

Are Americans Saving “Optimally” for Retirement?

John Karl Scholz

University of Wisconsin–Madison and National Bureau of Economic Research

Ananth Seshadri

University of Wisconsin–Madison

Surachai Khitatrakun

Urban Institute

We solve each household’s optimal saving decisions using a life cycle model that incorporates uncertain lifetimes, uninsurable earnings and medical expenses, progressive taxation, government transfers, and pension and social security benefits. With optimal decision rules, we compare, household by household, wealth predictions from the life cycle model using a nationally representative sample. We find, making use of household-specific earnings histories, that the model accounts for more than 80 percent of the 1992 cross-sectional variation in wealth. Fewer than 20 percent of households have less wealth than their optimal targets, and the wealth deficit of those who are under-saving is generally small.

We thank colleagues at the University of Michigan for developing the Health and Retirement Survey. We also are grateful to the National Institute of Aging, the Center for the Demography of Health and Aging at the University of Wisconsin–Madison, and the Russell Sage Foundation for financial support; Sheng-Kai Chang for sharing code that assisted in the earnings estimation; Amy Finkelstein for information on nursing home stays and expenditures; and numerous colleagues at Wisconsin, many seminar participants, Doug Bernheim, Bill Gale, Jagadeesh Gokhale, Jane Gravelle, Erik Hurst, Yuichi Kitamura, David Laibson, Anna Lusardi, Derek Neal, Jim Poterba, John Sabelhaus, Emmanuel Saez, Andrew Samwick, Jon Skinner, and the referees for helpful comments.

[*Journal of Political Economy*, 2006, vol. 114, no. 4]
© 2006 by The University of Chicago. All rights reserved. 0022-3808/2006/11404-0001\$10.00

There is considerable skepticism in public policy discussions and in the financial press that Americans are preparing adequately for retirement. A quotation from the *Wall Street Journal* captures a popular view:

A long time ago, New England was known for its thrifty Yankees. But that was before the baby boomers came along. These days, many New Englanders in their 30s and 40s, and indeed their counterparts all over America, have a different style: they are spending heavily and have sunk knee-deep in debt. . . . A recent study sponsored by Merrill Lynch & Co. showed that the average middle-aged American had about \$2,600 in net financial assets. Another survey by the financial-services giant showed that boomers earning \$100,000 will need \$653,000 in today's dollars by age 65 to retire in comfort—but were saving only 31 percent of the amount needed. In other words, the saving rate will have to triple. Experts say the failure to build a nest egg will come to haunt the baby boomers, forcing them to drastically lower standards of living in their later years or to work for longer, perhaps into their 70s. (Wysocki 1995, A1)

Assessing the adequacy or optimality of wealth accumulation is difficult since it requires some standard against which to measure observed behavior. Several authors use augmented life cycle models for this standard, simulating the expected distribution of wealth for representative household types (see, e.g., Hubbard, Skinner, and Zeldes 1995; Engen, Gale, and Uccello 1999).¹ While augmented life cycle models provide natural benchmarks, these researchers do not fully assess the adequacy (let alone the optimality) of wealth accumulation. They derive optimal *distributions* of wealth (or wealth-to-income ratios). But given underlying model parameters, expectations, and earnings realizations, the models have household-specific implications for optimal wealth accumulation. These household-specific implications have not been studied.

We examine the degree to which households are optimally preparing for retirement by constructing a stochastic life cycle model that captures the key features of a household's consumption decisions. Our model incorporates many behavioral features shown by prior work to affect

¹ Kotlikoff, Spivak, and Summers (1982), Moore and Mitchell (1998), and Gustman and Steinmeier (1999) examine saving adequacy by comparing data to financial planning rules of thumb. But a rule of thumb cannot describe optimal behavior for households with widely different patterns of earnings realizations, even if preferences are homogeneous. Banks, Blundell, and Tanner (1998) and Bernheim, Skinner, and Weinberg (2001) make inferences about adequacy from consumption changes around retirement. But, for the reasons given in Hurd and Rohwedder (2003), Blau (2004), Haider and Stevens (2004), and Aguiar and Hurst (2005), it is difficult to make inferences about adequacy or optimality from patterns of consumption changes around retirement.

consumption, including precautionary savings and buffer stock behavior (Deaton 1991; Aiyagari 1994; Carroll 1997). It incorporates asset-tested public transfers (Hubbard et al. 1995), though in our model benefits vary over time and by household size. Our model incorporates end-of-life uncertainty and medical shocks (Palumbo 1999). We also incorporate a stylized, time-varying progressive income tax that reflects the evolution of average effective federal income tax rates over the period spanned by our data. Households in the model form realistic expectations about earnings; about social security benefits, which depend on lifetime earnings; and about pension benefits, which depend on earnings in the final year of work. We incorporate detailed data from the Health and Retirement Study (HRS) on family structure and age of retirement (treating both as exogenous and known from the beginning of working life) in calculating optimal life cycle consumption profiles.

Our approach has other distinctive features. Most important, we calculate household-specific optimal wealth targets using data from the HRS. A crucial input to our behavioral model is 41 years of information on earnings realizations drawn from restricted-access social security earnings records. The timing of earnings shocks can cause optimal wealth to vary substantially, even for households with identical preferences, demographic characteristics, and lifetime income. Hence, it is essential for life cycle models of wealth accumulation to incorporate earnings realizations, at least to the extent model implications are compared to actual behavior.

We find that over 80 percent of HRS households have accumulated more wealth than their optimal targets. These targets indicate the amounts of private saving households should have acquired at the time we observe them in the data, given their life cycle planning problem and social security and defined-benefit pension expectations and realizations. For those not meeting their targets, the magnitudes of the deficits are typically small. In addition, the cross-sectional distribution of wealth in 1992 closely matches the predictions of our life cycle model.

I. The Health and Retirement Study

The HRS is a national panel study with an initial sample (in 1992) of 12,652 persons in 7,702 households.² It oversamples blacks, Hispanics, and residents of Florida. The baseline 1992 study consisted of in-home, face-to-face interviews of the 1931–41 birth cohort and their spouses, if they were married. Follow-up interviews were given by telephone in

² An overview of the HRS is given in a supplementary issue of the *Journal of Human Resources* (1995, vol. 30). There, 22 authors discuss and assess the data quality of many dimensions of the initial wave of the HRS.

1994, 1996, 1998, 2000, 2002, and 2004. For the analyses in this paper we exclude 379 married households in which one spouse did not participate in the 1992 HRS, 93 households that failed to have at least one year of full-time work, and 908 households in which the highest earner began working full-time prior to 1951.³ Our resulting sample has 10,523 respondents in 6,322 households.

The survey covers a wide range of topics, including batteries of questions on health and cognitive conditions, retirement plans, subjective assessments of mortality probabilities and the quality of retirement preparation, family structure, employment status and job history, demographic characteristics, housing, income and net worth, and pension details.

A. *Wealth Measures in the HRS*

Households typically maintain living standards in retirement by drawing on their own (private) savings, employer-provided pensions, and social security wealth. To study the degree to which households optimally accumulate wealth, therefore, we need accurate measures of these wealth components.

Net worth (private savings) is a comprehensive measure that includes housing assets less liabilities, business assets less liabilities, checking and saving accounts, stocks, bonds, mutual funds, retirement accounts including defined-contribution pensions, certificates of deposit, the cash value of whole life insurance, and other assets, less credit card debt and other liabilities. It excludes defined-benefit pension wealth, social security wealth, and future earnings.⁴ The concept of wealth is similar (and in many cases identical) to those used in other studies of wealth and saving adequacy.

We use the Pension Present Value Database that Bob Peticolas and Tom Steinmeier made available through the HRS to calculate the value

³ We drop the first group because we do not have information on spousal, and hence household, income. We drop the second group because we do not have information on transfer payments in years prior to the HRS survey, and therefore we cannot model the lifetime budget constraint. We drop households in which the highest earner started working before 1951 for computational reasons. Our procedures to impute missing and top-coded data are more complicated when initial values of the earnings process are missing. Details for the earnings imputations are given in online App. B.

⁴ We account rigorously for defined-benefit pensions, social security, and future earnings in the household decision problem. We do not model the purchase and service flows that result from consumer durables since we do not have data on durables. We do not include term life policies in net worth because we do not have information on premium payments; more important, a comparison of data from the HRS and Assets and Health Dynamics of the Oldest Old (AHEAD) suggests that a substantial fraction of term life policies are dropped in retirement. Those that are not dropped have a median value of less than \$7,000.

of defined-benefit pensions and, as described below, estimate the household's expectations of future pension benefits.⁵ The program makes present-value calculations of pensions for HRS respondents in 1992 for nine different scenarios, corresponding to the Social Security Administration's low, intermediate, and high long-term projections for interest rates, wage growth rates, and inflation rates. We use the intermediate values (6.3 percent for interest rates, 5 percent for wage growth, and 4 percent for inflation) when calculating defined-benefit pension wealth.

B. HRS Earnings Data

Restricted-access social security earnings data provide a direct measure of earnings realizations and lifetime income, and, as described below, they are used to estimate a household's expectations of future earnings. They also allow us to simulate accurately social security benefits for the respondent and spouse or for the couple, if the benefit would be higher.

Two issues arise in using earnings information. First, social security earnings records are not available for 23 percent of the respondents included in the analysis. Second, the social security earnings records are top-coded—respondents earned more than the social security taxable wage caps—for 13 percent of earnings observations between 1951 and 1979. From 1980 through 1991 censoring is much less of an issue, because we have access to W-2 earnings records, which are very rarely censored.

We impute earnings histories for those individuals with missing or top-coded earnings records assuming the individual log earnings process

$$\begin{aligned} y_{i,0}^* &= \mathbf{x}'_{i,0} \beta_0 + \epsilon_{i,0}, \\ y_{i,t}^* &= \rho y_{i,t-1}^* + \mathbf{x}'_{i,t} \beta + \epsilon_{i,t}, \quad t \in \{1, 2, \dots, T\}, \\ \epsilon_{i,t} &= \alpha_i + u_{i,t} \end{aligned} \tag{1}$$

where $y_{i,t}^*$ is the log of latent earnings of the individual i at time t in 1992 dollars, $\mathbf{x}_{i,t}$ is the vector of i 's characteristics at time t , and the error term $\epsilon_{i,t}$ includes an individual-specific component α_i , which is constant over time, and an unanticipated white-noise component, $u_{i,t}$. We employ

⁵ See the Pension Present Value Database at <http://hrsonline.isr.umich.edu/data/avail.html>. The programs use detailed plan descriptions along with information on employee earnings. We use self-reported defined-benefit pension information for households not included in the Peticolas and Steinmeier file. The assumptions used in the program to calculate the value of defined-contribution pensions—particularly the assumption that contributions were a constant fraction of income during years worked with a given employer—are likely inappropriate. Consequently, we follow others in the literature (e.g., Engen et al. 1999, 159) and use self-reported information to calculate defined-contribution pension wealth.

TABLE 1
DESCRIPTIVE STATISTICS FOR THE HEALTH AND RETIREMENT STUDY (Dollar Amounts in
1992 Dollars)

Variable	Mean	Median	Standard Deviation
1991 earnings	\$35,958	\$28,976	\$39,368
Present discounted value of lifetime earnings	\$1,718,932	\$1,541,555	\$1,207,561
Defined-benefit pension wealth	\$106,041	\$17,327	\$191,407
Social security wealth	\$107,577	\$97,726	\$65,397
Net worth	\$225,928	\$102,600	\$464,314
Mean age (years)	55.7		4.7
Mean education (years)	12.7		3.4
Fraction male	.70		.46
Fraction black	.11		.31
Fraction Hispanic	.06		.25
Fraction couple	.66		.48
No high school diploma	.22		.41
High school diploma	.55		.50
College graduate	.12		.33
Postcollege education	.10		.30
Fraction self-employed	.15		.35
Fraction partially or fully retired	.29		.45

SOURCE.—Authors' calculations from the 1992 HRS. The table is weighted by the 1992 HRS household analysis weights.

random-effect assumptions with homoskedastic errors to estimate equation (1).

We estimate the model separately for four groups: men without a college degree, men with a college degree, women without a college degree, and women with a college degree. In the online Appendix B we present details of the empirical earnings model and coefficient estimates from that model, and we describe our Gibbs sampling procedure that we use to impute earnings for individuals who refuse to release or who have top-coded social security earnings histories.⁶ Our approach is appealing in that it uses information from the entire sequence of individual earnings, including uncensored W-2 data from 1980–91, to impute missing and top-coded earnings.

Table 1 provides descriptive statistics for the HRS sample. Mean (median) earnings in 1991 of HRS households are \$35,958 (\$28,976), though note that 29 percent of the sample was partially or fully retired when interviewed in the 1992 HRS. The mean (median) present discounted value of lifetime household earnings is \$1,718,932

⁶ We repeated our central empirical analyses dropping individuals who refused to release their social security records and generated results nearly identical to those reported in the paper. Brief details are given in the sensitivity analysis, Sec. IV E.

(\$1,541,555).⁷ Retirement consumption will be financed out of defined-benefit pension wealth (mean is \$106,041, median is \$17,327), social security wealth (mean is \$107,577, median is \$97,726), and nonpension net worth (mean is \$225,928, median is \$102,600). The mean age of the household head is 55.7.⁸

Social security replacement rates are defined as equaling annual social security benefits divided by the average of the final five years of earned income (prior to retirement), multiplied by 100. The median for our sample of married couples is 37.6 percent. Those with less than a high school diploma have a median rate of 41.7 percent. Those with a high school diploma or some college have a median rate of 38.7 percent. College graduates have a median rate of 31.1 percent, and those with more than a college degree have a median rate of 28.2 percent. Grad (1990) reports that median replacement rates for newly retired couples were between 49 and 62 percent in 1982.⁹ Because we use social security earnings records and a close approximation to the social security benefit rules, our measure compared to those in Grad's article shows how replacement rates changed over time.

Figure 1, which shows the median levels of defined-benefit pension wealth, social security wealth, and net worth (excluding defined-benefit pensions) in each lifetime earnings decile, highlights the reason we account rigorously for social security in our model. Social security exceeds the combined value of pension and nonpension net worth in the bottom three deciles of the lifetime earnings distribution. Private net worth exceeds the value of social security only in the top three deciles of the lifetime earnings distribution. The metaphor of the "three-legged stool," in which retirement income security is supported by the three legs of social security, employer-provided pensions, and private wealth accumulation, appears to apply only to households in the top 70 percent of the lifetime earnings distribution because low-income workers typically do not have employer-provided pensions.

⁷ When calculating present discounted values of earnings and social security wealth, we discount the constant-dollar sum of earnings (social security or pensions) by a real interest rate measure (prior to 1992, we use the difference between the three-month Treasury bill rate and the year-to-year change in the Consumer Price Index for Urban Wage Earners and Clerical Workers; for 1992 and after we use 4 percent). For the defined-benefit pension wealth, we assume that the real interest rate is 2.21 percent, consistent with the 6.3 percent interest rates and 4 percent inflation assumed under the intermediate scenarios of the Pension Present Value Database.

⁸ The head of the household is defined throughout the paper as the person in the household with the largest share of lifetime earnings. When we refer to the age or retirement date of the household, we are referring to the age or retirement date of the household head.

⁹ Engen et al. (1999) and figures from the Social Security Administration suggest that replacement rates likely fell substantially since 1982.

