

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

Eric French

John Bailey Jones*

Federal Reserve Bank of Chicago

University at Albany, SUNY

June 5, 2008

Abstract

This paper provides an empirical analysis of the effects of employer-provided health insurance, Medicare, and Social Security on retirement behavior. Using data from the Health and Retirement Study, we estimate the first dynamic programming model of retirement that accounts for both saving and uncertain medical expenses. Our results suggest that uncertainty and saving are both important for understanding the labor supply responses to Medicare. Furthermore, we find evidence that individuals with stronger preferences for leisure self-select into jobs that provide post-retirement health insurance coverage. Properly accounting for this self-selection reduces the estimated effect of Medicare on retirement behavior. Nevertheless, we find that health insurance is an important determinant of retirement—the labor supply responses to the Medicare eligibility age are as large as the responses to the Social Security normal retirement age.

*Comments welcome at efrench@frbchi.org and jbjones@albany.edu. We thank Joe Altonji, Peter Arcidiacono, Gadi Barlevy, David Blau, John Bound, Chris Carroll, Mariacristina De Nardi, Tim Erikson, Hanming Fang, Donna Gilleskie, Lars Hansen, John Kennan, Spencer Krane, Hamp Lankford, Guy Laroque, John Rust, Dan Sullivan, Chris Taber, the editor (Costas Meghir) and referees, students of Econ 751 at Wisconsin, and participants at numerous seminars for helpful comments. We received advice on the HRS pension data from Gary Englehardt and Tom Steinmeier, and excellent research assistance from Kate Anderson, Olesya Baker, Diwakar Choubey, Phil Doctor, Ken Housinger, Kirti Kamboj, Tina Lam, and Santadarshan Sadhu. The research reported herein was supported by the Center for Retirement Research at Boston College (CRR) and the Michigan Retirement Research Center (MRRC) pursuant to grants from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium. The opinions and conclusions are solely those of the authors, and should not be construed as representing the opinions or policy of the SSA or any agency of the Federal Government, the CRR, the MRRC, or the Federal Reserve System. Recent versions of the paper can be obtained at <http://www.albany.edu/~jbjones/papers.htm>.

1 Introduction

One of the most important social programs for the rapidly growing elderly population is Medicare. In 2005, Medicare had 42.5 million beneficiaries and \$330 billion of expenditures, making it only slightly smaller than Social Security.¹

Prior to receiving Medicare at age 65, many individuals receive health insurance only if they continue to work. At age 65, however, Medicare provides health insurance to almost everyone. Thus an important work incentive disappears at age 65. An important question, therefore, is whether Medicare significantly affects the labor supply of the elderly. This question is particularly important when considering changes to the Medicare or Social Security programs; the fiscal cost of changing the programs depends critically on labor supply responses. Although there is a great deal of research on the labor supply responses to Social Security, there is much less research on the responses to Medicare.

This paper provides an empirical analysis of the effect of employer-provided health insurance and Medicare in determining retirement behavior. Using data from the Health and Retirement Study, we estimate the first dynamic programming model of retirement that accounts for both saving and uncertain medical expenses. Our results suggest that uncertainty and saving are both important for understanding the labor supply responses to Medicare. Furthermore, we find evidence that individuals with stronger preferences for leisure self-select into jobs that provide post-retirement health insurance coverage. Properly accounting for this self-selection reduces the estimated effect of Medicare on retirement behavior. Nonetheless, we find that health insurance is an important determinant of retirement—the Medicare eligibility age is as important for understanding retirement as the Social Security normal retirement age. For example, shifting forward the Medicare eligibility age from 65 to 67 would delay the average age of retirement by 0.07 years, whereas shifting forward the Social Security retirement age from 65 to 67 would delay retirement by 0.09 years.

Our work builds upon, and in part reconciles, several earlier studies. Assuming that in-

¹Figures taken from 2006 Medicare Annual Report (The Boards of Trustees of the Hospital Insurance and Supplementary Medical Insurance Trust Funds, 2006).

dividuals value health insurance at the cost paid by employers, 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 modest, and it declines by a relatively small amount after age 65.² If individuals are risk-averse, however, and if health insurance allows them to smooth consumption in the face of volatile medical expenses, they could value employer-provided health insurance well beyond the cost paid by employers. Medicare's age-65 work disincentive thus comes not only from the reduction in average medical costs paid by those without employer-provided health insurance, but also from the reduction in the volatility of those costs.

Addressing this point, Rust and Phelan (1997) and Blau and Gilleskie (2006a, 2006b) estimate dynamic programming models that account explicitly for risk aversion and uncertainty about out-of-pocket medical expenses. Their estimated labor supply responses to health insurance are larger than those found in studies that omit medical expense risk. Rust and Phelan and Blau and Gilleskie, however, assume that an individual's consumption equals his income net of out-of-pocket medical expenses. In other words, they ignore an individual's ability to smooth consumption through saving. If individuals can self-insure against medical expense shocks by saving, prohibiting saving will overstate the consumption volatility caused by medical cost volatility. It is therefore likely that Rust and Phelan and Blau and Gilleskie overstate the value of health insurance, and thus the effect of health insurance on retirement.

The first major contribution of this paper is that we construct a life-cycle model of labor supply that not only accounts for medical expense uncertainty and health insurance, but also has a saving decision. Moreover, we include the coverage provided by means-tested social insurance to account for the fact that Medicaid provides a close substitute for other forms of health insurance. To our knowledge, ours is the first study of its kind. While van der Klaauw and Wolpin (2006) also estimate a retirement model that accounts for both savings and uncertainty, they do not focus on the role of health insurance, and thus use a much

²Gustman and Steinmeier (1994) find that the average employer contribution to employee health insurance is about \$2,500 per year before age 65. (Data are from the 1977 NMES, adjusted to 1998 dollars with the medical component of the CPI.)

simpler model of medical expenses.

The second major contribution of this paper is that it measures self-selection into jobs with different health insurance plans, along both observable and unobservable dimensions.

Identifying the effect of Medicare on retirement is difficult because virtually everyone is eligible for Medicare at age 65. Furthermore, the Social Security system and pensions also provide retirement incentives at age 65. Thus it is difficult to distinguish whether the high job exit rates at age 65 are due to Medicare, Social Security, or pensions. We circumvent this problem by exploiting variation in employer-provided health insurance. Some individuals receive employer-provided health insurance only while they work, so that their coverage is *tied* to their job. Other individuals have *retiree* coverage, and receive employer-provided health insurance even after they retire. We find that individuals with *retiree* coverage tend to retire about a half year earlier than individuals with *tied* coverage. This suggests that employer-provided health insurance is an important determinant of retirement. Because employer-provided health insurance and Medicare are close substitutes, Medicare could be an important determinant of retirement as well.

One concern with using employer-provided health insurance to identify Medicare's effect on retirement is that individuals potentially choose to work for a firm because of its post-retirement benefits. Individuals with strong preferences for early retirement might choose firms that provide generous *retiree* health insurance benefits. The fact that early retirement is common for individuals with *retiree* coverage may not reflect the effect of health insurance on retirement. Instead, individuals with preferences for early retirement may be self-selecting into jobs that provide *retiree* coverage.

In order to better understand whether self-selection is important, we model preference heterogeneity, using the approach found in Keane and Wolpin (1997). We allow the value of leisure and the time discount factor to vary across individuals, and find evidence of preference heterogeneity along both dimensions. Furthermore, we find that individuals with strong preferences for leisure self-select into firms that provide *retiree* health insurance. We identify the extent of self-selection in part by using a series of self-reports of preferences for work, such

as the response to the question, “Even if I didn’t need the money, I would keep on working.” We find that these responses are highly correlated with future retirement decisions, even after controlling for all of the other incentives in the model. Furthermore, these responses are correlated with the type of health insurance the individual has. Individuals with *retiree* coverage are more likely to self-report that they would like to stop working than individuals whose health insurance is *tied* to their job. This is an important finding because many studies that measure the effect of health insurance on retirement have assumed that preferences for leisure are uncorrelated with the health insurance plan. Our findings on self-selection are corroborated by the fact that those with *tied* coverage still have high labor force participation rates after age 65, when Medicare severs the link between labor supply and health insurance. If health insurance were the sole driving factor, differences across health insurance types should vanish after age 65.

Estimating the model by the Method of Simulated Moments, we find that the model fits the data well with reasonable parameter values. Furthermore, the model fits the data well out-of-sample. Because the Social Security benefit rules vary with year of birth, the Health and Retirement Survey (HRS) contains households that face different Social Security incentives. This allows us to perform an out-of-sample validation exercise: we estimate the model on households with earlier birth years, and then use it to predict the retirement behavior of households with later birth years. We find that the model does a good job of predicting the differences between the two groups observed to date.

Next, we simulate the labor supply response to 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 increase the labor force participation of workers aged 60-67 by 0.07 years. Failure to account for self-selection into health insurance plans results in a larger estimated effect. Nevertheless, even after allowing for both savings and self-selection into health insurance plan, the effect of the Medicare eligibility on labor supply is as large as the effect of the Social Security normal retirement age on labor supply. An important reason why we find that Medicare is important is that we find that medical expense risk is

important. We find that even when we allow individuals to save, they value the consumption smoothing benefits of health insurance.

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 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 an assessment of the model’s performance, both within and outside of the estimation sample. In Section 7, we conduct several policy experiments. In Section 8 we consider a few important robustness checks. Section 9 concludes.

2 The Model

In order to capture the richness of retirement incentives, our model is very complex and has many parameters. Appendix A provides definitions for all the variables used in the main text.

2.1 Preferences and Demographics

Consider a household head seeking to maximize his expected discounted (where the subjective discount factor is β) lifetime utility at age t , $t = 59, 60, \dots, 95$. Each period that he lives, the individual derives utility from consumption, C_t , and hours of leisure, L_t , so that the within-period utility function is of the form

$$U(C_t, L_t) = \frac{1}{1-\nu} (C_t^\gamma L_t^{1-\gamma})^{1-\nu}. \quad (1)$$

We allow both β and γ to vary across individuals. Differences in β reflect differences in patience, while differences in γ represent differences in preferences for leisure. Furthermore, we allow these parameters to be correlated with differences in employer-provided health insurance plans. Thus allow for the possibility that workers with strong preferences for leisure select jobs that provide generous post-retirement health insurance.

The quantity of leisure is

$$L_t = L - H_t - \phi_P P_t - \phi_{HS} HS_t, \quad (2)$$

where L is the individual's total annual time endowment. Participation in the labor force is denoted by P_t , a 0-1 indicator equal to zero when hours worked, H_t , equal zero. The fixed cost of work, ϕ_P , is treated as a loss of leisure. Including fixed costs helps us capture the empirical regularity that annual hours of work are clustered around 2000 hours and 0 hours (Cogan, 1981). 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. Finally, the quantity of leisure depends on an individual's health status through the 0-1 indicator $HS_t = 1\{\text{health}_t = \text{bad}\}$, which equals one when his health is bad. A positive value of the parameter ϕ_{HS} implies that people in bad health find it more painful to work.

Following De Nardi (2004), workers that die value bequests of assets, A_t , according to the function $b(A_t)$:

$$b(A_t) = \theta_B \frac{(A_t + \kappa)^{(1-\nu)\gamma}}{1 - \nu}. \quad (3)$$

When the parameter $\kappa > 0$, the marginal utility of a zero dollar bequest is finite, and bequests are a luxury good.

Health status, HS_t , and the probability of being alive at age t conditional on being alive at age $t - 1$, s_t , both depend on age and previous health status.

2.2 Budget Constraints

The individual holds three forms of wealth: assets (including housing); pensions; and Social Security. He receives 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

asset accumulation equation is

$$A_{t+1} = A_t + Y(rA_t + W_tH_t + ys_t + pb_t, \tau) + ss_t + tr_t - m_t - C_t. \quad (4)$$

m_t denotes medical expenses, and post-tax income, $Y(rA_t + W_tH_t + ys_t + pb_t, \tau)$, is a function of taxable income and the vector τ , described in Appendix B, that captures the tax structure.

Individuals face the borrowing constraint³

$$A_t + Y_t + ss_t + tr_t - C_t \geq 0. \quad (5)$$

Because it is illegal to borrow against future Social Security benefits and difficult to borrow against many forms of future pension benefits, 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.

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

$$tr_t = \max\{0, C_{min} - (A_t + Y_t + ss_t)\}. \quad (6)$$

Equation (6) implies that government transfers bridge the gap between an individual’s “liquid resources” (the quantity in the inner parentheses) and the consumption floor. Equation (6) 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-

³We assume time t medical expenses are realized after time t labor decisions are made. Thus the borrowing constraint excludes medical expenses. 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. Given the borrowing constraint and timing of medical expenses, an individual with extremely high medical expenses this year could have negative net worth next year. Thus we allow for the fact that many people in our data still have unresolved medical expenses, medical expense debt seems reasonable. Because debt cannot legally be bequeathed in the US, we assume that all debts are erased at time of death when calculating the value of the bequest in equation (3).

provided 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

We define m_t as the sum of all out-of-pocket medical expenses (including those covered by the consumption floor), including insurance premia. We assume that an individual's medical expenses depend upon five different components. First, medical expenses depend on the individual's employer-provided health insurance, HI_t . Second, they depend on whether the person is working, P_t , because workers who leave their job often pay a larger fraction of their insurance premiums. Third, they depend on the individual's self-reported health status, HS_t . Fourth, medical expenses depend on age. At age 65, individuals become eligible for Medicare, which is a close substitute for employer-provided coverage.⁴ Offsetting this, as people age their health declines (in a way not captured by HS_t), raising medical expenses. Finally, medical expenses depend on person-specific effects, which we capture in the variable ψ_t , yielding:

$$\ln m_t = m(HS_t, HI_t, t, P_t) + \sigma(HS_t, HI_t, t, P_t) \times \psi_t. \quad (7)$$

Note that health insurance affects both the expectation of medical expenses, through $m(\cdot)$ and the variance, through $\sigma(\cdot)$.⁵

Even after controlling for health status, French and Jones (2004a) find that medical expenses are very volatile and persistent. Thus we model the idiosyncratic component of

⁴Individuals 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/>.

⁵We follow the existing literature and impose the simplifying assumption that medical expenditures are exogenous. To our knowledge, Blau and Gilleskie (2006b) is the only structural retirement study to have endogenous medical expenditures.

medical expenses, ψ_t , as

$$\psi_t = \zeta_t + \xi_t, \quad \xi_t \sim N(0, \sigma_\xi^2), \quad (8)$$

$$\zeta_t = \rho_m \zeta_{t-1} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma_\epsilon^2), \quad (9)$$

where ξ_t and ϵ_t are serially and mutually independent. ξ_t is the transitory component of medical expense uncertainty, while ζ_t is the persistent component, with autocorrelation ρ_m .

Differences in labor supply behavior across health insurance categories, HI_t , are an important part of identifying our model. We assume that there are three mutually exclusive categories of health insurance coverage. The first is *retiree* coverage, where workers keep their health insurance even after leaving their jobs. The second category is *tied* health insurance, where workers receive employer-provided coverage as long as they continue to work. If a worker with tied health insurance leaves his job, he can keep his health insurance coverage for that year. This is meant to proxy for the fact that most firms must provide “COBRA” health insurance to workers after they leave their job. After one year of *tied* coverage and not working, the individual’s insurance ceases.⁶ The third category consists of individuals whose potential employers provide no health insurance at all, or *none*. Workers move between these insurance categories according to⁷

$$HI_t = \begin{cases} retiree & \text{if } HI_{t-1} = retiree \\ tied & \text{if } HI_{t-1} = tied \text{ and } H_{t-1} > 0 \\ none & \text{if } HI_{t-1} = none \text{ or } (HI_{t-1} = tied \text{ and } H_{t-1} = 0) \end{cases} . \quad (10)$$

This transition rule implies those with *tied* coverage work not only for labor income, but also

⁶Although 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 one year. We do not model the option to take up COBRA, assuming that the take-up rate is 100%. Although the actual take-up rate is around $\frac{2}{3}$ (Gruber and Madrian, 1996), we have simulated the model assuming that the rate was 0% so that individuals transition from *tied* to *none* as soon as they stop working, and found very similar labor supply patterns. Thus assuming 100% take up does not seem to drive our results.

