Why Do the Elderly Save? The Role of Medical Expenses
Author(s): Mariacristina De Nardi, Eric French, and John B. Jones
Source: *Journal of Political Economy*, Vol. 118, No. 1 (February 2010), pp. 39-75
Published by: The University of Chicago Press
Stable URL: http://www.jstor.org/stable/10.1086/651674
Accessed: 22/10/2015 16:58

Why Do the Elderly Save? The Role of Medical Expenses

Mariacristina De Nardi

Federal Reserve Bank of Chicago and National Bureau of Economic Research

Eric French

Federal Reserve Bank of Chicago

John B. Jones

University at Albany, State University of New York

This paper constructs a model of saving for retired single people that includes heterogeneity in medical expenses and life expectancies, and bequest motives. We estimate the model using Assets and Health Dynamics of the Oldest Old data and the method of simulated moments. Out-of-pocket medical expenses rise quickly with age and permanent income. The risk of living long and requiring expensive medical care is a key driver of saving for many higher-income elderly. Social insurance programs such as Medicaid rationalize the low asset holdings of

This paper was previously circulated under the title "Differential Mortality, Uncertain Medical Expenses, and the Saving of Elderly Singles." For helpful comments and suggestions, we thank an editor and two referees, Jerome Adda, Kartik Athreya, Gadi Barlevy, Marco Bassetto, Marco Cagetti, Jeff Campbell, Chris Carroll, Michael Hurd, Helen Koshy, Nicola Pavoni, Monika Piazzesi, Luigi Pistaferri, Victor Rios-Rull, Tom Sargent, Karl Scholz, and seminar participants at many institutions. Olga Nartova, Kenley Pelzer, Phil Doctor, Charles Doss, and Annie Fang Yang provided excellent research assistance. De Nardi gratefully acknowledges financial support from National Science Foundation grant SES-0317872. Jones gratefully acknowledges financial support from National Institute on Aging grant 1R03AG026299. French thanks the Social Security Administration for hospitality while writing the paper. The views expressed in this paper are those of the authors and not necessarily those of the Federal Reserve Bank of Chicago, the Federal Reserve System, the Social Security Administration, the National Science Foundation, or the National Institute on Aging.

[*Journal of Political Economy*, 2010, vol. 118, no. 1]

the poorest but also benefit the rich by insuring them against high medical expenses at the ends of their lives.

I. Introduction

Many elderly keep large amounts of assets until very late in life. Furthermore, the income-rich run down their assets more slowly than the income-poor. Why is this the case?

To answer this question, we estimate a life cycle model of saving on a sample of single, retired elderly individuals. The key elements in our framework are risky and heterogeneous medical expenses, risky and heterogeneous life expectancy, a government-provided minimum consumption level (or "consumption floor," which is income and asset tested), and bequest motives.

Our main result is that medical expenditures are important in explaining the observed savings of the elderly, especially the richer ones. For example, our baseline model predicts that, as in the data, between ages 74 and 84 median assets for those in the top permanent income quintile fall from \$170,000 to \$130,000. When we eliminate all out-of-pocket medical expenses, median assets for this group fall much more, from \$170,000 at age 74 to \$60,000 at age 84.

This result is due to an important feature of out-of-pocket medical expenses: average medical expenditures rise very rapidly with age and income. For example, our estimates predict that average annual out-of-pocket medical expenditures rise from \$1,100 at age 75 to \$9,200 at age 95. While a 95-year-old in the bottom quintile of the permanent income distribution expects to spend \$1,700 a year on medical expenses, a person of the same age in the top quintile expects to spend \$15,800. Medical needs that rise with age provide the elderly with a strong incentive to save, and medical expenses that rise with permanent income encourage the rich to be more frugal.

We show that the consumption floor also has an effect on saving decisions at all levels of income. When we reduce old-age consumption insurance by 20 percent, median assets for 90-year-olds in the highest permanent income quintile increase from \$100,000 to \$120,000, and median assets for 90-year-olds in the second-highest quintile increase from \$40,000 to \$50,000. The net worth of those in the third and fourth income quintiles also increases. The consumption floor thus matters for wealthy individuals as well as poorer ones. This is perhaps unsurprising given the size of our estimated medical needs for the old and income-rich; even wealthy households can be financially decimated by medical needs in very old age.

We further find that heterogeneity in mortality is large and is important for understanding the savings patterns of the elderly. In particular, differential mortality gives rise to a bias that makes the surviving elderly seem more thrifty than they actually are. While a 70-year-old man in poor health in the bottom income quintile expects to live only 6 more years, a 70-year-old woman in good health and in the top income quintile expects to live 17 more years. Thus, unhealthy and low-income men drop out of our sample more quickly. Failing to account for the mortality bias would lead us to understate asset decumulation by over 50 percent for the 74-year-old people in our sample. Consistently with the data, our model allows people who are rich, healthy, and female to live longer and generates asset profiles consistent with these observations.[1]

Finally, we find that bequests are luxury goods and that bequest motives are potentially quite important for the richest retirees. Our estimates of the bequest motive, however, are very imprecise, and for most of our sample, savings barely change when the bequest motive is eliminated. One reason why the bequest motive is weakly identified is that even in the top permanent income quintile, median assets in our sample of elderly singles never exceed $200,000; hence we do not have enough "super-rich" individuals to pin down the bequest motive.

The above results are based on a life cycle model that takes out-of-pocket medical expenses as exogenous. That is, we first use the Assets and Health Dynamics of the Oldest Old (AHEAD) data set to estimate stochastic processes for mortality and out-of-pocket medical expenditures as functions of sex, health, permanent income, and age. We then estimate our model using the method of simulated moments, where the model's preference parameters are chosen so that the permanent income–conditional median age-asset profiles simulated from the model match those in the data, cohort by cohort.

The additional sources of heterogeneity that we consider allow the model to match important aspects of the data: our estimated structural model is not rejected when we test its overidentifying restrictions. In addition, the distribution of deceased persons' estates generated by our model matches up closely with that observed in the data.

Despite the good fit of the model to the data, one might think that some of the results, such as the responses of savings to changes in government insurance, might not be robust to allowing the retirees to adjust both savings and medical expenditure. As a robustness check, we construct a model in which retirees choose medical expenditure as well as savings.

In this version of the model, retirees derive utility from consuming

[1] See Hurd, McFadden, and Merrill (2001) and Attanasio and Emmerson (2005) for evidence on permanent income and mortality.

medical care as well as other goods. The magnitude of this utility depends on "medical needs" shocks, which in turn depend on age and health. Retirees optimally choose medical and nonmedical consumption while taking into account the cost sharing provided by insurance and their own resources. We estimate this version of the model by requiring it to fit observed out-of-pocket medical expenditures as well as observed asset holdings.

Importantly, although our medical needs shocks do not depend on permanent income, in our model, out-of-pocket medical expenditures rise with both permanent income and age, as in the data. Medicaid plays an important role in generating this correlation because it pays for most of the medical expenses of the poor, who thus have few or no out-of-pocket medical expenses, but not for those of the rich, who pay out of pocket for their own, higher-quality medical care. Moreover, we find that medical expenditures and the consumption floor have large effects on saving even when medical expenses are a choice variable.

The intuition for why medical needs are so important, even when people can adjust their medical expenditure, is that out-of-pocket medical expenditures found in the model have to match those in the data. This implies that high-income 70-year-olds anticipate that if they live into their 90s, they will probably choose to make large medical expenditures—like the 90-year-olds in our sample—and will probably save to pay for them. Once this feature of the data is taken into account, it is not surprising that medical expenditures have large effects on savings, whether they are exogenous or chosen.

Making medical expenses endogenous reduces the effects of social insurance on savings. The effects that remain, however, are still stronger at higher income levels than at lower ones because the out-of-pocket medical expenses of the richest people are much higher than those of the poorest.

The rest of the paper is organized as follows. In Section II, we review the most closely related literature. In Section III, we introduce our version of the life cycle model, and in Section IV we discuss our estimation procedure. In Sections V and VI, we describe the data and the estimated shock processes that elderly individuals face. We discuss parameter estimates and model fit in Section VII. Section VIII contains some decomposition exercises that gauge the forces affecting saving behavior. In Section IX we develop a version of the model in which medical expenses are a choice variable, estimate the model, and use it to perform some robustness checks. We present conclusions in Section X.

II. Related Literature

Our paper is related to a number of papers in the savings literature that consider either uncertain medical expenditures or bequest motives. In an early study, Kotlikoff (1988) finds that out-of-pocket medical expenditures are potentially an important driver of aggregate saving. However, Kotlikoff also stresses the need for better data on medical expenses and for more realistic modeling of this source of risk.

Subsequent works by Hubbard, Skinner, and Zeldes (1994) and Palumbo (1999) find that medical expenses have fairly small effects compared with the ones we find. Their effects are smaller because their data understate the extent to which medical expenditures rise with age and income. As an example, the average medical expense for a 100-year-old generated by Hubbard et al.'s medical expenditure model is about 16 percent of the average medical expense for a 100-year-old found in our data. Our data set contains detailed information for a large number of very old individuals. This richness allows us to provide a more precise picture of how medical expenses rise at very advanced ages. Furthermore, our data include nursing home expenses, which previous studies had to impute.

Hubbard, Skinner, and Zeldes (1995) and Scholz, Seshadri, and Khitatrakun (2006) argue that means-tested social insurance programs (in the form of a minimum consumption floor) provide strong incentives for low-income individuals not to save. Their simulations, however, indicate that reducing the consumption floor has little effect on the consumption of college graduates. In contrast, we find that the consumption floor has an effect on saving decisions at all levels of income. Because out-of-pocket medical expenditures rise rapidly with income, rich individuals value social insurance as a safeguard against catastrophic expenses, even if they often end up not using it. This finding is consistent with Brown and Finkelstein's (2008) work, which finds that Medicaid has large effects on the decisions of fairly rich people.

Scholz et al. (2006) find that a life cycle model, augmented with realistic income and medical expense uncertainty, does a good job of fitting the distribution of wealth at retirement. We add to their paper by showing that a realistic life cycle model can do a good job of fitting the patterns of asset decumulation observed after retirement.

In his seminal paper, Hurd (1989) estimates a simple structural model of savings and bequest motives in which bequests are normal goods and does not find support for large bequest motives. De Nardi's (2004) calibration exercise shows that modeling bequests as a luxury good is important to explain the savings of the richest few. Kopczuk and Lupton (2007) find that a majority of elderly singles have a bequest motive. However, whether the motive is active or not depends on the individual's financial resources because, consistently with De Nardi, they estimate bequests to be luxury goods. While none of the preceding papers accounted for medical expenses, Dynan, Skinner, and Zeldes (2004) argue that the same assets can be used to address both precautionary and bequest concerns. Using responses from an attitudinal survey to separate bequest and medical expense motives, Ameriks et al. (2007) find that bequests are important for many people. In this paper we allow bequests to be luxury goods and we let the AHEAD data speak to both the intensity of bequest motives and the level of wealth at which they become operative.

