THE EFFECTS OF HEALTH INSURANCE AND SELF-INSURANCE ON RETIREMENT BEHAVIOR

Eric French*
John Bailey Jones

About the Sandell Grant Program

About the Center for Retirement Research

The *Center for Retirement Research at Boston College*, part of a consortium that includes a parallel centers at the University of Michigan and the National Bureau of Economic Research, was established in 1998 through a grant from the Social Security Administration. The goals of the Center are to promote research on retirement issues, to transmit new findings to the policy community and the public, to help train new scholars, and to broaden access to valuable data sources. Through these initiatives, the Center hopes to forge a strong link between the academic and policy communities around an issue of critical importance to the nation's future.

Center for Retirement Research at Boston College
550 Fulton Hall
140 Commonwealth Ave.
Chestnut Hill, MA 02467
phone: 617-552-1762 fax: 617-552-1750
e-mail: crr@bc.edu
http://www.bc.edu/crr

Affiliated Institutions:

American Enterprise Institute
The Brookings Institution
Massachusetts Institute of Technology
Syracuse University
Urban Institute

Abstract

Using the first estimable dynamic programming model of retirement behavior that accounts for both savings and uncertain medical expenses, we assess the importance of employer-provided health insurance and Medicare in determining retirement behavior. Including both of these features allows us to determine whether workers value employer-provided health insurance because the subsidy contained in the insurance lowers their average medical expenses, or because health insurance also reduces their medical expense risk. Using data from the Health and Retirement Study, we find that the reduction in expected medical expenses explains about 60% of a typical individual's valuation of health insurance, with the reduction in volatility explaining the remaining 40%. We find that for workers whose insurance is tied to their job, shifting the Medicare eligibility age to 67 will significantly delay retirement. However, we find that the plan to shift the Social Security normal retirement age to 67 will cause an even larger delay.

1   Introduction

One of the most important social programs for the rapidly growing elderly population is Medicare; in 2002, Medicare had 41 million beneficiaries and \$266 billion of expenditures.[1] Medicare provides health insurance to individuals that are 65 or older. Prior to receiving Medicare, many individuals receive health insurance only if they continue to work. An important question, therefore, is whether Medicare significantly affects the labor supply of the elderly, especially around age 65. This question is particularly important to those considering changes to the Medicare eligibility age; the fiscal impact of such changes depends critically on their labor supply effects.

Several studies have developed structural models that can be used for such policy experiments. These studies of retirement behavior, however, have arrived at very different conclusions about the importance of Medicare. The different conclusions seem to result from differences in how the studies treat market incompleteness and uncertainty, which affect how much individuals value Medicare. In this paper, we construct and estimate a structural retirement model that includes not only medical expense risk and risk-reducing health insurance, but also a saving decision that allows workers to self-insure through asset accumulation. Including both of these features—to our knowledge, ours is the first paper to do so—yields a more general model that can reconcile the earlier results.[2]

Assuming that individuals value health insurance at the cost paid by employers, both Lumsdaine et al. (1994) and Gustman and Steinmeier (1994) find that health insurance has a small effect on retirement behavior. One possible reason for their results is that the average employer contribution to health insurance is relatively modest—Gustman and Steinmeier (1994) find that the average employer contribution to employee health insurance is about \$2,500 per year before age 65—and it declines by a relatively small amount after age 65.[3] In short, if health insurance is valued at the cost paid for by employers, the work disincentives

---

[1] Figures taken from 2003 Medicare Annual Report (The Boards of Trustees of the Hospital Insurance and Supplementary Medical Insurance Trust Funds, 2003).

[2] van der Klaauw and Wolpin (2002) and Rust, et al. (2003) are currently engaged in similar projects.

[3] Data are from the 1977 NMES, adjusted to 1998 dollars with the medical component of the CPI.

of Medicare are fairly small.

If individuals are risk-averse, however, and large out-of-pocket Medical costs are possible, individuals could value health insurance well beyond the cost paid by employers. If individuals are uninsured, they could face volatile medical expenses, which in turn could lead to volatile life-cycle consumption paths. If individuals are risk-averse, they will value the consumption smoothing that health insurance provides. Therefore, Medicare's age-65 work disincentive comes not only from the reduction in average medical costs paid by those without employer-provided health insurance, but from also the reduction in the volatility of those costs.[4]

Addressing this point, Rust and Phelan (1997) estimate a dynamic programming model that accounts explicitly for risk aversion and uncertainty about out-of-pocket medical expenses. They find that because of health cost uncertainty, Medicare has large effects on retirement behavior. Using newer and more inclusive data, Blau and Gilleskie (2003) find similar, though smaller, effects. Rust and Phelan and Blau and Gilleskie, however, all assume that an individual's consumption equals his income net of out-of-pocket medical expenses. In other words, these studies ignore an individual's ability to self-insure against out-of-pocket medical expenses through saving. Several empirical results suggest that savings might be important. Smith (1999) finds that out-of-pocket medical expenses generate large declines in wealth. Cochrane (1991) finds that short-term illnesses generate only small declines in food consumption. To the extent that Rust and Phelan and Blau and Gilleskie overstate the consumption volatility caused by out-of-pocket medical cost volatility, they overstate the value of health insurance, and thus the effect of health insurance on retirement.

Lumsdaine et al. (1994) and Gustman and Steinmeier (1994) potentially underestimate the effects of Medicare, while Rust and Phelan (1997) and Blau and Gilleskie (2003) potentially overestimate it. A major goal of this paper, therefore, is to reconcile these results by using a more general model of retirement behavior. In particular, we construct a life-cycle model of labor supply that not only accounts for health cost uncertainty and health insurance,

---

[4]While individuals can usually buy private health insurance, high administrative costs and adverse selection problems can make it prohibitively expensive. Moreover, private coverage often does not cover pre-existing medical conditions, whereas employer-provided coverage typically does.

but also has a savings decision. Moreover, we include the coverage provided by means-tested social insurance, by assuming that the government guarantees each individual a minimum level of consumption. All of this allows us to consider whether uncertainty and self-insurance greatly affects the value of health insurance. We also model two other important sources of retirement incentives, Social Security and private pensions, in some detail. Although Medicare, Social Security and pensions often generate contemporaneous incentives, our approach allows us to disentangle their effects.

Estimating the model by the Method of Simulated Moments, we find that the model fits the data well with reasonable parameter values. The model predicts that workers whose health insurance is tied to their job leave the labor force about 0.47 years later than workers whose coverage extends into retirement. This result, being consistent with several reduced-form estimates, also supports the model.

Next, we measure the changes in labor supply induced by raising the Medicare eligibility age to 67 and by raising the normal Social Security retirement age to 67. We find that shifting the Medicare eligibility age to 67 will significantly increase the cumulative labor force participation of workers whose insurance is tied to their job. We also find, however, that the incremental effect of raising the Social Security retirement age is even bigger, even for workers whose insurance is tied to their jobs. In order to understand why Social Security is more important, we evaluate how much individuals value health insurance. We find that around 60% of the value of health insurance comes from the reduction in average medical expenses, with the remaining 40% coming from the reduction in medical expense uncertainty.

We then re-estimate the model with saving prohibited. We find that eliminating the ability to self-insure through saving significantly increases the effects of Medicare and the value of health insurance. We also find, however, that this restricted model provides a worse fit to the data along several key dimensions. These results suggest that self-insurance significantly reduces the effects of health cost uncertainty, so that the effects of Medicare are modest.

The rest of paper proceeds as follows. Section 2 develops our dynamic programming model of retirement behavior. Section 3 describes how we estimate the model using the Method of

Simulated Moments. Section 4 describes the Health and Retirement Study (HRS) data that we use in our analysis. Section 5 presents life cycle profiles drawn from these data. Section 6 contains preference parameter estimates for the structural model and policy experiments. In Section 7 we assess how allowing workers to save affects our results. In Section 8 we consider a few important robustness checks. Section 9 concludes.

2 The Model

2.1 Preferences and Demographics

Consider a household head seeking to maximize his expected lifetime utility at age (or year) t, $t = 1, 2, ...$. Each period that he lives, the individual derives utility, U_t, from consumption, C_t, hours worked, H_t, and health (or medical) status, $M_t \in \{good, bad\}$, so that $U_t = U(C_t, H_t, M_t)$. When he dies, he values bequests of assets, A_t, according to the bequest function $b(A_t)$. Let s_t denote the probability of being alive at age t conditional on being alive at age $t - 1$, and let $S(j, t) = (1/s_t) \prod_{k=t}^{j} s_k$ denote the probability of living to age $j \geq t$, conditional on being alive at age t. Let $T = 95$ denote the terminal period, so that $s_{T+1} = 0$. The parameter β is the time discount factor. We assume that individuals maximize

$$U(C_t, H_t, M_t) + E_t \left(\sum_{j=t+1}^{T+1} \beta^{j-t} S(j-1, t) \left[s_j U(C_j, H_j, M_j) + (1 - s_j) b(A_j) \right] \right), \quad (1)$$

by choosing the contingency plans $\{C_j, H_j, B_j\}_{j=t}^{T+1}$, subject to the constraints on $\{C_j, H_j, B_j\}$ described below. In addition to choosing hours and consumption, eligible individuals can choose whether to apply for Social Security benefits; let the indicator variable $B_t \in \{0, 1\}$ equal one if the individual has applied for benefits.

The individual's within-period utility function is of the form

$$U(C_t, H_t, M_t) = \frac{1}{1 - \nu} \left(C_t^\gamma (L - H_t - \theta_P P_t - \phi \times 1\{M_t = bad\})^{1-\gamma} \right)^{1-\nu}, \quad (2)$$

where the total time endowment per year is L and the quantity of leisure consumed is $L -$

$H_t - \theta_P P_t - \phi \times 1\{M_t = bad\}$. The individual's utility from leisure depends on his health status through the 0-1 indicator $1\{M_t = bad\}$, which equals one when his health is bad. Participation in the labor force is denoted by P_t, a 0-1 indicator equal to zero when hours worked, H_t, equals zero. The fixed cost of work, θ_P, is measured in hours worked per year. Including fixed costs allows us to capture the empirical regularity that annual hours of work are clustered around 2000 hours and 0 hours (Cogan, 1981).[5] We treat retirement as a form of the participation decision, and thus allow retired workers to reenter the labor force; as stressed by Rust and Phelan (1997) and Ruhm (1990), reverse retirement is a common phenomenon.

The parameter γ is between 0 and 1 if utility is increasing in leisure and consumption. The parameter ν, the coefficient of relative risk aversion for total utility, is positive if individuals are risk averse. ν has two purposes. First, as ν increases individuals become less willing to substitute consumption and leisure across time. Second, ν measures the non-separability between consumption and leisure. Under perfect foresight and interiority, $\nu > 1$ implies that consumption and leisure are Frisch substitutes (Low, 2003). French (2003) shows that with his estimates of γ and ν, a model with this non-separable specification can replicate the consumption declines that are observed at retirement.

The bequest function takes the form

$$b(A_t) = \theta_B \frac{(\max\{A_t, 0\} + K)^{(1-\nu)\gamma}}{1 - \nu}, \tag{3}$$

where K is a constant that affects the strength of the bequest motive across wealth levels. In particular, as K grows, the marginal utility of bequests for poor, small-bequest individuals decreases, absolutely and relative to the marginal utility of bequests for the rich. The $\max\{A_t, 0\}$ term appears because debts cannot be bequeathed.

The individual's utility depends on two random demographic variables. One is the indi-

---

[5]Outside the mass points at 0 and 2000 hours, work hours in our HRS data are more or less uniformly distributed between 750 and 3300 hours (also see Rust and Phelan, 1997). We include fixed costs, rather than discretize the choice set for hours (into, say, full-time, half-time and none), because it provides a better way to capture this dispersion.

vidual's health status, M_t, which follows an exogenous Markov process. We assume that the transition probabilities for health status depend on the individual's current health status and age, so that the elements of the health status transition matrix are

$$\pi_{ij}(t) = \Pr(M_{t+1} = j | M_t = i, t), \quad i, j \in \{good, bad\}. \tag{4}$$

A second source of uncertainty is mortality. Mortality rates depend upon age and previous health status:

$$s_{t+1} = s(M_t, t+1). \tag{5}$$

2.2 Budget Constraints

The individual has several sources of income: asset income, rA_t, where r denotes the constant pre-tax interest rate; labor income, $W_t H_t$, where W_t denotes wages; spousal income, ys_t; pension benefits, pb_t; Social Security benefits, ss_t; and government transfers, tr_t. The individual's income is allocated between: taxes; consumption; health care expenses, hc_t; and asset accumulation. This implies the following accumulation equation:

$$A_{t+1} = A_t + Y(rA_t + W_t H_t + ys_t + pb_t, \tau) + ss_t + tr_t - hc_t - C_t. \tag{6}$$

where post-tax income, $Y(rA_t + W_t H_t + ys_t + pb_t, \tau)$, is a function of taxable income and the vector τ, described in Appendix A, which captures the tax structure.

In addition to these "financial" assets (which include housing), the individual also accumulates pension and Social Security benefits, which we discuss in some detail below.

Associated with this budget rule is the borrowing constraint

$$A_t + Y_t + ss_t + tr_t - C_t \geq 0. \tag{7}$$

Because it is illegal to borrow against Social Security benefits and difficult to borrow against most forms of pension wealth, individuals with low non-pension, non-Social Security wealth

may not be able to finance their retirement before their Social Security benefits become available at age 62. It is worth noting that this borrowing constraint excludes medical expenses, which we assume are realized after labor decisions are made. We view this assumption as more reasonable than the alternative, namely that the time-t medical expense shocks are fully known when workers decide whether to hold on to their employer-provided health insurance.[6]

Following Hubbard et al. (1994, 1995), government transfers provide a consumption floor:

$$tr_t = max\{0, C_{min} - (A_t + Y_t + ss_t)\}, \tag{8}$$

Equation (8) implies that government transfers bridge the gap between an individual's "liquid resources" (the quantity in the inner parentheses) and the consumption floor. Equation (8) also implies that if transfers are positive, $C_t = C_{min}$. Our treatment of government transfers implies that individuals can always consume at least C_{min}, even if their out-of-pocket medical expenses have exceeded their financial resources. With the government effectively providing low-asset individuals with health insurance, these people may place a low value on employer-provided health insurance. This of course depends on the value of C_{min}; if C_{min} is low enough, it will be the low-asset individuals who value health insurance most highly. Those with very high asset levels should be able to self-insure.

2.3   Medical Expenses, Health Insurance, and Medicare

Medical expenses, hc_t, which are the focus of this paper, are defined as the sum of out-of-pocket costs and insurance premia. We assume that an individual's health costs depend upon: health insurance status, HI_t; health status, M_t; age, t; whether the person is working, P_t; and a person-specific effect ψ_t:

$$\ln hc_t = hc(M_t, HI_t, t, P_t) + \sigma(M_t, HI_t, t, P_t) \times \psi_t. \tag{9}$$

---

[6]Given the timing of medical expenses, under this borrowing constraint an individual with extremely high medical expenses this year could have negative net worth next year. Given that many people in our data still have unresolved medical expenses, medical expense debt seems reasonable.

Note that health insurance affects both the expectation of medical expenses, through $hc(.)$ and the variance, through $\sigma(.)$. These differences across health insurance types usually shrink at age 65, when Medicare becomes the primary insurer for most individuals.

Differences in labor supply behavior across categories of health insurance coverage, HI_t, are an important part of identifying our model. We assume that there are four mutually exclusive categories of health insurance coverage. The first is retiree coverage, ret, where workers keep their health insurance even after leaving their jobs.[7] The second category is tied health insurance, $tied$, where workers receive employer-provided coverage as long as they continue to work. If a worker with tied health insurance leaves his job, however, he enters the third category and receives "COBRA" coverage, $COBRA$, which allows him to purchase insurance at his employer's group rate. After one year of $COBRA$ coverage, the worker's insurance ceases.[8] The fourth category consists of individuals whose potential employers provide no health insurance at all, or $none$.[9] Workers move between these insurance categories according to

$$
HI_t = \begin{cases} ret & \text{if} \quad HI_{t-1} = ret \\ tied & \text{if} \quad HI_{t-1} = tied \quad \text{and} \quad H_t > 0 \\ COBRA & \text{if} \quad HI_{t-1} = tied \quad \text{and} \quad H_t = 0 \\ none & \text{if} \quad HI_{t-1} = none \quad \text{or} \quad HI_{t-1} = COBRA \end{cases} . \tag{10}
$$

In imposing this transition rule, we are assuming that people out of the work force are never offered jobs with insurance coverage, and that workers with $tied$ coverage never upgrade to ret coverage. Restricting access to insurance in this way most likely leads us to overstate the value of employer-provided health insurance.

---

[7]If they leave their job, however, their medical expenses may rise, as those with retiree coverage often pay for their insurance, albeit at lower group rates, after they retire.

[8]Although there is some variability across states as to how long individuals are eligible for employer-provided health insurance coverage, by Federal law most individuals are covered for 18 months (Gruber and Madrian, 1995). Given a model period of one year, we approximate the 18-month period as a one-year term.

[9]Workers in the $none$ category buy insurance on their own, receive some sort of government coverage, or simply go uncovered. For simplicity, we assume that the three groups share a common medical expense distribution.

An individual's medical expenses depend not only on his private insurance coverage, HI_t, but also on his access to Medicare health insurance. Almost all individuals that are 65 or older are eligible for Medicare, which supplements the insurance coverage described in equation (10).[10] In particular, individuals without employer-provided insurance can receive Medicare coverage once they turn 65.

Following Feenberg and Skinner (1994) and French and Jones (2004), we assume that the idiosyncratic component of medical expenses ψ_t can be decomposed as

$$\psi_t = \zeta_t + \xi_t, \quad \xi_t \sim N(0, \sigma_\xi^2), \tag{11}$$

$$\zeta_t = \rho_{hc}\zeta_{t-1} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma_\epsilon^2), \tag{12}$$

where ξ_t and ϵ_t are serially and mutually independent. ξ_t is the transitory component of health cost uncertainty, while ζ_t is the persistent component, with autocorrelation ρ_{hc}.

2.4   Wages and Spousal Income

We assume that the logarithm of wages at time t, $\ln W_t$, is a function of health status (M_t), age (t), hours worked (H_t) and an autoregressive component, ω_t:

$$\ln W_t = W(M_t, t) + \alpha \ln H_t + \omega_t, \tag{13}$$

The inclusion of hours, H_t, in the wage determination equation captures the empirical regularity that, all else equal, part-time workers earn relatively lower wages than full time workers. The autoregressive component ω_t has the correlation coefficient ρ_W and the normally-distributed innovation η_t:

$$\omega_t = \rho_W \omega_{t-1} + \eta_t, \quad \eta_t \sim N(0, \sigma_\eta^2). \tag{14}$$

---

[10]Individuals who have paid into the Medicare system for at least 10 years become eligible at age 65. A more detailed description of the Medicare eligibility rules is available at `http://www.medicare.gov/`.

Because spousal income can serve as insurance against medical shocks, we include it in the model. In the interest of computational simplicity, we assume that spousal income is a deterministic function of an individual's age and the exogenous component of his wages:

$$ys_t = ys(W(M_t, t) + \omega_t, t). \tag{15}$$

These features allow us to capture assortive mating and the age-earnings profile.

2.5 Pensions and Social Security

Because pensions and Social Security both generate potentially important retirement incentives, we model the two programs in detail.

Pension benefits, pb_t, are a function of the worker's age and pension wealth. Pension wealth in turn depends on pension accruals, which are themselves a function of a worker's age, labor income, and health insurance type. Computational concerns lead us to use a stylized pension accrual formula, which we construct from HRS data. The formula is for a weighted average of defined benefit, defined contribution, and combination plans. This pension accrual formula captures the fact that: high income workers have higher pension accrual rates than low income workers; accrual rates are higher for workers in their 50s than in other ages; and workers with tied and (especially) retiree coverage tend to have higher accrual rates than those without coverage. The last feature of our accrual formula is particularly important in isolating the effects of employer-provided health insurance.[11] When finding an individual's decision rules, we assume further that the individual's existing pension wealth is a function of his Social Security wealth, age, and health insurance type. Details are in Appendix B.

Individuals receive no Social Security benefits until they apply, i.e., $ss_t = 0$ until $B_t = 1$. Individuals can first apply for benefits at age 62. Upon applying the individual receives benefits until death, i.e., $B_{t+1} = 1$ if $B_t = 1$. Social Security benefits depend on his Average

---

[11]After controlling for health insurance type and other factors, pensions still contain a fair bit of idiosyncratic variation. The relevant issue, however, is not whether this omitted heterogeneity is important for retirement, but whether it significantly affects the role of health insurance. We proceed under the assumption that it does not.

Indexed Monthly Earnings ($AIME$), which is roughly his average income during his 35 highest earnings years in the labor market.

