

NBER WORKING PAPER SERIES

SOCIAL SECURITY, PENSIONS AND RETIREMENT BEHAVIOR WITHIN THE FAMILY

Alan L. Gustman
Thomas L. Steinmeier

Working Paper 8772
<http://www.nber.org/papers/w8772>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2002

This is a revised version of a paper presented at a Conference on Social Insurance and Pension Research, Aarhus Denmark, November 16-18, 2001. The research reported herein using data from the Health and Retirement Study was performed pursuant to a grant from the U.S. Social Security Administration (SSA) to the Michigan Retirement Research Center, with a subcontract to the National Bureau of Economic Research. The estimates using data from the National Longitudinal Survey of Mature Women were performed under a grant from the U.S. Department of Labor, Bureau of Labor Statistics, to the National Bureau of Economic Research and to Dartmouth College. This study has also benefitted from work done while estimating a structural retirement model in a project for the National Institute on Aging (1R01AG13913-01A1), and from the work on pensions under NIA grant (1R03AG15224-01). The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research, SSA, NIA, BLS or any other agency of the Federal Government, or the Michigan Retirement Research Center.

© 2002 by Alan L. Gustman and Thomas L. Steinmeier. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Social Security, Pensions and Retirement Behavior Within the Family
Alan L. Gustman and Thomas L. Steinmeier
NBER Working Paper No. 8772
February 2002
JEL No. J26, H55, D91, J14, J16, J32

ABSTRACT

This paper estimates a structural model of family retirement using U.S. data from the Health and Retirement Study (HRS) and from the National Longitudinal Survey of Mature Women. Estimates using the HRS benefit from having, for each spouse, earnings histories provided by the respondent and the Social Security Administration, and employer provided pension plan descriptions. We find that a measure of how much each spouse values being able to spend time in retirement with the other accounts for a good portion of the apparent interdependence of the retirement decisions of husbands and wives. When we include this measure, the simulations almost double the frequency of predicted joint retirements. Once estimated, we use the model to investigate the labor supply effects of alternative social security policies, examining the effect of dividing credit for earnings evenly between spouses, or of basing social security benefits on the amounts accumulated in private accounts. Both policies change the relative importance of spouse and survivor social security benefits within the household and both raise the relative reward to work later in the life cycle. The incentives created are modest, and retirement responds accordingly. Nevertheless, at some ages, such as 65, there may be as much as a 6 percent increase in the old age work force under privatized accounts.

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

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

I. Introduction

This paper estimates a structural model of family retirement using U.S. data from the Health and Retirement Study (HRS). It provides further insight into household retirement decision making and the reasons for interdependence in the retirement decisions of each spouse. Improvements in HRS data and matched employer provided pension histories allow more precise identification of key parameters governing interdependent behavior within the household. Once estimated, we use the model to investigate the labor supply effects of alternative social security policies, examining the effect of dividing credit for earnings evenly between spouses, or of basing social security benefits on the amounts accumulated in private accounts.

There are a number of reasons why economists are interested in the process of retirement decision making within the family. Most fundamentally, it is not possible to understand the retirement decision of one spouse without considering the behavior of the other. Aspects of consumption are joint and so is family wealth. The valuation of one's own leisure may depend on the amount of the spouse's leisure. Moreover, the reward to work for one spouse may depend on the labor market activities and work history of the other.¹ More generally, analysis of joint retirement decisions may further our understanding of other dimensions of behavior within and by the family unit.

Yet much of the research on retirement behavior has focused on the individual.² To the

¹For example, the increase in the value of the social security benefits that accrues to the family from the work of one spouse depends on the work history and employment of the other spouse (Gustman and Steinmeier, 2001b). We analyze the effects of interactions within the social security system and some alternative structures in Section VIII.

²Among the studies focusing on the individual's retirement decision are: Burtless and Moffitt (1984), Fields and Mitchell (1984), Gustman and Steinmeier (1986a and b), Stock and

extent that important interactions between spouses have been ignored, the retirement decision will be misunderstood, and so will the roles played by public policies and private retirement programs.

One reason retirement studies may have focused on individuals rather than families is that estimation of many models of family behavior may be clouded by a lack of identifying instruments. Fortunately, however, in the case of the retirement decision, the reward structure is shaped not only by the wage, but by the accrual rates in the values of the pension and social security. Thus each spouse may face a sharply different retirement incentive. As a result, it is easier to identify the elements fostering interdependence in the family's retirement decision making than it is to identify the interactions governing other decisions reached within the family. Nevertheless, an increase in our understanding of how each spouse's retirement decision interacts with the other's will promote a clearer understanding of other decisions made within the household, including bargaining between spouses in the course of household production, saving and even the formation and break up of the household.

The location and size of each spike in the pension benefit accrual profile depends idiosyncratically on the date of hire and age at hire, and varies among pension plans, and thus differs between spouses.³ Spikes in accrual profiles for defined benefit plans may be very large,

Wise (1987), Berkovec and Stern (1991) and Lumsdaine, Stock and Wise (1990, 1992, 1994, 1996). A few studies, such as Pozzebon and Mitchell (1989), examine the retirement decisions of both spouses together. Most, such as recent examples by Coile (1999) and Johnson and Favreault (2001), are reduced form. For a survey, see Lundberg (1999).

³Typically, a defined benefit pension is a plan that provides a benefit based on the length of tenure on the job, annual earnings in the last few or highest few years of work, and the chosen age of retirement. At normal retirement age, a plan might provide a benefit equal to say 1.5 percent of the average of last three years of earnings times years of service. Most often, such

equaling or exceeding the wage for working another year. In contrast to wages and other benefits that are closely related to the wage, which accrue relatively smoothly over time, the sharp spikes in pension accruals break the close relationship between the substitution and wealth effects, and hence facilitate identification of these effects. If a person responds strongly to economic rewards, he or she is unlikely to retire in the few years before becoming eligible for an early retirement benefit, at least in the absence of a strong outside influence such as a bout of ill health.

Further aiding identification, there is reason to believe the incentives created by pensions and social security are truly exogenous to the individual decision maker. Those with pensions have only limited turnover from their jobs (Gustman and Steinmeier, 1993, 1995). Accordingly, by the time they retire, pension covered workers typically have a long tenure. There have been many large changes in pensions during the course of the work lives of those pension covered workers in the U.S. who are now approaching retirement.⁴ Thus for those who are now within a

plans allow individuals to retire early, but only if they have met tenure and age requirements. Moreover, these plans often reduce benefits for those retiring at the early retirement age, but not on an actuarially fair basis. The effect of such reductions is to create a spike in the pension accrual profile at the early retirement age. Many such plans also provide other incentives to retire early, further enhancing the size of the spike. In contrast, a defined contribution plan provides benefits based on contributions to an account, by the employer and perhaps also by the employee. The reward structure is much smoother. Although some DC plans offer special early retirement enhancements, most DC plans do not, and thus do not generate a spike in the benefit accrual profile. Gustman and Steinmeier (1989) discuss the relevant literature and the factors determining the shape of the benefit accrual profile in the case of covered workers surveyed by the Survey of Consumer Finances. Gustman and Steinmeier (2000c) present analogous results using pension data from the National Longitudinal Survey of Mature Women and the Health and Retirement Study.

⁴Consider some of the major changes in defined benefit pensions that have occurred over the work life of those cohorts now approaching retirement. The size of the spike associated with early retirement has increased. There has been a sharp decline in the age of eligibility for early

decade of retirement age, the incentives from pensions in place at the time of retirement are very different from what they were when they first accepted their jobs. Having made their decisions to join their firms decades earlier, these strong trends mean they could not foresee what their pensions would look like at the time they were hired. The implication is that causality does not run from leisure preference to the opportunity set, but from the pension to retirement choice⁵. Analogously, the incentives from social security are largely exogenous to the decision making of the family.

In an earlier study (Gustman and Steinmeier, 2000a), we used respondent self reports describing their pension plans to estimate a structural model in which the decisions of the two spouses were combined in a possibly non-cooperative bargaining model of retirement.⁶ There is an important concern about this earlier work. At the time we wrote that paper, there was no longitudinal survey that combined information on work history and current work effort with information from respondents' employers describing the pension plans that they offered. For those covered by defined benefit plans, we had to rely on the respondent's description as to the location and size of the early retirement spike. We also had to impute benefits and the accrual

retirement, falling eight years between the late 1960s and the early 1980s (Anderson, Gustman and Steinmeier, 1999). Today, three fourths of HRS respondents with a defined benefit plan are eligible for an early retirement benefit by age 55 (Gustman and Steinmeier, 2000c). There also have been analogous changes in the normal retirement age. In addition, there has been the rise of the 401(k) plan, with the predominant plan type shifting from defined benefit to defined contribution (Gustman and Steinmeier, 1992, 2000b).

