NBER WORKING PAPER SERIES

MEDICAID INSURANCE IN OLD AGE

Mariacristina De Nardi
Eric French
John Bailey Jones

Working Paper 19151
http://www.nber.org/papers/w19151

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2013

For useful comments and suggestions we thank Norma Coe, Amy Finkelstein, Victor Rios-Rull, Karl Scholz, Jon Skinner, and Gianluca Violante. Jones gratefully acknowledges financial support from the Social Security Administration through the Michigan Retirement Research Center (MRRC grants UM10-16 and UM12-12). The views expressed in this paper are those of the authors and not necessarily those of the Federal Reserve Bank of Chicago, the Social Security Administration or the MRRC. Jones gratefully acknowledges the hospitality of the Federal Reserve Bank of Chicago. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Medicaid Insurance in Old Age
Mariacristina De Nardi, Eric French, and John Bailey Jones
NBER Working Paper No. 19151
June 2013
JEL No. D11,D14,D31,E21,H2,I14

ABSTRACT

The old age provisions of the Medicaid program were designed to insure poor retirees against medical expenses. However, it is the rich who are most likely to live long and face expensive medical conditions when very old. We estimate a rich structural model of savings and endogenous medical spending with heterogeneous agents, and use it to compute the distribution of lifetime Medicaid transfers and Medicaid valuations across single retirees.

We find that retirees with high lifetime incomes can end up on Medicaid, and often value Medicaid's insurance features the most, as they face a larger risk of catastrophic medical needs at old ages, and face the greatest consumption risk. Finally, our compensating differential calculations indicate that retirees value Medicaid insurance at more than its actuarial cost, but that most would value expansions of the current Medicaid program at less than cost.

Mariacristina De Nardi
Federal Reserve Bank of Chicago
230 South LaSalle St.
Chicago, IL 60604
and NBER
denardim@nber.org

Eric French
Research Department
Federal Reserve Bank of Chicago
230 South LaSalle Street
Chicago, IL 60604
efrench@frbchi.org

John Bailey Jones
Department of Economics
BA-113B
University at Albany
State University of New York
Albany, NY 12222
jbjones@albany.edu

1 Introduction

Large and persistent government deficits have made it clear that most entitlement programs in the United States will be scrutinized for cost-saving reforms. One of the most debated programs is Medicaid, a means-tested, public health insurance program that covers medical expenses not covered by other insurance programs.

Despite the increasing importance (and cost) of Medicaid in the presence of an aging population and rising medical costs, very little is known about how the benefits of Medicaid are distributed among the elderly and about their valuation. Which elderly households receive Medicaid transfers? How redistributive are these transfers? What is the insurance value of these transfers? Is Medicaid of about the right size? How much would people lose if it were cut? These are important questions to answer before reforming the programs currently in place. This paper seeks to fill this gap.

It has been argued that Medicaid has outgrown its initial mandate, (e.g. Brown and Finkelstein [7]) and is now insuring middle- and higher-income retirees as well as lower-income ones. In fact, although Medicaid assists the lifetime poor, it also assists richer people impoverished by nursing home and other medical expenses not covered by other public or private insurance. This is an important feature of the program because it is the rich who are more likely to live long and face expensive medical conditions when very old.

In this paper, we first document who in the Assets and Health Dynamics of the Oldest Old (AHEAD) data receives Medicaid. We find that even high income people become Medicaid recipients if they live long enough and are hit by expensive medical conditions. The Medicaid recipiency rate in the bottom income quintile stays around 60%-70% throughout retirement. In contrast, the recipiency rate of higher-income retirees is initially very low, but it increases by age, reaching 20% by age 95. In addition, data from the Medicare Current Beneficiary Survey (MCBS) data set shows that high income individuals, conditional on receiving Medicaid transfers, receive larger payments than low income individuals.

Then, taking life expectancy and other important dimensions of heterogeneity into account, we estimate a structural model of savings and endogenous medical expenses for single retirees. Consistent with the institutions, we explicitly model two separate ways to become Medicaid eligible: having low income and assets, and becoming impoverished by high medical needs. We require our model to match some key aspects of the data, such as savings, out-of-pocket medical expenses, and Medicaid recipiency rates. Including Medicaid recipiency in the moments being matched adds an unexpected identification angle to bequest motives: to match Medicaid recipiency rates, Medicaid payments cannot be too low. If Medicaid payments are of a reasonable size compared to the data, retirees face less risk. To reconcile observed assets with reduced medical expenses, a bequest motive is needed.