There are three major incentives provided by the Social Security System, each of which tend to induce exit from the labor market when old. First, while income earned by workers with less than 35 years of earnings automatically increases their $AIME$, income earned by workers with more than 35 years of earnings increases their $AIME$ only if it exceeds earnings in some previous year of work. Because Social Security benefits increase in $AIME$, this causes work incentives to drop after 35 years in the labor market. We describe the computation of $AIME$ in more detail in Appendix C.

Second, the age at which the individual applies for Social Security affects the level of benefits. Recall that individuals can first apply at age 62. For every year before age 65 the individual applies for benefits, benefits are reduced by 6.67% of the age-65 level. This is roughly actuarially fair. But for every year after age 65 that benefit application is delayed, benefits rise by 5.0% up until age 70. This is less than actuarially fair, and encourages people to apply for benefits by age 65.[12]

Third, the Social Security Earnings Test is imposed on beneficiaries younger than age 70. For individuals aged 62-64, each dollar of labor income above the "test" threshold of $9,120 leads to a 1/2 dollar decrease in Social Security benefits, until all benefits have been taxed away. For individuals aged 65-69, each dollar of labor income above a threshold of $14,500 leads to a 1/3 dollar decrease in Social Security benefits, until all benefits have been taxed away. Although benefits taxed away by the earnings test are credited to future benefits, after age 64 the crediting rate is less than actuarially fair, so that the Social Security Earnings Test effectively taxes the labor income of beneficiaries aged 65-69.[13] When combined with the aforementioned incentives to draw Social Security benefits by age 65, the Earnings Test

---

[12]We use tax and benefit formulas from the *Social Security Handbook Annual Statistical Supplement* for the year 1998. The Social Security crediting formula depends on the individual's year of birth, with the formulae for later birth years providing smaller incentives to retire at age 65. We use the formula for individuals born in 1932, who turned 65 in 1997.

[13]If a year's worth of benefits are taxed away between ages 62 and 64, benefits in the future are increased by 6.67%. If a year's worth of benefits are taxed away between ages 65 and 69, benefits in the future are increased by 5.0%.

discourages work after age 65. This incentive is incorporated in the calculation of ss_t, which is defined to be net of the earnings test.

2.6 Recursive Formulation

In recursive form, the individual's problem can be written as

$$
V_t(X_t) = \max_{C_t, H_t, B_t} \left\{ \frac{1}{1-\nu} \left(C_t^\gamma (L - H_t - \theta_P P_t - \phi \times 1\{M_t = bad\})^{1-\gamma} \right)^{1-\nu} + \beta(1 - s_{t+1})b(A_{t+1}) \right.
$$
$$
\left. + \beta s_{t+1} \int V_{t+1}(X_{t+1})dF(X_{t+1}|X_t, t, C_t, H_t, B_t) \right\}, \tag{16}
$$

subject to equations (7) and (8). The vector $X_t = (A_t, B_{t-1}, M_t, AIME_t, HI_t, \omega_t, \zeta_{t-1})$ contains the individual's state variables, while the function $F(\cdot|\cdot)$ gives the conditional distribution of these state variables.[14] In doing so, $F(\cdot|\cdot)$ incorporates the budget constraints and stochastic processes described in equations (4) through (15).

An individual's decisions thus depend on his state variables, X_t, his preferences, θ, and his beliefs, χ, where

$$
\theta = (\gamma, \nu, \theta_P, \theta_B, \phi, L, \beta),
$$
$$
\chi = \left(r, W(M_t, t), \alpha, \sigma_\eta^2, \rho_W, hc(M_t, HI_t, t, B_t, P_t), \sigma(M_t, HI_t, t, B_t), \sigma_\xi^2, \sigma_\epsilon^2, \rho_{hc}, \right.
$$
$$
\left. \{prob(M_{t+1}|M_t, t)\}_{t=1}^T, \{S_t\}_{t=1}^T, Y(\cdot, \cdot), \{ss_t\}_{t=1}^T, \{pb_t\}_{t=1}^T, \{tr_t\}_{t=1}^T \right).
$$

It follows that the solution to the individual's problem consists of the set of consumption $\{C_t(X_t, \theta, \chi)\}_{1 \le t \le T}$, work $\{H_t(X_t; \theta, \chi)\}_{1 \le t \le T}$ and benefit application $\{B_t(X_t; \theta, \chi)\}_{1 \le t \le T}$ rules that solve equation (16). The labor force participation rule $P_t(X_t; \theta, \chi)$ is a 0-1 indicator equal to zero when $H_t(X_t; \theta, \chi) = 0$. Inserting these decision rules into the asset accumulation equation yields next period's assets, $A_{t+1}(X_t, \psi_t; \theta, \chi)$.

Given that the model lacks a closed form solution, these decision rules are found numer-

---

[14]Spousal income and pension benefits (see Appendix B) depend only on the other state variables and are thus not state variables themselves.

14

ically using value function iteration. To reduce the computational burden, we assume that all workers retire and apply for Social Security benefits by age 70: for $t \geq 70$, $B_t = 1$ and $H_t = P_t = 0$. Appendix D describes our numerical methodology.

3   Estimation

Our goal is to estimate preferences, θ, and beliefs, χ. Computational concerns lead us to use a two-step strategy, similar to the ones used by Gourinchas and Parker (2002) and French (2003). In the first step we estimate some belief parameters and calibrate others. In doing this we assume that individuals have rational expectations, so that the belief parameters can be found by estimating the data generating process for the exogenous state variables. We describe the belief parameters in Section 4. In the second step we estimate preference parameters using the method of simulated moments (MSM). In the next two subsections, we describe our MSM methodology in more detail.

3.1   Moment Conditions

Because some of the variables in the state vector X_t are probably mismeasured, traditional estimators, such as maximum likelihood or non-linear least squares, are unlikely to be consistent. For example, wages are notoriously mis-measured in virtually all datasets.[15] Moreover, our data contain no measure of consumption.[16] Although measurement error can be incorporated into the standard maximum likelihood framework, doing so tends to be computationally costly. We instead estimate the model by the MSM, an approach that places fewer demands on the data.[17]

---

[15] Because we use earnings divided by hours as the wage measure, measurement error in hours affects both measured wages and measured hours, creating the well-known "division bias" problem.

[16] Rust and Phelan (1997) point out that one could impute consumption using asset accumulation equations and measures of assets and income, but imputed consumption is often negative, suggesting measurement error.

[17] Gourinchas and Parker (2002) develop this point in more detail in the context of a model of life-cycle consumption. A related approach that explicitly incorporates measurement error is to use simulated age-conditional likelihood functions, as in Keane and Wolpin (2001). By working with unconditional (or age-conditional) distributions, rather than the conditional (on X_t) distribution used in traditional likelihood estimation, Keane and Wolpin's approach avoids problems caused by measurement error in X_t, in much the same way as does the MSM.

The objective of MSM estimation is to find the preference vector $\hat{\theta}$ that yields simulated life-cycle decision profiles that "best match" (as measured by a GMM criterion function) the profiles from the data. A key part of the MSM approach is selecting which features of the data distribution—which profiles—to match. The model predicts that labor supply behavior should differ by age, health insurance status, health status, medical expenses, and asset level. We therefore require our model to match labor force participation conditional on asset grouping and health insurance status. Moreover, because one's ability to self-insure against medical expense shocks depends critically upon one's asset level, we match asset quantiles. We also match participation rates and hours by overall health status.

Under the MSM approach, the matches between the model and the data are expressed as a collection of moment conditions. To construct the moment conditions for the asset quantiles, we assume that assets, A_{it}, have a continuous density. Let $j \in \{1, 2, ..., J\}$ index quantiles, and let $g_{\pi_j}(t; \theta_0, \chi_0)$ denote the value of the π_j-th asset quantile predicted by the model. This means, for example, that if $\pi_1 = 1/3$ and $g_{\pi_1}(53; \theta_0, \chi_0) = \$50,000$, the model predicts that $1/3$ of all individuals have assets of $\$50,000$ or less at age 53. In Appendix E we show that the moment condition for the jth quantile can be written as

$$E\left(1\{A_{it} \leq g_{\pi_j}(t; \theta_0, \chi_0)\} - \pi_j | t\right) = 0, \tag{17}$$

for $j \in \{1, 2, ..., J\}$, $t \in \{1, ..., T\}$. Since $J = 2$, equation (17) generates $2T$ moment conditions.

We compute $g_{\pi_j}(t; \theta, \chi)$ by finding the model's decision rules for consumption, hours, and benefit application, using the decision rules to generate artificial histories for many different simulated individuals, and finding the quantiles of the collected histories. Equation (17) therefore says that the data sample and the simulated sample have the same age-conditional asset quantiles. It is worth stressing that the distribution used to derive g_{π_j} is found by evaluating the model-generated decision rule, $A_{t+1}(X_{it}, \psi_{it}; \theta, \chi)$, over the simulated distributions of the state vector X_{it} and health cost shock ψ_{it}, rather than the empirical distributions; this

16

avoids the aforementioned problems in measuring X_{it}. This does not imply, however, that A_{it} cannot vary for reasons, such as measurement error, that are not incorporated into the simulations; the only requirement is that these effects on the data "average out," so that the expectation in equation (17), which is taken over the observed data, continues to hold.

Next, consider how a worker's asset quantile and health insurance status affects his participation. Let $\overline{P}_j(HI, t; \theta_0, \chi_0)$ denote the model-predicted labor force participation rate conditional upon assets being in the jth-quantile interval and health insurance being of type HI. In Appendix E we derive the following moment condition:

$$E\left(P_{it} - \overline{P}_j(HI, t; \theta_0, \chi_0) \mid HI_{it} = HI, g_{\pi_{j-1}}(t; \theta_0, \chi_0) \leq A_{it} \leq g_{\pi_j}(t; \theta_0, \chi_0), t\right) = 0, \quad (18)$$

for $j \in \{1, 2, ..., J+1\}$, $HI \in \{none, ret, tied\}$, $t \in \{1, ..., T\}$.[18] Equation (18) says that within each asset grouping, the data sample and the simulated sample have the same conditional mean. With 2 quantiles (generating 3 quantile-conditional means) and 3 health insurance types, equation (18) generates $9T$ moment conditions.

Finally, consider health-conditional hours and participation. Let $\overline{\ln H}(M, t; \theta_0, \chi_0)$ and $\overline{P}(M, t; \theta_0, \chi_0)$ denote the conditional expectation functions for hours (when working) and participation generated by the model; let $\ln H_{it}$ and P_{it} denote measured hours and participation. The model conditions are

$$E\left(\ln H_{it} - \overline{\ln H}(M, t; \theta_0, \chi_0) \mid P_{it} > 0, M_{it} = M, t\right) = 0, \quad (19)$$

$$E\left(P_{it} - \overline{P}(M, t; \theta_0, \chi_0) \mid M_{it} = M, t\right) = 0, \quad (20)$$

for $t \in \{1, ..., T\}$, $M \in \{good, bad\}$. Equations (19) and (20) yield $4T$ moment conditions. Combined with the $2T$ moment conditions for the asset quantiles and the $9T$ moment conditions for asset- and insurance-conditional participation, this generates a total of $15T$ moment

---

[18]Because we are interested in participation given one's opportunity set, we combine individuals who work and receive *tied* insurance with those who do not work and receive *COBRA* coverage, as those two groups had the same insurance opportunities.

conditions.

3.2 Estimation Mechanics

The mechanics of our MSM procedure are as follows. First, we estimate life cycle profiles from the data for hours, participation and assets. Second, using the same data used to estimate the profiles, we generate an initial distribution for health, health insurance status, wages, medical expenses, AIME and assets.[19] We also use this data to estimate many of the parameters contained in the belief vector χ, although we calibrate some of these parameters as well. Using χ, we generate matrices of random health, wage and medical expense shocks. The matrices hold shocks for 40,000 simulated individuals over their entire lives. Third, we compute the decision rules for an initial guess of the parameter vector θ, using χ and the numerical methods described in Appendix D.

The fourth step is to simulate profiles for the decision variables. Each simulated individual receives a draw of assets, health, wages and medical expenses from an initial distribution, and is assigned one of the simulated sequences of health, wage and health cost shocks. With the initial distributions and the sequence of shocks, we then use the decision rules to generate that person's decisions over the life cycle. Each period's decisions determine the conditional distribution of next period's states, and the simulated shocks pin the states down exactly.

Fifth, we aggregate the simulated data into profiles in the same way we aggregated the true data. Sixth, we compute moment conditions, i.e., we find the distance between the simulated and true profiles. Finally, we pick a new value of θ and repeat the whole process. The value of θ that minimizes the distance between the true data and the simulated data, as described in equations (17)-(20), $\hat{\theta}$, is the estimated value of θ_0.[20] We discuss the asymptotic distribution of the parameter estimates, the weighting matrix and the overidentification tests in Appendix E.

---

[19]As described in Appendices B and C, the starting values for AIME and pension wealth are imputed from the initial draws of assets and wages.

[20]Because the GMM criterion function is discontinuous, we search over the parameter space using a simplex algorithm. It usually takes around 2 days to estimate the model on a 40-node supercomputer, with each iteration (of steps 3-6) taking around 10 minutes.

4   Data and Calibrations

4.1   HRS Data

We estimate the model using data from the first five waves of the Health and Retirement Survey (HRS). The HRS is a sample of non-institutionalized individuals, aged 51-61 in 1992, and their spouses. With the exception of assets and health costs, which are measured at the household level, our data are for male household heads. The HRS surveys individuals every two years, so that we have 5 waves of data covering the period 1992-2000. The HRS also asks respondents retrospective questions about their work history that allow us to infer whether the individual worked in non-survey years. Details of this, as well as variable definitions, selection criteria, and a description of the initial joint distribution, are in Appendices F and H.

With the exception of wages, we do not adjust the data for cohort effects. Because the HRS covers a fairly narrow age range, this omission should not generate much bias.

4.2   Health Insurance Status and Health Costs

We assign individuals to one of four mutually exclusive health insurance groups: *ret*, *tied*, *COBRA*, and *none*, as described in section 2. Because of small sample problems, the *none* group includes those with no insurance as well as those with private insurance. Neither type receives employer-provided coverage. Because the model includes a consumption floor to capture the insurance provided by Medicaid, the *none* group also includes those whose only form of health insurance is Medicaid. We assign those who have health insurance provided by their spouse to the *ret* group, along with those who report that they could keep their health insurance if they left their jobs. Neither of these types has their health insurance tied to their job. We assign individuals who would lose their employer-provided health insurance after leaving their job to the *tied* group.

Unfortunately, the HRS has information on health insurance outcomes, not choices. This is an important problem for individuals out of the labor force with no health insurance; it

19

is unclear whether these individuals could have purchased COBRA coverage but elected not to do so. To circumvent this problem we use health insurance in the initial wave and the transitions implied by equation (10) to predict health insurance options. For example, if an individual has health insurance that is tied to his job and was working in the previous wave, that individual's choice set is *tied* health insurance and working or *COBRA* insurance and not working.

The HRS has data on self-reported medical expenses. Medical expenses are the sum of insurance premia paid by the household, drug costs, and out of pocket costs for hospital, nursing home care, doctor visits, dental visits, and outpatient care. We are interested in the medical expenses that households face. Unfortunately, we observe only the medical expenses that these households actually pay for. This means that the observed medical expense distribution for low-wealth households is censored, because programs such as Medicaid pay much of their medical expenses. Because our model explicitly accounts for government transfers, the appropriate measure of medical expenses includes medical expenses paid by the government. Therefore, we assign Medicaid payments to households that received Medicaid benefits. The *2000 Green Book* (Committee on Ways and Means, 2000, p. 923) reports that in 1998 the average Medicaid payment was \$10,242 per beneficiary aged 65 and older, and \$9,097 per blind or disabled beneficiary. Starting with this average, we then assume that Medicaid payments have the same volatility as the medical care payments made by uninsured households. This allows us to generate a distribution of Medicaid payments.

We fit these data to the health cost model described in Section 2. Because of small sample problems, we allow the mean, $hc(.)$, and standard deviation, $\sigma(.)$, to depend only on the individual's Medicare eligibility, health insurance type, health status, labor force participation and age. Following the procedure described in French and Jones (2004), $hc(.)$ and $\sigma(.)$ are set so that the model replicates the mean and 95th percentile of the cross-sectional distribution of medical expenses (in levels, not logs) in each of these categories. We found that this procedure did an extremely good job of matching the top 20% of the medical expense distribution. Details are in Appendix G.

Table 1 presents some summary statistics, conditional on health status. Table 1 shows that for healthy individuals who are 64 years old, and thus not receiving Medicare, average annual medical costs are $2,950 for those with tied coverage and $5,140 for those with no employer-provided coverage, a difference of $2,190. With the onset of Medicare at age 65, the difference shrinks to $410. For individuals in bad health, the difference shrinks from $2,810 at age 64 to $530 at age 65.[21]

It is not just differences in mean medical expenses, however, that determine the value of health insurance, but also differences in variance and skewness. If health insurance reduces health cost volatility, risk-averse individuals may value health insurance at well beyond the cost paid by employers. To give a sense of the volatility, Table 1 also presents the standard deviation and 99.5th percentile of the health cost distributions. Table 1 shows that for healthy individuals who are 64 years old, average annual medical costs have a standard deviation of $7,150 for those with tied coverage and $19,060 for those with no employer-provided coverage. With the onset of Medicare at age 65, average annual medical costs have a standard deviation of $5,370 for those with tied coverage and $8,090 for those with no employer-provided coverage. Therefore, Medicare not only reduces average health costs for those without employer provided health insurance. It also reduces the volatility of health costs.

The parameters for the idiosyncratic process ψ_t, $(\sigma_\xi^2, \sigma_\epsilon^2, \rho_{hc})$, are estimated in French and Jones (2004). Table 2 presents the parameters, which have been normalized so that that overall variance, σ_ψ^2, is one. Table 2 reveals that at any point in time, the transitory component generates almost 67% of the cross-sectional variance in medical expenses. The results in French and Jones (2004) reveal, however, that most of the variance in cumulative *lifetime* medical expenses is generated by innovations to the persistent component. Given the autocorrelation coefficient ρ_{hc} of 0.925, this is not surprising.

---

[21]The pre-Medicare cost differences are roughly comparable to EBRI's (1999) estimate that employers on average contribute $3,288 to their employees' health insurance.

	Retiree - Working	Retiree - Not Working	Tied	COBRA	None
Age = 64, without Medicare, Good Health					
Mean	$2,930	$3,360	$2,950	$3,670	$5,140
Standard Deviation	$6,100	$7,050	$7,150	$8,390	$19,060
99.5th Percentile	$35,530	$41,020	$40,210	$47,890	$91,560
Age = 65, with Medicare, Good Health					
Mean	$2,590	$2,800	$3,420	$2,750	$3,830
Standard Deviation	$4,700	$4,700	$5,370	$5,420	$8,090
99.5th Percentile	$28,000	$28,240	$32,460	$31,880	$47,010
Age = 64, without Medicare, Bad Health					
Mean	$3,750	$4,300	$3,770	$4,690	$6,580
Standard Deviation	$7,970	$9,220	$9,330	$10,960	$24,840
99.5th Percentile	$46,240	$53,380	$52,210	$62,240	$118,400
Age = 65, with Medicare, Bad Health					
Mean	$3,310	$3,580	$4,380	$3,520	$4,910
Standard Deviation	$6,150	$6,150	$7,040	$7,080	$10,570
99.5th Percentile	$36,530	$36,890	$42,460	$41,520	$61,180

Table 1: Medical Expenses, by Medicare and Health Insurance Status

Parameter	Variable	Estimate
σ_ϵ^2	innovation variance of persistent component	0.04811
ρ_{hc}	autocorrelation of persistent component	0.925
σ_ξ^2	innovation variance of transitory component	0.6668

Table 2: Variance and Persistence of Innovations to Medical Expenses

4.3 Wages

Recall from equation (13) that $\ln W_t = \alpha \ln(H_t) + W(M_t, t) + \omega_t$. Following Aaronson and French (2004), we set $\alpha = 0.415$, which implies that a 50% drop in work hours leads to a 25% drop in the offered hourly wage. This is in the middle of the range of estimates of the effect of hours worked on the offered hourly wage. Because the wage information in the HRS varies from wave to wave, we take the second term, $W(M_t, t)$, from French (2003), who estimates a fixed effects wage profile using data from the Panel Study of Income Dynamics. We rescale the level of wages to match the average wages observed in the HRS at age 53. Because fixed-effects estimators estimate the growth rates of wages of the same individuals, the fixed-effects estimator accounts for cohort effects—the cohort effect is the average fixed effect for all members of that cohort. Moreover, the fixed-effects estimator avoids composition

bias problems—questions of whether high wage or low wage individuals drop out of the labor market—as long as wage growth rates for workers and non-workers are the same.[22]

The parameters for the idiosyncratic process ω_t, (σ_η^2, ρ_W) are the non-measurement-error component of the model estimated in French (2003). The results indicate that the auto-correlation coefficient ρ_W is 0.977; wages are almost a random walk. The estimate of the innovation variance σ_η^2 is 0.0141; one standard deviation of an innovation in the wage is 12% of wages. These estimates imply a high degree of long-run wage uncertainty.