⁵For a contrasting view arguing that selection into pension plans is related to the propensity to save and to leisure preference, see Ippolito (1998).

⁶In related studies, we are extending this work to include the effects of imperfect capital markets. These studies are being conducted separately for individuals and married couples. We also are dropping the assumption of perfect foresight, using dynamic programming models to allow for unforeseen disturbances in each period.

spike using a generic formula and information on average actuarial adjustments in typical plans. Available evidence now establishes that respondents do a poor job of reporting the key pieces of information necessary to locate and determine the size of spikes in their pension accrual profile (Gustman and Steinmeier, 2001a, 2002). Indeed, respondents even report plan type with considerable error (Mitchell, 1988; Gustman and Steinmeier, 1989).⁷

The present paper takes advantage of new longitudinal surveys that link employer provided pension plan descriptions with panel data following the household through the retirement decision. These data provide a precise picture of the location and size of the spikes in the accrual profiles of defined benefit pensions, while allowing retirement behavior to be recorded in a timely fashion in the relevant wave following retirement. We rely primarily on the Health and Retirement Study (HRS). However, later in the analysis we also use data from the National Longitudinal Survey of Mature Women (NLS-MW), which now includes employer provided pension plan descriptions. Despite the availability of employer provided pension data, certain features of the NLS-MW make it less satisfactory for estimating a model of joint retirement behavior than the data from the HRS. A major shortcoming is that the husband's work history is reported by the wife, with most of the husband's work history also reported retrospectively⁸. However, the estimates based on the NLS-MW help to build a bridge to the

⁷This is also a major problem because reported plan type is used in the HRS and other surveys to determine whether the respondent is asked questions about the characteristics of a defined benefit plan or a defined contribution plan.

⁸More specifically the respondent to the NLS-MW is the woman in the household. There is an excellent earnings history, based on interviews from 1967 to the last survey, available for the woman. She also reports, but only incompletely, and for much of the period retrospectively, on the labor market activities and earnings of her spouse. Another problem with the NLS-MW data is that although we update the results based on the NLS-MW to incorporate employer

findings from our earlier study.

Once estimated, the model is used to investigate three potential channels which might generate the elevated level of instances in which the spouses retire at around the same time. The estimation also is extended to incorporate information on the preferences of each spouse regarding the value of time spent together.

A joint retirement model is suitable for analyzing the effects of alternative social security policies regarding the distribution of spouse and survivor benefits. If spouses coordinate their retirement, then policies modifying the reward to each spouse will have effects beyond those suggested by the incentives they create for each individual. The effects of social security policies will be filtered through the household decision making process, suggesting that the impact of these policies can be understood only if they are analyzed with a model of family retirement decision making. In a final section we simulate the effects of adopting individual social security accounts and benefit splitting, policies that would change the distribution of social security benefits within the household as well as the time profile of the reward to work.⁹

II. Overview of the Model

We estimate a joint retirement model that mixes noncooperation with some elements of

provided plan descriptions, pension plan descriptions were matched in the NLS-MW after a three year delay that greatly reduces the number of successful matches. In contrast, the Health and Retirement Study interviews each spouse separately, asking individually about their own labor market activities, including current and previous work. Moreover, the match process is more contemporaneous. There is also a further advantage to the HRS. It provides earnings records from the Social Security Administration, permitting more precise measurement of the earnings history of each spouse. Thus the HRS data allow more precise identification of parameters for each spouse within the structural model.

⁹A caveat should be noted. Because our analysis applies only to couples, it does not project the effects that private accounts would have on the behavior of single individuals.

cooperation, and selfish utility maximization with joint utility maximization.¹⁰ The two spouses share household consumption. They do not consume goods according to own income. For each spouse, utility is a function of own leisure, which in part may be determined by spouse's leisure, and household consumption. Although each spouse acts to maximize own utility, at a given level of own utility, each would choose any feasible alternative that improves their spouse's utility.

The utility functions for the two spouses are specified symmetrically. The subscript or superscript h signifies a variable that pertains to the husband; w signifies a variable pertaining to the wife.

For the husband we have:

$$U_h = \sum_{t=0}^T \left[\frac{1}{\alpha} C_t^\alpha + e^{X_t^h \beta_h + \gamma_h L_t^w + \epsilon_h L_t^h} \right]$$

For the wife, the utility function is:

$$U_w = \sum_{t=0}^T \left[\frac{1}{\alpha} C_t^\alpha + e^{X_t^w \beta_w + \gamma_w L_t^h + \epsilon_w L_t^w} \right]$$

C_t is family consumption, and L_t^h and L_t^w are the leisure of the husband and wife. L_t is a dichotomous variable taking on a value of 0 if the individual is working and 1 if retired at time t .¹¹ Each individual lives T years, and t is time since household formation. The terms

¹⁰This model is developed in Gustman and Steinmeier (2000a). A more complete description can be found there.

¹¹Primarily to keep the model simple enough to estimate, part-time work is ignored and retirement is considered to be an absorbing state; once retired, one cannot return to work.

$e^{X_t^h \beta_h + \gamma_h L_t^w + \epsilon_h}$ and $e^{X_t^w \beta_w + \gamma_w L_t^h + \epsilon_w}$ determine the relative values of retirement to the husband and wife. X_t is a vector of variables that includes a constant term, age, and health. ϵ_h is an individual fixed effect, where higher values of ϵ_h indicate higher values of retirement to the individual. As age increases, so does the value of leisure. When the value of retirement outweighs the value of the wages from working, the individual retires.

Each spouse's utility may be linked to the other's through three possible channels. Most directly, consumption is family consumption, financed by a joint budget constraint which is described below. In addition, the spouse's utility appears in the exponential expression affecting the value of one's own leisure. Lastly, the fixed effects in their respective utility functions may be correlated for husbands and wives. In an extension below, we also include a direct measure of how each spouse values the opportunity to share leisure with the other.

Both the husband and wife maximize their respective utility functions subject to the constraint that lifetime family consumption cannot exceed family income:

$$\sum_{t=0}^T d^t C_t \leq Y = \sum_{t=0}^T d^t (1 - L_t^w) W_t^w + \sum_{t=0}^T d^t (1 - L_t^h) W_t^h$$

In this budget constraint, both consumption and wages are expressed in real terms, and d is the real interest rate. W_t^h and W_t^w are the husband's and wife's compensation amounts when employed. In addition to wages, compensation includes annual accruals to the present values of pensions and social security, due both to own and spouse and survivor benefits.¹²

¹²The opportunity set does not include employer provided health insurance while working or retiree health insurance. In this type of model, which assumes perfect capital and insurance markets, on the job and retiree health insurance have little effect on retirement incentives. Insurance provided on the job has the same effect as an increase in the wage. However, the wage level has little effect on retirement. After age 65 retiree health insurance consists of

The sequence of decisions is straightforward. Because there is a common consumption parameter β , both spouses can agree on how to spend a given amount of lifetime family income. Each spouse then chooses own labor supply to maximize his or her own utility function. In choosing own labor supply, we assume that each spouse knows the leisure preferences of the other, and so bases their choice of own labor supply on the labor supply that the other spouse will choose as a result. With each spouse's labor supply entering the utility function of the other spouse, there is the possibility of two or more Nash equilibria. Should there be more than one Nash equilibrium, the one that is advantageous to both spouses will be chosen. When the spouses prefer different Nash equilibria, we assume that the spouse who retires first chooses the retirement date which is advantageous to that spouse, taking into account the retirement date that the second spouse will subsequently choose. There is no uncertainty in the model. Since both spouses know each others' preferences from the start, consumption and labor supply decisions are planned at the beginning of the life cycle with perfect foresight.

