INSTITUT NATIONAL DE LA STATISTIQUE ET DES ETUDES ECONOMIQUES
Série des Documents de Travail du CREST
(Centre de Recherche en Economie et Statistique)

n° 2003-41

Health Insurance and Retirement of Married Couples

D. M. BLAU[1]
D. B. GILLESKIES[2]

Thanks to the National Institute on Aging for funding, and to John Bound, Charlie Brown, Brigitte Madrian, Jon Skinner, Ken Wolpin, three referees, Costas Meghir and seminar participants for helpful comments.
Comments welcome at david_blau@unc.edu or donna_gilleskie@unc.edu
[1] University of North Carolina at Chapel Hill, USA and CREST.
[2] University of North Carolina at Chapel Hill, USA.

Health Insurance and Retirement of Married Couples

David M. Blau and Donna B. Gilleskie

University of North Carolina at Chapel Hill

September 2003

Thanks to the National Institute on Aging for funding, and to John Bound, Charlie Brown, Brigitte Madrian, Jon Skinner, Ken Wolpin, three referees, Costas Meghir, and seminar participants for helpful comments. Comments welcome at david_blau@unc.edu or donna_gilleskie@unc.edu.

Health Insurance and Retirement of Married Couples

Abstract

Health insurance is a potentially important determinant of employment behavior at older ages in the United States. Most health insurance in the U.S. is provided by employers until eligibility for Medicare begins at age 65. A key link between health insurance and employment is that some employer health insurance plans continue to provide coverage for retirees while others do not. Retiring before age 65 exposes individuals and their spouses who lack retiree health insurance coverage to the risk of potentially catastrophic medical expenditure. We build a dynamic behavioral model of the employment decisions of older couples that includes risky medical expenditure and health insurance. Estimates of the model allow us to determine the empirical importance of health insurance constraints in explaining the timing of retirement of couples. Estimates using data from the Health and Retirement Survey indicate that couples have a moderate degree of aversion to the risk of large medical expenditure. The risk-reducing feature of health insurance can fully account for the relatively modest observed association between retiree health insurance and employment for married men, but can account for only about one third of the large observed association for married women.

JEL: J26

1. Introduction

Most health insurance in the United States is provided by employers until eligibility for public health insurance for the elderly (Medicare) begins at age 65. Many employer health insurance plans provide coverage for retired workers, while others do not. The absence of retiree health insurance (RHI) coverage creates a link between employment decisions and health insurance coverage for older workers that may affect the incentive to retire at specific ages before 65. This link does not exist for workers who are eligible for RHI coverage from their employers, or who have coverage from other sources. The association between health insurance and employment of older individuals has been investigated in several studies, and has been found to be quite strong. For example, Blau and Gilleskie (2001a) report that older men with employer-provided health insurance are about two percentage points more likely to retire in a given year if they have retiree coverage than if they lack such coverage, controlling for many other factors.[1] This is an important policy issue because reform proposals that would make health insurance coverage independent of employment status could increase the already-high rate of retirement before age 65, thus worsening the financial condition of Social Security and Medicare.

In this study, we analyze the effect of health insurance on retirement in a family context. In section 2 we describe the association between the labor force behavior of older couples and their health insurance coverage, using data from the Health and Retirement Study (HRS). These data show that couples who face employment incentives arising from shared health insurance appear to respond to those incentives. A man whose health insurance coverage depends on him

[1]See also Gruber and Madrian (1995), Madrian (1994), and Rust and Phelan (1997).

remaining employed is substantially more likely to be employed than is a man whose health insurance coverage is independent of the individual's employment status. For women, the association between RHI and employment is even larger than for men.

In section 3 we specify a forward-looking dynamic model of the employment decisions of older couples that can be solved numerically and estimated empirically. The goal of the model is to determine the extent to which aversion to uninsured medical expenditure risk can explain the employment patterns we observe in the data. The model places no restrictions on the degree of risk aversion: this key aspect of preferences is identified by variation in medical expenditure risk and health insurance constraints and the behavior of couples in response to these constraints. Estimates of the model therefore allow us to determine the empirical importance of health insurance as a determinant of retirement in a context in which the estimates have clear economic interpretations.[2] Other recent efforts to estimate structural models of health insurance and retirement include Gustman and Steinmeier (1994), Lumsdaine, Stock, and Wise (1994), and Rust and Phelan (1997). Our analysis is most similar to that of Rust and Phelan, whose results strongly suggest that health insurance incentives affected employment behavior of older men in the 1970s. Our analysis builds on their innovative approach, and extends it in the following ways.

First, health insurance is often a family affair, so we model the joint employment decisions of couples. Second, we include in our sample individuals who are covered by employer-provided pension plans. This makes our sample more representative than the sample

---

[2]See Berkovec and Stern (1991), Blau and Gilleskie (2003), Gustman and Steinmeier (1986), Rust and Phelan (1997), Stock and Wise (1990), and van der Klaauw and Wolpin (2002) for other dynamic models of retirement.

analyzed by Rust and Phelan, who were forced to exclude men with private pensions as a result of lack of information on the provisions of the pension plans. The HRS contains the detailed information on pension plan parameters needed in order to accurately measure the benefit available as a function of alternative ages of exit from a firm. Third, Rust and Phelan modeled medical expenditure as a random draw from a health-insurance-specific distribution of out-of-pocket expenditure. This approach worked well in their case, but in our data it fails to adequately capture empirically a crucial feature of our model: the medical expenditure distribution that an employed individual without access to retiree health insurance *would* face were he to leave employment before age 65. Thus, we model out-of-pocket medical expenditure as determined by health insurance plan rules, such as the premium, deductible, and coinsurance rate, applied to a random draw from a distribution of *total* medical expenditure. This approach provides a realistic and tractable link between employment decisions and medical expenditure risk. Finally, we analyze data from the 1990s, which provide a more up-to-date basis for policy analysis than the data from the 1970s used by Rust and Phelan. Section 4 describes the HRS data that we use to estimate the model.

Section 5 presents estimates of the model and simulations of behavior generated from the estimated model. The estimated model fits the observed employment choices well in most dimensions. The model is able to fully account for the observed association between RHI and employment for married men. Simulations imply that policy reforms that break the tight link between employment and health insurance will have modest impacts on employment decisions of men at older ages. However, the model significantly underpredicts the observed association between RHI and employment for married women. Health insurance enters the model only through the budget constraint, so this suggests that the risk-reducing feature of health insurance may not be the main driving factor behind the observed differences in employment choices of women by health insurance. Women may value health insurance for other reasons not captured in our model. Section 6 offers concluding thoughts.

2. Descriptive Overview

The HRS sampled men and women aged 51-61 in 1992 and their spouses, and has interviewed them every two years since 1992. In this section, we describe the relationships among employment status, defined by whether an individual is employed at the survey date, employment transitions between survey dates, and health insurance coverage, using data from the first four waves of the HRS (1992, 1994, 1996, and 1998). Table I shows differences in employment status and employment transitions between individuals who are covered by employer-provided health insurance with and without retiree coverage, predicted from logit models.[3] Panel A shows that husbands with coverage from their own employer are 12.9 percentage points more likely to be employed if they lack RHI. This difference falls to 4.4 points with controls for age, earnings, health, nonwage income, future pension benefits, and future Social Security benefits. The difference for wives, shown in panel B, is 32.3 percentage points (17.1 with controls). We and others have previously documented such patterns for men, and

[3]The logit models are estimated separately for husbands and wives, using the sample of 1,752 couples described in detail below. There are up to four observations per individual in the employment status models, and up to three observations per individual in the employment transition models. Observations with all possible sources of health insurance are included, but the discussion focuses only on individuals with health insurance coverage provided by the individual's employer.

interpreted them as evidence that absence of retiree health insurance is a strong deterrent to labor force exit before the age of 65.[4] These results suggest an important role for health insurance in early retirement decisions for women, and a modest role for men.

Panels C and D show that the exit rate from employment is higher with RHI than without it. But the association is not significantly different from zero, and for men the magnitude is a very small 0.4 percentage points per two years. The extensive controls for future Social Security and pension benefits in the models on which these estimates are based may explain why the observed association between RHI and employment for men is smaller than in other analyses.[5] This highlights the importance of modeling the budget constraint as accurately as possible, a task to which we devote considerable attention below.

3. The Model and Estimation

3.1 Overview

The key features of the model are that health insurance can help couples smooth the marginal utility of consumption across states of the world with different levels of medical

---

[4]The results are very similar if the sample is limited to individuals under the age of 63 or 65. See Blau and Gilleskie (2001a) and Madrian (1994) for similar findings. If the spouse is covered by the individual's health insurance, the effect of retiree coverage on the individual's probability of non-employment is only slightly stronger than if the spouse is not covered by the individual's insurance. This suggests that most of the effect of retiree coverage operates through the association between the individual's own coverage and his employment status.

[5]Madrian and Beaulieu (1998) compare the retirement decisions of married men who have a spouse 65 or older with those of men with younger spouses, in order to estimate the impact of Medicare eligibility on retirement. Using data from the 1980 and 1990 Censuses, they find that men aged 55-69 who worked in the previous calendar year are more likely to be retired at the census date if the spouse is Medicare-eligible than if she is not Medicare-eligible. This is consistent with our findings reported in Table I, but is not directly comparable.

expenditure, and for some couples health insurance is linked to employment. There is no saving in the model, so a shock that results in large uninsured medical expenditure forces a reduction in consumption in the period in which the expenditure is incurred. Health insurance helps smooth consumption by reducing consumption in good states of the world (those with low medical expenditure) via the insurance premium, and increasing consumption in bad states via reimbursement of medical expenses. Because health insurance coverage is valuable to risk averse consumers, a worker whose own or spouse's insurance coverage is tied to his employment status may make different employment decisions than an otherwise similar worker whose insurance coverage is independent of his employment status. If couples are highly risk-averse, then these differences could be substantial.

Three key assumptions are imposed for reasons of computational feasibility: there is no saving, health insurance coverage is not subject to choice (except via the choice of whether to continue in a job with employer-provided insurance without retiree coverage), and medical expenditure is not a choice variable. By forcing agents to satisfy a series of period-specific budget constraints, and by restricting choice over health insurance and medical expenditure, the estimated model will provide an upper bound estimate of the effect of health insurance on employment decisions. If couples could self-insure by saving in anticipation of the possibility of large medical expenses, then health insurance would be less important and the incentive to be employed in order to retain health insurance would be weaker. If individuals could acquire health insurance coverage by changing jobs or purchasing private non-group insurance, then the employment decisions of individuals with and without retiree coverage might not be very different. And if agents could choose to forego medical treatment, then the absence of retiree coverage might not be a deterrent to retirement before 65. Thus by limiting other mechanisms for smoothing consumption to account for medical expenditure risk, the model forces individuals who wish to avoid exposure to such risk and who lack retiree health benefits to remain employed until they become eligible for Medicare. If the estimates of this model imply little impact of health insurance on employment, then we would expect that relaxing these assumptions would also yield small impacts. On the other hand, if the estimates imply a large impact of health insurance on employment, then it would be important to determine whether the findings are robust to allowing savings, health insurance, and medical care choices.[6]

We specify a discrete state, discrete time model of the employment choices of married couples. At the beginning of a period, a couple learns the realizations of all stochastic processes for the period, *except* the process governing medical expenditure. Given these realizations and the values of deterministic state variables such as work experience and job tenure, the husband and wife make their employment choices for the period and these choices remain fixed for the duration of the period. Total medical expenditure for the period is then realized, and together with the already-committed employment choices and the resulting health insurance coverage, this determines out-of-pocket medical expenditure and consumption net of out-of-pocket medical expenditure. State variables are updated at the end of the period, and the process repeats until the terminal date, T^*. The length of a period is two years, corresponding to the time between

---

[6]Starr-McCluer (1996) presents evidence suggesting that precautionary saving by individuals without health insurance is rare. Blau and Gilleskie (2003) model medical care decisions jointly with employment behavior of older men. Also, our upper-bound argument applies only to the effect of health insurance as it operates through the budget constraint. If agents derive utility directly from health insurance, then our estimates are not necessarily upper bounds.

interviews in the HRS.

3.2. Employment and Health Insurance

The employment status of member $a = m$ (male), f (female), of a couple at the end of period $t\text{-}1$ is e_{at-1}, where e_{at-1}=0 denotes not employed, and e_{at-1}=1 denotes employed. With probability δ an individual who was employed at the end of period t-1 is laid off at the beginning of period t, an event denoted by L_{at}=1, with L_{at}=0 denoting no layoff. With probability one, individuals receive a job offer from a new employer at the beginning of each period.[7] Jobs are assumed to be identical except for the pension insurance coverage they offer. Let j_{at} represent the job choice of spouse a in period t. If an individual was employed in period t-1 and not laid off at the beginning of period t, she chooses among non-employment (j_{at} = 1), the new job offer (j_{at} = 2), and the old job (j_{at} = 3). If the individual was laid off at the beginning of t or was not employed during t-1, she chooses between non-employment and the new job offer. Allowing for the possibility of changing jobs is important because pension income is tied to specific jobs rather than employment status *per se*. Allowing re-entry to employment is important because it is quite common in the age range of our sample (on average, 13 percent of individuals in our HRS sample who were not employed at wave t are employed two years later at wave t+1).

