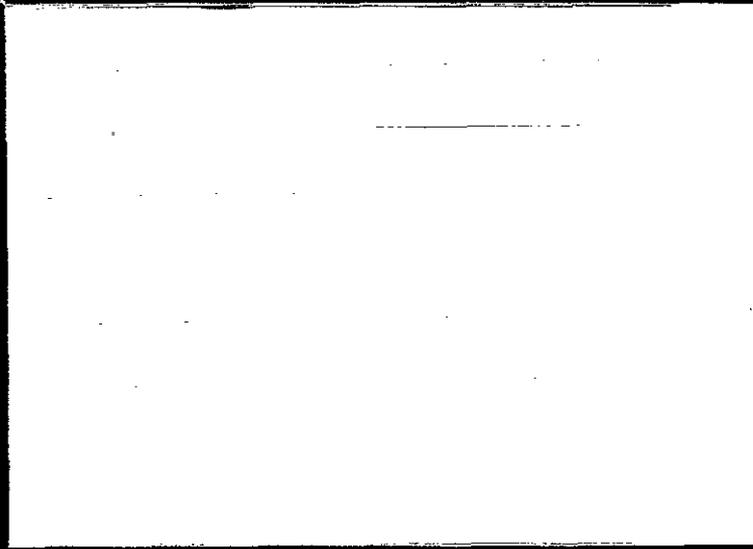


National Longitudinal Surveys

U.S. Department of Labor
Bureau of Labor Statistics

Discussion
Paper



The National Longitudinal Surveys (NLS) program supports many studies designed to increase understanding of the labor market and methods of data collection. The *Discussion Papers* series incorporates some of the research coming out of the program. Many of the papers in this series have been delivered as final reports commissioned as part of an extramural grant program, under which grants are awarded through a competitive process. The series is intended to circulate the findings to interested readers within and outside the Bureau of Labor Statistics.

Persons interested in obtaining a copy of any *Discussion Paper*, other NLS reports, or general information about the NLS program and surveys should contact the Bureau of Labor Statistics, Office of Economic Research, Washington, DC 20212-0001 (202-606-7405).

Material in this publication is in the public domain and, with appropriate credit, may be reproduced without permission.

Opinions and conclusions expressed in this document are those of the author(s) and do not necessarily represent an official position of the Bureau of Labor Statistics or the U.S. Department of Labor.

Retirement in a Family Context: A Structural Model for Husbands and Wives

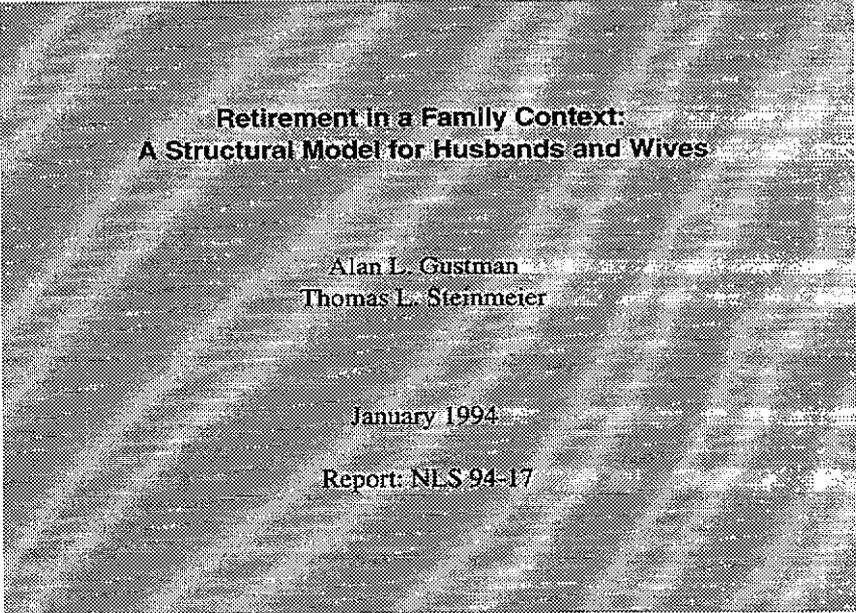
Alan L. Gustman
Dartmouth College and National Bureau of Economic Research

Thomas L. Steinmeier
Texas Tech University

January 1994

This paper was funded by the U.S. Department of Labor, Bureau of Labor Statistics, under Grant Number E-9-J-1-0068. Opinions stated in this paper do not necessarily represent the official position or policy of the U.S. Department of Labor. Research support was also provided by the National Institute on Aging. Any opinions expressed also do not represent the opinions of the NIA or the NBER. This paper is part of the NBER programs in Labor Studies and Aging. The authors wish to thank Andrew Samwick, Olivia Mitchell, and participants in seminars held at the Bureau of Labor Statistics and the NBER Summer Institute for providing helpful comments.

NLS Discussion Papers



**Retirement in a Family Context:
A Structural Model for Husbands and Wives**

Alan L. Gustman
Thomas L. Steinmetz

January 1994

Report: NLS 94-17

I. Introduction.

This paper specifies and estimates a structural model of the retirement decisions of husbands and wives. The feature of the data that is of central interest to us is the tendency of husbands and wives to retire together. An econometric approach is developed for estimating preferences of both spouses jointly and is implemented using data from the National Longitudinal Survey of Mature Women (NLS), a survey that provides the most recent data available for a joint retirement study. Alternative specifications of joint decision making are tested, and the importance of various sources of interdependence in decision making are investigated.

The present analysis joins together two branches of the retirement literature. In one, structural retirement models are estimated from data on individuals, usually men, while ignoring the retirement decisions and retirement status of their spouse.¹ In the other, retirement decisions of husbands and wives, and the tendency of their retirement dates to cluster, are analyzed in the context of a reduced form model.²

To join these strands of the literature, this paper extends structural retirement modeling to incorporate the joint determination of retirement decisions of husbands and wives.³ In

1. See, for example, Burtless and Moffitt (1984), Fields and Mitchell (1984), Gustman and Steinmeier (1986a and b), Stock and Wise (1990a and b), Berkovec and Stern (1991) and Lumsdaine, Stock and Wise (1990, 1992a and b).

2. Estimates of systems of reduced form retirement equations, such as those in Clark and Johnson (1980) and Hurd (1990), suggest the importance of the spouse's retirement status, and of the joint determination of the retirement decisions of husbands and wives.

3. One paper which comes closest to bridging these two strands of the literature is Pozzebon and Mitchell (1989). That study fits a version of the Fields and Mitchell structural model of retirement to data for married women using the observations for working wives that are available in the Retirement History Study. It assumes that the retirement decision of the

particular, the model is designed to recognize a number of potential sources of interdependence in the retirement decisions of both spouses. In the opportunity set, jobs may be selected with peaks in pension accrual profiles that encourage joint retirement. On the preference side, each spouse's utility may depend on the retirement status of the other spouse; the preferences of each spouse may be correlated; and each spouse may not make a retirement decision independently, but may collude to insure that the retirement decision of each is jointly optimal.

In a world where both the husband and wife are more and more likely to be working until the retirement years, this kind of a model will be increasingly necessary to assess how pension, social security and other retirement related policies affect retirement outcomes, including the question of whether policy measures which affect the retirement decision of one family member can indirectly influence the retirement of the remaining spouse.

A second purpose of this paper is to provide a structural retirement analysis using much more recent data for a nationally representative panel study than has been used in the past.⁴ The women in the National Longitudinal Survey of Mature Women (NLS) were born between 1923 and 1937. In the last year for which the survey is available, 1989, they were 52 to 66

husband is predetermined from the perspective of the wife's decision.

4. Most structural analyses which are based on nationally representative panel data sets, even recently completed studies, have made use of the Retirement History Study (RHS), a longitudinal survey with cohorts born between 1906 and 1911. Only very limited information was provided in the RHS on the labor market activities of wives. In addition, even if sufficient information had been provided in the RHS for a family labor market study, the cohorts in the RHS would be outdated because the participation patterns of wives have changed so drastically over the past two decades. Studies which use data from the past five or ten years have not employed data which is nationally representative. See, for example, Lumsdaine, Stock and Wise (1990, 1992a and b).

years old, about as young as feasible for a retirement study.

The paper also addresses a number of econometric and behavioral issues. One issue is selection bias. In an old panel such as the RHS, all couples have retired, so there is complete information on retirement dates and on the characteristics of all individuals included in the sample. In a more recent panel such as the National Longitudinal Survey of Mature Women, many of the couples are too young for both to have retired. Accordingly, in the process of the analysis, we deal explicitly with truncation of continuing employment spells. A related selection problem occurs in data sets which are restricted to retirees, such as the Survey of Newly Retired Beneficiaries analyzed by Hurd (1990).

In the course of specifying and estimating the retirement model for couples, we will address the following specific questions: How does the retirement behavior of each spouse compare to that of the other? To what extent is the wife's retirement decision influenced by the husband's, the husband's decision influenced by the wife's, and to what extent are their unmeasured tastes correlated? How does interdependence in the opportunity set and in preferences affect the coordination of retirement by the spouses? What are the effects on model parameters of entirely ignoring interactions in preferences? What are the effects of treating each spouse's retirement decision as exogenous in estimating the retirement behavior of the other?

The next section will present evidence from the NLS that in fact, in the raw data, there is a noticeable tendency among couples who have both retired to retire together. A family labor supply model is developed in Section III. Section IV details the data preparation and presents alternative estimates of the model. Section V presents some simulations based upon

the estimated model, concentrating mainly on the extent to which the husband's and wife's retirement decisions are coordinated. A final section presents some concluding thoughts.

II. Evidence of Joint Retirement.

There is not much point in investigating the cause of couples tending to retire together if the phenomenon is not evident empirically. The purpose of this section is to verify that the phenomenon does exist in the NLS older women's data, and to document the sample which is used in this section and in the remainder of the project.

Table 1
Sample Inclusion Criteria For The NLS Data Used In This Study

Selection Criteria	Number
Number of women in the 1967 NLS survey	5083
With observations in all NLS surveys	2715
With the same husband in all surveys	1520
Who worked full-time at age 50	654
With husband who worked full time at age 50	594
With at least one full-time wage observation for the wife	578
With at least one full-time wage observation for the husband	564

Note: The numbers for each line represent the number of respondents in the previous line with the additional characteristic

The sample we use from the National Longitudinal Survey of Mature Women (NLS) is restricted in several ways in order to meet the objectives of this paper. Table 1 indicates the effects on the number of observations applying various screens for inclusion in the sample.

First, the sample is restricted to women who remained in the survey.⁵ Second, since the project focuses on the joint retirement decisions of husbands and wives, the sample is restricted to women who were married to the same man in each year of the sample, and to those couples both of whose members survived.⁶ Third, the sample is restricted to couples for whom the idea of retirement is a meaningful concept. Specifically, it excludes any couple if either the husband or wife quit full-time work for good prior to age 50.⁷ Additionally, it excludes any couple if the wife did not have at least three consecutive surveys of full-time work after age 40, or if the husband worked full-time for less than two-thirds of the sample years before he left full-time work for good.⁸ Thus the sample does not include women who tried working for a year or two and then dropped out of the labor force. For these women, it would be stretching things to construe the fact that they dropped out of the labor force as

5. Most of the women who do not meet this criterion were women who attrited permanently from the survey. Typically, they were interviewed in the early years and then dropped out. Cases where one or two interviews are missing in the middle of the survey are relatively rare. Since the initial age is 30 to 44 in 1967, most of the women who dropped out in the early years did so before reaching retirement age, and hence these women would not shed much light on a retirement analysis in any case.

6. The survey does not directly ask the question "Are you married to the same man as you were during the previous survey", but we do make sure that the woman is never observed to be separated or divorced, and that the answers to questions in 1977 indicate that the marriage to her husband was before 1967.

7. For wives, full-time means at least 25 hours per week at the time of the survey. For husbands, for whom usual weekly hours is not always available, it means at least 1250 hours in the past year, with one exception. If, on the last instance in which the annual hours exceeded 1250, the weeks worked were less than 48, and the weeks worked on the following survey were zero, we conclude that the husband retired before the survey date.

8. The women in the NLS have been surveyed periodically, sometimes at yearly intervals, but mostly at two-year intervals.

retirement. Lastly, to be included in the sample, at least one full-time wage observation must have been collected both for the husband and for the wife.

In this study retirement is defined as no longer working full-time.⁹ Some wives work full-time up to a given survey date, and after that date she does not work full-time (she may report she is retired, or that she is a housewife).¹⁰ A smaller number of wives have a period where they work full-time continuously for a period of time, and then have a period of time where they work full-time some years but not others. In this section, the retirement date is taken to be the year following the last year of full-time work, unless the last year is 1989 (the last year of the survey). If they are still working full-time in 1989, they are treated as retiring sometime after the survey ends, with the date being left unspecified. Retirement of the husbands is defined in a similar manner.

Table 2 shows the distribution of retirement ages of wives and husbands. Those with a recorded retirement age are a minority of the sample, since most of the couples had one or both partners not retiring by 1989.¹¹ The diagonal dotted path indicates retirement at the

9. The analysis does not distinguish between full-retirement and partial retirement, counting the partially retired and the retired in a single category. For discussions of partial retirement, see Gustman and Steinmeier (1983, 1984, 1985a, 1986a).

10. The number of retired wives is greater than the number of wives who explicitly said that they were retired in 1989, since many wives report themselves as housewives rather than as retired once they stop working.

11. The number of couples included in Table 2 in the category where both spouses have retired includes a larger number of early joint retirees than will be observed for the whole sample. Bias from this source will cause the gap in retirement ages to be understated in that table. From the last column and bottom row of that table it can be determined that those in the sample who have yet to both retire, 228 couples have both husbands and wives still working full-time, 61 have only the husband working full-time, and 115 have only the wife working full-time.

Table 2
Retirement Ages

		Retirement Age of Husband															nr	
		51	52	53	54	55	56	57	58	59	60	61	62	63	64	65		>65
Retirement Age of Wife	51	•		2	1		2	1		1	1	3						11
	52	1	•	1		1	2		1		1		1			1		5
	53			•		1	1		1			1				1		5
	54				•		1		1	2		4	3		2	1	1	7
	55		1		2	•	1	2		1	2		1	1				4
	56					1	1	3				2	2		2		1	9
	57						5	•	2		2	2	4	1		1		3
	58		1		1				1	2	2	1		2	3		1	5
	59				1	2	1	1	1	2	2	2	3	3	2	1	4	5
	60					1	1				2	2	1	2	1		2	2
	61									1	1	3	2	3	1			5
	62				1			1		1	1	1	1	1	1	1		3
	63						1		2			2	1	•	1			2
	64									1					•	1		2
	65						1			1					1	•		
nr	4	3	5	6	7	6	9	8	6	11	10	12	13	6	5	4	228	

Note: nr = not retired

same age for both partners. The fact that the midpoint of the data is above the dotted line suggests that, on average, husbands retire at ages that are two or three years later than the ages their wives retire, on average.