⁷In 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 *retiree* coverage.

for health insurance. Thus, differences in labor supply patterns among those with *tied* and *retiree* coverage are useful for inferring the effect of health insurance on retirement.

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 (HS_t), age (t), hours worked (H_t) and an autoregressive component, ω_t :

$$\ln W_t = W(HS_t, t) + \alpha \ln H_t + \omega_t, \quad (11)$$

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 :

$$\omega_t = \rho_W \omega_{t-1} + \eta_t, \quad \eta_t \sim N(0, \sigma_\eta^2). \quad (12)$$

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:

$$ys_t = ys(W(HS_t, t) + \omega_t, t). \quad (13)$$

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

2.5 Social Security and Pensions

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

Individuals receive no Social Security benefits until they apply. Individuals can first apply for benefits at age 62. Upon applying the individual receives benefits until death.

The individual's Social Security benefits depend on his Average Indexed Monthly Earnings (*AIME*), which is roughly his average income during his 35 highest earnings years in the labor market.

The Social Security System provides three major retirement incentives.⁸ 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 D.

Second, the age at which the individual applies for Social Security affects the level of benefits. 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 5.5% up until age 70. This is less than actuarially fair, and encourages people to apply for benefits by age 65.

Third, the Social Security Earnings Test taxes labor income of beneficiaries at a high rate. 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 before 2000, 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 65-69.⁹ When combined with the aforementioned incentives to draw Social Security benefits by age 65, the Earnings Test discourages work after age 65. In 2000, the Social Security Earnings Test was abolished

⁸A description of the Social Security rules can be found in recent editions of the *Green Book* (Committee on Ways and Means). Some of the rules, such as the benefit adjustment formula, depends on an individual's year of birth. Because we fit our model to a group of individuals that on average were born in 1933, we use the benefit formula for that birth year.

⁹The credit rates are based on the benefit adjustment formula. 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 66, benefits in the future are increased by 5.5%.

for those 65 and older. Because those born in 1933 (the average birth year in our sample) turned 67 in 2000, we assume that the earnings test was repealed at age 67. These incentives are incorporated in the calculation of ss_t , which is defined to be net of the earnings test.

Pension benefits, pb_t , are a function of the worker's age and pension wealth. Pension wealth (the present value of pension benefits) in turn depends on pension accruals. We assume that pension accruals are a function of a worker's age, labor income, and health insurance type, using a formula estimated from confidential HRS pension data. The data show that pension accrual rates differ greatly across health insurance categories; accounting for these differences is essential in isolating the effects of employer-provided health insurance. 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 of our pension model are described in Section 4.3 and Appendix C.

2.6 Recursive Formulation

In addition to choosing hours and consumption, eligible individuals decide whether to apply for Social Security benefits; let the indicator variable $B_t \in \{0, 1\}$ equal one if an individual has applied. In recursive form, the individual's problem can be written as

$$V_t(X_t) = \max_{C_t, H_t, B_t} \left\{ \frac{1}{1-\nu} \left(C_t^\gamma (L - H_t - \phi_P P_t - \phi_{HS} H S_t)^{1-\gamma} \right)^{1-\nu} + \beta(1-s_{t+1})b(A_{t+1}) + \beta s_{t+1} \int V_{t+1}(X_{t+1}) dF(X_{t+1}|X_t, t, C_t, H_t, B_t) \right\}, \quad (14)$$

subject to equations (5) and (6). The vector $X_t = (A_t, B_{t-1}, H S_t, A I M E_t, H I_t, \omega_t, \zeta_{t-1})$ contains the individual's state variables, while the function $F(\cdot|\cdot)$ gives the conditional distribution of these state variables, using equations (4) and (7) - (13).¹⁰ The solution to the individual's problem consists of the consumption rules, work rules, and benefit application rules that solve equation (14). Given that the model lacks a closed form solution, these de-

¹⁰Spousal income and pension benefits (see Appendix C) depend only on the other state variables and are thus not state variables themselves.

cision rules are found numerically using value function iteration. Appendix E describes our numerical methodology.

3 Estimation

To estimate the model, we adopt a two-step strategy, similar to the one used by Gourinchas and Parker (2002) and French (2005). In the first step we estimate or calibrate parameters that can be cleanly identified without explicitly using our model. For example, we estimate mortality rates and health transitions straight from demographic data. In the second step, we estimate the preference parameters of the model, as well as the consumption floor, using the method of simulated moments (MSM).

3.1 Moment Conditions

The objective of MSM estimation is to find the preference vector that yields simulated life-cycle decision profiles that “best match” (as measured by a GMM criterion function) the profiles from the data. The moment conditions that comprise our estimator are:

1. Because an individual’s ability to self-insure against medical expense shocks depends critically upon his asset level, we match 1/3rd and 2/3rd asset quantiles by age.¹¹ We match these quantiles in each of T periods (ages), for a total of $2T$ moment conditions.
2. We match job exit rates by age for each health insurance category. With three health insurance categories (*none*, *retiree* and *tied*), this generates $3T$ moment conditions.
3. Because the value a worker places on employer-provided health insurance may depend on his wealth, we match labor force participation conditional on the combination of asset quartile and health insurance status. With 2 quantiles (generating 3 quantile-conditional means) and 3 health insurance types, this generates $9T$ moment conditions.

¹¹Our approach to constructing quantile-related moment conditions follows Manski (1988), Powell (1994), or Buchinsky (1998). Chernozhukov and Hansen (2002) provide additional, earlier, references. Related approaches appear in Epple and Seig (1999) and Cagetti (2003).

4. To help identify preference heterogeneity, we utilize a series of questions in the HRS that ask workers about their preferences for work. We combine the answers to these questions into a time-invariant index, $pref \in \{high, low, out\}$, which is described in greater detail in Section 4.4. Matching participation conditional on each value of this index generates another $3T$ moment conditions.
5. Finally, we match hours of work and participation conditional on our binary health indicator. This generates $4T$ moment conditions.

Combined, the five preceding items result in $21T$ moment conditions. Appendix F provides a detailed description of the moment conditions, the mechanics of our MSM estimator, the asymptotic distribution of our parameter estimates, and our choice of weighting matrix.

3.2 Initial Conditions and Preference Heterogeneity

A key part of our estimation strategy is to compare the behavior of individuals with different forms of employer-provided health insurance. If access to health insurance is an important factor in the retirement decision, we should find that individuals with *tied* coverage should retire later than those with *retiree* coverage. In making such a comparison, however, we must account for the possibility that individuals with different health insurance options differ systematically along other dimensions as well. For example, individuals with retiree coverage tend to have higher wages and more generous pensions.

We control for this “initial conditions” problem in three ways. First, the initial distribution of simulated individuals is drawn directly from the data. Because wealthy households are more likely to have retiree coverage in the data, wealthy households are more likely to have retiree coverage in our initial distribution. Second, we model carefully the way in which pension and Social Security accrual varies across individuals and groups.

Finally, we control for unobservable differences across health insurance groups by introducing permanent preference heterogeneity, using the approach introduced by Heckman and Singer (1984) and adapted by (among others) Keane and Wolpin (1997) and van der Klaauw

and Wolpin (2006). Each individual is assumed to belong to one of a finite number of preference “types”, with the probability of belonging to a particular type a logistic function of the individual’s preference index, initial wealth, wages and health insurance type. Our approach allows for the possibility that people with different preferences systematically self-select into different types of health insurance coverage. We estimate the type probability parameters jointly with the preference parameters and the consumption floor.

4 Data and Calibrations

4.1 HRS Data

We estimate the model using data from 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 medical expenses, 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 7 waves of data covering the period 1992-2004. 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 Appendix G.

As noted above, the Social Security benefit adjustment formula depends on an individual’s year of birth. To ensure that workers in our sample face a similar set of Social Security retirement incentives, we fit our model to the decision profiles of the cohort of individuals aged 57-61 in 1992. However, when estimating the stochastic processes that individuals face, we often use the full sample in order to increase sample size.

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

4.2 Health Insurance and Medical Expenses

We assign individuals to one of three mutually exclusive health insurance groups: *retiree*, *tied*, and *none*, as described in Section 2. Because of small sample problems, the *none* group includes those with private health insurance as well as those with no insurance at all. Both face high medical expenses because they lack employer-provided coverage. Private health insurance is usually not a substitute for employer-provided coverage, as high administrative costs and adverse selection problems can result in prohibitively expensive premiums. Moreover, private coverage often does not cover pre-existing medical conditions, whereas employer-provided coverage typically does. Because the model includes a consumption floor to capture the insurance provided by Medicaid, the *none* group also includes those who receive health care through Medicaid. We assign those who have health insurance provided by their spouse to the *retiree* 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.

Although the HRS's insurance-related data are detailed, they are never completely consistent with our definitions of *tied* or *retiree* coverage. Appendix H shows, however, that the health-insurance-specific job exit rates are not very sensitive to the assumptions we imposed in interpreting the data.

The HRS has data on self-reported medical expenses. Medical expenses are the sum of insurance premia paid by households, 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 medical expense model described in Section 2. Because of small sample problems, we allow the mean, $m(\cdot)$, and standard deviation, $\sigma(\cdot)$, 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 (2004a), $m(\cdot)$ and $\sigma(\cdot)$ 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 I.

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 expenses are \$2,950 for workers 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 65.¹² Thus, the value of having employer provided health insurance coverage largely vanishes at age 65.

As Rust and Phelan (1997) emphasize, it is not just differences in mean medical expenses that determine the value of health insurance, but also differences in variance and skewness. If health insurance reduces medical expense 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 medical expense distributions. Table 1 shows that for healthy individuals who are 64 years old, average

¹²The 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 - Working	Tied - Not Working	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

annual medical expenses have a standard deviation of \$7,150 for workers 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 medical expenses for those without employer-provided health insurance. It reduces medical expense volatility as well.

Relative to other research on the cross sectional distribution of medical expenses, we find higher medical expenses at the far right tail of the distribution. For example, Blau and Gilleskie (2006a) use different data and methods to find average medical expenses that are comparable to our estimates. However, they find that medical expenses are less volatile than our estimates suggest. For example, they find that for households in good health and younger than 65, the maximum expense levels (which seem to be slightly less likely than 0.5% probability events) were \$69,260 for those without coverage, \$6,400 for those with retiree coverage, and \$6,400 for those with tied coverage. Table 1 shows that our estimates of the 99.5th percentile (i.e., the top 0.5 percentile of the distribution) of the distributions

for healthy individuals are \$91,560 for those with no coverage, \$41,020 for those with retiree coverage, and \$40,210 for those with tied coverage.

Berk and Monheit (2001) use data from the MEPS, which arguably has the highest quality medical expense data of all the surveys. Using a measure of medical expenses that should be comparable to our estimates for the uninsured,¹³ Berk and Monheit find that those in the top 1% of the medical expense distribution have average medical expenses of \$57,900 (in 1998 dollars). Again, this is below our estimate of \$91,560 for the uninsured. This discrepancy is not surprising. Berk and Monheit’s estimates are for all individuals in the population, whereas our estimates are for older households (many of which include two individuals). Furthermore, Berk and Monheit’s estimates exclude all nursing home expenses, while the HRS, although initially consisting only of non-institutionalized households, captures the nursing home expenses these households incur in later waves.

The parameters for the idiosyncratic process $\psi_t, (\sigma_\xi^2, \sigma_\epsilon^2, \rho_m)$, are taken from French and Jones (2004a). Table 2 presents the parameters, which have been normalized so that the 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 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_m of 0.925, this is not surprising.

Parameter	Variable	Estimate
σ_ϵ^2	innovation variance of persistent component	0.04811
ρ_m	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

¹³Berk and Monheit use data on total billable expenses. The uninsured should pay all billable expenses, so Berk and Monheit’s estimated distribution should be comparable to our distribution for the uninsured.

4.3 Pension Accrual

Appendix C gives details on how we use the confidential HRS pension data to construct an accrual rate formula. Figure 1 shows the average pension accrual rates generated by this formula, conditional on having average income.




Figure 1: AVERAGE PENSION ACCRUAL RATES, BY AGE AND HEALTH INSURANCE COVERAGE

Workers with *retiree* coverage are the most likely to have defined benefit plans (which often have sharp drops in pension accrual after age 60), workers with *tied* coverage are the most likely to have a defined contribution plan, and workers with no coverage are the most likely no have no pension plan. As a result, those with *retiree* coverage face the sharpest drops in pension accrual after age 60.¹⁴ Furthermore, the confidential pension data show that, conditional on having a defined benefit pension plan, those with *retiree* coverage face the sharpest drops in pension accrual after age 60. In short, not only does *retiree* coverage in and of itself provide an incentive for early retirement, but the pension plans associated with *retiree* coverage provide the strongest incentives for early retirement. Modeling the association between pension accrual and health insurance coverage is thus critical; failing to

¹⁴Because Figure 1 is based on our estimation sample, it does not show accrual rates for earlier ages. Estimates that include the validation sample show, however, that those with retiree coverage have the highest pension accrual rates in their early and middle 50s.

capture this link will lead the econometrician to overstate the importance of *retiree* coverage on retirement.

4.4 Preference Index

In order to better measure preference heterogeneity in the population (and how it is correlated with health insurance), we estimate a person’s “willingness” to work using three questions from the first (1992) wave of the HRS. The first question asks the respondent the extent to which he agrees with the statement, “Even if I didn’t need the money, I would probably keep on working.” The second question asks the respondent, “When you think about the time when you will retire, are you looking forward to it, are you uneasy about it, or what?” The third question asks, “How much do you enjoy your job?”

To combine these three questions into a single index, we regress wave 5-7 (survey year 2000-2004) participation on the response to the three questions along with polynomials and interactions of all the state variables in the model: age, health status, wages, wealth, and AIME, medical expenses, and health insurance type. Multiplying the numerical responses to the three questions by their respective estimated coefficients and summing yields an index. We then discretize the index into three values: *high*, for the top 50% of the index for those working in wave 1; *low*, for the bottom 50% of the index for those working in wave 1; and *out* for those not working in wave 1. Appendix J provides additional details on the construction of the index. Figure 7 below shows that the index has great predictive power: at age 65, participation rates are 56% for those with an index of *high*, 39% for those with an index of *low*, and 12% for those with an index of *out*.

The discretized preference index plays a major role in our strategy for estimating preference heterogeneity. The preference index is included—and, as shown below, has significant effects—in the preference type prediction equations. Furthermore, in our MSM estimation procedure we require the model to match participation at each value of the preference index. As a result, our estimated distribution of unobserved preference heterogeneity depends on the distribution of observed preference heterogeneity, which should improve the precision of

our estimates.

4.5 Wages

Recall from equation (11) that $\ln W_t = \alpha \ln(H_t) + W(HS_t, t) + \omega_t$. 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, $W(HS_t, t)$, from French (2005), 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 59.

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. However, if individuals leave the market because of a sudden wage drop, such as from job loss, wage growth rates for workers will be greater than wage growth rates for non-workers. This will bias estimated wage growth upward. To correct for this problem, our baseline analysis uses the selection-adjusted wage profiles estimated by French (2005).

The parameters for the idiosyncratic process ω_t , (σ_η^2, ρ_W) are also estimated by French (2005). The results indicate that the autocorrelation 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.

4.6 Remaining Calibrations

We set the interest rate r equal to 0.03. 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. We estimate the Markov transition matrices using data from the HRS and Assets and Health Dynamics of the Oldest Old.