Until the age of Medicare eligibility, we assume that an individual can be covered by up to two of the following types of health insurance during a given period: (0) none; (1) own-employer, with retiree coverage; (2) own-employer, without retiree coverage; (3) spouse-employer, with coverage available to the individual after the spouse retires; (4) spouse-employer,

[7]We assume that a new job is always available because identification of job offer arrival rates would be difficult given the assumption (discussed below) that earnings are not stochastic.

without retiree coverage; (5) private (non-employer); and (6) Medicare. Medicare is available to individuals under the age of 65 only if they receive Social Security Disability Insurance (SSDI).[8] Upon reaching age 65, individuals are assumed to enroll in Medicare.[9] We assume that if an individual is covered by his employer's health insurance plan without retiree coverage, he expects to become uninsured if he leaves employment before he turns 65.[10] He expects to remain uninsured until he becomes employed again or reaches age 65 and receives Medicare coverage. Similarly, if an individual is covered by his spouse's employer's plan without retiree coverage, he expects to become uninsured when his spouse leaves employment, and to remain uninsured until *she* gets a new job or he turns 65. Health insurance coverage of individuals covered by employer plans with retiree insurance or by private plans is unaffected by employment decisions. We allow for two sources of health insurance coverage because this is fairly common in the data.

The key implication of these assumptions is that health insurance coverage is fixed before age 65. Individuals who lack employer-provided health insurance cannot acquire such

[8]Medicare is available only after two years on SSDI, but this does not affect our analysis because we take SSDI enrollment as given and do not model applications to SSDI. Following SSDI policy, we assume that individuals younger than 65 who receive Medicare through SSDI would lose their Medicare coverage and SSDI benefit if they re-entered employment.

[9]In the model, it is assumed that employer and private health insurance convert to Medigap coverage at age 65, as required by Medicare rules. Medigap insurance is secondary coverage that pays for expenses not covered by Medicare. Medicare rules require that employer and private plans must be the primary payer even after age 65 if the individual is employed. If the individual is retired, then Medicare is the primary payer.

[10]The Consolidated Omnibus Budget Reconciliation Act of 1986 (COBRA) requires firms with 20 or more employees to allow workers who leave the firm to continue their health insurance coverage for up to 18 months after the date of separation. Thus a worker who retires at age 63.5 can use COBRA coverage as a bridge to Medicare coverage at age 65. The firm can require the former worker to pay the full cost of the coverage. We ignore COBRA coverage because the time period in our model is two years.

coverage by changing jobs, and individuals who have employer coverage without RHI cannot acquire RHI by changing jobs. Individuals who do not have private (non-employer) coverage cannot purchase such coverage. The only agents who can change their health insurance coverage are those who are covered by employer-provided insurance without RHI: they lose coverage if they leave employment, and regain coverage if they reenter employment. It is the behavior of this group relative to the other groups that drives the model, as discussed below.

3.3. Health and Medical Expenditure

Health status of spouse *a*, h_{at}, is a discrete variable that can take on values 0, 1, 2, where 0 denotes good health, 1 denotes bad health, and 2 denotes death. The probability of observing spouse *a* in health status h' during period t+1 given that he or she was in health status h during period t is

$$\pi_{at}^{hh'} = \Pr(h_{at+1}=h' | h_{at}=h) = \frac{\exp(\gamma_{0hh'}^{a} + \gamma_{1hh'}^{a} A_{at})}{\sum_{h''=0}^{2} \exp(\gamma_{0hh''} + \gamma_{1hh''}^{a} A_{at})} \tag{1}$$

where A is age, and the γ's are parameters. These first-order Markov logit processes depend only on age, gender, and previous health.

Total medical expenditure for each spouse is modeled as a random draw from a known distribution that is realized *after* the individual has committed to an employment choice for the period.[11] The individual's out-of-pocket expenditure is then computed by applying the rules of

[11]As a result of the assumptions that there is no saving, no possibility of acquiring health insurance, and a new job offer is available every period, the medical expenditure realization in a given period affects consumption in that period but has no future consequences. This is why it is important to model period t medical expenditure as being realized *after* the period t employment decision is made. Otherwise, after observing the period t medical expenditure realization, an individual without RHI could simply choose to be employed each time a bad realization occurs, and to be non-employed otherwise.

his health insurance plan(s). The rules include the premium, deductible, co-payment rate, and maximum out-of-pocket expenditure. Out-of-pocket expenditure for spouse *a* is denoted $m_{at}*$, and is given by $m_{at}* = \mathrm{F}(m_{at}, HI_{at})$, where m_{at} is total medical expenditure and HI_{at} is an indicator of the type (or types) of health insurance plan. We specify a discrete approximation of the underlying continuous distribution of medical expenditure for each spouse, instead of a parametric continuous distribution. This makes computation faster and produces results very similar to those based on parametric distributions. Let m_{ahtk} represent the k^{th} masspoint, k=1,...,K, of the distribution of *total* medical expenditure facing spouse *a* in health status *h* at age *t*, and let p_{ahtk} be the probability of realizing this outcome. The distributions are allowed to differ by age, health, and gender.[12]

3.4. Income and Consumption

Earnings per period, W_a, are fixed and non-stochastic. We make this assumption because Social Security benefits depend on average lifetime earnings (specifically, Average Indexed Monthly Earnings, or AIME), and allowing for earnings uncertainty results in a drastic increase in the size of the state space.[13] We do allow for the most important form of earnings risk, namely

---

[12]In this paper, medical expenditure is not a choice. See Blau and Gilleskie (2003) for a joint model of employment and medical care decisions. Rust and Phelan (1997) assumed that individuals receive a random draw from a health-insurance-specific Pareto distribution for out-of-pocket medical expenditure. The differences between our approach and the approach of Rust and Phelan are discussed below.

[13]In the analysis, we condition on AIME observed in 1991, the year before the first HRS survey. This is treated as an initial condition, and because it varies across households, we must solve the model for each couple. By treating future earnings as known with certainty, we can update AIME for all possible numbers of periods worked and treat the resulting Social Security benefit conditional on the number of periods worked as data in the solution of the model, with the number of periods worked as the state variable. In contrast, if future earnings are unknown, then AIME must be treated as a state variable and the model must be solved for all possible

layoffs. Given that earnings are non-stochastic, it follows that non-earned income b_t is also non-stochastic and depends on the age, experience, tenure, and employment status of the husband and wife according to known rules summarized by the function

$$b_t = b(j_{\mathrm{mt}}, x_{1mt}, x_{2mt}, A_{\mathrm{mt}}, j_{\mathrm{ft}}, x_{1ft}, x_{2ft}, A_{\mathrm{ft}}, g_{\mathrm{t}}) \tag{2}$$

where j_{at} is the employment choice, x_{1at} and x_{2at} are work experience and job tenure, respectively, A is age, and g_{t} is other nonwage income. This function is shorthand notation for the rules of the Social Security system and private pension plans. Employment status matters because of the Social Security earnings test. Years of work experience affects the AIME. Job tenure affects benefits in many pension plans. Age affects eligibility for Social Security and pension benefits and the level of benefits.[14] Specific details are described in the next section and the Appendix. Asset accumulation is not modeled, so the asset income component of other nonwage income is treated as given.

Given the absence of saving, consumption in period t equals total family income net of out-of-pocket medical expenses:

$$C_{\mathrm{t}} = E_{\mathrm{jt}} + b_t - m_t^* - \Upsilon(E_{\mathrm{jt}}, b_t, m_t^*) \tag{3}$$

where $E_{\mathrm{jt}} = W_m 1[\mathrm{j}_{\mathrm{mt}}>1] + W_f 1[\mathrm{j}_{\mathrm{ft}}>1]$ is total earnings of the couple, 1[] is an indicator function

---

combinations of earnings realizations that result in distinct values of AIME. The resulting state space is far too large for computational feasibility.

[14]Tenure on jobs that began after the initial survey date is not a state variable because, as discussed below, pensions on jobs held after the first job are not modeled. An individual is assumed to become entitled for Social Security benefits in the first period in which he or she is not employed beginning at age 62. Thus, unlike Rust and Phelan (1997), we do not model the decision to apply for benefits. The Appendix provides details on some additional employment-related state variables included in the model to help capture important features of Social Security and private pension plans.

equal to one if the statement inside the brackets is true, and equal to zero otherwise, $m_t^* = m_{mt}^* + m_{ft}^*$ is out-of-pocket medical expenditure for the couple in period t, and $\Upsilon(\ )$ is a tax function, incorporating income and payroll taxes and accounting for the medical expense deduction.

3.5. Utility Function

Utility of the couple in period t depends on each member's choice of employment, and on health, consumption, age, and a preference shock. Let w_{at} indicate whether spouse a chooses to be employed in period t, regardless of whether the job is new or old: $w_{at} = 0$ if $j_{at} = 1$ (not employed); $w_{at} = 1$ if $j_{at} > 1$ (employed). Utility in period t is (suppressing a household subscript)

$$U_t = \bar{U}_t(C_t, j_{mt}, j_{ft}, h_{mt}, h_{ft}) + \epsilon_t(j_{mt}, j_{ft}, h_{mt}, h_{ft}) = \alpha_{0h_{mt}h_{ft}w_{mt}w_{ft}} + \frac{1}{1-\alpha_1}C_t^{1-\alpha_1}$$

$$+ \alpha_{2h_{mt}}e_{mt-1}1[j_{mt}=2] + \alpha_{3h_{mt}}e_{mt-1}1[j_{mt}=1] + \alpha_{4h_{ft}}e_{ft-1}1[j_{ft}=2] + \alpha_{5h_{ft}}e_{ft-1}1[j_{ft}=1]$$

$$+ \alpha_6(1-e_{mt-1})1[j_{mt}=1](1-e_{ft-1})1[j_{ft}=1] + \alpha_7 e_{mt-1}1[j_{mt}>1]e_{ft-1}1[j_{ft}>1]$$

$$+ \alpha_{8h_{mt}w_{mt}}A_{mt} + \alpha_{9h_{ft}w_{ft}}A_{ft} + \epsilon_t(j_{mt}, j_{ft}, h_{mt}, h_{ft}) \qquad \text{if } C_t>0 \qquad (4)$$

$$= \alpha_{10} + \epsilon_t(j_{mt}, j_{ft}, h_{mt}, h_{ft}) \qquad \text{if } C_t \leq 0 \qquad (5)$$

The utility function incorporates the following features. (1) The intercept (α_0) varies freely by employment and health of both spouses. Hence, employment preferences are allowed to be non-separable within a period across spouses.[15] (2) The utility of consumption is separable from

[15]One of the 16 intercepts is normalized to zero. The α's are defined for employed and not employed (w=0 or 1) and for bad and good health (h=0 or 1), not for death (h=2). It is important to allow for non-separability of preferences for leisure of spouses. For example, wives are 23 percentage points more likely to be employed if the husband is employed than if he is not employed. Wives are 13.4 percentage points more likely to exit employment between interviews if the husband also exits during the period than if he remains employed. These patterns are robust to controls for other factors, and are similar for the husband's employment choices conditional on the wife's employment status. See Baker (2002), Blau (1998), Blau and Riphahn (1999), Hurd (1990), Gustman and Steinmeier (2000), and Zweimüller, Winter-Ebmer, and

employment and takes the isoelastic form with coefficient of relative risk aversion α_1 independent of choices and states. (3) Employment preferences are dynamic and depend on health. There are utility costs (or benefits) of changing jobs (α_2 and α_4) and exiting employment (α_3 and α_5). (4) Interactions in utility between the employment dynamics of husbands and wives are allowed (α_6 and α_7). (5) Age affects utility with health-and-employment-specific parameters (α_8 and α_9). (6) There are combinations of employment choices and realizations from the medical expenditure distributions that would result in negative income net of out-of-pocket medical expenditure, and therefore "negative consumption" ($C_t \leq 0$). Such outcomes are not literally possible, of course, but carefully modeling the mechanisms that allow households to avoid such outcomes is too complex to attempt here. The parameter α_{10} is included as a summary measure of the disutility from this outcome.[16] (7) The utility shocks (ϵ) are assumed to be independently and identically Type I Extreme Value distributed within periods and over time. This rules out both serial correlation and contemporaneous correlation in the utility shocks, as well as correlation with the medical expenditure and health shocks.

This function characterizes the utility of a couple in which both members are alive. Since we do not model behavior following the end of a marriage, we do not specify a utility function for individual members of the couple.[17] Given that the choices we model are all discrete, the

---

Falkinger (1996) for similar findings in other data sources.

[16]In order to explain why we do not observe outcomes involving "negative consumption," a more detailed model might include an insurer of last resort such as a public hospital, an option to file for bankruptcy, the option to forego medical treatment, and the option to self-insure by saving. Any of these approaches would complicate the model considerably. The approach we adopt here is based on Rust and Phelan (1997).