Table 3 rearranges the data on the basis of retirement dates, and addresses the central

Table 3
Distributions of Retirement Dates for Couples
With Both Husband and Wife Retired

		Difference in Retirement Dates (Husband - Wife)															
		-7	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6	7	8
Age Difference (Husband - Wife)	-5								1								
	-4																
	-3																
	-2									1				1			
	-1							1	1			1					
	0			3	1		1	3	5	1		2		3			
	1		2	1			4		3	2	2	1		1		1	
	2			2	1	1	2		9		4	4		2	1	1	2
	3			1	1	2	1		4	1	3	2		2			
	4			2	1	7	2		2	2	2	2	2				
	5								4	2	1			1			
6					3		1	2		1							
7					1	2		1									
8	1							3			1						
9			1			1											
10	1					1		2									
Column Total		2	2	10	4	14	14	5	38	7	13	13	2	10	1	2	2

Note: Retirement dates are calculated as the year immediately following the last observation of full-time work

concern of this paper, the coordination of retirement dates by husbands and wives.¹² Down

12. The reader should remember that these data refer only to those couples who have both retired at the time of the 1989 survey. This means that for wives of a given age, a larger proportion of couples with a husband who is much older than his wife will be retired, and thus included in the table. In addition, those individuals who have a stronger preference for leisure are disproportionately more likely to be reported among those couples who have both retired. The later estimation will take account of this selection problem.

the left side are the age differences between the husbands and wives. A value of 7, for instance, indicates that the husband is seven years older than the wife. Most of the entries are in the part of the table lower than the line corresponding to an age difference of zero, indicating that in most cases the husbands are somewhat older than the wives. Across the top are the difference in retirement dates between the husband and wife. A value of 5 indicates that the husband retired five years after the wife.

One might expect a negative relationship in these data, because if the husband is considerably older than the wife, it is not unreasonable to expect him to retire earlier relative to the wife's retirement. There does not appear to be a dominating relationship, though.

The interesting feature of this table is the distinct concentration of retirement when the husband and wife retire simultaneously. Given the way the retirement dates were constructed, what this really means is a concentration of couples reporting the last date of full-time work in the same survey year. Because the surveys were frequently conducted at two-year intervals, the actual retirements could have taken place a year apart one way or the other. Even so, the concentration of retirements at dates so close to one another suggests that couples do tend to retire at about the same time, and that this phenomenon is evident in the data to be used in this study.

III. A Model of Family Labor Supply and Retirement.

In this section we will develop a model of family labor supply and retirement. In the course of investigating the characteristics of this model, we will also find an approach to estimating the model empirically.

A. Relation Of The Present Analysis To The Retirement Literature.

Before engaging in an effort to bridge the structural retirement literature and studies of family retirement, it is appropriate to comment on the state of the structural retirement literature, and where the approach taken in this paper fits into that literature. There are a number of elements which should be included in a fully dynamic specification of a structural retirement model. No single structural retirement model incorporates all of these elements, and they are not all incorporated here either.

To be more specific, Lazear and Moore (1988), Stock and Wise (1990a and b) and Lumsdaine, Stock and Wise (1990, 1992a and b), have emphasized the importance of including the option value of the pension in the opportunity set. Gustman and Steinmeier (1983, 1984, 1985a, 1986a and b) have emphasized the importance of modeling hours constraints on the main job, and the availability of partial retirement only at a lower wage. Berkovec and Stern (1991) and Rust (1990) have emphasized the role of reverse flows in a dynamic model. There is no model which incorporates all of these features. Analyses of retirement which take account of the pension and the option value of the pension focus on the decision to leave the main job, but do not model either the decision to engage in part-time work, or the dynamic of flows among retirement states. Models which take account of flows among retirement states, including reverse flows, ignore the existence of the pension, let alone the option value of the pension.¹³

13. Lack of detailed information on the pension has presented problems for these studies. Gustman and Steinmeier use self reported pension information to fashion the opportunity set. A few studies do have matched, employer provided information on the details of the pension. Fields and Mitchell (1984) use a longitudinal sample of retirees from fourteen firms, while Stock and Wise (1990a and b) and Lumsdaine Stock and Wise (1990, 1992a and b) use one or two firms. Because the data sets used in the studies based on a few firms are not representative of the entire universe of pension-covered workers, empirical findings cannot be

The model used in the present study both extends a structural analysis to incorporate the retirement decisions of households with two full-time earners, and also specifies and estimates the model to incorporate the incentives from the option value of the pension. However, the model focuses only on the decision to reduce work effort below full time. It does not analyze the decision to partially retire, nor does it analyze reverse flows among retirement states. Accordingly, the work presented in this paper falls short of an ideal structural retirement analysis in a family setting. This ideal has yet to be reached in any structural retirement analysis, despite students of the subject having confined such studies to the analysis of behavior of individuals considered in isolation.

B. Model Specification And Estimation Strategy.

The model begins with a fairly standard utility function for the wife which depends on lifetime consumption and labor supply:

$$U_w = \sum_{t=0}^{T-1} \left[\frac{1}{\alpha} C_t^\alpha + e^{X_w \beta_w + \epsilon_w} L_{wt} \right]$$

In this utility function, C is consumption and L_w is a variable which indicates whether the

generalized. More importantly, the data sets which use detailed employer provided information on the pension do not have information on the opportunities or activities of the individual after leaving the firm offering the pension. Other structural analyses using more sophisticated dynamic specifications have ignored the pension (Berkovec and Stern, 1991), or have eliminated pension-covered workers from the sample (Rust, 1990). The present paper utilizes a nationally representative sample, but approximates the incentives from the pension by using worker reported information on the plan.

wife has retired from full-time work.¹⁴ The term $e^{X\beta+\varepsilon}$ determines the relative value of retirement to the wife. The variables in X_w include, among other things, age and health. As the wife becomes older, $e^{X\beta+\varepsilon}$ increases because of the effect of age. Eventually the value of retirement outweighs the value of the wages from working, and the individual retires. The value of ε is an individual effect which determines the relative value of retirement for different women. The higher the value of ε , the more the wife values retirement, and the sooner she will retire, all other things constant.

There are three ways in which this utility function can be construed to be part of a family labor supply model.¹⁵ First, the consumption in this function is not the consumption from the wife's own earnings, but the family consumption financed by the earnings of both husband and wife.¹⁶ Second, one of the variables in X_w is L_h , the retirement status of the husband. If the coefficient of L_h is positive, the wife will value her leisure more highly if the husband is already retired. Finally, the value of ε_w may be correlated with the corresponding value ε_h for the husband. This is the means by which the retirement preferences of the husband and wife may be correlated.

The utility function for the husband is symmetric:

14. A time preference term of the form $e^{-\phi t}$ could also be included in the utility function. However, since no data on consumption (which is poorly measured in datasets such as the NLS) is used in the estimation, the value of ϕ and the value of the constant term in the linear expression βX are not separately identified. As a result, the value of ϕ will be subsumed in the value of the constant term in βX .

15. The specification we employ is a version of what Killingsworth (1983, p. 34) calls the individual utility, family budget constraint model of labor supply of family members.

16. If one prefers to think of the wife consuming a fixed percentage of the total family budget, it would be a simple matter to insert that percentage in front of the family consumption in the utility function. The nature of the model would not change materially.

$$U_h = \sum_{t=0}^{t=T} \left[\frac{1}{\alpha} C_t^\alpha + e^{X_h \beta_h + \epsilon_h} L_{ht} \right]$$

The terms in this function are analogous to the terms in the wife's function, with the term X_h in the husband's function containing a variable L_w indicating the wife's retirement status. Both husband and wife maximize their respective utility functions subject to the constraint that lifetime family consumption cannot exceed family income:

$$\sum_{t=0}^{t=T} e^{-rt} C_t = \sum_{t=0}^{t=T} e^{-rt} W_{wt} (1 - L_{wt}) + \sum_{t=0}^{t=T} e^{-rt} W_{ht} (1 - L_{ht})$$

In this budget constraint, both consumption and wages are expressed in real terms, and r is a real interest rate.¹⁷

It is perhaps easier to analyze the model if we start with a simplified version. This simplified version includes only the wife's age in the vector X . We initially concentrate on the wife's retirement decision. In the simplified model, her utility function is given by:

$$U_w = \sum_{t=0}^{t=T} \left[\frac{1}{\alpha} C_t^\alpha + e^{\beta_0 + \beta_1 A_w + \epsilon} L_{wt} \right]$$

The first problem is to calculate the range of values of ϵ which will induce her to retire at some given age R .

To begin the analysis, we first calculate the marginal utility of income in the model.

17. Some forms of altruism can also be accommodated within this model. For instance, if the husband values the wife's leisure time, and the wife wants to take this into account in choosing her retirement date, the values of β_w can be interpreted as including both her and her husband's value of her leisure. However, if the husband values the wife's leisure time only if he is retired himself, the wife cannot take this into account in the present model, and a more complicated model is required.

Given her husband's income stream and retirement date R_h , and given that the wife is to retire at R_w , the total family discounted income can be denoted as:

$$y = \sum_{t=0}^{t=R_w-1} e^{-\pi t} W_{wt} + \sum_{t=0}^{t=R_h-1} e^{-\pi t} W_{ht}$$

This income is to be divided up among consumption at various points in time so as to maximize $\alpha \sum C_t^\alpha$. The Lagrangian for the consumption decision is:

$$\mathcal{L} = \frac{1}{\alpha} \sum_{t=0}^{t=T} C_t^\alpha + \lambda \left[y - \sum_{t=0}^{t=T} e^{-\pi t} C_t \right]$$

Note that λ in this problem is the marginal utility of income.

Taking the first-order condition for C_t yields:

$$C_t^{\alpha-1} - \lambda e^{-\pi t} = 0$$

Solving the first order condition for C_t and summing over all the periods gives:

$$y = \sum_{t=0}^{t=T} e^{-\pi t} C_t = \sum_{t=0}^{t=T} e^{-\pi t} [\lambda e^{-\pi t}]^{\frac{1}{\alpha-1}}$$

Solving for λ implies:

$$\lambda = \left[\frac{y}{\sum_{t=0}^{t=T} e^{-\frac{\alpha\pi t}{\alpha-1}}} \right]^{\alpha-1} = \left[\frac{y}{k} \right]^{\alpha-1}$$

where k is defined as the denominator of the middle term.

In any given year, the wife would want to work if the value of the wages exceeds the value of the leisure foregone, and otherwise she would want to retire. In terms of the values

we have derived, she would want to work if

$$\lambda e^{-r_t w_{wt}} = \left[\frac{y}{k} \right]^{\alpha-1} e^{-r_t w_{wt}} > e^{\beta_0 + \beta_1 A_{wt} + \epsilon}$$

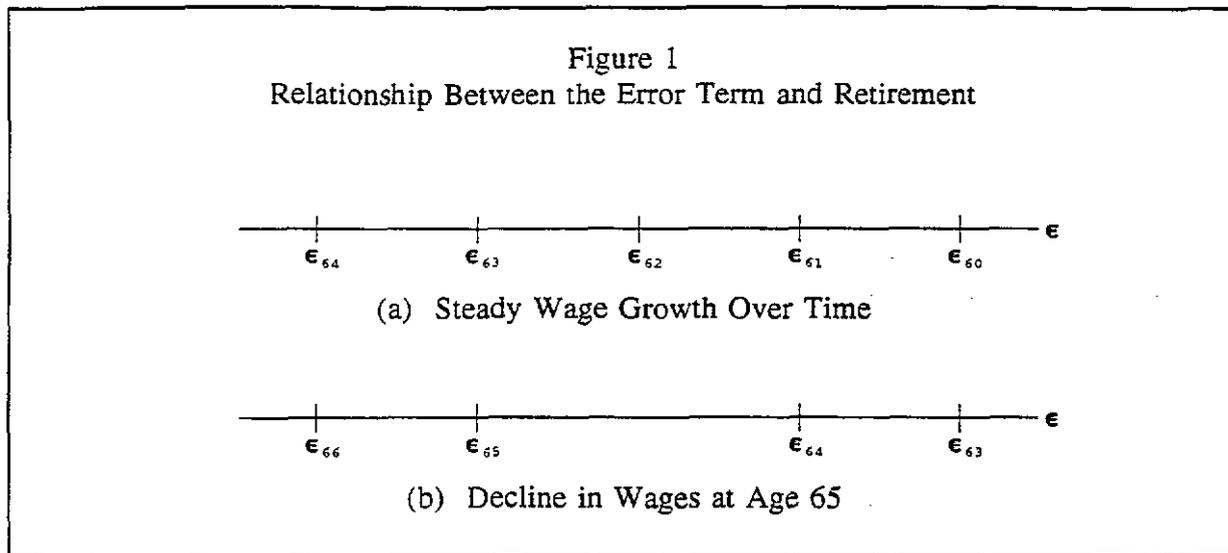
and otherwise she would want to retire. A_{wt} in this equation is the age of the wife at time t . Taking the log of both sides of the above inequality and rearranging implies that the wife would like to continue working as long as:

$$\epsilon < \log(e^{-r_t w_{wt}}) + (\alpha-1) \log \left[\frac{y}{k} \right] - \beta_0 - \beta_1 A_{wt}$$

Another way of looking at this equation is that the right-hand side defines the value of ϵ which makes the wife just indifferent between working and retiring at age R .

Consider a series of potential retirement ages between 60 and 64, and for the time being let discounted wages be constant (i.e., wages grow by the discount factor each year). At age 60, the right-hand side of the equation above will yield some value of ϵ which will make the wife just indifferent between working and retiring. Denote this value of ϵ as ϵ_{60} . At age 61, the right-hand side of the equation will yield another, presumably lower, value of ϵ which will make the wife just indifferent between working and retiring during that year. Denote this value as ϵ_{61} . How much lower is ϵ_{61} than ϵ_{60} ? If wages are growing at approximately the real interest rate, the term $\log(e^{-r_t w_{wt}})$ is approximately constant. Since the intervening year's wages form only a small fraction of the lifetime family income, the term $\log(y/k)$ should also be approximately the same. If wages are growing at approximately the real interest rate, the main difference between ϵ_{60} and ϵ_{61} arises from the term $-\beta_1 A_{wt}$. Since A_{wt} increases by one between age 60 and 61, ϵ_{61} should be lower than

ϵ_{60} by an amount about equal to β_1 .