5 Data Profiles and Initial Conditions

5.1 Data Profiles

Figure 2 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 \$80,000 in assets, and about one third live in households with over \$270,000 of assets. The asset profiles also show that assets grow with age. This growth is higher than that reported in other studies (for example, Cagetti, 2003, and French, 2005). Earlier drafts of this paper showed that the run-up in asset prices during the sample period can explain some, although not all, of this run-up.




Figure 2: ASSET QUANTILES, DATA

The first panel of Figure 3 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* (18.4%), *retiree* (16.1%), or *none* (16.9%), those with *retiree* coverage have significantly higher exit rates at

62 (21.9%) than those with *tied* (14.2%) or *none* (14.5%).¹⁵ At almost every age other than 65, 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).

The low job exit rates before age 65 and the relatively high job exit rates at age 65 for those with *tied* coverage suggests that some people with *tied* coverage are working until age 65, when they become eligible for Medicare. The profiles therefore provide evidence that there is a causal effect of health insurance on retirement. On the other hand, job exit rates for those with *tied* coverage are lower than those with *retiree* coverage for every age other than 65, and are not much higher at age 65. This suggests that differences in health insurance coverage may not be the only reason for differences in job exit rates between those with *tied* coverage and those with *retiree* coverage.

Although our health insurance classifications probably contain measurement error, Appendix H shows that the estimated job exit rates are not very sensitive to different coding decisions for health insurance.

The bottom panel of Figure 3 presents observed labor force participation rates. In comparing participation rates across health insurance categories, 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.

¹⁵The differences across groups are not statistically different at the 5% level at age 62. However, when we include our validation sample of younger individuals, the differences are statistically different at age 62. Furthermore, F -tests reject the hypothesis that the three groups have identical exit rates at all ages at the 5% level.





Figure 3: JOB EXIT AND PARTICIPATION RATES, DATA

5.2 Initial Conditions

Each artificial individual in our model begins its simulated life with the year-1992 state vector of an individual, aged 57-61 in 1992, observed in the data. Table 3 summarizes this initial distribution, the construction of which is described in Appendix G. Table 3 shows that individuals with *retiree* coverage tend to have the most asset and pension wealth, while individuals in the *none* category have the least—the median individual in the none category has no pension wealth at all. Individuals in the *none* category are also more likely to be in bad health, and not surprisingly, less likely to be working. In contrast, individuals with *tied* coverage have high values of the preference index, suggesting that their delayed retirement reflects differences in preferences as well as in incentives.

	<i>Retiree</i>	<i>Tied</i>	<i>None</i>
Age			
Mean	58.7	58.6	58.7
Standard deviation	1.5	1.5	1.5
<i>AIME</i> (in thousands of 1998 dollars)			
Mean	25.1	25.3	16.5
Median	27.2	26.9	16.4
Standard deviation	9.1	8.6	9.2
Assets (in thousands of 1998 dollars)			
Mean	229.7	203.5	201.6
Median	146.0	112.3	55.6
Standard deviation	246.1	254.3	306.7
Pension Wealth (in thousands of 1998 dollars)			
Mean	129.2	80.0	18.7
Median	65.1	14.5	0.0
Standard deviation	181.2	213.4	100.8
Wage (in 1998 dollars)			
Mean	17.2	17.7	12.6
Median	14.7	14.6	8.5
Standard deviation	12.2	12.3	14.4
Preference index			
Fraction <i>out</i>	0.27	0.04	0.48
Fraction <i>low</i>	0.42	0.44	0.18
Fraction <i>high</i>	0.32	0.52	0.33
Fraction in bad health	0.20	0.13	0.41
Fraction working	0.73	0.96	0.52
Number of observations	1,022	225	454

Table 3: SUMMARY STATISTICS FOR THE INITIAL DISTRIBUTION

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 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 can be thought of as an “exclusion restriction”, which changes the incentives for work and savings but does not change preferences.

Table 4 presents preference parameter estimates. The first 3 rows of Table 4 show the parameters that vary across the preference types. We assume that there are three types of individuals, and that the types differ in the utility weight on consumption, γ , and their time discount factor, β . Individuals with high values of γ have stronger preferences for work. Individuals with high values of β are more patient and thus more willing to defer consumption and leisure.

Table 4 reveals significant differences in γ and β across preference types. To understand these differences, it is useful to consider Table 5, which shows simulated summary statistics for each of the preference types. Table 5 reveals that Type-0 individuals have the lowest value of γ , i.e., they place the highest value on leisure. 93% of Type-0 individuals were *out* of the labor force in wave 1. Type-2 individuals, in contrast, have the highest value of γ . 99.8% of Type-2 individuals have a preference index of *high*, meaning that they were working in wave 1 and self-reported having a low preference for leisure. Type-1 individuals fall in the middle, valuing leisure less than Type-0 individuals, but more than Type-2 individuals. 59% of Type-1 individuals have a preference index value of *low*. Figure 7 shows that *low*-index individuals are also an intermediate case: although they initially work as much as *high*-index individuals, *low*-index workers leave the labor force much more quickly.

Including preference heterogeneity allows us to control for the possibility that workers with different preferences select jobs with different health insurance packages. Table 5 suggests

that some self-selection is occurring, as it reveals that workers with *tied* coverage are more likely to be Type-2 agents, who have the strongest preference for work. This suggests that workers with *tied* coverage might be more willing to retire at later dates simply because they have a lower disutility of work.¹⁶

<i>Parameters that vary across individuals</i>			
γ : consumption weight		β : time discount factor	
Type 0	0.438 (0.080)	Type 0	0.828 (0.072)
Type 1	0.620 (0.011)	Type 1	1.115 (0.016)
Type 2	0.907 (0.028)	Type 2	0.971 (0.077)
<i>Parameters that are common to all individuals</i>			
ν : coefficient of relative risk aversion, utility	7.49 (0.421)	θ_B : bequest weight [†]	0.0320 (0.0009)
L : leisure endowment, in hours	3,863 (51.9)	κ : bequest shifter, in thousands	449 (31.7)
ϕ_P : fixed cost of work, in hours	835 (27.4)	c_{min} : consumption floor	4,118 (159.5)
ϕ_{HS} : hours of leisure lost, bad health	445 (38.8)		
χ^2 statistic = 1,677; Degrees of freedom = 181			
Method of simulated moments estimates. Diagonal weighting matrix used in calculations. See Appendix F for details. Standard errors in parentheses. [†] Parameter expressed as marginal propensity to consume out of final-period wealth. Parameters estimated jointly with type probability prediction equation. See Appendix K for estimated coefficients of the type probability prediction equation.			

Table 4: ESTIMATED STRUCTURAL PARAMETERS

The bottom line of Table 5 shows the fraction of each preference type. Averaging over the three preference types reveals that the average value of β implied by our model is 1.02, which is slightly higher than most estimates. There are two reasons for this. 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

¹⁶Interestingly, Type-2 agents also include wealthy individuals who have no health insurance coverage. Given that many of these individuals are entrepreneurs, it is not surprising that they are often placed in the “motivated” group.

	Type 0	Type 1	Type 2
<i>Key preference parameters</i>			
γ^*	0.438	0.620	0.907
β^*	0.828	1.115	0.971
<i>Means by preference type</i>			
Assets (\$1,000s)	164	236	239
Pension Wealth (\$1,000s)	92	103	56
Wages (\$/hour)	10.5	19.4	11.7
<i>Probability of health insurance type, given preference type</i>			
Health insurance = none	0.405	0.232	0.184
Health insurance = retiree	0.578	0.642	0.461
Health insurance = tied	0.018	0.126	0.355
<i>Probability of preference index value, given preference type</i>			
Preference Index = out	0.930	0.098	0.0002
Preference Index = low	0.004	0.591	0.0018
Preference Index = high	0.066	0.311	0.998
Fraction with preference type	0.250	0.600	0.150
*Values of β and γ are from Table 4.			

Table 5: MEAN VALUES BY PREFERENCE TYPE, SIMULATIONS

marginal tax rate.¹⁷ 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 taxes—implicitly setting s_{t+1} and $1-\tau_t$ 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 by the intertemporal substitution of leisure, as embodied in the labor supply profiles.¹⁸ Models of labor supply and savings, such as MaCurdy (1981) or French (2005), often suggest that agents are very patient.

Another 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, it can be approximated as $-\frac{(\partial^2 U/\partial C^2)C}{\partial U/\partial C} = -(\gamma(1-\nu)-1)$. As we move across preference types, the coefficient increases from 3.8 to 5.0 to

¹⁷Note that this equation does not hold exactly when individuals value bequests.

¹⁸This restriction is often relaxed by adding a time trend to leisure- (or consumption-) related utility parameters. See, e.g., Rust and Phelan, 1997, Blau and Gilleskie, 2006a and 2006b, and Gustman and Steinmeier, 2005, Rust et al., 2003, van der Klaauw and Wolpin, 2006.

6.9. These values are within the range of estimates found in recent studies by Cagetti (2003) and French (2005), but they are larger than the values of 1.1, 1.8, and 1.0 reported by Rust and Phelan (1997), Blau and Gilleskie (2006a), and Blau and Gilleskie (2006b) respectively, in their studies of retirement.

The consumption floor c_{min} and ν are identified in large part by the asset quantiles, which reflect precautionary motives. 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 medical expense shock, which over a lifetime could equal over \$100,000 (see French and Jones (2004a)), can generate strong precautionary incentives. Conversely, as emphasized by Hubbard, Skinner and Zeldes (1995), a high consumption floor discourages saving among the poor, since the consumption floor effectively imposes a 100% tax on the saving of those with high medical expenses and low income and assets.

Our estimated consumption floor of \$4,118 is similar to other estimates of social insurance transfers for the indigent. For example, when we use Hubbard, Skinner and Zeldes’s (1994, Appendix A) procedures and more recent data, we found that the average benefits available to a childless household with no members aged 65 or older was \$3,500.¹⁹ A value of \$3,500 understates the benefits available to individuals over age 65; in 1998 the Federal SSI benefit for elderly (65+) couples was nearly \$9,000 (Committee on Ways and Means, 2000, p. 229). On the other hand, about half of eligible households do not collect SSI benefits (Elder and Powers, 2006, Table 2), possibly because transactions or “stigma” costs outweigh the value of public assistance. Low take-up rates, along with the costs that probably underly them, suggest that the effective consumption floor need not equal statutory benefits.

The bequest parameters θ_B and κ are identified largely from the top asset quantile. It follows from equation (3) that when the shift parameter κ is large, the marginal utility of bequests will be lower than the marginal utility of consumption unless the individual is rich.

¹⁹Our treatment of the consumption floor differs markedly from that of Rust and Phelan (1997) and Blau and Gilleskie (2006a, 2006b), 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.

In other words, the bequest motive mainly affects the saving of the rich; for more on this point, see De Nardi (2004). Our estimate of θ_B implies that the marginal propensity to consume out of wealth in the final period of life (which is a nonlinear function of θ_B , β , γ , ν and κ) is 1 for low income individuals and 0.032 for high-income individuals.

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 0.535 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 0.404 at age 60, implying that wage changes cause relatively small hours changes for workers. For example, the Frisch labor supply elasticity of a type-1 individual working 2000 hours per year is approximated as $-\frac{L-H_t-\phi_P}{H_t} \times \frac{1}{(1-\gamma)(1-\nu)-1} = 0.21$.

The fixed cost of work, ϕ_P , 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, 50 weeks per year) 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 time endowment L is identified by the combination of the participation and hours profiles. The time cost of bad health, ϕ_{HS} , is identified by noting that unhealthy individuals work fewer hours than healthy individuals, even after conditioning on wages.²⁰

6.2 Simulated Profiles

The bottom of Table 4 displays the overidentification test statistic. 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 4 shows that the model fits both asset quantiles fairly well. The model is able to

²⁰In the current specification, we have not imposed any re-entry costs. Adding previous employment as a state variable doubles the computational burden, and the current specification matches observed re-entry patterns very well.




Figure 4: ASSET QUANTILES, DATA AND SIMULATIONS

fit the lower quantile in large part because of the consumption floor of \$4,118; the predicted lower asset quantile rises dramatically when the consumption floor is lowered. This result is consistent with Hubbard, Skinner, and Zeldes (1995). They show that if the government guarantees a minimum consumption level, those with low income will tend not to save because their savings will reduce the transfers they receive from the government. It is therefore not surprising that within the model the consumption floor reduces saving by individuals with low income and assets.

The three panels in the left hand column of Figure 5 show that the model is able to replicate the two key features of how labor force participation varies with 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 5.3 percentage point decline in the data versus a 5.8 percentage point decline predicted by the model), although it underpredicts the decline in participation at 62 (a 10.6 percentage point decline in the data versus a 3.5 percentage point decline predicted by the model).