[17]Divorce is modeled as the result of an exogenous stochastic process. A marriage ends in divorce or separation at the end of period *t*, an event denoted by $D_t=1$, with probability $\Pr(D_t=1)$.

implications of the family utility function model specified here cannot be easily distinguished from those of a collective model in which each spouse has his or her own utility function and the couple reaches a Pareto-efficient bargain.[18]

3.6. Value Function and Solution

The expected present discounted value (EPDV) of a couple's remaining lifetime utility in period $t<T^*$ resulting from a given joint employment choice, conditional on the vector of state variables ($\mathbf{s}_t$), health, and the vector of period t utility shocks ($\boldsymbol{\epsilon}_t$), but *not* conditional on medical expenditure, is

$$V_{j_{mt}j_{ft}}^{h_{mt}h_{ft}}(\mathbf{s}_t,\ \boldsymbol{\epsilon}_t) = E_t\bar{U}_t(C_t, j_{mt}, j_{ft}, h_{mt}, h_{ft}) + \epsilon_t(j_{mt}, j_{ft}, h_{mt}, h_{ft}) + \beta V_{t+1}(\mathbf{s}_{t+1}),$$

where

$$V_{t+1}(\mathbf{s}_{t+1}) = (1-\Pr(D_{t+1}=1))\Big[\sum_{h'_m=0}^{1}\sum_{h'_f=0}^{1}\pi_{mt}^{h_m h'_m}\pi_{ft}^{h_f h'_f}V^{h'_m h'_f}(\mathbf{s}_{t+1}) + \sum_{a=1}^{2}\pi_{at}^{h_a 2}V(h_{at+1}=2)$$

$$+\ \pi_{mt}^{h_m 2}\pi_{ft}^{h_f 2}V(h_{mt+1}=h_{ft+1}=2)\Big] + \Pr(D_{t+1}=1)V(D_{t+1}=1) \qquad (6)$$

In practice, we simply set $\Pr(D_t=1)$ to the sample average divorce rate observed between waves. In order to avoid the complications of modeling behavior following divorce, we assign a "terminal value" to the event of the marriage ending as a result of divorce, and do not model behavior following divorce. The terminal value is denoted $V(D_t=1)$, and is normalized to -200. We follow the same approach for widowhood: a terminal value is assigned and the behavior of the surviving spouse is not modeled. The terminal value function in the event of the death of spouse a is given by $V(h_{at}=2)$ and is also normalized to -200. The terminal value function in the event of the death of both spouses in the same period is treated similarly: $V(h_{mt} = h_{ft} = 2) \equiv -200$.

[18]Blundell et al. (1998) show that a static collective model imposes different restrictions on the data than a static unitary utility function model when there is a continuous choice variable such as hours of work. In the absence of a continuous choice variable, the sharing rule parameters in the collective model are nonparametrically unidentified. As shown in Blundell et al. (p. 12), the sharing rule parameters are identified from the labor supply function, and there is no labor supply function in a pure discrete choice model. Van der Klaauw and Wolpin (2002) identify the sharing rule by functional form.

and β is the discount factor. The $\bar{U}_t$ term is unknown at the time the employment choice is made because medical expenditure is realized only after the employment choice is made. The value function accounts for the divorce probability $\Pr(D_t=1)$, the probability of death of one or both spouses ($\pi_{at}^{h_a 2}$), and the probability of arriving in period t+1 with the marriage intact and both spouses alive and in any combination of good and bad health. $V_{j_{mt}j_{ft}}^{h_{mt}h_{ft}}(\mathbf{s}_t,\ \boldsymbol{\epsilon}_t)$ is the expected value of $V_{j_{mt}j_{ft}}^{h_{mt}h_{ft}}(\mathbf{s}_t,\ \boldsymbol{\epsilon}_t,\ m_{mt},\ m_{ft})$ with respect to the distribution of medical expenditure:

$$V_{j_{mt}j_{ft}}^{h_{mt}h_{ft}}(\mathbf{s}_t,\ \boldsymbol{\epsilon}_t) = \sum_{k=1}^{K}\sum_{\ell=1}^{K} p_{mh_m tk}\, p_{fh_f t\ell}\, V_{j_{mt}j_{ft}}^{h_{mt}h_{ft}}(\mathbf{s}_t,\ \boldsymbol{\epsilon}_t,\ m_{mh_m tk},\ m_{fh_f t\ell})$$

$$= \left(\sum_{k=1}^{K}\sum_{\ell=1}^{K} p_{mh_m tk}\, p_{fh_f t\ell}\, \bar{U}_t(C_{jtk\ell},\ j_{mt},\ j_{ft},\ h_{mt},\ h_{ft})\right) + \epsilon_t(j_{mt},\ j_{ft},\ h_{mt},\ h_{ft}) + \beta V_{t+1}(\mathbf{s}_{t+1}),$$

where $C_{jtk\ell}$ is given by (3) with m_{mhtk} and $m_{fht\ell}$ substituted into the expression for m^* in (3). This expression integrates (by summation) over the distributions of medical expenditures, and illustrates the fact that period-t medical expenditure has no future consequences. The maximal EPDV of remaining lifetime utility from being in health states h_{mt} and h_{ft} unconditional on period t choices is

$$V^{h_{mt}h_{ft}}(\mathbf{s}_t) = E_{t-1}\max V_{j_{mt}j_{ft}}^{h_{mt}h_{ft}}(\mathbf{s}_t,\ \boldsymbol{\epsilon}_t) \tag{7}$$

where the expected value is taken with respect to $\boldsymbol{\epsilon}_t$, and the max is taken with respect to the choices j_{mt}, j_{ft}.

The model is solved numerically by backward recursion with two approximations used in the solution instead of solving the model exactly. The recursive solution must be computed for every couple and every trial value of all parameters. The state space is very large, and the exact solution takes too long to compute for this to be feasible. Thus we truncate the decision period at

$T < T^*$. The value function at T+1 is approximated as

$$V^{h_{mT+1}h_{fT+1}}(\mathbf{s}_{T+1}) \approx f(\mathbf{s}_{T+1}, h_{mT+1}, h_{fT+1} \mid \varsigma), \tag{8}$$

where f( ) is a function of the state space at T, and ς is a vector of parameters estimated jointly with the other parameters of the model. In addition to reducing the computational burden, by setting T so that the oldest age to which we model behavior is relatively young (71 in what follows below), we avoid solving the model for ages far beyond those observed in the data. The second approximation is to compute the value function for a randomly selected subset of points in the state space at each period instead of for all points in the state space. Following Keane and Wolpin (1994) the value function is regressed on the state variables using the sample of randomly selected points. The estimated regression parameters are used to approximate the expected value of next period's value function.

As part of the recursive solution, choice probabilities are computed for the observed choices for each couple. As a consequence of the assumptions about the ϵ's, the choice probabilities have the multinomial logit form:

$$p(d^t_{j_m j_f} = 1 \mid \mathbf{s}_t, h_{mt}, h_{ft}) = \frac{\exp(\bar{V}^{h_{mt}h_{ft}}_{j_{mt}j_{ft}}(\mathbf{s}_t))}{\sum_{j'_{mt}=1}^{J(\mathbf{s}_t)} \sum_{j'_{ft}=1}^{J(\mathbf{s}_t)} \exp(\bar{V}^{h_{mt}h_{ft}}_{j'_{mt}j'_{ft}}(\mathbf{s}_t))} \tag{9}$$

where $d^t_{j_m j_f} = 1$ if the couple chooses alternatives j_m, j_f in period t, and equals zero otherwise, J( ) is the number of employment alternatives available (J=3 if the individual was employed in the previous period and not laid off at the beginning of the period; otherwise J=2), and

$$\bar{V}^{h_{mt}h_{ft}}_{j_{mt}j_{ft}}(\mathbf{s}_t) = V^{h_{mt}h_{ft}}_{j_{mt}j_{ft}}(\mathbf{s}_t, \boldsymbol{\epsilon}_t) - \epsilon_t(j_{mt}, j_{ft}, h_{mt}, h_{ft}).$$

3.7. Identification

The two key sets of identifying assumptions are that (1) saving, health insurance, and medical expenditure are not subject to choice, and (2) there is no unobserved individual heterogeneity. Under the first set of assumptions, consumption equals current income net of medical expenses, and this is determined by employment choices and medical expense draws given an individual's health insurance coverage. With a normalization of one of the utility function intercepts (and of the terminal value functions in the event of divorce or death as described above), the remaining parameters are identified. The observed employment choices directly identify all of the utility function parameters except α_1, the risk-aversion coefficient, and α_{10}, the parameter determining the disutility of "negative consumption." The former is identified by the employment choices of couples whose health insurance depends on employment, relative to the choices of couples whose health insurance is independent of employment. Given the medical expenditure risk they face, the employment choices of the former group relative to the latter group identify willingness to bear medical expenditure risk. α_{10} is identified by the employment choices couples make in order to avoid "negative consumption." The health transition parameters are identified by the observed health transitions. The terminal value function parameters (ς) are identified by the nonlinear form in which they enter the choice probabilities.

The assumptions of no saving, no control over health insurance, and no control over medical expenditure are clearly quite strong. We argued in section 3.1 that our estimates are nevertheless useful because we can predict the direction of the bias caused by these assumptions. By limiting other mechanisms for smoothing consumption to account for medical expenditure risk, the model forces individuals who wish to avoid exposure to risk and who lack retiree health benefits to remain employed until they become eligible for Medicare. Thus our estimates provide an upper bound on the magnitude of the effect of health insurance operating through the budget constraint on retirement.

However, this line of reasoning does not account for unobserved heterogeneity in preferences. There are a number of possible biases that could be induced by such heterogeneity, but we believe the most plausible story would result in an upward bias in the estimated impact of health insurance on retirement. If preferences for leisure vary in the population, then individuals who value leisure highly will expect to want to retire before age 65, giving them an incentive to seek employment at a firm that offers retiree health insurance coverage. Individuals who place less value on leisure will be more likely to expect to remain employed until age 65, and therefore do not have an incentive to seek retiree health insurance coverage. In this scenario, differences in retiree health insurance coverage do not cause differences in employment behavior - both differences result from unobserved heterogeneity in preferences. The observed association will yield an upward biased estimate of risk aversion and a spuriously large impact of health insurance on retirement. Thus the argument that our estimates are an upper bound on the true effect of health insurance on retirement via the budget constraint is robust to the presence of at least one important form of unobserved heterogeneity in preferences.[19]

3.8. Estimation

---

[19]The most convincing sources of identification of health insurance effects in the literature to date have been the introduction or expansion of government programs such as Medicaid (Currie and Gruber, 1996) and COBRA (Gruber and Madrian, 1996), and random assignment of insurance in the RAND health insurance experiment (Manning, 1988). Analyses based on policy or social experiments generally do not attempt to identify and estimate parameters of economic models with a clear behavioral interpretation, which is the goal of our analysis. See Todd and Wolpin (2003) for a recent innovative analysis of a structural model using data from a social experiment.

Let t=1 denote the period in which we first observe a couple (the 1992 survey), and let t=0 denote the period prior to the first observation. We observe employment choices as well as health for t=1-4, and employment-related state variables for periods t=0-3. We treat employment in t=0 and health in t=1 as initial conditions that are not to be explained by the model, an approach that is consistent with the assumption of no unobserved heterogeneity. We take employment choices for t=1-4 and health transitions from t=1-4 as quantities to be explained by the model. The likelihood contribution for a couple that does not divorce or attrit from the sample and in which neither spouse dies by period 4, is

$$\mathcal{L} = \Big(\prod_{t=1}^{3}\Big[\prod_{a=m}^{f}\pi_{at}^{h_{at}h_{at+1}}\Big](1-\Pr(D_{t+1}=1))^{1-D_{t+1}}\Big)*\Big(\prod_{t=1}^{4}\Big[p(d_{j_m j_f}^{t}=1|\mathbf{s}_t,\ h_{mt},\ h_{ft})\delta^{L_{at}}(1-\delta)^{1-L_{at}}\Big]\Big), \qquad (10)$$

where h_{at}, j_{at}, D_t, and L_{at} are the observed health, employment, divorce, and layoff outcomes, and the employment choice probability p( ) is given by (9). The likelihood contribution for a couple that divorces before the beginning of period t+1 omits the choice and layoff probabilities for periods t+1 and later, and replaces $(1-\Pr(D_{t+1}=1))^{1-D_{t+1}}$ with $\Pr(D_{t+1}=1)^{D_{t+1}}$. The likelihood contribution for a couple that experiences a death at the end of period t omits the choice, layoff, and divorce probabilities for periods t+1 and later. The likelihood contribution for a couple that attrits between periods k and k+1 is the same as (10) with k replacing “3" and “4" in the product terms.

4. Data

4.1. Sample

We use data from the first four waves of the HRS, fielded in 1992, 1994, 1996, and 1998. The survey includes extensive sections on employment, pensions, health insurance, Social Security, earnings, assets, income, and health, among other topics. Two additional sources of information have been matched to the survey responses. The Social Security Earnings Records (SSER) of individuals who agreed to sign release forms were made available by the Social Security Administration. And individuals who reported being covered by a pension or by employer-provided health insurance in the first interview were asked to provide the names and addresses of the firms that provide the coverage. These firms were surveyed by telephone as part of the Health Insurance and Pension Provider Survey (HIPPS), and were asked to provide details of health insurance plans over the telephone and to provide written descriptions of their pension plans. These supplementary sources of data provide crucial pieces of information that allow us to measure the budget constraint facing each couple more accurately than would be possible with information provided by the HRS respondents alone. However, they also limit the sample that we can use because there are many cases in which the supplementary information is unavailable.