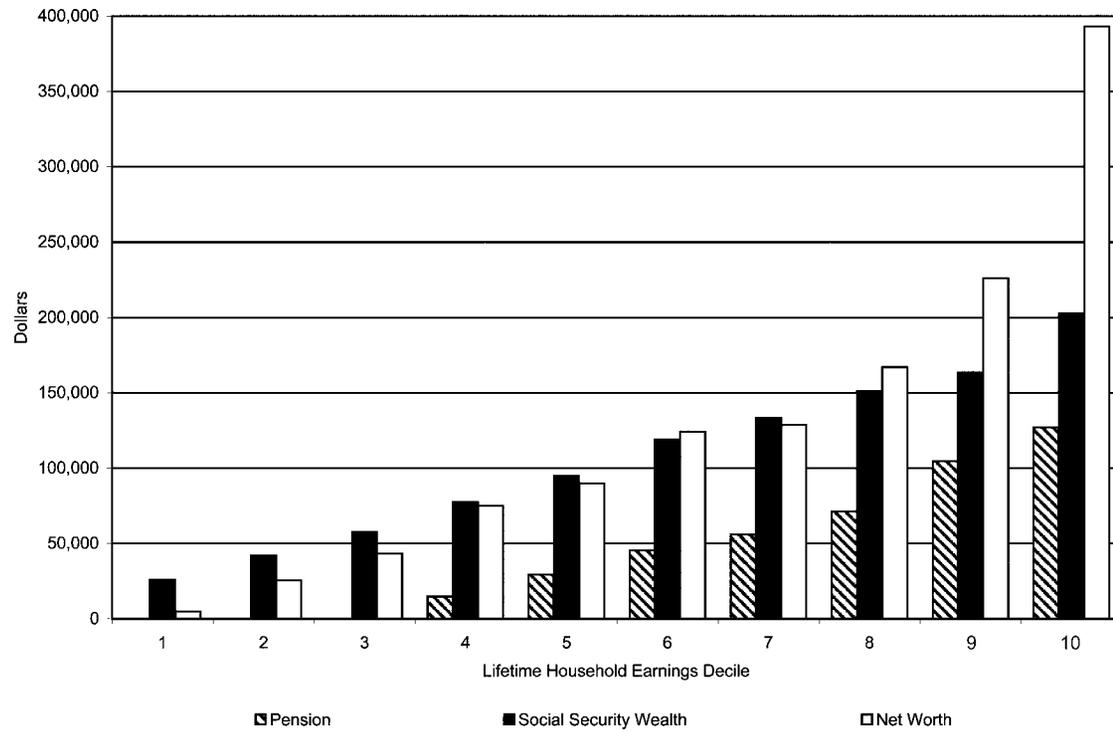


FIG. 1.—Median defined-benefit pension wealth, social security wealth, and net worth (excluding defined-benefit pensions) by lifetime earnings decile (1992 dollars).

II. A Model of Optimal Wealth Accumulation

We solve a simple life cycle model, augmented to incorporate uncertain lifetimes, uninsurable earnings, uninsurable medical expenses, and borrowing constraints. The unit of analysis is a household, which can be married or single.¹⁰ Individuals within a household live to a maximum age D . Between ages 0 and $S - 1$, individuals are children and make no consumption decisions. Adults start working at age S , have exogenous labor supply, and give birth to as many as n children at ages B_1, B_2, \dots, B_n . Earnings depend on age (which affects work experience) and a random shock that can be correlated across time. Each period, adults decide how much to consume and how much to save for the future.

Households retire exogenously at the end of age R and face a probability of death in each remaining year of life. In retirement, they start receiving health shocks that can be correlated across ages. They receive income from social security, defined-benefit plans (if covered), and assets. Social security benefits depend on total earnings during the preretirement period. Defined-benefit pension receipts are a function of the household's earnings in the period prior to retirement.

A. A Household's Maximization Problem

A household derives utility $U(c)$ from period-by-period consumption in equivalent units, where n_j adjusts consumption for the number of adults A_j and children K_j in the household at age j . Let c_j and a_j represent consumption and assets at age j . With probability p_j , the household survives into the next period, so the household survives until age j with probability $\prod_{k=S}^{j-1} p_k$, where $\prod_{k=S}^{j-1} p_k = 1$ if $j - 1 < R$. At age D , $p_D = 0$. The discount factor on future utilities is β . Expected lifetime utility is then

$$E \left[\sum_{j=S}^D \beta^{j-S} n_j U \left(\frac{c_j}{n_j} \right) \right].$$

The expectation operator E denotes the expectation over future earnings uncertainty, uncertainty in health expenditures, and uncertainty over the life span.

¹⁰ We do not model marriage or divorce. The head of a married household in 1992 is modeled as making the household's life cycle consumption decisions jointly with his or her partner throughout their working lives. The head or partner becomes single only if a spouse dies. Similarly, single households in 1992 are modeled as making their life cycle consumption decisions as if they were single throughout their working lives. They are assumed to remain single until death.

Consumption and assets are chosen to maximize expected utility subject to the constraints¹¹

$$\begin{aligned}
 y_j &= e_j + ra_j + T(e_j, a_j, n_j), \quad j \in \{S, \dots, R\}, \\
 y_j &= SS\left(\sum_{j=S}^R e_j\right) + DB(e_R) + ra_j \\
 &\quad + T_R\left(e_R, \sum_{j=S}^R e_j, a_j, n_j, m_j\right), \quad j \in \{R+1, \dots, D\}, \\
 c_j + a_{j+1} &= y_j + a_j - \tau(e_j + ra_j), \quad j \in \{S, \dots, R\}, \\
 c_j + a_{j+1} + m_j &= y_j + a_j - \tau\left(\Psi\left(SS\sum_{j=S}^R e_j\right) + DB(e_R) + ra_j\right), \\
 &\quad j \in \{R+1, \dots, D\}.
 \end{aligned}$$

The first two equations define taxable income for working and for retired households. The last two equations show the evolution of resources available for consumption. In these constraints e_j denotes labor earnings at age j ; $SS(\cdot)$ are social security benefits, which are a function of aggregate lifetime earnings; and $DB(\cdot)$ are defined-benefit receipts, which are a function of earnings received at the last working age. The functions $T(\cdot)$ and $T_R(\cdot)$ denote means-tested transfers for working and retired households. Transfers depend on earnings, social security benefits, and defined-benefit pensions, assets, the year, and the number of equivalent adults in the household, n . Medical expenditures are denoted by m_j , and the interest rate is denoted by r . The function $\tau(\cdot)$ is a tax function that depicts total tax payments as a function of earned and capital income for working households and as a function of pension and capital income plus a portion of social security benefits for retired households.¹²

¹¹ The economic environment implies a borrowing constraint in the sense that asset balances are nonnegative in every period. The intuition is the following: for the problem to be well specified, the household should not be allowed to die with debt, regardless of the stochastic sequence of earnings (and medical) shocks. Since earnings shocks in every period can get arbitrarily close to zero, the household should be in a position to repay debt even if it gets a long sequence of near-zero earnings draws; when this fails, consumption goes to zero and marginal utility of consumption goes to infinity, which is clearly not optimal (since the utility function satisfies the Inada condition). Consequently, the household will maintain a nonnegative asset position at every age. The same logic applies in retirement, with the exception that rather than earnings uncertainty, the individual now faces uncertainty in medical expenses and life span.

¹² Specifically, taxable social security benefits for single taxpayers are calculated from the expression

$$\begin{aligned}
 \Psi(\text{SS benefits}) &= \max(0, \min(0.5 \times \text{SS benefits}, \text{income} - 0.5 \\
 &\quad \times \text{SS benefits} - 25,000)).
 \end{aligned}$$

Taxable benefits for married couples are calculated similarly, but replacing 25,000 with 32,000. This approach approximates the law in effect in 1992.

B. Recursive Formulation of the Life Cycle Problem

We solve the life cycle problem backward from age D , given the terminal condition at that age. There are two sources of uncertainty in retirement—life span and medical expenses. We start by describing the problem for retired married households. The problem for retired single households is dealt with in a similar fashion.

1. The Retired Household's Problem¹³

A retired household between the ages of $R + 1$ and D obtains income from social security, defined-benefit pensions, and preretirement assets. The dynamic programming problem at age j for a retired married household with both members alive at the beginning of age j is given by

$$\begin{aligned}
 V_R(e_R, E_R, a_j, j, m_j, 3) = & \max_{c_j, a_{j+1}} \left\{ n_j U\left(\frac{c_j}{n_j}\right) \right. \\
 & + \beta p_{hj} p_{wj} \int V_R(e_R, E_R, a_{j+1}, j+1, m_{j+1}, 3) d\Omega_j(m_{j+1}|m_j) \\
 & + \beta p_{hj}(1 - p_{wj}) \int V_R\left(e_R, E_R, a_{j+1}, j+1, \frac{m_{j+1}}{2}, 1\right) d\Omega_j(m_{j+1}|m_j) \\
 & \left. + \beta p_{wj}(1 - p_{hj}) \int V_R\left(e_R, E_R, a_{j+1}, j+1, \frac{m_{j+1}}{2}, 2\right) d\Omega_j(m_{j+1}|m_j) \right\} \quad (2)
 \end{aligned}$$

subject to

$$\begin{aligned}
 y &= \text{SS}(E_R) + \text{DB}(e_R) + ra_j + T_R(e_R, E_R, a_j, n_j, m_j), \\
 c_j + a_{j+1} + m_j &= y_j + a_j - \tau(\Psi(\text{SS}(E_R)) + \text{DB}(e_R) + ra_j). \quad (3)
 \end{aligned}$$

In equation (2), $V_R(e_R, E_R, a_j, j, m_j, 3)$ denotes the present discounted value of maximized utility from age j until the date of death; $V_R(e_R, E_R, a_{j+1}, j+1, m_{j+1}, 3)$ denotes the corresponding value in the following year; β is the discount factor on future utilities; and p_{hj} and p_{wj} denote the probability of survival between ages j and $j+1$ for the household head and the spouse, respectively. Medical expenses are drawn from the Markov processes $\Omega_j(m_{j+1}|m_j)$ for married households.¹⁴ Total earnings up to the retirement period are denoted by $E_R \equiv \sum_{j=S}^R e_j$ and the last

¹³ To define a household's retirement date for those already retired, we use the actual retirement date for the head of the household. For those not retired, we use the expected retirement date of the person who is the head of the household.

¹⁴ Medical expenses for married households that become single are assumed to be half of those drawn from the distribution for married couples.

earnings draw at the age of retirement is e_R . Note that E_R and e_R do not change once the household is retired. The integers in the last argument of the value function signify that only the husband is alive (1), only the wife is alive (2), or both the husband and wife are alive (3) at the beginning of the period.

2. The Working Household's Problem

We assume that households incur no out-of-pocket medical expenses prior to retirement and face no preretirement mortality risk. Therefore, the dynamic programming problem for working households has two fewer state variables than it does for retired households. Between ages S and R , the household receives an exogenous earnings draw e_j . Given earnings and assets, the household decides how much to consume and save. The decision problem is

$$V(e_j, E_{j-1}, a_j, j) = \max_{c_j, a_{j+1}} \left\{ n_j U\left(\frac{c_j}{n_j}\right) + \beta \int V(e_{j+1}, E_j, a_{j+1}, j+1) d\Phi_j(e_{j+1}|e_j) \right\} \quad (4)$$

subject to

$$y_j = e_j + ra_j + T(e_j, a_j, n_j),$$

$$c_j + a_{j+1} = y_j + a_j - \tau(e_j + ra_j),$$

and

$$E_j = E_{j-1} + e_j.$$

Note that during working years, the earnings draw for the next period comes from the distribution Φ_j conditional on the household's age and current earnings draw. The solution to this problem yields the decision rule that we denote $a_{j+1} = G(e_j, E_{j-1}, a_j, j)$. We assume that each household begins life with zero assets.

At age R the household knows that in the next period it will cease working and begin receiving income from social security and defined-benefit pensions. The recursive representation of the optimization problem at age R must reflect the fact that the future utility value is given by $V_R(e_R, E_R, a_R, R, m_R, n_R)$. See Scholz, Seshadri, and Khitatrakun (2004, sec. II.2.2) for details.

III. Model Parameterization and Estimation of Exogenous Processes

In this section we specify functional forms and parameter values that we use to solve the model.

Preferences.—The utility function for consumption of final goods is assumed to exhibit constant relative risk aversion:

$$U(c) = \begin{cases} \frac{c^{1-\gamma}}{1-\gamma} & \text{if } \gamma \neq 1 \\ \log c & \text{if } \gamma = 1. \end{cases}$$

The discount factor is set as $\beta = 0.96$, and the coefficient of relative risk aversion (the reciprocal of the intertemporal elasticity of substitution) is set as $\gamma = 3$. These parameters are similar to those used by Hubbard et al. (1995) and Engen et al. (1999), who use 0.97 for the discount factor and 3 for the coefficient of relative risk aversion.¹⁵ We describe sensitivity analyses on these key parameters below.

Equivalence scale.—This is obtained from Citro and Michael (1995) and takes the form $n_j = g(A_j, K_j) = (A_j + 0.7K_j)^{0.7}$, where again A_j indicates the number of adults in the household and K_j indicates the number of children in the household.

Survival probabilities.—These are based on the 2002 life tables of the Centers for Disease Control and Prevention, U.S. Department of Health and Human Services (http://www.cdc.gov/nchs/data/nvsr/nvsr53/nvsr53_06.pdf).

Rate of return.—We assume an annualized real rate of return of 4 percent. This assumption is consistent with that in the article by McGrattan and Prescott (2003), who find that the real rate of return for both equity and debt in the United States over the last 100 years, after accounting for taxes on dividends and diversification costs, is about 4 percent.¹⁶ We include sensitivity analysis on this parameter below.

Taxes.—We model an exogenous, time-varying, progressive income tax that takes the form

$$\tau(y) = a_0[y - (y^{-a_1} + a_2)^{-1/a_1}],$$

where y is in thousands of dollars. Parameters are estimated by Gouveia and Strauss (1994, 2000) and characterize U.S. effective, average house-

¹⁵ Samwick (1998) allows discount rates to vary when calibrating a life cycle model with uncertain incomes to household wealth data from the 1992 Survey of Consumer Finances. He uses the model to examine the effects of Social Security privatization on saving.

¹⁶ Four percent is also the difference (rounded to the nearest percentage point) of the average real stock market return between 1947 and 1996 (7.6 percent) and the average real return on three-month Treasury bills (0.8 percent).

hold income taxes between 1966 and 1989.¹⁷ We use the 1966 parameters for years before 1966 and the 1989 parameters for 1990 and 1991.

Transfers.—We model the cumulative benefits from public income transfer programs using a specification suggested by Hubbard et al. (1995). Specifically, the transfer that a household receives while working is given by

$$T(e_j, a_j, n_j) = \max \left\{ 0, \underline{c} \times \frac{n_j}{g(1, 2)} - [e_j + (1 + r)a_j] \right\},$$

whereas the transfer that the household will receive upon retiring is

$$T_R(e_R, E_R, a_j, n_j, m_j) = \max \left\{ 0, \underline{c} \times \frac{n_j}{g(1, 2)} + m_j - [\text{SS}(E_R) + \text{DB}(e_R) + (1 + r)a_j] \right\}.$$

This transfer function guarantees a pretax income of \underline{c} , which we set on the basis of parameters drawn from Moffitt (2002).¹⁸ Subsistence benefits (\underline{c}) for a one-parent family with two children increased sharply, from \$5,992 in 1968 to \$9,887 in 1974 (all in 1992 dollars). Benefits have trended down from their 1974 peak: in 1992 the consumption floor was \$8,159 for the one-parent, two-child family. We assume through this formulation that earnings, retirement income, and assets reduce public benefits dollar for dollar.