Figure 5: PARTICIPATION AND JOB EXIT RATES, DATA AND SIMULATIONS

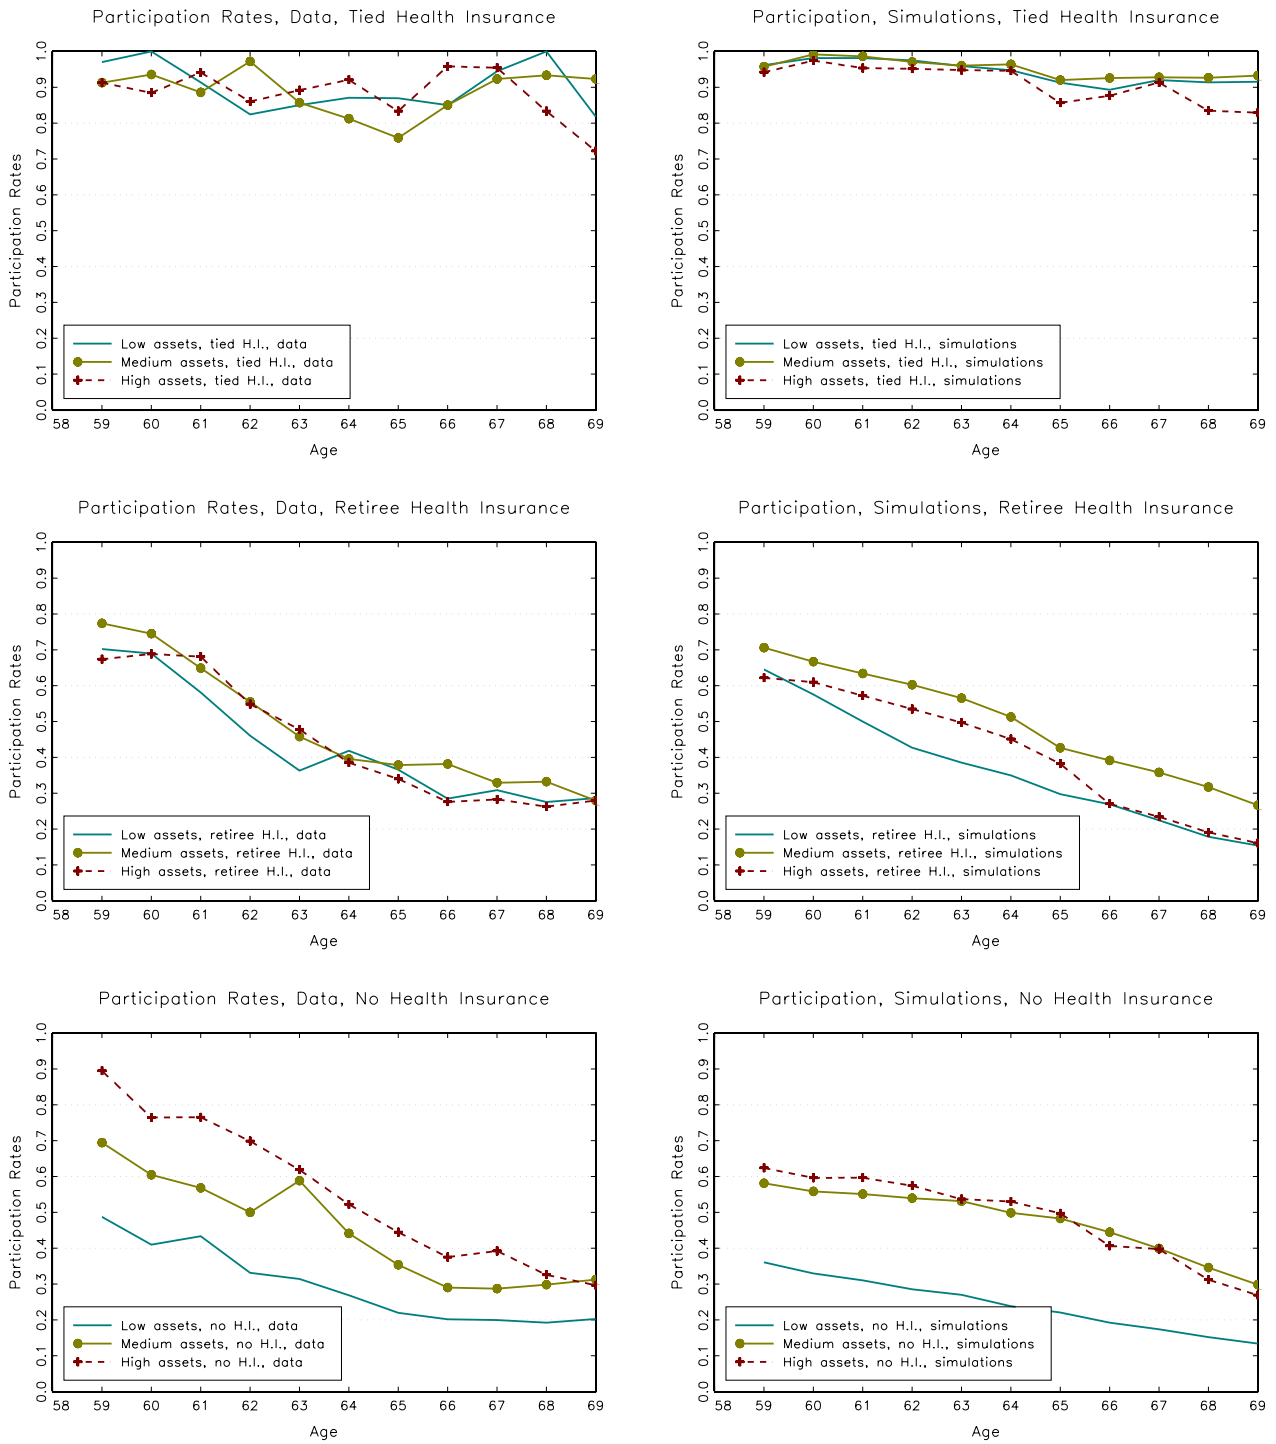


Figure 6: LABOR FORCE PARTICIPATION RATES BY ASSET GROUPING, DATA AND SIMULATIONS

The second key feature is that there are large differences in participation and job exit rates across health insurance types. The model does a good job of replicating observed differences in participation rates. For example, the model matches the low participation levels of the uninsured. Turning to the lower left panel of Figure 6, the data show that the group with the lowest participation rates are the uninsured with low assets. The model is able to replicate this fact because of the consumption floor. 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 5 compare observed and simulated job exit rates for each health insurance type. They show that the model correctly predicts that workers with retiree coverage and no health insurance have fairly high exit rates after age 62. In contrast, the model under-predicts exit for workers with *tied* health insurance.




Figure 7: LABOR FORCE PARTICIPATION RATES BY PREFERENCE INDEX, DATA AND SIMULATIONS

Figure 7 shows how participation differs across the three values of our discretized preference index. The model does a good job of replicating the observed differences in participation. Recall that an index value of *out* implies that the individual was not working in 1992. Not surprisingly, participation for this group is always low. Individuals with positive values of the preference index differ primarily in the rate at which they leave the labor force, i.e., the slopes of their participation profiles. As noted in our discussion of the preference parameters, the

model replicates these differences by allowing the taste for leisure (γ) and the discount rate (β) to vary across preference types. When we do not allow for preference heterogeneity, the model is unable to replicate the patterns observed in Figure 7. This highlights the importance of the preference index in identifying preference heterogeneity in the population.

6.3 The Effects of Employer-Provided Health Insurance

The labor supply patterns shown in Figures 3 and 5 show a correlation between health insurance and labor supply. However, they do not identify the effects of health insurance on retirement, for three reasons. First, as shown in Table 3, the distributions of wages and wealth in our sample differ across health insurance types. For example, those with retiree coverage have greater pension wealth than other groups. Second, as shown in Figure 1, pension plans for workers with retiree coverage provide stronger incentives for early retirement than the pension plans held by other groups. Third, as shown in Table 5, preferences for leisure vary by health insurance type. In short, retirement incentives differ across health insurance categories for reasons unrelated to health insurance incentives.

To isolate the effects of employer-provided health insurance on labor supply, we conduct some additional simulations. We fix pension accrual rates so that they are identical across health insurance types. We then simulate the model twice, assuming first that all workers *retiree* health insurance coverage at age 59, then *tied* coverage at age 59. Across the two simulations, households face different medical expense distributions, but in all other dimensions the distribution of incentives faced by individuals is identical.

This exercise reveals that the job exit rate at age 60 would be 2.6 percentage points higher if all workers had *retiree* coverage rather than *tied* coverage. The gap is 3.8 percentage points at age 61 and 3.2 percentage points at age 62, then declines slightly to 1.0 percentage points at age 65. These differences in exit rates across health insurance types are smaller than the raw differences in exit rates observed in the data (see Section 5) and the raw differences predicted by the model. Such results are consistent with Tables 3 and 5, which show that workers with retiree coverage have more generous pension plans and stronger preferences for leisure than

those with tied coverage. Failing to account for these effects will lead the econometrician to overstate the effect of health insurance on exit rates.

The effect of health insurance can also be measured by comparing participation rates. We find that the labor force participation rate for ages 60-67 would be 6.0 percentage points lower if workers had retiree, rather than tied, coverage at age 59. Yet another way to measure the effect of health insurance is consider 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 0.41 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. They also 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 lower bound of 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 larger than Blau and Gilleskie's (2006a, 2006b), who find that retiree coverage reduces average labor force participation 1.7 and 1.6 percentage points, respectively,²¹ but are smaller than the effects

²¹Blau and Gilleskie (2006a) consider the retirement decision of couples, and allow husbands and wives to

found by Rust and Phelan (1997).

6.4 Model Validation

In order to better understand whether structural models produce accurate predictions, it has become increasingly common to subject them to out-of-sample validation exercises (see, e.g., Keane and Wolpin, 2006, and the references therein). Recall that we estimate the model on a cohort of individuals aged 57-61 in 1992. We test our model by considering the HRS cohort aged 51-55 in 1992; we refer to this cohort as our validation sample. These individuals faced different Social Security incentives than did the estimation cohort. The validation sample did not face the Social Security earnings test after age 65, had a slightly later full retirement age, and faced a benefit adjustment formula that more strongly encouraged delayed retirement. In addition to facing different Social Security rules, the validation sample possessed different endowments of wages, wealth, and employer benefits. A valuable test of our model, therefore, is to see if it can predict the behavior of the validation sample.

Age	Data			Model		
	1933 (1)	1939 (2)	Difference [†] (3)	1933 (4)	1939 (5)	Difference* (6)
60	0.694	0.700	0.006	0.621	0.654	0.033
61	0.656	0.652	-0.003	0.589	0.619	0.030
62	0.551	0.549	-0.002	0.554	0.574	0.021
63	0.487	0.509	0.023	0.522	0.537	0.016
64	0.433	0.475	0.042	0.484	0.501	0.017
65	0.379	0.427	0.048	0.426	0.444	0.018
66	0.338	0.429	0.091	0.370	0.400	0.030
67	0.327	0.484	0.157	0.338	0.343	0.005
Total, 60-65	3.198	3.312	0.114	3.195	3.330	0.135
Total, 60-67	3.863	4.225	0.362	3.903	4.073	0.171
	† Column (2) – Column (1)			*Column (5) –Column (4)		

Table 6: PARTICIPATION RATES BY BIRTH YEAR COHORT

retire at different dates. Blau and Gilleskie (2006b) allow workers to choose their medical expenses. Because these modifications provide additional mechanisms for smoothing consumption over medical expense shocks, they could reduce the effect of employer-provided health insurance. Furthermore, eliminating the ability to save reduces a worker's willingness to substitute labor across time; as current consumption becomes more closely linked to current earnings, labor supply becomes less flexible. Domeij and Floden (2006) find that borrowing constraints reduce the effective intertemporal elasticity of substitution by 50 percent.

Columns (1)-(3) of Table 6 show the participation rates observed in the data for each cohort, and the difference. The data suggest that the change in the Social Security rules coincides with increased labor force participation, especially at later ages. The estimated increase in labor supply at ages 62-65 is similar, and the estimated increase at ages 66-67 larger than the increases in labor supply reported in Song and Manchester (2007).²²

Columns (4)-(6) of Table 6 show the differences predicted by the model. The simulations for the validation sample use the initial distribution observed for the validation cohort, but use the decision rules estimated on the older estimation cohort.²³ Comparing Columns (3) and (6) shows that although the model does not always match the data year-by-year, it predicts that total labor supply over ages 60-65 will increase by 0.135 years, compared to the difference of 0.114 years years in the data. We conclude that the model does a good job of fitting the data out of sample.

7 Policy Experiments

The preceding section showed that the model fits the data very well, given plausible preference parameters. In this section, we use the model to predict 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 7.

The first column of Table 7 shows model-predicted labor market participation at ages 60 through 67 under the 1998 Social Security rules. Under the 1998 rules, the average person works a total of 3.90 years over this eight-year period. The fifth column of Table 7 shows that this is close to the total of 3.86 years observed in the data.

The second column shows the average hours that result when the 1998 Social Security rules

²²Our participation rates for ages 66 and 67 are imprecisely estimated because at later ages we observe a decreasing fraction of the validation sample; at age 66, for example, we observe only the individuals born in 1937 and 1938—roughly two fifths of the sample—and at age 67 we observe only the individuals born in 1937.

²³We do not adjust for business cycle conditions.

Age	1998 rules:	2030 rules:		Data	
	SS = 65 MC = 65 (1)	SS = 67 MC = 65 (2)	SS = 65 MC = 67 (3)		SS = 67 MC = 67 (4)
60	0.621	0.628	0.624	0.633	0.694
61	0.589	0.597	0.594	0.602	0.656
62	0.554	0.559	0.558	0.566	0.550
63	0.522	0.527	0.527	0.534	0.487
64	0.484	0.492	0.491	0.499	0.432
65	0.426	0.435	0.445	0.454	0.379
66	0.370	0.401	0.389	0.416	0.338
67	0.338	0.355	0.342	0.358	0.327
Total 60-67	3.903	3.994	3.971	4.062	3.863
SS = Social Security normal retirement age					
MC = Medicare eligibility age					

Table 7: EFFECTS OF CHANGING THE SOCIAL SECURITY RETIREMENT AND MEDICARE ELIGIBILITY AGES

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 at the normal retirement age or older. The second column shows that imposing the 2030 rules leads the average worker to increase years worked between ages 60 and 67 from 3.90 years to 3.99 years, an increase of 0.09 years.²⁴

The third column of Table 7 shows participation when the Medicare eligibility age is increased to 67.²⁵ This change increases total years of work by 0.07 years. Averaged over an 8-year interval, a 0.07-year increase in total years of work translates into a 0.9-percentage-point increase in annual participation rates. This amount is larger than the changes found by

²⁴In addition to changing the benefit accrual rate, raising the normal retirement age from 65 to 67 effectively eliminates two years of Social Security benefits. Therefore, raising the normal retirement age has both substitution and wealth effects, both of which cause participation to increase. 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. Using this configuration to eliminate wealth effects, we find that total years of work increase by 0.048 years, implying that 0.043 years of the 0.091-year increase is due to wealth effects.

²⁵By 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.

Blau and Gilleskie (2006a), whose simulations show that increasing the Medicare age reduces the average probability of non-employment by about 0.1 percentage points, but is smaller than the effects suggested by Rust and Phelan’s (1997) analysis. 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 0.16 years, 0.07 more than that generated by raising the Social Security normal retirement age in isolation.

In short, the model predicts that raising the normal Social Security retirement age will have a slightly 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 most workers in our sample do not have *tied* coverage at age 59.²⁶ Medicare provides smaller retirement incentives to workers in the *retiree* or *none* categories. Simulations reveal that for those with *tied* coverage at age 60, shifting forward the Social Security age to 67 increases years in the labor force by 0.09 years, whereas shifting forward the Medicare eligibility age to 67 would increase years in the labor force by 0.11 years.

To understand better the incentives generated by Medicare, we compute the value that Type-1 individuals place on employer-provided health insurance, by finding the increase in assets that would make an uninsured Type-1 individual as well off as a person with retiree coverage. In particular, we find the compensating variation $\lambda_t = \lambda(A_t, B_t, HS_t, AIME_t, \omega_t, \zeta_{t-1}, t)$, where

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

Table 8 shows the compensating variation $\lambda(A_t, 0, \textit{good}, \$32000, 0, 0, 60)$ at several different

²⁶Only 13% of the workers in our sample had *tied* coverage at age 59. This figure, however, is probably too low. For example, Kaiser/HRET (2006) estimates that about 50% of large firms offered *tied* coverage in the mid-1990s. One potential reason that we may be understating the share with *tied* coverage is that, as shown in the Kaiser/HRET (2006) study, the fraction of workers with *tied* (instead of *retiree*) coverage grew rapidly in the 1990s, and our health insurance measure is based on wave-1 data collected 1992. In fact, the HRS data indicate that later waves had a higher proportion of individuals with *tied* coverage than in wave 1. We may also be understating the share with *tied* coverage because of changes in the wording of the HRS questionnaire; see Appendix G for details.

asset (A_t) levels.²⁷ The first column of Table 8 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 through much of the wealth distribution. Even though richer 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 over the asset range of $-\$2,300$ to $\$149,000$. However, individuals with asset levels of $\$600,000$ place less value on retiree coverage, because they can better self-insure against medical expense shocks.

Asset Levels	Compensating Assets		Compensating Annuity	
	With Uncertainty	Without Uncertainty	With Uncertainty	Without Uncertainty
	(1)	(2)	(3)	(4)
Baseline Case				
$-\$2,300$	\$51,150	\$18,710	\$4,580	\$1,970
$\$54,400$	\$55,860	\$18,700	\$4,070	\$1,960
$\$149,000$	\$57,170	\$19,690	\$3,980	\$1,950
$\$600,000$	\$40,000	\$19,500	\$2,830	\$1,780
No-Saving Case				
$-\$2,150$	463,700	\$28,100	\$11,100	\$1,910
Compensating variation between <i>retiree</i> and <i>none</i> coverages				
Calculations described in text				
Baseline results are for agents with type-1 preferences				

Table 8: 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 medical expense uncertainty, by setting the variance shifter $\sigma(HS_t, HI_t, t, B_t, P_t)$ to zero, and recompute λ_t , using the same state variables and mean medical expenses as before. Without medical expense uncertainty, λ_t is approximately $\$19,000$. Comparing the two values of λ_t shows that for the typical worker (with $\$150,000$ of assets) about one-third

²⁷In 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 of HS_t and ζ_{t-1} considered here, the conditional differences in expected medical expenses are smaller than the unconditional differences shown in Table 1.

of the value of health insurance comes from the reduction of average medical expenses, and two-thirds is due to the reduction of medical expense volatility.

The first two columns of Table 8 measure the lifetime value of health insurance as an asset increment that can be consumed immediately. An alternative approach is to express the value of health insurance as an illiquid annuity comparable to Social Security benefits. Columns (3) and (4) show this “compensating annuity”.²⁸ When the value of health insurance is expressed as an annuity, the fraction of its value attributable to reduced medical expense volatility falls from two-thirds to about one-half. In most other respects, however, the asset and annuity valuations of health insurance have similar implications.