Of the 4,704 married couples who were surveyed in 1992, 3,005 meet our age criteria. The age criteria are that the husband is 51-63 at the 1992 survey, his wife is not older than him, and his wife is at most ten years younger than him. We restrict the sample to couples in which the wife is not older than the husband because the age of the husband is (arbitrarily) used to specify the terminal period. We lose about 300 cases due to missing data on key variables, and about 1,000 observations as a result of missing data on Social Security. The estimation sample is 1,752 couples. The estimation sample is similar to the full sample in most respects (see the Appendix).

4.2. Employment

Employment status is recorded at each survey date. The survey dates are on average two years apart. A job history collected at wave 1 allows us to determine employment status two years prior to the date of the wave 1 interview. This gives us a measure of employment status at t=0. Table II displays the employment distributions at the survey dates for the estimation sample. The rate of non-employment increases by about 6 percentage points on average between surveys for men and by 4.5 points for women. The rate of joint non-employment increases by about four points on average between surveys.[20]

4.3. Medical Expenditure

In order to estimate the model, it is necessary to characterize the distribution of medical expenditure facing older individuals. There is a substantial amount of information on medical expenditure in the HRS, but it is problematic for our purposes for several reasons. First, medical expenditure information was collected beginning in HRS wave 2, but the design of this module changed significantly in wave 3. Thus, there is no data for wave 1 and the wave 2 data are not comparable to the data from waves 3-4. Second, information is collected only on out-of-pocket expenditure, rather than total expenditure. We used data from waves 3 and 4 to construct a measure of out-of-pocket medical expenditure for each individual. Following Rust and Phelan (1997), the sample was classified by type of health insurance and the distribution of medical expenditure within each health insurance category was calculated. We were surprised to find that the distributions were quite similar across different health insurance categories. We had expected

[20]Note the large increase in job changing in periods 3 and 4 for men. This may indicate a "seam" problem, but we were not able to find any HRS documentation on this issue.

the distribution for individuals without health insurance to be substantially riskier (i.e., show a longer right tail) than those for individuals with health insurance, but this was not the case. It seems that uninsured individuals either are relatively healthy, are able to economize on expensive medical care, or have access to "insurance of last resort" at public hospitals. In order to solve the model, we need to determine the distribution of out-of-pocket medical expenditure facing an individual who would lose his health insurance coverage if he left employment. If we were to assume that this distribution is equal to the distribution of expenditure among those who are actually uninsured in the HRS, then individuals in our model would face little consequence from going uninsured. This would make it essentially impossible to identify the degree of risk aversion.

Hence, we took a different approach to constructing the distribution of medical expenditure. We used the 1987 National Medical Expenditure Survey (NMES) to estimate the distribution of *total* medical expenditure for older married individuals (by age, sex, and health status), including expenditure on hospital stays, outpatient visits, prescription drugs, and all other medical items. This is one of the few surveys that collected information on total medical expenditure rather than just the out-of-pocket portion. We assume that all individuals face the distribution of total medical expenditure corresponding to their age, sex, and health. After considerable experimentation with alternative distributional assumptions, including the Pareto as in Rust and Phelan (1997), we found that a three-point discrete approximation to the distribution of total medical expenditure provided an adequate characterization of the distribution, and resulted in much faster speed of computation than in the case of the Pareto.[21] Multinomial logit models were used to estimate the distribution across the following three categories of total annual medical expenditure (in 1992 dollars): $0-1,999, $2,000-14,999, and $15,000 and above. The models were estimated separately by sex and health, including only an intercept and a linear age term. The coefficient estimates were then used to predict the probability of receiving a draw in each category at each age. Linear regressions were used to predict mean expenditure by age within each category, and the predicted mean by age is used as the value of total medical expenditure assigned to a given category, conditional on receiving a draw from the category.

Table III summarizes the estimated distributions at ages 51, 61, and 71. The probability of a bad draw is much larger when in bad health, and increases with age at a faster rate in bad health than in good health. For individuals in bad health, the probability of catastrophic medical expenditure (15K+) is .06 at age 51, so going uninsured is quite risky for such individuals.

Out-of-pocket expenditure (m_{at}*) is determined by the rules of the health insurance plan (or plans; the function F(m_{at}, HI_{at}) defined above) in which an individual is enrolled, for any given draw m_{at} from the distribution of total medical expenditure. Uninsured individuals, including those with employer-provided insurance without retiree coverage who leave the labor force before eligibility for Medicare, pay the full amount of their realized draw. The rules governing employer, private, and Medicare insurance are determined as follows. The HRS

---

[21]Rust and Phelan use data on realized medical expenditure of the individuals in their sample, but we do not do this for reasons explained above. Accordingly, we must integrate over the medical expenditure distribution in order to solve the model and form the likelihood function. In the context of our model, there is no closed form solution to the integral for the Pareto or any other continuous distribution, so time-consuming monte carlo integration would be required. A discrete approximation provides much faster computation time than monte carlo integration.

collected data on health insurance plan rules from the wave 1 employers of HRS respondents, but there was a substantial amount of missing data. Thus, instead of using the actual rules for each observed employer-provided plan, which would drastically reduce the sample size, we constructed a "generic" plan based on the median characteristics of employer-provided plans. The generic plan includes a premium, deductible, co-payment rate, and maximum out-of-pocket expenditure. There is no information about the rules of private health insurance plans in the HRS, so we used the NMES to construct a generic private plan which is assumed to apply to all individuals with private coverage. Finally, Medicare rules are known. The parameters of the insurance rules are described in the Appendix.

4.4. Health

We take a very simple approach to measuring health in order to focus on the economic aspects of the analysis and avoid the proliferation of parameters that would result from exploiting the richness of the HRS health data.[22] A dichotomous measure of health is constructed from responses to the question "Would you say your health is excellent, very good, good, fair, or poor?" by combining excellent, very good, and good (good), and poor and fair (bad). The distribution of health and health changes is shown in Table IV. Of the men who report good health at wave 1, 8.5 percent are in bad health by wave 2 and 1.1 percent have died. The corresponding figures for women are 7.1 and 0.4 percent. Of those who report bad health at wave 1, 24.1 percent of men and 29.9 percent of women are in good health at wave 2, and 7.3 percent of men and 3.9 percent of women have died. The figures for transitions between other waves are

[22] See Blau and Gilleskie (2001b), Bound et al. (1999), and Dwyer and Mitchell (1999) for detailed analyses of the effect of health on employment in the HRS.

similar. The large differential in death rates by self-assessed health status illustrates the fact that self-reported health contains useful "objective" health information.[23]

4.5. Health Insurance

Each individual is classified into one of the 7 mutually exclusive and exhaustive health insurance categories shown in Table V for both primary and secondary coverage. Individuals with multiple sources of health insurance were assigned to categories in the following order: own-employer with retiree health insurance (RHI); spouse employer with RHI; own-employer without RHI; spouse employer without RHI; Medicare (before age 65); private. For example, a man with both employer-provided coverage and privately purchased coverage is assigned to employer coverage as the primary category, and private coverage as the secondary category. At age 65, each individual is assumed to enroll in Medicare and retain up to two other types of coverage.

There was a major redesign of the health insurance section of the HRS in wave 3, and other significant changes in wave 4. There was also a significant error in the skip logic in waves 3 and 4 for some key health insurance questions. As a result, the wave 3 and 4 health insurance data are not directly comparable to the wave 1 and 2 data, and we do not use them here. We assign individuals their observed health insurance coverage at waves 1 and 2, and we assume that they expect to continue to be covered by their wave 2 source of health insurance in the future, except for those periods in which individuals with employer coverage without retiree benefits choose to be not employed.

[23]There is a strong positive association between health of husbands and wives: of men in good health, 85 percent have wives who are also in good health, while for men in bad health only 65-67 percent have wives in good health (not shown in the table).

4.6. Income

We treat earnings as deterministic for reasons discussed above. Aside from the risk of layoff, which we do model, we view earnings fluctuations as a relatively minor source of risk at older ages compared to medical risk. Consequently, the main issue for modeling earnings is how to obtain good forecasts to include in the model as a measure of expectations about future earnings. After considerable experimentation, we decided to use the most recent measure of earnings from the Social Security Earnings Record (SSER) file, denoted *lastearn*, assuming no growth in real earnings. There are two exceptions to this rule. First, earnings records in the SSER file are truncated at the maximum taxable annual earnings, so individuals with *lastearn* equal to the taxable maximum were assigned their self-reported earnings from the HRS. Second, there were many cases in which *lastearn* was implausibly small as a measure of full-time year-round earnings, most likely because the individual worked only part of the year. In these cases, we assigned either the self-reported value from the HRS, if it was larger than a specified threshold, or the predicted value from a regression otherwise. Details are provided in the Appendix.

The HRS collects detailed data from respondents on pensions for all jobs held on or before the wave 1 interview date that provide pension coverage. These data provide a rich source of information, but do not include the actual formula used to determine the pension benefit as a function of age of exit from the firm, tenure, earnings, and so forth. This formula is needed in order to compute the benefit to which the respondent would be entitled at different ages of exit from the firm. The written plan descriptions included as part of the HIPPS interview provide the information needed to construct the formula. We used these data together with the HRS survey responses to compute the benefit from the pension on the job held at period t=1 (if any) for every possible quit date from 1992 until the respondent reaches the terminal age (71). For pensions provided by employers on jobs that ended before 1992, we compute the benefit to which the individual would be entitled at the earliest age at which he is eligible for a benefit under the plan. The Appendix summarizes characteristics of pensions.

We use the SSER earnings history from 1951 through 1991 to construct each individual's AIME and Primary Insurance Amount (PIA) as of 1991, using the formula in effect for 1991. The AIME is a deflated average of earnings from age 21 to the current age, excluding the lowest five years of earnings. The PIA is a piecewise linear, highly progressive function of the AIME, and is the basis for computing the Social Security Benefit (SSB). We then use the earnings measure described above to update the AIME and PIA for each of the possible total number of additional years of work experience the individual could accumulate from 1992 through the terminal age. We use these to compute the SSB for which an individual would be eligible upon exiting the labor force for each possible number of years of experience from 1992 through age 71. These benefit measures are based on the exact formulas used by the Social Security Administration (which differ by cohort as the 1983 Social Security reforms are phased in), accounting for permanently reduced benefits for early retirement and increased benefits for delayed retirement. The spouse benefit is also computed and added to the family benefit if appropriate. We do not model the decision to apply for Social Security benefits. Instead, we assume that an individual begins to receive Social Security benefits in the first period in which he or she is not employed from age 62 on. Further details and summary statistics on the SSB are provided in the Appendix. Rust and Phelan (1997) model the choice of when to claim benefits. This is an important issue, but is not central to our focus, and accounting for this choice would increase the choice set substantially.

Other sources of nonwage income include income from assets, rent, alimony, veteran's benefits, and earnings of family members other than the husband and wife. We treat these sources of income as given and certain. Benefits from SSDI are included if an individual is covered by Medicare while under age 65 and chooses to remain non-employed. Income from means-tested government programs such as Supplemental Security Income (SSI) is included if the individual chooses non-employment and reported receiving such income. We summed these sources to create a measure of nonwage income for the calendar year prior to each interview. The average of these measures is assigned as the value of nonwage income for all periods after wave 4.

We use the 1992 Federal income tax and payroll tax schedules to compute measures of after-tax income. The computations account for taxation of Social Security benefits and the medical expense deduction. Because income is measured on an annual basis and the length of a period is two years in our model, we set consumption equal to twice the value of after-tax income net of out-of-pocket medical expenses.

5. Estimation Results

5.1. Estimates

Table VI shows estimates of the utility function parameters. The coefficient of relative risk aversion (CRRA; α_1) is estimated to be 2.144. This is substantially higher than some other estimates using micro data, such as Rust and Phelan's (1997) estimate of 1.072 and Hurd's (1989) estimates of 0.73 and 1.12, but it is lower than the values often assumed in simulation studies of savings behavior. A CRRA of 2.144 implies that the certainty-equivalent value of a gamble that pays $20,000 and $40,000 with equal probability is $26,500, compared to $30,000 under risk neutrality. This is a fairly substantial degree of risk aversion, and suggests that health insurance should be of considerable value. The utility intercepts (α_0) are usually lower in a bad health state compared to an otherwise similar good health state. The utility intercepts are almost always lower when the wife is employed than when she is not employed, other things equal, as would be expected if leisure is a normal good. But this pattern appears for husbands only when health is bad.

Changing jobs and exiting employment are both estimated to have utility costs (α_2, α_3, α_4, $\alpha_5 < 0$), and there are utility gains if both spouses remaining in the same employment state (α_6, $\alpha_7 > 0$). The utility of being not employed rises with age for both spouses but at a faster rate when in good health than when in bad health. The utility of being employed rises with age for the husband if he is in good health and declines with age if he is in bad health. The utility of being employed rises with age for the wife at a faster rate if she is in good health than if she is in bad health. The utility associated with negative consumption is positive and small.