III. The Model

Consider a retired single person, seeking to maximize his or her expected lifetime utility at age t, $t = t_{r+1}, \ldots, T$, where t_r is the retirement age. These individuals maximize their utility by choosing consumption c. Each period, the individual's utility depends on consumption and health status, h, which can be either good ($h = 1$) or bad ($h = 0$).

The flow utility from consumption is

$$u(c, h) = \delta(h)\frac{c^{1-\nu}}{1-\nu}, \tag{1}$$

with $\nu \geq 0$. Following Palumbo (1999), we model the dependence of utility on health status as

$$\delta(h) = 1 + \delta h, \tag{2}$$

so that when $\delta = 0$, health status does not affect utility.

When the person dies, any remaining assets are left to his or her heirs. We denote with e the estate net of taxes. The utility the household derives from leaving the estate e is

$$\phi(e) = \theta\frac{(e+k)^{(1-\nu)}}{1-\nu}, \tag{3}$$

where θ is the intensity of the bequest motive, and k determines the curvature of the bequest function and hence the extent to which bequests are luxury goods.

We assume that nonasset income, y_t, is a deterministic function of sex, g, permanent income, I, and age t:

$$y_t = y(g, I, t). \tag{4}$$

The individual faces several sources of risk, which we treat as exogenous. This seems a reasonable simplification since we focus on older people who have already shaped their health and lifestyle.

1. *Health status uncertainty.*—The transition probabilities for health status depend on previous health, sex, permanent income, and age:

$$\pi_{j,k,g,I,t} = \Pr(h_{t+1} = k | h_t = j, g, I, t), \quad j, k \in \{1, 0\}. \tag{5}$$

2. *Survival uncertainty.*—Let $s_{g,h,I,t}$ denote the probability that an individual of sex g is alive at age $t+1$, conditional on being alive at age t, having time t health status h, and enjoying permanent income I.

3. *Medical expense uncertainty.*—Medical expenses, m_t, are defined as out-of-pocket expenses. The mean and the variance of the log of medical expenses depend on sex, health status, permanent income, and age:

$$\ln m_t = m(g, h, I, t) + \sigma(g, h, I, t) \times \psi_t. \tag{6}$$

Following French and Jones (2004), we assume that the idiosyncratic component ψ_t can be decomposed as

$$\psi_t = \zeta_t + \xi_t, \quad \xi_t \sim N(0, \sigma_\xi^2), \tag{7}$$

$$\zeta_t = \rho_m \zeta_{t-1} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma_\epsilon^2), \tag{8}$$

where ξ_t and ϵ_t are serially and mutually independent. In practice, we discretize ξ and ζ using quadrature methods.

The timing is as follows: At the beginning of the period the individual's health status and medical expenses are realized. Then the individual consumes and saves. Finally, the survival shock hits. Individuals who die leave any remaining assets to their heirs.

Next period's assets are given by

$$a_{t+1} = a_t + y_n(ra_t + y_t, \tau) + b_t - m_t - c_t, \tag{9}$$

where $y_n(ra_t + y_t, \tau)$ denotes posttax income, r denotes the risk-free, pretax rate of return, the vector τ describes the tax structure, and b_t denotes government transfers. In addition, assets have to satisfy a borrowing constraint: $a_t \geq 0$.

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

$$b_t = \max\{0, \underline{c} + m_t - [a_t + y_n(ra_t + y_t, \tau)]\}. \tag{10}$$

Equation (10) says that government transfers bridge the gap between an individual's "total resources" (i.e., assets plus income less medical expenses) and the consumption floor. To make these transfers consistent with public insurance programs, we impose that if transfers are positive, $c_t = \underline{c}$ and $a_{t+1} = 0$.

To save on state variables, we redefine the problem in terms of cash on hand, x_t:

$$x_t = a_t + y_n(ra_t + y_t, \tau) + b_t - m_t. \tag{11}$$

Note that assets and cash on hand follow:

$$a_{t+1} = x_t - c_t, \tag{12}$$

$$x_{t+1} = x_t - c_t + y_n(r(x_t - c_t) + y_{t+1}, \tau) + b_{t+1} - m_{t+1}. \tag{13}$$

To enforce the consumption floor, we impose

$$x_t \geq \underline{c} \quad \forall t, \tag{14}$$

and to ensure that assets are always nonnegative, we require

$$c_t \leq x_t \quad \forall t. \tag{15}$$

Note that all the variables in x_t are given and known at the beginning of period t. We can thus write the individual's problem recursively, using the definition of cash on hand. Let β denote the discount factor. Then the value function for a single individual is given by

$$\begin{aligned} V_t(x_t, g, h_t, I, \zeta_t) = \max_{c_t, x_{t+1}} \{ & u(c_t, h_t) + \beta s_{g,h,I,t} E_t V_{t+1}(x_{t+1}, g, h_{t+1}, I, \zeta_{t+1}) \\ & + \beta(1 - s_{g,h,I,t})\phi(e_t)\} \end{aligned} \tag{16}$$

subject to

$$e_t = (x_t - c_t) - \max\{0, \tilde{\tau} \cdot (x_t - c_t - \tilde{x})\} \tag{17}$$

and equations (13)–(15). The parameter $\tilde{\tau}$ denotes the tax rate on estates in excess of $\tilde{x}$, the estate exemption level.

IV. Estimation Procedure

A. The Method of Simulated Moments

To estimate the model, we adopt a two-step strategy, similar to the one used by Gourinchas and Parker (2002), Cagetti (2003), and French and Jones (2007). In the first step we estimate or calibrate those parameters that can be cleanly identified without explicitly using our model. For example, we estimate mortality rates from raw demographic data. Let χ denote the collection of these first-step parameters.

In the second step we estimate the vector of parameters $\Delta = (\delta, \nu, \beta, \underline{c}, \theta, k)$ with the method of simulated moments (MSM), taking as given the elements of χ that were estimated in the first step. In particular, we find the vector $\hat{\Delta}$ yielding the simulated life cycle decision profiles that "best match" (as measured by a generalized method of moments [GMM] criterion function) the profiles from the data. Because our underlying motivations are to explain why elderly individuals retain so many assets and to explain why individuals with high income save at a higher rate, we match median assets by cohort, age, and permanent income.

Consider individual i of birth cohort p in calendar year v. Let a_{iv} denote individual i's assets. Sorting the sample by permanent income, we assign every individual to one of Q quantile-based intervals. In practice, we split the sample into five permanent income quintiles, so that $Q = 5$. Suppose that individual i of cohort p falls in the qth interval of the sample income distribution. Let $a_{pqv}(\Delta, \chi)$ be the model-predicted median asset level in calendar year v for an individual of cohort p who was in the qth income interval. Assuming that observed assets have a continuous density, at the "true" parameter vector (Δ_0, χ_0), exactly half of the individuals in group pqv will have asset levels of $a_{pqv}(\Delta_0, \chi_0)$ or less. This leads to the following set of moment conditions:

$$E(1\{a_{iv} \leq a_{pqv}(\Delta_0, \chi_0)\} - \tfrac{1}{2}\,|p,\ q,\ v,\ \text{individual } i \text{ alive at } v) = 0, \qquad (18)$$

for each p, q, and v triple. In other words, for each permanent income–cohort grouping, the model and the data have the same median asset levels. Our decision to use conditional medians rather than means reflects sample size considerations; in some pqv cells, changes in one or two individuals can lead to sizable changes in mean wealth. Sample size considerations also lead us to combine men and women in a single moment condition.

The mechanics of our MSM approach are fairly standard. We compute life cycle histories for a large number of artificial individuals. Each of these individuals is endowed with a value of the state vector $(t, x_t, g, h_t, I, \zeta_t)$ drawn from the data distribution for 1996, and each is assigned a series of health, medical expense, and mortality shocks consistent with the stochastic processes described in Section II above. We give each simulated person the entire health and mortality history realized by a person in the AHEAD data with the same initial conditions. (Since the data provide health and mortality only during interview years, we simulate it in off years using our estimated models and Bayes' rule.) The simulated medical expenditure shocks ζ and ξ are Monte Carlo draws from discretized versions of our estimated shock processes.

We discretize the asset grid, and using value function iteration, we solve the model numerically. This yields a set of decision rules, which, in combination with the simulated endowments and shocks, allows us to simulate each individual's assets, medical expenditures, health, and mortality. We then compute asset profiles (values of a_{pqv}) from the artificial histories in the same way that we compute them from the real data. We adjust until the difference between the data and simulated profiles—a GMM criterion function based on equation (18)—is minimized.[2]

We discuss the asymptotic distribution of the parameter estimates, the weighting matrix, and the overidentification tests in the online Appendix. (See Gouriéroux and Monfort [1996] for a primer on the method of simulated moments.)

B. *Econometric Considerations*

In estimating our model, we face two well-known econometric problems. First, in a cross section, older individuals will have earned their labor income in earlier calendar years than younger ones. Because wages have increased over time (with productivity), this means that older individuals are poorer at every age, and the measured saving profile will overstate asset decumulation over the life cycle. Put differently, even if the elderly do not run down their assets, our data will show that assets decline with age, since older individuals will have lower lifetime incomes and assets at each age. Not accounting for this effect will lead us to estimate a model that overstates the degree to which elderly people run down their assets.

Second, wealthier people tend to live longer, so that the average survivor in each cohort has higher lifetime income than the average deceased member of that cohort. This "mortality bias" tends to overstate asset growth in an unbalanced panel. In addition, as time passes and people die, the surviving people, relative to the deceased, will be healthy and female. These healthy females, knowing that they will live longer, will tend to be more frugal than their deceased counterparts and hence have a flatter asset profile in retirement. Not accounting for mortality bias will lead us to understate the degree to which elderly people run down their assets.

A major advantage of using a structural approach is that we can address these biases directly by replicating them in our simulations. We address the first problem by giving our simulated individuals age, wealth, health, gender, and income endowments drawn from the distribution observed in the data. If older people have lower lifetime incomes in our data, they will have lower lifetime incomes in our simulations. Fur-

[2] The programs solving for the value functions and simulating household histories are written in C. We use GAUSS for the econometrics. The GAUSS programs call the C programs, send them the necessary inputs—including parameter values and initial values of the state variables—and retrieve the simulated histories. The GAUSS programs then use the simulated histories and the data to compute the GMM criterion function. A copy of the codes and the data needed to run them are available on the *Journal* Web site.

thermore, we match assets at each age, conditional on cohort and income quintile. We address the second problem by allowing mortality to differ with sex, permanent income, and health status. As a result, our estimated decision rules and our simulated profiles incorporate mortality effects in the same way as the data.