The same arguments also apply to ϵ_{62} , ϵ_{63} , and ϵ_{64} . If wages are growing at approximately the real interest rate, each of these ϵ 's should be lower than the previous one by about β_1 . The situation is depicted in the top panel of Figure 1. Each of the ϵ 's in this figure depicts the value of ϵ for which the wife is just indifferent between working another year and retiring. The wife's actual behavior depends on the value of ϵ she actually has, that is, on how strong her preferences are for leisure over consumption. If, for instance, the wife has a value of ϵ between ϵ_{62} and ϵ_{63} , she will find it advantageous to work at age 62 but not at age 63. That is, she will retire at age 63. Note that the higher the value of ϵ is for the wife, the earlier she will retire.

What if wages do not grow over time at the real interest rate? The lower panel in Figure 1 illustrates the case where compensation drops at age 65 because the delayed retirement credit in the social security program is not actuarially fair. The differences

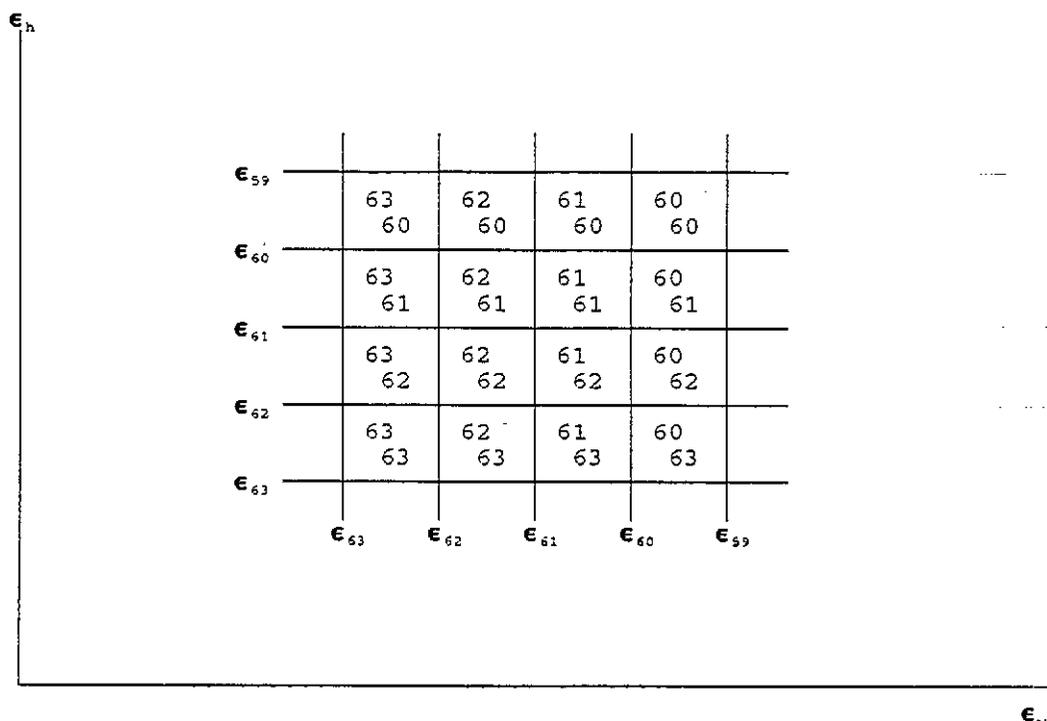
between ϵ_{63} and ϵ_{64} and between ϵ_{65} and ϵ_{66} are approximately equal to β_1 , for the reasons indicated in the previous paragraph. Between ϵ_{64} and ϵ_{65} , however, there is another factor at work. Not only is ϵ_{65} lower than ϵ_{64} because t in the term $-\beta_1 A_{wt}$ is incremented by one unit, but also because the term $\log(e^{-r} w_{wt})$ is lower at age 65 than at age 64. This means that the gap between ϵ_{64} and ϵ_{65} is larger than is the gap between the ϵ 's in other pairs of years, as is shown in the lower panel of Figure 1. If the values of ϵ for different individuals are coming from a relatively smooth distribution, this means that the probability that ϵ will fall between ϵ_{64} and ϵ_{65} will be enhanced relative to other pairs of ϵ . This in turn implies that the individual has a higher probability of retiring at age 65 than at other nearby ages.

The same general line of reasoning applies to the husband. For each potential retirement age, there is a critical value of ϵ for which the husband will be indifferent between retiring and working another year.

Combining the results for husbands and wives leads to a diagram along the lines of Figure 2. In this figure, potential values of ϵ_w are measured along the horizontal axis. The vertical lines in the diagram denote the values of ϵ_w for which the wives are just indifferent between working and retiring at the indicated age, much as in Figure 1. Potential values of ϵ_h are measured along the vertical axis, and the horizontal lines in the diagram are the values of ϵ_h for which the husbands are indifferent between working and retiring at the indicated age. The cells in the figure are combinations of ϵ_w and ϵ_h for which the wife will retire at

Figure 2

Relationship Between Error Term Values and Retirement
When Utility of Leisure Is Not Affected by the Retirement of the Spouse



Upper number in each box represents the retirement age of the wife;
lower number represents the retirement age of the husband.

the upper age indicated in the cell and the husband will retire at the lower age.¹⁸ For

18. The dividing lines in the figure are in fact not quite horizontal and vertical, as they are drawn. The reason is that as you move up one of the vertical lines, the husband's retirement age is decreasing, and this causes a small percentage decline in the lifetime family income. This will cause a small decrease in the term $\log(y/k)$ in the equation defining the critical value of ϵ_w along that vertical line. This means that as you go up the line, the values of ϵ_w

instance, if the values of ϵ_w and ϵ_h are in the upper left cell pictured in the table, the wife will retire at age 63, and the husband will retire at age 60.

In the figure, note that the wife's retirement age decreases as you go to the right in the diagram, reflecting the fact that the higher ϵ_w is, the more the wife values leisure, and the earlier her retirement is likely to be. Similarly, the husband's retirement age decreases as you move up the diagram, for similar reasons. Only a small fraction of the cells are actually plotted in the figure, but the other cells are located in a similar manner. Note also that values of ϵ_w and ϵ_h can be negative as well as positive, so that the cells may be located in the non-positive quadrants of the figure.

Now let us reintroduce the variables reflecting spouse retirement into the model. For the wife, this means that the critical value of ϵ which makes her indifferent between working and retiring at age A_{wt} is given in the following equation:

$$\epsilon_w = \log(e^{-r_w w_{wt}}) + (\alpha - 1) \log\left[\frac{y}{k}\right] - \beta_0 - \beta_1 A_{wt} - \beta_2 L_{ht}$$

Note that in addition to the terms previously discussed, there is now the term $-\beta_2 L_{ht}$ which reflects the retirement status of the husband. If the husband is retired at time t , the critical value of ϵ will be lower. In other words, if the husband is retired at time t , the wife is more likely to want also to be retired at time t . A symmetric relationship also exists for the husband.

increase slightly (recall that $\alpha - 1$ is negative), implying that the line shifts slightly to the right as you go up it. Similar arguments imply that as you go to the right along any of the horizontal lines, the lines tilt slightly upward. However, the effects should not be large as long as the value of α is not excessively negative.

cell gives the wife's retirement age, and the lower number gives the husband's retirement age.

Let us concentrate initially on the cells marked with the dots. These are all cells for which the wife retires at age 60. Among these cells, the lower two correspond to the husband retiring at age 63 or 64; that is, these cells correspond to the husband retiring after the wife. If this is the case, the wife will want to retire if her value of ϵ_w falls between ϵ_{60} and ϵ_{59} , as calculated in the formula above and illustrated in the figure. The upper two cells in this group correspond to the husband retiring at age 60 or 61, that is, before the wife retires. If the husband retires at age 61 or earlier, then the wife will want to retire if her value of ϵ falls between ϵ'_{60} and ϵ'_{59} . Note that the critical values ϵ_{60} and ϵ_{59} will be greater than ϵ'_{60} and ϵ'_{59} by an amount approximately equal to β_2 . That is, if the husband retires at age 61 or before, the wife is willing to retire at age 60 with a lower value of ϵ_w than if the husband retires at age 63 or later. Recall that a lower value of ϵ_w indicates that the wife places a lower value on leisure and retirement than if ϵ_w is high.

Under what circumstances will the wife want to retire at age 60 and the husband at age 62? If the husband retires at age 62, the critical values for the wife are ϵ_{59} (because the husband is not retired when the wife is age 59) and ϵ'_{60} (because the husband is retired when the wife is age 60). That is, if the husband retires at age 62, the wife would wish to retire at age 60 if her value of ϵ_w falls between ϵ'_{60} and ϵ_{59} . Similarly, if the wife were to retire at age 60, the husband would want to retire at age 62 if his value of ϵ_h falls between ϵ'_{62} and ϵ_{61} , as indicated on the horizontal lines in the figure. Thus, it might appear that the wife would want to retire at age 60 and the husband at age 62 if her value of ϵ_w falls between ϵ'_{60} and ϵ_{59} and his value of ϵ_h falls between ϵ'_{62} and ϵ_{61} , which would be a

rectangle in the ϵ_w, ϵ_h space.

However, the diagram indicates that the wife will retire age 60 and the husband at age 62 only in an L-shaped area, not the complete rectangle. What about the remainder of the rectangle? If the values of ϵ_w and ϵ_h fall within the bounds listed above, then it is true that the husband will want to retire at age 62 if the wife retires at age 60, and the wife will want to retire at age 60 if the husband retires at age 62. On the other hand, if within this rectangle, the value of ϵ_w is above ϵ'_{59} and the value of ϵ_h is above ϵ'_{61} , the combination will also fall within the rectangle for the wife retiring at age 59 and the husband at age 61. For points falling within both rectangles, utility is higher for both husband and wife if they retire at the earlier ages. Thus, for such points, although it is true that the wife would retire at age 60 if the husband retires at age 62, and the husband would retire at age 62 if the wife retires at age 60, they would both be better off if they agreed to retire at 59 and 61, respectively. This means that where the rectangles for two retirement age combinations overlap, the rectangle for the younger retirement age combination dominates. The result is the L-shaped areas when both the husband and wife retire at the same time.

This analysis also provides the key to estimating a model of family labor supply and retirement. Given the observed retirement ages for a husband and wife, the preceding analysis shows how to construct the area in the ϵ_w - ϵ_h plane which would result in retirement at those ages. If we assume that values of ϵ_w and ϵ_h over the population come from a parameterized distribution, then integrating the probability density over that area gives the probability that a family with observed characteristics would retire at those ages. This probability, in turn, can form the basis of a likelihood function.

To be more specific, the available parameters in the model are α , the β_w vector for the wife, and the β_h vector for the husband. For specific values of these parameters, and using the observed compensation stream of the two partners and their retirement ages, the appropriate cell boundaries in the ϵ_w - ϵ_h plane can be calculated. ϵ_w and ϵ_h may be regarded as random variables coming from a bivariate normal distribution with zero means, variances given by σ_w^2 and σ_h^2 , and a correlation given by ρ . Given these parameters of the distribution, the probability of retirement at R_w and R_h ages can be calculated as:

$$Pr(R_w, R_h | \alpha, \beta_w, \beta_h, \sigma_w^2, \sigma_h^2, \rho) = Pr(\epsilon_w, \epsilon_h | \sigma_w^2, \sigma_h^2, \rho)$$

These probabilities can be calculated feasibly from a bivariate normal distribution function.

The log-likelihood function is constructed from these probabilities as:

$$\ln \mathcal{L} = \sum_{i=1}^{i=N} Pr(A_{wi}, A_{hi} | \alpha, \beta_w, \beta_h, \sigma_w^2, \sigma_h^2, \rho)$$

where the i subscript indicates individual couples in the sample. The likelihood function can then be maximized by any standard function maximizer, and the resulting values of the parameters will be maximum likelihood estimates. Standard errors for these estimates can be calculated by the Berndt-Hall-Hall-Hausman method.

In the NLS data, it is often not possible to tell exact dates of retirement by looking at the activity on the survey dates, since in most instances the surveys were taken two years apart. For example, it may be possible to tell that the wife was working at age 57 in 1982 and was retired at age 59 in 1984, and that the husband was working at age 60 in 1986 and retired at age 61 in 1987. In this instance, the probability used in the likelihood function is the sum of two of the areas in Figure 3: one area for the wife's retirement at age 58 and the

husband's retirement at age 61, and another area for the wife's retirement at age 59 and the husband's retirement at age 61. This approach is also used if both retirement dates occur after a two-year interview interval; in that case, the probability would be calculated over four of the areas in Figure 3.

To implement the estimation scheme, it is necessary to adjust the estimation procedure to include the information on incomplete employment spells for those who have not yet retired. Dropping the many couples with one or both partners still working at the end of the observation period would subject the estimates to right censoring, so the following method is applied. Suppose that the wife of the couple was observed to retire at age 54, but the husband was still working at the end of the observed period, when he was 63. This implies that the values of ϵ_w and ϵ_h must lie somewhere in the cells defined by the wife's retirement age of 54 and the husband's retirement age of more than 64. In terms of Figure 3, these cells would form a column starting at the cell (54, 64) and extending downward. Note that the starting cell must lie below the L-shaped cells, since the wife is observed to retire before the husband. It is possible to calculate a probability over this area, and it is this probability which is included in the likelihood function.

Similarly, it is possible that at the end of the sample, neither spouse has retired. Suppose that at the end of the sample, the wife was 57 and the husband 63, and that both were still working. This must mean that the wife retired at age 58 or later, and the husband retired at age 64 or later. For such a family, ϵ_w and ϵ_h must lie somewhere on or to the southwest of the L-shaped cell (58, 64). This cell is L-shaped because if both retired the year after the survey ended, they would have retired together. Again, the probability of this area

can be calculated, and that probability can be included in the likelihood function.

One last estimation issue arises due to spikes in pension accruals. Sometimes an individual will retire just after becoming eligible for a more favorable pension. Such occurrences have become known as pension "spikes," and they may arise because the pension formula becomes more favorable if the individual waits until at least the early retirement age to leave the firm. If a person works the year she becomes eligible for a spike and then retires, the critical value of ϵ as calculated above may overstate the individual's actual preferences.