Social security and defined-benefit functions.—We calculate a close approximation of each household's social security entitlement making use of the social security earnings records. Households in the model expect the social security rules in 1992 to prevail and develop expectations of social security benefits that are consistent with their earnings expectations. Details concerning the social security calculations are given in online Appendix B.

¹⁷ Estimated parameters, e.g., in 1989 are $a_0 = 0.258$, $a_1 = 0.768$, and $a_2 = 0.031$. In this framework, $a_1 = -1$ corresponds to a lump-sum tax with $\tau(y) = -a_0 a_2$; when $a_1 \rightarrow 0$, the tax system converges to a proportional tax system with $\tau(y) = a_0 y$. For $a_1 > 0$, we have a progressive tax system.

¹⁸ The \underline{c} in the model reflects the consumption floor that is the result of all transfers (including, e.g., supplemental security income). Moffitt (2002) provides a consistent series for average benefits received by a family of four. To proxy for the effects of all transfer programs, we use his "modified real benefit sum" variable, which roughly accounts for the cash value of food stamps, Aid to Families with Dependent Children (AFDC), and Medicaid guarantees. We weight state-level benefits by population to calculate an average national income floor. We use 1960 values for years prior to 1960 and use the equivalence scale described above to adjust benefits for families with different configurations of adults and children. We confirm that the equivalence scale adjustments closely match average benefit patterns for families with different numbers of adults and children using data from the Green Book (1983, 259–60, 301–2; 1988, 410–12, 789).

Defined-benefit pension expectations are formed on the basis of an empirical pension function that depends in a nonlinear way on union status, years of service in the pension-covered job, and expectations about earnings in the last year of work. We estimate the function with HRS data. Details are given in Section A of Appendix A.

Earnings process.—The basic unit of analysis for our life cycle model is the household. We aggregate individual earnings histories into household earnings histories. Earnings expectations are a central influence on life cycle consumption decisions, both directly and through their effects on expected pension and social security benefits. The household model of log earnings (and earnings expectations) is

$$\begin{aligned}\log e_j &= \alpha^i + \beta_1 \text{AGE}_j + \beta_2 \text{AGE}_j^2 + u_j, \\ u_j &= \rho u_{j-1} + \epsilon_j,\end{aligned}$$

where e_j is the observed earnings of household i at age j in 1992 dollars, α^i is a household-specific constant, AGE_j is the age of the head of the household, u_j is an AR(1) error term of the earnings equation, and ϵ_j is a zero-mean independently and identically distributed, normally distributed error term. The estimated parameters are α^i , β_1 , β_2 , ρ , and σ_ϵ .

We divide households into six groups according to marital status, education, and number of earners in the household, giving us six sets of household group-specific parameters.¹⁹ Estimates are given in Section B of Appendix A, where the group averages of the household-specific effects are given as the constants.

Estimates of the persistence parameters range from 0.58 for single households without college degrees to 0.76 for married households with two earners in which the highest earner has at least a college degree. The variance of earnings shocks ranges from 0.08 for married households with either one or two earners and in which the highest earner has at least a college degree to 0.21 for single households without college degrees.

We set the persistence parameter to 0.90 for all groups as part of our sensitivity analysis.

¹⁹ The six groups are (1) single without a college degree; (2) single with a college degree or more; (3) married, head without a college degree, one earner; (4) married, head without a college degree, two earners; (5) married, head with a college degree, one earner; and (6) married, head with a college degree, two earners. A respondent is an earner if his or her lifetime earnings are positive and contribute at least 20 percent of the lifetime earnings of the household.

Out-of-pocket medical expenses.—The specification for household medical expense profiles for retired households is given by

$$\log m_t = \beta_0 + \beta_1 \text{AGE}_t + \beta_2 \text{AGE}_t^2 + u_t$$

$$u_t = \rho u_{t-1} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma_\epsilon^2),$$

where m_t is the household's out-of-pocket medical expenses at time t (the medical expenses are assumed to be \$1.00 if the self-report is zero or if the household has not yet retired), AGE_t is the age of the household head at time t , u_t is an AR(1) error term, and ϵ_t is white noise. The parameters to be estimated are β_0 , β_1 , β_2 , ρ , and σ_ϵ .

We estimate the medical expense specification for four groups of households: (1) single without a college degree, (2) single with a college degree, (3) married without a college degree, and (4) married with a college degree. We use the 1998 and 2000 waves of the HRS, which provide medical expense information on households aged 27–106.²⁰ We use the age and education of the head of the household in the empirical model. Results are given in Section C of Appendix A. The persistence parameters for medical shocks cluster tightly between 0.84 and 0.86 across groups. The variance of shocks is lower for households with greater education within a given household type (married or single), presumably reflecting higher rates of insurance coverage for households with college degrees relative to others.

Some readers might be concerned that medical price inflation or perhaps high anticipated future substantial out-of-pocket, end-of-life medical or nursing home expenses would make *current* HRS data poor measures for needed *future* out-of-pocket medical expenses.²¹ In the sensitivity analysis, therefore, we also model the possibility that all households face a 5 percent chance of having four consecutive years (prior to death) of incurring \$60,000 of out-of-pocket medical expenses, keeping all other aspects of the model the same.²²

²⁰ Older cohorts from the AHEAD and two new cohorts were added to the HRS in 1998, which gives us a broader range of ages to estimate medical expense profiles after retirement. These new cohorts were not matched to their social security earnings records, so they cannot be used for our baseline analysis.

²¹ Kopcuk (2005) examines exit interviews from the HRS/AHEAD and finds that end-of-life expenses average 12.4 percent of the value of the estate (45.5 percent with funeral expenses). For those with estates greater than \$295,888 (in 1992 dollars), the corresponding numbers were just 1.3 percent and 2.8 percent.

²² According to the 2002 MetLife Market Survey on Nursing Home and Home Care Costs, \$60,000 is roughly the national average for a full year of nursing home care in a private room. Brown and Finkelstein (2004, table 1) estimate that men (women) have a 27 (44) percent probability of using a nursing home. The average stay for men (women) is 1.3 (2.0) years. For men (women) who enter a nursing home, 5 (12) percent of the stays last five or more years.

A. Model Solution

We solve the dynamic programming problem by iterating on the value function. Recall that the state space is composed of six variables for retired households: earnings drawn at $j = R$, e_R ; cumulative earnings at the time of retirement, E_R ; assets, a ; age, j ; medical expenses, m ; and whether the head of the household, the spouse, or both are alive at the beginning of the period (as noted earlier, we assume that there are no mortality risks and out-of-pocket medical expenses for working households). We begin by “discretizing” the state space. The 25-point grid for earnings is constructed using the procedure discussed in Tauchen (1986). The 100-point grid for assets is chosen to be denser at lower levels of assets and progressively coarser so as to account for nonlinearities in the decision rules for assets induced by the borrowing constraint. We start at age D , assumed to be 100, and compute the value function $V_R(e_R, E_R, a_D, D, m_D, n_D)$ associated with all possible states in the discretized set. (The problem at this stage is trivial, since the household will simply consume all income.) We move backward to the previous period and solve for the value function and the decision rule for assets. We go all the way to the starting age S and consequently recover the decision rules $a_{j+1} = G(e_j, E_{j-1}, a_j, j)$ for all $j = S, \dots, R$ and $a_{j+1} = G_R(e_R, E_R, a_j, j, m_j, n_j)$ for all $j = R + 1, \dots, D$.

In summary, for each household in our sample we compute optimal decision rules for consumption (and hence asset accumulation) from the oldest possible age (D) to the beginning of its working life (S) for any feasible realizations of the random variables: earnings, health shocks, and mortality. These decision rules differ for each household, since each faces stochastic draws from different earnings distributions (recall that α^i is household specific). Household-specific earnings expectations also directly influence expectations about social security and pension benefits. Other characteristics also differ across households: for example, birth years of children affect the “adult equivalents” in a household at any given age. Consequently, it is not sufficient to solve the life cycle problem for just a few household types.

Once optimal decision rules are solved for each household, we calculate optimal consumption (and therefore wealth) each period for each household using data on the observed realizations of earnings. Specifically, we start at age S , the first working age, where the household is assumed to begin with zero assets. Earnings to date are also zero at S . Given observed earnings at age S , \hat{e}_S , wealth (saving) is given by $a_{S+1} = G(\hat{e}_S, 0, 0, S)$. In the next period, the household receives an observed earnings draw \hat{e}_{S+1} , so aggregate earnings are given by $\hat{E}_S = \hat{e}_S$. Consequently, wealth is given by $a_{S+1} = G(\hat{e}_{S+1}, \hat{E}_S, a_S, S + 1)$. We move

forward in this fashion until we reach the age at which wealth data are available for that particular household.²³

IV. Model Predictions and Their Correspondence to HRS Data

Table 2 summarizes the distribution of optimal net worth for HRS households. These targets include resources that could be accumulated in real and financial assets, the current value of defined-contribution pensions, including 401(k)'s, and housing net worth (for now, we assume that households are willing to reduce housing in retirement to maintain consumption standards).

The optimal wealth target for the median households in the lowest decile of the lifetime earnings distribution is very low, at \$2,050 (including housing wealth). The mean target for households in the bottom decile is \$48,445. These low targets are a consequence of four factors. First, lifetime earnings are low for bottom decile households, and social security is mildly progressive. Second, the number of children in this cohort is inversely related to lifetime earnings (married couples in the bottom decile of lifetime earnings had 5.3 children and couples in the highest decile had 3.2 children). This has an effect similar to increasing the discount rate for low-income households, hence reducing optimal wealth accumulation, all else being equal. Third, the average age of households is 55.7, so the average household will work (and accumulate wealth) for many additional years before retiring. Fourth, means-tested transfer programs have income and asset tests, which lower net worth relative to a world without a safety net (Hubbard et al. 1995).²⁴

Optimal wealth targets are \$63,116 for the median household and \$238,073 for the median household in the highest decile of the lifetime earnings distribution. The respective optimal mean targets are \$157,246 for the sample and \$463,807 for households in the top decile. The targets increase monotonically with lifetime earnings and with educational attainment.

A central feature of our work that distinguishes it from earlier papers is that we can compare optimal levels of wealth with actual wealth for each household in the HRS.

²³ Solving the optimal decision rules is computationally demanding, so we do not estimate the fit-maximizing discount rate and the coefficient of relative risk aversion. Computational considerations also limit our ability to add more features to the model.

²⁴ Empirical work on the effects of asset tests and asset accumulation comes to mixed conclusions. Gruber and Yelowitz (1999) find significant negative effects of Medicaid on asset accumulation, but Hurst and Ziliak (2006) find only very small effects of AFDC and food stamp asset limits.

TABLE 2
OPTIMAL NET WORTH (Excluding Social Security and Defined-Benefit Pensions) AND THE PERCENTAGE OF POPULATION FAILING TO MEET
OPTIMAL TARGETS (Dollar Amounts in 1992 Dollars)

Group	Median Optimal Wealth Target (1)	Mean Optimal Wealth Target (2)	Percentage below Optimal Target (3)	Median Deficit (Conditional) (4)	Median Net Worth (5)	Median Social Security Wealth (6)	Median Defined-Benefit Pension Wealth (7)
All households	\$63,116	\$157,246	15.6%	\$5,260	\$102,600	\$97,726	\$17,327
No high school diploma	20,578	70,711	18.6	2,632	36,800	72,561	0
High school diploma	63,426	139,732	15.6	5,702	102,885	97,972	21,290
College degree	128,887	243,706	12.7	14,092	209,616	127,704	60,752
Postcollege education	158,926	333,713	13.2	23,234	253,000	129,320	152,781
Earnings decile:							
Lowest	2,050	48,445	30.4	2,481	5,000	26,202	0
2nd	13,781	55,898	28.7	3,328	25,500	42,159	0
3rd	26,698	84,582	21.8	5,948	43,485	57,844	0
4th	43,566	123,441	19.4	4,730	75,000	77,452	14,830
Middle	53,709	128,285	16.9	6,979	90,000	94,929	29,497
6th	76,462	131,565	10.8	10,000	124,348	119,011	45,613
7th	80,402	154,891	9.9	11,379	128,580	133,451	56,033
8th	101,034	180,643	5.5	21,036	167,000	151,397	71,373
9th	136,075	238,186	4.4	5,206	226,000	163,639	104,657
Highest	238,073	463,807	5.4	25,855	393,000	202,659	126,998

SOURCE.—Authors' calculations as described in the text.

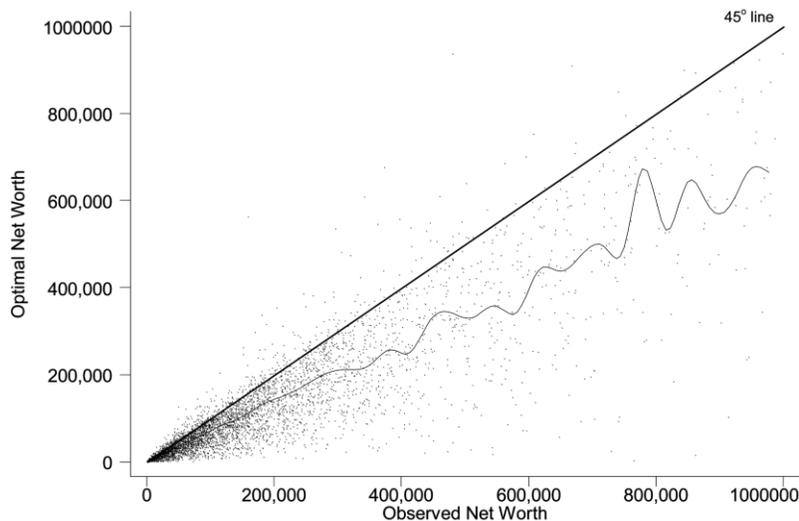


FIG. 2.—Scatter plot of optimal and actual wealth. Observed net worth is constructed from the 1992 HRS. Optimal net worth comes from solving the baseline model described in the text.

A. *Are Households Preparing Optimally for Retirement?*

Figure 2 gives a scatter plot of optimal net worth against actual net worth, for HRS households with optimal and actual net worth between \$0 and \$1,000,000. The curved line gives a cubic spline of the median values of observed and optimal net worth.²⁵ If household net worth was exactly the same as optimal net worth, the ordered pairs of actual and optimal net worth for the HRS sample would map out the 45-degree line. In fact, the ordered pairs cluster just below the 45-degree line. The scatter plot gives striking visual evidence that most HRS households have saved at or above their optimal retirement targets.

A second striking aspect of figure 2 is that it illustrates how a well-specified life cycle model can closely account for variation in cross-sectional household wealth accumulation. A linear regression of actual net worth against predicted net worth and a constant shows that the model explains 86 percent of the cross-household variation in wealth (i.e., the R^2 is 86 percent).