V. Data

The AHEAD is part of the Health and Retirement Survey conducted by the University of Michigan. It is a survey of individuals who were noninstitutionalized and were aged 70 or older in 1994. A total of 8,222 individuals in 6,047 households (in other words, 3,872 singles and 2,175 couples) were interviewed for the AHEAD survey in late 1993/early 1994, which we refer to as 1994. These individuals were interviewed again in 1996, 1998, 2000, 2002, 2004, and 2006. The AHEAD data include a nationally representative core sample as well as additional samples of blacks, Hispanics, and Florida residents.

We consider only single retired individuals in the analysis. This leaves us with 3,259 individuals, of whom 592 are men and 2,667 are women. Of these 3,259 individuals, 884 are still alive in 2006. The online Appendix provides details on the data.

Our measure of net worth (or assets) is the sum of all assets less mortgages and other debts. The AHEAD has information on the value of housing and real estate, autos, liquid assets (which include money market accounts, savings accounts, Treasury bills, etc.), individual retirement accounts, Keoghs, stocks, the value of a farm or business, mutual funds, bonds, and "other" assets. We do not use 1994 assets because they were underreported (see the online Appendix).

Nonasset income includes the value of Social Security benefits, defined-benefit pension benefits, annuities, veteran's benefits, welfare, and food stamps. We measure permanent income as the individual's average income over all periods during which he or she is observed. Because Social Security benefits and (for the most part) pension benefits are a monotonic function of average lifetime labor income, this provides a reasonable measure of lifetime or permanent income.

Medical expenses are the sum of what the individual spends out of pocket on insurance premia, drug costs, hospital stays, nursing home care, doctor visits, dental visits, and outpatient care. They include medical expenses during the last year of life. They do not include expenses covered by insurance, either public or private. French and Jones (2004) show that the medical expense data in the AHEAD line up with the aggregate statistics. For our sample, mean medical expenses are \$3,712 with a standard deviation of \$13,429 in 1998 dollars. Although this figure is large, it is not surprising because Medicare did not cover prescription drugs for most of the sample period, requires copays for services, and caps the number of reimbursed nursing home and hospital nights.

In addition to constructing moment conditions, we also use the AHEAD data to construct the initial distribution of permanent income, age, sex, health, and cash on hand that starts off our simulations. Each simulated individual is given a state vector drawn from the joint distribution of state variables observed in 1996.

VI. Data Profiles and First-Step Estimation Results

In this section, we describe the life cycle profiles of the stochastic processes (e.g., medical expenditures) that are inputs to our dynamic programming model and the asset profiles we want our model to explain.

A. *Asset Profiles*

We construct the permanent income–conditional age-asset profiles as follows. We sort individuals into permanent income quintiles, and we track birth year cohorts. We work with five cohorts. The first cohort consists of individuals who were aged 72–76 in 1996; the second cohort contains those aged 77–81; the third contains those aged 82–86; the fourth contains those aged 87–91; and the final cohort, for sample size reasons, contains those aged 92–102. We use asset data for six different years: 1996, 1998, 2000, 2002, 2004, and 2006. To construct the profiles, we calculate cell medians for the survivors for each year assets are observed.

To fix ideas, consider figure 1, which plots median assets by age and income quintile for the members of two birth year cohorts who are still alive at each point in time. The solid lines at the far left of the graph are for the youngest cohort, whose members in 1996 were aged 72–76, with an average age of 74. The dashed set of lines are for the cohort aged 82–86 in 1996.

There are five lines for each cohort because we have split the data into permanent income quintiles. However, the fifth, bottom, line is hard to distinguish from the horizontal axis because households in the lowest permanent income quintile hold few assets.

The members of the first cohort appear in our sample at an average age of 74 in 1996. We then observe them in 1998, when they are on average 76 years old, and then again every two years until 2006. The other cohorts start from older initial ages and are followed for 10 years, until 2006. The graph reports median assets for each cohort and permanent income grouping for six data points over time.

Unsurprisingly, assets turn out to be monotonically increasing in income, so that the bottom line shows median assets in the lowest income

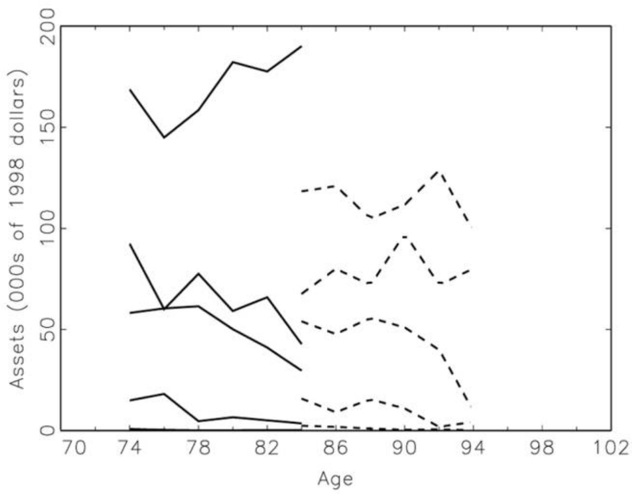

FIG. 1.—Median assets by cohort and permanent income quintile: data. Solid line, cohort aged 74 in 1996. Dashed line, cohort aged 84 in 1996.

quintile, and the top line shows median assets in the top quintile. For example, the top left line shows that for the top permanent income quintile of the cohort aged 74 in 1996, median assets started at \$170,000 and then stayed rather stable over time: \$150,000 at age 76, \$160,000 at age 78, \$180,000 at ages 80 and 82, and \$190,000 at age 84.

For all permanent income quintiles in these cohorts, the assets of surviving individuals neither rise rapidly nor decline rapidly with age. If anything, those with high income tend to have increases in their assets, whereas those with low income tend to have declines in assets as they age. (The profiles for other cohorts, which are shown in the online Appendix, are similar.) Our finding that the income-rich elderly dissave more slowly complements and confirms those by Dynan et al. (2004).

Figure 2 compares asset profiles that are aggregated over all the income quintiles. The solid line shows median assets for everyone observed at a given point in time, even if they died in a subsequent wave, that is, the unbalanced panel. The dashed line shows median assets for the subsample of individuals who were still alive in the final wave, that is, the balanced panel. It shows that the asset profiles for those who were alive in the final wave—the balanced panel—have much more of a downward slope. The difference between the two sets of profiles confirms that people who died during our sample period tended to have lower assets than the survivors.

The first pair of lines in figure 2 shows that failing to account for mortality bias would lead us to understate the asset decumulation of

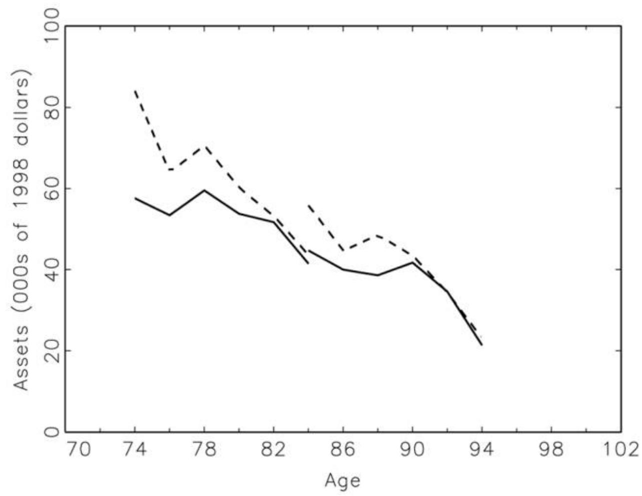

FIG. 2.—Median assets by birth cohort: everyone in the data (solid lines) versus survivors (dashed lines).

those who were 74 years old in 1996 by over 50 percent. In 1996 median assets of the 74-year-olds who survived to 2006 were \$84,000. In contrast, in 1996 median assets for all 74-year-olds were \$60,000. Median assets of those who survived to 2006 were \$44,000. The implied drops in median assets between 1996 and 2006 therefore depend on which population we look at: only \$16,000 for the unbalanced panel but \$40,000 for the balanced panel of those who survived to 2006. This is consistent with the findings of Love, Palumbo, and Smith (2009). Sorting the data by permanent income reduces, but does not eliminate, mortality bias.

Since our model explicitly takes mortality bias and differences in permanent income into account, it is the unbalanced panels that we use in our MSM estimation procedure. This greatly increases the size of our estimation sample.

B. *Medical Expense and Income Profiles*

The mean of logged medical expenses is modeled as a function of a quartic in age, sex, sex interacted with age, current health status, health status interacted with age, a quadratic in the individual's permanent income ranking, and permanent income ranking interacted with age. We estimate these profiles using a fixed-effects estimator.

We use fixed effects rather than ordinary least squares for two reasons. First, differential mortality causes the composition of our sample to vary with age, whereas we are interested in how medical expenses vary for

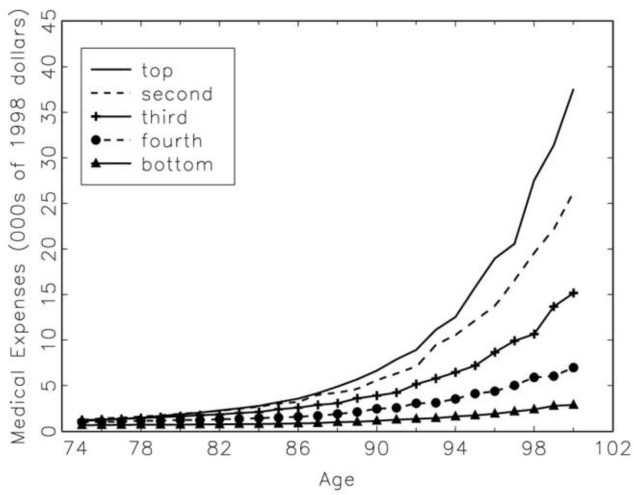

FIG. 3.—Average medical expenses, by permanent income quintile

the same individuals as they grow older. Second, cohort effects are likely to be important for both of these variables. Failing to account for the secular increase in medical expenses will lead to understating the growth of medical expenses by age. Cohort effects are captured in a fixed-effect estimator since they are merely the average fixed effect for all members of a given cohort.

The combined variance of the medical expense shocks ($\zeta_t + \xi_t$) is modeled with the same variables and functional form as the mean (see eq. [6]).

Our estimates indicate that average medical expenses for men are about 20 percent lower than for women, conditional on age, health, and permanent income. Average medical expenses for healthy people are about 50 percent lower than for unhealthy people, conditional on age, sex, and permanent income. These differences are large, but the differences across permanent income groups are even larger.

To better interpret our estimates, we simulate medical expense histories for the AHEAD birth year cohort whose members were aged 72–76 (with an average age of 74) in 1996. We begin the simulations with draws from the joint distribution of age, health, permanent income, and sex observed in 1996.

