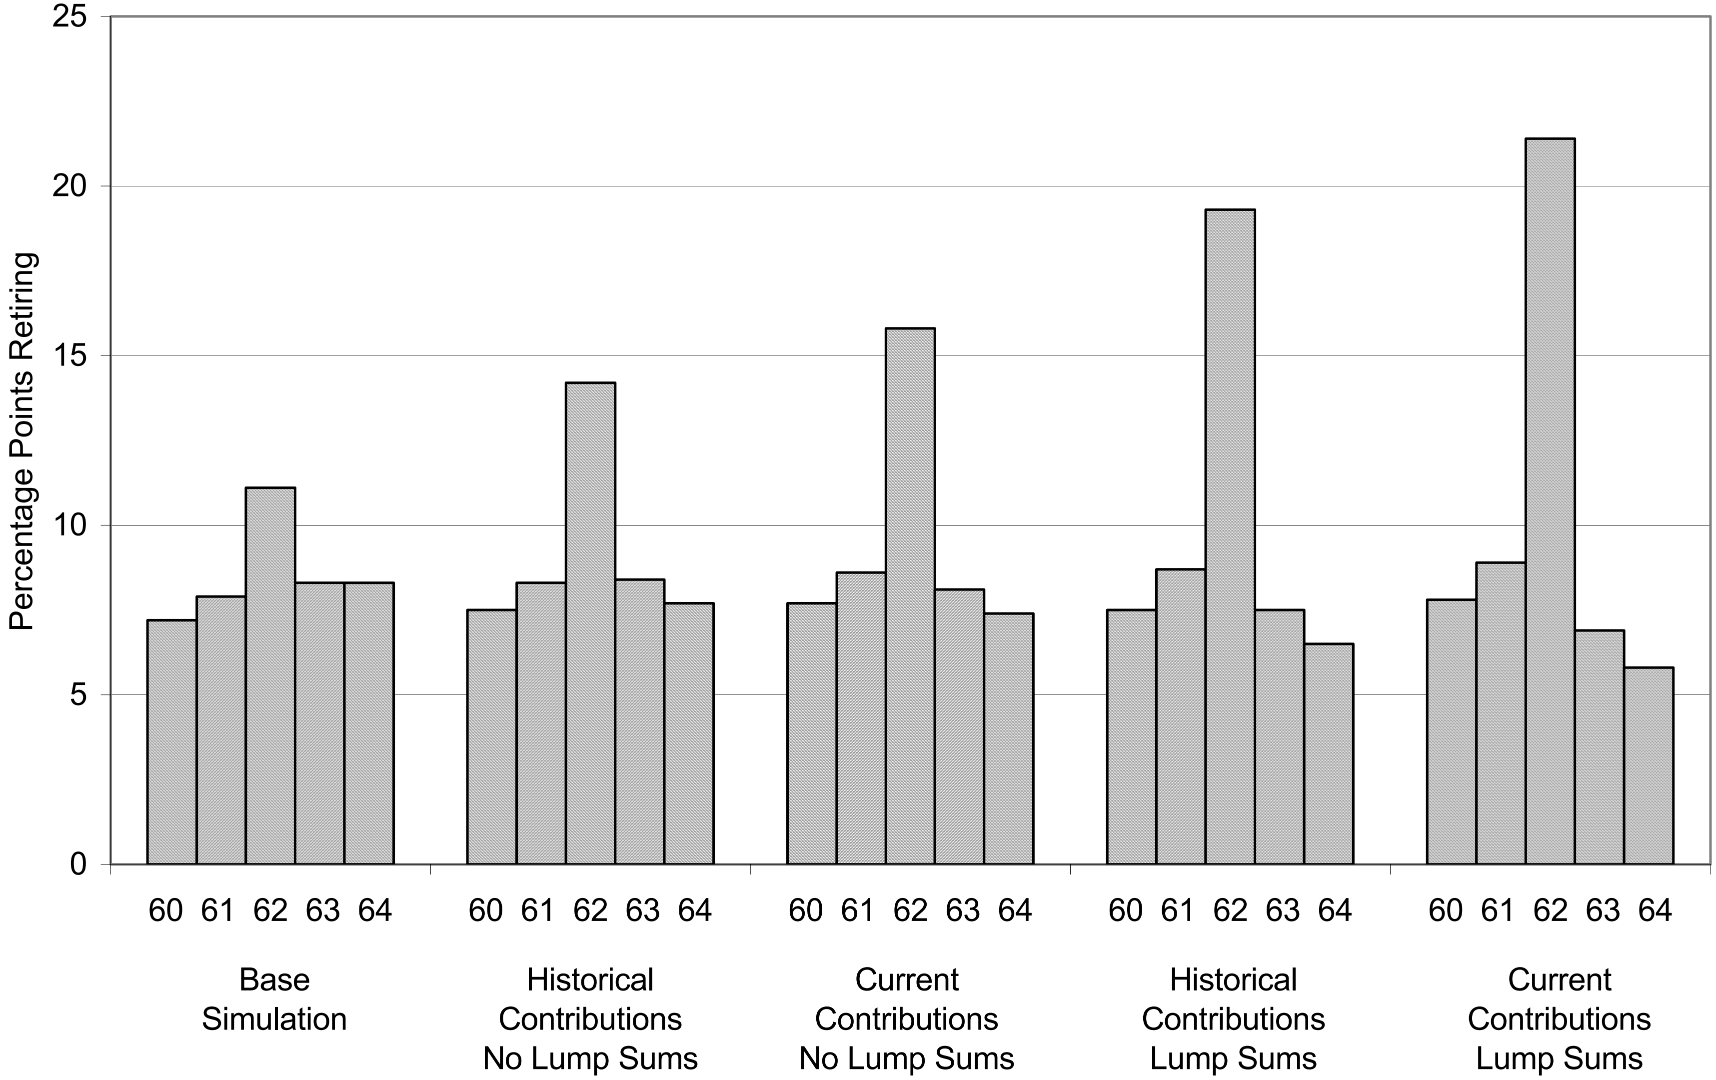

Figure 7
Percent Retired by Age, Base