Column 3 of table 2 shows the fraction of HRS households with wealth deficits, broken out by educational attainment and lifetime earnings deciles. Overall, 15.6 percent of the HRS sample has deficits (their net worth, excluding defined-benefit pensions and social security, is less than

²⁵ The median band is smoothed by dividing households into 30 groups on the basis of observed net worth. We use Stata's "connect(s) bands(30)" option for the figure.

the optimal target).²⁶ Moreover, the median magnitude (conditional on having a deficit) of \$5,260 is very small. Although some households are approaching retirement with significant wealth deficits, the data in table 2 suggest that HRS households overwhelmingly are well prepared for retirement.

There are three additional features of the data in table 2 that are noteworthy. First, across education groups and lifetime earnings deciles, median net worth substantially exceeds the optimal median targets. Later in the paper we discuss a few reasons why households might accumulate more than the optimal level of wealth. Second, recall that the targets refer to net worth *excluding* social security and defined-benefit pension wealth. Columns 6 and 7 of table 2, which show the median social security wealth and the median defined-benefit pension entitlement, help put the magnitude of the wealth targets and deficits in perspective. The median deficits for those who have them are a very small percentage of accumulated social security wealth, pension entitlements, and net worth outside of defined-benefit pensions and social security.

Third, table 2 suggests that the probability of failing to meet the target is 30.4 percent in the decile with the lowest lifetime earnings and falls to 5.4 percent in the decile with the highest lifetime earnings. This is a case, however, in which the simple cross-tabulation is particularly misleading. In a probit regression, lifetime earnings do not have a statistically significant effect on the probability that an HRS household failed to meet its optimal wealth target once we condition on other covariates.²⁷ The only factor that is strongly correlated with having a wealth deficit is being single: married households are 21.9 percentage points less likely to have a deficit than single households.²⁸ These results suggest that undersaving is approximately randomly distributed throughout the population; it is not a phenomenon disproportionately affecting poor households or households with low levels of education. The strong income gradient shown in table 2 is purely a composition effect: single house-

²⁶ This result is broadly consistent with those of other studies. Hurst (2003), e.g., shows that between 10 and 20 percent of the population appears not to be following the permanent income hypothesis in the Panel Study of Income Dynamics (PSID). In brief, he splits the sample into low-residual undersavers (the bottom 10–20 percent of wealth residuals) and other households on the basis of a log wealth regression estimated from the 1989 wealth supplement of the PSID. He then shows that undersavers violate Euler equation excess sensitivity tests, whereas other households do not.

²⁷ The income coefficients are not jointly significant at even the 25 percent level of confidence, nor are the coefficients for the three highest lifetime income deciles. See Scholz et al. (2004) for details.

²⁸ We also find no quantitatively significant correlates with failing to meet optimal targets in probit regressions estimated separately on a sample of married couples. In a similar regression estimated on single households, coefficients for the three highest lifetime income deciles are negative and jointly significant at 5 percent.

holds are much more likely than married households not to meet their wealth targets. Since single households are more likely to have lower incomes than married households, they are disproportionately represented in the lower deciles of the lifetime earnings distribution.²⁹

B. *Are Americans Oversaving?*

To this point we have presented figures only for the median household in the population or the median household within education groups or lifetime earnings deciles. Figure 3 shows selected percentiles—tenth, twenty-fifth, fiftieth, seventy-fifth, and ninetieth—of the *distribution* of the difference between actual and optimal wealth targets across lifetime earnings deciles. Two things are striking from this figure. First, only very small percentages of households accumulated less than their optimal wealth target. Undersavers are concentrated in the bottom half of the lifetime earnings distribution. And the magnitudes of the shortfalls, conditional on having a shortfall, are small. Second, the most striking aspect of figure 3 is the degree to which people are saving too much. We probe this result in the remaining part of this subsection.

There is some question about the degree to which the elderly are willing to reduce housing equity to sustain consumption in retirement. Venti and Wise (2004, 170), for example, write that “in considering whether families have saved enough to maintain their preretirement standard of living after retirement, housing equity should not, in general, be counted on to support nonhousing consumption.” This conclusion is controversial. Hurd (2003) shows that elderly households decumulate housing wealth as they age in the AHEAD sample. Engen et al. (1999) make a forceful case for including at least a significant portion of housing wealth when measuring resources households can draw on to maintain living standards in retirement.³⁰ Nevertheless, we do not

²⁹ We were also concerned that the results for singles could be driven by divorce. If one-earner, married households divorced prior to the HRS survey, we would likely treat single earners as undersaving (they had income that before was supporting a family, yet following the divorce they would be expected to have only half the assets). Similarly, the nonworking partners would appear to have oversaved: they earned no income but are observed to have half the previously accumulated assets. However, this concern appears to be misplaced: the fraction of singles failing to meet their wealth targets is stable as we drop recently divorced individuals from the sample or when we drop ever-divorced individuals.

³⁰ Engen et al. make four points. First, existing work suggesting that the elderly do not decumulate is flawed; housing should be the last asset to tap since it is illiquid and tax-preferred, and because some evidence is based on cohorts that were considerably less mobile than the HRS cohort. Second, households have vigorously extracted equity from houses in the 1980s and 1990s. Third, tax consequences of selling housing have fallen in recent years, making it difficult to make inferences about people’s willingness to downsize from earlier data. Fourth, housing provides consumption services and thus represents wealth. Conceptually and from a policy perspective, it seems odd to ignore one important source of wealth when considering economic well-being among households in retirement.

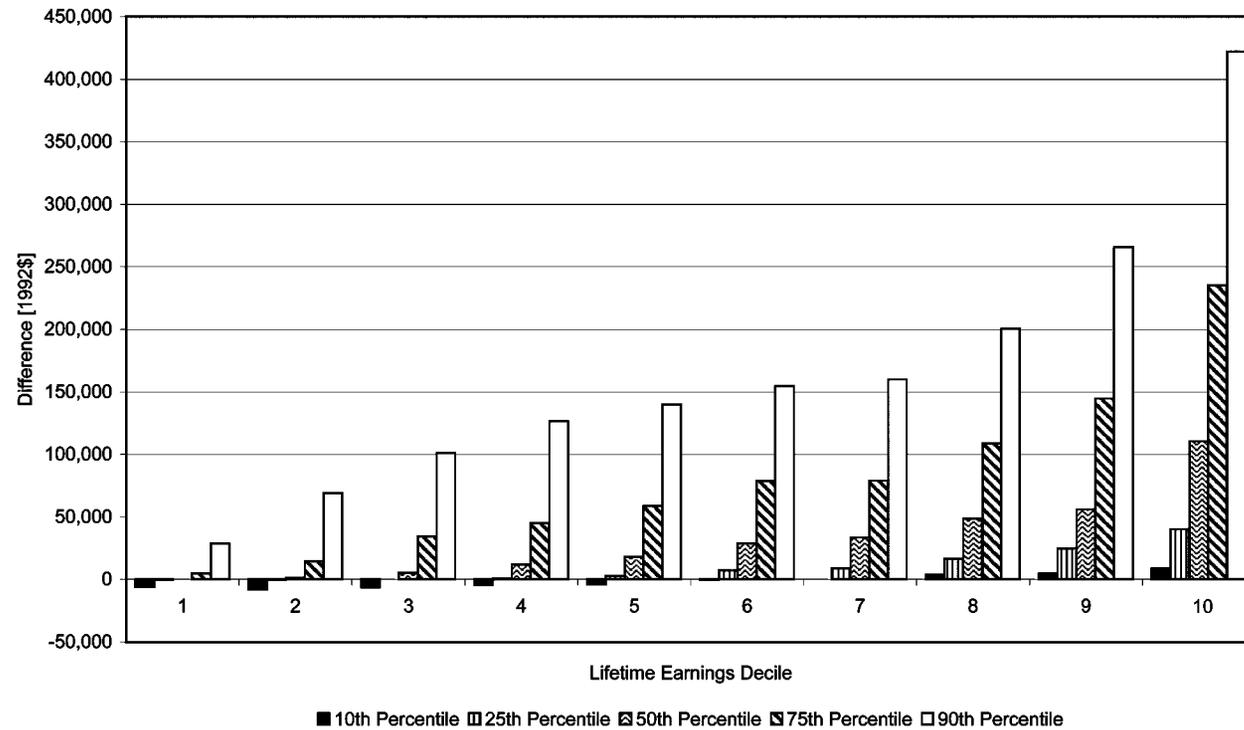


FIG. 3.—Distribution of “saving adequacy” (observed minus simulated optimal net worth, excluding defined-benefit pensions, by lifetime earnings decile, 1992 dollars).

want our conclusion, that a substantial majority of Americans are preparing well for retirement, to be driven by our treatment of home equity.

To explore the consequences of altering the treatment of housing in our calculations, we also examine the distribution of wealth deficits excluding half of housing from the resources available to meet the wealth target. When we exclude half of housing equity, 61.2 percent of all households meet or exceed their wealth targets. The twenty-fifth percentile of the saving surplus distribution (net worth minus optimal targets) has a deficit of \$7,692, implying that 75 percent of households are exceeding or are within \$7,692 of their optimal (nonpension, non-social security) wealth target, even when half their net home equity is excluded. Households in the lowest decile of the lifetime earnings distribution have an average deficit of \$33,178.³¹ We report results when using all net worth for the remainder of the paper, but note that our results are qualitatively similar if we exclude half of housing equity.

There are at least three features of the model that could account for the fact that many households appear to be accumulating significantly more than their optimal life cycle targets.³² First, we assume that households expected and received a real rate of return of 4 percent. To the extent that perceptions or realizations of real rates of return differ from our assumption, households will accumulate less or more than the target. There is little we can do to address this concern beyond including broad geographic indicators to crudely account for potential differences in house price appreciation.³³

Second, households may intend to leave bequests. We use HRS questions probing the subjective likelihood of households leaving bequests of \$10,000 and \$100,000 to explore whether households with a high (or certain) likelihood of leaving a bequest of these magnitudes are more likely than other households to exceed their wealth targets.³⁴

Third, households might expect to live longer than suggested by the life tables published by the Centers for Disease Control and Prevention. Households with expectations of greater longevity will optimally accumulate more resources than predicted by our model. We use HRS ques-

³¹ A figure similar to fig. 3 but excluding half of housing wealth is given in Scholz et al. (2004).

³² A fourth factor that could result in incorrect analyses would be that our chosen preference parameters are incorrect. We explore the model sensitivity to preference parameters in Sec. IV E.

³³ HRS rules prohibit using restricted-access geocoded data with the restricted-access earnings data. The repeat-sales house price index increased 381 percent between 1975 and 1992 in the Pacific, 278 percent in New England, and less than 120 percent in the East South Central, West South Central, and West North Central divisions. See <http://www.huduser.org/periodicals/ushmc/fall02/histdat10.htm>.

³⁴ The specific questions come from the 1994 wave of the HRS and read, "What are the chances that you [or your husband/wife/partner] will leave an inheritance totaling \$10,000 (or \$100,000) or more?"

tions probing subjective expectations of life expectancy to explore the importance of this factor in explaining oversaving.³⁵

We examine the importance of rate of return differences, bequest intentions, and longevity risk along with other factors in an empirical median regression model of “saving adequacy,” defined as the difference between actual net worth (excluding defined-benefit pensions and social security wealth) and optimal net worth. The results are shown in table 3. There is a sharply increasing, positive relationship between the net worth surplus and lifetime earnings that begins to significantly differ from the reference (bottom) category in the sixth lifetime earnings decile. There is also a strong positive relationship between the net worth surplus (with social security wealth and defined-benefit pensions excluded) and age, being self-employed, being retired at the time of the survey, and being married. The indicator variables for male and for holding an employer-provided pension are negatively correlated with oversaving in the form of nonpension wealth. There is no evidence that region of the country or subjective life expectancies have any relationship with saving adequacy (or oversaving). Bequest intentions, however, are positively, significantly related to acquiring more wealth than the optimal target. This result is consistent with purposeful bequest intentions affecting life cycle wealth accumulation, though the direction of causality associated with this correlation is, of course, unclear.

C. *Alternative, “Naive” Models*

In the conclusion of their paper on variation in retirement wealth, Bernheim et al. (2001, 855) write that “the empirical patterns in this paper are more easily explained if one steps outside the framework of rational, far-sighted optimization. If, for example, households follow heuristic rules of thumb to determine saving prior to retirement” Indeed, naive or rule-of-thumb models of consumption have had an important place in the consumption literature at least since the Keynesian consumption function.

The exceptionally rich data we have on household earnings contain a great deal of information. Health shocks prior to retirement, unemployment, and changes in labor demand and supply, among other things, will be reflected in the 41-year series of earnings we have for most households. Given the rich earnings data, it is natural to ask how much of the variation in HRS wealth can be explained by applying simple, rule-of-thumb saving behavior to the household-specific earnings trajectories. Our results are summarized in panel A of table 4.

³⁵ The specific questions come from the 1992 wave of the HRS and read, “What do you think are the chances that you will live to be 75 (or 85) or more?”

TABLE 3
CORRELATES OF THE MEDIAN WEALTH SURPLUS

	MEDIAN REGRESSION OF "SAVING ADE- QUACY" (Observed – Optimal Net Worth)	
	Coefficient Estimates	Standard Error
Constant	-22,721.7***	6,512.0
Lifetime income decile:		
2nd	-654.2	1,016.9
3rd	523.4	1,533.8
4th	139.3	1,765.4
5th	2,305.3	2,457.5
6th	8,032.6**	3,326.0
7th	12,978.1***	3,870.5
8th	19,828.8***	4,114.7
9th	26,526.0***	5,185.3
10th	60,023.1***	8,446.7
Retired	2,047.4**	868.1
Has pension	-2,362.5**	985.8
Social security wealth	.043*	.023
Age	352.6***	99.9
Male	-2,583.7**	1,032.3
Black	-1,746.3**	801.7
Hispanic	-1,227.9	1,029.9
Married	10,375.1***	1,401.4
High school diploma	276.5	814.2
College degree	4,876.5**	2,365.9
Graduate degree	4,556.8	3,283.1
Self-employed	14,226.7***	4,037.4
Number of children	-30.8	224.2
Number of grandchildren	-98.2	88.6
Subjective probability of living:		
≥ 75	11.7	16.3
≥ 85	-11.4	16.0
Subjective probability of bequest:		
≥ \$10,000	16.0*	9.6
≥ \$100,000	296.9***	31.3
Mid-Atlantic	1,692.6	3,747.3
East North Central	1,117.2	3,807.2
West North Central	4,079.0	3,936.5
South Atlantic	1,336.2	3,759.7
East South Central	2,166.4	3,908.3
West South Central	-282.1	3,963.4
Mountain	2,925.1	3,961.4
Pacific	3,736.3	3,941.6

NOTE.—Standard errors are bootstrapped in the median regression with 1,000 replications. The pseudo R^2 for the median regression is 0.1096, and the sample size is 4,952 because of missing values.

* Statistically significant at the 10 percent level.