Figure 3 presents average medical expenses conditional on age and permanent income quintile for a balanced sample of people. We rule out attrition in these simulations because it is easier to understand how medical expenses evolve over time when tracking the same individuals. The picture with mortality bias, however, is similar. Permanent income

TABLE 1
Persistence and Variance of Innovations to Medical Expenses
(Variances as Fractions of Total Cross-Sectional Variance)

Parameter	Variable	Estimate (Standard Error)
ρ_m	Autocorrelation, persistent part	.922 (.010)
σ_ϵ^2	Innovation variance, persistent part	.050 (.008)
σ_ξ^2	Innovation variance, transitory part	.665 (.014)

has a large effect on average medical expenses, especially at older ages. Average medical expenses are less than $1,000 a year at age 75 and vary little with income. By age 100, they rise to $2,900 for those in the bottom quintile of the income distribution and to almost $38,000 for those at the top of the income distribution. Mean medical expenses at age 100 are $17,700.

The mean medical expenses implied by our estimated processes line up with the raw data. We have 58 observations on medical expenses for 100-year-old individuals, averaging $15,603 (with a standard deviation of $33,723 and a standard error of $4,428) per year, and 72 percent of these expenses are from nursing home care. Between ages 95 and 100, we have 725 person-year observations on medical expenses, averaging $9,227 (with a standard deviation of $19,988 and a standard error of $737). Therefore, the data indicate that average medical expenses for the elderly are high.

Medical expenses for the elderly are volatile as well as high. We find that the average variance of log medical expenses is 2.53. This implies that medical expenses for someone with a two-standard-deviation shock to medical expenses are 6.8 times the average, conditional on the observables.[3] The variance of medical expenses rises with age, bad health, and income.

French and Jones (2004) find that a suitably constructed lognormal distribution can match average medical expenses and the far right tail of the distribution. They also find that medical expenses are highly correlated over time. Table 1 shows estimates of the persistent component ζ_t and the transitory component ξ_t. The table shows that 66.5 percent of the cross-sectional variance of medical expenses is from the transitory component; the remainder, 33.5 percent, is from the persis-

[3] We assume that medical expenses are lognormally distributed, so the predicted level of medical expenses is $\exp(m + \frac{1}{2}\sigma^2)$, where m denotes predicted log medical expenses and σ^2 denotes the variance of the idiosyncratic shock ψ_t. The ratio of the level of medical expenses two standard deviations above the mean to average medical expenses is

$$\frac{\exp(m + 2\sigma)}{\exp(m + \sigma^2/2)} = \exp(2\sigma - \sigma^2/2) = 6.80$$

if $\sigma = \sqrt{2.53}$.

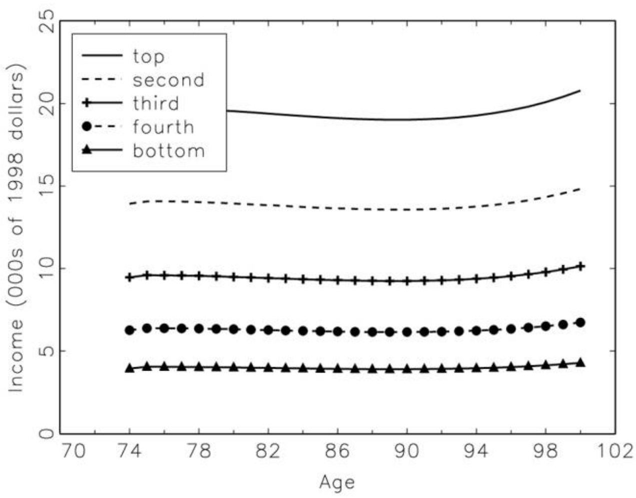

FIG. 4.—Average income, by permanent income quintile

tent component. The persistent component has an autocorrelation coefficient of 0.922, however, so that innovations to the persistent component of medical expenses have long-lived effects. Most of a household's *lifetime* medical expense risk comes from the persistent component.

Our estimates of medical expense risk indicate greater risk than found in other studies (see Hubbard et al. 1994; Palumbo 1999). However, our estimates still potentially understate both the level and risk of the medical expenses faced by older Americans because our measure of medical expenditures does not include the value of Medicaid contributions. As equation (10) shows, some of the medical expenses (m_t) in our model may be paid for by the government through the provision of the consumption floor. Therefore, the ideal measure of m_t drawn from the data would include both the out-of-pocket expenditures actually made by the consumer and the expenditures covered by Medicaid. The AHEAD data, however, do not include Medicaid expenditures. In this respect, the medical expense process we feed into our benchmark model is a conservative one.

We model mean nonasset income in the same way as mean medical expenses, using the same explanatory variables and the same fixed-effects estimator. Figure 4 presents income, conditional on age and permanent income. For those in the top permanent income quintile, annual income averages \$20,000 per year. Median wealth for the youngest cohort in this income group is slightly under \$200,000, or about 10 years' worth of income for this group.

TABLE 2
LIFE EXPECTANCY IN YEARS, CONDITIONAL ON REACHING AGE 70

Income Quintile	Healthy Male	Unhealthy Male	Healthy Female	Unhealthy Female	All*
Bottom	7.6	5.9	12.8	10.9	11.1
Second	8.4	6.6	13.8	12.0	12.4
Third	9.3	7.4	14.7	13.2	13.1
Fourth	10.5	8.4	15.7	14.2	14.4
Top	11.3	9.3	16.7	15.1	14.7
By gender:†					
Men	9.7				
Women	14.3				
By health status:‡					
Unhealthy	11.6				
Healthy	14.4				

NOTE.—Life expectancies are calculated through simulations using estimated health transition and survivor functions.
* Calculations use the gender and health distributions observed in each permanent income quintile.
† Calculations use the health and permanent income distributions observed for each gender.
‡ Calculations use the gender and permanent income distributions observed for each health status group.

C. Mortality and Health Status

We estimate the probability of death and bad health as logistic functions of a cubic in age, sex, sex interacted with age, previous health status, health status interacted with age, a quadratic in permanent income rank, and permanent income rank interacted with age.

Using the estimated health transitions, survival probabilities, and the initial joint distribution of age, health, permanent income, and sex found in our AHEAD data, we simulate demographic histories. Table 2 presents predicted life expectancies.[4] Rich people, women, and healthy people live much longer than their poor, male, and sick counterparts. Two extremes illustrate this point: an unhealthy male at the bottom quintile of the permanent income distribution expects to live only 6 more years, that is, to age 76. In contrast, a healthy woman at the top quintile of the permanent income distribution expects to live 17 more years, thus making it to age 87. Such significant differences in life expectancy, all else being equal, should lead to significant differences in saving behavior. In complementary work (De Nardi, French, and Jones 2009), we show that this is in fact the case.

We also find that for rich people, the probability of living to very old ages, and thus facing very high medical expenses, is significant. For example, using the same simulations used to construct table 2, we find that a healthy 70-year-old woman in the top quintile of the permanent

[4] Our predicted life expectancy at age 70 is about 3 years less than what the aggregate statistics imply. This discrepancy stems from using data on singles only: when we reestimate the model for both couples and singles, we find that predicted life expectancy is within half a year of the aggregate statistics for both men and women.

income distribution faces a 14 percent chance of living 25 years, to age 95.

VII. Second-Step Estimation Results

Table 3 presents preference parameter estimates for several specifications. Column 1 refers to a parsimonious model with no bequest motives and no health preference shifter. Column 2 refers to a model in which health can shift the marginal utility of consumption. In column 3, the bequest motive is activated, and in column 4, both the bequest motive and the preference shifter are active. In all cases, we set the interest rate to 2 percent.

Table 3 shows that the bequest parameters are never statistically significant and, as shown by the overidentification statistics, have little effect on the model's fit. When considered in isolation, the health preference parameter is not significant either. Column 4 shows that this parameter is statistically significant when bequest motives are included. A more appropriate test, however, is the joint test based on the change in the overidentification statistic. Comparing columns 1 and 4 of table 3 shows that the test statistic decreases by 4.8, and three degrees of freedom are removed. With a $\chi^2(3)$ distribution, this change has a p-value of 18.7 percent, implying that we cannot reject the parsimonious model.

In short, the bequest and health preference parameters are (collec-

TABLE 3
ESTIMATED STRUCTURAL PARAMETERS

	Benchmark (1)	With Health-Dependent Preferences (2)	With Bequests (3)	All (4)
ν: relative risk aversion coefficient	3.81	3.75	3.84	3.66
	(.50)	(.47)	(.55)	(.55)
β: discount factor	.97	.97	.97	.97
	(.04)	(.05)	(.05)	(.04)
δ: preference shifter	.0	−.21	.0	−.36
	NA	(.18)	NA	(.14)
θ: bequest intensity	.0	.0	2,360	2,419
	NA	NA	(8,122)	(1,886)
k: bequest curvature (in 000s)			273	215
			(446)	(150)
$\underline{c}$: consumption floor ($)	2,663	2,653	2,665	2,653
	(346)	(337)	(353)	(337)
Overidentification test	82.3	80.6	81.5	77.5
Degrees of freedom	98	97	96	95
p-value	87.4%	88.5%	85.4%	90.5%

NOTE.—Standard errors are in parentheses below estimated parameters. NA refers to parameters fixed for a given estimation.

tively) not statistically significant, do not help improve the fit of the model, and, moreover, have little effect on any of the other parameter estimates. We thus use the parsimonious model as our benchmark specification and only briefly discuss the other configurations.

A. *The Benchmark Model*

Column 1 of table 3 shows that the estimated coefficient of relative risk aversion is 3.8, the discount factor is .97, and the consumption floor is \$2,663. These estimates are within the range of parameter estimates provided in the previous literature.

Our estimated coefficient of relative risk aversion, 3.8, is higher than the coefficients found by fitting nonretiree consumption trajectories, either through Euler equation estimation (e.g., Attanasio et al. 1999) or through MSM (Gourinchas and Parker 2002). It is, however, at the lower end of the estimates found by Cagetti (2003), who matched wealth profiles with MSM over the whole life cycle. It is much lower than those produced by Palumbo (1999), who matched consumption data for retirees using maximum likelihood estimation. Given that our out-of-pocket medical expenditure data indicate more risk than that found by Palumbo, it is not surprising that we can match observed precautionary and life cycle savings with a lower level of risk aversion.

The consumption floor that we estimate (\$2,700 in 1998 dollars) is similar to the value that Palumbo (1999) uses (\$2,000 in 1985 dollars). However, our estimate is about half the size of the value that Hubbard et al. (1994) use and is also about half the average value of Supplemental Security Income (SSI) benefits. Our consumption floor proxies for Medicaid health insurance (which almost eliminates medical expenses to the financially destitute) and SSI. Given the complexity of these programs and the fact that many potential recipients do not participate in them (Elder and Powers 2006), it is tricky to establish a priori what the consumption floor should be. Our estimates likely provide an "effective" consumption floor, one that combines the complexity and variety of the statutory rules with people's perceptions and attitudes toward welfare eligibility. In the online Appendix, we show that fixing the consumption floor at \$5,000 significantly worsens the model's fit.