Other parameter estimates are shown in the Appendix (Table XV). The health transition parameters were estimated separately, as in Rust and Phelan (1997). Given the assumed absence of unobserved heterogeneity, this provides consistent estimates, although the standard errors of the utility function parameters have not been adjusted to reflect the separate estimation of the health transition parameters. The parameters for the probability of divorce and layoff were estimated outside the model and fixed. The discount factor was fixed at .925 for two year periods, implying a biannual rate of time preference of 8.1 percent. After considerable experimentation with alternative forms for the terminal value function, we found that the other parameters of the model were insensitive to the form of the latter. Hence, we specified the terminal value function as $-e^{\varsigma}$, where ς is a constant. The estimated value of ς is 2.725, which implies a terminal value at age 71 of the husband of -15.3 in the event that both spouses are alive at that date. This is higher than the terminal values for death and divorce, which were normalized to -200.

5.2. Model Fit

The fit of the model to the individual employment choices shown in the first row of panel A of Table VII is generally good: the predicted non-employment and job changing probabilities differ from the observed rates by 2-5 percentage points for men and 1-2 points for women. The model also fits reasonably well in most cases when choices are classified by employment status in the previous period, as shown in the next two rows of the table. The exit rate from employment is underpredicted by about 5 percentage points, and the predicted job-changing rate is off by two points. The predicted distribution of joint employment choices and transitions shown in panel B fits the actual distribution very well in most cases, with or without conditioning on employment status in the previous period. The age-specific fit of the model to the non-employment rate is shown in Figure 1. The model fits quite well at ages for which the sample size is large, but shows some significant departures from the data at the oldest ages.

Table VIII illustrates the fit of the model to employment choices by health insurance category. Recall that health insurance enters the model only through the budget constraint - there are no utility function or other parameters on health insurance. In view of this, the model does a reasonably good job of fitting employment choices by health insurance category. The model captures the higher rate of non-employment of individuals with own-employer insurance with RHI compared to individuals without RHI (men: actual .320 vs. .123; predicted .285 vs. .159; women: actual .214 vs. .098; predicted, .220 vs. .146). The model also reproduces the higher non-employment rate of the spouse when an individual's insurance coverage is from the spouse's employer with versus without RHI (men: actual .346 vs. .149; predicted .304 vs. .173; women: actual .240 vs. .079; predicted .246 vs. .148). In most cases, the model underpredicts the magnitude of the difference between the RHI and no-RHI cases, but the predictions are of the right order of magnitude.

Table IX provides a more exacting analysis of the model fit by comparing predicted to actual employment choices by health insurance category *conditional* on employment status in the previous period. The model fits the qualitative features of the data in most cases, but is sometimes far off quantitatively. The data show a 22.2 percentage point difference for men in the rate of remaining non-employed as a function of retiree health insurance (.800 vs. .578), while the model predicts a 15.6 percentage point difference (.845 vs. .689). This prediction is of the right order of magnitude, though quite a bit smaller than the actual value. However, the data show an 8.4 percentage point higher rate of exit from employment for men with retiree health insurance (.143 vs. .059), while the model predicts essentially no difference (.078 vs. .085). Similar patterns appear for women. Thus there are some dimensions of the observed health-insurance-employment relationship that are not captured by our model. We discuss possible explanations for this below.

5.3. Simulations

In order to assess the implications of the estimates, we used them to simulate employment behavior under alternative scenarios. The simulations were computed for each couple in the sample conditional on the couple's initial state observed at the beginning of period 1. Couples were randomly assigned an alternative from their choice set in period 1, with the probability of a particular alternative equal to the choice probability for that alternative computed from the model. The estimated health transition and layoff probabilities were used to randomly assign health and layoffs in period 2. The assigned choice along with health and layoff status were then used to update the state for the next period. This was repeated through t=6 or T, whichever came first.[24] The results were then averaged over the sample. Table X shows the results for the choice of non-employment, averaged over all ages.

Column one of Table X shows results from a baseline simulation of employment choices and transitions using the observed health insurance. Columns two and three show results from simulations in which all individuals are assigned no health insurance and universal health insurance independent of employment status. Universal health insurance was implemented by re-assigning all observations to the generic own-employer plan with retiree benefits. The no-health-insurance scenario removes all health insurance prior to age 65, but retains Medicare eligibility at age 65. These are extreme policy changes, but they are useful for bracketing the magnitudes of the employment responses implied by the model. If medical expenditure risk causes couples to remain employed longer than they might otherwise prefer in order to avoid being uninsured, then universal health insurance should increase the rate of non-employment compared to the baseline case and to the no-health-insurance scenario. The results in Table X show that this is in fact the

[24]In principle, this could be done for up to 11 periods, the maximum number of periods for which the model is solved for the youngest cohort (a man who is 51 in t=1 is 71, the terminal age, in t=11). However, in simulation, unlike in estimation, choice probabilities must be stored for each alternative at every point in the state space, and the state space is too large to make this feasible beyond period 6. In the simulations we assume that there are no divorces or deaths, so the sample composition does not change over time.

case, but the magnitude of the effect is modest. Averaging over the entire sample, universal health insurance increases the non-employment rate for both men and women by about 3.5 percentage points compared to no health insurance (.336 vs. .301 for men; .420 vs. .384 for women). The bi-annual exit rate from employment increases by only half a percentage point for both men and women. Comparing the universal health insurance case to the baseline yields even smaller effects, presumably because the majority of the sample already has health insurance in the baseline.

To get a better sense of how the availability of health insurance affects employment behavior, the remaining rows of Table X show simulation results classified by baseline health insurance. For those with own-employer health insurance with RHI at baseline, taking away health insurance reduces the rate of non-employment by .037 for men (.322 vs. .285) and .027 for women (.283 vs. .256). For individuals with own-employer insurance without RHI at baseline, making insurance independent of employment increases non-employment by .019 for men (.207 vs. .236) and by .060 for women (.219 vs. .279). The effects on the exit rate from employment are again very small. Thus only in the case of women without RHI would universal health insurance increase the rate of non-employment by more than 3.5 percentage points. It is also worth noting that the welfare effects of these large policy changes are estimated to be quite small. Using the period 1 value function as the metric, welfare increases by less than half of one percent on average between the no-health-insurance and universal-health-insurance scenarios (not shown in the table).

Columns 4 and 5 show the simulated employment response to policies that add RHI to all employer plans that lack it and remove RHI from all employer plans that have it. Removing RHI causes the non-employment rate to fall by 4.7 percentage points for men compared to the baseline

(.322 vs. .275), and by 4.2 points for women (.283 vs. .241). Adding RHI to plans that lack it increases the non-employment rate by 3.0 percentage points for men (.207 vs. .237) and by 6.1 points for women (.219 vs. 280).

The final column in Table X shows the results of a simulation in which the age of Medicare eligibility is changed from 65 to 67. This is a more realistic policy change than the others shown in the table, since it matches the scheduled increase in the age at which the full retired worker Social Security benefit can be collected.[25] This policy change would have very small effects on employment according to these results. The largest impact on the non-employment rate is a decrease of one percentage point for women with own-employer insurance with RHI, and most of the effects are on the order of 0.1-0.2 percentage points. Figure 2 plots the results of this simulation by age, and shows that there is only a tiny impact on men, even at age 65. For women, increasing the age of Medicare eligibility to 67 results in about a ten percentage point reduction in non-employment at age 65, no effect at age 66, and another large decrease in non-employment at age 67. This odd pattern for women may be due to the fact that the sample size for women at ages 65 and above is small, and the results therefore may not be reliable.

5.4. Discussion

In Table I, we saw that controlling for a host of other factors, men with RHI were 4.4 percentage points less likely to be employed than men without RHI. In Table X, we find that the simulated impact of removing RHI from men who have it is a 4.7 percentage point increase in the

[25] The 1983 amendments to the Social Security Act scheduled increases in the age at which the full Social Security retirement benefit (100 percent of the PIA) can be collected from 65 to 66 for individuals who turn 62 in 2000 to 2005, and from 66 to 67 for individuals who turn 62 from 2017 to 2022.

employment rate, and the simulated impact of providing RHI to men who lack it is a 3.0 percentage point decrease in the employment rate. Thus, our model is able to account for the relatively modest impact of RHI for men that we observe in the data by attributing it to aversion to medical expenditure risk. For women, Table I shows a 17.1 percentage point difference in the employment rate between those with and without RHI, other things equal, while Table X shows simulated impacts of removing and adding RHI of 4.2 and 6.1 percentage points, respectively. Thus the model can account for only about one third of the large association between employment and RHI for women observed in the data.

We examined several possible explanations for the relatively modest impact of health insurance on employment implied by our model for both men and women. First, the estimated CRRA of 2.144 indicates that couples are moderately but not highly averse to the risk of income fluctuations associated with health shocks. If couples are not highly risk averse, then insurance availability may not be a very important determinant of employment. In order to determine whether absence of substantial risk-aversion is responsible for the moderate impact of health insurance in our model, we computed simulations in which the risk aversion coefficient was set to 0 and 20, while all other parameters and variables remained unchanged. A CRRA of 20 implies a certainty-equivalent value of $20,743 for a gamble that pays $20,000 and $40,000 with equal probability, compared to $26,500 with the estimated value of the CRRA of 2.144. This is obviously an extreme degree of risk aversion, while a CRRA of zero implies risk neutrality. We re-computed the baseline simulation in column one of Table X and the column five simulation that eliminates all retiree health insurance. The results for a CRRA of 20 are quite similar to those in Table X, showing moderate sensitivity of employment choices to availability of RHI. The results for a CRRA of zero were quite different: in this case RHI makes virtually no difference at all to employment choices. Hence, *some* risk aversion is clearly necessary in order for RHI to matter, but even extreme risk-aversion would not increase the impact of RHI beyond the relatively modest effect that we find.

Second, we respecified our model to include health insurance indicators in the health transition functions. By allowing health insurance to affect health transition rates, we can determine whether health insurance affects employment behavior by influencing health, perhaps through (unmodeled) effects on medical care consumption, in addition to the effects that operate via the budget constraint. We added to the health transition model three binary indicators of health insurance: any employer-provided insurance, any private insurance, and any public insurance, with uninsured as the omitted category. This adds 12 new parameters to the health transition functions, and the results indicated that we could strongly reject the hypothesis that the parameters on these health insurance indicators are jointly equal to zero. The utility function parameter estimates and policy simulations with this specification were quite different. For example, the simulations indicate that removing RHI would have large effects on employment behavior. Indeed, in this case removing RHI closes most of the gap in both the non-employment rate and the exit rate to non-employment between those covered by employer insurance with and without RHI. This contrasts with the results in Table X, which show that removing RHI has very little impact on the exit rate to non-employment, and only moderate effects on the non-employment rate overall. The health transition estimates indicate that health insurance is good for your health, but because we have not accounted for individual heterogeneity we do not have a great deal of confidence that these estimates are reliable. We view these results therefore mainly as a specification test, and the results clearly indicates that at least for women our simpler specification fails to capture the full effect of RHI on employment behavior.

Third, we estimated a specification in which health insurance enters both the health transition functions and the utility function directly. We added to the utility function three binary indicators of health insurance: whether the individual has retiree coverage through his own or spouse's employer, whether he has coverage from his own or spouse's employer without retiree benefits, and whether he has private coverage or Medicare (before 65). Allowing the parameters on these variables to vary by employment choice and spouse, this adds 12 new parameters to the utility function. We do not have a specific argument as to why or through what mechanism health insurance would directly affect utility, net of its effects via the budget constraint and health transition functions. Hence, we refrain from giving an economic interpretation to this specification, and simply treat it as a test of our original specification. A likelihood ratio test strongly rejects the hypothesis that the health insurance parameters in the utility function are jointly equal to zero. The model fits the employment data significantly better with these additional parameters than without them. This indicates that health insurance likely affects behavior through some other channel in addition to the budget constraint and health transition functions. The precise nature of that channel cannot be determined from these estimates.

Our approach to modeling retiree health insurance and employment is most similar to that of Rust and Phelan (1997), whose results imply a larger impact of RHI on employment behavior than do our results. We use more recent data and have access to more accurate measures of health insurance coverage than Rust and Phelan. The HRS data also allow us to measure pension benefits accurately, and as a result our sample is more representative of the U.S. population than was their sample (which excluded workers with pension coverage). However, there is no straightforward way to determine the exact source of the difference between our results and theirs.

6. Conclusions

Our findings here indicate that health insurance has fairly modest effects on the labor force behavior of older married couples. This finding contrasts with previous estimates from reduced form and approximation models, which suggest that retiree health insurance has a strong effect on employment behavior at older ages. Studies using those approaches cannot determine whether health insurance affects employment behavior through aversion to the consequences of health risk or through some other mechanism. In the model analyzed here, aversion to the consequences of health risk is the *only* channel through which health insurance is allowed to affect behavior, and the effects we estimate are modest. Furthermore, we argue that these are almost surely upper bound estimates of the effect of health insurance on retirement operating through the budget constraint. Our results imply that policy changes that alter the connection between health insurance and employment will have relatively little impact on retirement decisions in the U.S. *via the budget constraint*. If there are other mechanisms through which health insurance affects retirement, as suggested by our findings, then there could be larger policy effects. It is important to do more research before accepting this conclusion, however. As computing costs fall, models that allow for saving, choice of health insurance, and unobserved heterogeneity should be explored. Our conjecture is that allowing for these factors would reinforce the conclusion that health insurance is not a major determinant of retirement decisions of older married couples, but further empirical analysis is required to determine if this is correct.