For instance, consider an individual whose wages are \$20,000 and who gains a pension spike worth another \$20,000 if she works her 54th year (this might occur if the early retirement age were 55). The critical value of ϵ calculated from the formula above would be fairly high at age 54, since $\log(e^{-\alpha}w)$ would be high. In other words, looking only at age 54, the individual would want to work unless the value of leisure were fairly high. But consider ages 53 and 54 combined. The individual would want to retire at age 53 if:

$$\lambda[e^{-53r}w_{53} + e^{-54r}w_{54}] < e^{X_{53}\beta + \epsilon} + e^{X_{54}\beta + \epsilon}$$

that is, if the sum of the value of wages at ages 53 and 54 is less than the sum of the value of retirement during those same two years. Substituting in for λ and factoring ϵ out of the right-hand expression, the individual will want to retire at age 53 if:

$$\epsilon > \log[e^{-53r}w_{53} + e^{-54r}w_{54}] + (\alpha - 1)\log\left[\frac{y}{k}\right] - \log[e^{X_{53}\beta} + e^{X_{54}\beta}]$$

If w_{53} is \$20,000 and w_{54} is \$40,000, the value of ϵ calculated from this expression may well be lower than the value calculated from w_{54} and X_{54} alone. In such a case, the

critical value of ϵ for retirement at 54 is the value calculated for the two-year combination of 53 and 54, rather than the value calculated for age 54 alone. The same procedure can be extended for three-year periods, or further if these yield still lower values of ϵ . Thus, for individuals who retire just after they become eligible for a pension spike, the critical values of ϵ are calculated by using a longer period, if necessary.¹⁹

IV. Estimation of the Model.

A. Data Preparation.

The model is estimated using data from the National Longitudinal Survey of Mature Women, 1968-1989. Section II outlined the sample inclusion restrictions.

The data requirements for estimating the model fall into three categories. First, the dependent variable (retirement) must be measured. Determining the last survey of full-time work is described in Section II, and in the estimation, the first date of retirement is taken to be the date of the next survey. For instance, if the wife is last working full-time at age 59 in 1982 and is not working full-time at age 61 in 1984 or any successive survey, she is considered to have retired at either age 60 or age 61 -- it makes no difference to the estimation procedure since the dates are bracketed. The same procedure is used for the

19. If an individual is observed to retire just before a spike (e.g., just before she becomes eligible for a more advantageous pension formula), the area in Figure 3 may become zero unless the coefficient of the age variable in $X\beta$ is large enough. If such observations are taken at face value, they dominate the model by forcing the estimated coefficient of the age variable to be quite large. However, such observations may well be the result of inaccuracies in reporting the date of the pension spike or of some other factor which is not considered in the model. We tried to accommodate this by including a trivariate observation error term, but the resulting model failed to converge. Therefore, in order to avoid having these anomalous cases unduly influence the estimated coefficients, the pension spikes are omitted for 8 wives and 5 husbands who retire just before they are eligible to receive them.

husbands.

The other data requirements are the elements of the X vector in the utility function for the husband and wife, year by year, and the compensation opportunities available for the husband and wife for each year of potential work. The X vectors contain four elements (besides a constant) for both the husband and wife. These are: age, spouse retirement, health, and vintage. The age variable is the individual's reported age in 1967, incremented by 1 for each year. Vintage is simply 1967 minus the age in 1967, and spouse retirement in a given year is whether the spouse is retired at the time. For the binary variable for a health problem, we examined in each year answers to questions about whether health prevented the individual from working or whether health limited the amount or kind of work that could be done. If health was reported as a problem in the first two surveys, we use questions about the length of the health problem to ascertain when the problem began. Otherwise, we look for the first instance in which health was reported as a problem in two consecutive surveys, and the health variable is set to one in all years on or after these two surveys. The idea is to record long-term health problems, and the two-survey requirement is imposed to screen out instances in which there is an isolated survey with reported health problems (a fairly common occurrence in the data). If the problem lasts for two surveys, it is usually apparent in most of the subsequent surveys as well.

The major difficulty in preparing the data set for estimation is to impute compensation streams. The compensation streams consist of three components: wages, pensions, and social security. Details of the construction of the wage offer and pension in the opportunity set are presented in an appendix which is available from the authors on request. Key highlights of

the estimation of the opportunity set are as follows.

Wages

For all the survey years except 1968, the survey asks about annual income from wages and salaries, as well as enough information to construct an hourly wage rate. For those years in which the wife is working full-time at the time of the survey and for which usable annual wage and salary information is available, the annual wage and salary information is used. For other years, annual wages must be imputed. The imputation process uses the tenure, experience, and health coefficients from an hourly wage regression. Hourly wage regressions, computed with and without fixed individual effects, are reported in the appendix. Those with fixed individual effects are used.

Pensions

The next element of compensation is the pension profile. Information on pensions comes mainly from questions that were asked in 1982, 1986, and 1989. In the first two of these years, the survey asked about whether the respondent was eligible for a pension from the current or past job. If the individual was eligible for a pension from either the current or past job, the survey inquired about the years of service in that job. If the pension was for the current job, the survey also asked about the age of initial eligibility for benefits. The 1989 survey expanded the questions so that they include the early and normal retirement ages on the present job, the early and normal retirement ages for pensions on previous jobs, and the amounts that the individual would be eligible to receive on those dates. In addition, the survey asked about the actual amount of benefits received if the respondent reported that benefits were currently being received.

The full pension is assumed to be determined by the formula $P = g W T$, where P is the annual pension benefit, W is the final wage, T is years of service, and g is a generosity factor.²⁰ For individuals who retire after the early retirement date but before the normal retirement date, the full pension is reduced by 4.9% per year. This is the average weighted reduction rate in Hatch et al., Table 4-8. For individuals who retire before the early retirement date, the reductions are actuarial. This arrangement, which is the basis of the early retirement spike described by Lumsdaine, Stock, and Wise, also appears to be fairly common in the plans in the 1983 Survey of Consumer Finances. The generosity figure is calculated on the basis of the individual's expected or actual pension benefits; if no figure is given, we use 1.6%, which is the median figure of plans for which we do have enough information to calculate it. We assume that, once begun, firms increase the value of the pensions by 37.9% of the inflation rate, the figure found in Allen, Clark, and Sumner (1986).²¹

Pension compensation (pension accrual) is the amount by which an additional year's work increases the expected present value of the pension. Several factors come into play in determining the change in present value from an additional year's work. An additional year's

20. Given the incidence of severe reporting error in self reported plan type which we found in Gustman and Steinmeier (1989), we have not attempted to distinguish those with primary defined benefit plans from those with primary defined contribution plans. Instead, we have treated all workers as if they were covered by a primary defined benefit plan.

21. An alternative procedure to impute pensions would be to match each observation in the NLS to pensions from the Survey of Consumer Finance pension file on the basis of occupation and industry. The problem with this approach is that it does not insure that the early retirement dates would be correctly aligned with the date the respondents gave. This is a serious drawback, since the early retirement age is probably the most important date in a pension profile, marking as it does a sharp drop in the accumulation rate of pension benefits. Another idea would be to match with the SCF on the basis of occupation, industry, and early retirement dates, but the SCF is too thin for this detailed a match.

tenure will increase the annual pension through the benefit formula. If the nominal wage increases during the course of the year, the annual pension amount increases proportionately. If the work is before the normal retirement date, the reduction for early retirement will be less. The present value of these increases is obtained by multiplying the pension amounts times the appropriate survival probabilities in future years and discounting back to the present.²² Offsetting these factors is the fact that an individual who is already eligible for the pension will lose the current year's benefits because pensions are rarely paid to continuing employees. The net effect is to cause pension accruals to change relatively smoothly except at the early and normal retirement dates, when the amount of the accruals drops sharply.

Further details of the pension calculation are presented in the appendix. Although the imputations are necessarily approximate, they do capture the main effects of pensions on compensation, which are the declines at the early and normal retirement ages. Especially important are the declines at the early retirement age. We do know, for most individuals, when the early retirement age is, and these calculations do capture, however approximately, the decline in compensation that usually occurs at that age.²³

Social Security

22. The real discount rate is taken to be equal to the real growth rate of adjusted hourly earnings. Both this and the expected inflation rate are the average values observed over the ten years previous to the year in question. The growth rate of adjusted hourly earnings is less than the interest that could be obtained with 10-year treasury bonds over the post-war period, but is about equal to the interest on 3-month treasury bills and probably exceeds what most households could obtain with banks or money market funds, especially if after-tax interest is considered.

23. For discussion of misreporting of early and normal retirement dates in the SCF, see Gustman and Steinmeier (1989).

The derivation of social security accruals is less subject to error than are pension accruals, primarily because all elements of the social security formula are known and apply to everyone. The computations do take into account the increases in the normal retirement age and in the delayed retirement credit scheduled to take effect, as specified in the 1983 Social Security Amendments. For the husbands, the computations take into account the effect of his retirement age on the eventual widow's benefits.

The present value calculations are similar to those for the pensions. Working another year can change the present value of social security benefits for three reasons. First, higher real wages during prime working years are likely to displace earlier wages from the average earnings calculations, increasing benefits. This effect is enhanced for those over sixty years of age because under the calculation of average indexed monthly earning (AIME), the basis for calculating monthly benefits, wages are indexed only to age 60. Second, if the individual is over the early retirement age, working another year will reduce the early retirement reduction factor or increase the delayed retirement credit. Finally, if the individual is over the early retirement age, working another year may cause benefits to be lost to the earnings test. In general, the effects of the social security system are to cause reductions in compensation primarily at the normal retirement age, since the delayed retirement credit is not sufficient to offset the loss of benefits from continuing working.²⁴

The overall compensation paths for the wife and the husband are calculated by adding to the wage the accruals for social security and, if appropriate, the pension. These are then

24. For further discussion of the social security rules and estimates of the incentives they create, see Gustman and Steinmeier (1985b, 1991).

combined with the retirement ages of both spouses and the values of the X vectors to provide the data set for estimation. The final data set has 564 observations of couples who satisfied all of the inclusion criteria and who had at least one wage and salary figure on which to base the compensation paths.

B. Estimation.

The estimates for various forms of the model are given in Tables 4 and 5. Table 4 reports on results when the spouse's retirement is ignored or is treated as exogenous. Table 5 reports on estimation of the model on the assumption that retirement is jointly determined, applying the estimation procedure outlined in Section III above. The estimated parameters in these tables include the exponent of the consumption term in the utility function (α), the coefficients of the linear forms for both the wife and the husband (the β 's), the variances in retirement tastes for both the wife and the husband (the σ 's), and when retirement is treated as jointly determined, the correlation of these tastes (ρ). The β vectors for both the wife and the husband include a constant, age, whether or not the spouse is retired at the time, whether or not there is a health problem at the time, and the vintage of the individual.

The first two columns of Table 4 contain estimates for the wife and the husband separately, sans any variables which might link the two retirement dates. In both cases, the earnings of the remaining spouse are taken as exogenous. All of the variables are significant at conventional levels. The coefficients on the consumption terms indicate, for both partners, a moderate desire to smooth consumption over time. For this parameter, a value of unity would indicate no desire to smooth consumption over time, while a large negative value would indicate a great desire to smooth consumption. Since the difference between the wife's

Table 4
Model Estimates Treating Spouse Retirement as Exogenous

		No Interdependence		Interdependence	
		Wife's Equation (1)	Husband's Equation (2)	Wife's Equation (3)	Husband's Equation (4)
α	Consumption	-1.549 (-3.13)	-1.355 (-2.35)	-1.563 (-3.15)	-1.256 (-2.27)
β_0	Constant	-19.042 (-21.83)	-18.350 (-11.32)	-19.401 (-20.82)	-17.489 (-10.83)
β_1	Age	0.578 (4.52)	0.65 (3.24)	0.528 (4.40)	0.599 (3.25)
β_2	Spouse Retired			0.953 (3.05)	1.302 (2.71)
β_3	Health Problem	1.050 (2.82)	1.974 (3.04)	1.036 (2.79)	2.005 (3.09)
β_4	Vintage	0.101 (2.20)	0.127 (2.28)	0.094 (2.08)	0.108 (2.08)
σ_ε^2	Std. Dev. of ε	2.878 (4.79)	3.427 (3.72)	2.864 (4.73)	3.427 (3.72)
	Sample Size	564	564	564	564
	Log Likelihood	-668.2	-743.4	-658.67	-732.72

Notes:

(a) Figures in parentheses are asymptotic t-statistics.

(b) For age and vintage, the actual variables are (age-55) and (vintage-30).

parameter and the husband's parameter is less than their standard errors, it appears unlikely *a priori* whether these two values are significantly different.

For the wife, the coefficient of the age term is 0.578. This indicates that the utility of retirement is rising by about over 78% ($e^{0.578}-1$) per year during the period of retirement. The procedure estimates this effect primarily by looking at how much retirement is concentrated in years where compensation falls considerably. In terms of Figure 3, if the coefficient of age is

smaller, then the areas in the figure will be generally smaller, except where there are large declines in compensation. This means that years having larger declines in compensation will have a relatively larger area in the figure, and more individuals should be retiring in those years. If the coefficient of age is larger, then the areas in Figure 3 are much larger, and a decline in compensation, though increasing the area by the same absolute amount, will not result in nearly as large a proportionate increase in area. The result is that if this is the case, a decline in compensation should increase retirement more moderately.

Among the other parameters, having a health problem has the same effect on retirement as being a not quite two years older, and the effect of vintage is modest. The standard error of the retirement taste term ε is 2.878, which amounts to over five times the coefficient of the age term. It is fairly clear that the taste term plays a very large role in accounting for the large range of ages at which various individuals retire.

For the husband, the value of the health coefficient amounts to a little over three times the magnitude of the coefficient of the age term, and the standard deviation in tastes for retirement are a little over five times the size of the coefficient of the age term. The age term itself is somewhat larger than it is for the wives. This may reflect two possibilities. First, the husbands may be less sensitive to monetary incentives to retire. From the discussion above, a higher coefficient on the age term implies that the husbands are more likely to retire at ages determined by their values of ε , regardless of changes in compensation incentives. Alternatively, the high coefficient may arise if the ages of compensation decreases are mismeasured, which would reduce the correlation between retirement and compensation decreases. It is plausible that such mismeasurement may be more severe for husbands than

for wives, since the wives and not the husbands were the primary respondents in this data set.

The last two columns of Table 4 continue to treat the spouse's retirement behavior as exogenously determined, but for each spouse allow the preference for leisure to depend on the retirement status of the other. The husband's retirement status has a significant effect on the valuation of leisure by the wife, with the effect of the husband being retired increasing the value of leisure by an amount equivalent to aging two years. There is only a small effect of including husband's retirement status on the other parameters. With regard to the estimate of the utility function of the husband, the effect of wife being retired is also significant, and equal to the effect of an additional two years of age. As in the wife's equation, there is only a modest effect on including spouse retirement status on the coefficients of other variables.