Simulations

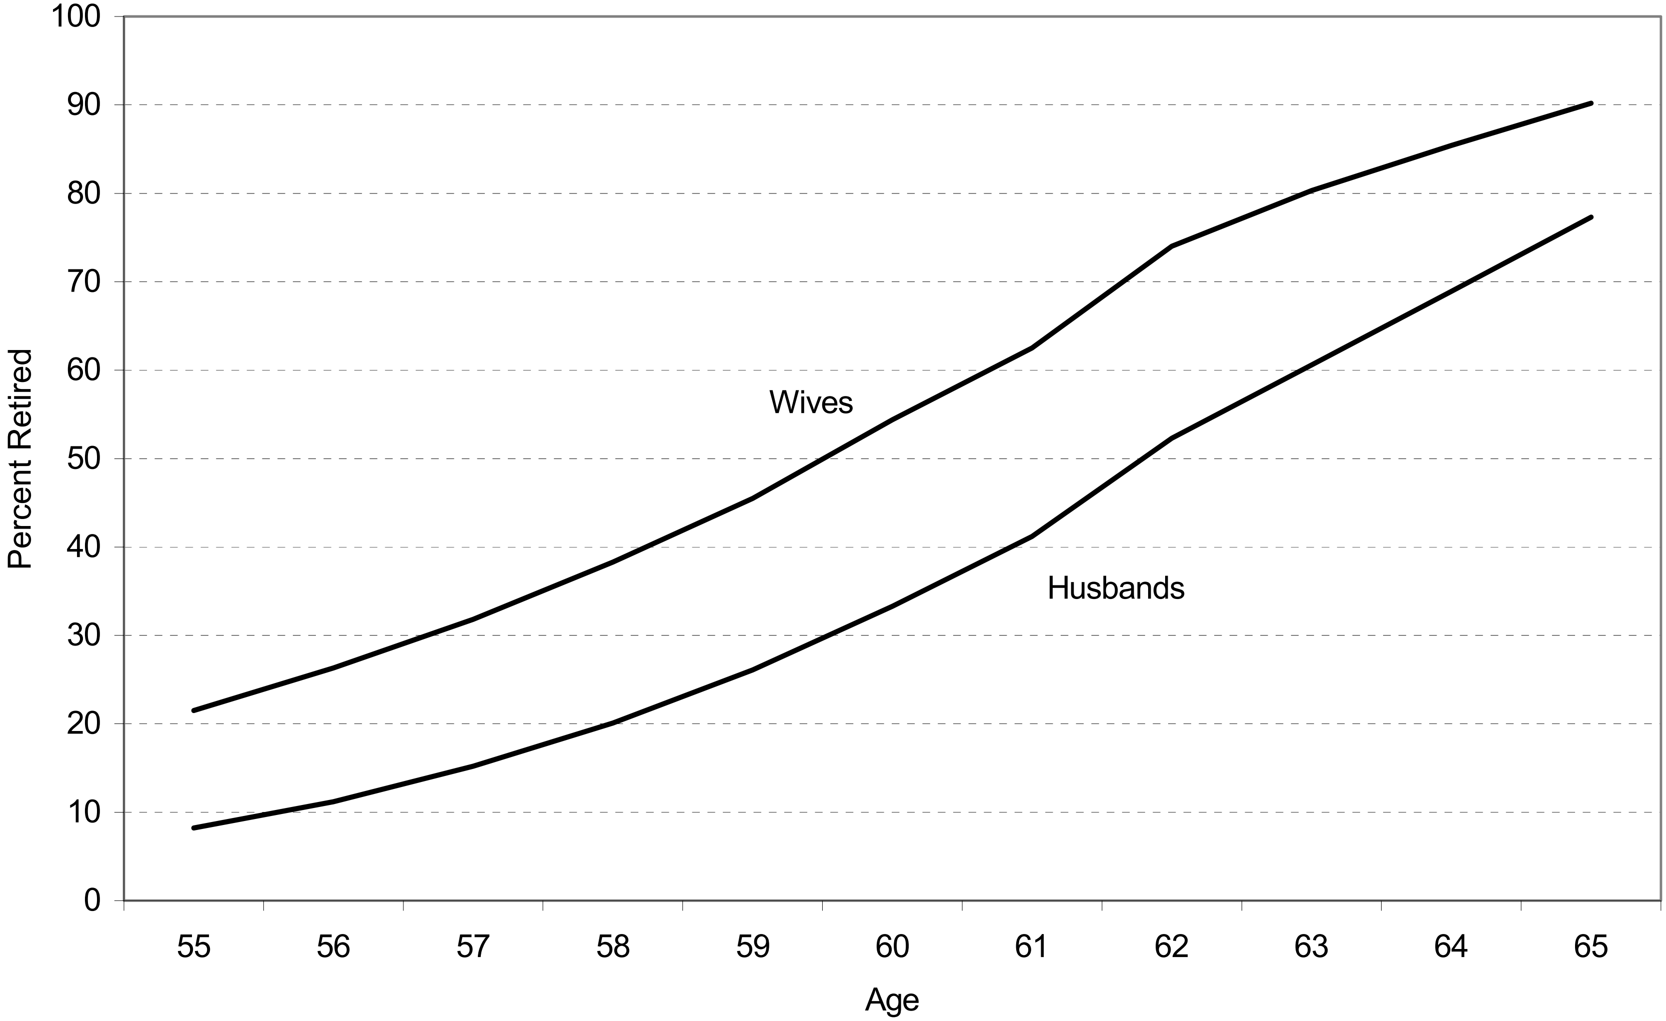

Figure 8