Appendix

A. Sample Selection and Characteristics

Table XI describes how we obtain the sample of 1,752 couples used in estimation. We lose 279 of the age-eligible couples as a result of missing information on employment, demographic variables, health insurance, and health, leaving 2,726 cases. Social Security records are missing for either the husband or the wife (or both) in 974 of these cases, leaving 1,752 cases. Most of the cases without Social Security records are the result of the absence of a signed release, but some cases may be due to the fact that an individual was never employed in a job covered by Social Security. This is difficult to determine from the HRS so we drop all cases without a Social Security record. After excluding cases with missing HRS or Social Security data, we do not lose any additional cases due to missing pension data. As described below, we were able to fill in a significant amount of missing pension provider data using information provided by the respondent. Table XII compares the estimation sample to the full sample of couples, and reveals few substantial differences.

B. Pensions

The HIPPS Survey obtained written plan descriptions for 6,381 pension plans. The plan characteristics were coded by the Institute for Social Research (ISR) at the University of Michigan into a computer program that calculates benefits under alternative scenarios. For jobs held at the wave 1 survey, we used the program to compute the benefit to which an individual would be entitled for every possible year in which he or she could quit the firm through age 71. The program takes as input the individual's age and tenure with the firm as of wave 1, and annual earnings for 1991 as reported by the respondent in the wave 1 survey. For jobs held prior to wave 1, we used the program to compute the benefit available at the earliest age of eligibility, taking as input the respondent's tenure and annual earnings at the time of exit from the firm. Benefits are computed for both defined benefit and defined contribution plans, with benefits for the latter expressed in the form of an annuity. Benefits are computed for up to three different plans from a wave 1 job and three different plans from previous jobs.

There was a substantial amount of missing data on pension benefits due to absence of written descriptions, and written descriptions that lacked some of the information needed to compute benefits. The HRS asks respondents to report the age at which they expect to start receiving benefits and the benefit amount for every pension plan for which they are or will be eligible for a benefit. We used these data to fill in missing values for pension benefits and age of eligibility for jobs held prior to wave 1, since the respondent's employment decisions made on or after the wave 1 interview date do not affect the pension benefit from jobs held prior to wave 1. These data are not sufficient to fill in missing information for pensions on jobs held at wave 1, since benefits from such jobs depend on the individual's employment decisions via the benefit formula, which we do not have in such cases.

The HIPPS survey covers wave 1 employers and previous employers but does not include any new employers after wave 1. If an individual took a job that provides pension coverage after wave 1 we have information from subsequent waves about characteristics of the pension but no information on the benefit formula, since the HIPPS survey was not repeated until recently and data from the new survey are not yet available. Thus we ignore pensions on jobs that begin after wave 1. Table XIII summarizes pension benefit information for the sample.

C. Nonwage Income

Nonwage income (other than pension and Social Security benefits) consists of income from interest, dividends, rent, alimony, private disability insurance, trust funds, royalties, veterans benefits, and the earnings of household members other than the husband and wife. Employment-conditioned sources of income such as welfare, unemployment compensation, and workmen's compensation are included if the individual chooses non-employment and reported receiving income from such sources. The mean of a couple's nonwage income from all available survey waves is assigned to the couple for periods beyond the end of the data in solution of the model. The sample mean of annual non-work-conditioned nonwage income in 1992 is $9,429, and the sample mean of work-conditioned income in 1992 is $600.

D. Social Security Benefits

The Social Security Benefit (SSB) is computed using the exact formula for an individual of his cohort and his age at the time of claiming benefits (which is a state variable), using the Primary Insurance Amount (PIA) computed from the Social Security Earnings Record in 1991 updated with the earning measure described below. The formula is cohort-specific as a result of the 1983 reforms that gradually increase the normal age of retirement to 67 and phase in other changes as well. However, if an individual experiences a non-employment spell at age 62 or above and then reenters the labor force, the SSB for which he is eligible upon exiting employment again can be computed exactly only by making the sequence of employment choices from age 62 onward a state variable. This makes the state space too large for solution of the dynamic programming problem. Instead we proceed as follows. First, we used the exact formula to calculate the benefit for which an individual would be eligible for every possible employment sequence involving reentry after age 62. We then regressed the benefit on the PIA corresponding to the cumulative years of experience associated with the sequence at the time of re-exit, with separate regressions by gender and each age of re-exit. Recall that cumulative experience is a state variable, and the PIA associated with each possible level of cumulative experience is part of the data set. We use the fitted values from these regressions to assign the SSB for individuals in their second nonemployment spell after age 61. The regression is of the form SSB = a + b*PIA, and the results are as follows:

	Men		Women	
Age	a	b	a	b
63	12.481	0.779	0.709	0.798
64	13.171	0.811	0.175	0.833
65	12.876	0.844	-0.861	0.870
66	14.465	0.884	-1.878	0.916
67	14.909	0.915	-3.013	0.951
68	15.528	0.944	-4.491	0.987
69	14.805	0.974	-6.428	1.023
70	13.294	1.005	-8.791	1.058

In order to follow this approach, whether an individual has reentered employment following a nonemployment spell after age 61 becomes a state variable. This increases the size of the state space, but not by as much as keeping track of the exact employment sequence.

We compute benefits conditional on employment as well as non-employment, applying the Social Security earnings test to determine the benefit entitlement conditional on being employed (only for re-entry to employment, i.e. after initially claiming benefits). This test, which is also cohort-specific, results in zero benefits for most individuals, but some low-earnings individuals have a positive benefit while employed. Finally, the spouse benefit is calculated and added to the family benefit if it is larger than the spouse's benefit based on his or her own work history. Table XIV displays summary statistics on the PIA for alternative amounts of work experience.

E. Earnings

Individuals are assumed to expect their earnings to remain unchanged in real terms as long as they remain employed, either on the same job or a new job. The 1991 value of annual earnings from the SSER (*lastearn*) is used for earnings except in the following cases. First, earnings records in the SSER file are truncated at the maximum taxable annual earnings, so individuals with *lastearn* equal to the taxable maximum were assigned their self-reported earnings from wave 1 of the HRS. Second, there were many cases in which *lastearn* was small, possibly because the individual worked only part of the year. We used the following rules to deal with such cases, defined in practice as *lastearn* less than $10,000. The average Social Security replacement rate for income less than $10,000 is .64. If the SSB implied by the PIA is less than .64**lastearn*, then we used HRS self-reported earnings instead of *lastearn* if the SSB implied by the PIA is less than .64*HRS earnings. Otherwise we used the (exponentiated) predicted value from the following regressions using HRS wave 1 data:

Men: ln(earnings) =	8.947	- .076*Black	-.105*Hispanic	+.102*Education.	$R^2 = .12$
	(.105)	(.081)	(.102)	(.008)	N = 1458
Women: ln(earnings) =	7.637	+ .188*Black	+.141*Hispanic	+.141*Education.	$R^2 = .11$
	(.157)	(.095)	(.136)	(.012)	N = 1163

If the SSB implied by the PIA is greater than .64**lastearn*, then we used *lastearn* unless it was zero, in which case we used HRS earnings if positive or the regression prediction otherwise. The mean of *lastearn* for men is $27,356 and for women is $12,432. The means of self-reported earnings in the HRS wave 1 are $32,208 for men and $12,337 for women.

F. Health Insurance and Medical Expenditure

The generic employer health insurance plan has a premium that varies by whether the spouse is covered and whether retiree health insurance (RHI) is provided. The values we use, based on medians reported in the HIPPS survey, are as follows: worker only with RHI: $966; worker and spouse with RHI: $1,019; worker only without RHI: $472; worker and spouse without RHI: $429. The deductible is $200, the coinsurance rate is .20, and the maximum out-of-pocket expenditure (excluding the deductible and premium) is $1,000 for non-inpatient expenses, and $1,200 for hospital inpatient expenses. The generic private plan has a premium of $2,222, a deductible of $100, a coinsurance rate of .20, and maximum out-of-pocket expenditure of $1,000. Some individuals are covered by VA-CHAMPUS, a public health insurance program for veterans and civilian employees of the military. It has no premium, a deductible of $150, a coinsurance rate of .25, and no maximum out-of-pocket expenditure. Medicare (including hospital (part A) and Supplementary Medical Insurance (part B)) has a premium of $493.20, an office visit deductible of $100, a hospital stay deductible of $696, and coinsurance rates of .20 for visits and $174 per hospital night in excess of 60. Given the maximum possible values of total medical expenditure shown in Table III, there is no effective maximum out-of-pocket expenditure in Medicare. Note also that unlike the other plans, Medicare does not cover prescription medication expenses.

In order to allocate expenditure between hospital, prescriptions, and other expenses for employer health insurance and Medicare, we used the median ratio of each of these expenses to total expenses, calculated from the NMES separately by sex, health, and expenditure category. Finally, applying the Medicare rules requires a measure of the number of hospital nights. We computed the average expense per hospital night in NMES, and used it to estimate the number of hospital nights as total hospital expenses divided by price (and rounded down).

Finally, secondary insurance coverage provided by an employer or private plan is applied to the out-of-pocket expenditure calculated from the primary plan, excluding the premium of the primary plan. Cases in which Medicare is the secondary coverage are more difficult to handle, because of very complicated rules. Officials familiar with Medicare rules told us that the practical impact of these rules is that Medicare is effectively useless as secondary coverage for medical expenses. Hence we do not include Medicare as secondary coverage.

References

Baker, M. (2002): "The Retirement Behavior of Married Couples: Evidence from the Spouse's Allowance," *Journal of Human Resources*, 37, 1-34.

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

Blau, D.M. (1998): "Labor Force Dynamics of Older Married Couples," *Journal of Labor Economics*, 16, 595-629.

Blau, D.M., and D.B. Gilleskie (2001a): "Retiree Health Insurance and the Labor Force Behavior of Older Men in the 1990's," *Review of Economics and Statistics*, 83, 64-80.

______ (2001b): "The Effect of Health on Employment Transitions of Older Men," in *Research in Labor Economics*, ed. By Solomon Polacheck. Greenwich CT, JAI Press.

_______(2003): " Health Insurance and Retirement," Unpublished manuscript, University of North Carolina at Chapel Hill.

Blau, D.M., and R. Riphahn (1999): "Labor Force Transitions of Older Married Couples in Germany," *Labour Economics* 6, 229-251.

Bound, J., M. Schoenbaum, T. Stinebrickner, and T. Waidmann (1999): "The Dynamic Effects of Health on the Labor Force Transitions of Older Workers," *Labour Economics*, 6, 179-202.

Blundell, R., P-A. Chiappori, T. Magnac, and C. Meghir (1998): "Collective Labor Supply: Heterogeneity and Nonparticipation," Unpublished manuscript, University College London.

Currie, J., and J. Gruber (1996): "Health Insurance Eligibility, Utilization of Medical Care, and Child Health," *Quarterly Journal of Economics*, 111, 431-466.

Dwyer, D. and O. Mitchell (1999): "Health Problems as Determinants of Retirement: Are Self-Rated Measures Endogenous?" *Journal of Health Economics*, 18,173-193.

Gruber, J., and B.C. Madrian (1995): "Health Insurance Availability and the Retirement Decision," *American Economic Review*, 85, 938-48.

Gustman, A., and T. Steinmeier (1986): "A Structural Retirement Model," *Econometrica*, 54, 555-584.

_____ (1994): "Employer-Provided Health Insurance and Retirement Behavior," *Industrial and Labor Relations Review*, 48, 124-140.

______ (2000): “Retirement in a Family Context: A Structural Model for Husbands and Wives.” *Journal of Labor Economics*, 18, 503-545.

Hurd, M.D. (1989): “Mortality Risk and Bequests,” *Econometrica*, 57, 779-814.

______ (1990): “The Joint Retirement Decisions of Husbands and Wives,” in *Issues in the Economics of Aging*, ed. by D.A. Wise. Chicago: University of Chicago Press.

Keane, M., and K. Wolpin. (1994): “The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence,” *Review of Economics and Statistics*, 77, 648-672.

Lumsdaine, R.L., J.H. Stock, and D.A. Wise (1994): “Pension Plan Provisions and Retirement: Men and Women, Medicare and Models,” in *Studies in the Economics of Aging*, ed. by D.A. Wise. Chicago: University of Chicago Press.

Madrian, B.C. (1994): “The Effect of Health Insurance on Retirement,” *Brookings Papers on Economic Activity*, 181-232.

Madrian, B.C., and N.D. Beaulieu (1998): “Does Medicare Eligibility Affect Retirement?” in *Inquiries in the Economics of Aging*, ed. by D.A. Wise. Chicago: University of Chicago Press.

Manning, W.G., Jr. (1988) “Health Insurance and the Demand for Medical Care - Evidence from a Randomized Experiment,” Rand Corporation Report: R-3476-HHS, Santa Monica, CA.

Rust, J., and C. Phelan (1997): “How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets,” *Econometrica*, 65, 781-831.

Starr-McCluer, M. (1996): “Health Insurance and Precautionary Savings,” *American Economic Review*, 86, 285-95.

Stock, J., and D.A. Wise (1990): “Pensions, the Option Value of Work, and Retirement,” *Econometrica*, 58, 1151-80.

Todd, P., and K.I. Wolpin (2003): “Using Experimental Data to Validate a Dynamic Behavioral Model of Child Schooling and Fertility: Assessing the Impact of a School Subsidy Program in Mexico,”Unpublished manuscript, University of Pennsylvania.