Details to the solution of the model are presented in our earlier paper. For family i , let $S_i(\beta, \beta_h, \beta_w, \beta_h, \beta_w)$ be the set of values of β_h and β_w in the utility maximization problem

medigap insurance. It has a low present value and thus has little effect on retirement behavior (Gustman and Steinmeier, 1994, 2000b). In contrast, Rust and Phalen (1997) find that health insurance is an important determinant of retirement behavior. However, they assume that the market for private health insurance is not working, and that those who are not seen to have purchased health insurance on the private market are unable to do so. There is no mechanism in their model for people to decide not to purchase health insurance in the private market because they are willing to self insure rather than pay the premium. This means that anyone who is still working in their study is assumed to be unable to purchase retiree health insurance in the private market, creating an artificial relationship between employment and availability of health insurance coverage. Other studies that do not rely on the assumption of a perfect insurance market nevertheless find only a small influence of retiree health insurance on retirement (Blau and Gilleskie, 2001; French and Jones, 2001).

which are consistent with retirement between the observed dates. If the retirement age for either spouse is not observed within the survey period (1992-2000), the set will not be bounded; this effectively is how the estimation procedure accommodates cases where a respondent has already retired before the survey starts or still has not retired when last observed.¹³ Note that the boundaries of the set depend on the values of the utility function parameters. Further suppose that the values of ϵ_h and ϵ_w come from a bivariate normal distribution with density $f(\epsilon_h, \epsilon_w | \sigma_h^2, \sigma_w^2, \rho)$, where F_h^2 and F_w^2 are the variances of ϵ_h and ϵ_w , and D is the correlation. Using this notation, the log-likelihood function is

$$\ln \mathcal{L} = \sum_{i=1}^{i=N} \ln \left[\int_{S_i(\alpha, \beta, \gamma_h, \gamma_w)} f(\epsilon_h, \epsilon_w | \sigma_h^2, \sigma_w^2, \rho) d\epsilon_h d\epsilon_w \right]$$

The integrals in the log-likelihood function are evaluated with a standard routine for cumulative joint probabilities of bivariate normal distributions. The likelihood function is maximized using a standard maximization routine, and standard errors for the estimates are calculated by the Berndt-Hall-Hall-Hausman method.

III. The Data

Our central focus is on results using the Health and Retirement Study (HRS).¹⁴ These results pertain to couples with a long term marriage, where each spouse also has a long term

¹³In a similar fashion, if either spouse begins to collect social security disability insurance benefits, the observation is treated as right censored at the time the disability benefits begin. This effectively means that for such couples, the last survey used in determining S_i is the survey before the disability payments start.

¹⁴The HRS is funded primarily by the National Institute on Aging, with additional support from the Social Security Administration and others.

commitment to the labor market. Table 1 describes the derivation of the sample. The HRS has 4767 couples for whom both spouses completed interviews in 1992. Of these, 1424 couples had at least one partner who had changed spouses, either through divorce or widowhood, after age 35, and thus do not qualify as long-term marriages.¹⁵ This means that approximately 30% of the couples are deleted because a lifetime planning model is probably not appropriate. More importantly, of those couples with a long term marriage, about 56% were dropped because one spouse was not a career worker. Specifically, 1876 couples out of 3343 did not meet the criteria for both spouses to be career workers.¹⁶

Next we face a trade-off between bias due to missing data and bias due to instrumentation. In a nonlinear model like ours, the choice is fairly clear: include only observations for which the required data are available. To illustrate, in our nonlinear model, certain types of incorrect information, such as the wrong date for the location of the spike in the

¹⁵If we were to include those who changed spouses after age 35, it would be necessary to determine how much wealth each spouse brought into the marriage, how they split obligations to children and facts that are not available in the HRS data.

¹⁶We use fairly liberal criterion in defining who is a career worker. Career workers are those who have worked full-time (30 hours or more) more than 50% of the time between age 40 (or 1982, whichever came earlier) and the last year of observed full-time work, as determined by the jobs in the job history and the full-time work answers in wave 3. The last year of full-time work must be no earlier than age 50, or if the worker was not 50 in 1992, he or she must have been working full-time in 1992. This was cross-checked with the social security records if those were available; a worker would not be considered to be a career worker if he or she had zero social security earnings in more than 50% of the years in the above-mentioned interval, unless the individual indicated that they worked on either government jobs or non-social security jobs. Also, an individual would always be considered to be a career worker if the social security earnings record indicated that he/she earned at least 60% of the real wages earned in the final full-time job for more than 50% of the years in the interval, even if the job history did not indicate enough years. This should catch instances of a series of short jobs which would be missed in the job history.

pension accrual profile due to early retirement provisions, will create a severe bias. Suppose that the detailed description of a defined benefit pension plan is missing, and we impute a value for the date of eligibility for early retirement benefits that is after the actual date of eligibility. Further suppose the person retired at the time he became eligible under his actual plan. In this circumstance, the model will find the respondent leaving just before becoming eligible for the (imputed) early retirement eligibility date. Consequently, the estimation will indicate that the person is not at all sensitive to economic incentives, since the foregone benefit accrual might have amounted to a year's pay or more from working for an additional few weeks or months. However, the error is not symmetric: if the imputed early retirement date is too low and the respondent retires at a later date, we would not necessarily conclude that the respondent is highly sensitive to economic incentives. By confining the estimation to those observations where a full set of information is available, we avoid this very strong bias that may result from imputation. That is why we have decided to omit observations for which we do not have an exact description of the pension, and to extend this choice to estimate findings only for the portion of the sample for which a complete data set is available.

Of the 1467 couples who are long-term families of career workers, data problems with the respondent reduce the number of couples by about a half ($143 + 24 + 116 + 7 + 462 = 752$; $752/1467 = 0.51$). Of this decline in the sample, about 61 percent of the loss is due to missing pension data ($462/752$).¹⁷ From the perspective of the entire sample, about two thirds of these

¹⁷The pension plan descriptions are missing disproportionately for employees in small firms, college grads, those with more than \$100,000 in assets, long tenure workers, those in manufacturing and management, those earning more than \$100,000 per year, those with defined contribution plans only, those with DB plans paying low benefits, and those with \$25,000 to \$100,000 in DC plans. Regressions are reported in Gustman and Steinmeier (2002).

older workers have pensions, and the provider profile is missing for about a third of them, meaning that over a fifth of these workers are dropped because of a missing pension. Again assuming relatively little duplication of missing pensions within a family, this means that almost 40% of two-worker families would be dropped for this reason. As seen in Table 1, in the end the sample used in the estimation amounts to 715 couples and represents not quite half of the original couples with career jobs and a long term marriage.¹⁸

We are going to compare the findings from the HRS with results from two studies that we undertook with data from the NLS Mature Women's Survey (NLS-MW). The first, Gustman and Steinmeier (2000a), used data through the 1989 wave. The women in the NLS-MW were born between 1923 and 1937 and thus were 52 to 66 years old in 1989. Pension characteristics used in that analysis were self reported. Plan descriptions were not available at that time from the respondents' employers. We know from other work (e.g., Gustman and Steinmeier, 1989, 2000a) that there are substantial errors in pension self reports. In addition, the respondent can provide only a brief list of determinants of pension benefits. Without a detailed description of the pension from the firm, we had to apply population averages for some pension plan features.¹⁹

¹⁸There are two other minor yet not completely non-trivial deletions. The first is instances where the number of full-time years is ambiguous. These are primarily cases where the social security record is missing and either the respondent was not interviewed at wave 3 or the wave 3 information about full-time years is missing. The second reason is instances where the age in one survey was greater than or equal to the age reported in a subsequent survey. This calls into question which age is correct and throws into doubt whether we have the correct age for the timing of retirement.

¹⁹Thus in this first paper using NLS-MW data, the pension is assumed to be a simple DB plan, where the benefit is the product of a generosity coefficient times years of service times final salary. The generosity coefficient (the term in the benefit formula determining the replacement rate) is calculated from the respondent's report of expected pension benefits, or if expected pension benefits are not reported, a figure of 1.6 percent is used, which is the median

After all exclusions, there were 564 couples in the NLS-MW sample using self reported pension incentives.

Bridging the gap between our earlier study and the present one, we have an updated version of our earlier study.²⁰ The later study uses the NLS-MW data through 1992 rather than 1989. Importantly, the later study does include information from employer provided pension plan descriptions. The pension summary plan descriptions were collected from employers after the 1989 survey, but there was a delay that adversely affected the rate at which employer plan descriptions could be successfully matched with the employer names provided by respondents. Moreover, a number of the husbands had retired by 1989. Although it was possible to match some employer provided pension plan descriptions, they were available for only one fifth of the husbands with a pension. In contrast, pension plan descriptions from current or last jobs are available for roughly two thirds of HRS respondents who report pensions on those jobs.