Percentage Point Increase Over Base Simulation, Husbands

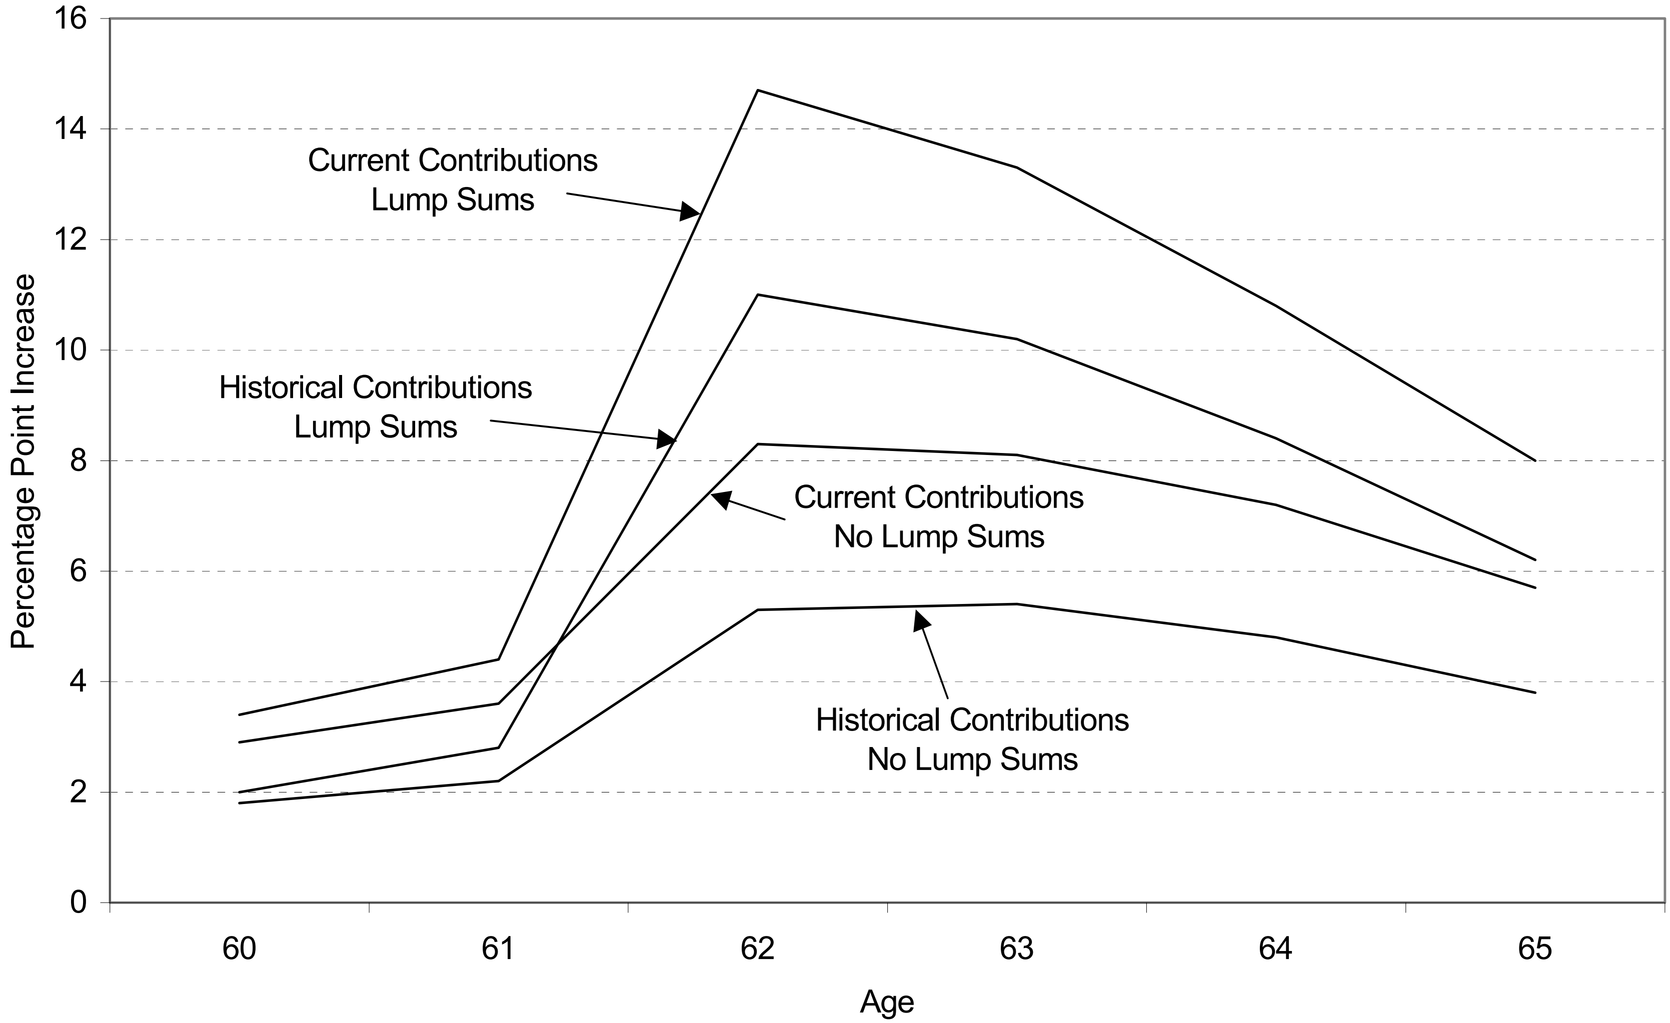

Figure 9
Percentage Point Increase Over Base Simulation, Wives