The Euler equation (3) and asset accumulation equation (9) show that the asset profiles generated by our model depend on expected and realized interest rates. As a robustness check, we solve a model with independent and identically distributed (i.i.d.) interest rate shocks, we use the realized rates of returns over the 1996–2005 period in the simulations, and we reestimate the model under these assumptions. Our main findings hold in this case as well (see the online Appendix).

The Euler equation also gives some intuition for the estimates and their identification. If we ignore taxes, borrowing constraints, and bequest motives, it is given by

$$(1 + \delta h_t)c_t^{-\nu} = \beta(1 + r)s_t E_t(1 + \delta h_{t+1})c_{t+1}^{-\nu}. \tag{19}$$

Log-linearizing it, we obtain that expected consumption growth follows:

$$E_t(\Delta \ln c_{t+1}) = \frac{1}{\nu}\{\ln [\beta(1 + r)s_t] + \delta E_t(h_{t+1} - h_t)\} + \frac{\nu + 1}{2}\text{Var}_t\,(\Delta \ln c_{t+1}). \tag{20}$$

The coefficient of relative risk aversion is identified by differences in saving rates across the income distribution, in combination with the consumption floor. Low-income households are relatively more protected by the consumption floor and will thus have lower values of $\text{Var}_t\,(\Delta \ln c_{t+1})$ and hence weaker precautionary motives. The parameter helps the model explain why individuals with high permanent income typically display less asset decumulation. The online Appendix discusses the identification of the coefficient of relative risk aversion and the consumption floor in more detail.

Figure 5 shows how the baseline model fits a subset of the data profiles, using unbalanced panels. (The model fits equally well for the cells that are not shown.) Both in the model and in the observed data, individuals with high permanent income tend to increase their wealth with age, whereas individuals with low income tend to run down their wealth with age.

The visual evidence shown in figure 5 is consistent with the test statistics shown in table 3. The p-value of the overidentification statistic for our baseline specification is 87.3 percent. Hence, our model is not rejected by the overidentification test at any standard level of significance. Both in the model and in the data, individuals with high permanent income do not run down their wealth with age, whereas those with low income do.

With regard to the mortality bias, figure 6 shows simulated asset profiles, first for all simulated individuals alive at each date and then for the individuals surviving the entire simulation period. As in the data, restricting the profiles to long-term survivors reveals much more asset decumulation. The mortality bias generated by the model is large, reflecting heterogeneity in both saving behavior and mortality patterns.

Given that the survival rate, s_t, is often much less than one, it follows from equation (19) that the model will generate downward-sloping, rather than flat, consumption profiles, unless the discount factor is fairly large. Figure 7 shows simulated consumption profiles for ages 74–100. Except for the last two years of life, consumption falls with age. This

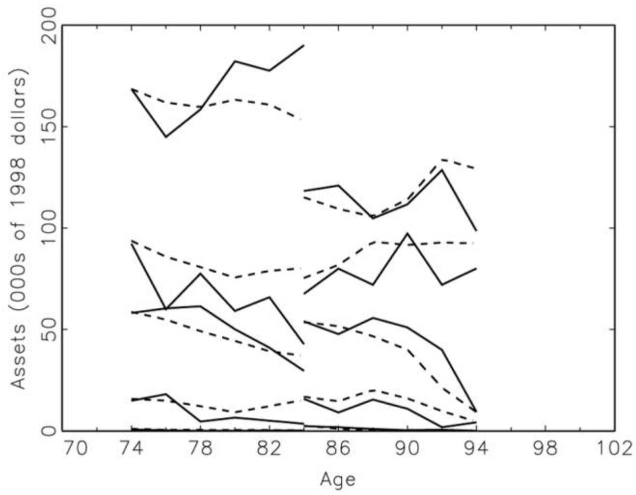

FIG. 5.—Median assets by cohort and permanent income quintile: data (solid lines) and model (dashed lines).

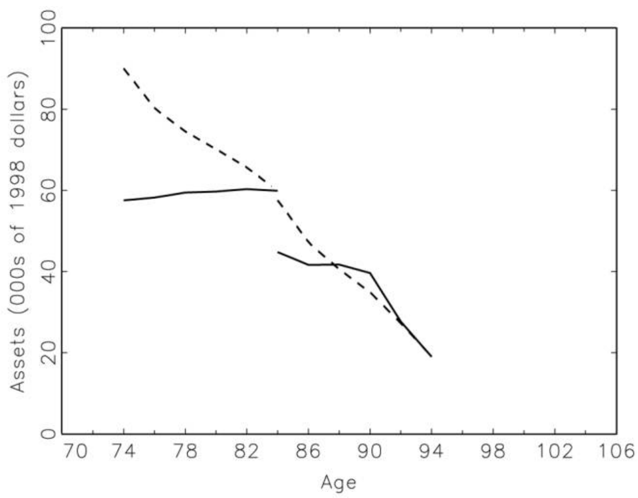

FIG. 6.—Median assets by birth cohort: everyone in the simulations (solid lines) versus survivors (dashed lines).

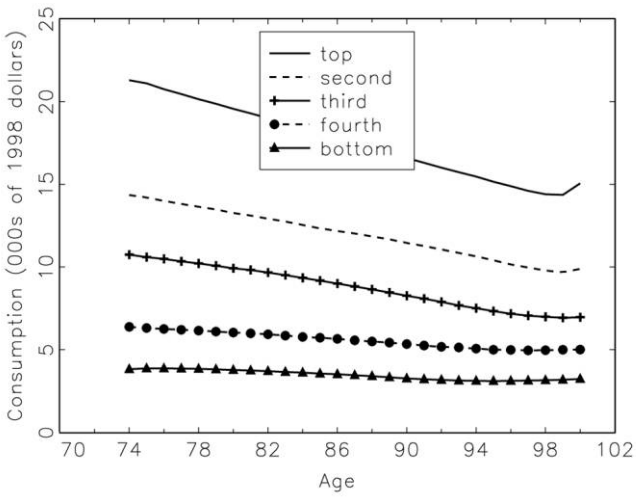

FIG. 7.—Median consumption by permanent income quintile: simulations

general tendency is consistent with many empirical studies of older-age consumption. For example, Fernandez-Villáverde and Krueger (2007) find that nondurable consumption declines about 1 percent per year between ages 70 and 90 (see also Banks, Blundell, and Tanner 1998).

B. *The Model with Health-Dependent Preferences*

Columns 2 and 4 of table 3 show point estimates of $\delta = -.21$ or $\delta = -.36$: when consumption is held fixed, being in good health lowers the marginal utility of consumption by 21–36 percent. This implies that an anticipated change from good to bad health leads consumption to increase by 6–10 percent, depending on the specification (see eq. [19]). This parameter, however, is not statistically significant. Moreover, none of the other parameter estimates are affected by its inclusion.

C. *The Model with Bequest Motives*

Columns 3 and 4 of table 3 show parameter estimates for models that include a bequest motive. Because the two specifications deliver similar parameter estimates, we focus on the results in column 3.

The point estimates of $\theta = 2{,}360$ and $k = 273{,}000$ indicate that the consumption level above which the bequest motive becomes operative is about \$36,000 per year. (See the online Appendix for a derivation.) By way of comparison, individuals in the top permanent income quintile have an average annuity income of about \$20,000 and hold less than

$200,000 of assets. This suggests that most people in our sample do not have a strong bequest motive. Not surprisingly, we find that bequest motives are not very important for fitting our data; none of the estimated bequest parameters are statistically significant, and adding bequest motives does not significantly improve the model's fit.

For those sufficiently rich, however, the marginal propensity to bequeath above that consumption level is also high, with 88 cents of every extra dollar above the threshold being left as bequests. Hence one can interpret our estimates as suggesting that the bequest motive could be present for the richest people in our sample.

To show the model's implications for bequests, figure 8 displays the distribution of assets that individuals hold one period before their deaths.[5] Comparing the two panels highlights that the models with and without the estimated bequest motive generate very similar distributions of assets at that time and that both match the actual data closely.

Our findings should not be interpreted as a rejection of bequest motives in general. Our sample is composed of elderly singles, who are poorer than couples, and other evidence indicates that bequests are a luxury good (De Nardi 2004; Dynan et al. 2004). Our sample of singles may not contain enough rich households to reveal strong bequest motives. Moreover, and importantly, a significant fraction of our sample is composed of people who have already lost their partner, and it is possible that some of the estate was already split between the surviving spouse and other heirs.

D. Summary

Our main results from the second-step estimation can be summarized as follows. First, our estimates of the coefficient of relative risk aversion, time discount factor, and consumption floor are within the range used in the previous literature. Second, we do not find that health-dependent preferences are important for understanding retirees' saving behavior. Third, we do not find that bequest motives significantly affect the savings of most households in our sample. Fourth, the model fits the data closely, both in terms of the moment conditions that we match and in terms of the distribution of bequests that it generates. Finally, the model's consumption implications are consistent with previous empirical evidence. Put together, these findings give us confidence that we can use our benchmark model (without bequest motives and health preference

[5] When AHEAD respondents die, their descendants are asked about the value of the estate. However, problems of nonresponse are severe: 49 percent of all estate values are imputed. Furthermore, it is unclear whether reported estates also include the value of the deceased individual's home. For these reasons, we report the distribution of assets 1 year before death rather than estates.

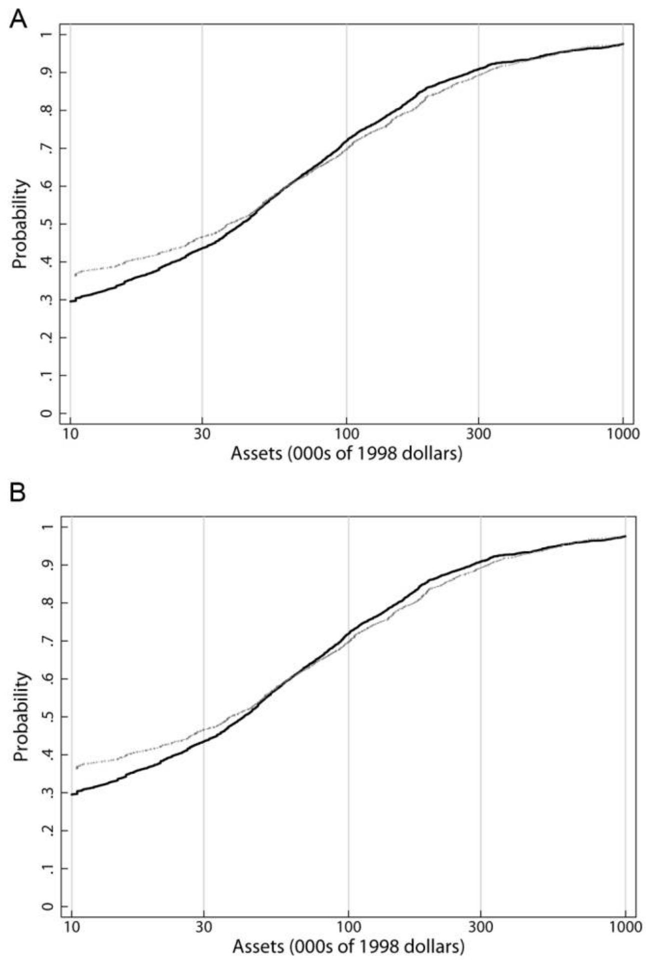

FIG. 8.—CDF assets held one period before death: data (solid line) and model (lighter line): *A*, results for models with bequest motives; *B*, results for models without bequest motives.

shocks) to study how savings depend on medical expenses and the consumption floor.