Van der Klaauw, H.W., and K.I. Wolpin (2002): “Social Security, Pensions, and the Savings and Retirement Behavior of Households,” Unpublished manuscript, University of North Carolina at Chapel Hill.

Zweimüller, J., R. Winter-Ebmer, and J. Falkinger (1996): “Retirement of Spouses and Social Security Reform,” *European Economic Review*, 40, 449-472.

Table I: Employment Status and Transitions by Retiree Health Insurance Status

	No controls	With controls
A. Percentage of Husbands Employed		
Own-employer HI, no RHI	73.8	74.3
Own-employer HI, with RHI	60.9	69.9
Difference	**12.9***	**4.4***
B. Percentage of Wives Employed		
Own-employer HI, no RHI	75.1	75.9
Own-employer HI, with RHI	42.8	58.8
Difference	**32.3***	**17.1***
C. Exit rate from employment: Percentage of Husbands *not* employed in period t, given employed at t-1		
Own-employer HI, with RHI	15.5	14.4
Own-employer HI, no RHI	11.4	14.0
Difference	**4.1**	**0.4**
D. Exit rate from employment: Percentage of Wives *not* employed in period t, given employed at t-1		
Own-employer HI, with RHI	14.6	8.2
Own-employer HI, no RHI	7.7	5.2
Difference	**6.9***	**3.0**

Source: Logit estimates and simulations using waves 1-4 of the HRS.

Note: HI = Health Insurance, RHI = Retiree Health Insurance. Sample size is 1,752 couples. The specification in column 1 includes only binary indicators for health insurance coverage of both spouses. The additional control variables included in the specification reported in the second column are (for both spouses) age, health, non-wage income, and earnings; the Social Security Primary Insurance Amount at three different levels of potential experience, and the pension benefit at five different levels of potential job tenure, for the spouse whose employment behavior is the dependent variable.

* indicates the difference is significantly different from zero at the 5% level.

Table II: Employment Status by Wave (Percent Distribution)

Husband	Wave 1 (1992)	Wave 2 (1994)	Wave 3 (1996)	Wave 4 (1998)
Not employed	23.3	29.6	35.2	41.4
Employed: New job	8.0	11.3	26.1	25.6
Employed: Same job	68.6	59.1	38.7	33.0
Wife				
Not employed	36.5	39.4	45.3	47.4
Employed: New job	8.9	11.1	16.6	17.5
Employed: Same job	54.8	49.5	38.1	35.0
Joint Distribution				
Husband: Not Employed				
Wife: Not employed	12.3	16.1	21.2	24.9
Wife: Employed, New job	1.4	2.3	3.4	5.3
Wife: Employed, Same job	9.6	11.1	10.2	11.0
Husband: Employed, new job				
Wife: Not employed	2.6	3.4	10.3	9.9
Wife: Employed, New job	1.1	2.1	5.7	6.2
Wife: Employed, Same job	4.3	5.7	10.1	9.2
Husband: Employed, same job				
Wife: Not employed	21.5	19.8	13.8	12.3
Wife: Employed, New job	6.2	6.4	6.9	5.5
Wife: Employed, Same job	40.9	33.0	18.5	15.6
Sample size	1,752	1,569	1,517	1,419

Table III: Distribution of Total Medical Expenditure

	Medical expenditure category (000 $)					
	0 to 1.99		**2 to 14.99**		**15+**	
	Prob-ability	Mean expenditure	Prob-ability	Mean expenditure	Prob-ability	Mean expenditure
Men, bad health						
Age 51	0.70	0.6	0.24	6.2	0.06	30.0
Age 61	0.62	0.7	0.28	5.7	0.09	38.9
Age 71	0.54	0.7	0.33	5.3	0.13	47.7
Men, good health						
Age 51	0.88	0.4	0.10	4.9	0.02	32.8
Age 61	0.82	0.5	0.15	5.1	0.03	32.8
Age 71	0.75	0.6	0.22	5.4	0.05	32.8
Women, bad health						
Age 51	0.67	0.7	0.27	5.7	0.06	24.5
Age 61	0.60	0.8	0.32	5.5	0.07	28.8
Age 71	0.52	0.8	0.39	5.3	0.09	33.1
Women good health						
Age 51	0.83	0.5	0.15	4.7	0.01	26.9
Age 61	0.81	0.6	0.17	4.6	0.02	30.3
Age 71	0.78	0.6	0.19	4.5	0.02	33.7

Source: Tabulations from the 1987 National Medical Expenditure Survey.

Table IV: Health Transitions and Health Status (Percent Distribution)

	Husband			Wife		
wave *t*+1→ wave *t* ↓	To good health	To bad health	To death	To good health	To bad health	To death
From good health						
wave 1 to 2	90.4	8.5	1.1	92.5	7.1	0.4
wave 2 to 3	89.7	9.0	1.3	93.4	6.1	0.5
wave 3 to 4	86.1	11.9	2.1	89.2	10.5	0.3
From bad health						
wave 1 to 2	24.1	68.7	7.3	29.9	66.1	3.9
wave 2 to 3	34.5	57.9	7.6	29.5	66.9	3.5
wave 3 to 4	22.2	68.1	9.7	22.7	72.1	5.2
Percent in good health (of those alive)						
Wave 1	81.4			84.6		
Wave 2	79.8			83.9		
Wave 3	80.8			83.9		
Wave 4	76.4			79.1		

Table V: Wave 1 Health Insurance (Percent Distribution)

	Husband		Wife	
	Primary	Secondary	Primary	Secondary
None	10.0	80.0	11.8	80.4
Own employer with RHI	53.5	1.9	23.8	0.8
Own employer without RHI	11.3	1.3	6.3	1.1
Spouse employer with RHI	11.0	5.0	40.7	8.8
Spouse employer without RHI	3.6	0.9	8.4	1.8
Private	7.4	8.5	7.4	5.8
Medicare	3.2	2.4	1.6	1.3

Note: RHI = Retiree Health Insurance. Sample size is 1,752.

Table VI: Utility Function Parameter Estimates

Consumption				α_1	2.144 (.004)
Intercepts					
Hus. health	Wife health	Hus. empl.	Wife empl.		
good	good	no	no	α_{00000}	0.000[a]
good	good	no	yes	α_{00001}	-1.060 (.011)
good	good	yes	no	α_{00010}	1.754 (.017)
good	good	yes	yes	α_{00011}	0.504 (.005)
good	bad	no	no	α_{00100}	-1.014 (.012)
good	bad	no	yes	α_{00101}	-2.079 (.023)
good	bad	yes	no	α_{00110}	0.964 (.009)
good	bad	yes	yes	α_{00111}	-0.748 (.009)
bad	good	no	no	α_{01000}	0.606 (.007)
bad	good	no	yes	α_{01001}	0.680 (.006)
bad	good	yes	no	α_{01010}	2.627 (.024)
bad	good	yes	yes	α_{01011}	-0.019 (.0002)
bad	bad	no	no	α_{01100}	0.748 (.008)
bad	bad	no	yes	α_{01101}	-0.669 (.009)
bad	bad	yes	no	α_{01110}	0.996 (.010)
bad	bad	yes	yes	α_{01111}	-2.054 (.017)

Employment effects[b]		
Change job, good health	α_{20},α_{40}	-1.241 (.009)
Change job, bad health	α_{21},α_{41}	-2.146 (.019)
Exit empl, good health	α_{30},α_{50}	-1.857 (.016)
Exit empl, bad health	α_{31},α_{51}	-2.411 (.021)
Both remain not employed	α_6	0.988 (.012)
Both remain employed	α_7	1.160 (.012)
Age effects[c]		
Hus., good health, not empl.	α_{800}	1.377 (.011)
Hus., good health, empl.	α_{801}	0.411 (.005)
Hus., bad health, not empl.	α_{810}	0.482 (.005)
Hus., bad health, empl.	α_{811}	-0.709 (.008)
Wife, good health, not empl.	α_{900}	0.819 (.008)
Wife, good health, empl.	α_{901}	1.218 (.011)
Wife, bad health, not empl.	α_{910}	0.574 (.006)
Wife, bad health, empl.	α_{911}	0.630 (.008)
Consumption ≤ 0	α_{10}	0.644 (.007)

Notes: Asymptotic standard errors estimated by the BHHH method are in parentheses. Log Likelihood = -32,803.5

a. Fixed.

b. Constrained to be equal for husbands and wives.

c. Age is measured as (age-41)/10.

Table VII: Model Fit, Overall and by Employment Status

	Husband						Wife					
A. Employment												
Employed in previous period	Not empl.		New job		Old job		Not empl.		New job		Old job	
	A	P	A	P	A	P	A	P	A	P	A	P
All	.329	.313	.135	.105	.535	.583	.427	.417	.103	.091	.470	.492
No	.785	.843	.215	.157			.840	.889	.160	.111		
Yes	.138	.090	.102	.083	.760	.827	.135	.084	.063	.084	.802	.839
B. Joint Employment												
Employed in previous period	H olf W olf		H olf W nj		H olf W oj		H nj W olf		H nj W nj		H nj W oj	
	A	P	A	P	A	P	A	P	A	P	A	P
All	.192	.176	.026	.023	.112	.114	.052	.037	.024	.011	.059	.057
Both spouses	.029	.024	.006	.002	.080	.048	.014	.004	.017	.008	.064	.075
Neither spouse	.764	.817	.074	.080			.113	.090	.049	.013		
	H oj W olf		H oj W nj		H oj W oj							
All	.183	.205	.054	.034	.299	.311						
Both spouses	.079	.048	.040	.072	.670	.720						
Neither spouse												

Notes: A = actual; P = predicted; H = husband; W = wife; olf = out of the labor force (not empl); nj = new job; oj = old job.

Table VIII: Employment Fit by Health Insurance Category

	Not employed		New job		Old job	
	Actual	Predicted	Actual	Predicted	Actual	Predicted
Husband's Employment Choice						
Husband's health insurance						
Own employer, RHI	.320	.285	.099	.103	.581	.612
Own employer, no RHI	.123	.159	.122	.113	.755	.728
Wife's health insurance						
Spouse employer, RHI	.346	.304	.091	.099	.563	.596
Spouse employer, no RHI	.149	.173	.070	.100	.782	.728
Wife's Employment Choice						
Wife's health insurance						
Own employer, RHI	.214	.220	.077	.088	.710	.691
Own employer, no RHI	.098	.146	.085	.084	.817	.769
Husband's health insurance						
Spouse employer, RHI	.240	.246	.057	.084	.703	.671
Spouse employer, no RHI	.079	.148	.067	.082	.854	.771

Note: RHI = Retiree Health Insurance

Table IX
Employment Fit by Health Insurance Conditional on Previous Period Employment Status

	Not employed		New job		Old job	
	Actual	Predicted	Actual	Predicted	Actual	Predicted
Husband	Not employed in previous period					
Own employer HI, RHI	.800	.845	.200	.155		
Own employer HI, no RHI.	.578	.689	.422	.312		
	Employed in previous period					
Own employer, RHI	.143	.078	.061	.084	.796	.838
Own employer, no RHI.	.059	.085	.080	.085	.861	.830
Wife	Not employed in previous period					
Own employer, RHI	.750	.856	.250	.144		
Own employer, no RHI.	.639	.814	.361	.186		
	Employed in previous period					
Own employer, RHI	.098	.084	.040	.076	.862	.840
Own employer, no RHI	.042	.078	.057	.074	.901	.848

Note: HI = Health Insurance; RHI = Retiree health insurance.

Table X: Simulated Probability of Choosing Non-employment

Husband	1. Base line	2. No Health Insurance	3. Universal Health Insurance	4. Add RHI to employer plans	5. Eliminate RHI from employer plans	6. Medi-care age = 67
All	.333	.301	.336	.340	.306	.331
Employed at t-1	.099	.094	.101	.101	.093	.099
HI = Own-Employer-provided with RHI	.322	.285	.330	.327	.275	.318
HI = Own-Employer-provided, no RHI	.207	.198	.236	.237	.196	.205
HI = Own-Employer-provided with RHI, employed at t-1	.100	.093	.104	.102	.089	.100
HI = Own-Employer-provided, no RHI, employed at t-1	.082	.079	.087	.086	.079	.082
Wife						
All	.409	.384	.420	.416	.391	.411
Employed at t-1	.064	.061	.066	.065	.060	.066
HI = Own-Employer-provided with RHI	.283	.256	.279	.284	.241	.285
HI = Own-Employer-provided, no RHI	.219	.226	.279	.280	.210	.209
HI = Own-Employer-provided with RHI, employed at t-1	.063	.048	.061	.063	.056	.064
HI = Own-Employer-provided, no RHI, employed at t-1	.061	.053	.070	.069	.056	.059

Note: RHI = Retiree Health Insurance

Table XI: Sample Selection

1. Number of married couples at wave 1	4,704
2. Age eligible	3,005
3. No missing data on key variables from the HRS survey	2,726
4. With Social Security records for both spouses	1,752
5. With pension records for both spouses, or no pension coverage	1,752

Note: Couples are age eligible if the husband is 51-63 at the 1992 survey, his wife is not older than him, and his wife is at most ten years younger than him. A few additional cases in which the wife was a lot younger than the husband (though less than 10 years) were deleted because there were no other couples with similar age patterns. This was done to make the programming easier.