To sum, allowing for medical expense uncertainty greatly increases the value of health insurance. It is therefore unsurprising that we find larger effects of health insurance on retirement than do Gustman and Steinmeier (1994) and Lumsdaine et al. (1994), who assume that workers value health insurance at its actuarial cost.

8 Alternative Specifications

To consider whether our findings are sensitive to our modelling assumptions, we re-estimate the model under two alternate specifications.²⁹ Table 9 shows model-predicted participation rates under the different specifications, along with the data. Column (1) of Table 9 presents our baseline case. Column (2) presents the case where individuals are not allowed to save. Column (3) presents the case with no preference heterogeneity. Column (4) presents the data. In general, the different specifications generate similar profiles.

²⁸To do this, we first find compensating $AIME$, $\hat{\lambda}_t$, where

$$V_t(A_t, B_t, HS_t, AIME_t, \omega_t, \zeta_{t-1}, retiree) = V_t(A_t, B_t, HS_t, AIME_t + \hat{\lambda}_t, \omega_t, \zeta_{t-1}, none).$$

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.

²⁹We have also estimated a specification where housing wealth is illiquid. As described in Appendix L, parameter estimates and model fit for this case were somewhat different than our baseline results, although the policy simulations looked similar. In earlier drafts of this paper (French and Jones, 2004b) we estimated a specification where we added measurement error to the simulated asset histories. Adding measurement error, however, had little effect on either the preference parameter estimates or policy experiments, and we thus dropped this case.

Age	Baseline (1)	No Saving (2)	Homogeneous Preferences (3)	Data (4)
60	0.621	0.605	0.676	0.695
61	0.589	0.581	0.639	0.656
62	0.554	0.536	0.570	0.550
63	0.522	0.512	0.513	0.486
64	0.484	0.480	0.464	0.433
65	0.426	0.449	0.440	0.379
66	0.370	0.406	0.372	0.339
67	0.338	0.322	0.367	0.327
Total 60-67	3.903	3.891	4.041	3.866

Table 9: MODEL PREDICTED PARTICIPATION, BY AGE: ALTERNATIVE SPECIFICATIONS

The parameter estimates behind these simulations can be found in Table 14 of Appendix K, which presents parameter estimates for all the specifications.

8.1 No Saving

We have argued that the ability to self-insure through saving significantly affects 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 require workers to consume their income net of medical expenses, as in Rust and Phelan (1997) and Blau and Gilleskie (2006a, 2006b).

The second column of Table 9 shows participation rates. The baseline case fits the labor supply profiles slightly better, and obviously fits the asset profiles much better, than the no-savings case.³⁰

The compensating annuity calculations in Table 8 show that eliminating the ability to save greatly increases the value of retiree coverage: when assets are -\$2,000, the compensating annuity increases from \$4,600 in the baseline case (with savings) to \$11,100 in the no-savings case. When there is no medical expense uncertainty, the comparable figures are \$1,970 in

³⁰Because the baseline and no-savings cases are estimated with different moments, the overidentification statistics shown in the first two columns of Table 4 are not comparable. However, inserting the decision profiles generated by the baseline model into the moment conditions used to estimate the no-savings case produces an overidentification statistic of 958, while the no-saving specification produces an overidentification statistic of 1,211.

the baseline case and \$1,910 in the no-savings case. Thus, the ability to self-insure through saving significantly reduces the value of employer-provided health insurance.

Simulating the responses to policy changes, we find that raising the Medicare eligibility age to 67 leads to an additional 0.05 years of work, an amount close to that of the baseline specification. Moving the Social Security normal retirement age to 67 generates an almost identical response, which is also consistent with the baseline results.

8.2 No Preference Heterogeneity

To assess the importance of preference heterogeneity, we estimate and simulate a model where individuals have identical preferences. Table 14 in the Appendix contains the revised parameter estimates.

Comparing columns (1), (3) and (4) of Table 9 shows that the model without preference heterogeneity matches aggregate participation rates as well as the baseline model. However, the no-heterogeneity specification does much less well in replicating the way in which participation varies across the asset distribution, and, not surprisingly, does not replicate the way in which participation varies across our discretized preference index.³¹

When preferences are homogenous the simulated response to delaying the Medicare eligibility age, 0.11 years, is larger than the response in the baseline specification, and it exceeds the effect of increasing the Social Security normal retirement age. This provides further evidence that failure to account for preference heterogeneity and self-selection into health insurance plan likely lead us to overstate the effect of health insurance on retirement.

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. To see if Medicare benefits have

³¹We do not include the index-related moments in the revised GMM criterion function.

a large effect on retirement behavior, we construct 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 estimate the structural parameters of our model. The model fits the data well, with reasonable preference parameters. In addition, the model does a good job of predicting the behavior of individuals who, by belonging to a younger cohort, faced different Social Security rules than the individuals upon which the model was estimated.

We find that health care uncertainty significantly affects the value of employer-provided health insurance. Our calculations suggest that about two thirds of the value workers place on employer-provided health insurance comes from its ability to reduce medical expense risk. We also find, however, that the ability to save significantly reduces the value of health insurance: when saving is prohibited, the value of insurance doubles. Furthermore, we find evidence of self-selection into employer-provided health insurance plans, which also reduces the estimated effect of Medicare on retirement. Nevertheless, we find that the labor supply effects of raising the Medicare eligibility age from 65 to 67 (0.07 years) are just as important as the effects of raising the Social Security retirement age from 65 to 67.

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, 22(2), 329-352.
- [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,” *Journal of Applied Econometrics*, 2006a, 21(7), 935-953.
- [6] **Blau, D. and D. Gilleskie**, “The Role of Retiree Health Insurance in the Employment Behavior of Older Men,” mimeo, 2006b.
- [7] **The Boards of Trustees of the Hospital Insurance and Supplementary Medical Insurance Trust Funds**, *2006 Annual Report of the Boards of Trustees of the Hospital Insurance and Supplementary Medical Insurance Trust Funds*, 2006.
- [8] **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.
- [9] **Bound, J., T. Stinebrickner, and T. Waidmann**, “Health, Economic Resources, and the Work Decisions of Older Men,” mimeo, 2003.
- [10] **Buchinsky, M.**, “Recent Advances in Quantile Regression Models: A Practical Guide-line for Empirical Research,” *Journal of Human Resources*, 1998, 33, 88-126.
- [11] **Cagetti, M.**, “Wealth Accumulation Over the Life Cycle and Precautionary Savings,” *Journal of Business and Economic Statistics*, 2003, 21(3), 339-353.
- [12] **Cochrane, J.**, “A Simple Test of Consumption Insurance,” *The Journal of Political Economy*, 1991, 99(5), 957-976.
- [13] **Chernozhukov, V., and C. Hansen**, “An IV Model of Quantile Treatment Effects,” MIT Working Paper 02-06, 2002.
- [14] **Committee On Ways And Means, U.S. House Of Representatives**, *Green Book*, United States Government Printing Office, various years.
- [15] **Cogan, J.**, “Fixed Costs and Labor Supply,” *Econometrica*, 1981, 49(4), 945-963.
- [16] **David, M., R. Little, M. Samuhel, and R. Triest**, “Alternative Methods for CPS Income Imputation,” *Journal of the American Statistical Association*, 1986, 81(393), 29-41.
- [17] **De Nardi, C.**, “Wealth Distribution, Intergenerational Links and Estate Taxation,” *Review of Economic Studies*, 2004, 71(3), 743-768.

- [18] **Domeij, D., and M. Floden**, “The Labor-supply Elasticity and Borrowing Constraints: Why Estimates are Biased,” *Review of Economic Dynamics*, 2006, 9(2), 242-262.
- [19] **Duffie, D. and K. Singleton**, “Simulated Moments Estimation of Markov Models of Asset Prices,” *Econometrica*, 1993, 61(4), 929-952.
- [20] **Elder, T. and E. Powers**, “The Incredible Shrinking Program: Trends in SSI Participation of the Aged,” *Research on Aging* 2006, 28(3), 341-358.
- [21] **Employee Benefit Research Institute**, *EBRI Health Benefits Databook*, EBRI-ERF, 1999.
- [22] **Epple, D. and H. Seig**, “Estimating Equilibrium Models of Local Jurisdictions,” *Journal of Political Economy* 1999, 107(4), 645-681.
- [23] **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.
- [24] **French, E.**, “The Labor Supply Response to Predictable (but Mismeasured) Wage Changes”, *Review of Economics and Statistics*, 2004, 86(2), 602-613.
- [25] **French, E.**, “The Effects of Health, Wealth and Wages on Labor Supply and Retirement Behavior,” *Review of Economic Studies*, 2005, 72(2), 395-427.
- [26] **French, E., and J. Jones**, “On the Distribution and Dynamics of Health Care Costs,” *Journal of Applied Econometrics*, 2004a, 19(4), 705-721.
- [27] **French, E., and J. Jones**, “The Effects of Health Insurance and Self-Insurance on Retirement Behavior,” Center for Retirement Research Working Paper 2004-12, 2004b.
- [28] **Gourieroux, C., and A. Monfort**, *Simulation-Based Econometric Methods*, Oxford University Press, 1997.
- [29] **Gourinchas, P. and Parker, J.**, “Consumption Over the Life Cycle,” *Econometrica*, 2002, 70(1), 47-89.
- [30] **Gruber, J., and B. Madrian**, “Health Insurance Availability and the Retirement Decision,” *American Economic Review*, 1995, 85(4), 938-948.
- [31] **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.
- [32] **Gustman, A., and T. Steinmeier**, “A Structural Retirement Model,” *Econometrica*, 1986, 54(3), 555-584.
- [33] **Gustman, A., and T. Steinmeier**, “Employer-Provided Health Insurance and Retirement Behavior,” *Industrial and Labor Relations Review* 1994, 48(1), 124-140.
- [34] **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.

- [35] **Gustman, A., and T. Steinmeier**, “The Social Security Early Entitlement Age in a Structural Model of Retirement and Wealth,” *Journal of Public Economics*, 2005, 89, 441-463.
- [36] **Gustman, A., O. Mitchell, A. Samwick and T. Steinmeier**, “Evaluating Pension Entitlements,” 1998, mimeo.
- [37] **Heckman, J., and B. Singer**, “A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data,” *Econometrica*, 1984, 52(2), 271-320.
- [38] **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.
- [39] **Hubbard, R., J. Skinner, and S. Zeldes**, “Precautionary Saving and Social Insurance,” *Journal of Political Economy*, 1995, 103(2), 360-399.
- [40] **Judd, K.**, *Numerical Methods in Economics*, MIT Press, 1998.
- [41] **Kaiser/HRET**, *The 2006 Kaiser/HRET Employer Health Benefit Survey*. <http://www.kff.org/insurance/7527/upload/7527.pdf>, 2006.
- [42] **Kahn, J.**, “Social Security, Liquidity, and Early Retirement,” *Journal of Public Economics*, 1988, 35, 97-117.
- [43] **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.
- [44] **Keane, M., and K. Wolpin**, “The Career Decisions of Young Men,” *Journal of Political Economy*, 1997, 105(3), 473-522.
- [45] **Keane, M., and K. Wolpin**, “The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment,” *International Economic Review*, 2001, 42(4), 1051-1103.
- [46] **Keane, M., and K. Wolpin**, “Exploring the Usefulness of a Non-Random Holdout Sample for Model Validation: Welfare Effects on Female Behavior,” PIER Working Paper 06-006 and forthcoming *International Economic Review*.
- [47] **Little, R.**, “Missing Data Adjustments in Large Surveys,” *Journal of Business and Economic Statistics*, 1988, 6(3), 287-301.
- [48] **Low, H.**, “Self-insurance in a Life-cycle Model of Labour Supply and Savings”, *Review of Economic Dynamics*, 2005, 8, 945-975.
- [49] **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.
- [50] **MaCurdy, T.**, “An Empirical Model of Labor Supply in a Life-Cycle Setting,” *Journal of Political Economy*, 1981, 89(6), 1059-1085.

- [51] **Madrian, B.**, “The Effect of Health Insurance on Retirement,” *Brookings Papers on Economic Activity*, 1994, 181-252.
- [52] **Manski, C.**, *Analog Estimation Methods in Econometrics*, 1988.
- [53] **Newey, W.**, “Generalized Method of Moments Specification Testing,” *Journal of Econometrics*, 1985, 29(3), 229-256.
- [54] **Newey, W. and D. McFadden**, “Large Sample Estimation and Hypothesis Testing” in R. Engle and D. McFadden (eds.) *Handbook of Econometrics*, Vol. 4., 1994.
- [55] **Pakes, A., and D. Pollard**, “Simulation and the Asymptotics of Optimization Estimators,” *Econometrica*, 1989, 57(5), 1027-1057.
- [56] **Palumbo, M.**, “Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle,” *Review of Economic Studies*, 1999, 66(2), 395-421.
- [57] **Pischke, J-S.**, “Measurement Error and Earnings Dynamics: Some Estimates From the PSID Validation Study,” *Journal of Business & Economic Statistics*, 1995, 13(3), 305-314.
- [58] **Powell, J.**, “Estimation of Semiparametric Models” in R. Engle and D. McFadden (eds.) *Handbook of Econometrics*, Vol. 4., 1994.
- [59] **Ruhm, C.**, “Bridge Jobs and Partial Retirement,” *Journal of Labor Economics*, 1990, 8(4), 482-501.
- [60] **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.
- [61] **Rust, J., Buchinsky, M., and H. Benitez-Silva**, “An Empirical Model of Social Insurance at the End of the Life Cycle,” mimeo, 2003.
- [62] **Smith, J.**, “Healthy Bodies and Thick Wallets: The Dual Relationship between Health and Economic Status,” *Journal of Economic Perspectives*, 13(2), 1999, 145-166.
- [63] **Song, J., and J. Manchester**, “New Evidence on Earnings and Benefit Claims Following Changes in the Retirement Earnings Test in 2000,” *Journal of Public Economics*, 91(3), 2007, 669-700.
- [64] **Tauchen, G.**, “Finite State Markov-chain Approximations to Univariate and Vector Autoregressions.” *Economics Letters*, 20, 1986, 177-181.
- [65] **United States Social Security Administration**, *Social Security Bulletin: Annual Statistical Supplement*, United States Government Printing Office, selected years.
- [66] **van der Klaauw, W., and K. Wolpin**, “Social Security, Pensions and the Savings and Retirement Behavior of Households,” mimeo, 2006.