VIII. What Are the Important Determinants of Savings?

To determine the importance of the key mechanisms in our model, we fix the estimated parameters at their benchmark values and then change one feature of the model at a time. For each of these different economic environments, we compute the optimal saving decisions, simulate the

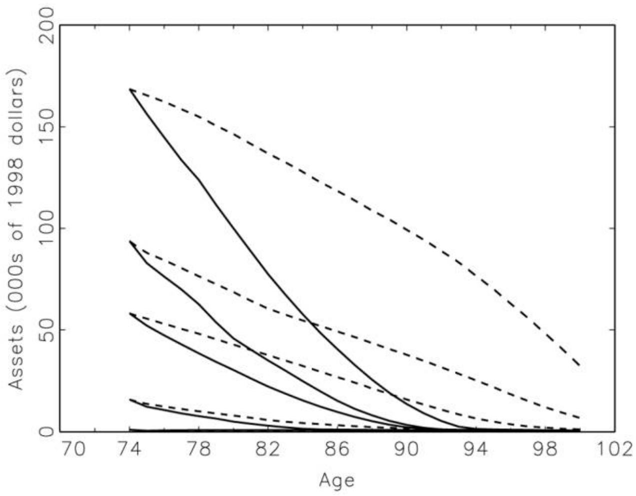

FIG. 9.—Median assets by cohort and permanent income quintile: baseline model (dashed lines) and model with no medical expenses (solid lines).

model, and compare the resulting asset accumulation profiles to the asset profiles generated by the baseline model.

We display asset profiles for the AHEAD birth year cohort whose members were aged 72–76 (with an average age of 74) in 1996. To focus on underlying changes in saving, we rule out attrition and assume that every individual lives to age 100. People will face the same risks and have the same expectations as before (except for the particular aspect changed in each experiment), but in their realized lives, they do not die and drop out of the simulations until age 100. None of our conclusions would change if we were to also allow for mortality bias.

First, we ask whether the out-of-pocket medical expenditures that we estimate from the data are important drivers of old-age savings. To answer this question, we zero out all out-of-pocket medical expenditures for everyone and look at the corresponding profiles. This could be seen as an extreme form of insurance provided by the government. Figure 9 shows that medical expenses are a big determinant of the elderly people's saving behavior, especially for those with high permanent income, for whom those expenses are especially high, and for those who are relatively less insured by the government-provided consumption floor. These retirees are reducing their current consumption in order to pay for the high out-of-pocket medical expenses they expect to bear later in life. For a given level of initial wealth, if there were no out-of-pocket medical expenses, individuals in the highest permanent income quintile would deplete their net worth by age 94. In the baseline model

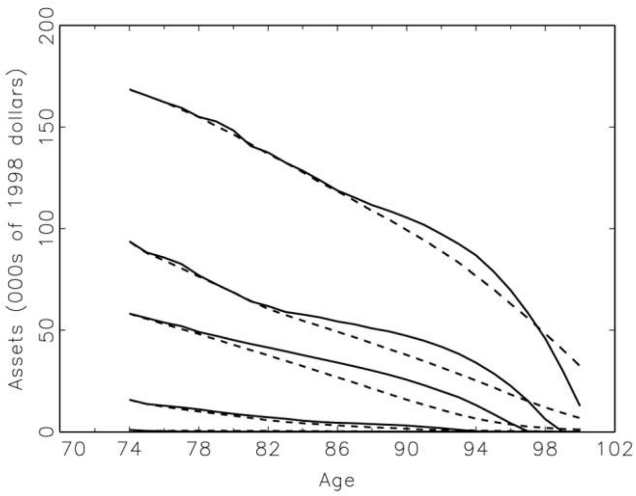

FIG. 10.—Median assets by cohort and permanent income quintile: baseline model (dashed lines) and model with no medical expense risk (solid lines).

with medical expenses, at age 100, they keep almost $40,000 to pay for out-of-pocket medical expenses in the last few years of life. The risk of living long and having high medical expenses late in life significantly increases savings. Our results indicate that modeling out-of-pocket medical expenses is important in evaluating policy proposals that affect the elderly.

We next shut down out-of-pocket medical expense risk (the shocks ζ and ξ), while keeping average medical expenditure constant (conditional on all the relevant state variables). Figure 10 shows the results. Interestingly, and consistently with Hubbard et al. (1994), we find that, conditional on average medical expenses, the risk associated with the volatility of medical expenses has only a small effect on the profiles of median wealth. Our results are also consistent with Palumbo's (1999) finding that eliminating medical expense risk has only small effects on consumption and assets.[6]

One reason why medical expense risk might not have a large effect is that the consumption floor limits the effects of catastrophic medical expenses. To explore this effect further, we reduce the consumption floor to 80 percent of its value, that is, from $2,663 to $2,100. One could

[6] Figure 10 shows that eliminating medical expense risk sometimes causes assets to increase. One reason for this increase is that a decrease in the variance of medical expenses decreases the frequency of large medical expenses, which in turn reduces the fraction of expenses covered by Medicaid, thus raising the average medical expense borne by retirees. Because this cost increase will be highest at oldest ages, individuals will respond by accumulating more assets.

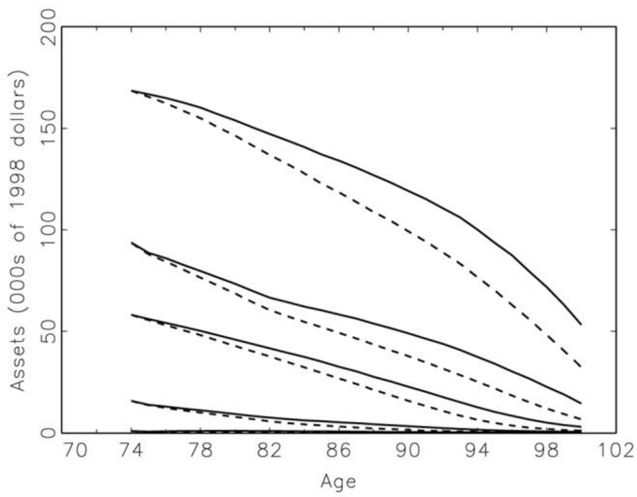

FIG. 11.—Median assets by cohort and permanent income quintile: baseline model (dashed lines) and model with 80 percent of the consumption floor (solid lines).

interpret this as a reform reducing the government-provided consumption safety net. Individuals respond to the increase in net income uncertainty by keeping more assets to self-insure. Figure 11 shows that reducing the consumption floor affects the savings profiles of singles with both low and high permanent income. The assets of 90-year-olds in the highest permanent income quintile increase from $100,000 to $120,000, whereas those of people of the same age but in the second-highest income quintile increase from $40,000 to $50,000. The net worth of those in the third and fourth income quintiles also displays some increases. The consumption floor thus matters for wealthy individuals as well as poor ones. This is perhaps unsurprising given the size of our estimated medical expenses for the old and income-rich; even wealthy households can be financially decimated by medical expenses in very old age.

IX. Endogenous Medical Spending

To check the sensitivity of our findings to the assumption that medical expenses are exogenous, we consider a more complex model in which retirees optimally choose how much to spend on medical goods and services as well as on nonmedical consumption. Our findings are robust to this extension. We assume that retirees derive utility from consumption of both nonmedical and medical goods, with the relative weights on the two goods varying with age, health, and an idiosyncratic "medical needs" shock.

A complementary approach is one in which medical expenses represent investments in health capital, which in turn decreases mortality (see, e.g., Yogo 2009). While this is an appealing mechanism, the existing empirical literature suggests that these effects are particularly small for the U.S. elderly for two reasons. First, the expenditures that we are considering are supplementing Medicaid, Medicare, and insurance-provided medical goods and services, which cover most life-threatening conditions. Second, the stock of health carried by an elderly person is in large part determined by the health investments that were made in the past, including those made by the person's parents in his or her childhood and even before birth. Hence, for our sample of people aged 70 and older, the effects of additional health investments are not as large as, for example, in infancy.

A key piece of empirical evidence comes from the RAND Health Insurance Experiment, where a random set of individuals were given copayment-free health insurance over a 3–5-year period, and a control group faced standard copayments. Brook et al. (1983) found that even though the group with free health care utilized medical services much more intensively than the control group, the additional treatments had only a "minimal influence" on subsequent health outcomes.

Surprisingly, some empirical studies show that even programs such as Medicare, which sometimes help pay for critical treatments, do not significantly increase life expectancy (Finkelstein and McKnight 2005). More in line with what one might expect, Card, Dobkin, and Maestas (2009) find that Medicare caused a small reduction in mortality among 65-year-olds admitted through emergency rooms for "nondeferrable" conditions.

Khwaja (forthcoming) studies a structural model in which medical expenditures both improve health and provide utility. He finds that "medical utilization would only decline by less than 20% over the life cycle if medical care was purely mitigative and had no curative or preventive components" (see also Davis 2005; Blau and Gilleskie 2008).

Given that the existing evidence indicates that the effect of additional medical spending on life expectancy is small, we focus on the utility effects of medical expenditures.

A. *The Endogenous Medical Expenditure Model*

At the beginning of each period, the individual's health status and medical needs shocks are realized and need-based transfers are given. The individual chooses consumption and medical expenditures and saves. Finally, the survival shock hits. If the person dies, the estate is passed on to his or her heirs.

The flow utility from consumption is given by

$$u(c_t,\ m_t,\ h_t,\ \zeta_t,\ \xi_t,\ t)\ =\ \frac{1}{1-\nu}\,c_t^{1-\nu} + \mu(t,\ h_t,\ \zeta_t,\ \xi_t)\,\frac{1}{1-\omega}\,m_t^{1-\omega}, \qquad (21)$$

where t is age, c_t is consumption of nonmedical goods, m_t is total consumption of medical goods, h_t is health status, and $\mu(\cdot)$ is the medical needs shock, which affects the marginal utility of consuming medical goods and services. The consumption of both types of goods is expressed in dollar values.