Given that each spouse's retirement decision is influenced by the retirement status of the other, it becomes important to recognize in the estimation the joint determination of retirement, and not to assume that the retirement status of one spouse is an exogenous determinant of the retirement status of the other. For if the decisions are truly joint, then treating the decisions as exogenous to one another will lead to an overstatement of the effect of the retirement status of one spouse on the retirement decision of the other. For instance, the wife may tell the husband "If you retire early, I will too, and we can enjoy our retirement years together." The husband, meanwhile, may be thinking the same thing, and both husband and wife of this couple will retire early. To the researcher studying only the wife and taking the husband's retirement as exogenous, it will appear that an early retirement by the husband is associated with an early retirement by the wife. In fact, the early retirement of the husband is partly the result of the wife's own preferences, which will be mistakenly attributed to the

husband's retirement.²⁵

Moreover, the preferences of the two spouses may be correlated. If individuals who have high preferences for early retirement tend to marry one another, then the retirement age of the spouse will be correlated with the individual's own preferences. Again, the investigator who is taking the husband's retirement as exogenous will observe that if the husband retires early, the wife is likely also to retire early. Thus it may mistakenly be concluded that the husband's retirement has a large influence on the retirement date of the wife, when in fact it is the correlation in retirement preferences that is driving both retirements.

The estimates in Table 5 are designed to deal with these issues. The first column of Table 5 provides a basis for comparison with equations reported in columns 2 through 5. It imposes the restriction that the consumption parameter is the same between husbands and wives, but as in columns 1 and 2 of Table 4, excludes the indicator of whether the spouse is retired. This is necessary if the couple is to share a common consumption stream. The resulting coefficient of α is about halfway between the estimates for wives and husbands in columns 1 and 2 of Table 4. The values of the likelihood statistics indicate that the

25. Notice also that the results of the retirement decisions made by each spouse, as estimated in columns 3 and 4 of Table 4, should not be taken as additive. For if the preferences are interdependent, each system of equations will be estimated to explain whatever joint peak there is in the retirement outcomes of both spouses. Together, the estimates are likely to "overexplain" any joint retirement behavior.

Table 5
Model Estimates With Retirement Jointly Determined

Consumption	(1)	(2)	(3)	(4)
α	-1.465 (-4.13)	-1.427 (-4.00)	-1.612 (-4.39)	-1.534 (-4.04)
Wife's Parameters:				
β_0 Constant	-18.151 (-23.63)	-17.762 (-23.42)	-19.410 (-32.75)	-18.615 (-26.78)
β_1 Age	0.566 (4.91)	0.555 (4.89)	0.542 (5.94)	0.531 (5.19)
β_2 Husband Retired		0.092 (0.27)		0.095 (0.27)
β_3 Health Problem	1.019 (2.96)	1.003 (2.96)	1.005 (3.30)	0.981 (3.06)
β_4 Vintage	0.099 (2.22)	0.096 (2.16)	0.088 (2.21)	0.084 (2.01)
σ_w^2 Std. Dev. of ϵ_w	2.821 (5.33)	2.787 (5.40)	2.761 (6.16)	2.713 (5.78)
Husband's Parameters:				
β_0 Constant	-19.560 (-13.41)	-19.383 (-12.71)	-20.744 (-17.60)	-20.031 (-15.67)
β_1 Age	0.670 (3.67)	0.633 (3.63)	0.632 (4.32)	0.613 (4.07)
β_2 Wife Retired		1.256 (2.15)		0.582 (1.12)
β_3 Health Problem	2.043 (3.41)	2.119 (3.38)	2.017 (3.91)	2.046 (3.71)
β_4 Vintage	0.130 (2.43)	0.114 (2.19)	0.126 (2.69)	0.109 (2.29)
σ_h^2 Std. Dev. of ϵ_h	3.530 (4.25)	3.572 (4.21)	3.411 (5.06)	3.394 (4.76)
Correlation:				
ρ			0.312 6.46	0.236 4.13
Sample Size	564	564	564	564
Log Likelihood	-1411.80	-1401.11	-1397.43	-1394.47

Notes to Table 5: (a) Figures in parentheses are asymptotic t-statistics. (b) For age and vintage, the actual variables are (age-55) and (vintage-30).

imposition of this constraint is not statistically significant.²⁶

Column 2 of Table 5 looks again at the possibility that spouses might find retirement more attractive if the other spouse is retired by including spouse retirement variables in the β vectors. The difference between these results, and those in columns 3 and 4 of Table 4, is that the estimation treats spouse's retirement status as endogenously determined.

In Table 4, with spouse retirement taken as exogenous, the spouse retirement coefficients depend on the degree to which the individual retires on the same date as his or her spouse. In Table 5, the spouse retirement coefficients depend on whether, when one spouse retires at a date when compensation drops substantially, the other spouse retires also. Comparing spouse coefficients for the husband and wife, the coefficients suggest that when wives retire on their early retirement dates, the husbands are a lot more likely to retire at that time also than if the situation is reversed.

When the spouse's retirement status is treated as endogenously determined, the coefficient of the husband's retirement status affecting the wife's retirement becomes small and insignificant. The coefficient for the wife's retirement status in explaining the husband's retirement, however, is of the proper sign and has an estimated t-statistic of over two. This coefficient indicates that, for husbands, having a wife that is retired has the same effect on the desirability of retirement as does being about two years older. These estimates suggest that the husband's retirement does not affect the retirement preferences of the wife, but the wife's

26. The log-likelihood when this constraint is not imposed (-1411.60, which is the sum of -668.20 and -743.40) is only 0.20 above the log-likelihood of the constrained equation. This implies that the likelihood ratio test statistic is 0.40, considerably less than the $\chi^2(1)$ critical value of 3.84 at a 5% significance level. Hence, the test does not reject the hypothesis that the value of α is the same between wives and husbands.

retirement has a notable effect on the retirement preferences of the husband. An exposte rationale for this result is that some husbands do not wish to find themselves with more time to take care of the household while the wife is still devoting a substantial part of her time to market work. But stronger confirmation of this result and investigation of further implications would be required before any conclusion is drawn.

Column 3 looks at an alternative explanation for couples retiring at around the same time. If men and women who prefer to retire early tend to marry one other, and men and women who prefer to retire later (and consume more over their lifetimes) tend to marry each other, the values of ε_w and ε_h in the model will be correlated with each other. This column introduces the correlation term ρ into the estimation (and omits the β_2 's). As can be seen from the table, the estimated value of ρ is about 0.3, and it is highly significant.

Finally, Column 4 in Table 5 allows these two competing explanations to be included in the same equation. Compared to Column 3, the value of ρ is now only about three-quarters as large, but it is still highly significant. The spouse retired coefficients now both have the proper sign, but neither one has a t-statistic above 1.2. However, the combined effect of the two of them is marginally significant, as can be seen by comparing the log-likelihood of Columns 3 and 4. Adding these two variables to a specification which already includes ρ increases the log-likelihood by 2.97. This implies a likelihood ratio test statistic of 5.94, which is approximately equal to the $\chi^2(2)$ critical value of 5.99 at a 5% confidence level. Although the estimation procedure is unable to separate out the spouse retirement coefficients, they are jointly significant. Moreover, the point estimate of the coefficient on the spouse retirement variable is over six times as large for the husband than it is for the wife. Because

the relative effects of the β_2 's and ρ cannot be inferred simply by comparing the relative size of the parameters, that issue will be investigated by simulation.²⁷

Before that, we should comment on the magnitudes of the other parameters in this final specification. The magnitude of α in this specification is slightly greater (in absolute value) than in the specification without the correlation of the spouse retirement variables (-1.53 vs -1.46). The coefficients in the β vector for the wife look very similar to the coefficients for the specification with no spouse interactions included (Column 1 of Table 5). The same is true for the coefficients in the β vector for the husband.

V. Simulations Of The Effects Of Husband Wife Retirement Interactions.

The model parameters do not indicate the relative importance of the interdependence of retirement decisions and correlation of tastes in shaping retirement outcomes. The figures in this section, which are based on simulations for the sample, will shed more light on this issue.

The simulations are performed as follows. First, the paths of the X vectors and the compensation stream for both husband and wife are calculated as indicated in the first part of the preceding section. The model parameters are taken from the last column of Table 5, that is, from the full model. Given the estimated variances of the error terms and the correlation

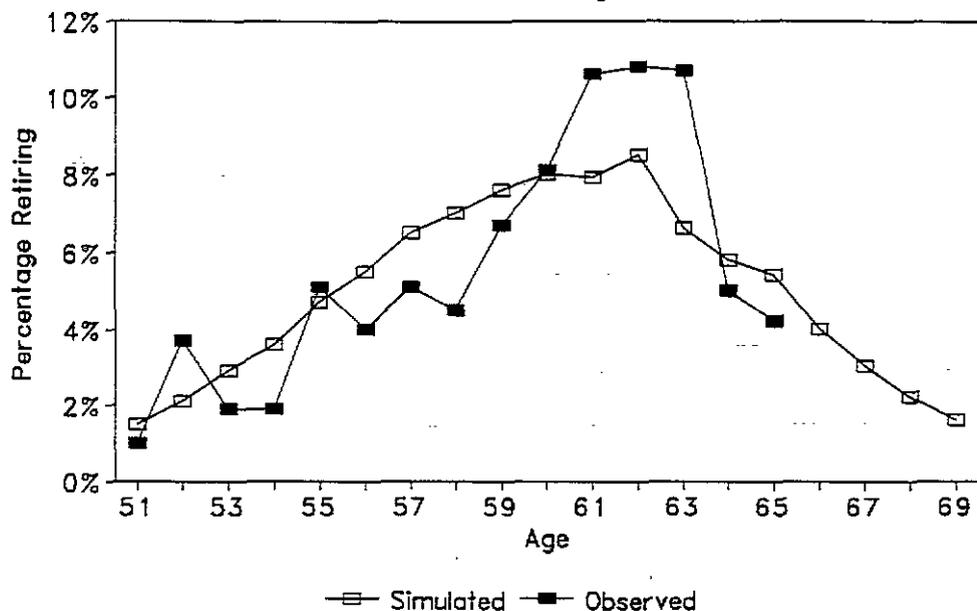
27. Identification of these effects comes by examining the peak of joint retirement. To the extent that the cause of the peak is the spouse retirement variables, there should be a peak of couples retiring at exactly the same date (See Figure 3). To the extent that the cause is a correlation of preferences, there should be a broad peak of couples who retire within a few years of each other, but not necessarily at exactly the same time. Separate identification of the two spouse retirement variables comes by examining when the joint retirement tends to occur. If joint retirement occurs at the time the wife (but not the husband) has strong financial incentives to retire, this would indicate that it is the husband who values having the wife at home. If it occurs when the husband has incentives to retire, the reverse would be true.

between them, a random draw is made to obtain ε_w and ε_h . These values of the X vectors, the compensation streams, α , the β vectors, and the two ε 's then imply specific retirement ages for the wife and the husband.

The solution method can be visualized easily with the aid of Figure 3. We begin by setting the potential retirement age of the husband at 75 and the potential retirement age of the wife at 73. This correspond to an L-shaped area far down and to the left in Figure 3. We then calculate whether the values of ε_w and ε_h are sufficiently high to move to the L-shaped area immediately up and to the right. If they are, we then repeat the process until we reach an L-shaped area for which either ε_w or ε_h (or both) is insufficiently high to move to the next L-shaped area. If both are too low, this L-shaped area is the solution. If ε_h is too low but ε_w is not, we see whether ε_w is sufficiently high to move to the rectangular area immediately to the right of the L-shaped block, and we continue this process until ε_w is no longer high enough to move any further to the right. At that point, we check to see whether ε_h is in the proper range for that rectangle, and if not we move up or down and repeat the process. A symmetrical process is done if ε_w is too low to move to the next L-shaped box, but ε_h is not. The objective of this procedure is to reach that area which corresponds to the random draw for ε_w and ε_h . The ages for that area are the solutions for the simulation, given the observed opportunity set and the draw for ε_w and ε_h . In the results reported below, there are 1000 simulations per individual.

Figure 4a presents the distribution of retirement ages among wives from the base simulation, and Figure 4b presents the corresponding distribution for husbands. For

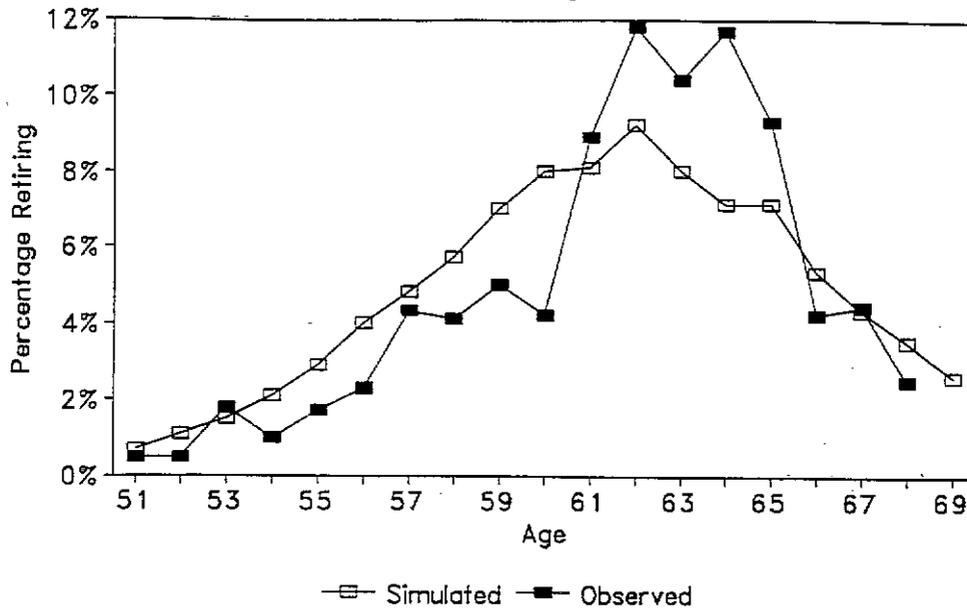
Figure 4a
Retirement Among Wives



comparison, the distributions of observed rates are also plotted in these figures.²⁸ Note from Figure 1 that if tastes for leisure (the ϵ 's) are distributed normally, then the distribution of retirement ages should be approximately normal also. The observed distributions have large humps in the 61-63 age range for wives and in the 62-64 age range for husbands. However, there is no noticeable spike at age 62, which is the Social Security early retirement age and is also a common retirement age in many pension plans. As a result, the model overpredicts

28. The observed rates are calculated by taking percentages of the sample who were retired at different ages and first differencing the results. While the rates in the early 50's are based on sample sizes of over 500, the sample sizes decline in the late 50's and 60's until they reach 54 at age 65 for wives and 164 at the same age for husbands.

Figure 4b
Retirement Among Husbands



retirement in the late 50's and underpredicts it in the 60's. The model predicts faint peaks at ages 62 and 65, reflecting the role of pensions and the downturn in the social security actuarial adjustment rates at age 65, but the lack of distinct and large peaks simply mirrors the lack of distinct peaks in the observed data at those ages.²⁹ Finally, both the simulated and actual distributions are shifted to the left for wives, reflecting the fact that wives in general retire at an earlier age than their husbands. (Gustman and Steinmeier, 1986a and b).