Appendix A: Cast of Characters

<i>Preference Parameters</i>		<i>Health-related Parameters</i>	
γ	consumption weight	HS_t	health status
β	time discount factor	m_t	out-of-pocket medical expenses
ν	coefficient of RRA, utility	HI_t	health insurance type
L	leisure endowment	$m(\cdot)$	mean shifter, logged medical expenses
ϕ_P	fixed cost of work	$\sigma(\cdot)$	volatility shifter, logged medical expenses
ϕ_{HS}	leisure cost of bad health	ψ_t	idiosyncratic medical expense shock
θ_B	bequest weight	ζ_t	persistent medical expense shock
κ	bequest shifter	ϵ_t	innovation, persistent shock
C_{min}	consumption floor	ρ_m	autocorrelation, persistent shock
<i>Decision Variables</i>		σ_ϵ^2	innovation variance, persistent shock
C_t	consumption	ξ_t	transitory medical expense shock
H_t	hours of work	σ_ξ^2	variance, transitory shock
L_t	leisure	<i>Wage-related Parameters</i>	
P_t	participation	W_t	hourly wage
A_t	assets	$w(\cdot)$	mean shifter, logged wages
B_t	Social Security application	α	coefficient on hours, logged wages
<i>Financial Variables</i>		ω_t	idiosyncratic wage shock
$Y(\cdot)$	after-tax income	ρ_W	autocorrelation, wage shock
τ	tax parameter vector	η_t	innovation, wage shock
r	real interest rate	σ_η^2	innovation variance, wage shock
ys_t	spousal income	<i>Miscellaneous</i>	
$ys(\cdot)$	mean shifter, spousal income	s_t	survival probability
ss_t	Social Security income	$pref$	discrete preference index
$AI ME_t$	Social Security wealth	X_t	state vector, worker's problem
pb_t	pension benefits	$\lambda(\cdot)$	compensating variation

Table 10: VARIABLE DEFINITIONS, MAIN TEXT

Appendix B: 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 11: AFTER TAX INCOME

Appendix C: Pensions

Although the HRS pension data and pension calculator allow one to estimate pension wealth with a high degree of precision, Bellman’s curse of dimensionality prevents us from including in our dynamic programming model the full range of pension heterogeneity found in the data. Thus we thus use the HRS pension data and calculator to construct a simpler model. The fundamental equation behind our model of pensions is the accumulation equation for pension wealth, pw_t :

$$pw_{t+1} = \begin{cases} (1/s_{t+1})[(1+r)pw_t + pacc_t - pb_t] & \text{if living at } t+1 \\ 0 & \text{otherwise} \end{cases} \quad (15)$$

where $pacc_t$ is pension accrual and pb_t is pension benefits. Two features of this equation bear noting. First, a pension is worthless once an individual dies. Therefore, in order to be actuarially fair, surviving workers must receive an above-market return on their pension balances. Dividing through by the survival probability s_{t+1} ensures that the expected value of pensions $E(pw_{t+1}|pw_t, pacc_t, pb_t)$ equals $(1+r)pw_t + pacc_t - pb_t$. 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.

Simulating equation (15) requires us to know pension benefits and pension accrual. We calculate pension benefits by assuming that at age t , the worker receives the *expected* pension

benefit

$$pb_t = pft \times pb_t^{\max}, \quad (16)$$

where pb_t^{\max} is the benefit received by individuals actually receiving pensions (given the earnings history observed at time t) and pft the probability that a person with a pension is currently drawing pension benefits. We estimate pft as the fraction of respondents who are covered by a pension that receive pension benefits 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} given the earnings history at time t (and assuming no further pension accruals so that $pacc_k = 0$ for $k = t, t + 1, \dots, T$), note first that recursively substituting equation (15) and imposing $pw_{T+1} = 0$ reveals that pension wealth is equal to the present discounted value of future pension benefits:

$$pw_t = \frac{1}{1+r} \sum_{k=t}^T \frac{S(k,t)}{(1+r)^{k-t}} p f_k p b_t^{\max}, \quad (17)$$

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 the maximum pension benefit is constant from time t forward, so that $pb_k^{\max} = pb_t^{\max}$, $k = t, t + 1, \dots, T$, this equation reduces to

$$pw_t = \Gamma_t p b_t^{\max}, \quad (18)$$

$$\Gamma_t \equiv \frac{1}{1+r} \sum_{k=t}^T \frac{S(k,t)}{(1+r)^{k-t}} p f_k. \quad (19)$$

Using equations (16) and (18), pension benefits are thus given by

$$pb_t = p f_t \Gamma_t^{-1} p w_t. \quad (20)$$

Next, we assume pension accrual is given by

$$pacc_t = \alpha_0(HI_t, W_t H_t, t) \times W_t H_t, \quad (21)$$

where $\alpha_0(\cdot)$ is the pension accrual rate as a function of health insurance type, labor income,

and age. We estimate $\alpha_0(\cdot)$ in two steps, estimating separately each component of:

$$\alpha_0 = E(pacc_t | W_t H_t, HI_t, t, pen_t = 1) \Pr(pen_t = 1 | HI_t, W_t H_t) \quad (22)$$

where $pacc_t$ is the accrual rate for those with a pension, and pen_t is a 0-1 indicator equal to 1 if the individual has a pension.

We estimate the first component, $E(pacc_t | W_t H_t, HI_t, t, pen_t = 1)$, from restricted HRS pension data. To generate a pension accrual rate for each individual, we combine the pension data with the HRS pension calculator to estimate the pension wealth that each individual would have if he left his job at different ages. The increase in pension wealth gained by working one more year is the accrual. Put differently, if pension benefits are 0 as long as the worker continues working, it follows from equation (15) that

$$pacc_t = s_{t+1}pw_{t+1} - (1+r)pw_t. \quad (23)$$

It bears noting that the HRS pension data have a high degree of employer- and worker-level detail, allowing us to estimate pension accrual quite accurately. With accruals in hand, we then estimate $E(pacc_t | W_t H_t, HI_t, t, pen_t = 1)$ on the subset of workers that have a pension on their current job. We regress accrual rates on a fourth-order age polynomial, indicators for age greater than 62 or 65, log income, log income interacted with the age variables, health insurance indicators, and health insurance indicators interacted with the age variables.

Figure 8 shows estimated pension accrual, by health insurance type and earnings. It shows that those with retiree coverage have the sharpest declines in pension accrual after age 60. It also shows that once health insurance and the probability of having a pension plan are accounted for, the effect of income on pension accrual is relatively small. Our estimated age (but not health insurance) pension accrual rates line up closely with Gustman et al. (1998), who also use the restricted firm based HRS pension data.

In the second step, we estimate the probability of having a pension, $\Pr(pen_t = 1 | HI_t, W_t H_t, t)$, using unrestricted self-reported data from individuals who are working and are ages 51-55.




Figure 8: PENSION ACCRUAL RATES FOR INDIVIDUALS WITH PENSIONS, BY AGE, HEALTH INSURANCE COVERAGE AND EARNINGS

The function $\Pr(\text{pen}_t = 1 | HI_t, W_t H_t, t)$ is estimated as a logistic function of log income, health insurance indicators, and interactions between log income and health insurance.

Table 12 shows the probability of having different types of pensions, conditional on health insurance. The table shows that only 8% of men with no health insurance have a pension, but 64% of men with tied coverage and 74% of men with retiree insurance have a pension. Furthermore, it shows that those with retiree coverage are also the most likely to have defined benefit (DB) pension plans, which provide the strongest retirement incentives at age 65.

Variable	Probability of Pension Type		
	No Insurance	Retiree Insurance	Tied Insurance
Defined Benefit	.026	.412	.260
Defined Contribution	.050	.172	.270
Both DB and DC	.006	.160	.106
Total	.082	.744	.636
Number of Observations	343	955	369

Table 12: PROBABILITY OF HAVING A PENSION ON THE CURRENT JOB, BY HEALTH INSURANCE TYPE, WORKING MEN, AGE 51-55

Combining the restricted data with the HRS pension calculator also yields initial pension balances as of 1992. Mean pension wealth in our estimation sample is \$93,300. Disaggre-

gating by health insurance type, those with retiree coverage have \$129,200, those with tied coverage have \$80,000, and those with none have \$18,700. With these starting values, we can then simulate pension wealth in our dynamic programming model with equation (15), using equation (21) to estimate pension accrual, and using equation (20) to estimate pension benefits. Using these equations, it is straightforward to track and record the pension balances of each simulated individual.

But even though it is straightforward to use equation (15) 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:

$$\begin{aligned} \widehat{pb}_t(PIA_t, HI_{t-1}, t) &= \sum_k \gamma_{0,k,t} 1\{HI_{t-1} = k\} + \gamma_3 PIA_t + \\ &\quad \gamma_{4,t} \max\{0, PIA_t - 9,999.6\} + \gamma_{5,t} \max\{0, PIA_t - 14,359.9\}, \end{aligned} \quad (24)$$

where PIA_t is the Social Security benefit the worker would get if he were drawing benefits at time t ; as shown in Appendix D below, PIA is a simple monotonic function of AIME. Using equations (18) and (24) yields imputed pension wealth, $\widehat{pw}_t = \Gamma_t \widehat{pb}_t$. The coefficients of this equation were estimated with regressions on simulated data generated by the model, with age effects captured by interacting the health insurance and PIA variables with a quadratic polynomial in age. Since these simulated data depend on the γ 's— \widehat{pw}_t affects the decision rules used in the simulations—the γ 's solve a fixed-point problem. Fortunately, estimates of the γ 's converge 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

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

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 (24). To correct for this, we increase non-pension wealth, A_{t+1} , by $s_{t+1}(1 - \tau_t)(\widehat{pw}_{t+1} - \widehat{pw}_{t+1})$. 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 that we reduce PIA if an individual draws benefits before age 65.) 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 (24) 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 D: Computation of AIME

We model several key aspects of Social Security benefits. First, 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.³² Third, AIME is capped. In 1998, the base year for the analysis, the maximum AIME level was \$68,400.

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

$$\begin{aligned} AIME_{t+1} &= (1 + g \times 1\{t \leq 60\})AIME_t \\ &+ \frac{1}{35} \max \{0, W_t H_t - \alpha_t(1 + g \times 1\{t \leq 60\})AIME_t\}, \end{aligned} \quad (25)$$

where the parameter α_t approximates the ratio of the lowest earnings year to $AIME_t$. 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 sequence $\{1\{\text{time-}t \text{ earnings do not increase } AIME_t\}\}_{t \geq 60}$ for each simulated wage history, and estimating α_t as the average of this indicator at each age. Linear interpolation yields α_{56} through α_{59} .

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

$$PIA_t = \begin{cases} 0.9 \times AIME_t & \text{if } AIME_t < \$5,724 \\ \$5,151.6 + 0.32 \times (AIME_t - 5,724) & \text{if } \$5,724 \leq AIME_t < \$34,500 \\ \$14,359.9 + 0.15 \times (AIME_t - 34,500) & \text{if } AIME_t \geq \$34,500 \end{cases} \quad (26)$$

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

³²After age 62, nominal benefits increase at the rate of inflation.

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%. 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 (26), this is easy to find.

Appendix E: 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 (14) 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 (14), having solved for V_T already. Continuing this backwards induction yields decision rules for times $T - 2, T - 3, \dots, 1$. The value function is directly computed at a finite number of points within a grid, $\{X_i\}_{i=1}^I$,³³ We use linear interpolation within the grid (i.e., we take a weighted average of the value functions of the surrounding gridpoints) 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, following the approach of Tauchen (1986);

³³In 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 states for the persistent component of medical expenses, 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, at 15,630 points before age 62 (when application is not an option) or at ages 70-71 (when we impose application), and at 7,680 points after age 71 (when we impose retirement as well).

using discretization rather than quadrature greatly reduces the number of times one has to interpolate when calculating $E_t(V(X_{t+1}))$. 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.³⁴ In particular, increasing the number of grid points seemed to have a small effect 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 (13), 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 the discretization of ζ_t and ω_t than when ζ_t and ω_t are drawn from Markov chains.

Finally, to reduce the computational burden, we assume that all workers apply for Social

³⁴The consumption grid has 100 points, and the hours grid is broken into 500-hour intervals. When this 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 $H_{t+1} = 0$. Once these values are found, we perform a quick, "second-pass" search in a neighborhood around them.

Security benefits by age 70, and retire by age 72: for $t \geq 70$, $B_t = 1$; and for $t \geq 72$, $H_t = 0$.

Appendix F: Moment Conditions, Estimation Mechanics, and the Asymptotic Distribution of Parameter Estimates

Following Gourinchas and Parker (2002) and French (2005), we estimate the parameters of the model in two steps. In the first step we estimate or calibrate parameters that can be cleanly identified without explicitly using our model. For example, we estimate mortality rates and health transitions from demographic data. As a matter of notation, we call this set of parameters χ . In the second step, we estimate the vector of “preference” parameters, $\theta = (\gamma_0, \gamma_1, \gamma_2, \beta_0, \beta_1, \beta_2, \nu, L, \phi_P, \phi_M, \theta_B, \kappa, C_{min}, \text{preference type prediction coefficients})$, using the method of simulated moments (MSM).

We assume that the “true” preference vector θ_0 lies in the interior of the compact set $\Theta \subset \mathbb{R}^{29}$. 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 $21T$ moment conditions. They are, for each age $t \in \{1, \dots, T\}$, two asset quantiles (forming $2T$ moment conditions), labor force participation rates conditional on asset quantile and health insurance type ($9T$), labor market exit rates for each health insurance type ($3T$), labor force participation rates conditional on the preference indicator described in the main text ($3T$), and labor force participation rates and mean hours worked conditional upon health status ($4T$).

Consider first the asset quantiles. As stated in the main text, let $j \in \{1, 2, \dots, J\}$ 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

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

In other words, the fraction of age- t 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. As is well known

(see, e.g., Manski, 1988, Powell, 1994 or Buchinsky, 1998; or the review in Chernozhukov and Hansen, 2002), the preceding equation can be rewritten as a moment condition. In particular, one can use the indicator function to rewrite the definition of the π_j -th conditional quantile as

$$E(1\{A_{it} \leq Q_{\pi_j}(A_{it}, t)\} | t) = \pi_j. \quad (27)$$

If the model is true then the data quantile in equation (27) can be replaced by the model quantile, and equation (27) can be rewritten as:

$$E(1\{A_{it} \leq g_{\pi_j}(t; \theta_0, \chi_0)\} - \pi_j | t) = 0, \quad j \in \{1, 2, \dots, J\}, \quad t \in \{1, \dots, T\}. \quad (28)$$

Since $J = 2$, equation (28) 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 (28) is a departure from the usual practice of minimizing a sum of weighted absolute errors in quantile estimation. The quantile restrictions just described, however, are part of a larger set of moment conditions, which means that we can no longer estimate θ by minimizing weighted absolute errors. Our approach to handling multiple quantiles is similar to the minimum distance framework used by Epple and Seig (1999).³⁵

The next set of moment conditions uses the quantile-conditional means of labor force participation. Let $\bar{P}_j(HI, t; \theta_0, \chi_0)$ 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, $\bar{P}_j(HI, t; \theta_0, \chi_0)$ should equal the conditional participation rates found in the data:

$$\bar{P}_j(HI, t; \theta_0, \chi_0) = E[P_{it} | HI, t, g_{\pi_{j-1}}(t; \theta_0, \chi_0) \leq A_{it} \leq g_{\pi_j}(t; \theta_0, \chi_0)], \quad (29)$$

³⁵Buchinsky (1998) shows that one could include the first-order conditions from multiple absolute value minimization problems in the moment set. However, his approach involves finding the gradient of $g_{\pi_j}(t; \theta, \chi)$ at each step of the minimization search.

with $\pi_0 = 0$ and $\pi_{J+1} = 1$. Using indicator function notation, we can convert this conditional moment equation into an unconditional one:

$$E([P_{it} - \bar{P}_j(HI, t; \theta_0, \chi_0)] \times 1\{HI_{it} = HI\} \times 1\{g_{\pi_{j-1}}(t; \theta_0, \chi_0) \leq A_{it} \leq g_{\pi_j}(t; \theta_0, \chi_0)\} | t) = 0, \quad (30)$$

for $j \in \{1, 2, \dots, J+1\}$, $HI \in \{none, retiree, tied\}$, $t \in \{1, \dots, T\}$. Note that $g_{\pi_0}(t) \equiv -\infty$ and $g_{\pi_{J+1}}(t) \equiv \infty$. With 2 quantiles (generating 3 quantile-conditional means) and 3 health insurance types, equation (29) generates $9T$ moment conditions.

The HRS asks workers about their willingness to work and/or their expectations about working in the future. We combine the answers to these questions into a time-invariant index, $pref \in \{high, low, out\}$. Because labor force participation differs significantly across values of $pref$, and because $pref$ significantly improves reduced-form predictions of employment, we interpret this index as a measure of otherwise unobserved preferences toward work. This leads to the following moment condition:

$$E(P_{it} - \bar{P}(pref, t; \theta_0, \chi_0) | pref_i = pref, t) = 0, \quad (31)$$

for $t \in \{1, \dots, T\}$, $pref \in \{0, 1, 2\}$. Equation (31) yields $3T$ moment conditions, which are converted into unconditional moment equations with indicator functions.