As before, we allow the need for medical services to have a temporary (ξ_t) and a persistent (ζ_t) component; we recycle the variable names to save on notation. We assume that these shocks follow the same processes as in equations (7) and (8), with potentially different parameters. It is worth stressing that we do not allow any of these shocks to depend on permanent income; income affects medical expenditures solely through the budget constraint.

We model two important features of the health care system.

1. Private and public insurance pay the share $1 - q(t,\ h_t)$ of the total medical costs incurred by the retiree. Its complement, $q(t,\ h_t)$, is the out-of-pocket share paid by the retiree. We estimate $q(t,\ h_t)$ as part of our first-stage estimation. (See the online Appendix for details.)

2. Social insurance programs, such as Medicaid and SSI, provide monetary transfers that vary with financial resources and medical needs. We model these transfers as providing a flow *utility* floor. The transfers thus depend on the retirees' state variables, including their medical needs shocks. For a given utility floor and state vector, we find the transfer $b^*(\cdot)\ =\ b^*(t,\ a_t,\ g,\ h_t,\ I,\ \zeta_t,\ \xi_t)$ that puts each retiree's utility at the floor. Transfers then kick in to provide the minimum utility level to retirees who lack the resources to afford it:

$$b(t,\ a_t,\ g,\ h_t,\ I,\ \zeta_t,\ \xi_t)\ =\ \max\{0,\ b^*(t,\ a_t,\ g,\ h_t,\ I,\ \zeta_t,\ \xi_t)\}. \qquad (22)$$

As before, we impose that if transfers are positive, the individual consumes all of his resources (by splitting them optimally between the two goods), so that $a_{t+1}\ =\ 0$.

The retiree's value function is given by

$$\begin{aligned} V_t(t, a_t, g, h_t, I, \zeta_t, \xi_t)\ =\ \max_{c_t, m_t, a_{t+1}} \Bigg\{ & \frac{1}{1-\nu}\,c_t^{1-\nu} + \mu(t, h_t, \zeta_t, \xi_t)\,\frac{1}{1-\omega}\,m_t^{1-\omega} \\ & + \beta s_{g,h,I,t} E_t(V_{t+1}(t+1, a_{t+1}, g, h_{t+1}, I, \zeta_{t+1}, \xi_{t+1}))\Bigg\}, \end{aligned} \qquad (23)$$

subject to equations (17), (22), $a_{t+1} \geq 0$, and

$$a_{t+1} = a_t + y_n(ra_t + y_t) + b(t, a_t, g, h_t, I, \zeta_t, \xi_t) - c_t - m_t q(t, h_t). \quad (24)$$

B. *Estimation*

The log of the medical needs shifter μ_t is modeled as a function of a cubic in age, current health status, and health status interacted with age. The combined variance of the shocks $\zeta_t + \xi_t$ is modeled as a quadratic in age, current health status, and health status interacted with age. To identify these parameters, we expand the moment conditions described by equation (18) to include mean medical expenses by age and birth cohort for each half of the permanent income distribution, the 90th percentile of medical expenses in the same cells, and the first and second autocorrelations for medical expenses in each cell.[7] Detailed moment conditions can be found in the online Appendix. In all other respects, our MSM procedure is the same as before.

C. *Results*

The preference parameter estimates for the endogenous medical expenditure model are shown in the online Appendix. The new estimate of the discount factor, β, is 0.99, slightly higher than the benchmark estimate of 0.97. The estimate of the coefficient of relative risk aversion for "regular" consumption, ν, is 2.15, whereas the estimate of the coefficient of relative risk aversion for medical goods, ω, is 3.19; hence the demand for medical goods is less elastic than the demand for consumption. Both coefficients are similar to, but somewhat smaller than, the benchmark estimate of 3.82.

The model also requires parameter estimates for the mean of the logged medical needs shifter $\mu(t, h_t, \zeta_t, \xi_t)$ and the process for the shocks ζ_t and ξ_t. The estimates for these parameters show that the demand for medical services rises rapidly with age.

Although the overidentification test statistic shows that the model with endogenous medical expenditure is rejected by the data, the model does match the main patterns of the asset profiles (see fig. 12). Furthermore, the model fits the medical expense distribution rather well (see the online Appendix).

Out-of-pocket (and total) medical expenses in the endogenous medical expenditure model can be eliminated by setting the medical needs

[7] The AHEAD medical expenses data are reported net of any Medicaid payment. We net out government transfers (b_t) from the model-generated medical expenses. Hence, AHEAD medical expenses and model-generated ones are both net of Medicaid payments.

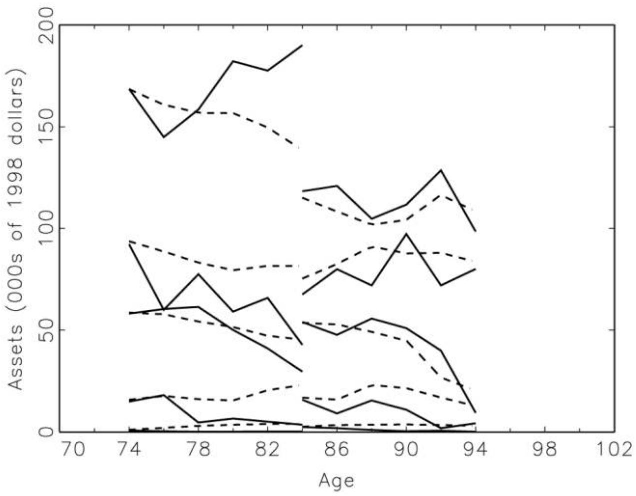

FIG. 12.—Median assets by cohort and permanent income quintile: data and endogenous medical expenditure model.

parameter μ_t to zero. Figure 13 shows that the effects of eliminating medical expenses are similar to those found in the exogenous medical expense model. Given that the consumer must prepare for the same pattern of medical expenditures in either model, this is not surprising. Retirees will save for high medical expenditures at old ages whether the expenditures are exogenous shocks estimated from the data or medical expenditure choices consistent with the same data.

Figure 14 shows the effects of reducing the generosity of social insurance by 50 percent for both the exogenous medical expenses model and the endogenous one with a utility floor.[8] As in the exogenous medical expenditure model, a change in social insurance affects the savings of all income groups, including the richest. The effects are smaller because retirees in the endogenous medical expense model can adjust medical expenditures as well as consumption.

In sum, the endogenous medical expense model confirms and reinforces our conclusion that medical expenses are a major driver of savings and that social insurance affects the saving of the income-rich as well as that of the income-poor.

[8] In the endogenous medical expense model, we index the estimated utility floor by the consumption level that provides the floor when $\mu = 1$. To run the counterfactual, we cut that consumption level in half.

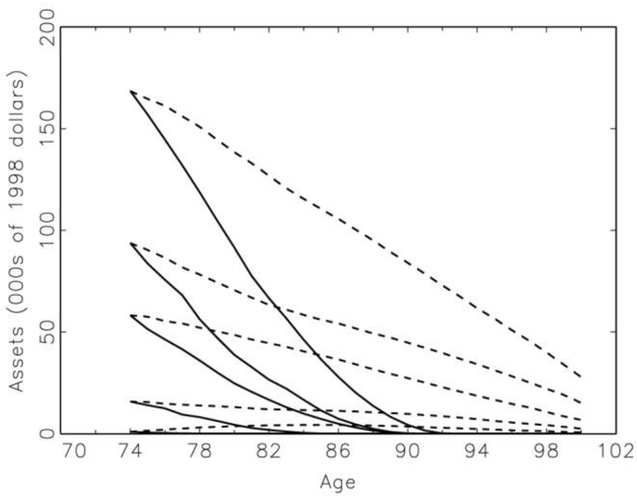

FIG. 13.—Median assets for the endogenous medical expense model: baseline model (dashed lines) and model without medical expenses (solid lines).

X. Conclusions

In this paper, we construct, estimate, and analyze a rich model of saving for retired single people. In doing so, we provide several contributions. First, we estimate the out-of-pocket medical expenses faced by the elderly using a larger data set (which includes nursing home expenses) and more flexible functional forms than in the previous literature. As a result, we find that medical expenses are much higher and more volatile than previously estimated and that they rise very fast with age. Also, at very advanced ages (i.e., starting from about age 80), medical expenses are very much a luxury good; that is, they are much higher for individuals with higher permanent income.

Second, we estimate mortality probabilities by age as a function of health, sex, and permanent income and find large variations along all three dimensions. We find that in an unbalanced panel, mortality bias—the tendency of rich people to live longer—is significant.

Third, we construct and estimate a structural model of saving using the method of simulated moments. As a result of our improved data and richer specifications, our model can explain many important features of the data. In particular, our estimated structural model fits saving profiles across the entire income distribution and reproduces the observation that elderly people with high permanent income dissave less than elderly people with low permanent income.

Fourth, our results imply that the pattern of out-of-pocket medical expenses by age and permanent income is a key determinant of savings.

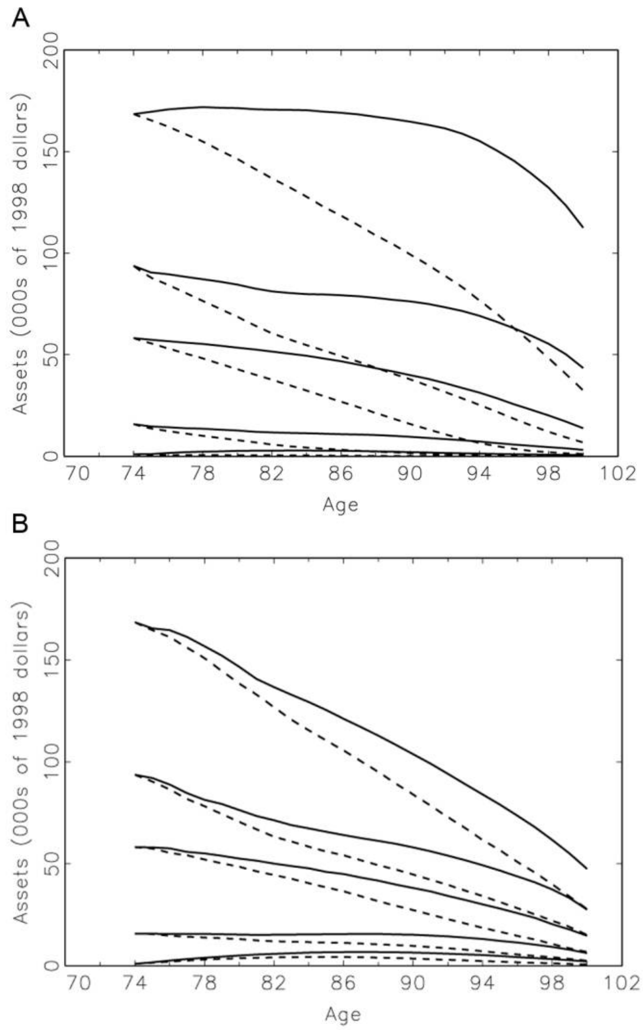

FIG. 14.—Median assets: baseline model (dashed lines) and model with 50 percent of the consumption floor (solid lines): *A*, results for models with exogenous medical expenses; *B*, results for models with endogenous medical expenses.