29. Note that the peaks in the hazard rate would be sharper at the relevant ages associated with a particular pension profile, but the results in Figure 4 are averages among people covered by pension plans with different requirements for early and normal retirement and with different employment histories.

Figure 5
Simulated Hazard Rates

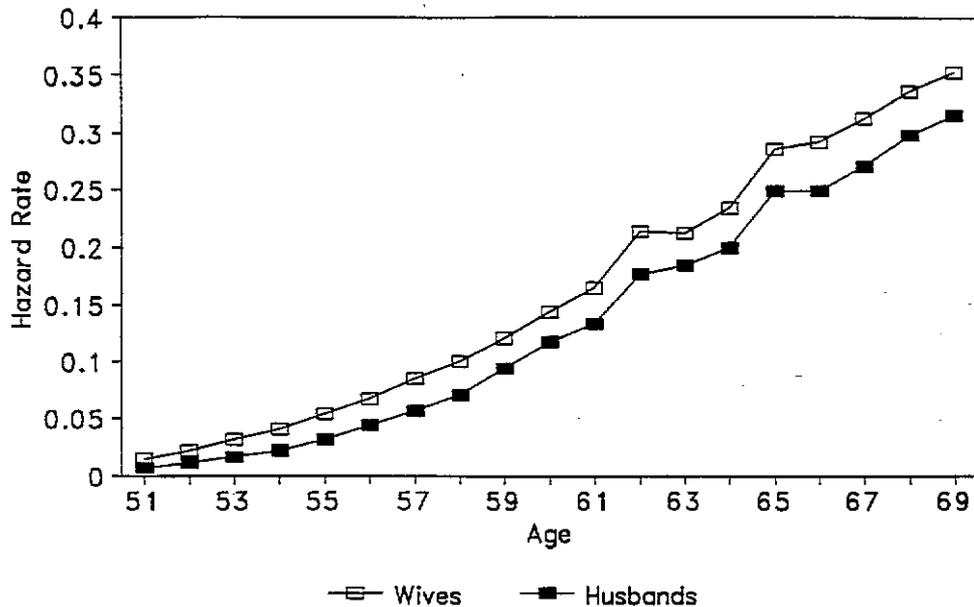
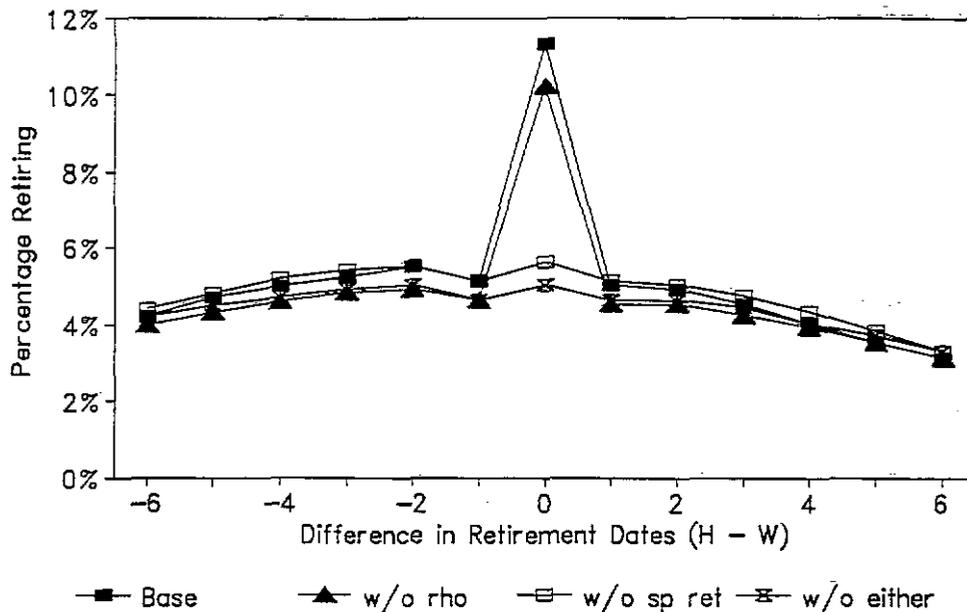


Figure 5 presents the simulated hazard rates for both husbands and wives. Because the retirements in Figure 4 are from a smaller base at older ages, the hazard rates are generally increasing. The rates are higher for women throughout the entire age range displayed. The small peaks at age 62 and 65 are upward bumps on the general trend.

Figure 6 investigates a central concern of this paper, the coordination of retirement by husbands and wives. The figure indicates the differences in retirement dates between the husband and wife under alternative specifications for the model. As can be seen, when the full model is estimated, there is a sharp peak at zero difference in retirement dates, with over 11 percent of the couples retiring at the same time, whereas for other age differences, the

Figure 6
Effects of Joint Retirement Variables



probability ranges around 4 to 5 percent.

It is clear from Figure 6 that the source of interdependence in retirement is in the coordination of preferences among spouses and is not due to the correlation of incentives between the opportunity sets facing husbands and wives. This result can be seen by comparing the probability of joint retirement under the full "base" model, where all sources of retirement are included, with the probability of joint retirement when all sources of interdependence in the utility function are suppressed (but the opportunity sets remain unchanged).

Investigating the source of the interdependence within the utility function, it can be seen

that the peak in the probability of retiring at the same age remains when ρ (the correlation between ε_w and ε_h) is set equal to zero, falling by only about one percentage point. In addition, setting $\rho = 0$ has virtually no effect on the distribution of retirement ages among either husbands or wives (result not shown). The conclusion is that the modest estimated correlation between ε_w and ε_h produces at most a modest increase in the tendency of couples to retire around the same time.

The remaining two lines in Figure 6 indicate that joint retirement is due almost entirely to the presence of spouse retirement status in the utility function.³⁰ The tendency for couples to retire together is reduced from about 11% to about 5% in the simulations where the β_2 's are set equal to zero, whether or not the correlation in the ε 's is included. In these simulations, retirement at exactly the same year is only a little above retirement at dates which are a year or two apart, in marked contrast to the situation in the previous two simulations. Moreover, retirement at dates which are a year or two apart (but not the same) is little affected as compared to the simulations in which the β_2 's are included.

In sum, from the simulations it is clear that most of the spike at zero years difference in retirement dates is due to the spouse retirement variables. Dropping ρ diminishes this peak only slightly, but dropping the β_2 's almost eliminates it. The reverse seems to be true for differences in retirement dates of one year or more: in this case, the

30. Simultaneously, the constants in the linear forms $X\beta$ are increased to offset the omission of the spouse retirement variables. Otherwise, omitting the spouse retirement variables will reduce $X\beta$ generally and will lead to an increase in the average retirement age. Since the spouse retirement variables average one half around the years of the retirement decision, the constant terms in the linear forms $X\beta$ are increased by one half of the corresponding β_2 parameters.

difference in retirement dates seems to be due mostly to the correlation between ε_w and ε_h . Dropping the β_2 's has very little effect on those differences, but dropping ρ lowers them to about the level found in the model in which all interdependence is suppressed.

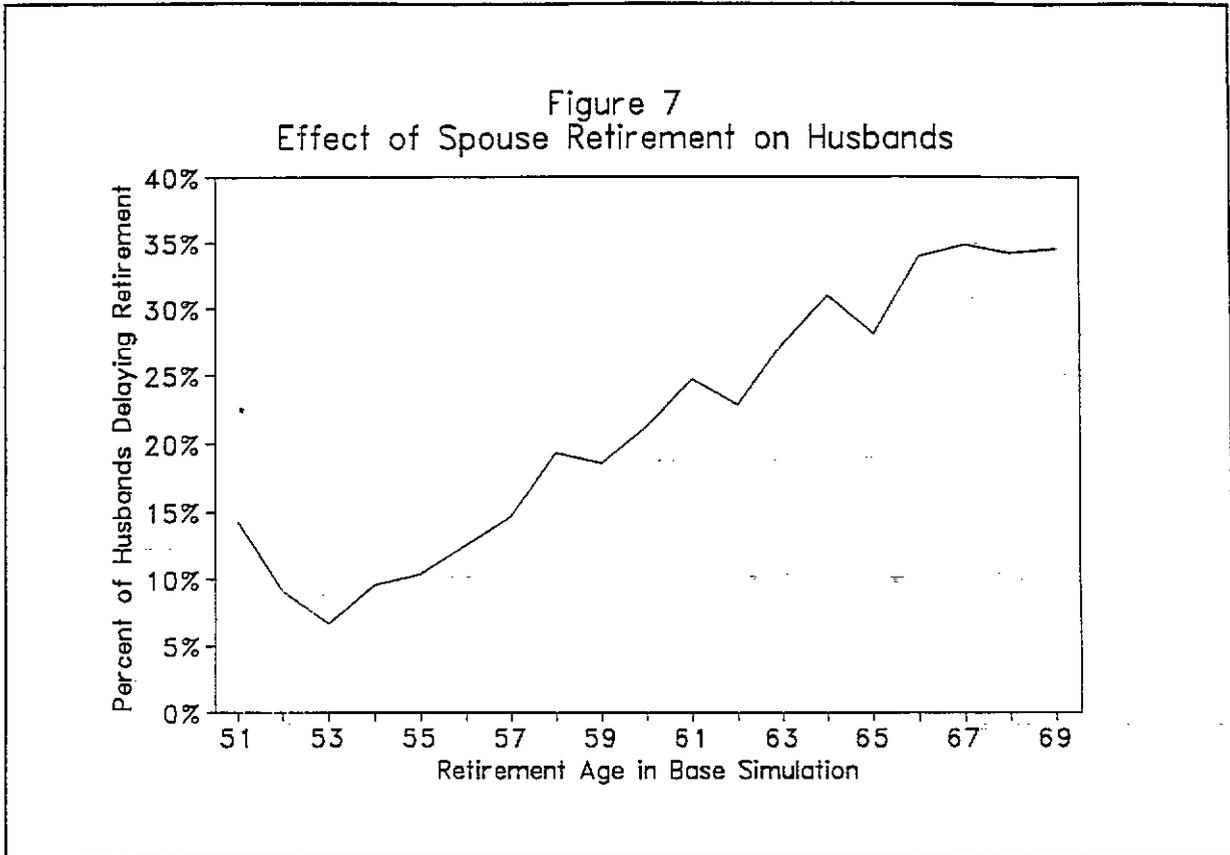


Figure 7 investigates the magnitude of the impact of the spouse retirement variables. Since the husband's coefficient is considerably larger than the wife's coefficient, the effect of that coefficient is presented. In the simulation, the wife's retirement age is held to be exactly the same as in the base simulation. The coefficient of the spouse retirement variable in the husband's utility function is then set equal to zero, and the husband's optimal retirement age is calculated. If the wife is still working, this will have no effect, but if the wife is already

retired, the husband may be induced to work a while longer.

The vertical axis on this figure measures the percentage of individuals retiring at that date who are induced to delay retirement by a year. For example, 8 percent of the husbands retire at age 60 in the base simulation. If the effect of spouse retirement is suppressed, 1.7 percent of the husbands work another year rather than retiring at age 60. The ratio of the two, which amounts to around 21 percent, is what is plotted in the figure.

After a modest decline in the early 50's, the chart turns up for the remainder of the age range. This means that as the age increases, the percentage who would want to work another year increases. One reason is that as the husband's retirement age increases, it is more probable that his wife has already retired, in which case setting the spouse retirement coefficient to zero will make a difference. The dips at age 62 and 65 occur because at these ages pensions and social security are larger factors in the husband's retirement decisions, and whether or not the spouse has retired is less of a factor.

The average over all ages is about 24 percent. In the base simulation, the fraction of cases where the wife is already retired when the husband retires is about 57 percent. This means that in cases where the wife retires first, the about 42 percent of the husbands would delay retirement by another year if the spouse retirement effect were inoperative. For the affected group, this translates into an increase in the average retirement age of about 5 months. While this is not an overwhelming effect, it is in the same order of magnitude as the average effects which have been reported in the literature for pensions and social security.

VI. Summary and Conclusions

This paper has specified and estimated a structural retirement model in which retirement

decisions by husbands and wives are jointly determined. The parameters estimated for the model generate behavior that brings the retirement dates of husbands and wives closer together than would otherwise be the case, with the probability of retirement at the same time being about 11 percent of couples.

Although the estimation procedure is unable to separate out the spouse retirement coefficients, they are jointly significant. Thus, both of the competing explanations (the spouse retirement explanation and the preference correlation explanation) are significant in the data. Simulation indicates the relative importance of each of these sources of interdependence in increasing coordination of the retirement outcomes of spouses. In contrast to the support found for the hypothesis that coordination of retirement outcomes is due to interdependence in preferences, support is not found for the hypothesis that coordination in retirement of spouses results from coordination in the incentives created by the spouse's opportunity sets.

There is some suggestion in the data that the wife's retirement decision is not strongly influenced by the husband's, but the husband's decision is more strongly influenced by the wife's. Although such a result is consistent with an interpretation that the husband would rather not face the house and related work load alone, this result is a preliminary suggestion that requires further support.

Finally, because the National Longitudinal Survey of Mature Women provides the most recent, nationally representative longitudinal sample available for a study of retirement behavior, the results reported in this study reflect the labor market trends shaping the retirement behavior of women. Although many members of the cohorts included in the NLS have not yet retired, the estimating technique has explicitly accounted for incomplete

employment spells.

References

- Allen, Steven G.; Clark, Robert L. and Sumner, Daniel. "Post Retirement Adjustments of Pensions", *Journal of Human Resources*, Winter 1986, 21(1), pp. 118-137.
- Berkovec, James C. and Stern, Steven. "Job Exit Behavior of Older Men", *Econometrica*, January 1991, 59 (1), pp. 189-210.
- Burtless, Gary and Moffitt, Robert A. "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.
- Clark, Robert. "Labor Market Responses of Divorced And Widowed Women", Report To The National Commission On Employment Policy, March 1984.
- Clark, Robert and Johnson, Thomas. "Retirement In The Dual Career Family", Final Report To The Social Security Administration, June 1980.
- Fields, Gary S. and Mitchell, Olivia S. *Retirement, Pensions and Social Security*. Cambridge: MIT Press, 1984.
- Gustman, Alan. L. and Steinmeier, Thomas L., " Minimum Hours Constraints And Retirement Behavior", in *Contemporary Policy Issues*, a Supplement to *Economic Inquiry*, 1983, pp. 77-91.
- _____. "Partial Retirement And The Analysis of Retirement Behavior", *Industrial And Labor Relations Review*, April 1984, 37(3), pp. 403-415.
- _____. "The Effects of Partial Retirement on Wage Profiles of Older Workers", *Industrial Relations*, Spring, 1985a, 24(2), pp. 257-265.
- _____. "The 1983 Social Security Reforms and Labor Supply Adjustments of Older Individuals in the Long Run," *Journal of Labor Economics*, Vol. 3, 1985b, pp. 237-253.
- _____. "A Structural Retirement Model," *Econometrica*, May, 1986a, 54(3), pp. 555-584.
- _____. "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.
- _____. "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, 1989. pp. 33-106.