** Statistically significant at the 5 percent level.

*** Statistically significant at the 1 percent level.

TABLE 4
ALTERNATIVE MODELS AND SENSITIVITY ANALYSIS

Parameter Value	Percentage Failing to Meet Optimal Target	Measure of Fit: R^2 (%)	Deficit Conditional on Failing to Meet Optimal Target (1992 Dollars)
Baseline: $\beta = 0.96$, $\gamma = 3$, $r = 4\%$	15.6	86.0	5,260
A. Alternative Models			
Naive:			
Save a constant fraction of Y_t	71.9	15.5	114,335
Save an income- and age-varying fraction of Y_t	75.7	15.8	160,676
Modigliani (annual consumption a function of lifetime resources)	48.7	6.5	89,129
Constant α	35.1	45.2	30,411
Regression with 41 years of earnings	59.4	29.2	109,212
Regression with quadratic terms for 41 years of earnings	60.2	35.3	101,229
Monte Carlo draws on earnings sequences	32.2	45.2	28,623
B. Parameter Sensitivity of Baseline Model			
$\beta = 1.0$	21.1	87.7	5,483
$\beta = 0.93$	11.9	83.6	5,919
$r = 5\%$	20.0	87.2	5,500
$r = 7\%$	35.9	76.7	15,955
$\gamma = 1.5$	11.8	91.9	4,699
$\gamma = 5$	31.6	85.9	9,087
$\rho = 0.9$	25.8	69.1	16,103
5% chance of 4 years of \$60,000 end-of-life medical expenses	20.5	85.1	4,800

SOURCE.—Authors' calculations as described in the text.

The simplest model we examined assumes that households save a constant fraction of their income, independent of their income or age. We iteratively sought the saving rate that maximized the goodness of fit measure, R^2 . The fit-maximizing saving rate is 14.7 percent, and the model explains 15.5 percent of the 1992 cross-sectional distribution of wealth in the HRS. A naive model with age-varying and income-varying saving rates, in this case drawn from the parameters estimated in Dynan, Skinner, and Zeldes (2004, table 3, col. 6), explains 15.8 percent of the variation in retirement wealth. The original formulation of the life cycle model (Modigliani and Brumberg 1954) in which households consume a constant real, discounted fraction of permanent income explains 6.5 percent of the variation in retirement wealth. It is clear that the augmented life cycle model presented in this paper, which explains 86.0

percent of the cross-sectional variation in wealth, does a vastly better job matching the cross-sectional distribution of wealth in the 1992 HRS than the rule-of-thumb models we examine.

Our augmented life cycle model includes many more parameters than the rule-of-thumb models. The household-specific intercept of the household age-earnings profiles, α^i , for example, adds 6,321 parameters to the model. We think that this is a sensible way to model earning expectations: households presumably have a reasonable understanding of their place in the ability distribution, given observable characteristics such as educational attainment and age. Nevertheless, we also consider an alternative, more parsimonious version of the baseline model using only the parameters shown in Appendix table A2, which restricts the α^i in the earnings expectations function to be the same within household types, and simulate the optimal decision rules. We find that the model can account for 45.2 percent of the observed variation in 1992 wealth. Thus our model, even with relatively few parameters, does a fairly good job of matching the 1992 cross-sectional distribution of wealth in the HRS.

Another useful benchmark for our augmented life cycle model is to compare its model fit to reduced-form regression models. We regress wealth in 1992 against earnings and a host of other household characteristics, in which each of the 41 years of earnings observations appear separately.³⁶ This regression accounts for 29.2 percent of the variation in wealth. If we add quadratic terms in each year of earnings, the resulting empirical model accounts for 35.3 percent of the variation in wealth. Even a parsimonious parameterization of the augmented life cycle model that includes many fewer parameters does a better job of explaining the observed variation in wealth than a regression that incorporates separately the linear and quadratic terms of annual earnings.

Our last experiment with alternative models attempts to clarify the importance of the augmented life cycle model's decision rules in explaining the model fit, relative to the unusually rich earnings histories we use. To examine this, consider the following thought experiment. Once we solve the model for each individual, we have the optimal decision rules. Now, rather than using the actual earnings draws to obtain wealth predictions for 1992, imagine that we were to obtain, for each household, 10,000 (sequences of) draws from the empirical earnings distribution and then use these draws to obtain predictions for wealth in 1992. The predicted wealth level for each household is the average

³⁶ The model also includes a quadratic in age, indicator variables for race and ethnicity, marital status, educational attainment, nine regions, 18 occupations, 14 industries, two-earner household, retired, self-employment, unemployment, health status, pension coverage, past marital history, and counts of the number of children, the number of children between 12 and 17, and the number of children between 18 and 21.

value implied by the potential realization of all such sequences. One can do the same for all the households in the sample and obtain the goodness of fit between the model and the data. The resulting R^2 is 45.2 percent. This suggests that although it is important to have the earnings realizations, the decision rules arising from the augmented life cycle model are equally critical in arriving at such a close correspondence between model and data.

D. Other Model Features

In this subsection we briefly examine three other features of the baseline model and show that they are consistent with several well-established facts about consumption. First, consumption is hump-shaped over the life cycle (see, e.g., Carroll 1997). Mean optimal consumption and income by age for our sample, after we net out children's consumption, is hump-shaped and peaks around age 46, whereas the peak in income occurs around age 52. These patterns are similar to those reported by Gourinchas and Parker (2002).³⁷

A second issue has to do with how our augmented life cycle model can match the well-known skewness of the wealth distribution. Our model predicts that the top 1 percent of our sample holds about 18.3 percent of the wealth (17.1 percent when sample weights are used). The corresponding figure in the data is 17.0 percent (15.9 percent with sample weights). With actual data on earnings realizations, our life cycle model does not need to add a bequest motive to match the skewness of the wealth distribution.

Third, a more stringent test of our model is how well it can match the change in wealth between 1992 and 2000. To put this in context, it is instructive to compare our model fit with a regression of the difference in wealth levels between 1992 and 2000 against household characteristics and earnings at every age.³⁸ The resulting R^2 is 3.8 percent. Introducing a quadratic term for earnings (at each age) increases the R^2 to 6.9 percent. In contrast, the baseline model (with a 4 percent real interest rate between 1992 and 2000) generates an R^2 of 22.6 percent. We conclude that our model does a good job of capturing observed behavior, not just of a cross section at a point in time, but also changes over time across households.

³⁷ In the early part of the life cycle, the variance of consumption growth is large and individuals are borrowing constrained, which leads consumption growth to be positive. The variance of consumption growth is smaller later in the life cycle. Since the after-tax discount rate is less than the discount factor, the growth rate of demographically adjusted household consumption is negative. The consumption growth rate changes sign around age 46 in the model.

³⁸ The regression model also includes the covariates listed in n. 36.

E. Sensitivity Analysis

There are three model parameters that we specify exogenously before solving the model: the discount factor (β), the coefficient of relative risk aversion (γ), and the real interest rate (r). In this subsection we analyze the sensitivity of the results to our choices of β , γ , and r . Panel B of table 4 shows the results. As expected, increases in β and r increase the incentive to save more in the future. In the life cycle model, this raises the optimal (or “target”) level of wealth. When these targets are matched to the observed HRS data, more households fail to save adequately for retirement. For example, raising the real interest rate from 4 percent to 7 percent increases the fraction of households with wealth less than the optimal level from 15.6 percent to 35.9 percent. An increase in γ has a similar effect because, as households become more risk averse, precautionary saving increases, increasing the optimal (or target) level of wealth accumulation and consequently undersaving. Nevertheless, the degree of undersaving is not particularly high: assuming that $\gamma = 5$, for example, increases the fraction of households with wealth less than the optimal from 15.6 percent to 31.6 percent and the median deficit is \$9,087.

Another parameter that plays an important role is the persistence in earnings across ages. Recall that these persistence parameters vary by type; they range from 0.58 for the single household without a college degree to 0.76 for the married, two-earner household in which the head has a college degree. These parameters were estimated directly from the 41 years of actual and imputed social security earnings data. But many life cycle models assume more persistence in the earnings process. The second to last line of table 4 arbitrarily sets all persistence parameters in earnings expectations to 0.9. Households understand that this dramatically increases the odds of retaining a bad or good draw if one is received. The resulting R^2 is 69.1 percent, and 25.8 percent of households fail to meet their optimal targets. As expected, increasing the persistence in earnings increases optimal wealth targets at retirement, and hence, more households will appear to undersave.

The last line of table 4 shows the results of estimates in which, in addition to facing out-of-pocket medical expense shocks, all households have a 5 percent chance of having four years of \$60,000 in end-of-life medical expenses. Households that are unable to privately finance this expenditure would still consume \underline{c} , which is the model proxy for poor elderly households receiving supplemental security income and having nursing home and health care expenses paid by Medicaid and Medicare. The R^2 of this specification is similar to that of the baseline model. But 20.5 percent of the households are failing to meet or exceed the optimal

target instead of 15.6 percent. The median conditional shortfall of \$4,800, however, is still a very small fraction of lifetime resources.

Our sensitivity analysis leads us to conclude that within the range of values considered, most households in the HRS appear to have saved well for retirement.³⁹ Moreover, within a reasonably broad range of parameter values, the model can explain at least 69.1 percent of the cross-sectional variation in wealth in the 1992 HRS. These results do not depend on the inclusion of households in the sample with fully imputed earnings histories. When we drop households that did not allow the HRS to have access to their social security earnings records, the results are nearly identical using our baseline parameters: 15.7 percent of households accumulate less wealth than their optimal targets. For the subsample with deficits, the median shortfall is \$5,000. And the model accounts for 86.3 percent of the cross-sectional distribution of wealth in this subsample.

V. Conclusions

In this paper we develop a rigorous approach for assessing the degree to which a representative sample of households nearing retirement have prepared financially for that event. We find strikingly little evidence that HRS households have undersaved. And because consumption requirements likely fall when households reach retirement (if for no other reason than work expenses fall), our standard may overstate required wealth. We also note that our primary data come from 1992 and 1993, well before the exceptionally strong stock market performance of the 1990s. Because 84.4 percent of households meet or exceed their wealth targets (and most of those who are below miss by a relatively small amount), we are skeptical that the consumption changes around retirement documented by Bernheim et al. (2001) are due to inadequate accumulation of retirement wealth.

We also find it striking how much of the variation in observed wealth accumulation can be explained by our life cycle model. We explain over 87 percent of the variation in wealth for married households and nearly 80 percent for single households. And the results presented reflect no tweaking or prior fitting of the model. If we had found major deviations from the model and behavior, it would be difficult to determine whether Americans were preparing poorly for retirement or we had constructed a poor behavioral benchmark. The fact that our predictions and data closely align suggests two things. First, as mentioned above, Americans

³⁹ If social security benefits are cut by 25 percent for all households, we find that 37.2 percent of households undersave. But we think that such benefit cuts are unlikely for the HRS cohort.

are saving enough to maintain living standards in retirement. And second, the life cycle model provides a very good representation of behavior related to the accumulation of retirement wealth. Of course, we still admit the possibility that Americans are preparing poorly for retirement, our underlying behavioral model is poor, and the errors, coincidentally, offset.

Although the specific measures of undersaving and model fit clearly depend on parameter values, our two main results—that the life cycle model is capable of closely matching the cross-sectional distribution of wealth in the HRS and that most HRS households are saving more than their optimal targets—are not affected significantly by parameter choices within the range commonly found in the related literature. We also find that the life cycle model does a much better job of matching the cross-sectional distribution of wealth in 1992 than a naive model in which households save an income- and age-varying fraction of income.

The title of this paper poses the following question: Are Americans saving optimally for retirement? The HRS covers a specific cohort of Americans—households aged 51–61 in 1992. Consequently, we need to be careful in generalizing our results for the HRS cohort to younger households. This is particularly true if the generosity of social security is reduced in the future. Moreover, saving too much has efficiency costs in the sense that, in the absence of preferences about intergenerational transfers or charitable contributions, reallocating consumption across time could increase lifetime utility. Because we cannot determine whether the systematic oversaving of HRS households reflects bequest motives, the expectation that social security will be reduced in the future, other failures in our characterization of the economic environment, or nonoptimal behavior on the part of HRS households, we cannot definitively answer the question posed. But the paper provides new, strong support for the life cycle model as a good characterization of the process governing retirement wealth accumulation. And more important, it adds considerably to our confidence that the Americans in the HRS cohort are preparing well for retirement.

Appendix A

Underlying Model Processes

A. *Defined-Benefit Pension*

The annual defined-benefit pension benefit is estimated as

$$\begin{aligned}
 db = & DB^h[\beta_0^h + \beta_1^h \text{UNION}^h + \beta_2^h \text{YRSV}^h + (\gamma_0^h + \gamma_1^h \text{UNION}^h \\
 & + \gamma_2^h \text{YRSV}^h)\phi_R^h e_R] + DB^w[\beta_0^w + \beta_1^w \text{UNION}^w + \beta_2^w \text{YRSV}^w \\
 & + (\gamma_0^w + \gamma_1^w \text{UNION}^w + \gamma_2^w \text{YRSV}^w)\phi_R^w e_R] + \beta_0^h DB^h DB^w + \xi, \quad (\text{A1})
 \end{aligned}$$

TABLE A1
COEFFICIENT ESTIMATES FOR ANNUAL DEFINED-BENEFIT PENSION BENEFITS

Variable	Coefficient Estimate	Standard Error
Husband's estimate of:		
Constant	1,903.7***	716.0
Union status	-457.610	612.5
Years in service	46.3	30.2
His last-period earnings	-.027	.024
His last-period earnings interacting with union status	.007	.022
His last-period earnings interacting with years in service	.004***	.001
Wife's estimate of:		
Constant	-249.394	534.3
Union status	1,128.7***	331.4
Years in service	67.5***	22.4
Her last-period earnings	.013	.032
Her last-period earnings interacting with union status	.002	.021
Her last-period earnings interacting with years in service	.004***	.001
Estimate of constant if both husband and wife have a pension	-176.902	421.1
R^2	.572	
Observations	2,203	

*** Statistically significant at the 1 percent level.

where the superscripts h and w indicate “husband” and “wife,” respectively; DB^i is a binary variable equal to one if i has a defined-benefit pension; $UNION^i$ is a binary variable equal to one if i belongs to a union at the defined-benefit job; $YRSV^i$ is the number of years that i stays in the defined-benefit job up to i 's retirement date; e_R is the household earnings in the last period of work; ϕ_R^h and ϕ_R^w indicate the fractions of e_R that belong to the husband and wife, respectively, with $\phi_R^h + \phi_R^w = 1$ by construction; and ξ is an error term that is assumed to be distributed as $N(0, \sigma_\xi^2)$.⁴⁰ Finally, the parameters to be estimated are $\beta_0^h, \beta_0^w, \beta_1^h, \beta_1^w, \beta_2^h, \beta_2^w, \gamma_0^h, \gamma_0^w, \gamma_1^h, \gamma_1^w, \gamma_2^h, \gamma_2^w$, and σ_ξ^2 .

The term db is calculated by assuming that the household receives annual defined-benefit pension benefits that are constant in *real* terms from the first period of retirement until none of the recipients survive. In particular, let $dbwealth$ be the observed present discounted value of db :

$$dbwealth = \sum_{j=R+1}^D \pi_j \frac{db}{\delta_j} \Rightarrow db = \frac{dbwealth}{\sum_{j=R+1}^D (\pi_j / \delta_j)},$$

where δ_j is the discount rate that converts pension benefits at age j into an equivalent value of 1992 dollars, π_j is the probability that the household will survive at age j conditional on surviving in the year that $dbwealth$ was reported, R is the last period of work, and D is a terminal age (the household will not live beyond this age). The estimation results are given in table A1.

⁴⁰ The specification is estimated with ordinary least squares using the White formula for the standard error.