There also are other problems with the NLS-MW data, creating problems for both of our studies based on the NLS-MW data. Because the NLS-MW was a study of women, there was very limited information collected about the employment and earnings of the husband. Moreover, a full work history is available for about 70 percent of the respondents to the HRS in the form of the social security earnings history. In contrast, although the NLS-MW data provides an extensive work history for wives, the work history for husbands is badly incomplete,

for those plans for which we did have information. We assumed that all pensions reduce benefits from the age of normal retirement by 4.9 percent per year, a figure found in earlier work by Hatch et al. (1981).

²⁰The later study was conducted for the U.S. Department of Labor, Bureau of Labor Statistics (Gustman and Steinmeier, 1998). Results from this study have not been published elsewhere.

and in the end relies on retrospective data rather than panel data recorded contemporaneously, or administrative data. Indeed, in the NLS-MW, key information about the husband's pension was collected from the wife, including the self reported pension information used in our first study, and the employer names which we rely on when matching an employer provided pension plan description with an interview record for the husband.

Because these problems undermined our ability to match employer provided plan descriptions in the NLS-MW, especially for husbands, we used a different procedure in our second NLS-MW study than we followed when using HRS data. Even though pension data were available for the 1992 NLS-MW sample, in cases of missing pension descriptions, we used the early and normal retirement dates reported on the respondent survey (by the wife for her own and for her husband's pension). From this we constructed the pension accrual profiles using the generosity and early retirement reduction factors calculated as the means of employer provided pensions in the same industry, occupation, and earnings category. The idea was to use as many observations as possible in the smaller NLS-MW sample by anchoring the pensions on the self-reported early and normal retirement ages and imputing the generosity and reduction factors as the averages for pensions in similar jobs. However, the fact that we employed imputation procedures for the pensions in the NLS-MW means that estimates of the responsiveness to economic incentives are probably understated in these results.

The NLS-MW contains 2,084 women who were married at the beginning of the survey and who participated in each of the surveys through 1992.²¹ Of these, in our second study using

²¹Since the initial age of respondents was 30 to 44 in 1967, women who dropped out in the early years of the survey did so before reaching retirement age, and hence these women would not shed much light on a retirement analysis in any case.

NLS-MW data there are 499 couples in the sample.²²

IV. Descriptive Analysis

Tables 2 and 3 provide some idea as to the timing of retirement of husbands and wives within the HRS sample. By comparing the two parts of each table, we can determine the similarity between the sample used in our later analysis, and the full sample including observations with missing data. Part A of each table is based two-career couples in long-term marriages, while part B of each table excludes couples for whom we are missing information critical to calculating the budget constraint. It is the sample in part B of each table that is estimated and analyzed in later in the paper.

Among those in Table 2A who meet the definition of couples with a lifetime commitment to the labor market, 514 wives and 476 husbands retire after the last wave of the survey. In addition, 235 wives and 284 husbands retired before the first wave of the survey. With 2,934 total wives and husbands in the sample in Table 2A (1,467 couples), that leaves 1,425 individuals, or 48.6% of the original sample of career workers married only to their current spouses, who retired in waves 2, 3, 4, or 5 of the survey. Summing the observations along the diagonal, for 607 out of 1467 couples, or 41.4 percent of the observations, the husband and wife both retired in the same period.

²²Career workers refer to those with substantial full-time work experience (at least three consecutive surveys of work after age 40 and at least one-half of the surveys before the last survey with full-time work for women, or at least two-thirds of the surveys before the last survey with full-time work for men), and at least one survey of full-time work after age 50. Full-time work means at least 25 hours of work per week for women or at least 1250 hours per year for men, for whom usual weekly hours are not always available. Using a 35 hour per week or 1500 hour per year definition results in slightly higher joint retirement, but at a cost of about 20 percent of the sample.

In Table 2B there are 715 couples. Thus approximately half of the couples in Table 2A will be lost for not having economic information available. Among the sample with all the information required for estimation of our structural model available, 662 out of 1,430, or 46.3 percent, retired in waves 2, 3, 4 or 5 of the survey. Again summing along the diagonal, for the sample with full economic information available, 290 out of 715 couples, or 40.6 percent of couples, had the husband and wife both retire in the same period, very close to the 41.4 percent of the sample found for Table 2A.

Table 3 examines the patterns of retirement among HRS couples, according to their age differences. Among the 435 couples in Table 3A, which includes observations whether budget constraint variables are available or not, the wife is older than the husband in only 49 of them, or in 11.3% of the cases. Similarly, among the 192 couples in Table 3B, which excludes observations with budget constraint variables missing, the wife is older than the husband in 24 of them, or in 12.5% of the cases. In an additional 36 households in Table 3A, the wife and husband are the same age. So in more than three quarters of the households in both samples, the husband is older than the wife. Nevertheless, in Table 3A the median difference in time of retirement is zero, with 205 couples (47 percent of couples) with spouses who retire within the same year. Similarly, the median difference in Table 3B is zero, with 86 couples (45 percent of couples) with spouses who retire in the same year. Moreover, the distributions of differences in retirement age are symmetric around zero in the two tables. In both samples, about three times as many couples retire in the same survey as found with the husband retiring one survey later than the wife, and three times as many couples retire in the same survey as are found with the husband retiring one survey earlier than the wife. This evidence suggests the two samples, those

with complete data and the full sample which also includes observations with missing data, are similar. These findings are also strong evidence of coordination of retirement among the two career couples in the HRS who have already retired.

The data in Table 3 do not describe the patterns of retirement that will ultimately be observed, however, since the couples in Table 3 are selected to include those who retired by the fifth wave of the survey, and thus who have a stronger preference for retirement. As indicated in Section II, the estimation procedure does not censor the sample if either spouse has yet to retire, and thus the analysis below will focus on explaining the distribution of retirements that will ultimately be observed for this cohort.

V. Estimates of the Structural Model

Columns 1 and 2 of Table 4 report the maximum likelihood estimates for the parameters of the joint utility function and the associated t statistics using data from the Health and Retirement Study. Following the methodology reported above, the equations for the status of each spouse are jointly determined, allowing for the underlying interaction of the decisions of each spouse in a noncooperative bargaining model. The estimation searches for the coefficients of each of the parameters appearing in the utility function(s) and the range of fixed effects that are most likely to be associated with the retirement outcomes observed for the couple, conditional on the constraints formed by the wage offer, any pension and social security. The dependent variable in the equation for each spouse is an indicator of the work-retirement decision in each wave of the survey for which the respondent was observed.

We estimate a parsimonious specification of the utility function, with only a few right hand side variables included in evaluating the utility for each spouse. First there is " , the

exponent on the measure of joint consumption. The remaining measures affect the utility of retirement and are different for each of the spouses. For each spouse, the measure of age is continuous, so that no special effects are built into the outcomes through a dummy variable corresponding to whatever age the retirement hazard happens to spike at. Spouse's retirement status is a qualitative binary variable defined as whether the spouse is contemporaneously retired. Health status is an indicator equal to one if the respondent has reported in two successive surveys that health status is fair or poor, or if self reported health status is fair or poor for the last observed survey. An indicator of vintage (year of birth) is also included.

The estimated coefficients are similar to those we found in our earlier study (Gustman and Steinmeier, 2000a), which are reported in column 3 of Table 4, with the associated t statistics in column 4. Column 5 reports the coefficients obtained from the expanded sample from the NLS-MW which also included employer provided pension plan descriptions, or matched descriptions that involved defined benefit plans with comparable early and normal retirement dates.

The easiest way to interpret the findings is to begin with the coefficient on the age measure. This parameter indicates that roughly speaking, utility of retirement is increasing for the husband by about 60 percent per year with each year of age ($e^{.47} - 1$), and by about 52 percent per year for the wife ($e^{.42} - 1$). The coefficients on the age variable are lower in the HRS than in the NLS-MW.²³ That suggests that policies will be found to be more effective when they are evaluated using utility function parameters from the HRS. As suggested earlier, the smaller

²³The age coefficients in column 3 translate into percentage effects of 85 percent for each year of age for men, and 70 percent for women; while the coefficients in column 5 translate into percentage effects of 96 and 108 percent for each year of age for men and women respectively.

effects of age in the HRS may be due to more precise estimation of the pension accrual profile in the HRS, where pension plan descriptions were exactly matched, in contrast to reliance on crude pension formulas and self reported plan descriptions as in column 3, or a mix between imputed and matched plan descriptions, as in column 5.