4.4 Remaining Calibrations

Proceeding analogously to Hubbard, Skinner and Zeldes (1994, Appendix A), we set the consumption floor C_{min} equal to $3,500. This figure is an estimate of the average benefits available to a childless household with no members aged 65 or older. This value may well be too low; in 1998 the Federal SSI benefit for elderly (65+) couples was nearly $9,000 (Committee on Ways and Means, 2000, p. 229). We have chosen to be conservative because, as discussed below, consumption floors can drastically reduce the value of health insurance. We also show below, however, that the model does match the data better when $C_{min} = \$3,500$ than when $C_{min} = \$100$.

We set the anticipated component of the interest rate r equal to 0.03, although we allow for rate of return shocks, as described below. Following De Nardi (2004), the value of K, which determines the curvature of the bequest function, is $500,000. Spousal income depends upon an age polynomial and the wage. Health status and mortality both depend on previous health status interacted with an age polynomial. Estimates are available from the authors.

---

[22]However, if individuals leave the market because of a sudden wage drop, such as from job loss, then wage growth rates for workers will be greater than wage growth for non-workers. This problem will bias estimated wage growth upward. French (2003) estimates the extent of selection, and finds that it does not seriously affect his results. The results shown below are based on wage profiles that do not account for selection.

5 Data Profiles

The top panel of Figure 1 shows the 1/3rd and 2/3rd asset quantiles at each age for the HRS sample. About one third of the men sampled live in households with less than $70,000 in assets, and about one third live in households with over $240,000 of assets.
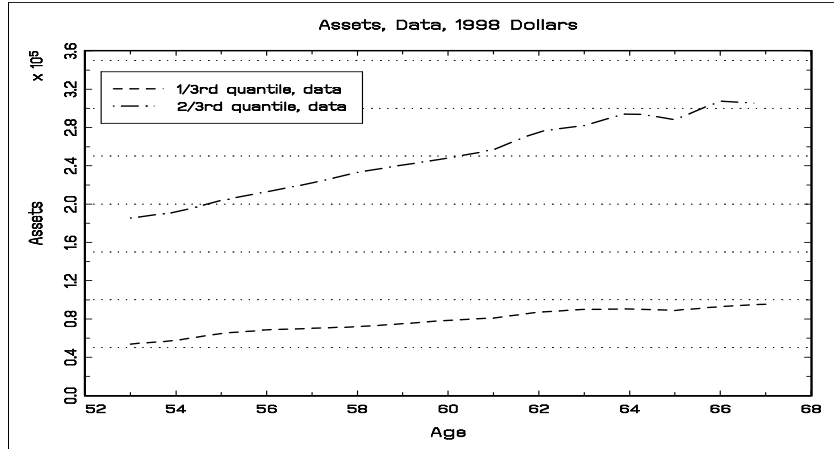

Figure 1: ASSET QUANTILES, DATA

The asset profiles show that assets grow rapidly with age. This rapid growth, which is higher than that reported in other studies (for example, Cagetti, 2003), is partly due to a run-up in asset prices during the 1990's. Recall that the core HRS sample was aged 51-61 in 1992. Therefore, an individual who is observed at age 51 is likely to have been observed in 1992, whereas an individual observed at age 67 is likely to have been observed in later waves—in fact, the oldest individuals in the core sample do not turn 67 until 1998, when the fourth wave is collected. Because older individuals are more likely to be observed in later waves, they are more likely to have enjoyed rapid asset growth. Failure to account for this will lead the econometrician to overstate the saving of sample members, which in turn will likely lead him to overstate their patience (β).

To account for the rapid run-up in asset prices that occurred over our sample period, we use the age-specific interest rate $r_t = r + \varepsilon_t$ when simulating artificial life cycle histories. The

estimates of the unanticipated component ε_t are described in Appendix I. When finding the decision rules, however, computational concerns led us to set $r_t = r$.[23]

The first panel of Figure 2 shows empirical job exit rates by health insurance type. Recall that Medicare should provide the largest labor market incentives for workers that have tied health insurance. If these people place a high value on employer-provided health insurance, they should either work until age 65, when they are eligible for Medicare, or they should work until age 63.5 and use COBRA coverage as a bridge to Medicare. The job exit profiles provide some evidence that those with tied coverage do tend to work until age 65. While the age-65 job exit rate is similar for those whose health insurance type is *tied* (20.5%), *ret* (21.2%), or *none* (21.2%), those with *ret* coverage have significantly higher exit rates at 62 (22.9%) than those with *tied* (16.9%) or *none* (14.0%). Although the hypothesis that the three groups have identical exit rates at age 65 cannot be rejected, the hypothesis that the three groups have identical exit rates at 62 is rejected at the 99% percent level and the hypothesis that the three groups have identical exit rates at all ages is rejected at the 99.5% level.[24] At every age between 56 and 64, those with retiree coverage have higher job exit rates than those with tied or no coverage. These differences across health insurance groups, while large, are smaller than the differences in the empirical exit profiles reported in Rust and Phelan (1997).

If individuals with tied coverage use COBRA coverage as a bridge to Medicare, we would expect that those with tied coverage would be more likely to exit the labor market at age 63.5. Those with tied coverage, however, have lower job exit rates at ages 63 and 64 than those with retiree coverage. Because COBRA coverage is costly, this is not evidence that people do not value retiree health insurance. It is evidence, however, that people place relatively little value on the insurance aspect of health insurance, as the option to buy actuarially fair insurance when not working appears to have a small effect on job exit rates.

The health insurance classifications generated by the HRS data probably contain measure-

---

[23]In applying this differential treatment, we are assuming both that the asset price run-up was unanticipated and that interest rate uncertainly has little effect on saving behavior.

[24]The *F-statistic* for the hypothesis that all profiles are equal to one another is 1.95, versus the 95% critical value of 1.41.
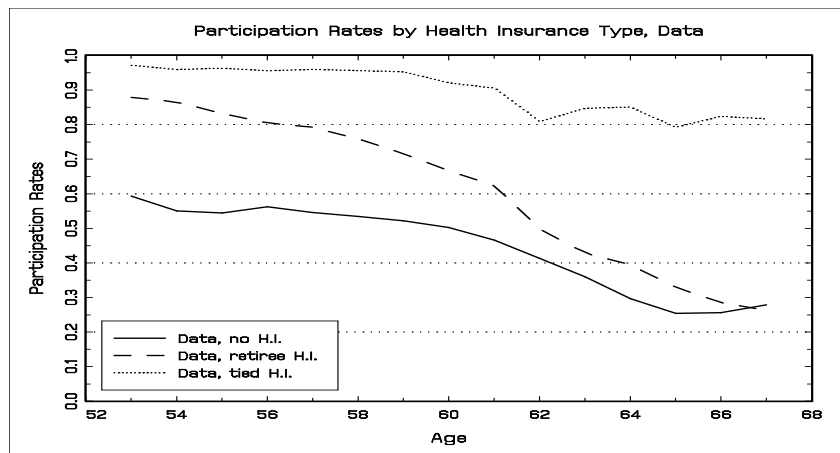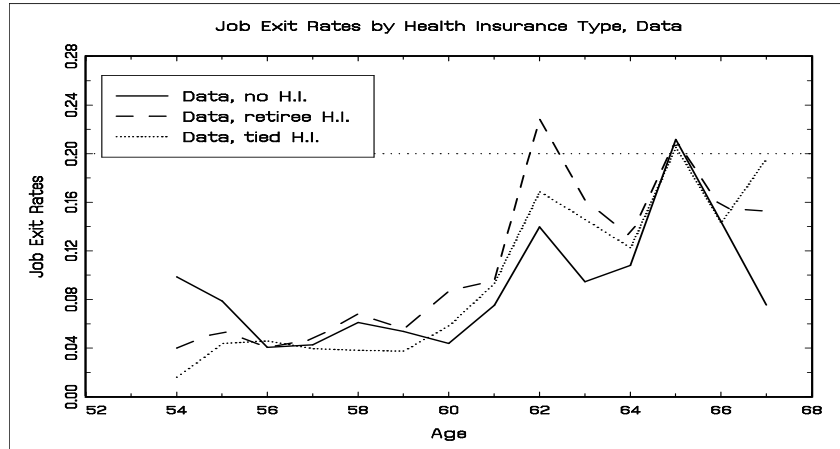
Figure 2: JOB EXIT AND PARTICIPATION RATES, DATA

ment error. Appendix H shows job exit rates generated under several alternative measures of health insurance type. All of the measures generate similar sets of profiles.

The bottom panel of Figure 2 presents observed labor force participation rates. The differences in participation across health insurance types are quite large. Even at age 53, labor force participation of the uninsured is 60%, well below the participation rate of those with either retiree or tied coverage. Moreover, the top panel of Figure 2 shows that before age 56, those with no health insurance coverage have the highest job exit rates. Nevertheless, when considering the bottom panel of Figure 2, it is useful to keep in mind the transitions

implied by equation (10): retiring workers in the *tied* insurance category transition into the *none* category. Because of this, the labor force participation rates for those with *tied* insurance are calculated over a group of individuals that were all working in the previous period. It is therefore unsurprising that the *tied* category has the highest participation rates. Conversely, it is not surprising that the *none* category has the lowest participation rates, given that category includes *tied* workers who retire. What is surprising, however, is the magnitude of the differences.

6 Baseline Results

6.1 Preference Parameter Estimates

The goal of our MSM estimation procedure is to match the life cycle profiles for assets, hours and participation rates found in the HRS data. In order to use these profiles to identify preferences, we make several identifying assumptions, the most important being that preferences vary with age only as a result of changes in health status. Therefore, age and health insurance can be thought of as "exclusion restrictions", which change the incentives for work and savings but do not change preferences.

Table 3 presents preference parameter estimates under several different specifications. In this section, we discuss the baseline specification, where $C_{min} = \$3,500$ and the household can use all of its wealth to insure against medical expense shocks. We return to the other specifications in Section 8.[25]

Perhaps the most important parameter is ν, the coefficient of relative risk aversion for flow utility. A more familiar measure of risk aversion is the coefficient of relative risk aversion for consumption. Assuming that labor supply is fixed and the value of bequests is close to zero, it can be approximated as $-\frac{(\partial^2 U/\partial C^2)C}{\partial U/\partial C} = -(\gamma(1-\nu)-1) = 2.97$. This value is within the range of estimates found in studies of consumption/savings (Gourinchas and Parker, 2002, Cagetti,

---

[25]In the interest of space, we do not discuss the standard errors, which might appear to be rather small. A complete analysis of this issue, as well as further insights into identification, can be found in French (2003).

Parameter and Definition	Baseline (1)	No Saving (2)	Illiquid Housing (3)	$C_{min} =$ \$100 (4)	Mismeasured Assets (5)
γ: consumption weight	0.610	0.498	0.708	0.662	0.611
	(0.0022)	(0.0042)	(0.0024)	(0.0032)	(0.0021)
ν: coefficient of relative risk aversion, utility	4.23	6.84	0.806	1.13	4.31
	(0.054)	(0.067)	(0.013)	(0.017)	(0.064)
β: time discount factor	0.994	0.994	0.959	0.964	0.995
	(0.0028)	(NA)	(0.0003)	(0.0008)	(0.0029)
L: leisure endowment	4,493	5,299	4,369	4,189	4,505
	(25.8)	(46.8)	(17.6)	(20.5)	(26.0)
ϕ: hours of leisure lost, bad health	242	341	407	134	237
	(12.7)	(20.9)	(10.8)	(6.4)	(13.6)
θ_P: fixed cost of work, in hours	1,261	1,334	1,530	1,272	1,279
	(14.0)	(13.1)	(8.4)	(7.8)	(14.5)
θ_B: bequest weight	1.56	2.57×10^{-5}	13.60	38.56	1.49
	(0.035)	(NA)	(0.401)	(0.523)	(0.042)
GMM Criterion	2,239	1,316	6,358	3,089	2,221
χ^2 statistic	1,876	744	7,655	2,106	1,766
Degrees of freedom	218	100	218	218	218
Standard errors in parentheses NA indicates parameters were fixed during estimation					

Table 3: ESTIMATED STRUCTURAL PARAMETERS

2003), but it is larger than the values of 1.07 and 2.14 reported by Rust and Phelan (1997) and Blau and Gilleskie (2003), respectively, in their studies of retirement. ν is identified largely by the asset quantiles. The bottom quantile in particular depends on the interaction of precautionary motives and the consumption floor. If the consumption floor is sufficiently low, the risk of a catastrophic health cost shock, which over a lifetime could equal tens of thousands of dollars, can generate strong precautionary incentives; we discuss this point in the robustness checks below.

Turning to labor supply, we find that individuals in our sample are willing to intertemporally substitute their work hours. In particular, simulating the effects of a 2% wage change reveals that the wage elasticity of average hours is 1.38 at age 60. This relatively high labor supply elasticity arises because the fixed cost of work generates volatility on the participation margin. The participation elasticity is 1.15 at age 60, implying that wage changes cause rel-

atively small hours changes for workers. For example, the Frisch labor supply elasticity of an individual working 2000 hours per year is approximated as $-\frac{L-H_t-\theta_P}{H_t} \times \frac{1}{(1-\gamma)(1-\nu)-1} = 0.27$, which is similar to MaCurdy's (1981) estimate.

The fixed cost of work is identified by the life cycle profile of hours worked by workers. Average hours of work (available upon request) do not drop below 1,000 hours per year (or 20 hours per week) even though labor force participation rates decline to near zero. In the absence of a fixed cost of work, one would expect hours worked to parallel the decline in labor force participation.

The parameter γ is identified by noting that the within-period utility function is Cobb-Douglas between consumption and leisure, so that γ (roughly) gives the share of resources spent on consumption rather than leisure. Therefore, the coefficient of relative risk aversion for consumption, the labor supply elasticity at the both the intensive (hours) and extensive (participation) margin, and the share of resources spent on consumption versus leisure identify the structural parameters γ, ν, L, and θ_P. The parameter ϕ is identified by noting that unhealthy individuals work fewer hours than healthy individuals, even after conditioning on the wage.

Another important parameter is the time discount factor β, with a value of 0.994. It is higher than most estimates of β for two reasons. The first reason is clear upon inspection of the Euler Equation: $\frac{\partial U_t}{\partial C_t} \geq \beta s_{t+1}(1 + r(1 - \tau_t))E_t \frac{\partial U_{t+1}}{\partial C_{t+1}}$, where τ_t is the marginal tax rate.[26] Note that this equation identifies the product $\beta s_{t+1}(1 + r(1 - \tau_t))$, but not its individual elements. Therefore, a lower value of s_{t+1} or $(1 + r(1 - \tau_t))$ results in a higher estimate of β. Given that many studies omit mortality risk and/or taxes—implicitly setting s_{t+1} and/or $1 - \tau_t$ to one—it is not surprising that they find lower values of β. The second reason is that β is identified not only by the intertemporal substitution of consumption, as embodied in the asset profiles, but also in the intertemporal substitution of leisure, as embodied in the labor supply profiles. Models of labor supply and savings, such as MaCurdy (1981) or French

---

[26]Note that this equation does not hold exactly when individuals value bequests. Also note that the Euler Equation holds with equality when assets are positive.

(2003), often suggest that agents are very patient.[27]

The bequest parameter θ_B is identified largely from the top asset quantile. It follows from equation (3) that when the shift parameter K is large, the marginal utility of bequests will be lower than the marginal utility of consumption unless the individual is rich. In other words, the bequest motive mainly affects the saving of the rich; for more on this point, see De Nardi (2004). The estimates of θ_B vary a great deal across specifications. However, the marginal propensity to consume out of wealth in the final period of life, when any savings will be bequeathed, is much more stable across specifications. For low-income individuals, this marginal propensity to consume—which is a nonlinear function of $\theta_B, \beta, \gamma, \nu$, and K—is 1. For high-income individuals, the marginal propensity to consume ranges from a baseline value of 0.155 to 0.29.

6.2   Simulated Profiles

The bottom of Table 3 displays GMM criterion values and overidentification test statistics; the two differ because we use a diagonal weighting matrix (see Appendix E). Even though the model is formally rejected, the life cycle profiles generated by the model for the most part resemble the life cycle profiles generated by the data.

Figure 3 shows that the model fits both asset quantiles fairly well. The model is able to fit the lower quantile in large part because of the consumption floor of $3,500; the predicted lower quantile rises dramatically when the consumption floor is lowered. This is consistent with the results found by Hubbard, Skinner, and Zeldes (1995). Hubbard, Skinner, and Zeldes show that if the government guarantees a minimum consumption level, those with low assets and income will tend not to save, because their consumption will never drop below a certain level, even in the presence of a large negative health cost shock. Put differently, if an individual is at the consumption floor, his savings will be taxed at a marginal rate of 100%. It is therefore not surprising that within the model the consumption floor reduces saving by

---

[27]As a sensitivity analysis, we fixed β to 0.95 and θ_B to 0, and re-estimated the model. The restricted model fits the data much more poorly: the χ^2 test statistic rises from 1,880 to 5,460.
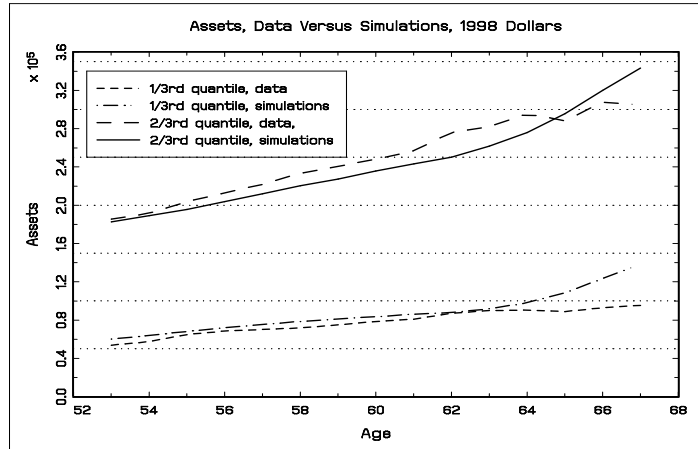
Figure 3: ASSETS, DATA AND SIMULATIONS

individuals with low income and assets.