TABLE A2
COEFFICIENT ESTIMATES FOR THE HOUSEHOLD AR(1) EARNINGS PROFILES

GROUP	COEFFICIENT ESTIMATES						Observations
	Group Constant	Age	.01 × Age ²	$\hat{\rho}$	$\hat{\sigma}$	R ²	
Single, no college	4.785*** (.022)	.231*** (.003)	-.259*** (.004)	.578	.456	.064	43,339
Single, college	3.803*** (.042)	.292*** (.007)	-.314*** (.009)	.682	.383	.179	8,677
Married, no college, one-earner	6.742*** (.018)	.174*** (.002)	-.195*** (.003)	.614	.318	.139	65,472
Married, no college, two-earner	5.165*** (.038)	.263*** (.006)	-.280*** (.007)	.699	.306	.287	15,779
Married, college, one-earner	6.743*** (.019)	.174*** (.003)	-.189*** (.004)	.672	.281	.182	56,482
Married, college, two-earner	5.014*** (.038)	.259*** (.007)	-.269*** (.009)	.759	.286	.258	14,626

NOTE.—The numbers of households for these groups are 1,873, 351, 2,076, 512, 1,821, and 519, respectively.
*** Statistically significant at the 1 percent level.

B. *Estimates of Household Earnings Expectations*

We construct household earnings as the summation of individual earnings for all adults in the household. The estimates for the model of household earnings described in the text are given in table A2.

C. *Out-of-Pocket Medical Expenses*

We construct household annual medical expenses on the basis of the (HRS-imputed) answers to the four medical expense questions asked in the 1998 and 2000 HRS. The four questions are as follows:

E10. About how much did you pay out-of-pocket for [nursing home/hospital/nursing home and hospital] bills [since R's LAST IW MONTH, YEAR/in the last two years]?

E18a. About how much did you pay out-of-pocket for [doctor/outpatient surgery/dental/doctor and outpatient surgery/doctor and dental/outpatient surgery and dental/doctor, outpatient surgery, and dental] bills [since R's LAST IW MONTH, YEAR/in the last two years]?

E21a. On the average, about how much have you paid out-of-pocket per month for these prescriptions [since R's LAST IW MONTH, YEAR/in the last two years]?

E24a. About how much did you pay out-of-pocket for [in-home medical care/special facilities or services/in-home medical care, special facilities or services] [since R's LAST IW MONTH, YEAR/in the last two years]?

We construct household annual medical expenses as one-half of E10 + E18a + 24 × E21a + E24a. The 1996 and 1997 household annual medical expenses are calculated from the information from the 1998 HRS; similarly, the 1998 and 1999 household annual medical expenses are calculated from the information

TABLE A3
COEFFICIENT ESTIMATES FOR THE AR(1) ANNUAL MEDICAL EXPENSES

GROUP	COEFFICIENT ESTIMATES						Observations
	Group Constant	Age	.01 × Age ²	$\hat{\rho}$	$\hat{\sigma}$	R ²	
Single, no college	-.003	.123***	-.074***	.859	2.081	.225	20,156
	.019	.003	.004				
Single, college	.114***	.140***	-.085***	.838	1.454	.355	3,136
	.041	.006	.009				
Married, no college	.056***	.175***	-.115***	.836	.930	.540	18,348
	.014	.002	.004				
Married, college	.027	.186***	-.123***	.840	.512	.705	5,704
	.018	.003	.005				

NOTE.—The numbers of households for these groups are 5,040, 784, 4,584, and 1,429, respectively.
*** Statistically significant at the 1 percent level.

from the 2000 HRS. The sample included is all households (HRS, AHEAD, CODA [Children of the Depression], and WB [War Baby]) that participated in and retained marital statuses between the 1998 and 2000 HRS. The estimation results, based on Prais-Winsten feasible generalized least squares, are given in table A3.

References

- Aguilar, Mark, and Erik Hurst. 2005. "Consumption versus Expenditure." *J.P.E.* 113 (October): 919–48.
- Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Q.J.E.* 109 (August): 659–84.
- Banks, James, Richard Blundell, and Sarah Tanner. 1998. "Is There a Retirement-Savings Puzzle?" *A.E.R.* 88 (September): 769–88.
- Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg. 2001. "What Accounts for the Variation in Retirement Wealth among U.S. Households?" *A.E.R.* 91 (September): 832–57.
- Blau, David. 2004. "Retirement and Consumption in a Life Cycle Model." Manuscript, Univ. North Carolina.
- Brown, Jeffrey R., and Amy Finkelstein. 2004. "The Interaction of Public and Private Insurance: Medicaid and the Long-Term Care Insurance Market." Working Paper no. 10989 (December), NBER, Cambridge, MA.
- Carroll, Christopher D. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Q.J.E.* 112 (February): 1–55.
- Chang, Sheng-Kai. 2002. "Simulation Estimation of Dynamic Panel Tobit Models with an Application to the Convergence of the Black-White Earnings Gap and Income Dynamics." Manuscript, Wayne State Univ.
- Citro, Constance F., and Robert T. Michael, eds. 1995. *Measuring Poverty: A New Approach*. Washington, DC: Nat. Acad. Press.
- Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica* 59 (September): 1221–48.
- Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes. 2004. "Do the Rich Save More?" *J.P.E.* 112 (April): 397–444.
- Engen, Eric M., William G. Gale, and Cori E. Uccello. 1999. "The Adequacy of Household Saving." *Brookings Papers Econ. Activity*, no. 2, pp. 65–165.

- Geweke, John, Michael P. Keane, and David Runkle. 1994. "Alternative Computational Approaches to Inference in the Multinomial Probit Model." *Rev. Econ. and Statis.* 76 (November): 609–32.
- Goldberger, Arthur S. 1991. *A Course in Econometrics*. Cambridge, MA: Harvard Univ. Press.
- Gourinchas, Pierre-Olivier, and Jonathan A. Parker. 2002. "Consumption over the Life Cycle." *Econometrica* 70 (January): 47–89.
- Gouveia, Miguel, and Robert P. Strauss. 1994. "Effective Federal Individual Income Tax Functions: An Exploratory Empirical Analysis." *Nat. Tax J.* 47 (June): 317–39.
- . 2000. "Effective Tax Functions for the U.S. Individual Income Tax: 1966–89." In *Proceedings of the 92nd Annual Conference on Taxation*, Atlanta, October 24–26. Washington, DC: Nat. Tax Assoc.
- Grad, Susan. 1990. "Earnings Replacement Rates of New Retired Workers." *Soc. Security Bull.* 53 (October): 2–19.
- Green Book. 1983. *Background Material and Data on Programs within the Jurisdiction of the Committee on Ways and Means*. Washington, DC: U.S. House of Representatives, Comm. Ways and Means.
- . 1988. *Background Material and Data on Programs within the Jurisdiction of the Committee on Ways and Means*. Washington, DC: U.S. House of Representatives, Comm. Ways and Means.
- Gruber, Jonathan, and Aaron Yelowitz. 1999. "Public Health Insurance and Private Savings." *J.P.E.* 107, no. 6, pt. 1 (December): 1249–74.
- Gustman, Alan L., and Thomas L. Steinmeier. 1999. "Effects of Pensions on Savings: Analysis with Data from the Health and Retirement Study." *Carnegie-Rochester Conf. Ser. Public Policy* 50 (June): 271–324.
- Haider, Steven J., and Melvin Stephens Jr. 2004. "Is There a Retirement-Consumption Puzzle? Evidence Using Subjective Retirement Expectations." Manuscript, Michigan State Univ.
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1995. "Precautionary Saving and Social Insurance." *J.P.E.* 103 (April): 360–99.
- Hurd, Michael D. 2003. "Bequests: By Accident or by Design?" In *Death and Dollars: The Role of Gifts and Bequests in America*, edited by Alicia H. Munnell and Annika Sunden. Washington, DC: Brookings Inst. Press.
- Hurd, Michael D., and Susann Rohwedder. 2003. "The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement." Working Paper no. 9586 (March), NBER, Cambridge, MA.
- Hurst, Erik. 2003. "Grasshoppers, Ants and Pre-retirement Wealth: A Test of Permanent Income Consumers." Manuscript, Univ. Chicago, Grad. School Bus.
- Hurst, Erik, and James P. Ziliak. 2006. "Do Welfare Asset Limits Affect Household Saving? Evidence from Welfare Reform." *J. Human Resources* 41 (Winter): 46–71.
- Kopczuk, Wojciech. 2005. "Bequest and Tax Planning: Evidence from Estate Tax Returns." Manuscript, Columbia Univ.
- Kotlikoff, Laurence J., Avia Spivak, and Lawrence H. Summers. 1982. "The Adequacy of Savings." *A.E.R.* 72 (December): 1056–69.
- McGrattan, Ellen R., and Edward C. Prescott. 2003. "Average Debt and Equity Returns: Puzzling?" *A.E.R. Papers and Proc.* 93 (May): 392–97.
- Modigliani, Franco, and Richard E. Brumberg. 1954. "Utility Analysis and the Consumption Function: An Interpretation of Cross Section Data." In *Post-*

- Keynesian Economics*, edited by Kenneth K. Kurihara. New Brunswick, NJ: Rutgers Univ. Press.
- Moffitt, Robert. 2002. "Documentation for Moffitt Welfare Benefits File." Manuscript (February 22), Johns Hopkins Univ. <http://www.econ.jhu.edu/People/Moffitt/DataSets.html>.
- Moore, James F., and Olivia S. Mitchell. 1998. "Projected Retirement Wealth and Savings Adequacy in the Health and Retirement Survey." Pension Research Council Working Paper no. 98-1, Univ. Pennsylvania, Wharton School. <http://prc.wharton.upenn.edu/prc/prc.html>.
- Palumbo, Michael G. 1999. "Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle." *Rev. Econ. Studies* 66 (April): 395–421.
- Samwick, Andrew A. 1998. "Discount Rate Heterogeneity and Social Security Reform." *J. Development Econ.* 57 (October): 117–46.
- Scholz, John Karl, Ananth Seshadri, and Surachai Khittrakun. 2004. "Are Americans Saving 'Optimally' for Retirement?" Working paper no. 10260 (February), NBER, Cambridge, MA.
- Tauchen, George. 1986. "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions." *Econ. Letters* 20 (2): 177–81.
- Venti, Steven F., and David A. Wise. 2004. "Aging and Housing Equity: Another Look." In *Perspectives on the Economics of Aging*, edited by David A. Wise. Chicago: Univ. Chicago Press (for NBER).
- Wysocki, Bernard, Jr. 1995. "Binge Buyers: Many Baby Boomer Save Little, May Run into Trouble Later On: They Don't Build Nest Eggs Nearly Rapidly Enough for an Easy Retirement." *Wall Street J.* (June 5), p. A1.

Appendix B from Scholz et al., “Are Americans Saving “Optimally” for Retirement?”

(JPE, vol. 114, no. 4, p. 607)

Imputing Earnings in the HRS

We have two problems with the earnings data that we address. For 77 percent of the 1992 HRS sample, we have access to each individual’s social security earnings records from 1951 to 1991. The social security earnings records report wage, salary, and self-employment income up to the earnings maximum (the earnings thresholds at which social security taxes are no longer taken from income). For 93 percent of the respondents with social security earnings records, we also have W-2 earnings records from 1980 to 1991. These W-2 records provide complete earnings information for wage and salary earners and the self-employed. The first difficulty is that 24 percent of positive social security earnings records are top-coded, and 40 percent of respondents with social security earnings records have at least one top-coded observation.

Our second problem is that 23 percent of respondents refused to grant access to social security earnings records. For these households we have self-reported earnings information for their current job (or the most recent job if not employed) and as many as three previous jobs. We need to estimate their earning profiles on the basis of their self-reported earnings information.

The goal is to use all available information to impute top-coded and missing earning observations and as a result obtain complete individual earnings histories. For the imputation, we proceed in two steps. First, on the basis of the social security and W-2 records, we estimate a dynamic-panel Tobit model to obtain individual earning processes. Then, conditional on all available earnings information, we use the estimates to impute the top-coded and missing observations.

A. Estimation

We start by describing our approach to estimating earnings for individuals with top-coded earnings.

For simplicity, suppose that we have earnings records of N individuals from time $t = 0$ to T , where 0 is the first period in which these individuals started working full-time. Assume for the moment that earnings are positive in each time period.⁴¹ Denote the logarithmic value of individual i ’s latent and observed earnings as $y_{i,t}^*$ and $y_{i,t}$, respectively. The relationship between the latent and observed earnings is

$$y_{i,t} = \begin{cases} y_{i,t}^* & \text{if } y_{i,t}^* < y_t^{tc} \\ y_t^{tc} & \text{if } y_{i,t}^* \geq y_t^{tc}, \end{cases}$$

where y_t^{tc} is the logarithmic value of the social security maximum taxable earnings at time t .

The individual log earnings process is specified as

$$\begin{aligned} y_{i,0}^* &= \mathbf{x}_{i,0}'\beta_0 + \epsilon_{i,0}, \\ y_{i,t}^* &= \rho y_{i,t-1}^* + \mathbf{x}_{i,t}'\beta + \epsilon_{i,t}, \quad t \in \{1, 2, \dots, T\}, \\ \epsilon_{i,t} &= \alpha_i + u_{i,t}, \end{aligned} \tag{B1}$$

where $\mathbf{x}_{i,t}$ is the vector of i ’s characteristics at time t , and the error term $\epsilon_{i,t}$ includes an individual-specific component α_i , which is constant over time and known to the individual before time 0, and the unanticipated

⁴¹ Generalizing this to the case in which the earnings series begins after time 0 and the case in which some earnings observations are zero is straightforward but detail-oriented, so we omit the discussion. We did treat these cases in practice, however.

white-noise component, $u_{i,t}$. Notice that parameters β_0 and β are allowed to be different. In the following analysis, we employ random-effect assumptions with homoskedastic errors.

ASSUMPTION 1. $\alpha_i | \underline{x}_i \sim \text{iid } N(0, \sigma_\alpha^2)$.

ASSUMPTION 2. $u_{i,t} | \underline{x}_i \sim \text{iid } N(0, \sigma_u^2)$ for all t .

ASSUMPTION 3.

$$E[u_{i,t} | \underline{x}_i, \alpha_i] = 0, E[u_{i,t}^2 | \underline{x}_i, \alpha_i] = \sigma_u^2, E[u_{i,t} u_{i,k} | \underline{x}_i, \alpha_i] = 0 \quad \forall t \in \{0, 1, \dots, T\}, t \neq k,$$

where $\underline{x}_i \equiv (x_{i,0}, x_{i,1}, \dots, x_{i,T})$.

These three assumptions imply that

$$\epsilon_i = (\epsilon_{i,0}, \epsilon_{i,1}, \dots, \epsilon_{i,T})' \sim N(0, \Sigma), \quad (\text{B2})$$

where

$$\Sigma = \begin{bmatrix} \sigma_{0,0}^2 & \sigma_{0,1} & \cdots & \sigma_{0,T} \\ \sigma_{1,0} & \sigma_{1,1}^2 & \cdots & \sigma_{1,T} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{T,0} & \sigma_{T,1} & \cdots & \sigma_{T,T}^2 \end{bmatrix},$$

with $\sigma_{j,k}^2 = \sigma_\alpha^2 + \sigma_u^2$ for $j = k$, and $\sigma_{j,k}^2 = \sigma_\alpha^2$ otherwise. Our goal here is to obtain consistent estimates of the true parameters $\theta^* = (\beta, \beta_0, \rho, \sigma_\alpha^2, \sigma_u^2)$. We do this by maximum likelihood.