_____. "Changing Social Security Rules For Workers Over 65: Proposed Policies And Their Effects", *Industrial And Labor Relations Review*, July 1991, Vol. 44, No. 4, pp. 733-745.

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

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

Killingsworth, Mark R. *Labor Supply*, Cambridge: Cambridge University Press, 1983.

Lazear, Edward P. and Moore, Robert L. "Pensions and Mobility," in *Pensions In The U.S. Economy*. Edited by Zvi Bodie, John Shoven and David Wise. National Bureau of Economic Research. Chicago: The University Of Chicago Press, 1988, pp. 163-188.

Lumsdaine, Robin; Stock, James and Wise, David. "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 David Wise, editor, *Topics In The Economics of Aging*, Chicago: University of Chicago Press, 1992a, pp. 19-57.

_____. "Pension Plan Provisions And Retirement: Men & Women, Medicare, And Models", Cambridge, Mimeo, June, 1992b.

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

Rust, John. "Behavior of Male Workers At The End Of The Life-Cycle: An Empirical Analysis Of States And Controls", in David A Wise, Editor, *Issues In The Economics of Aging*, Chicago: University of Chicago Press, 1990, pp. 317-379.

Stock, James H. and Wise, David A., "The Pension Inducement to Retire: An Option Value Analysis," in David A Wise, Editor, *Issues In The Economics of Aging*, Chicago: University of Chicago Press, 1990a, pp. 205-224.

_____. "Pensions, The Option Value Of Work, And Retirement", *Econometrica*, September, 1990b, 58(5), pp. 1151-1180.

Appendix

Construction of The Opportunity Set

The Wage.

To construct the wage profile, we begin at the last observation of full-time annual wages and salaries. This figure is projected forward on the basis of the tenure, experience, and if appropriate, health coefficients in the fixed effects wage regression. It is also projected backward on the basis of these coefficients until the next observation of full-time wages and salaries is encountered. That observation is used, and is projected backward until the next observation is encountered. This process is repeated for each successive annual wage and salary observation. For years before the beginning of the survey, there may be a time before the start of the job held at the beginning of the survey when we do not observe the sequence of jobs, and hence do not know the tenure levels. In this case, we assume that such years formed a single job.³¹ If the experience levels at the beginning of the survey indicate insufficient experience to fill in the wage profile back to age 25, the wife is assumed to have been out of the labor force in the remaining years.

Construction of the wage profile for husbands follows a similar strategy as for the wives, except that it is the age coefficients that are used to project wages between actual observations of annual wages and salaries. For the years before the beginning of the survey, the wages were projected backward until age 25, the beginning year of the period used in

31. This assumption probably does not matter a great deal, since it has only a minor effect on the total lifetime earnings (y) and no effect on annual compensation (w) in any year used in calculating the retirement values of ε .

both the estimation and simulations.³²

The hourly wage regression is estimated based on usable hourly wage observations in non-self-employed jobs. For each observation, variables are created for education, health status, residence (central city, suburban, rural, and south or non-south), industry (9 categories), and occupation (3 categories). These variables are measured at the time of the survey. In addition, imputed tenure and experience variables are used. For tenure, the NLS lists a unique employer number for each job, so that it is possible to see whether the wife worked with the same employer in previous years. The beginning date of the job is the survey date the wife first worked for that employer, unless the wife worked for that employer in the first survey (in which case the beginning date is calculated from a question in the first survey about the length of time the wife had been with that employer). Once the beginning of the job is established, tenure is calculated as the time from the beginning of the job to the survey in question.³³ A similar approach is taken to construct the experience variable. In the first survey, the wife was asked how many years of full-time work she had engaged in before that

32. There would be a problem with this strategy for calculating the wage profile if either the husband or the wife did not have any valid observations for full-time annual wages and salaries. This does not often happen if the couple met the other selection criteria, but in the relatively few cases where it does, the couple is excluded from the sample. It might have been possible to project the wages from the wage equations, but the estimated R^2 of the wage equations leaves plenty of room for large errors in such an imputation process. Basically, if there are no observations to pin down whether we are talking about a high-wage or low-wage worker, the entire profile may be very poorly located, and in such cases we deemed it advisable to omit the couple from the sample.

33. This method of calculating tenure will produce consistent numbers across different surveys. This eliminates inconsistencies where, for instance, the individual reports that he has worked for the employer 10 years in one survey and, when asked the same question two years later, responds that he has worked 9 years in the later survey.

survey. The experience variable simply increments this value by one or two (depending on the length of time between surveys) if the wife reported working full-time at the time of the survey.

The resulting (log) hourly wage regression for the wives is presented in Table A.1. Many of the variables are significant, and most have the expected sign. The second order terms of experience and tenure indicate that both have the expected concave shape. The effect of experience peaks after only 28 years, but this is counteracted to some extent by tenure peaking after 54 years. The R^2 of this regression is 50%, which is certainly respectable for a wage regression.

Since several observations for each individual are included in the regression, it is likely that the regression contains a fixed effect, meaning that some individuals consistently earn more due to factors not included in the wage regression, and others less. Such a fixed effect may cause the standard errors in the regression to be understated, and if it is correlated with other explanatory variables, it may lead to biased estimates as well. Table A.2 presents estimates which allow for the fixed effects among individuals. This regression drops out education, since any variable which is constant for a particular individual cannot be distinguished from the fixed effect. The results for the health, residence, industry, and occupation variables show widespread reductions from the values found in the previous regression. The linear term in experience is unchanged, but the second order term now suggests that the effect of experience peaks at around 40 years. Both the linear and second order tenure terms are reduced by about a third, but they still indicate a peak in the effect of tenure at a little over 50 years.

For the husbands, a similar strategy is followed, except that there is not as much information for the husband in the survey. In every survey year except 1968, the survey asks about the wage and salary income of the husbands over the previous 12 months (or sometimes the previous calendar year). This information is much the same as was collected from the wives. Also, the section on other household members collected information about the husband's occupation, if he was working. The surveys did not collect regular information about the industry of the husband's employer. Also, they did not regularly ask about the length of the husband's stay with an employer, nor were we able to tell whether the husband had been with the same employer as in the previous survey. This means that it is not generally possible to construct a tenure variable.

As a result, the wage regression for husbands omits experience, tenure, and industry and adds age instead.³⁴ The results are presented in Table A.3. Again many of the variables are significant and of the expected sign. The relative magnitudes of the linear and second-order age terms indicate that the effect of age peaks at around 50. The results of the fixed effects regression are presented in Table A.4. This regression sharply reduces the effects of health and occupation, but not of central city or suburban location. The age coefficients are about 10-20% lower, but they still indicate that the peak effect of age occurs at around 50 years old.

Pensions

We are particularly concerned with pensions on the last job before retirement, since these pensions contain incentives relevant for the retirement decision. We are in relatively

34. We could have used an experience variable rather than age, but if experience is defined as (age - education - 5), the two variables are picking up the same effect, given that education is also in the regression.

good shape if the individual worked until at least 1982. In that case, we can tell reasonably well whether the individual retired under a pension. The years of service in the job is available in most instances. In many cases we can tell the age they were first eligible for a pension, and we can get at least one figure as to the amount of the expected pension. However, the information on the normal retirement age appears in general to be either missing or unreasonable, and very few individuals provided enough information to be able to infer how much pensions are reduced for early retirement.

Faced with this situation, we proceed as follows. For individuals who were covered on their last jobs before retirement, we use the last figure they reported as the early retirement age. If no figures are reported, we use age 55.³⁵ If the individual reported that he or she was receiving benefits after the last year of full-time work but before the reported early retirement age, the early retirement age was adjusted downward to be consistent with the reported benefits. The normal retirement age, about which we have relatively little information, is taken to be 62 except if the individual retired around age 65, in which case the normal retirement age is taken to be 65. For women, the tenure is the tenure calculated for the wage profiles, and for men, if no tenure is given in the pension section, tenure is measured from age 25.

For individuals who retired before 1982, we have much the same information about the

35. Age 55 is the most common early retirement age among those who do report an early retirement age. The distribution of cases is as follows:

Total pensions--wives:	399
Of which early retirement date imputed:	48
Total pensions--husbands:	319
Of which early retirement date imputed:	133

pension, but the data do not indicate which job is associated with the pension. For women who retire before 1982 and report that they had a pension, the pension is considered to have come from the latest job if that job lasted for longer than five years.³⁶ For the husbands, we have little choice but to assume that if they retired before 1982 and had a pension, the pension came from the last full-time job. If it appears that the pension did come from the last job, the calculations of the value of accruals proceed in the same manner as outlined above.

If the pension is from a job previous to the last job, the incentives from the pension do not affect the retirement decision directly. The only effect of these pensions on retirement is through a wealth effect. The wealth effect can be calculated from the size of the pension and the length of time the individual has been collecting it, which we know for most such individuals. In cases of missing information, the pension amounts and/or the period of collection are constructed along the lines outlined for individuals who receive pensions on their last job.

36. The numbers for each of these cases are as follows:

Pensions before 1982	
Imputed to last full time job > 5 years:	22
Imputed to previous job on basis of years of service:	10
Imputed to longest previous job:	3

Note that the pension is considered to come from the latest job if that job lasted longer than five years. If the latest job lasted less than five years, previous jobs are searched for the one whose tenure came closest to matching the tenure given in the pension questions. If no tenure was given in the pension questions, the pension is taken to come from the longest previous job.

Table A.1
Wage Regression For Wives

Dependent Variable: Log of Hourly Wage

Mean 1.390

Standard Deviation .444

	Coefficient	t Statistic	Mean
Constant	.476	11.84	
Education			
Years of Education	.043	11.14	11.8
High School Graduate	-.007	-0.50	.694
Years of College	-.010	-1.41	.755
College Graduate	.057	2.13	.133
Health Problem	-.082	-7.73	.139
Residence			
Central City SMSA	.152	16.31	.327
Suburban SMSA	.158	17.25	.347
South	-.126	-15.70	.411
Years of Experience	.015	8.61	20.5
Square of Experience / 100	-.027	-6.93	505
Years of Tenure	.013	8.24	9.9
Square of Tenure / 100	-.012	-2.36	159
Industry			
Agriculture, Forestry & Fishing	-.233	-5.17	.007
Mining	.212	3.06	.003
Construction	.036	0.89	.009
Transport & Communication	.146	6.80	.036
Wholesale & Retail Trade	-.230	-17.87	.146
Finance, Insurance & Real Estate	-.061	-3.59	.070
Services	-.132	-12.61	.439
Public Administration	.074	4.28	.065
Occupation			
Management & Professional	.384	29.93	.273
White Collar	.173	16.22	.379

Number of Observations = 7317

R Square = 0.5036

Standard Error of Estimate = .313

Table A.2
Fixed Effects Wage Regression For Wives*

Dependent Variable: Log of Hourly Wage
Mean 1.390
Standard Deviation .444

	Coefficient	t Statistic	Mean
Health Problem	-.030	-3.58	.139
Residence			
Central City SMSA	.019	0.79	.327
Suburban SMSA	.043	1.88	.347
South	-.073	-2.22	.411
Years of Experience	.015	10.51	20.5
Square of Experience / 100	-.018	-5.88	505
Years of Tenure	.009	6.88	9.9
Square of Tenure / 100	-.008	-2.01	159
Industry			
Agriculture, Forestry & Fishing	-.123	-2.87	.007
Mining	.032	0.40	.003
Construction	-.039	-1.13	.009
Transport & Communication	-.018	-0.65	.036
Wholesale & Retail Trade	-.131	-8.47	.146
Finance, Insurance & Real Estate Services	-.042	-1.77	.077
Public Administration	-.119	-8.27	.439
Public Administration	-.048	-2.31	.065
Occupation			
Management & Professional	.061	4.73	.273
White Collar	.004	0.31	.372

Number of Observations = 7317

R Square = 0.8265

Standard Error of Estimate = .185

* Regression includes dummy variables for repeated observations on each individual.

Table A.3
Wage Regression for Husbands

Dependent Variable: Log of Hourly Wage
 Mean 1.809
 Standard Deviation .498

	Coefficient	t Statistic	Mean
Constant	.021	0.10	
Education			
Years of Education	.049	9.50	11.8
High School Graduate	.048	2.02	.656
Years of College	-.016	-1.55	1.04
College Graduate	-.028	-0.71	.178
Health Problem	-.113	-5.55	.104
Residence			
Central City SMSA	.184	11.42	.257
Suburban SMSA	.209	14.49	.407
South	-.168	-12.59	.370
Age	.047	5.82	49.2
Square of Age / 100	-.051	-6.09	2483
Occupation			
Management & Professional	.211	12.44	.325
White Collar	.073	3.61	.124

Number of Observations = 4575
 R Square = 0.3120
 Standard Error of Estimate = .413

Table A.4
Fixed Effects Wage Regression for Husbands*

Dependent Variable: Log of Hourly Wage
 Mean 1.809
 Standard Deviation .498

	Coefficient	t Statistic	Mean
Health Problem	-.054	-3.18	.104
Residence			
Central City SMSA	.220	6.31	.257
Suburban SMSA	.221	7.00	.407
South	-.120	-2.89	.370
Age	.042	7.33	49.2
Square of Age / 100	-.042	-7.07	2483
Occupation			
Management & Professional	.004	.26	.325
White Collar	-.022	-1.07	.124

Number of Observations = 4575
 R Square = 0.7224
 Standard Error of Estimate = .2622

* Regression includes dummy variables for repeated observations on each individual.

Retirement In A Family Context: A Structural Model For Husbands and Wives

Alan L. Gustman
Dartmouth College and National Bureau of Economic Research
and

Thomas L. Steinmeier
Texas Tech University

Executive Summary

I. Introduction.

This paper specifies and estimates a structural model of the retirement decisions of husbands and wives. The feature of the data that is the focus of this paper is the tendency of husbands and wives to retire together. An econometric approach is developed for estimating preferences of both spouses jointly and is implemented using data from the National Longitudinal Survey of Mature Women (NLS), a survey that provides the most recent data available for a joint retirement study. Alternative specifications of joint decision making are tested, and the importance of various sources of interdependence in decision making are investigated.