The three panels in the left hand column of Figure 4 show that the model is able to replicate the two key features of labor force participation across age and health insurance. The first key feature is that participation declines with age, and the declines are especially sharp between ages 62 and 65. The model is also able to match the aggregate decline in participation at age 65 (a 6.2 percentage point decline in the data versus a 8.1 percentage point decline predicted by the model), although it underpredicts the decline in participation at 62 (a 10.4 percentage point decline in the data versus a 5.4 percentage point decline predicted by the model). We return to the age-62 decline in participation below.
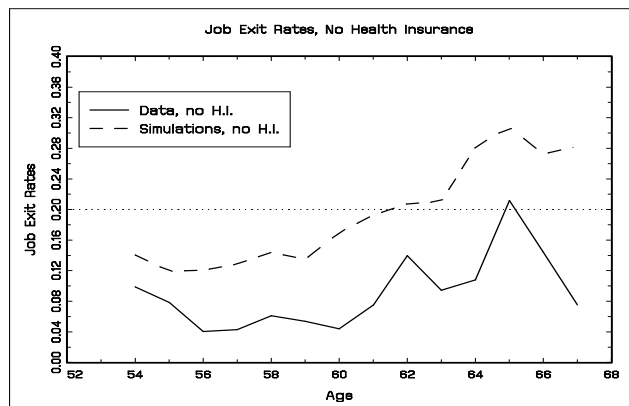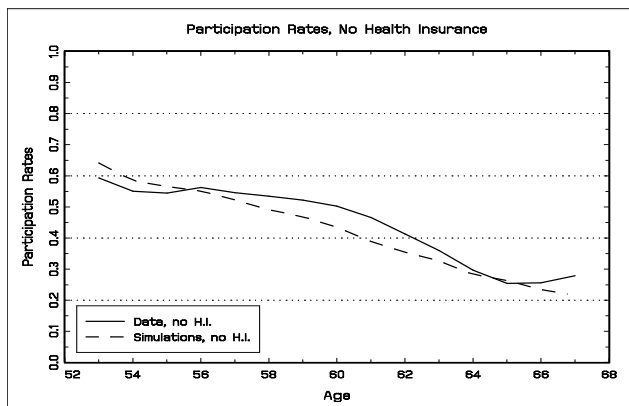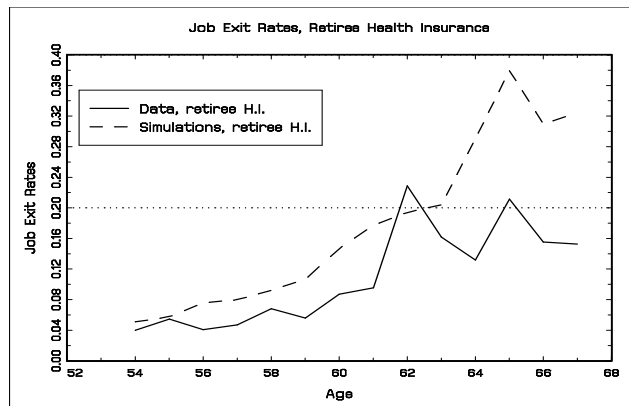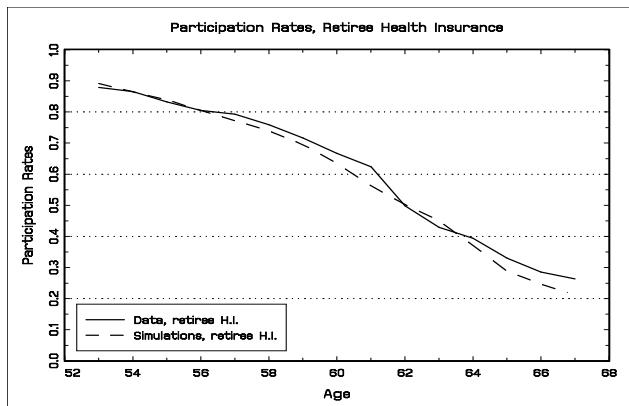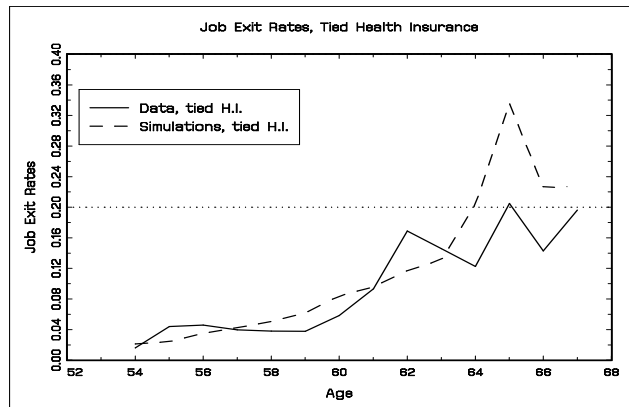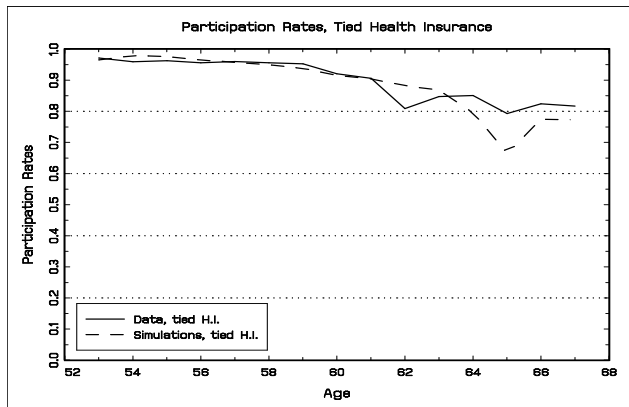
Figure 4: Participation and Job Exit Rates, Data and Simulations

The second key feature is that there are large differences in participation and job exit rates across health insurance types. The model is able to match the fact that exit rates for those with tied coverage are low until age 65. Recall that in the data, exit rates are 3, 6, and 0 percentage points higher at ages 54, 62, and 65 for those with retiree coverage than for those with tied coverage. In the simulations, exit rates are 3, 8, and 4 percentage points higher at ages 54, 62, and 65 for those with retiree coverage than for those with tied coverage. Moreover, the model also matches the low participation levels of the uninsured. Turning to the lower left panel of Figure 5, the data show that the group with the lowest participation rates are the uninsured with low assets. Although the model is not fully able to replicate this fact, the consumption floor greatly reduces participation of the uninsured with low assets. Without a high consumption floor, the risk of catastrophic medical expenses, in combination with risk aversion, would cause the uninsured to remain in the labor force and accumulate a buffer stock of assets.

The panels in the right hand column of Figure 4 compare observed and simulated job exit rates for each health insurance type. They show that the model over-predicts the job exit rates of workers with either retiree coverage or no health insurance. The poor fit is in part an artifact of our MSM estimation procedure: because participation contains information on the level of labor supply as well as year-to-year changes, we match participation rates rather than job exit rates. The model's participation rate profiles, shown in the left hand column of Figure 4, match the data much better. Taken together, the two columns imply that the model over-predicts the amount of labor market exit and re-entry for workers in the *ret* or *none* categories.[28] In contrast, the model predicts very little exit and re-entry for workers with *tied* health insurance. This reflects our assumption that once an elderly worker with tied coverage leaves his job, he will never have a job with tied coverage again.

---

[28]The model lacks tenure effects, which would imply that workers who exited the labor market would usually re-enter with a lower wage. Adding tenure effects would likely reduce the amount of exit and re-entry predicted by the model. We omit tenure effects largely for computational reasons: adding them would require us to include the worker's previous employment status as a state variable.
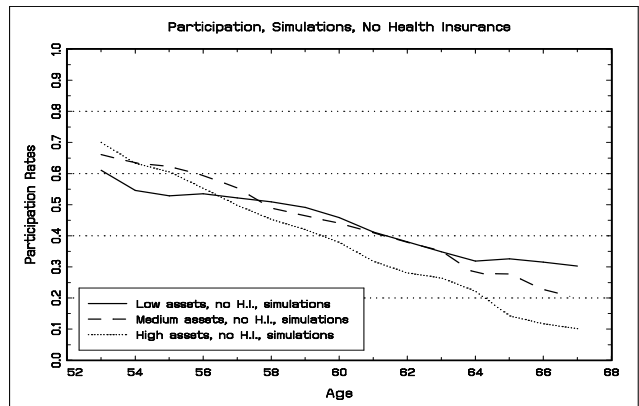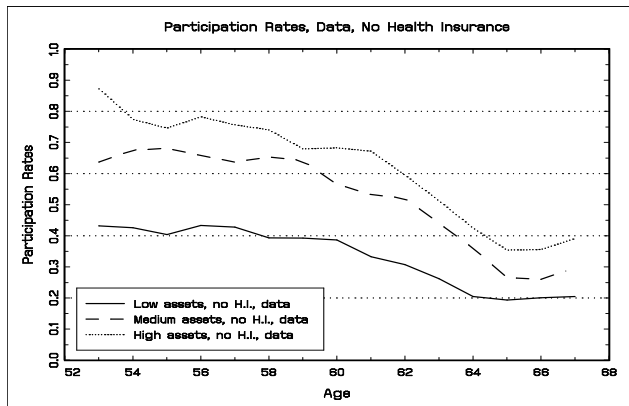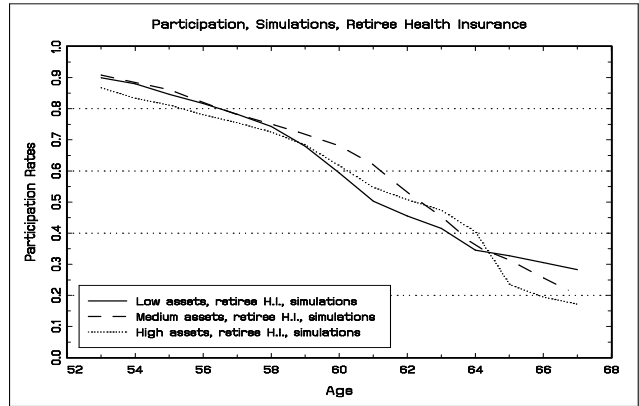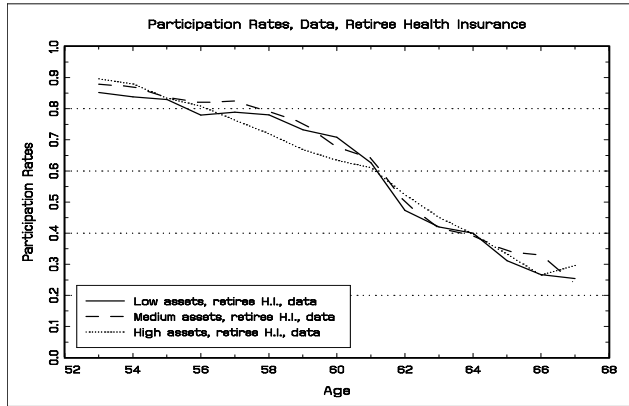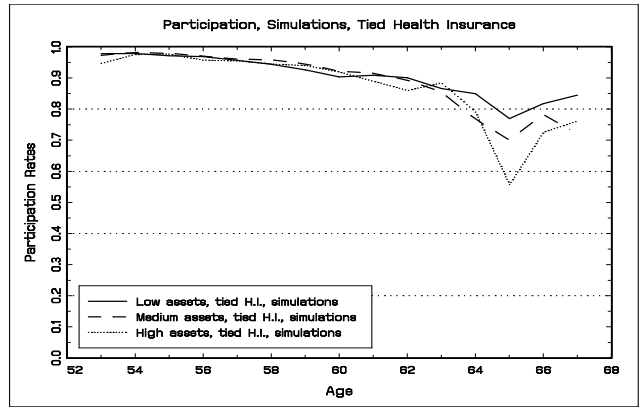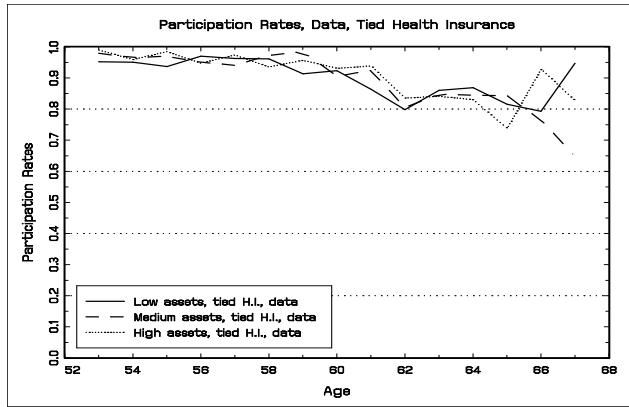
Figure 5: LABOR FORCE PARTICIPATION RATES BY ASSET GROUPING, DATA AND SIMULATIONS

6.3   The Effects of Employer-Provided Health Insurance

The empirical profiles discussed above are informative, but do not identify the effects of health insurance on retirement, for two reasons. First, the distributions of wages and wealth in our sample differ across health insurance types; for example, workers with retiree coverage tend to be wealthier. Second, holding everything else fixed, workers with retiree coverage have the highest pension accrual rates, while workers with no health insurance have the lowest accrual rates. Therefore, retirement incentives differ across health insurance categories for reasons unrelated to health insurance incentives. Our model can disentangle these different effects.

To isolate the effects of employer-provided health insurance, we conduct some additional simulations. Retaining the parameter values shown in the first column of Table 3, we fix pension accrual rates so that they are identical across health insurance types. We then simulate the model three times, assuming first that all workers have no health insurance, then retiree coverage, then tied coverage at age 53.

This exercise reveals that at age 54, the job exit rate would be 2.9 percentage points higher if all workers had retiree coverage rather than tied coverage at age 53. The gap rises to 5.8 percentage points at age 61 and 5.3 percentage points at age 62, then declines to -1.8 percentage points at age 65. These differences in exit rates across health insurance types are similar to, but smaller than, the raw differences in exit rates observed in both the data and the simulations. This indicates that not accounting for other observable differences in retirement incentives leads the econometrician to slightly overstate the effect of health insurance on exit rates.

The effect of health insurance can also be measured by calculating how it affects the retirement age, defined here as the oldest age at which the individual worked. Moving from retiree to tied coverage increases the average retirement age by 0.47 years.

A useful comparison appears in the reduced form model of Blau and Gilleskie (2001), who study labor market behavior between ages 51 and 62 using waves 1 and 2 of the HRS data. They find that having retiree coverage, as opposed to tied coverage, increases the job

exit rate around 1% at age 54 and 7.5% at age 61. Given that they use the same data as we do, this similarity is not surprising. Moreover, they find that accounting for selection into health insurance plans modestly increases the estimated effect of health insurance on exit rates. Other reduced form findings in the literature are qualitatively similar to Blau and Gilleskie. For example, Madrian (1994) finds that retiree coverage reduces the retirement age by 0.4 - 1.2 years, depending on the specification and the data employed. Karoly and Rogowski (1994), who attempt to account for selection into health insurance plans, find that retiree coverage increases the job exit rate 8 percentage points over a $2\frac{1}{2}$ year period. Our estimates, therefore, lie within the range established by previous reduced form studies, giving us confidence that the model can be used for policy analysis.

Structural studies that omit medical expense risk usually find smaller health insurance effects than we do. For example, Gustman and Steinmeier (1994) find that retiree coverage reduces years in the labor force by 0.1 year. Lumsdaine et al. (1994) find even smaller effects. In contrast, structural studies that include medical expense risk but omit self-insurance usually find effects that are at least as large as ours. Our estimated effects are roughly similar to Blau and Gilleskie's (2003), who find that retiree coverage reduces participation 3.4 percentage points, but are smaller than the effects found by Rust and Phelan (1997).

6.4   Policy Experiments

The model allows us to analyze how changing the Social Security and Medicare rules would affect retirement behavior. In particular, we increase both the normal Social Security retirement age and the Medicare eligibility age from 65 to 67, and measure the resulting changes in simulated work hours and exit rates. The results of these experiments are summarized in Table 4.

The first column of Table 4 shows model-predicted labor market participation at ages 60 through 67 under the current (1998) retirement and eligibility ages. Under the current rules, the average person works a total of 3.42 years over this eight-year period. The fifth column of Table 4 shows that this is close to the total of 3.59 years observed in the data.

36

	1998 rules:	2030 rules:			
	SS = 65	SS = 67	SS = 65	SS = 67	
	MC = 65	MC = 65	MC = 67	MC = 67	Data
Age	(1)	(2)	(3)	(4)	(5)
	Participation rates, all health insurance types				
60	0.633	0.648	0.641	0.651	0.657
61	0.574	0.590	0.578	0.594	0.617
62	0.520	0.530	0.527	0.535	0.513
63	0.474	0.485	0.476	0.492	0.450
64	0.394	0.475	0.398	0.484	0.404
65	0.313	0.372	0.349	0.400	0.342
66	0.273	0.304	0.305	0.322	0.308
67	0.239	0.411	0.243	0.410	0.295
Total 60-67	3.420	3.815	3.518	3.888	3.588
	Participation rates, workers with tied coverage at age 60				
60	0.912	0.925	0.920	0.927	0.921
61	0.842	0.848	0.845	0.858	0.833
62	0.773	0.773	0.775	0.788	0.704
63	0.698	0.698	0.697	0.714	0.609
64	0.576	0.648	0.592	0.672	0.567
65	0.417	0.492	0.519	0.575	0.468
66	0.355	0.409	0.434	0.460	NA
67	0.308	0.499	0.315	0.497	NA
Total 60-67	4.881	5.293	5.098	5.492	NA
SS = Social Security normal retirement age					
MC = Medicare eligibility age					
NA indicates insufficient number of observations					

Table 4: Effects of Changing the Social Security Retirement and Medicare Eligibility Ages: Baseline Parameters

The second column shows the average hours that result when the 1998 Social Security rules are replaced with the rules planned for the year 2030. Imposing the 2030 rules: (1) increases the normal Social Security retirement age, the date at which the worker can receive "full benefits", from 65 to 67; (2) significantly increases the credit rates for deferring retirement past the normal age; and (3) eliminates the earnings test for workers aged 67 and older. The second column shows that imposing the 2030 rules leads the average worker to increase years worked between ages 60 and 67 from 3.42 years to 3.82 years, an increase of 0.4 years. It is worth noting that in addition to changing the rate at which benefits accrue, raising the retirement age effectively eliminates two years of Social Security benefits. Therefore, raising the normal retirement age to 67 has both substitution and wealth effects, both of which cause

37

participation to increase.[29]

The third column of Table 4 shows participation when the Medicare eligibility age is increased 67.[30] This change increases total years of work by only 0.1 years. The fourth column shows the combined effect of raising both the Social Security retirement and the Medicare eligibility age. The joint effect is an increase of 0.47 years, 0.07 more than that generated by raising the retirement age in isolation. In short, the model predicts that raising the normal Social Security retirement age will have a much larger effect on retirement behavior than increasing the Medicare eligibility age.

One reason that Social Security has larger labor market effects than Medicare is that only 21% of our sample have *tied* coverage at age 53 and only 8.2% have *tied* coverage at 64. Medicare provides much smaller retirement incentives to workers in the *ret* or *none* categories. To focus on the incentives facing workers with *tied* coverage, the lower half of Table 4 shows labor force participation rates for those workers alone. These participation rates reveal that the effects of Medicare are significant. Table 4 shows that increasing the Medicare eligibility age from 65 to 67 in isolation increases total years worked by 0.22 years. This is not a trivial change, as it implies an annual increase in participation of over 2.5 percentage points. This amount is in fact larger than the changes found by Blau and Gilleskie (2003), whose simulations show that increasing the Medicare age reduces the average probability of non-employment by about 0.5 to 1 percentage points. Nevertheless, the effect of shifting forward the Social Security retirement age is larger still—it would increase years worked by 0.41 years.

To understand better the incentives generated by Medicare, we compute the value that individuals place on employer-provided health insurance, by finding the increase in assets that would make an uninsured individual as well off as a person with retiree coverage. In other

---

[29]To measure the size of the wealth effect, we raise the retirement age to 67 while increasing annual benefits at every age by 15.4%. The net effect of these two changes is to alter the Social Security incentive structure while keeping the present value of Social Security wealth (at any age) roughly equivalent to the age-65 level. Under this configuration, total years of work increase by 0.27 years, implying that 0.13 years of the 0.4-year increase is due to wealth effects.

[30]By shifting forward the Medicare eligibility age to 67, we increase from 65 to 67 the age at which medical expenses can follow the "with Medicare" distribution shown in Table 1.

words, we find the compensating variation $\lambda_t = \lambda(A_t, B_t, M_t, AIME_t, \omega_t, \zeta_{t-1}, t)$, where

$$V_t(A_t, B_t, M_t, AIME_t, \omega_t, \zeta_{t-1}, ret) = V_t(A_t + \lambda_t, B_t, M_t, AIME_t, \omega_t, \zeta_{t-1}, none).$$

Table 5 shows the compensating variation $\lambda(A_t, 0, good, \$32000, 0, 0, 60)$ at several different asset (A_t) levels.[31] The first column of Table 5 shows the valuations found under the baseline specification. One of the most striking features is that the value of employer-provided health insurance is fairly constant. Even though rich individuals can better self-insure, they also receive less protection from the government-provided consumption floor. In the baseline case, these effects more or less cancel each other out. Relative to their wealth, however, poor individuals do value health insurance more.

| | Baseline Specification | | |
| | With Uncertainty | Without Uncertainty | $C_{min} = \$100$ |
Asset Levels	(1)	(2)	(3)
-$2,300	$24,100	$14,000	$22,500
$54,400	$23,000	$14,200	$21,600
$149,000	$23,600	$14,300	$20,500
$600,000	$23,400	$14,700	$19,300
Compensating variation between *ret* and *none* coverages Calculations described in text			

Table 5: VALUE OF EMPLOYER-PROVIDED HEALTH INSURANCE

Part of the value of retiree coverage comes from a reduction in average medical expenses—because retiree coverage is subsidized—and part comes from a reduction in the volatility of medical expenses—because it is insurance. In order to separate the former from the latter, we eliminate health cost uncertainty, by setting the variance shifter $\sigma(M_t, HI_t, t, B_t, P_t)$ to zero, and recompute λ_t, using the same state variables and mean medical expenses as before. Without health cost uncertainty, λ_t is approximately $14,000. Comparing the two values of λ_t shows that about 60% of the value of health insurance comes from the reduction of average

---

[31]In making these calculations, we remove health-insurance-specific differences in pensions, as described in section 6.3. It is also worth noting that for the values considered, the conditional differences in expected health costs are smaller than the unconditional differences shown in Table 1.

medical expenses, and 40% is due to the reduction of medical expense volatility. This explains why Medicare has a small effect on retirement behavior. Individuals value health insurance not too much beyond the cost paid by employers.

Another way to assess the importance of health cost uncertainty is to redo the controlled experiments in Section 6.3 with health cost uncertainty eliminated. When workers face health cost uncertainty, participation drops 5.8 percentage points between ages 64 and 65 when all workers are presumed to have retiree coverage at age 53, and drops 11.3 percentage points when all workers are presumed to have tied coverage at age 53. In the absence of health cost uncertainty, the participation drop is 4.6 percentage points when workers are assumed to have retiree coverage and 9.3 percentage points when workers are assumed to have tied coverage. These results, too, suggest that workers value employer-provided health insurance mainly because it reduces average expenses, rather than risk.