To construct the likelihood function for each individual's earnings series, notice that we can write the joint probability density function of each pair of random variables $(y_{i,t}, y_{i,t}^*)$ as

$$g(y_{i,t}, y_{i,t}^* | y_{i,t-1}, y_{i,t-1}^*, y_{i,t-2}, y_{i,t-2}^*, \dots, y_{i,0}, y_{i,0}^*; \underline{x}_i, \theta).$$

From the AR(1) assumption on earnings made in (B1), it follows that

$$g(y_{i,t}, y_{i,t}^* | y_{i,t-1}, y_{i,t-1}^*, y_{i,t-2}, y_{i,t-2}^*, \dots, y_{i,0}, y_{i,0}^*; \underline{x}_i, \theta) = g(y_{i,t}, y_{i,t}^* | y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta).$$

In other words, of all the information about past realized and observed earnings, only information from the previous period matters. As a special case, the conditional likelihood of the pair $(y_{i,t}, y_{i,t}^*)$ is $g_0(y_{i,0}, y_{i,0}^* | \underline{x}_i, \theta)$ because there is no information about earnings before period 0.

Applying Bayes' rules to the density $g(\cdot)$, we have

$$g(y_{i,t}, y_{i,t}^* | y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta) = h(y_{i,t}^* | y_{i,t}, y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta) q(y_{i,t} | y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta), \quad (\text{B3})$$

where the density for the log of observed earnings conditional on the past information is a conventional Tobit likelihood function

$$q(y_{i,t} | y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta) = [f(y_{i,t} | y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta)]^{1(y_{i,t}^* < y_{i,t}^c)} [\Pr(y_{i,t}^* \geq y_{i,t}^c | y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta)]^{1(y_{i,t}^* \geq y_{i,t}^c)},$$

where $f(\cdot)$ and $\Pr(\cdot)$ are a probability density and a cumulative distribution function, respectively, and the conditional density $h(\cdot)$ for noncensored observations is the probability mass function

$$h(y_{i,t}^* | y_{i,t} < y_{i,t}^c, y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta) = \begin{cases} 1 & \text{if } y_{i,t}^* = y_{i,t} \\ 0 & \text{if } y_{i,t}^* \neq y_{i,t} \end{cases}$$

and the conditional density is simply $h(y_{i,t}^* | y_{i,t} = y_{i,t}^c, y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta)$ for censored observations.

Similarly, we can write

$$g_0(y_{i,0}, y_{i,0}^* | \underline{x}_i, \theta) = h_0(y_{i,0}^* | y_{i,0}; \underline{x}_i, \theta) q_0(y_{i,0} | \underline{x}_i, \theta),$$

where the conditional density $h_0(\cdot)$ for noncensored observations is the probability mass function

$$h_0(y_{i,0}^* | \underline{x}_i, \theta) = \begin{cases} 1 & \text{if } y_{i,0}^* = y_{i,0} \\ 0 & \text{if } y_{i,0}^* \neq y_{i,0} \end{cases}$$

and the conditional density is $h_0(y_{i,0}^* | y_{i,0} = y_0^{tc}; \underline{x}_i, \theta)$ for censored observations. In addition, the density for the time 0 log of observed earnings conditional on known information is

$$q_0(y_{i,0} | \underline{x}_i, \theta) = [f(y_{i,0} | \underline{x}_i, \theta)]^{1(y_{i,0}^* < y_0^{tc})} [\Pr(y_{i,0}^* \geq y_0^{tc} | \underline{x}_i, \theta)]^{1(y_{i,0}^* \geq y_0^{tc})}.$$

From (B1) and (B2), it is apparent that the functions $f(y_{i,0} | \underline{x}_i, \theta)$ and $f(y_{i,t} | y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta)$, $t > 0$, are normal probability density functions and $\Pr(y_{i,0}^* \geq y_0^{tc} | \underline{x}_i, \theta)$ and $\Pr(y_{i,t}^* \geq y_t^{tc} | y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta)$, $t > 0$, are normal cumulative distribution functions. For expositional convenience, define

$$\begin{aligned} h^{tc}(0; \underline{x}_i, \theta) &\equiv h_0(y_{i,0}^* | y_{i,0} = y_0^{tc}; \underline{x}_i, \theta), \\ h^{tc}(t; \underline{x}_i, \theta) &\equiv h(y_{i,t}^* | y_{i,t} = y_t^{tc}, y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta), \quad t > 0, \\ q(0; \underline{x}_i, \theta) &\equiv q_0(y_{i,0}; \underline{x}_i, \theta), \\ q(t; \underline{x}_i, \theta) &\equiv q(y_{i,t} | y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta), \quad t > 0. \end{aligned}$$

The likelihood function for i 's series of observed log earnings is

$$\begin{aligned} L_i(y_{i,T}, y_{i,T-1}, \dots, y_{i,1}, y_{i,0}; \underline{x}_i, \theta) &= \int_{y_{i,T}^*}^{\infty} \dots \int_{y_{i,0}^*}^{\infty} \left[\prod_{t=0}^T q(t; \underline{x}_i, \theta) \right] \cdot \left[h(y_{i,0}^*; \underline{x}_i, \theta) \prod_{t=1}^T h(y_{i,t}^* | y_{i,t}, y_{i,t-1}, y_{i,t-1}^*; \underline{x}_i, \theta) \right] dy_{i,0}^* \dots dy_{i,T}^* \\ &= \int_{y_{i,c_n}^*}^{\infty} \dots \int_{y_{i,c_1}^*}^{\infty} \left[\prod_{t=0}^T q(t; \underline{x}_i, \theta) \right] \cdot \left[\prod_{k=c_1}^{c_n} h^{tc}(k; \underline{x}_i, \theta) \right] dy_{i,c_1}^* \dots dy_{i,c_n}^* \\ &= E_{y_{i,c_1}^*, \dots, y_{i,c_n}^*} \left[\prod_{t=0}^T q(t; \underline{x}_i, \theta) \right], \end{aligned} \tag{B4}$$

where c_1, c_2, \dots, c_n are the periods in which the log observed earnings are censored, that is, equal to their corresponding top-coded limits. Notice that, since we do not observe $y_{i,t}^*$ when it is censored, we integrate out $y_{i,t}^*$ for censored observations. Unfortunately, the integration does not yield any analytical solution, nor is direct numerical evaluation of the integral computationally feasible in this case. As an alternative, Chang (2002) proposes using a Geweke-Hajivassiliou-Keane (GHK) (probit) simulator to deal with the computational burden of the integration.⁴² The estimation results are given in table B1.

B. Imputation

The idea is to impute top-coded and missing earnings observations with their conditional expectations, where the conditioning variables include both the individual's characteristics and observed earnings. The conditional expectations are calculated numerically on the basis of the dynamic earnings model (B1) and the distributional assumption (B2). The imputation scheme is similar for top-coded and missing observations; therefore, we discuss only the scheme for top-coded observations here.⁴³

⁴² The GHK simulator gives a numerical approximation of a probit probability of interest. It is a popular choice of probit simulators because of its relative accuracy; see Geweke, Keane, and Runkle (1994) for details.

⁴³ We can think of a missing earnings observation as an observation with a top-coded value of zero, which is equivalent to saying that we know nothing about earnings in that period (as opposed to the case in which we observe top-coded earnings and know that the actual earnings are at least as

To be concrete, notice that (B1) implies that

$$\begin{aligned} E[y_0^* | \underline{y}, \underline{x}, \theta] &= x'_0 \beta_0 + E[\epsilon_0 | \underline{y}, \underline{x}, \theta], \\ E[y_t^* | \underline{y}, \underline{x}, \theta] &= \rho E[y_{t-1}^* | \underline{y}, \underline{x}, \theta] + x'_t \beta + E[\epsilon_t | \underline{y}, \underline{x}, \theta], \quad t \in \{1, 2, \dots, T\}, \end{aligned} \quad (\text{B5})$$

where $\underline{y} = (y_0, y_1, \dots, y_T)$ is the series of individual i 's log of observed earnings (the individual subscript i is omitted throughout this subsection). By construction, given information about the individual's characteristics and observed earnings, $E[y_t^* | \underline{y}, \underline{x}, \theta]$ is on average the best guess for y_t^* . In other words, for top-coded observations, equation (B1) suggests the imputed values

$$y_t^{\text{imp}} = E[y_t^* | y_t = y_t^{tc}; \underline{y}, \underline{x}, \theta],$$

which requires knowledge of $E[\epsilon_t | y_t = y_t^{tc}; \underline{y}, \underline{x}, \theta]$ for every period t in which $y_t = y_t^{tc}$. The analytical form of $E[\epsilon_t | y_t = y_t^{tc}; \underline{y}, \underline{x}, \theta]$ is not available in our case; therefore, we calculate this object numerically using the Gibbs sampling procedure.⁴⁴

To facilitate the discussion about details of the procedure, denote $\underline{\epsilon}_{<t} = (\epsilon_0, \epsilon_1, \dots, \epsilon_{t-1})'$, $\underline{\epsilon}_{>t} = (\epsilon_{t+1}, \epsilon_{t+2}, \dots, \epsilon_T)'$, and $\underline{\epsilon}_{-t} = (\epsilon_0, \epsilon_1, \dots, \epsilon_{t-1}, \epsilon_{t+1}, \dots, \epsilon_T)'$ for any vector $\underline{\epsilon} = (\epsilon_0, \epsilon_1, \dots, \epsilon_T)'$. Here, we want to simulate R sets of $\underline{\epsilon}$ that are consistent with the observed \underline{y} and \underline{x} given θ . The Gibbs sampling procedure does this in two steps for each round of simulation.

1. In the r th round of simulation, $r = 1, 2, \dots, R$, generate a “random” initial value $\underline{\epsilon}_0^{(r)} = (\epsilon_{0,0}^{(r)}, \epsilon_{1,0}^{(r)}, \dots, \epsilon_{t,0}^{(r)}, \dots, \epsilon_{T,0}^{(r)})$ that satisfies (B1) given $\underline{y}, \underline{x}$, and θ . Notice that $\epsilon_{t,0}^{(r)}$ is not identified when $y_t = y_t^{tc}$. In this case, $\epsilon_{t,0}^{(r)}$ is chosen randomly under the restriction that $y_{t,0}^{*(r)} \equiv (y_t^* | \underline{\epsilon}_{<t} = \underline{\epsilon}_{<t,0}^{(r)}, \epsilon_t = \epsilon_{t,0}^{(r)}; \underline{x}, \theta) \geq y_t^{tc}$. If $y_{t-1} < y_{t-1}^{tc}$ and $y_t < y_t^{tc}$, $\epsilon_{t,0}^{(r)}$ is defined by (B1); that is, it is actually not random.

2. Starting with $m = 1$, draw a random number $\epsilon_{t,m}^{(r)}$ from $t = 0, \dots, T$ from the distribution of $\epsilon_t | \underline{\epsilon}_{-t}; \underline{x}, \theta$ such that

$$y_{t,m}^{*(r)} = (y_t^* | \underline{\epsilon}_{>t} = \underline{\epsilon}_{>t,m-1}^{(r)}, \epsilon_{<t} = \epsilon_{<t,m}^{(r)}, \epsilon_t = \epsilon_{t,m}^{(r)}; \underline{x}, \theta) = y_t \quad \text{if } y_t < y_t^{tc}$$

and

$$y_{t,m}^{*(r)} = (y_t^* | \underline{\epsilon}_{>t} = \underline{\epsilon}_{>t,m-1}^{(r)}, \epsilon_{<t} = \epsilon_{<t,m}^{(r)}, \epsilon_t = \epsilon_{t,m}^{(r)}; \underline{x}, \theta) \geq y_t \quad \text{if } y_t = y_t^{tc}.$$

(This is equivalent to drawing $\epsilon_{t,m}^{(r)}$ from $\epsilon_t | \underline{\epsilon}_{-t}; \underline{y}, \underline{x}, \theta$.) Then, continue from $m = 2$ to $m = M$.

With $\underline{\epsilon}_M^{(r)}$, $r = 1, 2, \dots, R$, an estimate of $E[\epsilon_t | \underline{y}, \underline{x}, \theta]$ is

$$\hat{E}[\epsilon_t | \underline{y}, \underline{x}, \theta] = \frac{1}{R} \sum_{r=1}^R \epsilon_{t,M}^{(r)}. \quad (\text{B6})$$

Given the estimate $\hat{E}[\epsilon_t | \underline{y}, \underline{x}, \theta]$, we calculate the imputed value of earnings as

$$y_0^{\text{imp}} = x'_0 \beta_0 + \hat{E}[\epsilon_0 | \underline{y}, \underline{x}, \theta]$$

and

$$y_t^{\text{imp}} = \rho y_{t-1}^{\text{imp}} + x'_t \beta + \hat{E}[\epsilon_t | \underline{y}, \underline{x}, \theta], \quad t \in \{1, 2, \dots, T\}. \quad (\text{B7})$$

Notice that, by construction,

large as the top-coded earnings).

⁴⁴ Briefly, Gibbs sampling is a procedure to draw a set of numbers randomly from a (valid) joint distribution. Then, the random draws are used to estimate properties of any marginal distribution of interest, which is difficult to derive analytically from the joint distribution. The procedure relies on the law of large numbers, i.e., that moments of a distribution can be estimated consistently from a set of random draws from that distribution.

$$y_0^{\text{imp}} = \frac{1}{R} \sum_{r=1}^R y_{0,M}^{*(r)}$$

and

$$y_t^{\text{imp}} = \frac{1}{R} \sum_{r=1}^R y_{t,M}^{*(r)}, \quad t \in \{1, 2, \dots, T\},$$

and that $y_t^{\text{imp}} = y_t$ if $y_t < y_t^{tc}$ and $y_t^{\text{imp}} \geq y_t$ if $y_t = y_t^{tc}$.

The remaining parts of this subsection (i) construct the functional form for the conditional distribution of $\epsilon_t | \underline{\epsilon}_{-t}; \underline{x}, \theta$ and (ii) show how to draw a random number $\epsilon_{t,m}^{(r)}$ from this conditional distribution to satisfy (B1) given $\underline{y}, \underline{x}$, and θ . More notation is required. For any matrix Σ , denote $\Sigma_{t,t}$ as the element of Σ on the t th row and t th column, $\Sigma_{t,-t}$ as the t th row of Σ with the element $\Sigma_{t,t}$ removed, $\Sigma_{-t,t}$ as the t th column of Σ with the element $\Sigma_{t,t}$ removed, and $\Sigma_{-t,-t}$ as the matrix Σ with the t th row and t th column removed.

Recall the property of a joint-normal vector that

$$\begin{aligned} \underline{\epsilon} &\sim N(E[\underline{\epsilon}], \Sigma) \Rightarrow \epsilon_t | \underline{\epsilon}_{-t} \sim N(\mu_{t|-t}, \Sigma_{t|-t}), \\ \mu_{t|-t} &= E[\epsilon_t] + \Sigma_{t,-t} \Sigma_{-t,-t}^{-1} \{\underline{\epsilon}_{-t} - E[\underline{\epsilon}_{-t}]\}, \\ \Sigma_{t|-t} &= \Sigma_{t,t} - \Sigma_{t,-t} \Sigma_{-t,-t}^{-1} \Sigma_{-t,t} \end{aligned} \quad (\text{B8})$$

(see, e.g., Goldberger 1991, 196–97). Recall from (B2) that $E[\epsilon_j] = 0$ and $\Sigma = (1 - \rho)\sigma_\epsilon^2 I_{T+1} + \rho\sigma_\epsilon^2 1_{T+1}$, where 1_{T+1} is a $(T + 1) \times (T + 1)$ matrix whose elements are all 1. Thus $\mu_{t|-t} = \Sigma_{t,-t} \Sigma_{-t,-t}^{-1} \underline{\epsilon}_{-t}$, and $\Sigma_{t,-t} \Sigma_{-t,-t}^{-1} = \Sigma_{s,-s} \Sigma_{-s,-s}^{-1}$ and $\Sigma_{t|-t} = \Sigma_{s|-s}$ for any $t = 0, \dots, T$ and $s = 0, \dots, T$.