We also match exit rates for each health insurance category. Let $\bar{EX}(HI, t; \theta_0, \chi_0)$ denote the fraction of time- $t-1$ workers predicted to leave the labor market at time t . The associated moment condition is

$$E([1 - P_{it}] - \bar{EX}(HI, t; \theta_0, \chi_0 | HI_{i,60} = HI, P_{i,t-1} = 1, t) = 0, \quad (32)$$

$[1 - P_{it}] \neq Prob(P = 0 | P_{t-1} = 1)$ for $HI \in \{none, retiree, tied\}$, $t \in \{1, \dots, T\}$. Equation (32) generates $3T$ moment conditions, which are converted into unconditional moments as

well.³⁶

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

$$E(\ln H_{it} - \overline{\ln H}(HS, t; \theta_0, \chi_0) \mid P_{it} > 0, HS_{it} = HS, t) = 0, \quad (33)$$

$$E(P_{it} - \overline{P}(HS, t; \theta_0, \chi_0) \mid HS_{it} = HS, t) = 0, \quad (34)$$

for $t \in \{1, \dots, T\}$, $HS \in \{0, 1\}$. Equations (33) and (34), once again converted into unconditional form, yield $4T$ moment conditions, for a grand total of $21T$ moment conditions.

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 $21T$ -element vector of moment conditions that was described in the main text and immediately above, and let $\hat{\varphi}_I(\cdot)$ 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 $\varphi(\cdot)$ to be "missing", as in French and Jones (2004a). Letting $\widehat{\mathbf{W}}_I$ denote a $21T \times 21T$ weighting matrix, the MSM estimator $\hat{\theta}$ is given by

$$\arg \min_{\theta} \frac{I}{1 + \tau} \hat{\varphi}_I(\theta, \chi_0)' \widehat{\mathbf{W}}_I \hat{\varphi}_I(\theta, \chi_0), \quad (35)$$

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

To find the solution to equation (35), we proceed as follows:

1. We aggregate the sample data into life cycle profiles for hours, participation, exit rates

³⁶Because exit rates apply only to those working in the previous period, they normally do not contain the same information as participation rates. However, this is not the case for workers with *tied* coverage, as a worker stays in the *tied* category only as long as he continues to work. To remove this redundancy, the exit rates in equation (32) are conditioned on the individual's age-60 health insurance coverage, while the participation rates in equation (29) are conditioned on the individual's current coverage.

and assets.

2. 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. See Appendix G for details. We also use these data to estimate many of the parameters contained in the belief vector χ , although we calibrate other parameters. The initial distribution also includes preference type, assigned using our type prediction equation.
3. Using χ , we generate matrices of random health, wage, mortality and medical expense shocks. The matrices hold shocks for 40,000 simulated individuals.
4. We compute the decision rules for an initial guess of the parameter vector θ , using χ and the numerical methods described in Appendix E.
5. We simulate profiles for the decision variables. Each simulated individual receives a draw of preference type, assets, health, wages and medical expenses from the initial distribution, and is assigned one of the simulated sequences of health, wage and medical expense 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 the next period's states, and the simulated shocks pin the states down exactly.
6. We aggregate the simulated data into life cycle profiles.³⁷
7. We compute moment conditions, i.e., we find the distance between the simulated and true profiles, as described in equation (35).
8. We pick a new value of θ , update the simulated distribution of preference types, and repeat steps 4-7 until we find the value of θ that minimize that minimizes the distance

³⁷Because the moments we match include asset quantiles and asset-quantile-conditional participation rates, measurement error could affect our results. In earlier drafts of this paper (French and Jones, 2004b) we added measurement error to the simulated asset histories. Adding measurement error, however, had little effect on either the preference parameter estimates or policy experiments. For the moments we fit, the measurement error is largely averaged out.

between the true data and the simulated data. This vector of parameter values, $\hat{\theta}$, is our estimated value of θ_0 .³⁸

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;

$$\mathbf{D} = \left. \frac{\partial \varphi(\theta, \chi_0)}{\partial \theta'} \right|_{\theta=\theta_0} \quad (36)$$

is the $21T \times 29$ Jacobian matrix of the population moment vector; and $\mathbf{W} = \text{plim}_{T \rightarrow \infty} \{\widehat{\mathbf{W}}_I\}$.

Moreover, Newey (1985) shows that if the model is properly specified,

$$\frac{I}{1 + \tau} \hat{\varphi}_I(\hat{\theta}, \chi_0)' \mathbf{R}^{-1} \hat{\varphi}_I(\hat{\theta}, \chi_0) \rightsquigarrow \chi_{21T-29}^2,$$

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

$$\mathbf{R} = \mathbf{P}\mathbf{S}\mathbf{P},$$

$$\mathbf{P} = \mathbf{I} - \mathbf{D}(\mathbf{D}'\mathbf{W}\mathbf{D})^{-1}\mathbf{D}'\mathbf{W}.$$

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

³⁸Because the GMM criterion function is discontinuous, we search over the parameter space using a Simplex algorithm written by Honore and Kyriazidou. It usually takes around 2 weeks to estimate the model on a 20-node cluster, with each iteration (of steps 4-7) taking around 30 minutes.

$(1 + \tau)(\mathbf{D}'\mathbf{S}^{-1}\mathbf{D})^{-1}$, 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 $21T \times 21T$ estimated variance-covariance matrix of the sample data. That is, a typical diagonal element of $\widehat{\mathbf{S}}_I$ is the variance estimate $\frac{1}{T} \sum_{i=1}^I [1\{A_{it} \leq Q_{\pi_j}(A_{it}, t)\} - \pi_j]^2$, 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 (36) is straightforward to estimate for hours worked and participation conditional upon age and health status; we merely take numerical derivatives of $\hat{\varphi}_I(\cdot)$. However, in the case of the asset quantiles and labor force participation, discontinuities make the function $\hat{\varphi}_I(\cdot)$ 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 (28) 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

$$\mathbf{D}_{jt} = f(g_{\pi_j}(t; \theta_0, \chi_0)|t) \frac{\partial g_{\pi_j}(t; \theta_0, \chi_0)}{\partial \theta'}, \quad (37)$$

where \mathbf{D}_{jt} is the row of \mathbf{D} corresponding to the π_j -th quantile at year t . 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. We use a kernel estimator for GAUSS written by Ruud Koning.

To find the component of the matrix \mathbf{D} for the asset-conditional labor force participation rates, it is helpful to write equation (30) as

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

which implies that

$$\begin{aligned} \mathbf{D}_{jt} = & \left[-\Pr(g_{\pi_{j-1}}(t; \theta_0, \chi_0) \leq A_{it} \leq g_{\pi_j}(t; \theta_0, \chi_0)|HI, t) \frac{\partial \bar{P}_j(HI, t; \theta_0, \chi_0)}{\partial \theta'} \right. \\ & + [E(P_{it}|g_{\pi_j}(t; \theta_0, \chi_0), HI, t) - \bar{P}_j(HI, t; \theta_0, \chi_0)] f(g_{\pi_j}(t; \theta_0, \chi_0)|HI, t) \frac{\partial g_{\pi_j}(t; \theta_0, \chi_0)}{\partial \theta'} \\ & \left. - [E(P_{it}|g_{\pi_{j-1}}(t; \theta_0, \chi_0), HI, t) - \bar{P}_j(HI, t; \theta_0, \chi_0)] f(g_{\pi_{j-1}}(t; \theta_0, \chi_0)|HI, t) \frac{\partial g_{\pi_{j-1}}(t; \theta_0, \chi_0)}{\partial \theta'} \right] \\ & \times \Pr(HI_{t-1} = HI), \end{aligned} \quad (38)$$

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

Appendix G: Data and Initial Joint Distribution of the State Variables

Our data are drawn from the HRS, a sample of non-institutionalized individuals aged 51-61 in 1992. The HRS surveys individuals every two years; we have 7 waves of data covering the period 1992-2004. We use men in the analysis.

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.³⁹ We treat the health status for non-survey years as missing. Appendix H describes how we construct the health insurance indicator.

We use Social Security Administration earnings histories to construct AIME. Approximately 74% of our sample released their Social Security Number to the HRS, which allowed them to be linked to their Social Security earnings histories. For those who did not release their histories, we use the procedure described below to impute AIME as a function of assets, health status, health insurance type, labor force participation, and pension type.

The HRS collects pension data from both workers and employers. The HRS asks individuals about their earnings, tenure, contributions to defined contribution (DC) plans, and their employers. HRS researchers then ask employers about the pension plans they offer their employees. If the employer offers different plans to different employees, the employee is matched to the plan based on other factors, such as union status. Given tenure, earnings, DC

³⁹Bound et al. (2003) consider a more detailed measure of health status.

contributions, and pension plan descriptions, it is then possible to calculate pension wealth for each individual who reports the firm he works for. Following Scholz et al. (2006), we use firm reports of defined benefit (DB) pension wealth and individual reports of DC pension wealth if they exist. If not, we use firm-reported DC wealth and impute DB wealth as a function of wages, hours, tenure, health insurance type, whether the respondent also has a DC plan, health status, age, assets, industry and occupation. We discuss the imputation procedure below.

Workers are asked about two different jobs: (1) their current job if working or last job if not working; (2) the job preceding the one listed in part 1, if the individual worked at that job for over 5 years. Both of these jobs are included in our measure of pension wealth. Below we give descriptives for our estimation sample (born 1931-1935) and validation sample (born 1936-1941). 41% of our estimation sample [and 52% of our validation sample] are currently working and have a pension (of which 56% [57% for the validation sample] have firm-based pension details), 6% [5%] are not working, and had a pension on their last job (of which 62% [62%] have firm-based pension details), and 32% [32%] of all individuals had a pension on another job (of which 35% [29%] have firm-based pension details).

We dropped respondents for the following reasons. First, we drop all individuals who spent over 5 years working for an employer who did not contribute to Social Security. These individuals usually work for state governments. We drop these people because they often have very little in the way of Social Security wealth, but a great deal of pension wealth, a type of heterogeneity our model is not well suited to handle. Second, we drop respondents with missing information on health insurance, labor force participation, hours, and assets. When estimating labor force participation by asset quantile and health insurance for those born 1931-35 for the estimation sample [and 1936-41 for the validation sample], we begin with 19,547 [30,890] person year observations. We lose 3,139 [5,227] observations because of missing participation, 1,930 [2,162] observations who worked over 5 years for firms that did not contribute to Social Security, 150 [384] observations due to missing wave 1 participation, and 1,967 [2,883] observations due to missing health insurance data observations due to missing

asset data. In the end, from a potential sample of 19,547 [30,890] person-year observations for those between ages 51 and 69, we keep 11,773 [19,407] observations.

To generate the initial joint distribution of assets, wages, AIME, pensions, participation, health insurance, health status and medical expenses, we draw random vectors (i.e., random draws of individuals) from the empirical joint distribution of these variables for individuals aged 57-61 in 1992, or 1,701 observations. We drop observations with missing data on labor force participation, health status, insurance, assets, and age. We impute values for observations with missing wages, medical expenses, pension wealth, and AIME.

To impute these missing variables, we follow David et al. (1986) and Little (1988) and use the following predictive mean matching regression approach. First, we regress the variable of interest y_i (e.g., pension wealth) on a vector of observable variables x_i , $y_i = x_i\beta + \epsilon_i$. Second, we generate a predicted value $\hat{y}_i = x_i\hat{\beta}$ and generate a residual $\epsilon_i = y_i - \hat{y}_i$ for every member of the sample. Third, we split the predicted value \hat{y}_i into deciles. Fourth, we impute a value of y_i by taking a residual for a random individual j with a value of \hat{y}_j that is in the same decile of the distribution as is \hat{y}_i . Thus the imputed value of y_i is $\hat{y}_i + \epsilon_j$.

As David et al. (1986) point out, our imputation approach is equivalent to hot-decking when the “ x ” variables are discretized and include a full set of interactions. The advantages of the above approach over hot-decking are two-fold. First, many of the “ x ” variables are continuous, and it seems unwise to discretize them. Second, we have very few observations for some variables (such as pension wealth on past jobs), and hot-decking is very data-intensive. Only a small number of “ x ” variables are needed to generate a large number of hot-decking cells, as hot-decking uses a full set of interactions. We found that the interaction terms are relatively unimportant, but adding extra variables were very important for improving goodness of fit when imputing pension wealth.

If someone is not working (and thus does not report a wage), we use the wage on their last job as a proxy for their current wage if it exists, and otherwise impute the log wage as a function of assets, health, health insurance type, labor force participation, AIME, and quarters of covered work. We predict medical expenses using assets, health, health insurance

type, labor force participation, AIME, and quarters of covered earnings.

Lastly, we must infer the persistent component of the medical expense residual from medical expenses. Given an initial distribution of medical expenses, we construct ζ_t , the persistent medical expense component, by first finding the normalized log deviation ψ_t , as described in equations (7) and (10), and then applying standard projection formulae to impute ζ_t from ψ_t .

Appendix H: Measurement of Health Insurance Type

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 *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 *retiree* coverage. Nevertheless, estimated health insurance specific job exit rates are not very sensitive of our definition of health insurance, as we show below.

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 *retiree* 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 *retiree* 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 *retiree* 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, *retiree* 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 *retiree* coverage. It is possible, however, to estimate the probability of mismeasurement in these waves. Consider first the problem of distinguishing the *retiree* 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 $HI = retiree$ —so that we can drop the conditioning statement in the analysis below. Recall that many individuals who report retiree coverage in waves 1 and 2 likely have tied coverage. We are therefore interested in the misreporting probability $\Pr(HI = tied | HI^* = retiree, wv < 3, t < 65)$, where wv denotes HRS wave and t denotes age. To find this quantity, note first that by the law of total probability:

$$\begin{aligned} \Pr(HI = tied | wv < 3, t < 65) = \\ \Pr(HI = tied | HI^* = tied, wv < 3, t < 65) \times \Pr(HI^* = tied | wv < 3, t < 65) + \\ \Pr(HI = tied | HI^* = retiree, wv < 3, t < 65) \times \Pr(HI^* = retiree | wv < 3, t < 65). \end{aligned} \quad (39)$$

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:

$$\begin{aligned}\Pr(HI = tied|wv < 3, t < 65) &= \Pr(HI = tied|wv \geq 3, t < 65) \\ &= \Pr(HI^* = tied|wv \geq 3, t < 65).\end{aligned}$$

Inserting these assumptions into equation (39) and rearranging yields the mismeasurement probability:

$$\begin{aligned}\Pr(HI = tied|HI^* = retiree, wv < 3, t < 65) \\ &= \frac{\Pr(HI^* = tied|wv \geq 3, t < 65) - \Pr(HI^* = tied|wv < 3, t < 65)}{\Pr(HI^* = retiree|wv < 3, t < 65)} \\ &= \frac{\Pr(HI^* = retiree|wv < 3, t < 65) - \Pr(HI^* = retiree|wv \geq 3, t < 65)}{\Pr(HI^* = retiree|wv < 3, t < 65)}.\end{aligned}\quad (40)$$

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 *retiree* insurance is the same as for individuals in waves 1 and 2 who are younger than 65. We can then use equation (40) 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 *retiree* 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 *retiree* 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 9 shows how our treatment of these measurement problems affects measured job exit rates. The top two graphs in Figure 9 do not adjust for the measurement error problems described immediately above. The bottom two graphs account for the measurement error problems, using the approach described by equation 40. 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 9 shows that the profiles are not very sensitive to these changes. Those with *retiree* 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.

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.⁴⁰ 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.⁴¹

⁴⁰For 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.

⁴¹Note that this particular assumption implies that 100% of those eligible for *COBRA* take up coverage. In practice, only about $\frac{2}{3}$ of those eligible take up coverage (Gruber and Madrian, 1996). In order to determine whether our failure to model the *COBRA* decision is important, we shut down the *COBRA* option (imposed a 0% take-up rate) and re-ran the model. Eliminating *COBRA* had virtually no effect on labor supply.