7   How Important Are Savings?

We have argued that the ability to self-insure through saving could significantly affect the value of employer-provided health insurance. One test of this hypothesis is to modify the model so that individuals cannot save, and examine how labor market decisions change. In particular, we cap assets at $1,000, effectively requiring workers to consume their income net of health costs, as in Rust and Phelan (1997) and Blau and Gilleskie (2003).

Eliminating the ability to save requires changing our estimation strategy. In the absence of saving, β and θ_B are both very weakly identified. We therefore follow Rust and Phelan and Blau and Gilleskie by fixing β, in this case to its baseline value of 0.994. Similarly, we fix θ_B to replicate the baseline marginal propensity to consume out of wealth. Since the asset distribution is effectively degenerate in this no-saving case, we no longer match asset quantiles or quantile-conditional participation rates, matching instead participation rates for each health insurance category. The second column of Table 3 shows the parameter estimates for this specification.

When individuals cannot save, it makes little sense to express the value of health insurance as an asset increment, because any assets in excess of $1,000 must be spent immediately.[32] We therefore alter the compensating variation calculations used in Table 5 to express the value of retiree coverage as a compensating annuity.[33] By this measure, eliminating the ability to save greatly increases the value of retiree coverage: when assets are -$2,300, the compensating annuity increases from $3,600 in the baseline case to $15,800 in the no-savings case. When there is no health cost uncertainty, the comparable figures are $1,820 in the baseline case and $2,670 in the no-savings case. The ability to self-insure through saving appears to be quite important.

Table 6 shows the effects of changing the Social Security and Medicare rules under the no-saving specification. The first column of this table shows participation under the current rules. One of the most striking results is that in the absence of saving, there is a pronounced drop in participation rates—and spikes in the underlying job exit rates—at age 62. Unable to hold assets, workers must save through Social Security and pensions. Because they cannot borrow against their Social Security benefits, many workers that would otherwise retire earlier cannot fund their retirement before age 62.[34] Not surprisingly, workers with *ret* coverage have the largest age-62 responses. Most *tied* workers are unwilling to forgo their coverage so early.

Although the no-saving specification better fits age-62 job exit rates, along several other dimensions it fits worse than the baseline case with savings. For example, eliminating saving significantly increases the participation of workers with *tied* coverage, relative to both the baseline case and the data, even at ages beyond age 65. Because the baseline and no-savings cases are estimated with different moments, the overidentification statistics shown in the first two columns of Table 3 are not comparable. Additional calculations suggest, however, that

---

[32]When initial assets are -$2,300, compensating assets in the no-savings case are $250,000, a 10-fold increase over the baseline value of $24,000.

[33]To do this, we first find the compensating $AIME$, $\widehat{\lambda}_t$, where

$$V_t(A_t, B_t, M_t, AIME_t, \omega_t, \zeta_{t-1}, ret) = V_t(A_t, B_t, M_t, AIME_t + \widehat{\lambda}_t, \omega_t, \zeta_{t-1}, none).$$

This change in $AIME$ in turn allows us to calculate the change in expected pension and Social Security benefits that the individual would receive at age 65, the sum of which can be viewed as a compensating annuity. Because these benefits depend on decisions made after age 60, the calculation is only approximate.

[34]See Kahn (1988), Rust and Phelan (1997), and Gustman and Steinmeier (2002) for similar arguments.

Age	1998 rules: SS = 65 MC = 65 (1)	2030 rules: SS = 67 MC = 65 (2)	SS = 65 MC = 67 (3)	SS = 67 MC = 67 (4)	Data (5)
	Participation rates, all health insurance types				
60	0.704	0.700	0.716	0.715	0.657
61	0.627	0.637	0.635	0.648	0.617
62	0.514	0.544	0.527	0.556	0.513
63	0.472	0.511	0.483	0.524	0.450
64	0.429	0.473	0.440	0.485	0.404
65	0.336	0.344	0.403	0.417	0.342
66	0.280	0.335	0.332	0.376	0.308
67	0.222	0.338	0.221	0.341	0.295
Total 60-67	3.584	3.882	3.757	4.063	3.588
	Participation rates, workers with tied coverage at age 60				
60	0.881	0.856	0.939	0.930	0.921
61	0.877	0.861	0.915	0.910	0.832
62	0.808	0.809	0.857	0.871	0.704
63	0.733	0.742	0.777	0.804	0.609
64	0.666	0.690	0.730	0.762	0.567
65	0.592	0.597	0.669	0.680	0.468
66	0.484	0.528	0.526	0.565	NA
67	0.371	0.555	0.379	0.565	NA
Total 60-67	5.414	5.638	5.792	6.087	NA
SS = Social Security normal retirement age					
MC = Medicare eligibility age					
NA indicates insufficient number of observations					

Table 6: Effects of Changing the Social Security Retirement and Medicare Eligibility Ages: No Saving

the model with no savings provides a poorer overall fit.[35]

Columns (2)-(4) of Table 6 show that eliminating the ability to save makes workers much more sensitive to changes in Medicare. While raising the Medicare age leads to an additional 0.1 years of work in the baseline case, in the no-saving case the increase is 0.17 years. Although the increased effect is in part due to bigger reactions by individuals with tied coverage, it also reflects bigger reactions by individuals with retiree coverage. The supplemental effects of Medicare on retiree coverage, as shown in Table 1, become more important when individuals

---

[35]Inserting the decision profiles generated by the baseline model into the moment conditions used to estimate the no-savings case produces a GMM criterion value of 900 and an overidentification statistic of 600. In contrast, the no-saving specification produces a GMM criterion value of 1,320 and an overidentification statistic of 740, suggesting that the model fits better when saving is allowed.

cannot save.

8   Robustness Checks

8.1   Illiquid Housing

It has often been argued (e.g., Rust and Phelan, 1997) that housing equity is considerably less liquid than financial assets. Since housing comprises a significant proportion of most individuals' assets, its illiquidity would greatly weaken their ability to self-insure through saving.

To account for this possibility, we re-estimate the model using "liquid assets", which excludes housing and business wealth.[36] The third column of Table 3 contains the revised parameter estimates. The most notable changes are that ν, the coefficient of relative risk aversion, drops from 4.2 to 0.8, and that β, the discount rate, drops from 0.994 to 0.959. Both changes—lower risk aversion and lower patience—help the model fit the bottom quantile of liquid asset holdings, which averages less than \$5,000. Even with these changes, the model tends to overstate the bottom quantile, leading to a large overidentification statistic.

Table 7 shows the effects of changing the Social Security and Medicare rules when housing assets are illiquid. The first column shows participation under the current rules. The most notable result is that when housing assets are illiquid, simulated participation drops markedly at age 62.[37] The underlying asset-conditional profiles reveal that the drop is more pronounced for workers in the bottom 1/3rd of the asset distribution. These workers accumulate so few liquid assets that, unable to borrow against their Social Security benefits, they cannot fund their retirement before age 62. This contrasts with the data, where the age-62 drops vary across the asset quantiles to a much smaller extent.

---

[36] A complete analysis of illiquid housing would require us to treat housing as an additional state variable, with its own accumulation dynamics, and to impute the consumption services provided by owner-occupied housing. This is not computationally feasible. In this paper, we simply allow these effects to be captured in the preference parameters.

[37] Because the definition of assets affects which observations are retained in our sample, the empirical participation rates presented here differ slightly from those shown in earlier tables.

Age	1998 rules: SS = 65 MC = 65 (1)	2030 rules: SS = 67 MC = 65 (2)	SS = 65 MC = 67 (3)	SS = 67 MC = 67 (4)	Data (5)
	Participation rates, all health insurance types				
60	0.685	0.690	0.686	0.691	0.660
61	0.605	0.611	0.606	0.613	0.616
62	0.480	0.459	0.479	0.458	0.508
63	0.460	0.416	0.458	0.413	0.460
64	0.352	0.514	0.355	0.515	0.420
65	0.254	0.312	0.279	0.346	0.352
66	0.203	0.232	0.236	0.255	0.298
67	0.184	0.444	0.186	0.442	0.298
Total 60-67	3.223	3.677	3.285	3.733	3.612
	Participation rates, workers with tied coverage at age 60				
60	0.864	0.873	0.869	0.874	0.931
61	0.778	0.779	0.783	0.791	0.861
62	0.680	0.673	0.681	0.680	0.699
63	0.587	0.596	0.592	0.600	0.637
64	0.460	0.580	0.487	0.592	0.587
65	0.313	0.374	0.390	0.458	0.455
66	0.237	0.293	0.293	0.343	NA
67	0.201	0.490	0.206	0.481	NA
Total 60-67	4.122	4.659	4.302	4.819	NA
SS = Social Security normal retirement age					
MC = Medicare eligibility age					
NA indicates insufficient number of observations					

Table 7: Effects of Changing the Social Security Retirement and Medicare Eligibility Ages: Illiquid Housing

Despite this strong liquidity effect, the policy experiments conducted with illiquid housing yield the same conclusion as the baseline policy experiments: Social Security has stronger effects than Medicare.

8.2 Lower Consumption Floor

The baseline results were based on a consumption floor (C_{min}) of $3,500. Although quite low, this total could still be too high. The total includes the average housing subsidy, but many poor households receive no housing subsidy at all. Moreover, many eligible households do not collect benefits, possibly because transactions or "stigma" costs outweigh the value of public assistance.

Therefore, we re-estimate the model with C_{min} set to $100, a value that exposes workers to considerably more risk.[38] The revised parameter estimates appear in the fourth column of Table 3. The most notable change is that ν, the coefficient of relative risk aversion, drops from 4.2 to 1.1; the increase in risk exposure is offset by a decrease in risk aversion.

Although the baseline and revised models have similar chi-square statistics, along some dimensions the revised model fits the data much more poorly. In particular, while the baseline model over-predicts participation by low-asset workers—see the bottom panels of Figure 5— the revised model over-predicts this participation to a much greater extent. Recall that in the absence of a consumption floor, the model would predict that those with low assets and without health insurance would have high participation rates, as they would face precautionary motives to work and accumulate assets. It is thus not surprising that reducing the consumption floor from $3,500 to $100 increases the participation of poor workers, even after the preference parameters have been re-estimated.

This conclusion is reinforced by the insurance valuations shown in Table 5. Comparing the first and third columns of Table 5 shows that even though the average valuation is higher in the baseline case—because it utilizes a much larger value of the risk coefficient ν—in the low-floor case poor workers have a higher relative valuation. If we retain the baseline preference parameters, reducing C_{min} to $100 increases the valuation as well: for example, $\lambda(\$149000, 0, good, \$32000, 0, 0, 60)$, rises from its Table-5 value of $23,400 to $166,000. All of this suggests that social insurance, as modelled by the consumption floor, significantly affects precautionary motives.

Using the revised parameter estimates, we repeat the policy experiments shown in Table 4 with $C_{min} = \$100$. The results of these experiments are virtually identical to those for the baseline case. Recall that in the baseline case, increasing the normal Social Security retirement age leads to an additional 0.4 years of work during ages 60-67, and increasing the

---

[38]Our treatment of consumption floors differs markedly from that of Rust and Phelan (1997), who simply impose a penalty when an individual's implied consumption is negative. Although Rust and Phelan's estimates do not translate into a consumption floor, they find the penalty to be large, implying a fairly low floor. Since Rust and Phelan assume that consumption equals income net of health costs, it is unclear what their penalties imply for asset accumulation.

Medicare eligibility age leads to an additional 0.1 years of work. When $C_{min} = \$100$, the increases are 0.36 and 0.08 years, respectively, for Social Security and Medicare. Given that changes in C_{min} lead to offsetting changes in the estimated value of ν, these similarities are not surprising.

8.3   Mismeasured Assets

Because the moments we match in estimation include asset quantiles and asset-quantile-conditional participation rates, measurement error could affect our results. To explore this possibility, we re-estimate the model while assuming that observed assets for individual i at time t, A_{it}^*, relate to his actual assets, A_{it}, through

$$A_{it}^* = A_{it} \exp(\vartheta_{it}),$$

where ϑ_{it} is a normally-distributed, zero-mean measurement noise variable that is independent across individuals and time and independent of A_{it}. In particular, we generate simulated asset histories as before, multiply them by simulated sequences of $\exp(\vartheta_{it})$, and use the modified asset histories to construct our moment conditions.[39]

The fifth column of Table 3 shows the parameter estimates that result when the standard deviation of ϑ_{it} is 20%, an amount in the middle of the estimates discussed in Bound, et al. (2001). These parameter estimates are very similar to the baseline estimates contained in the first column. Moreover, the simulated profiles and the policy experiments are all very similar to the baseline case, showing that our results are robust to a significant amount of measurement error.

---

[39]We also modify the initial distribution of assets, which is based on the HRS data, to account for measurement error, by taking logs and applying standard projection formulae.

9 Conclusion

Prior to age 65, many individuals receive health insurance only if they continue to work. At age 65, however, Medicare provides health insurance to almost everyone. Therefore, a potentially important work incentive disappears at age 65. If individuals place a high value on health insurance, the provision of Medicare benefits may have a large effect on retirement behavior. To see if this is the case, we construct and estimate a retirement model that includes health insurance, uncertain medical costs, a savings decision, a non-negativity constraint on assets and a government-provided consumption floor. Including all these features produces a general model that can reconcile previous results.

Using data from the Health and Retirement Study, we find that Medicare has limited effects. Empirically, we find that workers whose employer-provided coverage continues after they leave their jobs are more likely to exit the labor force at age 62. This suggests that *tied* access to employer-provided health insurance might compel some workers to stay on the job. We also find, however, that age-65 job exit rates differ little across health insurance types.

The results from our model yield a similar conclusion. Although we find that changing the Medicare eligibility age does have a fairly significant effect on workers whose health insurance is tied to their jobs, we find that changing the Social Security rules has an even bigger effect. Moreover, we find that workers that can save value employer-provided health insurance mainly because it reduces average medical expenses, not because it reduces health cost uncertainty. Given that the actuarial value of employer-provided health insurance is fairly small, this result also suggests that the retirement incentives of employer-provided health insurance, although important, are modest.

References

[1] **Aaronson, D., and E. French,** "The Effect of Part-Time Work on Wages: Evidence from the Social Security Rules," *Journal of Labor Economics,* 2004, forthcoming.

[2] **Altonji, J., and L. Segal**, "Small Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business and Economic Statistics*, 1996, 14(3), 353-366.

[3] **Blau, D.,** "Labor Force Dynamics of Older Men," *Econometrica*, 1994, 62(1), 117-156.

[4] **Blau, D. and D. Gilleskie,** "Retiree Health Insurance and the Labor Force Behavior of Older Men in the 1990's," *Review of Economics and Statistics*, 2001, 83(1), 64-80.

[5] **Blau, D. and D. Gilleskie,** "Health Insurance and Retirement of Married Couples," mimeo, 2003.

[6] **The Boards of Trustees of the Hospital Insurance and Supplementary Medical Insurance Trust Funds,** *2003 Annual Report of the Boards of Trustees of the Hospital Insurance and Supplementary Medical Insurance Trust Funds*, 2003.

[7] **Bound, J., C. Brown and N. Mathiowetz,** "Measurement Error in Survey Data," in J. Heckman and E. Leamer (eds.) *Handbook of Econometrics*, Vol. 5., 2001.

[8] **Bound, J., T. Stinebrickner, and T. Waidmann,** "Health, Economic Resources, and the Work Decisions of Older Men," mimeo, 2003.

[9] **Buchinsky, M.,** "Recent Advances in Quantile Regression Models: A Practical Guideline for Empirical Research," *Journal of Human Resources*, 1998, 33, 88-126.

[10] **Cagetti, M.,** "Wealth Accumulation Over the Life Cycle and Precautionary Savings," *Journal of Business and Economic Statistics,* 2003, 21(3), 339-353.

[11] **Cochrane, J.,** "A Simple Test of Consumption Insurance," *The Journal of Political Economy*, 1991, 99(5), 957-976.

[12] **Committee On Ways And Means, U.S. House Of Representatives,** *2000 Green Book*, United States Government Printing Office, 2000.

[13] **Cogan, J.,** "Fixed Costs and Labor Supply," *Econometrica*, 1981, 49(4), 945-963.

[14] **De Nardi, C.,** "Wealth Distribution, Intergenerational Links and Estate Taxation," *Review of Economic Studies,* 2004, forthcoming.

[15] **Duffie, D. and K. Singleton,** "Simulated Moments Estimation of Markov Models of Asset Prices," *Econometrica*, July 1993, 61(4), 929-952.

[16] **Employee Benefit Research Institute,** *EBRI Health Benefits Databook*, EBRI-ERF, 1999.

[17] **Epple, D. and Holger Seig,** "Estimating Equilibrium Models of Local Jurisdictions," *Journal of Political Economy* 1999. 107(4), 645-81.

[18] **Feenberg, D., and J. Skinner,** "The Risk and Duration of Catastrophic Health Care Expenditures," *The Review of Economics and Statistics* 1994. 76(4), 633-647.

[19] **French, E.,** "The Labor Supply Response to Predictable (but Mismeasured) Wage Changes", *Review of Economics and Statistics,* 2004, forthcoming.

[20] **French, E.,** "The Effects of Health, Wealth and Wages on Labor Supply and Retirement Behavior," manuscript, 2003.

[21] **French, E., and J. Jones,** "On the Distribution and Dynamics of Health Care Costs," *Journal of Applied Econometrics,* 2004, forthcoming.

[22] **Gourieroux, C., and A. Monfort,** *Simulation-Based Econometric Methods*, Oxford University Press, 1997.

[23] **Gourinchas, P. and Parker, J.,** "Consumption Over the Life Cycle," *Econometrica*, 2002, 70(1), 47-89.

[24] **Gruber, J., and B. Madrian,** "Health Insurance Availibility and the Retirement Decision," *American Economic Review*, 1995, 85(4), 938-948.

[25] **Gruber, J., and B. Madrian,** "Health Insurance and Early Retirement: Evidence from the Availability of Continuation Coverage," in D.A. Wise, ed., *Advances in the Economics of Aging* 1996, University of Chicago Press, 115-143.

[26] **Gustman, A., and T. Steinmeier,** "A Structural Retirement Model," *Econometrica*, 1986, 54(3), 555-584.

[27] **Gustman, A., and T. Steinmeier,** "Employer-Provided Health Insurance and Retirement Behavior," *Industrial and Labor Relations Review* 1994, 48(1), 124-140.

[28] **Gustman, A., and T. Steinmeier,** "Effects of Pensions on Savings: Analysis with Data from the Health and Retirement Study," *Carnegie-Rochester Conference Series on Public Policy*, 1999, 50, 271-324.

[29] **Gustman, A., and T. Steinmeier,** "The Social Security Early Entitlement Age in a Structural Model of Retirement and Wealth," NBER Working Paper 9183, 2002.

[30] **Gustman, A., O. Mitchell, A. Samwick and T. Steinmeier,** "Evaluating Pension Entitlements,", 1998, mimeo.

[31] **Hubbard, R., J. Skinner, and S. Zeldes,** "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving", *Carnegie-Rochester Series on Public Policy*, 1994, 40, 59-125.

[32] **Hubbard, R., J. Skinner, and S. Zeldes,** "Precautionary Saving and Social Insurance", *Journal of Political Economy*, 1995, 103(2), 360-399.

[33] **Judd, K.,** *Numerical Methods in Economics*, MIT Press, 1998.

[34] **Kahn, J.,** "Social Security, Liquidity, and Early Retirement", *Journal of Public Economics* , 1988, 35, 97-117.

[35] **Karoly, L., and J. Rogowski,** "The Effect of Access to Post-Retirement Health Insurance on the Decision to Retire Early," *Industrial and Labor Relations Review* 1994, 48(1), 103-123.

[36] **Keane, M., and K. Wolpin,** "The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment", *International Economic Review* 2001, 42(4), 1051-1103.

[37] **Low, H.,** "Self-Insurance and Unemployment Benefit in a Life-Cycle Model of Labour Supply and Savings", mimeo, 2002.

[38] **Lumsdaine, R., J. Stock, and D. Wise,** "Pension Plan Provisions and Retirement: Men, Women, Medicare and Models," in D. Wise (ed.) *Studies in the Economics of Aging* 1994.

[39] **MaCurdy, T.,** "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy* 1981, 89(6), 1059-1085.

[40] **Madrian, B.,** "The Effect of Health Insurance on Retirement," *Brookings Papers on Economic Activity,* 1994, 181-252.

[41] **Newey, W. and D. McFadden,** "Large Sample Estimation abd Hypothesis Testing" in R. Engle and D. McFadden (eds.) *Handbook of Econometrics*, Vol. 4., 1994.

[42] **Pakes,A., and D. Pollard,** "Simulation and the Aysmptotics of Optimization Estimators," *Econometrica*, 1989, 57(5), 1027-1057.

[43] **Palumbo, M.,** "Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle," *Review of Economic Studies*, 1999, 66(2), 395-421.