If single people live to very advanced ages, they are almost sure to need expensive medical care, and they thus choose to keep a large amount of assets (this amount increases in permanent income as medical expenses also increase) to self-insure against this risk. We also find that a publicly provided consumption floor has a large effect on the asset profiles of all people, even those with high income. Our findings are robust to endogenizing medical expenditures in an empirically realistic way.

In short, we find that out-of-pocket medical expenditures, and the way in which they interact with the consumption floor, go a long way toward explaining the elderly's saving decisions and should be accounted for when considering old-age policy reforms.

Following up on our findings, several recent papers have incorporated medical expenses and government insurance programs in general equilibrium, heterogeneous agent, life cycle settings and used these models to study policy reforms. Kopecky and Koreshkova (2009) consider nursing home expenses and Medicaid. İmrohoroğlu and Kitao (2009) and Attanasio, Kitao, and Violante (2010) evaluate Social Security reforms and Medicare, respectively, and Paschenko (2009), Pang and Warshawsky (2010), and Peijnenburg, Nijman, and Werker (2010) study why so few retirees purchase annuities.

References

Altonji, Joseph G., and Lewis M. Segal. 1996. "Small Sample Bias in GMM Estimation of Covariance Structures." *J. Bus. and Econ. Statis.* 14 (Fall): 353–66.

Ameriks, John, Andrew Caplin, Steven Laufer, and Stijn Van Nieuwerburgh. 2007. "The Joy of Giving or Assisted Living? Using Strategic Surveys to Separate Bequest and Precautionary Motives." Working Paper no. 13105 (May), NBER, Cambridge, MA.

Attanasio, Orazio P., James Banks, Costas Meghir, and Guglielmo Weber. 1999. "Humps and Bumps in Lifetime Consumption." *J. Bus. and Econ. Statis.* 17 (January): 22–35.

Attanasio, Orazio P., and Carl Emmerson. 2003. "Mortality, Health Status, and Wealth." *J. European Econ. Assoc.* 4 (March): 821–50.

Attanasio, Orazio P., Sagiri Kitao, and Gianluca Violante. 2010. "Financing Medicare: A General Equilibrium Analysis." In *Demography and the Economy*, edited by John B. Shoven. Chicago: Univ. Chicago Press, forthcoming.

Banks, James, Richard Blundell, and Sarah Tanner. 1998. "Is There a Retirement-Savings Puzzle?" *A.E.R.* 88 (September): 769–88.

Blau, David M. 2008. "Retirement and Consumption in a Life Cycle Model." *J. Labor Econ.* 26 (January): 35–71.

Blau, David M., and Donna Gilleskie. 2008. "The Role of Retiree Health Insurance in the Employment Behavior of Older Men." *Internat. Econ. Rev.* 49 (May): 475–514.

Brook, Robert, et al. 1983. "Does Free Care Improve Adults' Health? Results from a Randomized Trial." *New England J. Medicine* 309 (December): 1426–34.

Brown, Jeff, and Amy Finkelstein. 2008. "The Interaction of Public and Private Insurance: Medicaid and the Long Term Care Insurance Market." *A.E.R.* 98 (June): 1083–1102.

Buchinsky, Moshe. 1998. "Recent Advances in Quantile Regression Models: A Practical Guideline for Empirical Research." *J. Human Resources* 33 (Winter): 88–126.

Cagetti, Marco. 2003. "Wealth Accumulation over the Life Cycle and Precautionary Savings." *J. Bus. and Econ. Statis.* 21 (July): 339–53.

Card, David, Carlos Dobkin, and Nicole Maestas. 2009. "Does Medicare Save Lives?" *Q.J.E.* 124 (May): 597–636.

Chamberlain, Gary. 1992. "Comment: Sequential Moment Restrictions in Panel Data." *J. Bus. and Econ. Statis.* 10 (Winter): 20–26.

Davis, Morris. 2005. "The Insurance, Health, and Savings Decisions of Elderly Women Living Alone." Manuscript, Dept. Econ., Univ. Pennsylvania.

De Nardi, Mariacristina. 2004. "Wealth Inequality and Intergenerational Links." *Rev. Econ. Studies* 71 (July): 743–68.

De Nardi, Mariacristina, Eric French, and John B. Jones. 2009. "Life Expectancy and Old Age Savings." *A.E.R. Papers and Proc.* 99 (May): 110–15.

Duffie, Darrell, and Kenneth J. Singleton. 1993. "Simulated Moments Estimation of Markov Models of Asset Prices." *Econometrica* 61 (July): 929–52.

Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes. 2004. "Do the Rich Save More?" *J.P.E.* 112 (April): 397–444.

Elder, Todd E., and Elizabeth T. Powers. 2006. "Public Health Insurance and SSI Program Participation among the Aged." Working Paper no. 117 (May), Michigan Retirement Res. Center, Ann Arbor.

Fernández-Villaverde, Jesús, and Dirk Krueger. 2007. "Consumption over the Life Cycle: Facts from CEX Data." *Rev. Econ. and Statis.* 89 (August): 552–65.

Finkelstein, Amy, and Robin McKnight. 2005. "What Did Medicare Do (and Was It Worth It)?" Working Paper no. 11609 (September), NBER, Cambridge, MA.

French, Eric, Phil Doctor, and Olesya Baker. 2007. "Asset Rundown after Retirement: The Importance of Rate of Return Shocks." *Econ. Perspectives* 31 (Spring): 48–65.

French, Eric, and John Bailey Jones. 2004. "On the Distribution and Dynamics of Health Care Costs." *J. Appl. Econometrics* 19 (November): 705–21.

———. 2007. "The Effects of Health Insurance and Self-Insurance on Retirement Behavior." Working Paper no. 170 (October), Michigan Retirement Res. Center, Ann Arbor.

Gouriéroux, Christian, and Alain Monfort. 1996. *Simulation-Based Econometric Methods.* Oxford: Oxford Univ. Press.

Gourinchas, Pierre-Olivier, and Jonathan A. Parker. 2002. "Consumption over the Life Cycle." *Econometrica* 70 (January): 47–89.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1994. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." *Carnegie-Rochester Conf. Ser. Public Policy* 40:59–125.

———. 1995. "Precautionary Saving and Social Insurance." *J.P.E.* 103 (April): 360–99.

Hurd, Michael D. 1989. "Mortality Risk and Bequests." *Econometrica* 57 (July): 779–813.

Hurd, Michael D., Daniel McFadden, and Angela Merrill. 2001. "Predictors of Mortality among the Elderly." In *Themes in the Economics of Aging*, edited by David A. Wise, 171–97. Chicago: Univ. Chicago Press.

İmrohoroğlu, Selahattin, and Sagiri Kitao. 2009. "Social Security, Benefit Claiming and Labor Force Participation: A Quantitative General Equilibrium Approach." Manuscript, Marshall School Bus., Univ. Southern California.

Khwaja, Ahmed. Forthcoming. "Estimating Willingness to Pay for Medicare Using a Dynamic Life-Cycle Model of Demand for Health Insurance." *J. Econometrics.*

Koning, Ruud H. 1996. “Kernel: A Gauss Library for Kernel Estimation.” http://www.xs4all.nl/.

Kopczuch, Wojciech, and Joseph Lupton. 2007. “To Leave or Not to Leave: The Distribution of Bequest Motives.” *Rev. Econ. Studies* 74 (January): 207–35.

Kopecky, Karen, and Tatyana Koreshkova. 2009. “The Impact of Medical and Nursing Home Expenses and Social Insurance Policies on Savings and Inequality.” Manuscript, Dept. Econ., Univ. Western Ontario.

Kotlikoff, Laurence J. 1988. “Health Expenditures and Precautionary Savings.” In *What Determines Savings?* Cambridge, MA: MIT Press.

Love, David A., Michael Palumbo, and Paul Smith. 2009. “The Trajectory of Wealth in Retirement.” *J. Public Econ.* 93 (February): 191–208.

Manski, Charles F. 1988. *Analog Estimation Methods in Econometrics.* London: Chapman & Hall.

Newey, Whitney K. 1985. “Generalized Method of Moments Specification Testing.” *J. Econometrics* 29 (September): 229–56.

Newey, Whitney K., and Daniel L. McFadden. 1994. “Large Sample Estimation and Hypothesis Testing.” In *Handbook of Econometrics*, vol. 4, edited by Robert F. Engle and Daniel L. McFadden. Amsterdam: Elsevier.

Pakes, Ariel, and David Pollard. 1989. “Simulation and the Asymptotics of Optimization Estimators.” *Econometrica* 57 (September): 1027–57.

Palumbo, Michael G. 1999. “Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle.” *Rev. Econ. Studies* 66 (April): 395–421.

Pang, Gaobo, and Mark Warshawsky. 2010. “Optimizing the Equity-Bond-Annuity Portfolio in Retirement: The Impact of Uncertain Health Expenses.” *Insurance: Math. and Econ.* 46:198–209.

Paschenko, Svetlana. 2009. “Accounting for Non-annuitization. “Manuscript, Dept. Econ., Univ. Virginia.

Peijnenburg, Kim, Theo Nijman, and Bas J. M. Werker. 2010. “Health Cost Risk and Optimal Retirement Provision: A Simple Rule for Annuity Demand.” Manuscript, Tilburg Univ.

Pischke, Jorn-Steffen. 1995. “Measurement Error and Earnings Dynamics: Some Estimates from the PSID Validation Study.” *J. Bus. and Econ. Statis.* 13 (Summer): 305–14.

Powell, James. 1994. “Estimation of Semiparametric Models.” In *Handbook of Econometrics*, vol. 4, edited by Robert F. Engle and Daniel L. McFadden. Amsterdam: Elsevier.

Rohwedder, Susann, Steven J. Haider, and Michael Hurd. 2004. “Increases in Wealth among the Elderly in the Early 1990s: How Much Is Due to Survey Design?” Working Paper no. 10862 (October), NBER, Cambridge, MA.

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun. 2006. “Are Americans Saving Optimally for Retirement?” *J.P.E.* 114 (August): 607–43.

Yogo, Motohiro. 2009. “Portfolio Choice in Retirement: Health Risk and the Demand for Annuities, Housing, and Risky Assets.” Working Paper no. 15307 (September), NBER, Cambridge, MA.