For the husband, in the results using HRS data, having a retired wife is equivalent to the effect of being about a year older. This is similar to our published findings based on the NLS-MW seen in column 3, and a bit smaller than the NLS-MW results based on employer plan descriptions, which suggest that having a retired wife is equivalent to about another 1.8 years of age. For the wife, having a retired husband is equivalent to about another three quarters of a year of age, whereas there was almost no effect of having a retired husband in the NLS-MW sample. This finding, that in HRS data there is a stronger dependence of the wife's labor supply on the husband's retirement than in NLS-MW data, is in part traceable to better measurement of the opportunity set facing the husband in HRS data. Because in the HRS sample, the effect of each year of age is greater for the husband than the wife, having a retired spouse continues to have a larger effect for men than for women, but the difference between husbands and wives is narrower than we found in our earlier work using NLS-MW data.

In the HRS findings, for the husband the effect of ill health is equivalent to about an additional 1.5 years of age. This is considerably less than we found using the NLS-MW self reported data seen in columns 3 and 4, where poor health is equivalent to about three years of aging. For the wife, ill health has the same effect as about another 2.7 years of age, which is greater than the NLS-MW results, where ill health is equivalent to about another 1.6 years of age.

Vintage is also significant, as is the standard deviation of the fixed effects. The former result suggests that those in widely different vintages will have considerable differences in taste. However, we should note that both the NLS and HRS are focused on a fairly narrow range of vintages, and extrapolating very far outside this range may be unsound. As for unobserved differences in retirement preferences (the fixed effects), it is clear from the magnitude of the standard deviation of these preferences that variations in taste create a considerable difference in retirement behavior.

Lastly, the correlation of the fixed effect retirement preferences using the HRS data is almost identical to the value found in our earlier published data. This correlation is considerably weaker in the NLS results with employer provided pension data.

VI. Sources of Joint Retirement

In this model, it is difficult to compare directly the coefficients for the spouse retirement variables with the correlation coefficient for the unobserved part of preferences. Both the correlation coefficient and the coefficient of the wife retired variable in the husband's preferences are significant, and the coefficient of the husband retired variable in the wife's preferences is close to significant. By themselves, the sizes or even the significance of these measures do not establish which is more important as a determinant of joint retirement. To determine the relative importance of each effect, we conduct simulations of retirement behavior which include and exclude these effects.

To do the simulations, the procedure is as follows. The simulations are performed for the same couples who were used in the estimation, using the same values for the compensation streams and for the variables in the X vector as were used in the estimation. A random draw is

made from the bivariate normal distribution of ϵ_w and ϵ_h , allowing for the standard errors of the two ϵ 's and their correlation. This gives the retirement ages of the wife and husband corresponding to these values of the ϵ 's. This process is repeated 10,000 times for each couple in the sample.

Table 5 reports on the main results of these simulations for the HRS sample and for each of the NLS-MW samples. The fractions of households retiring together in each survey are reported in row 1 of the table. A great deal of caution is required in interpreting the results in row 1. Specifically, statistics on the baseline level of joint retirements should not be compared across surveys nor should they necessarily be compared to the simulation results. According to row 1 in Table 5, the proportion retiring together is much higher in the HRS than in the NLS-MW. But this may be due to two factors which make the numbers in this row to some degree non-comparable. First, these figures consider couples to retire together if they retire between the same two waves. But the waves are separated by different amounts of time in the two surveys. HRS waves are always two years apart. However, the NLS-MW waves are in some cases only one year apart. The longer period between survey waves will make the HRS figures on coincidence of retirement higher. Secondly, the percentages in row 1 of Table 5 use in the denominator only those couples for whom both retirements were observed. In the HRS, these cover only four periods between the five waves, while in the NLS-MW they cover a considerably larger number of waves. A perhaps more useful comparison across the surveys is that in the HRS, about three times as many couples retire together as at adjacent cells (see Table 3 above), and this is about the same proportion as in Figure 3 in our previous work using NLS-MW data (Gustman and Steinmeier, 2000a).

The simulations in the second row of Table 5 are for the full model, and in subsequent rows for the model with one or another source of interdependence in preferences suppressed. To clarify the measure of simultaneous retirement reported in the sample, if the simulation resulted in the husband retiring in 1994 at age 62 and the wife retiring in 1997 at age 58, the value of this variable would be -3. A value of 0 indicates that both spouses retired in the same year.

In contrast to the results in Row 1 of Table 5, which present the fraction retiring at the same time in the raw data, consisting only of those who had retired by the time the survey was taken, the simulations in the other rows of Table 5 report the retirement dates for all couples in the sample. Thus the results in rows 2 through 5 adjust for selectivity to incorporate the retirement dates for those who were not observed to retire by the last year of the survey.

Row 2 of Table 5 gives the results using the full model, including the spouse retirement variables and the correlation between unobserved preferences. In the HRS data, 9% of couples are simulated to retire in exactly the same year. Figure 1 shows the simulated distribution of relative retirement ages. The spike in the middle of the figure indicates the joint retirement. The part of the figure to the right refers to cases where the husband retires first, and the part to the left refers to cases where the wife retires first. The figure indicates that the incidence of joint retirement appears to be almost twice as great as the incidence of retirement one or two years apart.

Row 5 of Table 5 shows the results of simulations setting to zero the correlation in unobserved preferences and omitting the spouse retirement variables from the utility functions of

the two spouses.²⁴ These results lower the spike at joint retirement to the same level as the adjacent values in Figure 1 and thus exhibit no evidence of joint retirement. Note that this simulation eliminates any preferences for joint retirement, but does not eliminate incentives for joint retirement that operate through the opportunity set. For instance, if couples tended to choose jobs that had the same early retirement date in their pensions, the pensions might still induce a tendency toward joint retirement even if the couples otherwise had no particular preferences towards retiring at about the same time. This simulation, however, effectively rules out the possibility that a significant proportion of joint retirement arises because of coordinated retirement incentives in the compensation profiles.

The other two simulations reported in Table 5 examine separately the omission of the spouse retirement variables and setting the correlation of the unobserved preferences to zero. Row 3 omits the spouse retirement variables but keeps the correlation at the value found in the last row of Table 4. The correlation parameter has almost no effect on joint retirement. In contrast, when in row 4 we include the spouse retirement variables but omit the correlation, the spouse retirement variables alone account for almost all of the spike in joint retirement that is evident in the full model.

VII. Including A Direct Measure of Spouse Preferences for Joint Retirement

To further explore the role of preferences for joint retirement, we include a direct measure of the desire of each spouse to retire with the other. The Health and Retirement Study

²⁴In this simulation, the constant in the linear form $X\beta$ is increased to compensate for the omission of the spouse retirement variable. Otherwise, the omission of the spouse retirement variable would reduce the coefficient of leisure in the utility function and lead to an increase in retirement ages generally.

asked each respondent how much being with the other spouse is a positive point of retirement (questions K11d and K21d). This variable is defined to have a value of 1 if the respondent said that being with the spouse was a “very important” benefit of retirement. About half the respondents gave this response to the question. In the expanded model, the new variable is entered in the linear form $\$X$ as $\dots + \beta_e (\text{spouse retired}) (\text{enjoy time with spouse}) + \dots$. This has the effect of splitting the old coefficient of spouse retired into a part dependent on the new enjoy spouse variable and a remaining effect.

The model estimates with the new variable are presented in Table 6. This new variable picks up almost all of the effect of the original spouse retirement variable for wives, and around half for husbands. The effect is stronger for husbands than wives, and we still find that the husbands disproportionately prefer to have their wives retire with or before them. The wife’s parameter is significant at the 92% confidence level, and both variables are jointly significant at the 98% confidence level. Figure 2 indicates the relative retirement distributions implied by these results. Compared to Figure 1, this indicator of preference for joint retirement leads to a substantial increase in the share of joint retirements, from about 9 percent to almost 16 percent. Moreover, joint retirement now is about three times more common than retirement at adjacent values, which more closely approximates the retirement observed in Table 3B.