[44] **Pischke, J-S** , "Measurement Error and Earnings Dynamics: Some Estimates From the PSID Validation Study," *Journal of Business & Economics Statistics*, 1995, 13(3), 305-314.

[45] **Powell, J.,** "Estimation of Semiparametric Models" in R. Engle and D. McFadden (eds.) *Handbook of Econometrics*, Vol. 4., 1994.

[46] **Ruhm, C.,** "Bridge Jobs and Partial Retirement," *Journal of Labor Economics*, 1990, 8, 482-501.

[47] **Rust, J. and C. Phelan,** "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets," *Econometrica*, 1997, 65(4), 781-831.

[48] **Rust, J., Buchinsky, M., and H. Benitez-Silva,** "An Empirical Model of Social Insurance at the End of the Life Cycle," mimeo, 2003.

[49] **Smith, J.,** "Healthy Bodies and Thick Wallets: The Dual Relationship between Health and Economic Status," *Journal of Economic Perspectives*, 13(2), 1999, 145-166.

[50] **United States Social Security Administration,** *Social Security Bulletin: Annual Statistical Supplement*, United States Government Printing Office, selected years.

[51] **van der Klaauw, W., and K. Wolpin,** "Social Security, Pensions and the Savings and Retirement Behavior of Households," mimeo, 2002.

Appendix A: Taxes

Individuals pay federal, state, and payroll taxes on income. We compute federal taxes on income net of state income taxes using the Federal Income Tax tables for "Head of Household" in 1998. We use the standard deduction, and thus do not allow individuals to defer medical expenses as an itemized deduction. We also use income taxes for the fairly representative state of Rhode Island (27.5% of the Federal Income Tax level). Payroll taxes are 7.65% up to a maximum of $68,400, and are 1.45% thereafter. Adding up the three taxes generates the following level of post tax income as a function of labor and asset income:

Pre-tax Income (Y)	Post-Tax Income	Marginal Tax Rate
0-6250	0.9235Y	0.0765
6250-40200	5771.88 + 0.7384(Y-6250)	0.2616
40200-68400	30840.56 + 0.5881(Y-40200)	0.4119
68400-93950	47424.98 + 0.6501(Y-68400)	0.3499
93950-148250	64035.03 + 0.6166(Y-93950)	0.3834
148250-284700	97515.41 + 0.5640(Y-148250)	0.4360
284700+	174474.21 + 0.5239(Y-284700)	0.4761

Table 8: AFTER TAX INCOME

Appendix B: Pensions

The fundamental equation behind our calculation of pension benefits is the accumulation equation for pension wealth, pw_t:

$$pw_{t+1} = \begin{cases} (1/s_{t+1})[(1+r)pw_t + pacc_t - pb_t] & \text{if living at } t+1 \\ 0 & \text{otherwise} \end{cases} \tag{21}$$

where $pacc_t$ is pension accrual and pb_t is pension benefits. Two features of this equation bear noting. First, workers cannot bequeath their pensions. It immediately follows that in order to be actuarially fair, surviving workers must receive an above-market return on their pension balances. Therefore, we should divide next period's pension wealth by the survival probability s_{t+1} in equation (21). Note that $E[pw_{t+1}|pw_t, pacc_t, pb_t] = [(1+r)pw_t + pacc_t - pb_t]$ because with probability s_{t+1}, $pw_{t+1} = 0$. Second, since pension accrual and pension interest are not directly taxed, the appropriate rate of return on pension wealth is the pre-tax one. Pension

benefits, on the other hand, are included in the income used to calculate an individual's income tax liability.

We calculate pension benefits by assuming that at age k, the worker receives the fraction pf_t of the actuarially fair annuity, $pb_t^{\max}$. pf_t is estimated as the fraction of respondents receiving pensions at each age; the fraction increases fairly smoothly, except for a 23-percentage-point jump at age 62. To find the annuity $pb_t^{\max}$, note first that recursively substituting equation (21) and imposing $pw_{T+1} = 0$ reveals that

$$pw_t = \frac{1}{1+r} \sum_{k=t}^{T} \frac{S(k,t)}{(1+r)^{k-t}} (pf_k pb_t^{\max} - pacc_k),$$

where $S(k,t) = (1/s_t) \prod_{j=t}^{k} s_j$ gives the probability of surviving to age k, conditional on having survived to time t. If we assume further that there is no more pension accrual, so that $pacc_k = 0$ $k = t, t+1, ..., T$, and that the maximum pension benefit is constant from time t forward, so that $pb_k^{\max} = pb_t^{\max}$, $k = t, t+1, ..., T$, this equation reduces to

$$pw_t = \Gamma_t pb_t^{\max}, \tag{22}$$

$$\Gamma_t \equiv \frac{1}{1+r} \sum_{k=t}^{T} \frac{S(k,t)}{(1+r)^{k-t}} pf_k. \tag{23}$$

Pension benefits are thus given by

$$pb_t = pf_t \Gamma_t^{-1} pw_t. \tag{24}$$

Pension accrual is given by

$$pacc = \alpha_0(HI_t, W_t H_t, t) \times W_t H_t, \tag{25}$$

where $\alpha_0(.)$ is the pension accrual rate as a function of health insurance type, labor income, and age. To construct $\alpha_0(.)$, we first find $\alpha_1(pen_t, t)$, which gives the pension accrual rate as a function of pension type, pen_t, and age. We then find the distribution of pension types for

each health insurance type and labor income level. The function $\alpha_0(.)$ is simply the average of $\alpha_1(pen_t, t)$ across this distribution:

$$\alpha_0(HI_t, W_t H_t, t) = \sum_{pen \in PT} \alpha_1(pen, t) \times \Pr(pen_t = pen | HI_t, W_t H_t, t)$$

where $PT = \{$none, defined benefit, defined contribution, combined defined benefit and defined contribution$\}$ is the set of possible pension types. The function $\alpha_1(.)$ is taken from Gustman and Steinmeier's (1999) analysis of HRS data.[40] The function $\Pr(pen_t = pen | HI_t, W_t H_t, t)$, which measures the probability that pension type is pen, is estimated using the regression function

$$\Pr(pen_t = pen | HI_t, W_t H_t, t) = \gamma_{0,pen} + \gamma_{1,pen} t + \gamma_{2,pen} \ln(W_t H_t) + \sum_k \gamma_{k,pen} 1\{HI_t = k\} \quad (26)$$

and HRS data. Note that equation (26) is estimated separately for each pension type.

Figure 6 shows pension accrual by health insurance type. Those with retiree coverage are the most likely to have a defined benefit pension plan or a combined plan. As a result, those with retiree coverage have the highest pension accrual in their 50s and the sharpest drops in pension accrual at ages 62 and 65. Figure 6 also shows the effect of having log income one standard deviation above the mean. Note that the effect of income on pension accrual is relatively small, once health insurance is accounted for.

To recapitulate, pension wealth follows equation (21), with pension accrual given by equation (25) and pension benefits given by equation (24). Using these equations, it is straightforward to track and record the pension balances of each simulated individual. To start the simulations, we assume that initial pension balances are a function of assets and wages:

$$pw_{53} = \eta_1 + \eta_2 A_{53} + \eta_3 A_{53}^2 + \eta_4 W_{53} + \eta_5 W_{53}^2 + \eta_6 A_{53} W_{53}, \quad (27)$$

---

[40]We first adjust their pension accrual profile by their assumed rate of wage growth so that pension accrual is measured in rates. We then smooth their pension accrual profile using a 20th order polynomial with dummy variables for age greater than 61, 62, 63, 64 and 65. Predicted accrual rates that are negative are set to zero.
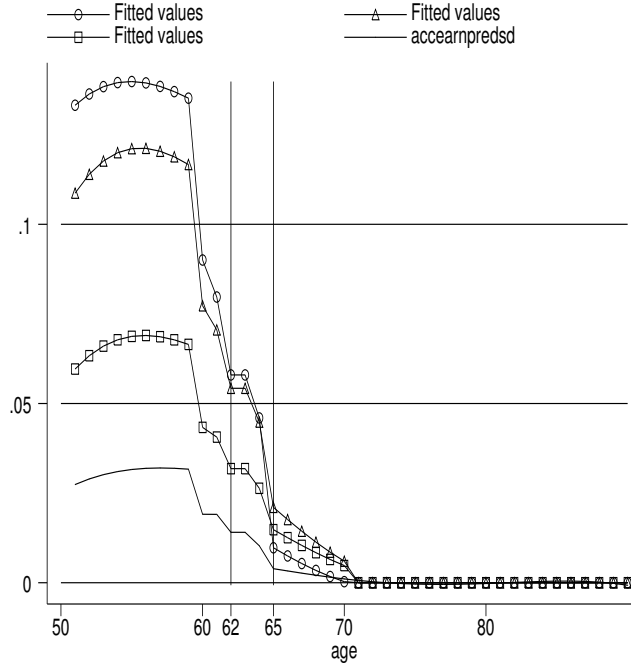
$\square$ = no health insurance $\bigcirc$ = retiree health insurance $\triangle$ = tied health insurance coverage
$-$ = log earnings one standard deviation above the mean

Figure 6: Pension Accrual Rates, by Age and Health Insurance Type

using a regression on data generated by the model in French (2003), where the simulated life histories begin at age 30. The combination of equations (25) and (27) appears to work fairly well; simulated mean pension wealth is roughly \$115,000 in 1998 dollars at age 57, which roughly coincides with Gustman and Steinmeier's (1999, Table 5) estimate for male heads of households of \$113,000.

But even though it is straightforward to use equation (21) when computing pension wealth in the simulations, it is too computationally burdensome to include pension wealth as a separate state variable when computing the decision rules. Our approach is to impute pension wealth as a function of age and AIME. In particular, we impute a worker's annual pension benefits as a function of his Social Security benefits:

$$\widehat{pb}_t(PIA_t, HI_t, t) = \sum_k \gamma_{0,k} 1\{HI_t = k\} + \gamma_1 t + \gamma_2 \max\{0, t - 70\} + \gamma_3 PIA_t + \qquad (28)$$
$$\gamma_4 \max\{0, PIA_t - 14,359.9\} + \gamma_5 t PIA_t + \gamma_6 \max\{0, t - 70\} PIA_t,$$

54

where PIA_t is the Social Security benefit the worker would get if he were drawing benefits at time t; as shown in Appendix C below, PIA is a simple monotonic function of AIME. Applying equation (22) yields imputed pension wealth, $\widehat{pw}_t = \Gamma_t \widehat{pb}_t$. The coefficients of this equation were estimated with a regression on data generated by the model in this paper. Since these simulated data depend on the γs—$\widehat{pw}_t$ affects the decision rules used in the simulations—the γs solve a fixed-point problem. Fortunately, estimates of the γs converged after a few iterations.

This imputation process raises two complications. The first is that we use a different pension wealth imputation formula when calculating decision rules than we do in the simulations. If an individual's time-t pension wealth is $\widehat{pw}_t$, his time-$t+1$ pension wealth (if living) should be

$$\widehat{\widehat{pw}}_{t+1} = (1/s_{t+1})[(1+r)\widehat{pw}_t + pacc_t - pb_t].$$

This quantity, however, might differ from the pension wealth that would be imputed using PIA_{t+1}, $\widehat{pw}_{t+1} = \Gamma_{t+1}\widehat{pb}_{t+1}$ where $\widehat{pb}_{t+1}$ is defined in equation (28). To correct for this, we increase non-pension wealth, A_{t+1}, by $s_{t+1}(1 - \tau_t)(\widehat{\widehat{pw}}_{t+1} - \widehat{pw}_{t+1})$. The first term in this expression reflects the fact that while non-pension assets can be bequeathed, pension wealth cannot. The second term, $1 - \tau_t$, reflects the fact that pension wealth is a pre-tax quantity —pension benefits are more or less wholly taxable—while non-pension wealth is post-tax—taxes are levied only on interest income.

A second problem is that while an individual's Social Security application decision affects his annual Social Security benefits, it should not affect his pension benefits. Recall, however, that we reduce PIA if an individual draws benefits before age. The pension imputation procedure we use, however, would imply that it does. We counter this problem by recalculating PIA when the individual begins drawing Social Security benefits. In particular, suppose that a decision to accelerate or defer application changes PIA_t to $rem_t PIA_t$. Our approach is to use equation (29) find a value PIA_t^* such that

$$(1 - \tau_t)\widehat{pb}_t(PIA_t^*) + PIA_t^* = (1 - \tau_t)\widehat{pb}_t(PIA_t) + rem_t PIA_t,$$

so that the change in the sum of PIA and imputed after-tax pension income equals just the change in PIA, i.e., $(1 - rem_t)PIA_t$.

Appendix C: Computation of AIME

Social Security benefits are based on the individual's 35 highest earnings years, relative to average wages in the economy during those years. The average earnings over these 35 highest earnings years are called Average Indexed Monthly Earnings, or AIME. It immediately follows that working an additional year increases the AIME of an individual with less than 35 years of work. If an individual has already worked 35 years, he can still increase his AIME by working an additional year, but only if his current earnings are higher than the lowest earnings embedded in his current AIME. To account for real wage growth, earnings in earlier years are inflated by the growth rate of average earnings in the overall economy. For the period 1992-1999, real wage growth, g, had an average value of 0.016 (Committee on Ways and Means, 2000, p. 923). This indexing stops at the year the worker turns 60, however, and earnings accrued after age 60 are not rescaled. [41] Lastly, AIME is capped. In 1998, the base year for the analysis, the maximum AIME level was $68,400 in 1998 dollars.

Precisely modelling these mechanics would require us to keep track of a worker's entire earnings history, which is computationally infeasible. As an approximation, we assume that (for workers beneath the maximum) annualized AIME is given by

$$AIME_{t+1} = (1 + g \times 1\{t \leq 60\})AIME_t \tag{29}$$
$$+ \frac{1}{35} \max\left\{0, \ W_t H_t - \alpha_t(1 + g \times 1\{t \leq 60\})AIME_t\right\},$$

where $1 - \alpha_t$ is the probability that time t earnings increases $AIME_t$. We assume that 20% of the workers enter the labor force each year between ages 21 and 25, so that $\alpha_t = 0$ for workers aged 55 and younger. For workers aged 60 and older, earnings only update $AIME_t$ if current earnings replace the lowest year of earnings, so we estimate α_t by simulating wage (not earnings) histories with the model developed in French (2003), calculating the

---

[41] After age 62, benefits increase at the rate of inflation.

sequence $\{\alpha_t\}_{t \geq 60}$ for each simulated wage history, and taking averages at each age. Linear interpolation yields α_{56} through α_{59}.

For the simulations, each person's AIME history is initialized as a function of assets and wages:

$$AIME_{53} = \eta_1 + \eta_2 A_{53} + \eta_3 A_{53}^2 + \eta_4 W_{53} + \eta_5 W_{53}^2 + \eta_6 A_{53} W_{53} + \nu_{AIME},$$

where ν_{AIME} is drawn from a random number generator. The η's and the standard deviation of ν_{AIME} are taken from a regression on 1996 PSID data where $AIME_{53}$ is calculated using PSID earnings histories, rescaled to real wage levels in 1992.

AIME is converted into a Primary Insurance Amount (PIA) using the formula

$$PIA_t = \begin{cases} 0.9 \times AIME_t & \text{if } AIME_t < \$5,724 \\ \$5,151.6 + 0.32 \times AIME_t & \text{if } \$5,724 \leq AIME_t < \$34,500 \\ \$14,359.9 + 0.15 \times AIME_t & \text{if } AIME_t \geq \$34,500 \end{cases} \tag{30}$$

Social Security benefits ss_t depend both upon the age at which the individual first receives Social Security benefits and the Primary Insurance Amount. For example, pre-Earnings Test benefits for a Social Security beneficiary will be equal to PIA if the individual first receives benefits at age 65. For every year before age 65 the individual first draws benefits, benefits are reduced by 6.67% and for every year (up until age 70) that benefit receipt is delayed, benefits increase by 5.0%.[42]

Appendix D: Numerical Methods

Because the model has no closed form solution, the decision rules it generates must be found numerically. We find the decision rules using value function iteration, starting at time T and working backwards to time 1. We find the time-T decisions by maximizing equation

---

[42]The effects of early or late application can be modelled as changes in AIME rather than changes in PIA, eliminating the need to include age at application as a state variable. For example, if an individual begins drawing benefits at age 62, his adjusted AIME must result in a PIA that is only 80% of the PIA he would have received had he first drawn benefits at age 65. Using equation (30), this is easy to find.

(16) at each value of X_T, with $V_{T+1} = b(A_{T+1})$. This yields decision rules for time T and the value function V_T. We next find the decision rules at time $T-1$ by solving equation (16) with V_T. Continuing this backwards induction yields decision rules for times $T - 2, T - 3, ..., 1$. The value function is directly computed at a finite number of points within a grid, $\{X_i\}_{i=1}^{I}$;[43] we use linear interpolation within the grid and linear extrapolation outside of the grid to evaluate the value function at points that we do not directly compute. Because changes in assets and AIME are likely to cause larger behavioral responses at low levels of assets and AIME, the grid is more finely discretized in this region.

At time t, wages, medical expenses and assets at time $t + 1$ will be random variables. To capture uncertainty over the persistent components of medical expenses and wages, we convert ζ_t and ω_{t+1} into discrete Markov chains, and calculate the conditional expectation of V_{t+1} accordingly.[44] We integrate the value function with respect to the transitory component of medical expenses, ξ_t, using 5-node Gauss-Hermite quadrature (see Judd, 1999).

Because of the fixed time cost of work and the discrete benefit application decision, the value function need not be globally concave. This means that we cannot find a worker's optimal consumption and hours with fast hill climbing algorithms. Our approach is to discretize the consumption and labor supply decision space and to search over this grid. Experimenting with the fineness of the grids suggested that the grids we used produced reasonable approximations.[45] In particular, increasing the number of grid points seemed to have a small effect

---

[43]In practice, the grid consists of: 32 asset states, $A_h \in [-\$55,000, \$1,200,000]$; 5 wage residual states, $\omega_i \in [-0.99, 0.99]$; 16 AIME states, $AIME_j \in [\$4,000, \$68,400]$; 3 health cost states, ζ_k, over a normalized (unit variance) interval of $[-1.5, 1.5]$. There are also two application states and two health states. This requires solving the value function at 30,720 different points for ages 62-69, when the individual is eligible to apply for benefits, and at 15,360 points before age 62 or after age 69 (when we impose application).

[44]Using discretization rather than quadrature greatly reduces the number of times one has to interpolate when calculating $E_t(V(X_{t+1}))$.

[45]We use two search grids. Initially, the consumption grid has 100 points, and the hours grid is broken into 300-hour intervals. When this initial, coarser, grid is used, the consumption search at a value of the state vector X for time t is centered around the consumption gridpoint that was optimal for the same value of X at time $t + 1$. (Recall that we solve the model backwards in time.) If the search yields a maximizing value near the edge of the search grid, the grid is reoriented and the search continued. We begin our search for optimal hours at the level of hours that sets the marginal rate of substitution between consumption and leisure equal to the wage. We then try 6 different hours choices in the neighborhood of the initial hours guess. Because of the fixed cost of work, we also evaluate the value function at $H_t = 0$, searching around the consumption choice that was optimal when $H_{t+1} = 0$. The second grid centers the search for C_t and H_t around the values of C_t and H_t that were optimal in some earlier search, and looks over a finer grid.

on the computed decision rules.

We then use the decision rules to generate simulated time series. Given the realized state vector X_{i0}, individual i's realized decisions at time 0 are found by evaluating the time-0 decision functions at X_{i0}. Using the transition functions given by equations (4) through (6), we combine X_{i0}, the time-0 decisions, and the individual i's time-1 shocks to get the time-1 state vector, X_{i1}. Continuing this forward induction yields a life cycle history for individual i. When X_{it} does not lie exactly on the state grid, we use interpolation or extrapolation to calculate the decision rules. This is true for ζ_t and ω_t as well. While these processes are approximated as finite Markov chains when the decision rules are found, the simulated sequences of ζ_t and ω_t are generated from continuous processes. This makes the simulated life cycle profiles less sensitive to decision rules at particular values of ζ_t and ω_t than when ζ_t and ω_t are drawn from Markov chains.

Appendix E: Moment Conditions and the Asymptotic Distribution of Parameter Estimates

We assume that the "true" preference vector θ_0 lies in the interior of the compact set $\Theta \subset \mathbb{R}^7$. Our estimate, $\hat{\theta}$, is the value of θ that minimizes the (weighted) distance between the estimated life cycle profiles for assets, hours, and participation found in the data and the simulated profiles generated by the model. We match $15T$ moment conditions. They are, for each age $t \in \{1, ..., T\}$, two asset quantiles (forming $2T$ moment conditions), labor force participatipon rates conditional on asset quantile and health insurance type (forming $9T$ moment conditions), labor force participation rates conditional upon health status (forming $2T$ moment conditions), and mean hours worked conditional upon health status (forming $2T$ moment conditions).