In the course of specifying and estimating the retirement model for couples, we address the following specific questions: How does the retirement behavior of each spouse compare to that of the other? To what extent is the wife's retirement decision influenced by the husband's, the husband's decision influenced by the wife's, and to what extent are their unmeasured tastes correlated? How does interdependence in the opportunity set and in preferences affect the coordination of retirement by the spouses? What are the effects on

model parameters of entirely ignoring interactions in preferences? What are the effects of using an inappropriate estimating technique, one which treats each spouse's retirement decision as exogenous, and therefore does not recognize the dependence of each spouse's retirement decision on the other's?

Although wives are younger, and women retire at younger ages, we find that husbands and wives coordinate their retirement. One reason is a coincidence of tastes for leisure. More importantly, each spouse, and perhaps husbands in particular, find leisure in retirement to be more valuable once their spouse has retired. Holding jobs with similar incentives is not the reason husbands and wives coordinate retirement dates.

The present analysis joins together two branches of the retirement literature. In one, structural retirement models, that is models that separate the roles of market opportunities and individual preferences, are estimated from data on individuals, usually men, while ignoring the behavior of their spouse. In the other, retirement decisions of husbands and wives, and the tendency of their retirement dates to cluster, are analyzed in the context of equations in which retirement outcomes for each spouse are related to a set of explanatory variables, but the influences of preferences and tastes are not separately analyzed.

In a world where both the husband and wife are more and more likely to be working until the retirement years, this kind of a model will be increasingly necessary to assess how pension, social security and other retirement related policies affect retirement outcomes, including the question of whether policy measures which affect the retirement decision of one family member can indirectly influence the retirement of the remaining spouse.

A second purpose of this paper is to provide a structural retirement analysis using much

more recent data for a nationally representative panel study than has been used in the past. The women in the National Longitudinal Survey of Mature Women (NLS) were born between 1923 and 1937. They have been surveyed since 1967. In the last year for which the survey is available, 1989, they were 52 to 66 years old, about as young as feasible for a retirement study.

The paper also addresses a number of econometric and behavioral issues. One issue is selection bias. In an old panel data set, all couples have retired, so there is complete information on retirement dates and on the characteristics of all individuals included in the sample. In a more recent panel such as the National Longitudinal Survey of Mature Women, many of the couples are too young for both to have retired. Accordingly, in the process of the analysis, we deal explicitly with truncation of continuing employment spells, that is, the fact that we do not observe the ultimate retirement date for one or both members of the couple.

Section II of the paper presents evidence from the NLS that in fact, in the raw data, there is a noticeable tendency among couples who have both retired to retire together. A family labor supply model is developed in Section III. Section IV details the data preparation and presents alternative estimates of the model. Section V presents some simulations based upon the estimated model, concentrating mainly on the extent to which the husband's and wife's retirement decisions are coordinated. A final section presents some concluding thoughts.

II. Evidence of Joint Retirement.

There is not much point in investigating the cause of couples tending to retire together if the phenomenon is not evident empirically. Section II presents data that verifies that the

phenomenon does exist in the NLS older women's data, and documents the sample which is used in this section and in the remainder of the project.

In this study retirement is defined as no longer working full-time. The sample we use from the National Longitudinal Survey of Mature Women (NLS) is restricted in several ways in order to meet the objectives of this paper. First, the sample is restricted to women who remained in the survey. Second, since the project focuses on the joint retirement decisions of husbands and wives, the sample is restricted to women who were married to the same man in each year of the sample, and to those couples both of whose members survived. Third, the sample is restricted to women for whom the idea of retirement is a meaningful concept because they have spent a prolonged period as a full-time member of the labor force. Specifically, to be included in the sample, the woman must have been working full-time (at least 25 hours per week) for at least three consecutive surveys after the age of 50.

Data on retirement date differences between husbands and wives is found to exhibit a distinct concentration of retirement when the husband and wife retire simultaneously. Given the way the retirement dates were constructed, what this really means is that there is a concentration of couples reporting the last date of full-time work in the same survey year. The concentration of retirements at dates so close to one another suggests that couples do tend to retire at about the same time, and that this phenomenon is evident in the data to be used in this study.

III. A Model of Family Labor Supply and Retirement.

Section III develops a model of family labor supply and retirement, and finds an approach to estimating the model empirically. The model begins with a fairly standard

preference function in which utility depends on lifetime consumption of the family together and the labor supply (leisure) of the husband and wife individually. Utility is gained from consumption and leisure at each moment in time, and is also influenced, among other things, by age and health. As each spouse becomes older, eventually the value of retirement outweighs the value of the wages from working, and the individual retires.

There are three ways in which this is a family labor supply model. First, the consumption in this utility function is not the consumption from the wife's or husband's own earnings, but the family consumption financed by the earnings of both husband and wife. Second, the value the wife places on leisure is influenced by the retirement status of the husband, in that she enjoys her retirement more when he is home. Similarly, the value the husband places on leisure is specified to be influenced by the retirement status of the wife. Finally, the retirement preferences of the husband and wife may be correlated because people who share the same tastes are more likely to marry.

Both husband and wife choose a path of consumption and work and ultimately retirement, consistent with their preferences, but subject to the constraint that lifetime family consumption cannot exceed family income. In this budget constraint, the value of consumption over the lifetime cannot exceed the value of wages, pensions and social security.

In any given year, the wife would want to work if the value of the wages exceeds the value of the leisure foregone, and otherwise she would want to retire, and the same is true for the husband. A methodology is developed for determining the range of values placed on leisure which would lead each individual in the sample to decide to retire, given the labor market opportunities they faced. The estimating routine searches for values to place on

parameters representing each individual's valuation of leisure, searching for parameter values that would maximize the likelihood of the individual retiring when he or she does given those opportunities. That is, in light of the opportunities available to each husband and wife, and in view of the retirement behavior they exhibited, the estimating routine asks: what parameters of their preferences are most likely to account for the behavior that is seen in the sample?

The methodology is designed to allow the preferences of the husband and wife to interact on a number of dimensions, and incorporates the fact that the value placed on leisure by one spouse is likely to be influenced by the retirement status of the other. To implement the estimation scheme, it is necessary to adjust the estimation procedure to include the information on incomplete employment spells for those who have not yet retired. Dropping the many couples with one or both partners still working at the end of the observation period would subject the estimates to censoring. Accordingly, the methodology is adjusted to incorporate the information on retirement behavior for all of those couples in which only the husband, or only the wife has retired, or indeed if neither has retired.

IV. Estimation of the Model.

The data requirements for estimating the model fall into three categories. First, the dependent variable (retirement) must be measured. The other data requirements are for measures of the elements which affect the value the individual places on leisure. These are: age, spouse retirement, health, and vintage. The age variable is the individual's reported age in 1967, incremented by 1 for each year. Vintage is simply 1967 minus the age in 1967, and spouse retirement in a given year is whether the spouse is retired at the time. For the health variable, we examine in each year the answers to questions about whether health prevented

the individual from working or whether health limited the amount or kind of work that could be done.

The major difficulty in preparing the data set for estimation is to calculate compensation streams. The compensation streams consist of three components: wages, pensions, and social security.

For all the survey years except 1968, the survey asks about annual income from wages and salaries, as well as enough information to construct an hourly wage rate. For those years in which the wife is working full-time at the time of the survey and for which usable annual wage and salary information is available, the annual wage and salary information is used. For other years, annual wages must be imputed. The imputation process uses the tenure, experience, and health coefficients from an hourly wage regression with fixed individual effects.

The next element of compensation is the pension profile. Information on pensions comes mainly from questions that were asked in 1982, 1986, and 1989. In the first two of these years, the survey asked about whether the respondent was eligible for a pension from the current or past job. If the individual was eligible for a pension from either the current or past job, the survey inquired about the years of service in that job. If the pension was for the current job, the survey also asked about the age of initial eligibility for benefits. The 1989 survey expanded the questions so that they include the early and normal retirement ages on the present job, the early and normal retirement ages for pensions on previous jobs, and the amounts that the individual would be eligible to receive on those dates. In addition, the survey asked about the actual amount of benefits received if the respondent reported that

benefits were currently being received. The pension is calculated on the basis of this information assuming that the individual faces a defined benefit pension plan.

The present value calculations for social security are similar to those for the pensions. Working another year can change the present value of social security benefits for three reasons. First, higher real wages during prime working years are likely to displace earlier wages from the average earnings calculations, increasing benefits. This effect is enhanced for those over sixty years of age because under the calculation of average indexed monthly earning (AIME), the basis for calculating monthly benefits, wages are indexed only to age 60. Second, if the individual is over the early retirement age, working another year will reduce the early retirement reduction factor or increase the delayed retirement credit. Finally, if the individual is over the early retirement age, working another year may cause benefits to be lost to the earnings test. In general, the effects of the social security system are to cause reductions in compensation primarily at the normal retirement age, since the delayed retirement credit is not sufficient to offset the loss of benefits from continuing working.

The overall compensation paths for the wife and the husband are calculated by adding to the wage the accruals for social security and, if appropriate, the pension. These are then combined with the retirement ages of both spouses and the values of the variables reflecting age, vintage, spouse retirement status and health, to provide the data set for estimation. The final data set has 564 observations of couples who satisfied all of the inclusion criteria and who had at least one wage and salary figure on which to base the compensation paths.

The estimation begins with a simple model in which each spouse is assumed to retire independently. The estimates indicate that having a health problem has the same effect on

retirement as being not quite two years older, and the effect of vintage is very small. The standard deviation of the retirement taste term amounts to over five times the coefficient of the age term. It is fairly clear that the taste term plays a very large role in accounting for the large range of ages at which various individuals retire.

For the husband, the value of the health coefficient also amounts to a little over three times the magnitude of the coefficient of the age term, and the standard deviation in tastes for retirement are a little over five times the size of the coefficient of the age term. The age term itself is somewhat larger than it is for the wives.

When interdependence in retirement is examined, but not in the context of a fully integrated model for both the husband and wife, there is some suggestion that the husband's retirement does not affect the retirement preferences of the wife, but the wife's retirement has a notable effect on the retirement preferences of the husband. One possible rationale for this result is that some husbands do not wish to find themselves with more time to take care of the household while the wife is still devoting a substantial part of her time to market work. But stronger confirmation of this result and investigation of further implications would be required before any conclusion is drawn.

An alternative explanation for couples retiring at around the same time is examined. If men and women who prefer to retire early tend to marry one other, and men and women who prefer to retire later (and consume more over their lifetimes) tend to marry each other, the tastes for leisure of husbands and wives will be correlated with each other. When this explanation is examined in isolation, the estimates indicate that the correlation in preferences is highly significant.

Given that each spouse's retirement decision is influenced by the retirement status of the other, it becomes important to recognize in the estimation the joint determination of retirement, and not to assume that the retirement status of one spouse is an exogenous determinant of the retirement status of the other. Moreover, the model should be specified to incorporate all influences affecting the decisions of the husband and wife.

When these two competing explanations are included in the same equation, the correlation in preferences is only about two-thirds as large, but it is still highly significant. The spouse retired coefficients now both have the proper sign, but neither one alone is significant. However, the combined effect of the two of them is marginally significant. Because the relative effects of these different dimensions of preferences cannot be inferred simply by comparing the relative size of the parameters, the issue is investigated by simulation.

V. Simulations Of The Effects Of Husband Wife Retirement Interactions.

The simulated retirement outcomes track the retirement outcomes in the data fairly well. The model also does a good job of simulating the coincidence of retirement dates among couples. When the full model is estimated, there is a sharp peak at zero difference in retirement dates, with over eleven percent of the couples retiring at the same time, whereas for other age differences, the probability ranges around 4 or 5 percent.

It is clear that the source of interdependence in retirement is in the coordination of preferences among spouses and is not due to the correlation of incentives between the opportunity sets facing husbands and wives. This result is seen by comparing the probability of joint retirement under the full model, where all sources of retirement are included, with the

probability of joint retirement when all sources of interdependence in the utility function are suppressed, but the opportunity sets remain unchanged.

Investigating the source of the interdependence within the utility function, the conclusion is that joint retirement is due almost entirely to the presence of spouse retirement status in the utility function. It appears that considering the spouse's retirement implies that couples have a considerably greater probability of retiring together. The modest estimated correlation between the preferences of husbands and wives produces a modest increase in the tendency of couples to retire around the same time.

In sum, from the simulations it is clear that most of the coincidence in retirement dates at zero years difference in retirement dates is due to the spouse retirement variables. Dropping the simple correlation between the error terms diminishes this peak only slightly, but dropping the retirement status of one spouse in the leisure valuation of the other almost eliminates it.

A final simulation shows that in cases where the wife retires first, about 42 percent of the husbands would delay retirement by another year if the spouse retirement effect were inoperative. For the affected group, this translates into an increase in the average retirement age of about 5 months. While this is not an overwhelming effect, it is in the same order of magnitude as the average effects which have been reported in the literature for pensions and social security.

VI. Summary and Conclusions

This paper specifies and estimates a structural retirement model in which retirement decisions by husbands and wives are jointly determined. The estimated model generates

behavior that brings the retirement dates of husbands and wives closer than their age differences, with the probability of retirement peaking at the same time for over eleven percent of husbands and wives.

Although the estimation procedure is unable to separate out the spouse retirement coefficients, both of the competing explanations (the spouse retirement explanation and the preference correlation explanation) are significant in the data. Simulation indicates the relative importance of each of these sources of interdependence in increasing coordination of the retirement outcomes of spouses. In contrast, there is no evidence that the coordination in retirement outcomes results from couples selecting jobs with pensions offering compatible incentives. Accordingly, in contrast to the support found for the hypothesis that coordination of retirement outcomes is due to interdependence in preferences, support is not found for the hypothesis that coordination in retirement of spouses results from coordination in the incentives created by the spouse's opportunity sets.

Moreover, the evidence suggests that ignoring the statistically significant effects of one spouse's retirement behavior on that of the other, by treating retirement outcomes of one spouse as an exogenous determinant of the other's, does affect the estimates of other key parameters in the model.

There is some suggestion in the data that the wife's retirement decision is not strongly influenced by the husband's, but the husband's decision is more strongly influenced by the wife's. Although such a result is consistent with an interpretation that the husband would rather not face the house and related work load alone, this result is a preliminary suggestion that requires further support.

Finally, because the National Longitudinal Survey of Mature Women provides the most recent, nationally representative longitudinal sample available for a study of retirement behavior, the results reported in this study take account of the labor market trends shaping the retirement behavior of women, at least to the extent that these trends have influenced the experience and behavior of the youngest generations of women yet to enter retirement. Although some members of the cohorts included in the NLS have not yet retired, there is likely to be only a small effect on key parameter estimates. The estimating technique has explicitly accounted for incomplete employment spells.