Table XII
Sample Characteristics

	Full Sample		Estimation Sample	
	Husband	Wife	Husband	Wife
Age in 1992	57.7	53.5	56.7	53.2
Education	12.1	12.2	12.4	12.4
Black	.12	.12	.09	.09
Hispanic	.08	.09	.06	.06
Employer health insurance (EHI)	.74	.74	.79	.79
Retiree coverage, of those with EHI	.80	.79	.81	.81
Employed at the 1992 survey date	.73	.61	.77	.63
Hourly wage in 1992 (if employed)	18.0	11.5	14.2	11.0
Pension in 1992	.48	.28	.53	.29
Good health in 1992	.79	.82	.81	.85
Died, divorced, or attrited by 1998	.30	.23	.25	.19
Net worth in 1992 ($000)	266		275	
Sample size	4,704		1,752	

Table XIII
Pension Characteristics

Wave 1 Job	Husband	Wife
Youngest age at which benefits could be collected	57.5	56.3
Age at which normal benefits could be collected	60.5	60.1
Annual Benefit		
If exit job in 1992	13,228	6,475
If exit job in 1997	15,164	7,745
If exit job in 2002	17,939	9,656
If exit job in 2007	20,428	11,236
Sample size	574	419
Previous Job		
Youngest age at which benefits could be collected	56.3	58.6
Annual Benefit	11,886	6,234
Sample Size	514	157

Notes: Dollar values are measured in 1992 dollars.

Table XIV
Social Security Monthly Primary Insurance Amount (PIA)

	Husband	Wife
Initial PIA	806	356
PIA after 5 additional years of work	857	416
PIA after 10 additional years of work	895	465
PIA after 15 additional years of work	927	511

Notes: Initial PIA is the value computed from the Social Security Earnings Record file. The sample in each row includes only those individuals who are 70 or less after the indicated number of additional years of work. Dollar values are measured in 1992 dollars. Sample size for the initial PIA is 1,752.

Table XV: Health Transition and Other Parameters

Health transition parameters							
Hus, good → good:	intercept	γ_{000}	6.11 (0.68)	Wife, good → good:	intercept	γ_{000}	6.21 (0.86)
	Age	γ_{100}	-1.10 (0.34)		Age	γ_{100}	-0.56 (0.54)
Hus, good → bad:	intercept	γ_{001}	3.12 (0.71)	Wife, good → bad:	intercept	γ_{001}	3.31 (0.87)
	Age	γ_{101}	-0.67 (0.36)		Age	γ_{101}	-0.25 (0.55)
Hus, bad → good:	intercept	γ_{010}	2.84 (0.66)	Wife, bad → good:	intercept	γ_{010}	1.75 (0.65)
	Age	γ_{110}	-0.88 (0.34)		Age	γ_{110}	0.08 (0.42)
Hus, bad → bad:	intercept	γ_{011}	3.58 (0.62)	Wife, bad → bad:	intercept	γ_{011}	2.90 (0.62)
	Age	γ_{111}	-0.81 (0.32)		Age	γ_{111}	-0.08 (0.41)
Other Parameters							
Terminal Value Function		ς	2.725 (.029)				
Probability of Divorce		ζ	-4.314[a]				
Probability of Layoff		δ	-3.892[a]				
Discount Factor		β	.925[a]				

Notes: a. Fixed. Probability of divorce is $e^{\zeta}/(1+e^{\zeta})$. Probability of layoff is $e^{\delta}/(1+e^{\delta})$. Terminal value function is $-e^{\varsigma}$. Age is measured as (age-41)/10. Standard errors are in parentheses.

Figure 1: Predicted versus Actual Non-employment rate by Age

Men

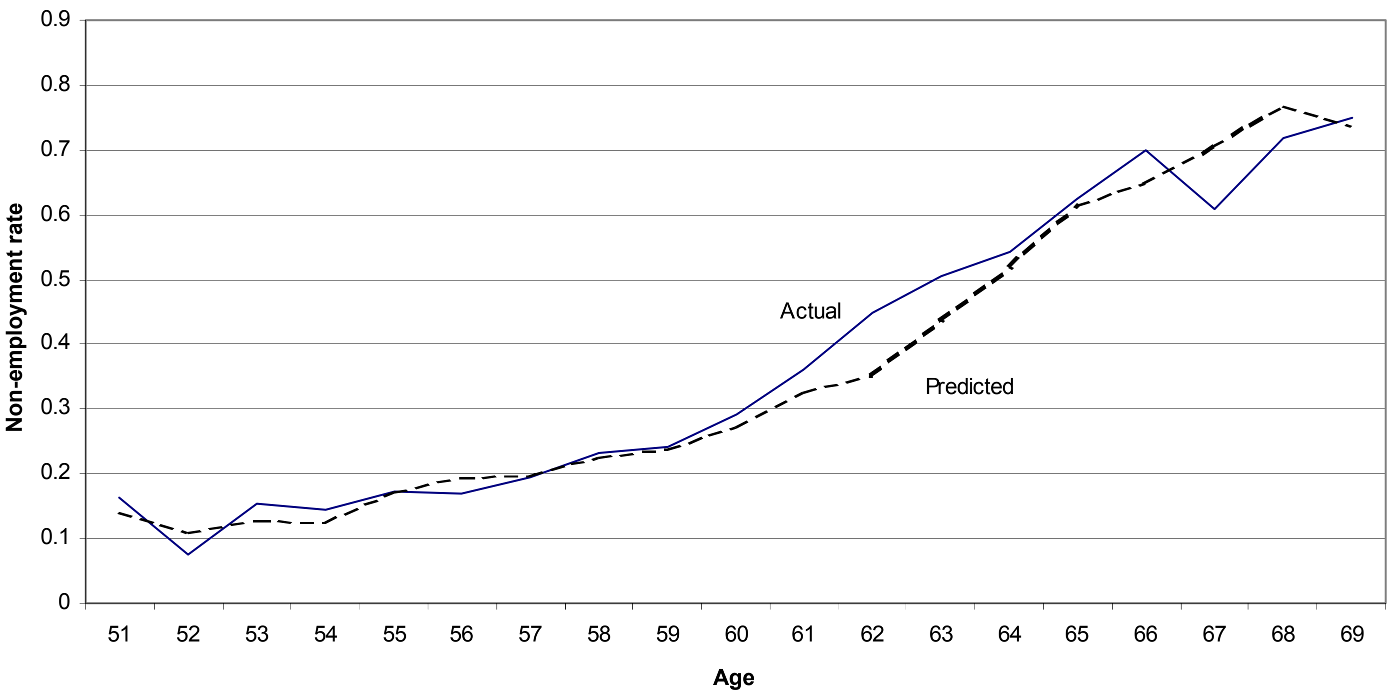

Women

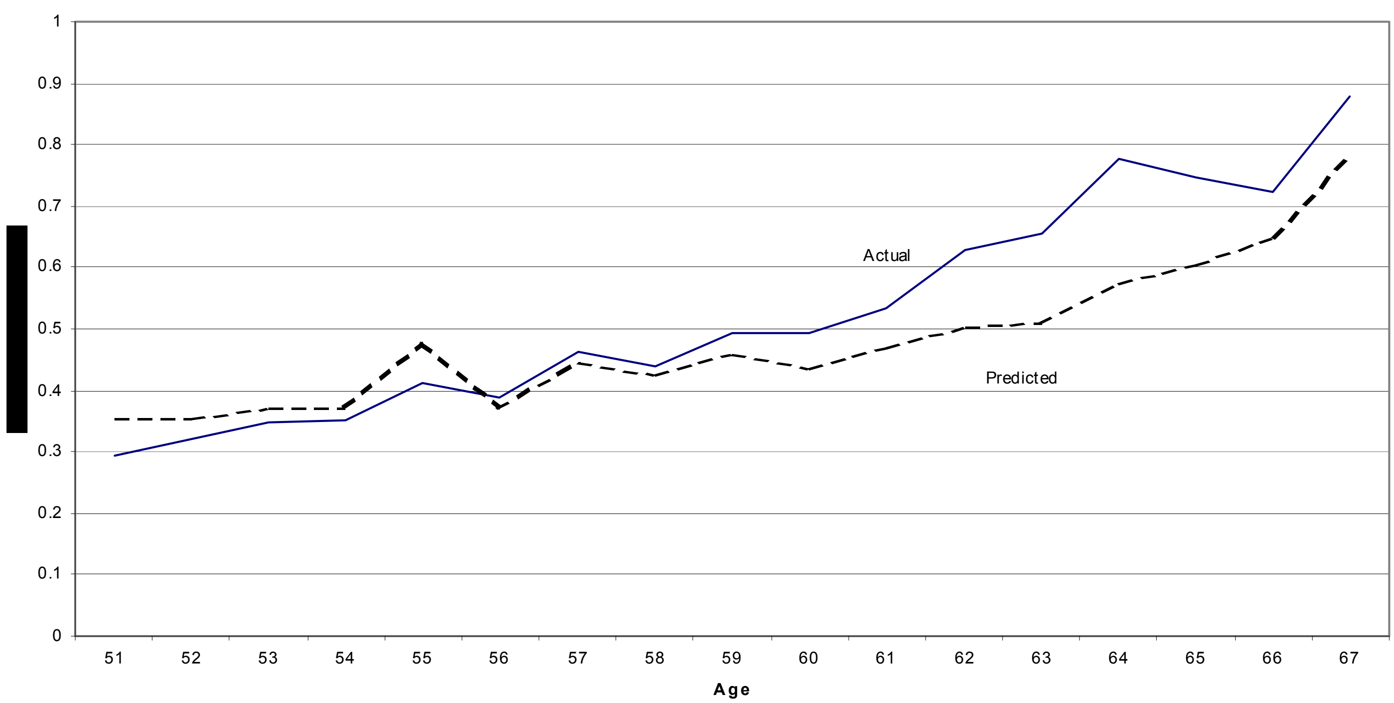

Figure 2: Simulated Impact on Non-employment of Changing the Age of Medicare Eligibility from 65 to 67: Individuals with Own-Employer Health Insurance Without Retiree Benefits

Men

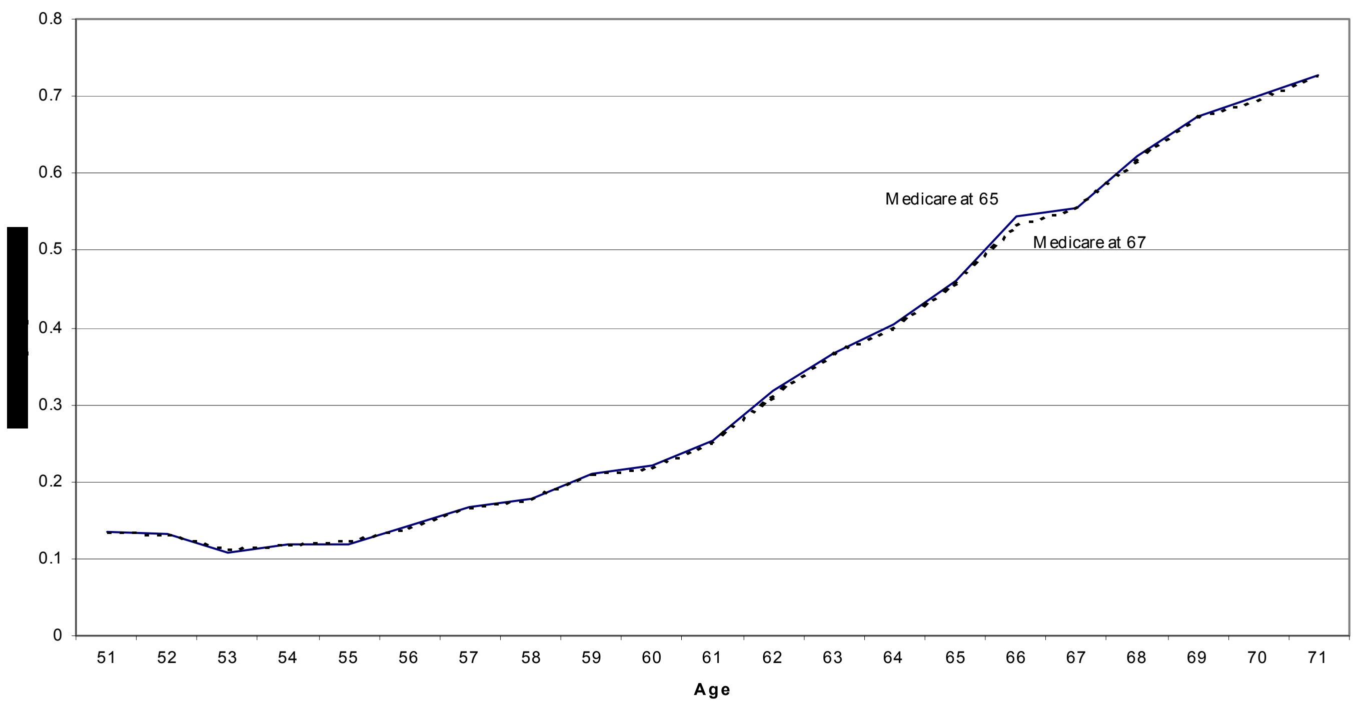

Women