Consider first the asset quantiles. As stated in the main text, let $j \in \{1, 2, ..., J\}$ index asset quantiles, where J is the total number of asset quantiles. Assuming that the age-conditional distribution of assets is continuous, the π_j-th age-conditional quantile of measured

assets, $Q_{\pi_j}(A_{it}, t)$, is defined as

$$\Pr\left(A_{it} \leq Q_{\pi_j}(A_{it}, t) | t\right) = \pi_j.$$

In other words, the fraction of individuals with less than Q_{π_j} in assets is π_j. Therefore, $Q_{\pi_j}(A_{it}, t)$ is the data analog to $g_{\pi_j}(t; \theta_0, \chi_0)$, the model-predicted quantile. Using the indicator function, the definition of π_j-th conditional quantile can be rewritten as

$$E\left(1\{A_{it} \leq Q_{\pi_j}(A_{it}, t)\} | t\right) = \pi_j. \tag{31}$$

If the model is true then the data quantile in equation (31) can be replaced by the model quantile, so that equation (31) can be rewritten as:

$$E\left(1\{A_{it} \leq g_{\pi_j}(t; \theta_0, \chi_0)\} - \pi_j | t\right) = 0, \quad j \in \{1, 2, ..., J\}, \ t \in \{1, ..., T\}. \tag{32}$$

Equation (32) is merely equation (17) in the text. While equation (32) is a departure from the usual practice of minimizing a sum of weighted absolute errors in quantile estimation, the quantile restrictions just described are part of a larger set of moment conditions. This means that we can no longer estimate θ by minimizing weighted absolute errors, if only because we are considering multiple quantiles.[46]

The next set of moment conditions use the quantile-conditional means of labor force participation. Let $\overline{P}_j(HI, t; \theta_0, \chi_0)$ denote the model's prediction of labor force participation given asset quantile interval j, health insurance type HI, and age t. If the model is true, $\overline{P}_j(HI, t; \theta_0, \chi_0)$ should equal the conditional participation rates found in the data:

$$\overline{P}_j(HI, t; \theta_0, \chi_0) = E[P_{it} \mid HI, t, g_{\pi_{j-1}}(t; \theta_0, \chi_0) \leq A_{it} \leq g_{\pi_j}(t; \theta_0, \chi_0), \tag{33}$$

---

[46]A slightly different approach to handling multiple quantiles is the minimum distance framework developed in Epple and Seig (1999). Buchinsky (1998) shows that one could include the first-order conditions from an absolute value minimization problem in the moment set. However, his approach involves finding the gradients of $g_{\pi_j}(t; \theta, \chi)$ at each step of the minimization search.

with $\pi_0 = 0$ and $\pi_{J+1} = 1$. Equation (33) is equivalent to equation (18) in the text. Using indicator function notation, we can convert the conditional moment equation given by equation (33) into an unconditional one:

$$E([P_{it} - \overline{P}_j(HI, t; \theta_0, \chi_0)] \times 1\{HI_{it} = HI\}$$
$$\times 1\{g_{\pi_{j-1}}(t; \theta_0, \chi_0) \leq A_{it} \leq g_{\pi_j}(t; \theta_0, \chi_0)\} \mid t) = 0, \tag{34}$$

for $j \in \{1, 2, ..., J+1\}$, $HI \in \{none, ret, tied\}$, $t \in \{1, ..., T\}$. Note that $g_{\pi_0}(t) \equiv -\infty$ and $g_{\pi_{J+1}}(t) \equiv \infty$.

Lastly, the moment conditions for labor force participation and hours worked conditional upon health status and age are those described in equations (19) and (20) of the text, converted into unconditional moment equations with indicator functions. Combining all the moment conditions described here is straightforward: we simply stack the moment conditions and estimate jointly.

Suppose we have a data set of I independent individuals that are each observed for T periods. Let $\varphi(\theta; \chi_0)$ denote the $15T$-element vector of moment conditions that was described in the main text and immediately above, and let $\hat{\varphi}_I(.)$ denote its sample analog. Note that we can extend our results to an unbalanced panel, as we must do in the empirical work, by simply allowing some of the individual's contributions to $\varphi(.)$ to be "missing", as in French and Jones (2003). Letting $\widehat{\mathbf{W}}_I$ denote a $15T \times 15T$ weighting matrix, the MSM estimator $\hat{\theta}$ is given by

$$\arg\min_{\theta} \frac{I}{1+\tau} \hat{\varphi}_I(\theta, \chi_0)' \widehat{\mathbf{W}}_I \hat{\varphi}_I(\theta, \chi_0),$$

where τ is the ratio of the number of observations to the number of simulated observations.

Under the regularity conditions stated in Pakes and Pollard (1989) and Duffie and Singleton (1993), the MSM estimator $\hat{\theta}$ is both consistent and asymptotically normally distributed:

$$\sqrt{I}(\hat{\theta} - \theta_0) \rightsquigarrow N(0, \mathbf{V}),$$

with the variance-covariance matrix $\mathbf{V}$ given by

$$\mathbf{V} = (1+\tau)(\mathbf{D}'\mathbf{W}\mathbf{D})^{-1}\mathbf{D}'\mathbf{W}\mathbf{S}\mathbf{W}\mathbf{D}(\mathbf{D}'\mathbf{W}\mathbf{D})^{-1},$$

where: $\mathbf{S}$ is the variance-covariance matrix of the data;

$$\mathbf{D} = \left.\frac{\partial\varphi(\theta,\chi_0)}{\partial\theta}\right|_{\theta=\theta_0} \tag{35}$$

is the gradient matrix of the population moment vector; and $\mathbf{W} = \text{plim}_{\to\infty}\{\widehat{\mathbf{W}}_I\}$. Moreover, Newey (1985) shows that if the model is properly specified,

$$\frac{I}{1+\tau}\hat{\varphi}_I(\hat{\theta},\chi_0)'\mathbf{R}^{-1}\hat{\varphi}_I(\hat{\theta},\chi_0) \rightsquigarrow \chi^2_{15T-7},$$

where $\mathbf{R}^{-1}$ is the generalized inverse of

$$\begin{aligned}
\mathbf{R} &= \mathbf{P}\mathbf{S}\mathbf{P}, \\
\mathbf{P} &= \mathbf{I} - \mathbf{D}(\mathbf{D}'\mathbf{W}\mathbf{D})^{-1}\mathbf{D}'\mathbf{W}.
\end{aligned}$$

The asymptotically efficient weighting matrix arises when $\widehat{\mathbf{W}}_I$ converges to $\mathbf{S}^{-1}$, the inverse of the variance-covariance matrix of the data. When $\mathbf{W} = \mathbf{S}^{-1}$, $\mathbf{V}$ simplifies to $(1+\tau)(\mathbf{D}'\mathbf{S}^{-1}\mathbf{D})^{-1}$, and $\mathbf{R}$ is replaced with $\mathbf{S}$. But even though the optimal weighting matrix is asymptotically efficient, it can be severely biased in small samples. (See, for example, Altonji and Segal, 1996.) We thus use a "diagonal" weighting matrix, as suggested by Pischke (1995). The diagonal weighting scheme uses the inverse of the matrix that is the same as $\mathbf{S}$ along the diagonal and has zeros off the diagonal of the matrix.

We estimate $\mathbf{D}$, $\mathbf{S}$ and $\mathbf{W}$ with their sample analogs. For example, our estimate of $\mathbf{S}$ is the $15T \times 15T$ estimated variance-covariance matrix of the sample data. That is, a typical diagonal element of $\widehat{\mathbf{S}}_I$ is the variance estimate $\frac{1}{I}\sum_{i=1}^{I}[1\{A_{it} \le Q_{\pi_j}(A_{it},t)\} - \pi_j]^2$, while a typical off-diagonal element is a covariance. When estimating preferences, we use

sample statistics, so that $Q_{\pi_j}(A_{it}, t)$ is replaced with the sample quantile $\widehat{Q}_{\pi_j}(A_{it}, t)$. When computing the chi-square statistic and the standard errors, we use model predictions, so that Q_{π_j} is replaced with its simulated counterpart, $g_{\pi_j}(t; \hat{\theta}, \hat{\chi})$. Covariances between asset quantiles and hours and labor force participation are also simple to compute.

The gradient in equation (35) is straightforward to estimate for hours worked and participation conditional upon age and health status; we merely take numerical derivatives of $\hat{\varphi}_I(.)$. However, in the case of the asset quantiles and labor force participation, discontinuties make the function $\hat{\varphi}_I(.)$ non-differentiable at certain data points. Therefore, our results do not follow from the standard GMM approach, but rather the approach for non-smooth functions described in Pakes and Pollard (1989), Newey and McFadden (1994, section 7) and Powell (1994). We find the asset quantile component of $\mathbf{D}$ by rewriting equation (32) as

$$F(g_{\pi_j}(t; \theta_0, \chi_0)|t) - \pi_j = 0,$$

where $F(g_{\pi_j}(t; \theta_0, \chi_0)|t)$ is the c.d.f. of time-t assets evaluated at the π_j-th quantile. Differentiating this equation yields:

$$\mathbf{D}'_{jt} = f(g_{\pi_j}(t; \theta_0, \chi_0)|t) \frac{\partial g_{\pi_j}(t; \theta_0, \chi_0)}{\partial \theta'}. \tag{36}$$

In practice we find $f(g_{\pi_j}(t; \theta_0, \chi_0)|t)$, the p.d.f. of time-$t$ assets evaluated at the π_j-th quantile, with a kernel density estimator.

To find the component of the matrix $\mathbf{D}'$ for the asset-conditional labor force participation rates, it is helpful to write equation (34) as

$$\Pr(HI_t = HI) \times \int_{g_{\pi_{j-1}}(t; \theta_0, \chi_0)}^{g_{\pi_j}(t; \theta_0, \chi_0)} \left[E(P_{it}|A_{it}, HI, t) - \overline{P}_j(HI, t; \theta_0, \chi_0) \right] f(A_{it}|HI, t) dA_{it} = 0,$$

which implies that

$$\mathbf{D}'_{jt} = \Bigg[- \Pr(g_{\pi_{j-1}}(t;\theta_0,\chi_0) \le A_{it} \le g_{\pi_j}(t;\theta_0,\chi_0)|HI,t)\frac{\partial \overline{P}_j(HI,t;\theta_0,\chi_0)}{\partial \theta'}$$

$$+ [E(P_{it}|g_{\pi_j}(t;\theta_0,\chi_0),HI,t) - \overline{P}_j(HI,t;\theta_0,\chi_0)]f(g_{\pi_j}(t;\theta_0,\chi_0)|HI,t)\frac{\partial g_{\pi_j}(t;\theta_0,\chi_0)}{\partial \theta'}$$

$$- [E(P_{it}|g_{\pi_{j-1}}(t;\theta_0,\chi_0),HI,t) - \overline{P}_j(HI,t;\theta_0,\chi_0)]f(g_{\pi_{j-1}}(t;\theta_0,\chi_0)|HI,t)\frac{\partial g_{\pi_{j-1}}(t;\theta_0,\chi_0)}{\partial \theta'} \Bigg]$$

$$\times \Pr(HI_t = HI), \tag{37}$$

with $f(g_{\pi_0}(t;\theta_0,\chi_0)|HI,t)\frac{\partial g_{\pi_0}(t;\theta_0,\chi_0)}{\partial \theta'} = f(g_{\pi_{J+1}}(t;\theta_0,\chi_0)|HI,t)\frac{\partial g_{\pi_{J+1}}(t;\theta_0,\chi_0)}{\partial \theta'} \equiv 0$.

Appendix F: Data Appendix

Our data are drawn from HRS, a sample of non-institutionalized individuals aged 51-61 in 1992. With the exception of assets and health costs, which are measured at the household level, our data are for male household heads. The HRS surveys individuals every two years; we have 5 waves of data covering the period 1992-2000.

The variables used in our analysis are constructed as follows. Hours of work are the product of usual hours per week and usual weeks per year. To compute hourly wages, the respondent is asked about how they are paid, how often they are paid, and how much they are paid. If the worker is salaried, for example, annual earnings are the product of pay per period and the number of pay periods per year. The wage is then annual earnings divided by annual hours. If the worker is hourly, we use his reported hourly wage. We treat a worker's hours for the non-survey (e.g. 1993) years as missing.

For survey years the individual is considered in the labor force if he reports working over 300 hours per year. The HRS also asks respondents retrospective questions about their work history. Because we are particularly interested in labor force participation, we use the work history to construct a measure of whether the individual worked in non-survey years. For example, if an individual withdraws from the labor force between 1992 and 1994, we use the 1994 interview to infer whether the individual was working in 1993.

The HRS has a comprehensive asset measure. It includes the value of housing, other real

estate, autos, liquid assets (which includes money market accounts, savings accounts, T-bills, etc.), IRAs, stocks, business wealth, bonds, and "other" assets, less the value of debts. For non-survey years, we assume that assets take on the value reported in the preceding year. This implies, for example, that we use the 1992 asset level as a proxy for the 1993 asset level. Given that wealth changes rather slowly over time, these imputations should not severely bias our results.

To measure health status we use responses to the question: "would you say that your health is excellent, very good, good, fair, or poor?" We consider the individual in bad health if he responds "fair" or "poor", and consider him in good health otherwise.[47] We treat the health status for non-survey years as missing. Appendix H describes how we construct the health insurance indicator.

Problems of missing information on health insurance and on retrospective labor force participation are severe in the HRS. When estimating labor force participation by asset quantile and health insurance, we lose 1,734 person-year observations due to missing participation, 6,317 observations due to missing health insurance data, and 292 observations due to missing asset data. Of a potential sample of 39,701 person-year observations for those between ages 51 and 69, we keep 31,358 observations.

To generate the initial joint distribution of assets, wages, AIME, pensions, health status and health costs, we draw random vectors from the empirical joint distribution of assets, wages,[48] and health status for individuals aged 51-53. Given an initial distribution of health costs, we construct ζ_t, the persistent health cost component, by first finding the normalized log deviation ψ_t, as described in equations (9) and (10), and then applying standard projection formulae to impute ζ_t from ψ_t. The way in which we construct the initial distributions for pensions and AIME is described in Appendices B and C.

---

[47]Bound et al. (2003) consider a more detailed measure of health status.
[48]We impute wages of non-workers using education and (if it exists) the most recent wage.

Appendix G: The Health Cost Model

Recall from equation (9) that health status, health insurance type, labor force participation and age affect health costs through the mean shifter $hc(.)$ and the variance shifter $\sigma(.)$. Health status enters $hc(.)$ and $\sigma(.)$ through 0-1 indicators for bad health, and age enters through linear trends. On the other hand, the effects of Medicare eligibility, health insurance and labor force participation are almost completely unrestricted, in that we allow for an almost complete set of interactions between these variables. This implies, for example, that mean health costs are given by

$$hc(M_t, HI_t, t, P_t) = \gamma_0 \times 1\{M_t = bad\} + \gamma_1 t + \sum_{h \in HI} \sum_{P \in \{0,1\}} \sum_{a \in \{t<65, t \geq 65\}} \gamma_{hPa}.$$

The one restriction is that $\gamma_{none,0,a} = \gamma_{none,1,a}, \forall a$. This implies that there are 10 γ_{hPa} parameters, for a total of 12 parameters apiece in the $hc(.)$ and the $\sigma(.)$ functions.

To estimate this model, we group the data into 10-year-age (55-64, 65-74, 75-84) × health status × health insurance × participation cells. For each of these 60 cells, we calculate both the mean and the 95th percentile of medical expenses. We estimate the model by finding the parameter values that best fit this 120-moment collection. One complication is that the medical expense model we estimate is an annual model, whereas our data are for medical expenses over two-year intervals. To overcome this problem, we first simulate a panel of medical expense data at the one-year frequency, using the dynamic parameters from French and Jones (2004) shown in Table 2 and the empirical age distribution. We then aggregate the simulated data to the two-year frequency; the means and 95th percentiles of this aggregated data are comparable to the means and 95th percentiles in the HRS. Our approach is similar to the one used by French and Jones (2004), who provide a detailed description.

Appendix H: Measurement of Health Insurance Type

Recall that much of the identification in this paper comes from differences in medical expenses and job exit rates between those with *tied* health insurance coverage and those with *ret* (retiree) coverage. Unfortunately, identifying these health insurance types is not

straightforward. The HRS has rather detailed questions about health insurance, but the questions asked vary from wave to wave. Moreover, in no wave are the questions asked consistent with our definitions of *tied* or *ret* coverage.

In all of the HRS waves (but not AHEAD waves 1 and 2), the respondent is asked whether he has insurance provided by a current or past employer or union, or a spouse's current or past employer or union. If he responds yes to this question, we code him as having either *ret* or *tied* coverage. We assume that this question is answered accurately, so that there is no measurement error when individual reports that his insurance category is *none*. All of the measurement error problems arise when we allocate individuals with employer-provided coverage between the *ret* and *tied* categories.

If an individual has employer-provided coverage in waves 1 and 2 he is asked "Is this health insurance available to people who retire?" In waves 3, 4, and 5 the analogous question is "If you left your current employer now, could you continue this health insurance coverage up to the age of 65?". For individuals younger than 65, the question asked in waves 3 through 5 is a more accurate measure of whether the individual has *ret* coverage. In particular, a "yes" response in waves 1 and 2 might mean only that the individual could acquire $COBRA$ coverage if he left his job, as opposed to full, *ret* coverage. Thus the fraction of individuals younger than 65 who report that they have employer-provided health insurance but who answer "no" to the follow-up question roughly doubles between waves 2 and 3. On the other hand, for those older than 65, the question used in waves 3, 4, and 5 is meaningless.

Our preferred approach to the misreporting problem in waves 1 and 2 is to assume that a "yes" response in these waves indicates *ret* coverage. It is possible, however, to estimate the probability of mismeasurement in these waves. Consider first the problem of distinguishing the *ret* and *tied* types for those younger than 65. As a matter of notation, let HI denote an individual's actual health insurance coverage, and let HI^* denote the measure of coverage generated by the HRS questions. To simplify the notation, assume that the individual is known to have employer-provided coverage—$HI = tied$ or $HI = ret$—so that we can drop the conditioning statement in the analysis below. Recall that many individuals who report

67

retiree coverage in waves 1 and 2 likely have tied coverage. We are therefore interested in the misreporting probability $\Pr(HI = tied | HI^* = ret, wv < 3, t < 65)$, where wv denotes HRS wave and t denotes age. To find this quantity, note first that by the law of total probability:

$$\Pr(HI = tied | wv < 3, t < 65) =$$

$$\Pr(HI = tied | HI^* = tied, wv < 3, t < 65) \times \Pr(HI^* = tied | wv < 3, t < 65) +$$

$$\Pr(HI = tied | HI^* = ret, wv < 3, t < 65) \times \Pr(HI^* = ret | wv < 3, t < 65). \qquad (38)$$

Now assume that all reports of *tied* coverage in waves 1 and 2 are true:

$$\Pr(HI = tied | HI^* = tied, wv < 3, t < 65) = 1.$$

Assume further that for individuals younger than 65 there is no measurement error in waves 3-5, and that the share of individuals with tied coverage is constant across waves:

$$\Pr(HI = tied | wv < 3, t < 65) = \Pr(HI = tied | wv \geq 3, t < 65)$$
$$= \Pr(HI^* = tied | wv \geq 3, t < 65).$$

Inserting these assumptions into equation (38) and rearranging yields the mismeasurement probability:

$$\Pr(HI = tied | HI^* = ret, wv < 3, t < 65)$$
$$= \frac{\Pr(HI^* = tied | wv \geq 3, t < 65) - \Pr(HI^* = tied | wv < 3, t < 65)}{\Pr(HI^* = ret | wv < 3, t < 65)}$$
$$= \frac{\Pr(HI^* = ret | wv < 3, t < 65) - \Pr(HI^* = ret | wv \geq 3, t < 65)}{\Pr(HI^* = ret | wv < 3, t < 65)}. \qquad (39)$$

To estimate the mismeasurement in waves 1 and 2 for those aged 65 and older, we make the same assumptions as for those who are younger than 65. We assume that all reports of *tied* health insurance are true and the probability of having *tied* health insurance given a report of *ret* insurance is the same as for individuals in waves 1 and 2 who are younger than

65. We can then use equation (39) to estimate this probability.

The second misreporting problem is that the "follow-up" question in waves 3 through 5 is completely uninformative for those older than 65. Our strategy for handling this problem is to treat the first observed health insurance status for these individuals as their health insurance status throughout their lives. Since we assume that reports of *tied* coverage are accurate, older individuals reporting *tied* coverage in waves 1 and 2 are assumed to receive *tied* coverage in waves 3 through 5. (Recall, however, that if an individual with *tied* coverage drops out of the labor market, his health insurance is *none* for the rest of his life.) For older individuals reporting *ret* coverage in waves 1 and 2, we assume that the misreporting probability—when we choose to account for it—is the same throughout all waves. (Recall that our preferred assumption is to assume that a "yes" response to the follow-up question in waves 1 and 2 indicates *ret* coverage.)