Table 7 reports the effects of the decomposition as to the reason for joint retirement when the “enjoy spouse in retirement” measure is included with the preference variables. Once again, virtually all of the explanation for joint retirements resides with the spouse retirement coefficient rather than with the correlation in preferences.

VIII. Simulating the Effects of Alternative Rules for Sharing Benefits Within the

Household

Policy makers are concerned with the rules governing the sharing of social security benefits between spouses. Under current provisions, when both spouses are alive each spouse is entitled to an amount equal to approximately half the benefits earned by the other, or to benefits based on own earnings, whichever is larger. When one dies, the other will receive either the survivor benefit (equal to the benefit the deceased was entitled to with some adjustment for early claiming), or the benefit based on their own earnings, whichever is larger. It can be shown that, because of the progressivity of the social security benefit formula, a lower earning spouse will have all benefits received while their spouse is alive based on own earnings if, very roughly, the AIME from own earnings is one third or more of the AIME of the high earning spouse. That is, one third of the AIME results in half of the higher earning spouse's PIA.

The incentive to continue to work depends in part on the marginal reward to continued work. Part of the marginal reward consists of any increase in social security benefits associated with an additional year of work. This, in turn, depends in a fairly complex way on whether the spouse is the higher or lower earning spouse, whether the spouses are currently eligible for benefits, and relative difference between the earnings amounts. At one extreme, an individual over 65 who is collecting spouse benefits would, by working, be giving up current benefits with no increase in future benefits at all. At the other extreme, an individual whose spouse has very low earnings can increase not only his or her own benefits, but also the spouse and potential survivor benefits of the spouse, by working an additional year. There are many cases in between these extremes. For instance, if both spouses are collecting benefits based on own earnings, an additional year of work by the lower income spouse will increase the future benefits of that

spouse, but only as long as the other spouse is alive.

The model we have estimated is structural and as a result allows us to isolate the effects both of current law and of alternative policies governing the crediting of benefits within the household. Some schemes for sharing benefits among spouses, including schemes that would simply divide credit for total earnings in a household evenly between the two spouses, will change the incentives for continued work for each spouse. A policy that would split the credit for earnings by either spouse evenly between both of them would increase the reward for work at older ages. For the higher income spouse, the reason is that after calculating the average indexed monthly earnings (AIME), a person's benefit in 2001 is 90 percent of the first \$6,732 of annualized AIME, 32 percent of the next \$33,840, and 15% of the remainder of AIME up to maximum covered earnings. When benefits are jointly credited, this spouse is more likely to be in the 90 percent bracket rather than the 32 percent bracket or the 32 percent bracket rather than the 15 percent bracket, and this will make benefits respond more strongly with increased earnings from further work. The lower earning spouse may also see an increased incentive to work if he or she would collect spouse benefits under the current system. Under the current system, increased work by such an individual generates no increase in future benefits at all, whereas with a scheme to split the earnings credit, increased work by the lower income spouse would generate increased benefits for both spouses.

As with earnings splitting, private accounts accrue benefits more evenly over the lifetime than under the current 90, 32, 15 percent brackets, again raising the reward to work later in life relative to the current system. The flatter accruals mean that the rewards for working later are relatively higher than the rewards to working in the early years, and this should delay retirement.

In addition to the work incentives, either earnings splitting or private accounts will redistribute benefits away from families where one spouse is the primary earner. Under the current system, a family where one spouse is the primary earner will collect more benefits, including spouse and survivor benefits, than will a family with the same total income but where the two spouses earn more nearly equal amounts. In other words, the current system redistributes benefits toward families where one spouse is the primary earner, and this redistribution would be nullified under either of the two alternatives. However, note that this does not mean that a family with two workers is as well off as a family with a single worker earning the same total amount, since the financial calculations do not value the leisure of the stay at home spouse.

Table 8 presents cumulative retirement probabilities by age from retirement simulations under three different programs. The first two columns present results under the current program. In the next two columns, the results are simulated for a program where the accruals are simply equal to the contributions. This corresponds roughly to a situation where the entire amount is placed in a private account and allowed to grow at the interest rate.²⁵ In the last two columns, the results pertain to a program where there is simple earnings splitting. That is, credited earnings are divided equally between the husband and wife each year.

Table 8B presents simulations based on the model in which each respondent indicates how much they value being with their spouse in retirement, whereas Table 8A runs the same

²⁵Note that whether or not the benefit is annuitized at retirement is irrelevant in this model, since the only thing that matters in the model is the expected present value of the accrual. This also means that any liquidity effects are not accounted for. Thus these findings are not the same as those that would be observed were liquidity constraints included in the model.

simulations where the indicator of spouse retirement status influences each spouse's valuation of retirement, but the variable indicating the valuation of spending time with one's spouse in retirement is missing.

As seen in Tables 8A, these alternative programs reduce the ranks of the retired by one to two percentage points. At age 55 men are about eight tenths of a percentage point less likely to have retired under a privatized system or one where credit for working is evenly split between spouses than under the current system. Wives are almost two percentage points less likely to have retired by age 55 under the alternative systems. By age 60 and 62, men are about 1.5 percentage points less likely to be retired under the alternative systems. Wives are two to three percentage points less likely to have retired under the alternative system. Smaller differences are found at age 65. With half the male labor force retired by age 62 and more than two thirds of the female labor force retired, these one to two percentage point differences in the share of the population retired translate into more than a two to four percent increase in the labor force around age 62. By age 65, given the lower base in number working, an almost two percentage point difference in the proportion retired translates into roughly a six percent increase in the size of the male labor force, and a 1.5 percentage point difference in the proportion retired translates into almost a 14 percent increase in the number of women working. Even when account is taken of the preference for having the spouse jointly retired, the implications of these two alternative programs are roughly the same as before, as is indicated in Table 8B.

To summarize, comparing outcomes between the two programs, the bigger increase in work effort is found under the private accounts plan rather than under the plan in which earnings are split between spouses.

IX. Conclusions

At the outset of this paper we emphasized the potential importance of having employer provided plan descriptions for identification of the factors shaping each spouse's retirement decision within the household. We find that key parameters are estimated with much greater precision when employer provided pension plan descriptions are matched for a large share of the pension covered workers. To be sure, we obtain the same qualitative message whether employer provided plan descriptions are available or not. Interdependence in retirement is due the appearance of spouse's retirement status in the preferences of both the husband and wife. But there is evidence of stronger interdependence in preferences with the improved HRS data, and the suggestion that social security policies changing the relative rewards to work by each spouse will have a larger effect on retirement outcomes. In addition, when labor supply histories are reported independently by each spouse, as they are in the HRS, we also obtain an improved understanding of retirement decision making within the household.

Using a measure of how much each spouse values being able to spend time in retirement with the other, we find that this direct measure of preferences accounts for much of the apparent interdependence in retirement within the household. When we include this measure, the simulations double the frequency of predicted joint retirements. Moreover, the wife's interdependence is due entirely to the difference between those who value spending time in retirement with their spouse and those who do not. Although it also remains true that husbands are more influenced by whether their spouse is retired than wives are, half the effect for the husband reflects whether he enjoys the idea of spending time in retirement with his wife.

Policy alternatives that would privatize social security, or divide benefits between

spouses encourage work at older ages. Compared to the current system, these policies will have a limited but not trivial effect on retirement outcomes. Due to the greater precision in estimating preference parameters with HRS data, the effects of age on retirement preferences are substantially lower in this paper than in our earlier work. This is particularly important because the responsiveness of retirement to the incentives created by pension and social security policies is greater the lower the coefficient on the age measure in the preference function. Thus we are better able to distinguish the retirement effects of changes in the allocation of benefits within the household under the social security changes. We find these effects to be noteworthy. At some ages, such as 65, there may be as much as a 6 percent increase in the old age work force under privatized accounts compared to the current social security program.

References

Anderson, Gustman and Steinmeier. 1999. "Trends in Male Labor Force Participation and Retirement: Some Evidence on the Role of Pensions and Social Security in the 1970s and 1980s." *Journal of Labor Economics* 17(4) part 1:757-783.

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

Blau, David and Donna Gilleskie. 2001. "Health Insurance and Retirement of Married Couples". Paper presented at Conference on Social Insurance and Pension Research, Aarhus Denmark.