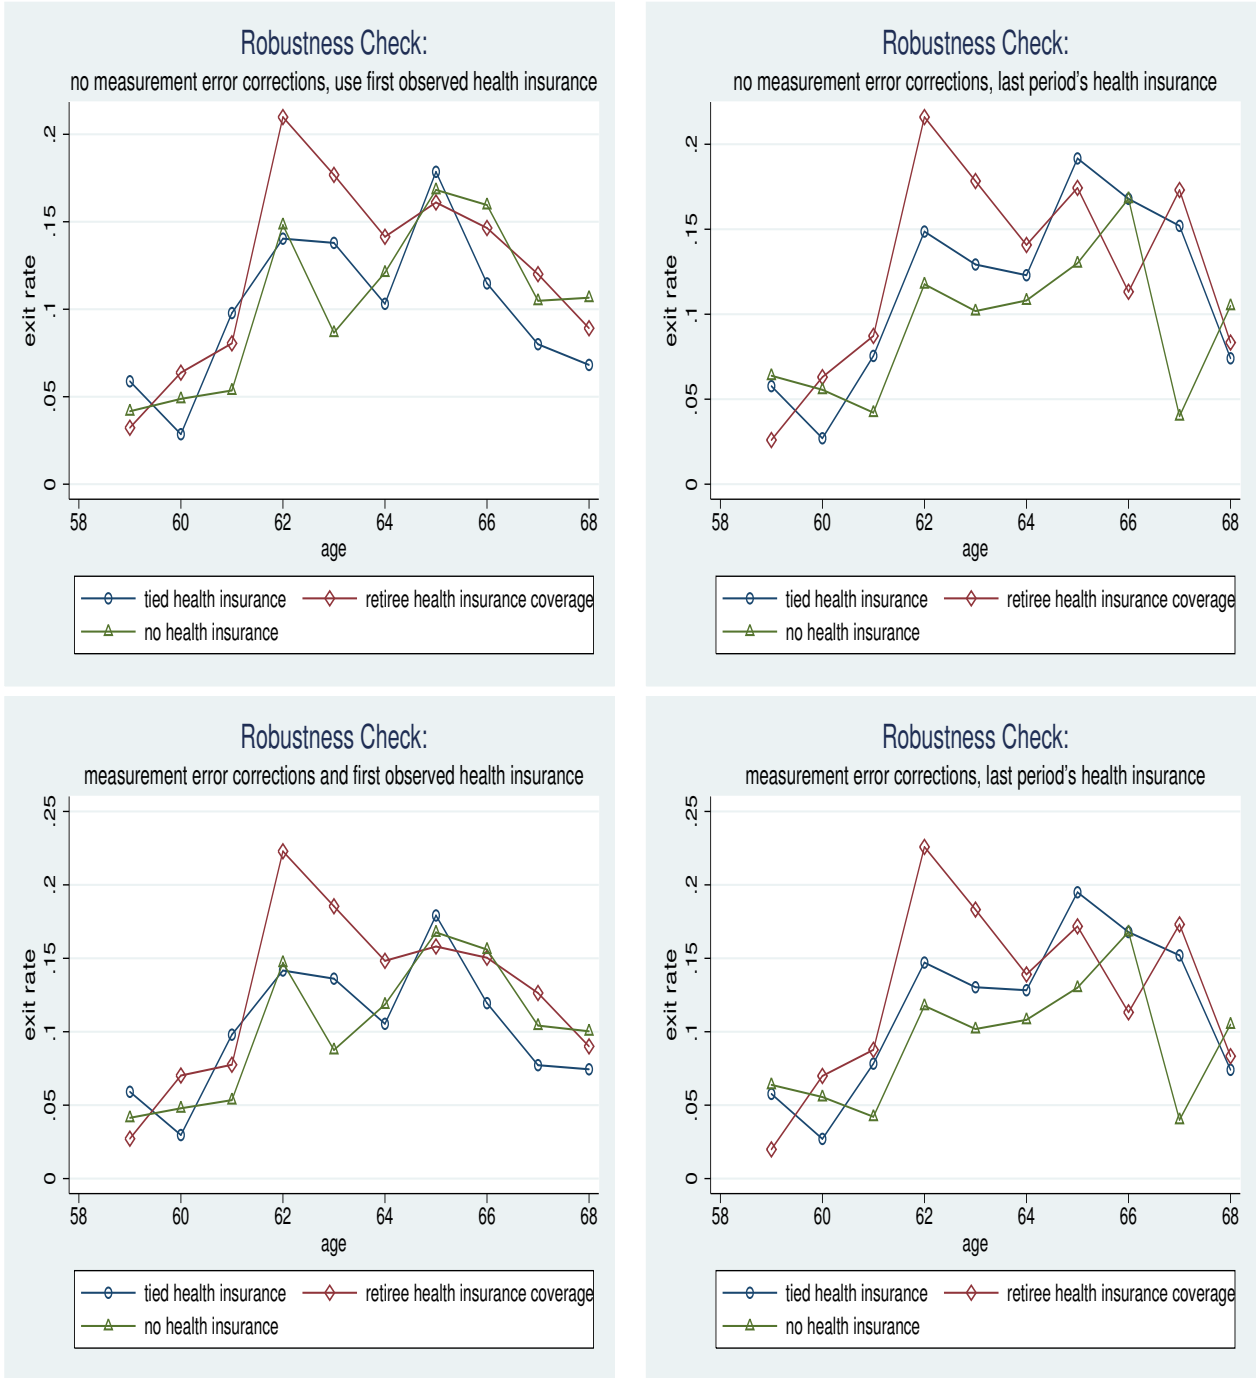


Figure 9: JOB EXIT RATES USING DIFFERENT MEASURES OF HEALTH INSURANCE TYPE

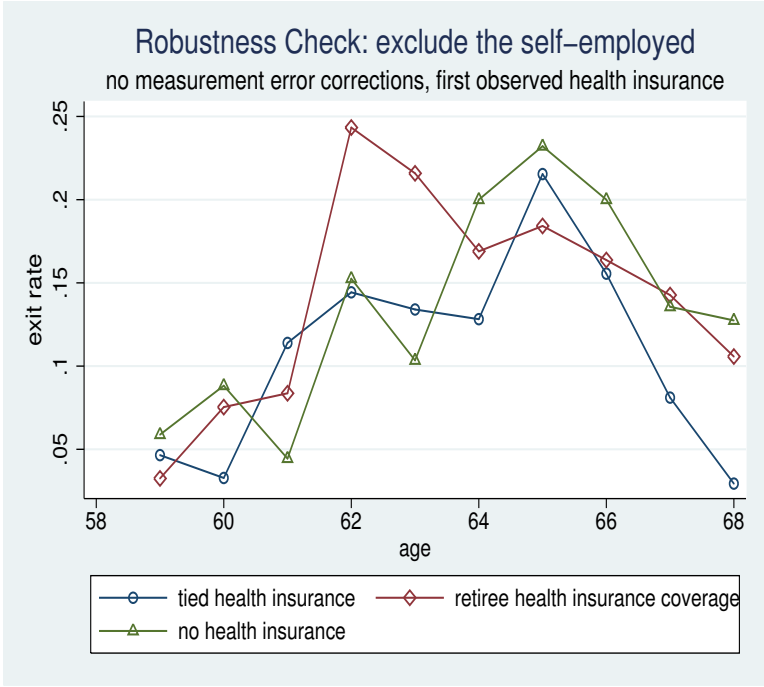



Figure 10: THE EFFECT OF DROPPING THE SELF-EMPLOYED ON JOB EXIT RATES

Another measurement issue is the treatment of the self-employed. Figure 10 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 *retiree* 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.

Because agents in our model are forward-looking, we need to know the health-insurance-conditional process for medical expenses facing the very old. The data we use to estimate medical expenses for those over age 70 comes from the Assets and Health Dynamics of the Oldest Old survey. French and Jones (2004a) discuss some of the details of the survey, as well as some of our coding decisions. The main problem with the AHEAD is that there is no question asked of respondents about whether they would lose their health insurance if they left their job, so it is not straightforward to distinguish those who have *retiree* coverage from those with *tied* coverage. In order to distinguish these two groups, we do the following. If the individual exits the labor market during our sample, and has employer-provided health insurance at least one full year after exiting the labor market, we assume that individual has *retiree* coverage. All individuals who have employer-provided coverage when first observed, but do not meet this criteria for having *retiree* coverage, are assumed to have *tied* coverage.

Appendix I: The Medical Expense Model

Recall from equation (7) that health status, health insurance type, labor force participation and age affect medical expenses through the mean shifter $m(\cdot)$ and the variance shifter $\sigma(\cdot)$. Health status enters $m(\cdot)$ and $\sigma(\cdot)$ 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 medical expenses are given by

$$m(HS_t, HI_t, t, P_t) = \gamma_0 HS_t + \gamma_1 t + \sum_{h \in HI} \sum_{P \in \{0,1\}} \sum_{a \in \{t < 65, t \geq 65\}} \gamma_{h,P,a}.$$

The one restriction we impose is that $\gamma_{none,0,a} = \gamma_{none,1,a}$ for both values of a , i.e., participation does not affect health care costs if the individual does not have insurance. This implies that there are 10 $\gamma_{h,P,a}$ parameters, for a total of 12 parameters apiece in the $m(\cdot)$ and the $\sigma(\cdot)$ functions.

To estimate this model, we group the data into 10-year-age (55-64, 65-74, 75-84) \times health status \times health insurance \times 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 (2004a) shown in Table 2 of this paper 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 (2004a), who provide a detailed description.

Appendix J: The Preference Index

We construct the preference index for each member of the sample using the wave 1 variables V3319, V5009, V9063. All three variables are self-reported responses to questions about preferences for leisure and work. In V3319 respondents were asked if they agreed with the statement (if they were working): “Even if I didn’t need the money, I would probably keep on working.” In V5009 they were asked: “When you think about the time when you [and your (husband/wife/partner)] will (completely) retire, are you looking forward to it, are you

uneasy about it, or what? In V9063 they were asked (if they were working): “On a scale where 0 equals dislike a great deal, 10 equals enjoy a great deal, and 5 equals neither like nor dislike, how much do you enjoy your job?”

Because it is computationally intensive to estimate the parameters of the type probability equations in our method of simulated moments approach, we combine these three variables into a single index that is simpler to use. To construct this index, we regress labor force participation on current state variables (age, wages, assets, health, etc.), squares and interactions of these terms, the wave 1 variables V3319, V5009, V9063, and indicators for whether these variables are missing. We then partition the $x\hat{\beta}$ matrix from this regression into: $x_1\hat{\beta}_1$, where the x_1 matrix includes V3319, V5009, V9063, and indicators for these variables being missing; and $x_2\hat{\beta}_2$, where the x_2 matrix includes all other variables. Our preference index is $x_1\hat{\beta}_1$.

Individuals who were not working in 1992 were not asked any of the preference questions, and are not included in the construction of our index. Because there is no variation in participation in 1992, we estimate the regression models with participation data from 1998-2004.

Finally, we discretize the index into three values: *out*, for those not employed in 1992; *low*, for workers with an index in the bottom half of the distribution; and *high* for the remainder.

Appendix K: Additional Parameter Estimates

We assume that the probability of belonging to a particular type follows a multinomial logit function. Table 13 shows the coefficients of the preference type prediction equation.

Table 14 shows the parameter estimates for the robustness checks. In the no-saving case, shown in column (2), 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 values of 0.83, 1.12, and 0.97 (for types 0, 1 and 2, respectively). Similarly, we fix θ_B to zero. Since the asset distribution is 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.

	Preference Type 1		Preference Type 2	
	Parameters	Std. Errors	Parameters	Std. Errors
	(1)	(2)	(3)	(4)
Preference Index = <i>out</i>	-4.51	0.69	-5.22	18.58
Preference Index = <i>low</i>	3.97	7.25	0.62	4.94
Preference Index = <i>high</i>	-0.07	0.46	5.55	0.82
No HI Coverage	1.69	0.67	-4.17	1.28
Retiree Coverage	0.23	0.33	-2.48	0.60
Initial Wages [†]	2.64	0.46	-0.85	0.45
Assets/Wages [†]	-0.44	0.45	-0.52	0.23
Assets [†] × (No HI Coverage)	0.20	0.30	1.85	0.82

[†]Variables expressed as fraction of average

Table 13: PREFERENCE TYPE PREDICTION COEFFICIENTS

Appendix L: Illiquid Housing

Although allowing for no savings seems extreme, it has often been argued (e.g., Rust and Phelan, 1997, Gustman and Steinmeier, 2005) 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.⁴² The third column of Table 14 contains the revised parameter estimates. The most notable changes are: (1) the coefficient of relative risk aversion, ν , drops from 7.5 to 6.5; (2) the type-1 value of β , the discount factor, drops from 1.115 to 0.858; (3) the consumption floor, c_{min} , increases from \$4,100 to \$6,300. All three changes—lower risk aversion, lower patience and more government protection—help the model fit the bottom third of liquid asset holdings, which averages less than \$5,000.

Column (3) of Table 9 shows participation when housing assets are illiquid. The most notable result is that simulated participation drops markedly at age 62. Several authors (Kahn, 1988, Rust and Phelan, 1997, and Gustman and Steinmeier, 2005) have argued that, because they cannot borrow against their Social Security benefits, many workers that would other-

⁴²A 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.

Parameter and Definition	Baseline (1)	No Saving (2)	Homogeneous Preferences (3)	Illiquid Housing (4)
γ : consumption weight				
Type 0	0.438 (0.080)	0.239 (9.760)	NA	0.403 (0.113)
Type 1	0.620 (0.011)	0.548 (0.006)	0.700 (0.007)	0.695 (0.006)
Type 2	0.907 (0.028)	0.928 (0.031)	NA	0.911 (0.026)
β : time discount factor				
Type 0	0.828 (0.072)	0.828 (NA)	NA	0.821 (0.074)
Type 1	1.115 (0.016)	1.115 (NA)	0.971 (0.010)	0.858 (0.012)
Type 2	0.971 (0.077)	0.971 (NA)	NA	0.957 (0.067)
ν : coefficient of relative risk aversion, utility	7.49 (0.421)	6.61 (0.166)	3.93 (0.202)	6.46 (0.196)
L : leisure endowment, in hours	3,863 (51.9)	4,052 (26.2)	4,101 (34.3)	3,960 (47.6)
ϕ_P : fixed cost of work, in hours	835 (27.4)	1,146 (30.5)	1,196 (16.3)	904 (20.4)
ϕ_{HS} : hours of leisure lost, bad health	445 (38.8)	432 (29.6)	432 (28.1)	412 (12.9)
θ_B : bequest weight [†]	0.0320 (0.0009)	0.00 (NA)	0.0241 (0.0005)	0.0338 (0.0022)
κ : bequest shifter, in thousands	449 (31.7)	0.00 (NA)	509 (12.6)	460 (32.3)
c_{min} : consumption floor	4,118 (159.5)	3,517 (159.1)	5,386 (141.1)	6,275 (215.7)
χ^2 statistic	1,677	1,211	1,009	3,081
Degrees of freedom	181	96	171	181
<p>Diagonal weighting matrix used in calculations. See Appendix F for details. Standard errors in parentheses. [†]Parameter expressed as marginal propensity to consume out of final-period wealth.</p>				

Table 14: ROBUSTNESS CHECKS

wise retire earlier cannot fund their retirement before age 62. Making housing illiquid, along with the large decrease in the estimated value of β , strengthens this effect. The underlying asset-conditional profiles reveal that the participation drop is most pronounced for simulated workers in the bottom 1/3rd of the asset distribution. This contrasts with the data, where the age-62 exit rates vary across the asset quantiles to a much smaller extent.

We find that in this framework delaying the Medicare eligibility age has a bigger effect than delaying the Social Security normal retirement age. Shifting forward the Medicare eligibility age to 67 increases total years in the labor force by 0.11 years (versus the 0.07 years for the baseline specification that we presented in Table 7).