A related problem is that individuals' health insurance reports often change across waves, in large part because of the misreporting problems just described. Our preferred approach for handling this problem is classify individuals on the basis of their first observed health insurance report. We also consider the approach of classifying individuals on the basis of their report from the previous wave, analogous to the practice of using lagged observations as instruments for mismeasured variables in an instrumental variables regression.

Figure 7 shows how our treatment of these measurement problems affects measured job exit rates. The top two graphs in Figure 7 do not adjust for the measurement error problems described immediately above. The bottom two graphs account for the measurement error problems, using the approached described by equation 39. The two graphs in the left column use the first observed health insurance report whereas the graphs in the right column use the previous period's health insurance report. Figure 7 shows that the profiles are not very sensitive to these changes. Those with *ret* coverage tend to exit the labor market at age 62, whereas those with *tied* and *no* coverage tend to exit the labor market at age 65.
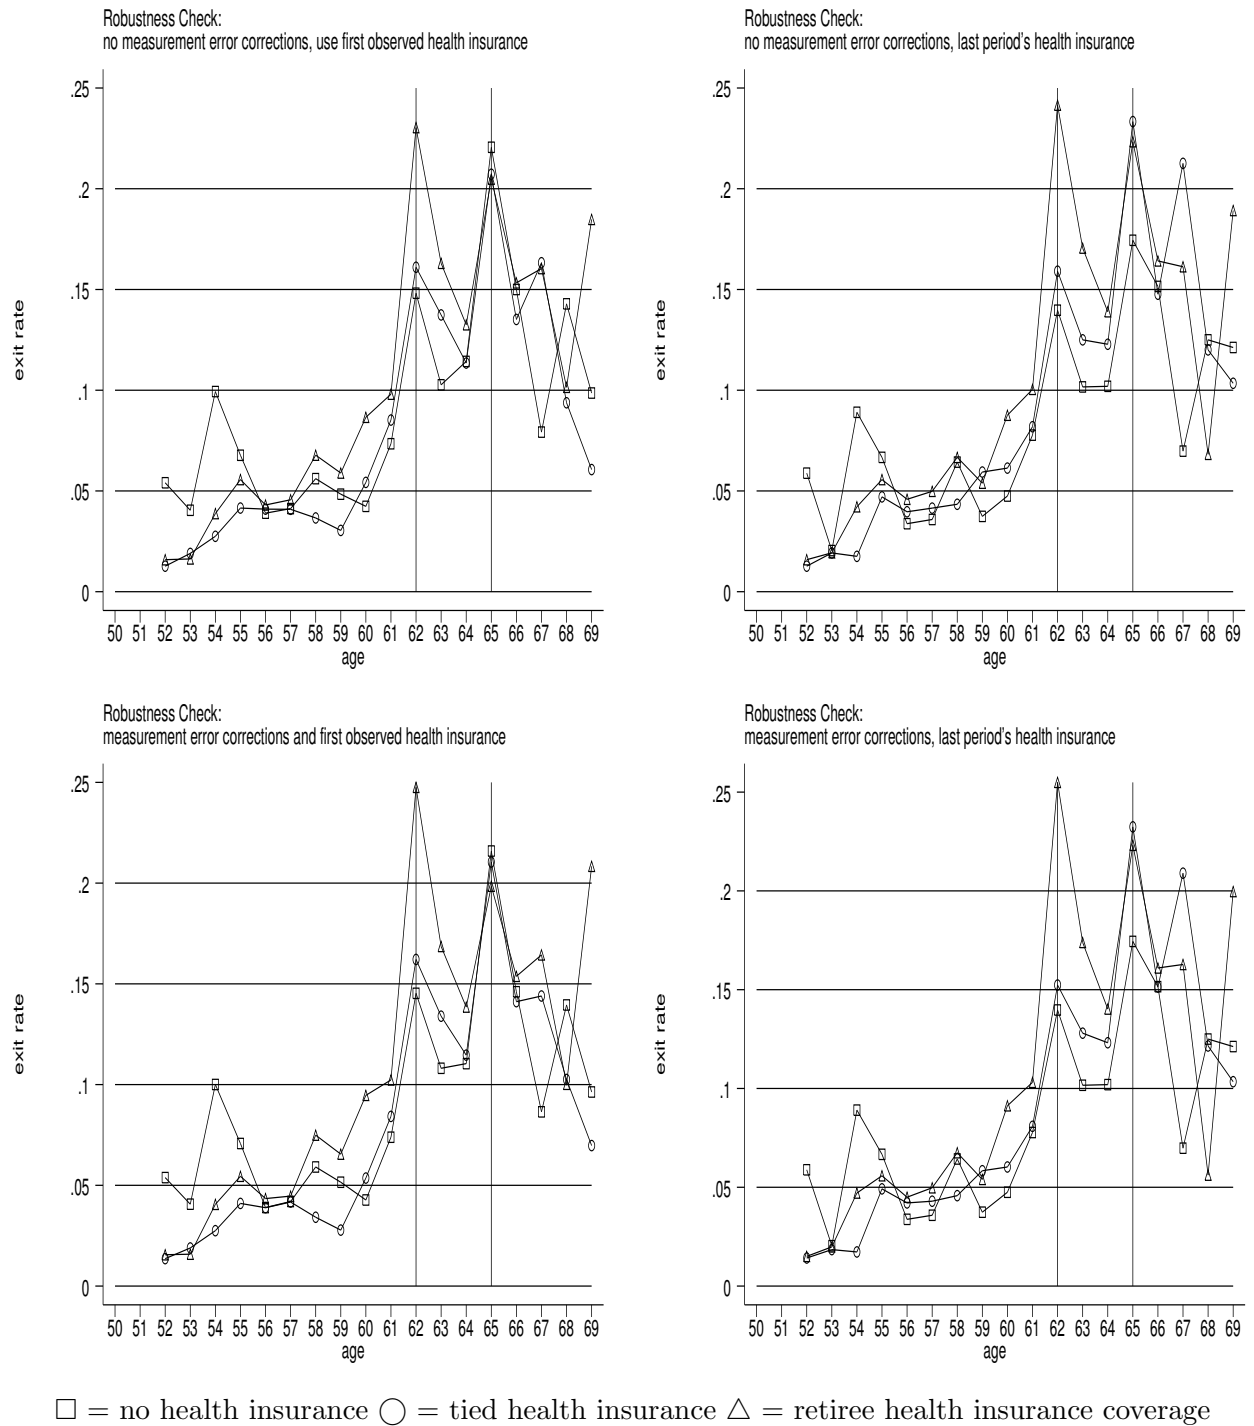
Robustness Check:
no measurement error corrections, use first observed health insurance

Robustness Check:
no measurement error corrections, last period's health insurance

Robustness Check:
measurement error corrections and first observed health insurance

Robustness Check:
measurement error corrections, last period's health insurance

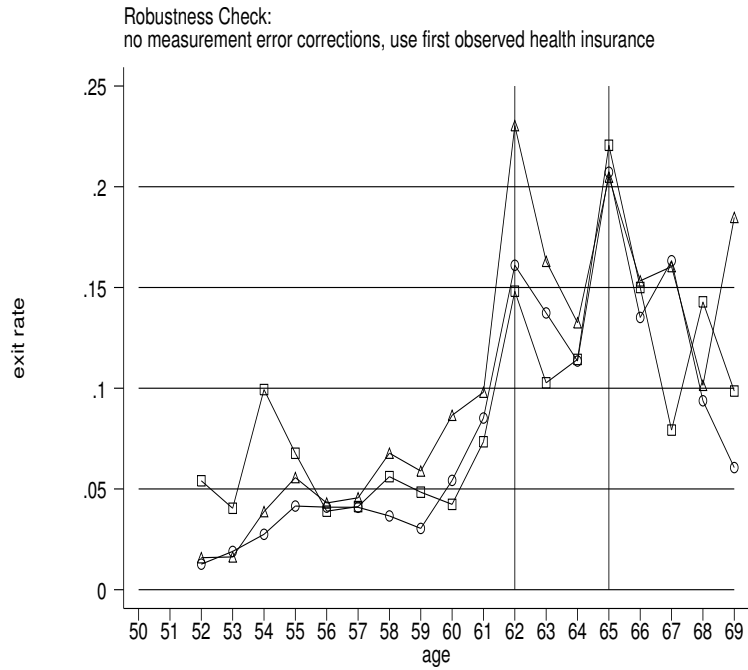□ = no health insurance ◯ = tied health insurance △ = retiree health insurance coverage

Figure 7: Job Exit Rates Using Different Measures of Health Insurance Type

70

Another, more conceptual, problem is that the HRS has information on health insurance outcomes, not choices. This is an important problem for individuals out of the labor force with no health insurance; it is unclear whether these individuals could have purchased COBRA coverage but elected not to do so.[49] To circumvent this problem we use health insurance in the previous wave and the transitions implied by equation (10) to predict health insurance options. For example, if an individual has health insurance that is tied to his job and was working in the previous wave, that individual's choice set is *tied* health insurance and working or *COBRA* insurance and not working.

A final measurement issue is the treatment of the self-employed. Figure 8 shows the importance of dropping the self-employed on job exit rates. The top panel treats the self-employed as working, whereas the bottom panel excludes the self-employed. The main difference caused by dropping the self-employed is that those with *no* health insurance have much higher job exit rates at age 65. Nevertheless, those with *ret* coverage are still most likely to exit at age 62 and those with *tied* and *no* health insurance are most likely to exit at age 65.

Our preferred specification, which we use in the analysis, is to include the self-employed, to use the first observed health insurance report, and to not use the measurement error corrections.

---

[49]For example, the model predicts that all HRS respondents younger than 65 who report having tied health insurance two years before the survey date, work one year before the survey date, and are not currently working should report having COBRA coverage on the survey date. However, 19% of them report having no health insurance.

Robustness Check:
no measurement error corrections, use first observed health insurance

Robustness Check: exclude the self–employed
no measurement error corrections, first observed health insurance
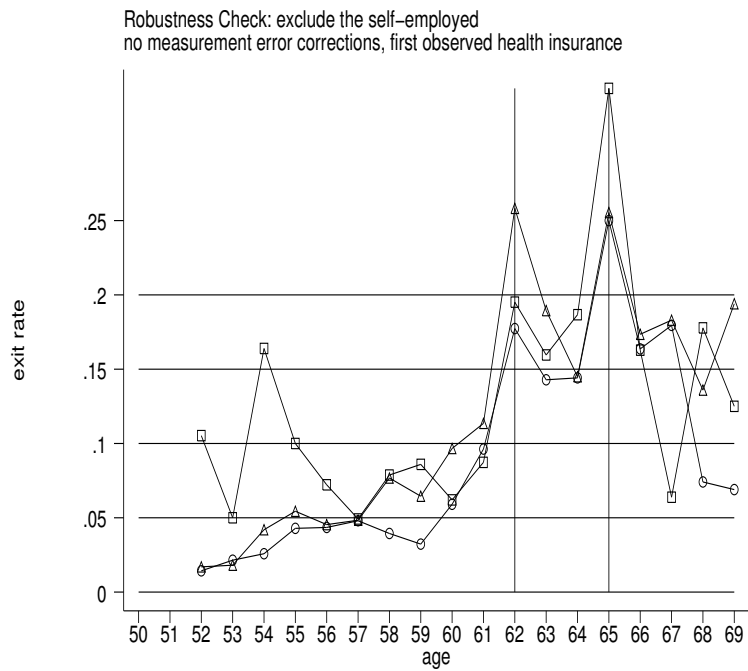
□ = no health insurance ○ = tied health insurance △ = retiree health insurance coverage

Figure 8: THE EFFECT OF DROPPING THE SELF-EMPLOYED ON JOB EXIT RATES

Appendix I: Estimating Age-Specific Rates of Return

During the sample period of 1992-2000, asset prices grew rapidly. This creates a problem because most of our observations come from the HRS core sample, who were ages 51-61 in 1992, 53-63 in 1994, and so on. Therefore, a 63-year old is more likely to be observed in 1994 than in 1992, whereas a 51-year old is more likely to be observed in 1992 than 1994. In other words, older individuals were typically interviewed in later years, when asset prices were considerably higher. This means that in our data, older individuals have more assets not just because of high savings rates, but because of high realized rates of return. These rates of return were likely unexpected. In order to control for this, we estimate the ex-post rate of return that individuals faced as they aged.

Let r_t be the year-specific interest rate faced by a household and let age_{it} be the age of individual i in year t. The average interest rate faced by households headed by an individual of age k is $r_k \equiv E[r_t | age_{it} = k]$. We estimate this object as

$$\hat{r}_k = \frac{1}{I_k} \sum_{i=1}^{I_k} \sum_{t=1}^{T} r_t \times 1\{age_{it} = k\}, \tag{40}$$

where $1\{age_{it} = k\}$ is a 0-1 indicator equal to one when the age of individual i in year t is equal to k, and I_k is the number of households of age k. The goal of this appendix is to estimate $\hat{r}_k$, which means that we must estimate r_t.

We use two approaches to estimating the year-specific interest rate r_t. In both cases, the idea is to first estimate the historical growth rates in asset prices (we use data over the 1952-2001 period), then take the difference in asset price growth in the 1990s and the historical growth rates. Let $r_t = E[r_t] + \varepsilon_t$, where ε_t represents the the unanticipated component of asset returns, which is white noise. We estimate the anticipated component of asset returns $E[r_t]$ using the sample mean $\frac{1}{T} \sum_{t=1}^{T} r_t$. The goal of this appendix is to estimate ε_t.

Our first approach to estimating asset returns is to combine estimates of the rate of return on stocks and housing with the shares of household wealth invested in stocks and housing. Let A_{lit} be the amount of household assets in asset l, $l \in \{1, ..., L\}$, and let r_{lt} be the return

on that asset. The total return on assets at time t is

$$r_t A_{it} = \sum_{l=1}^{L} A_{lit} r_{lt},$$

which implies that the average rate of return at time t is

$$r_t = \sum_{l=1}^{L} \frac{A_{lit}}{A_{it}} r_{lt}.$$

We can then obtain ε_t by estimating

$$\varepsilon_t = r_t - E[r_t] = \sum_{l=1}^{L} \left[\frac{A_{lit}}{A_{it}} \right] (r_{lt} - E[r_{lt}]).$$

We estimate $\frac{A_{lit}}{A_{it}}$ using the sample means of shares in different assets.

Most of the unanticipated component of the rate of return over the sample period came from high rates of return in stocks and housing. Assuming that $r_{lt} - E[r_{lt}] = 0$ for all assets other than stocks and housing, we estimate the share of all assets in stocks (19% in our HRS-AHEAD sample) and housing (32% in our sample) and multiply that by the difference between stock returns in year t, r_{lt}, less their sample average, $\frac{1}{T} \sum_{t=1}^{T} r_{lt}$. We estimate the share of stocks in different assets, such as IRAs, using data from the *Flow of Funds*. The *Flow of Funds* shows, for example, that in 1995 50% of wealth in Defined Contribution pension plans was in stocks. We estimate rates of return on stocks using data from the Center for Research in Security Prices (CRSP) for the period 1952-2001. Therefore, we multiply .5 by the share of wealth in IRAs to give the share of wealth in stocks held in IRA accounts. We estimate rates of return on housing using data from the Office of Federal Housing Enterprise Oversight for the period 1975-2001. (Earlier data are not available.) These estimated growth rates use data on repeat sales of single family housing that were originated by or subsequently purchased by either Freddie Mac or Fannie Mae.

Our second approach to estimating rates of return is to use aggregate data on savings and

asset growth. Aggregate wealth grows according to

$$A_{t+1} = (1 + (r_t(1 - \tau))A_t + S_t,$$

where S_t is savings between time t and time $t+1$ and τ is the average tax rate. Rearranging this equation yields

$$r_t = \frac{A_{t+1} - A_t - S_t}{(1 - \tau)A_t}.$$

We take values of aggregate assets and savings (where savings are defined as personal savings plus undistributed corporate profits) from the *Flow of Funds* for the period 1952-1991. We assume that $\tau = 0.2$.

The second, aggregate data approach is potentially better than the first one because it accounts for the fact that other forms of wealth, such as business wealth, also grew rapidly over the sample period. Moreover, estimating the share of stocks held by households is difficult because it is difficult to infer what share of a household's IRA wealth is in stocks. Data from the *Flow of Funds* indicate that we are potentially understating the share of total wealth in stocks by over 25% (6.3 percentage points). This would lead us to underestimate the wealth gains from holding stocks over our sample period and thus underestimate ε_t.

The drawback to the second approach is that it is very difficult to measure savings. For example, the *Flow of Funds* measure of savings is income minus consumption, where income includes rent, dividends and interest. Ideally, the savings measure would be free of rent, dividends and interest. Because firms reduced the share of earnings going to dividends between 1992 and 2000 (leading to higher growth in the value of firms), the data tend to overstate the decline in the savings rate over the sample period. In other words, part of the run-up in assets not explained by savings rates is merely firms buying new equipment instead of paying dividends. This will lead us to overstate the growth in assets not explained by savings and thus ε_t over the sample period. Therefore, the two procedures likely provide bounds on ε_t, which leads us to use the average of the two measures in our analysis. Figure 9 shows estimates of ε_t, by age, over the sample period. The approach using aggregate asset
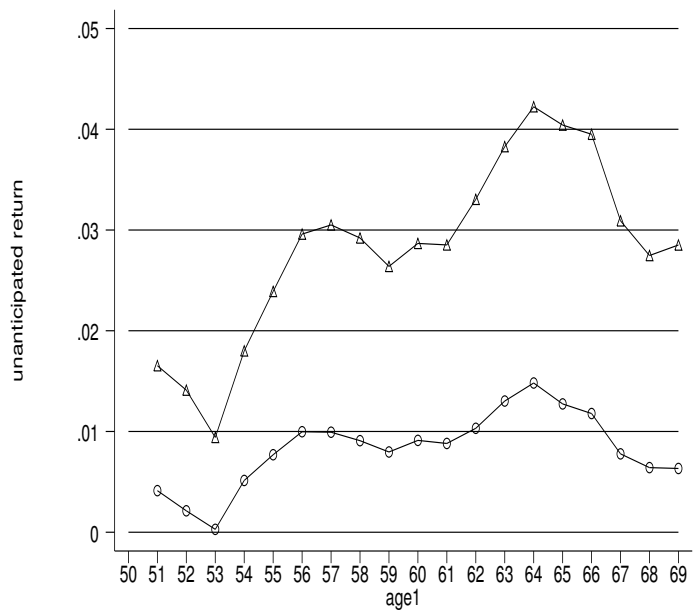
△ = Estimates using flow of funds data
◯ = Estimates using rates of return data and household wealth shares

Figure 9: UNANTICIPATED COMPONENT OF RATE OF RETURN, BY AGE

growth results in considerably higher rates of return than does the approach using stock and
housing returns.

RECENT WORKING PAPERS FROM THE

CENTER FOR RETIREMENT RESEARCH AT BOSTON COLLEGE

Living Arrangements and Supplemental Security Income Receipt Among the Aged
Melissa M. Favreault and Douglas A. Wolf, February 2004

Interactions Between Social Security Reform and the Supplemental Security Income Program for the Aged
Paul S. Davies and Melissa M. Favreault, February 2004

Supply-Side Consequences of Social Security Reform: Impacts on Saving and Employment
Barry Bosworth and Gary Burtless, January 2004

It's All Relative: Understanding the Retirement Prospects of Baby-Boomers
Barbara A. Butrica, Howard M. Iams, and Karen E. Smith, November 2003

The Notional Defined Contribution Model: An Assessment of the Strengths and Limitations of a New Approach to the Provision of Old Age Security
John B. Williamson and Matthew Williams, October 2003

Simulating the Distributional Consequences of Personal Accounts: Sensitivity to Annuitization Options
Cori E. Uccello, Melissa M. Favreault, Karen E. Smith, and Lawrence H. Thompson, October 2003

Aggregate Implications of Defined Benefit and Defined Contribution Systems
Francisco Gomes and Alexander Michaelides, September 2003

Can Unexpected Retirement Explain the Retirement-Consumption Puzzle? Evidence for Subjective Retirement Explanations
Melvin Stephens Jr. and Steven J. Haider, August 2003

Employment, Social Security and Future Retirement Outcomes for Single Mothers
Richard W. Johnson, Melissa M. Favreault, and Joshua H. Goldwyn, July 2003

The Outlook for Pension Contributions and Profits in the U.S.
Alicia H. Munnell and Mauricio Soto, June 2003

Social Security Reform and the Exchange of Bequests for Elder Care
Meta Brown, June 2003

All working papers are available on the Center for Retirement Research website (http://www.bc.edu/crr) and can be requested by e-mail (crr@bc.edu) or phone (617-552-1762).