Burtless, Gary and Robert A. Moffitt. "The Effect of Social Security Benefits On The Labor Supply Of The Aged," in Henry Aaron and Gary Burtless, Editors, *Retirement And Economic Behavior*, Washington, D.C., Brookings Institution, 1984, pp. 135 - 74.

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

Fields, Gary S. and Olivia S. Mitchell. *Retirement, Pensions and Social Security*, Cambridge: MIT Press, 1984.

French, Eric and John Jones. 2001. "The Effects of Health Insurance and Self Insurance on Retirement Behavior". Paper presented at Conference on Social Insurance and Pension Research, Aarhus Denmark.

Gustman, Alan L. and Thomas L. Steinmeier. 1986a. "A Structural Retirement Model," *Econometrica*, 54(3), May, pp. 555-584.

_____. 1986b. "A Disaggregated Structural Analysis of Retirement By Race, Difficulty of Work and Health," *Review of Economics And Statistics*, 67(3), August, 1986b, pp. 509-513.

_____. 1989. "An Analysis of Pension Benefit Formulas, Pension Wealth and Incentives From Pensions," in *Research In Labor Economics*, 10. Edited by Ronald Ehrenberg. Greenwich, Ct: JAI Press, pp. 33-106.

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

_____. 1993. "Pension Portability and Labor Mobility: Evidence from the Survey of Income and Program Participation". *Journal of Public Economics* 50 (March): 299-323.

_____. 1994. "Employer Provided Health Insurances and Retirement Behavior". *Industrial and Labor Relations Review* 47(3): 417-438.

_____. 1995. *Pension Incentives and Job Mobility*. Kalamazoo, Michigan: Upjohn Institute for Employment Research.

_____. 1998. "Retirement in Dual-Career Families: Estimates Using Firm-Reported Pensions", July. Report to the Bureau of Labor Statistics.

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

_____. 2000a. "Retirement In Dual Career Families". *Journal of Labor Economics* 18(3): 503-545.

_____. 2000b. "Pensions and Retiree Health Benefits In Household Wealth: Changes From 1969 to 1992", *Journal of Human Resources* 35(1): 30-50.

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

_____. 2001a. "Imperfect Knowledge, Retirement and Saving". National Bureau of Economic Research. Working Paper No. 8406.

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

_____. 2002. "What People Don't Know About Their Pensions and Social Security," in William G. Gale, John B. Shoven and Mark J. Warshawsky, editors, *Public Policies and Private Pensions*. Washington: D.C., Brookings Institution. Forthcoming.

Hatch, S.P. et al. 1981. "Financial Retirement Incentives in Private Pension Plans". Final report to the Department of Labor. Washington, D.C.: Urban Institute.

Ippolito, 1998. *Pension Plans and Employee Performance: Evidence, Analysis, and Policy*. Chicago: University of Chicago Press.

Johnson, Richard W. and Melissa M. Favreault. 2001. "Retiring Together or Retiring Alone: The Impact of Spousal Employment and Disability on Retirement Decisions". Center for Retirement Research at Boston College.

Lumsdaine, Robin, James Stock, and David Wise. "Efficient Windows And Labor Force

Reduction," *Journal of Public Economics* 43, 1990, pp. 131-159.

_____. "Three Models of Retirement: Computational Complexity versus Predictive Validity", in *Topics In The Economics of Aging*. Edited by David Wise. Chicago: University of Chicago Press, 1992, pp. 19-57.

_____. "Pension Plan Provisions And Retirement: Men & Women, Medicare, And Models," in *Studies in The Economics of Aging*. Edited by David Wise. Chicago: University of Chicago Press, 1994, pp. 183-220.

_____. "Why Are Retirement Rates So High At Age 65?". In David A. Wise, editor, *Advances In The Economics of Aging*, Chicago: University of Chicago Press, 1996, pp. 11- 82.

Lundberg, Shelly. 1999. In Henry Aaron. *Behavioral Dimensions of Retirement Economics*.

Mitchell, Olivia S. "Worker Knowledge of Pension Provisions". *Journal of Labor Economics* 6(1): 28-39.

Pozzebon, Silvana and Olivia S. Mitchell. "Married Women's Retirement Behavior," *Journal of Population Economics*, 2(1), 1989, pp. 301-53.

Rust, John and Christopher Phalen. 1997. "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets". *Econometrica* 65(4): 781-831.

Stock, James H. and David A. Wise. 1987. "Pensions, the Option Value of Work and Retirement". *Econometrica* 58(5):1151-1180.

Table 1
Reasons for deletions of observations from the HRS Sample

	Observations Deleted	Observations Remaining
Couples with both spouses interviewed		4767
Changed spouses after age 35	1424	3343
Not both career workers	1876	1467
Age not consistent among surveys	143	1324
Social security status ambiguous	24	1300
Number of full-time years ambiguous	116	1184
Earnings unclear from SS record alone	7	1177
No Pension Provider record in last job	462	715

Table 2
Retirement Tabulations From the HRS by Year

A. Including Observations With Missing Budget Constraint Data

	Retirement of Wife						Sum of Husbands
	Before 1992	1992-1994	1994-1996	1996-1998	1998-2000	After 2000	
Retirement of Husband							
Before 1992	107	43	37	25	27	45	284
1992-1994	26	59	15	19	15	29	163
1994-1996	24	21	46	21	21	40	173
1996-1998	22	17	27	53	24	36	179
1998-2000	26	8	17	25	47	69	192
After 2000	30	28	34	40	49	295	476
Sum of Wives	235	176	176	183	183	514	1467

B. Excluding Observations With Missing Budget Constrain Data

	Retirement of Wife						Sum of Husbands
	Before 1992	1992-1994	1994-1996	1996-1998	1998-2000	After 2000	
Retirement of Husband							
Before 1992	48	19	19	12	11	29	138
1992-1994	13	24	8	12	9	14	80
1994-1996	10	4	14	8	11	26	73
1996-1998	10	7	13	25	10	14	79
1998-2000	10	3	9	12	23	39	96
After 2000	12	16	18	24	23	156	249
Sum of Wives	103	73	81	93	87	278	715

Table 3
Retirement Differences Between Husbands and Wives In The HRS Among Couples
Who Have Already Retired, By Age Difference Between Husband and Wife

A. Including Observations With Missing Budget Constraint Data

Age Difference Husband - Wife	Difference in Retirement Surveys (Husband - Wife)							Sum
	-3	-2	-1	0	1	2	3	
-10	0	0	0	1	0	0	0	1
-9	0	0	0	0	0	0	0	0
-8	0	0	0	0	0	0	0	0
-7	0	0	0	1	0	0	0	1
-6	0	0	0	1	0	0	0	1
-5	0	0	1	1	0	0	0	2
-4	0	0	0	2	0	0	0	2
-3	0	0	0	4	1	0	0	5
-2	1	0	1	6	2	1	0	11
-1	1	2	3	12	5	3	0	26
0	1	2	3	17	9	4	0	36
1	1	1	8	26	8	3	1	48
2	3	6	1	24	12	5	2	53
3	0	7	11	21	8	6	1	54
4	2	6	11	25	8	2	1	55
5	1	8	3	13	4	3	1	33
6	2	0	7	10	5	0	1	25
7	0	1	5	14	4	5	0	29
8	1	0	0	11	2	1	0	15
9	1	4	2	5	3	0	0	15
10	1	3	4	11	2	1	1	23
Sum	15	40	60	205	73	34	8	435

Table 3
 Retirement Differences Between Husbands and Wives In The HRS Among Couples
 Who Have Already Retired, By Age Difference Between Husband and Wife

B. Excluding Observations With Missing Budget Constraint Data

Age Difference Husband - Wife	Difference in Retirement Surveys (Husband - Wife)							Sum
	-3	-2	-1	0	1	2	3	
-10	0	0	0	0	0	0	0	0
-9	0	0	0	0	0	0	0	0
-8	0	0	0	0	0	0	0	0
-7	0	0	0	1	0	0	0	1
-6	0	0	0	0	0	0	0	0
-5	0	0	1	0	0	0	0	1
-4	0	0	0	1	0	0	0	1
-3	0	0	0	2	0	0	0	2
-2	0	0	1	5	2	1	0	9
-1	0	1	0	7	1	1	0	10
0	1	1	0	7	5	2	0	16
1	1	1	6	11	3	1	0	23
2	2	3	1	8	6	3	1	24
3	0	4	7	8	3	4	1	27
4	1	5	6	15	3	1	0	31
5	0	3	2	4	1	2	0	12
6	1	0	1	1	0	0	1	4
7	0	0	1	5	1	1	0	8
8	1	0	0	5	1	0	0	7
9	1	2	0	1	2	0	0	6
10	1	3	0	5	1	0	0	10
Sum	9	23	26	86	29	16	3	192