Recall that we draw a value for $\epsilon_{t,m}^{(r)}$ randomly from the conditional distribution $\epsilon_t | \underline{\epsilon}_{-t}$ such that, given $\underline{\epsilon}_{>t,m-1}^{(r)}$, $\epsilon_{<t,m}^{(r)}$, $\epsilon_{t,m}^{(r)}$, \underline{x} , and θ ,

$$y_{t,m}^{*(r)} \begin{cases} = y_t & \text{if } y_t < y_t^{tc} \\ \in [y_t^{tc}, \infty) & \text{if } y_t = y_t^{tc}. \end{cases} \quad (\text{B9})$$

In practice, it is more convenient to work with the standard-normal transformation of $\epsilon_t | \underline{\epsilon}_{-t}$,

$$z_{t|-t} \equiv \frac{(\epsilon_t | \underline{\epsilon}_{-t}) - \mu_{t|-t}}{\sigma_{t|-t}} \sim N(0, 1), \quad \sigma_{t|-t} = \sqrt{\Sigma_{t|-t}}. \quad (\text{B10})$$

From (B1), $\epsilon_0 | \underline{\epsilon}_{-0} = (y_0^* | \underline{\epsilon}_{-0}) - x_0' \beta_0$ and $\epsilon_t | \underline{\epsilon}_{-t} = (y_t^* - \rho y_{t-1}^* | \underline{\epsilon}_{-t}) - x_t' \beta$, $t \in \{1, 2, \dots, T\}$. Also, since $(y_t^* | \underline{\epsilon}_{-t}; \underline{x}, \theta) = (y_t^* | \underline{\epsilon}_{<t}; \underline{x}, \theta)$, $(y_{t-1}^* | \underline{\epsilon}_{>t} = \underline{\epsilon}_{>t,m-1}^{(r)}, \underline{\epsilon}_{<t} = \underline{\epsilon}_{<t,m}^{(r)}; \underline{x}, \theta) = y_{t-1,m}^{*(r)}$.

Thus, with the transformation (B10), drawing $\epsilon_{t,m}^{(r)}$ from (B8) to satisfy (B9) is equivalent to drawing $z_{t|-t,m}^{(r)}$ such that

$$z_{0|-0,m}^{(r)} = \begin{cases} \frac{y_0 - x_0' \beta_0 - \mu_{0|-0,m}^{(r)}}{\sigma_{0|-0}} & \text{if } y_0 < y_0^{tc} \\ \Phi^{-1} \left(\xi_{0,m}^{(r)} + [1 - \xi_{0,m}^{(r)}] \Phi \left(\frac{y_0^{tc} - x_0' \beta_0 - \mu_{0|-0,m}^{(r)}}{\sigma_{0|-0}} \right) \right) & \text{if } y_0 = y_0^{tc} \end{cases}$$

for $t = 0$, and

$$z_{t|t-m}^{(r)} = \begin{cases} \frac{y_t - \rho y_{t-1,m}^{*(r)} - x_t' \beta - \mu_{t|t-m}^{(r)}}{\sigma_{t|t-m}} & \text{if } y_t < y_t^{tc} \\ \Phi^{-1} \left(\xi_{t,m}^{(r)} + [1 - \xi_{t,m}^{(r)}] \Phi \left(\frac{y_t^{tc} - \rho y_{t-1,m}^{*(r)} - x_t' \beta - \mu_{t|t-m}^{(r)}}{\sigma_{t|t-m}} \right) \right) & \text{if } y_t = y_t^{tc} \end{cases}$$

for $t > 0$,⁴⁵ with

$$y_{0,m}^{*(r)} = x_0' \beta_0 + [\sigma_{0|-0} z_{0|-0,m}^{(r)} + \mu_{0|-0,m}^{(r)}]$$

and

$$y_{t,m}^{*(r)} = \rho y_{t-1,m}^{*(r)} + x_t' \beta + [\sigma_{t|-t} z_{t|-t,m}^{(r)} + \mu_{t|-t,m}^{(r)}], \quad t \in \{1, 2, \dots, T\},$$

where

$$\mu_{t|t-m}^{(r)} = \Sigma_{t,-t} \Sigma_{-t,-t}^{-1} \begin{pmatrix} \epsilon_{<,m}^{(r)} \\ \epsilon_{>,m-1}^{(r)} \end{pmatrix},$$

$\epsilon_{t,m}^{(r)} = \sigma_{t|t-m} z_{t|t-m}^{(r)} + \mu_{t|t-m}^{(r)}$, $\sigma_{t|t-m} = \sqrt{\Sigma_{t|t-m}}$ and $\Sigma_{t|t-m} = \Sigma_{t,t} - \Sigma_{t,-t} \Sigma_{-t,-t}^{-1} \Sigma_{-t,t}$, and $\xi_{t,m}^{(r)}$ is a random draw from a $[0, 1]$ uniform distribution. Notice that $y_{t,m}^{*(r)} = y_t$ if $y_t < y_t^{tc}$ and $y_{t,m}^{*(r)} \geq y_t$ if $y_t = y_t^{tc}$ by construction.

C. Social Security Function

From the expected earnings profiles, we can calculate the *lifetime* summation of household earnings up to the year of retirement as $E_R \equiv \sum_{j=S}^R e_j$, where e_j denotes the household earnings at age j in a common base year unit, and S and R denote the first and the last working ages, respectively.⁴⁶ Denote $\bar{\phi}^h$ and $\bar{\phi}^w$ as the fractions of E_R that are contributed by the husband and wife of the household, respectively.⁴⁷ On the basis of E_R , $\bar{\phi}^h$, and $\bar{\phi}^w$, we can approximate the household annual social security benefits as follows.

a. Calculate *Individual Primary Insurance Amount (PIA)*⁴⁸

Individual i 's annual indexed monthly earnings (AIME) can be approximated as

$$\text{AIME}^i \approx \frac{\bar{\phi}^i E_R}{L^i}, \quad (\text{B11})$$

with

⁴⁵ To see how this works, note first that for $\epsilon \sim N(\mu, \sigma^2)$, $f(\epsilon) = (2\pi\sigma^2)^{-1/2} \exp[-0.5(\epsilon - \mu)/\sigma]$. Define $z \equiv (\epsilon - \mu)/\sigma$. It follows that $F(\epsilon) = \Phi(z(\epsilon))$, where Φ is the standard normal cumulative distribution function. Thus

$$\Phi(z(\epsilon^{(r)})) = F(\epsilon^{(r)}) = \xi_t^{(r)} + [1 - \xi_t^{(r)}] F(\epsilon^{tc}) = \xi_t^{(r)} + [1 - \xi_t^{(r)}] \Phi(z(\epsilon^{tc})).$$

In other words, drawing $\epsilon^{(r)}$ from a truncated distribution of $\{\epsilon | \epsilon \geq \epsilon^{tc}\}$ is equivalent to drawing $z^{(r)} = z(\epsilon^{(r)})$ from a truncated distribution of $\{z | z \geq z^{tc}\}$ and then transforming $z^{(r)}$ back to $\epsilon^{(r)}$.

⁴⁶ As opposed to a *discounted* present value of earnings, the summation is a straightforward summation of earnings in a common base year currency unit, which is the concept employed by the Social Security Administration.

⁴⁷ The terminologies “husband” and “wife” are not literal. In particular, we call a single male respondent “husband” and a single female respondent “wife.” Without this simplification, we need separate treatments for married and single households. Under this generalization, $\bar{\phi}_i^h = 1$ and $\bar{\phi}_i^w = 0$ for single-male households, and $\bar{\phi}_i^h = 0$ and $\bar{\phi}_i^w = 1$ for single-female households.

⁴⁸ Social security benefits derived from the calculations in this section are not precise because the calculated AIME may be smaller than the actual AIME and, conditional on AIME being correctly calculated, the calculated household benefits may be larger than the actual ones. For the former, the reasons are that (i) we do not exclude the five years of lowest earnings from the calculation, (ii) we use base year (i.e., real) values of earnings after age 60 instead of nominal values, and (iii) we do not take into account earnings in retirement if respondents work beyond their household retirement dates. For the latter, the reason is that we assume that both the husband and the wife of a married household are eligible to collect benefits at the household retirement date. If one of them is not eligible at the retirement date, the approximation will overstate the benefits. Nevertheless, by virtue of having complete earnings histories for most individuals, our calculations are considerably more accurate than those in other life cycle simulation models of wealth accumulation.

$$L^i = 12 \times \max \{R^i - 22, 40\},$$

where $i = h$ (husband) or w (wife), and L^i is the number of months of i 's covered period.⁴⁹ Without loss of generality, we set $L^w = 40$ for single-male households and $L^h = 40$ for single-female households.

Individual PIA can be calculated as

$$\begin{aligned} \text{PIA}^i &= 0.90 \times \min \{AIME^i, b_0\} + 0.32 \times \min \{\max \{AIME^i - b_0, 0\}, b_1 - b_0\} \\ &+ 0.15 \times \max \{AIME^i - b_1, 0\}, \end{aligned} \quad (\text{B12})$$

where b_0 and b_1 are the bend points. For the 1992 formula, $b_0 = \$387$ and $b_1 = \$2,333$.

b. Calculate Household Annual Social Security Benefits

First, the *individual* monthly social security benefits are calculated as

$$ssb^i = \max \{d_{\text{own}}^i \text{PIA}^i, d_{\text{spouse}}^i \text{PIA}^{i\text{spouse}}, ssx^i\}, \quad (\text{B13})$$

where i 's spouse = h (w) if $i = w$ (h), d_{own}^i is the fraction of i 's PIA that i would get if i collected benefits based on i 's PIA, d_{spouse}^i is the fraction of PIA of i 's spouse that i would get if i collected benefits based on PIA of i 's spouse, and ssx^i is the monthly benefit that i would get if i collected benefits based on the PIA of i 's ex-spouse.⁵⁰ Without loss of generality, for single-male households,

$$d_{\text{spouse}}^h = d_{\text{own}}^w = d_{\text{spouse}}^w = ssx^w = 0$$

and

$$d_{\text{spouse}}^w = d_{\text{own}}^h = d_{\text{spouse}}^h = ssx^h = 0$$

for single-female households. In addition, we set $ssx^h = ssx^w = 0$ for married households because we do not have any information to determine ssx^i . Similarly, $ssx^i = 0$ for any single households without information to determine their ex-spouses' PIA.

Finally, household i 's *annual* social security benefits can be approximated as

$$ss_i = 12 \times (ssb_i^h + ssb_i^w), \quad (\text{B14})$$

which, for a married household, is the benefits the household would get when both the husband and wife survive. When one of the spouses in a married household dies, the *annual* social security benefits of the surviving spouse are

$$ss_i^{\text{survive}} = 12 \times \max \{d_{\text{own}}^h \text{PIA}^h, d_{\text{own}}^w \text{PIA}^w\}. \quad (\text{B15})$$

In other words, we approximate the surviving spouse's benefits to be the higher of the husband's and wife's benefits that they would be able to collect on the basis of their own earning histories (which determine their PIAs) and the household retirement date (which determines the factors d). This approximates the actual guideline of the Social Security Administration.

⁴⁹ Without the lower bound of 40 years in the max operator ($\max \{R^i - 22, 40\}$), AIME would be too high for households whose members retire before age 62. In addition, notice that we use the *household* retirement date (R^i) rather than the *individual* retirement date.

⁵⁰ To recover the ex-spouse's PIA, we first compute the benefit amount that a single respondent would get based on her own earning history. Then, we compare the amount to the reported amount of social security benefits in the first wave in which the respondent reported collecting the benefits. If the reported benefit amount is higher, we assume that the single respondent collected benefits based on her ex-spouse's records, and the reported amount is used to recover her ex-spouse's PIA.

TABLE B1
Estimation Results of Individual Earnings Processes

VARIABLE	SAMPLE			
	Male without College	Male with College	Female without College	Female with College
A. Initial Observations ($t = 0$)				
Constant	5.691*** (.204)	3.910*** (.444)	7.407*** (.133)	4.738*** (.434)
Race (white = 1, 0 otherwise)	.282*** (.033)	.330 (.305)	.135*** (.029)	.006 (.087)
Years of schooling/professional postgraduate degree dummy	.044*** (.005)	-.135 (.209)	.017*** (.006)	.058 (.138)
No high school dummy/ postgraduate degree dummy	-.005 (.003)	-.120 (.149)	-.162*** (.041)	.065 (.086)
Marital status in 1992 HRS	.107*** (.030)	.034 (.042)	-.032* (.017)	.033 (.278)
Two-earner household dummy	-.047* (.025)	-.145** (.072)	.164*** (.028)	.187 (.155)
Age	.156*** (.013)	.291*** (.024)	.057*** (.009)	.247*** (.028)
.01 × Age ²	-.177*** (.020)	-.330*** (.038)	-.056*** (.013)	-.307*** (.042)
B. Subsequent Observations ($t > 0$)				
Constant	2.642*** (.042)	2.066*** (.156)	3.296*** (.071)	2.917*** (.141)
Race (white = 1, 0 otherwise)	.106*** (.010)	.061** (.025)	.040*** (.010)	-.009 (.091)
Years of schooling/professional postgraduate degree dummy	.020*** (.002)	.033 (.028)	.025*** (.003)	.066* (.035)
No high school dummy/ postgraduate degree dummy	-.037*** (.010)	.002 (.007)	-.069*** (.018)	.077*** (.019)
Marital status in 1992 HRS	.114*** (.012)	.129** (.056)	-.141*** (.014)	-.166*** (.048)
Two-earner household dummy	-.048*** (.007)	-.078*** (.016)	.210*** (.012)	.197*** (.032)
Age	.039*** (.002)	.032*** (.005)	.018*** (.002)	.019*** (.005)
.01 × Age ²	-.047*** (.002)	-.037*** (.006)	-.012*** (.003)	-.015*** (.006)
Earnings at the previous period	.629*** (.006)	.737*** (.012)	.568*** (.008)	.659*** (.017)
Variance of the individual-specific effect (σ_v^2)	.029*** (.002)	.012** (.005)	.041*** (.003)	.022*** (.004)
Variance of the gross error term (σ_ϵ^2)	.213*** (.004)	.223*** (.011)	.240*** (.005)	.213*** (.011)
Number of individual-year observations	86,382	21,286	47,145	8,609
Number of respondents	2,914	720	2,576	446

Note.— The dependent variable is the respondents’ natural log earnings. For samples with at most high school, the education variables are (i) years of schooling and (ii) no high school dummy. For samples with at least a bachelor’s degree, the education variables are (i) professional postgraduate degree dummy (master of business administration, doctor of law, doctor of medicine, or doctor of philosophy) and (ii) postgraduate degree dummy. Standard errors are in parentheses.

* Statistically significant at the 10 percent level.

** Statistically significant at the 5 percent level.

*** Statistically significant at the 1 percent level.