Table 4
Parameter Estimates for a Structural Model

	HRS		NLS-MW Self Reported Pensions		NLS-MW Firm Reported Pensions	
	coefficient	t-statistic	coefficient	t-statistic	coefficient	t-statistic
Joint consumption exponent	-0.59	-2.7	-1.53	-4.0	-1.21	-3.4
Husband's parameters						
Constant	-10.18	-12.0	-20.03	-15.7	-18.34	13.9
Age ^a	0.47	4.8	0.61	4.1	0.68	4.4
Wife's Retirement	0.50	2.2	0.58	1.1	1.19	2.9
Health	0.72	2.3	2.05	3.7	1.88	3.7
Vintage ^b	0.11	2.7	0.11	2.3	0.12	2.6
Std. dev. of fixed effects	2.75	5.5	3.41	3.4	3.51	5.0
Wife's parameters						
Constant	-9.23	-17.0	-18.62	-26.8	-17.28	-16.6
Age ^a	0.42	5.3	0.53	5.2	0.73	4.3
Husband's Retirement	0.31	1.7	0.10	0.3	0.00	— ^c
Health	1.12	3.4	0.98	3.1	1.05	2.5
Vintage ^b	0.12	3.5	0.08	2.0	0.11	2.4
Std. dev. of fixed effects	2.35	5.8	2.71	5.8	3.56	4.7
Correlation of fixed effects	0.24	4.2	0.24	4.1	0.09	1.5
Number of Observations	715		564		449	
Log likelihood	-1776.50		-1394.47		-1545.24	

Age is measured at the time of each survey. An individual is retired if not working full-time with no further observations of full-time work. Health equals one if in two consecutive surveys (or in the last observed survey) self-reported health is fair or poor.

a. The actual age variable is the observed age minus 55. This is done to facilitate the maximization routine, and it has no implications for the estimates other than affecting the constant terms in the linear forms.

b. The actual vintage is the year of birth minus 1930 for the NLS, and the year of birth minus 1936 for the HRS.

c. In the wife's retirement equation in column 5, the coefficient for the variable indicating the husband is retired is constrained to be zero.

Table 5
Proportion of Households With Husband and Wife Retiring Together

	HRS Data with Employer Reported Pensions	NLS-MW with Self Reported Pensions	NLS-MW with Employer Reported Pensions
Raw Data ^a	0.45 ^c	0.25	0.27
Full Model	0.09	0.11	0.14
Without Spouse Retirement ^b	0.05	0.05	0.06
Without Rho	0.08	0.10	0.14
Without Both Spouse Retirement and Rho	0.05	0.04	0.06

a. Results for raw data include only those observations where both spouses retired by the last wave of the survey. Results estimated with model adjust for selection and include those retiring after the age range observed for the survey.

b. When the spouse retirement variable is set equal to zero, the constants are increased so as not cause an increase in the average retirement age.

c. Computed from Table 3B as 86/192.

Table 6
The Structural Model Estimated With HRS Data
Including A Measure of Enjoyment of One's Spouse

	coefficient	t-statistic
Joint consumption exponent	-0.58	-2.7
Husband's parameters		
Constant	-10.02	-11.8
Age	0.45	4.7
Wife's Retirement	0.53	2.0
Wife's Retirement*Enjoy Time With Wife	0.34	1.2
Health	0.71	2.3
Vintage	0.09	2.5
Std. dev. of fixed effects	2.69	5.4
Wife's parameters		
Constant	-9.10	-16.7
Age	0.42	5.3
Husband's Retirement	0.06	0.3
Husband's Retirement*Enjoy Time With Husband	0.52	2.2
Health	1.07	3.3
Vintage	0.12	3.4
Std. dev. of fixed effects	2.36	0.4
Correlation of fixed effects	0.19	3.3
Number of Observations	715	
Log likelihood	-1772.3922	

Table 7
Proportion of HRS Households With Husband and Wife Retiring Together

	HRS Without “Enjoy Spouse” Variable	HRS With “Enjoy Spouse Variable”
Raw Data ^a	0.47 ^c	0.47 ^c
Full Model	0.09	0.16
Without Spouse Retirement ^b	0.05	0.05
Without Rho	0.08	0.15
Without Both Spouse Retirement and Rho	0.05	0.05

a. Results for raw data include only those observations where both spouses retired by the last wave of the survey. Results estimated with model adjust for selection and include those retiring after the age range observed for the survey.

b. When the spouse retirement variable is set equal to zero, the constants are increased so as not cause an increase in the average retirement age.

c. Computed from Table 3 as 205/435.

Table 8
Effects of Alternative Social Security Schemes on Cumulative Retirements by Age

A. Model Without Enjoy Spouse Variable

age	Current System		Private Accounts		Divide Earnings	
	husband	wife	husband	wife	husband	wife
50	1.3	6.9	1.2	5.9	1.1	5.6
51	2.1	9.4	1.8	8.2	1.8	7.9
52	3.2	12.3	2.8	10.9	2.9	10.6
53	4.7	15.9	4.2	14.4	4.2	14.1
54	6.6	20.4	6.0	18.6	6.1	18.4
55	9.5	25.8	8.7	23.8	8.8	23.7
56	12.9	31.5	11.9	29.4	12.1	29.4
57	17.2	38.0	16.0	35.9	16.2	36.0
58	22.5	45.2	21.2	43.1	21.4	43.3
59	28.8	52.2	27.3	50.1	27.6	50.4
60	36.2	59.7	34.5	57.8	34.9	58.1
61	43.4	66.6	41.7	64.9	42.1	65.3
62	50.9	74.0	49.1	71.6	49.6	72.0
63	58.3	79.7	56.6	77.6	57.1	78.0
64	65.4	84.7	63.8	83.0	64.7	83.3
65	72.7	89.1	70.9	87.5	72.4	88.1
66	78.6	92.2	77.0	90.9	78.6	91.6
67	84.0	94.6	82.4	93.6	84.0	94.2
68	88.3	96.3	86.9	95.6	88.4	96.2
69	91.7	97.6	90.5	97.1	91.7	97.5
70	93.6	98.2	93.3	98.1	93.4	98.2

Table 8
Effects of Alternative Social Security Schemes on Cumulative Retirements by Age

B. Model With Enjoy Spouse Variable

age	Current System		Private Accounts		Divide Earnings	
	husband	wife	husband	wife	husband	wife
50	1.4	6.6	1.3	5.9	1.3	5.7
51	2.2	9.2	2.0	8.2	2.0	7.9
52	3.3	12.1	3.0	10.9	3.0	10.7
53	4.8	15.7	4.3	14.4	4.3	14.1
54	6.8	20.2	6.2	18.6	6.2	18.4
55	9.7	25.6	8.9	23.8	8.9	23.7
56	13.2	31.3	12.1	29.4	12.3	29.4
57	17.5	37.9	16.3	35.9	16.4	36.0
58	22.9	45.0	21.4	42.9	21.7	43.1
59	29.3	52.0	27.7	49.9	28.0	50.2
60	36.8	59.5	35.0	57.5	35.4	57.9
61	44.1	66.0	42.2	64.6	42.7	65.0
62	51.8	73.8	49.8	71.3	50.3	71.7
63	59.4	79.5	57.4	77.4	58.1	77.8
64	66.5	84.5	64.6	82.7	65.6	83.0
65	74.0	89.0	72.1	87.2	73.7	87.9
66	79.7	92.1	77.9	90.7	79.5	91.4
67	85.1	94.5	83.4	93.4	85.0	94.1
68	89.2	96.3	87.7	95.5	89.1	96.0
69	92.5	97.6	91.2	96.9	92.4	97.4
70	94.2	98.1	93.8	98.0	94.0	98.1

Figure 1
Distribution of Differences in Retirement Dates Between Husband and Wife
Using the Model without the “Enjoy Spouse Retirement” Variable




Figure 2
Distribution of Differences in Retirement Dates Between Husband and Wife
Using the Model with the “Enjoy Spouse Retirement” Variable