Our model matches key aspects of the data well and produces parameter estimates within the bounds established by previous work. It also generates an elasticity of total medical expenditures to co-payment changes that is close to the one estimated by Manning et al. [39] using the RAND Health Insurance Experiment. Moreover, although our model was not required to match the distribution of out-of-pocket and total medical expenditures, and Medicaid payments, it turns out to match well the corresponding data from the MCBS survey.

Finally, we use our estimated model to assess the distribution and benefits of Medicaid. We compute how Medicaid payments vary by age, gender, permanent income, and health status. We find that the current Medicaid system provides different kinds of insurance to households with different resources. Households in the lower permanent income quintiles are much more likely to receive Medicaid transfers, but the transfers that they receive are on average relatively small. Households in the higher permanent income quintiles are much less likely to receive any Medicaid transfers, but when they do, these transfers are very big and correspond to severe and expensive medical conditions. Therefore, and consistent with the MCBS data, Medicaid is an effective insurance device for the poorest, but also offers valuable insurance to the rich, by insuring them against catastrophic medical conditions, which are the most costly in terms of utility and the most difficult to insure in the private market.

Our model also allows us to compute the value retirees place on Medicaid insurance, thus enabling us to perform a cost and benefit analysis. We find that, with moderate risk aversion and realistic lifetime and medical needs risk, the value most retirees place on Medicaid (the benefit) exceeds the actuarial value of their expected benefits (the cost). In many cases, it is the richer retirees, who have the most to lose, who value Medicaid most highly. On the other hand, we find that a Medicaid expansion would be valued by most retirees at less than its cost.

Our findings come from a life-cycle model of consumption and endogenous medical expenditure that accounts for Medicare, Supplemental Social Insurance (SSI), and Medicaid. Agents in the model are face uncertainty about their health, lifespan, and medical needs (including nursing home stays). This uncertainty is partially offset by the insurance provided by the government and private institutions. Agents choose whether they want to apply for Medicaid if they are eligible, how much to save, and how to split their consumption between medical and non-medical goods. Consistent with program rules, we model two pathways to Medicaid, the one for the lifelong poor, and that for people impoverished by large medical expenses.

To appropriately evaluate Medicaid redistribution, we allow for heterogeneity in wealth, permanent income (PI), health, gender, life expectancy, and medical needs. We also require our model to fit well across the entire income distribution, rather than simply explain mean or median behavior. Our model matches the life-cycle profiles of assets, out-of-pocket medical spending, and Medicaid recipiency rates for elderly singles in different cohorts and permanent income groups.

2 Literature review

This paper is related to several previous papers on savings, health risks, and social insurance. Hurd [30] and Hurd, McFadden, Merrill [31] highlight the importance of accounting for the link between wealth and mortality when estimating life-cycle models. Kotlikoff [37] stresses the importance of modeling health expenditures when studying precautionary savings.

Hubbard et al. [28] and Palumbo [49] solve dynamic programming models of saving under medical expense risk, and find that medical expenses have relatively small effects. These papers likely underestimated medical spending risk, however, because the data sets available at that time were missing late-in-life medical spending and had poor measures of nursing home costs. As a result, the data understated the extent to which medical expenses rise with age and income. De Nardi et al. [15] and Marshall, McGarry, and Skinner [41] find that late-in-life medical expenses are large and generate powerful savings incentives. Furthermore, Poterba, Venti, and Wise [52] show that those in poor health have considerably lower assets than similar individuals in good health. Lockwood [38], Nakajima and Telyukova [43], and Yogo [57] add to the literature by estimating life cycle models that include additional insurance choices, housing, and portfolio choices respectively.

De Nardi et al. [15] and [14] focus on the role of medical expense risk in shaping savings. This paper extends their endogenous medical spending framework and focuses on the role of Medicaid. Specifically, the paper assesses what groups of individuals benefit from the Medicaid program, and how much they value these benefits. In order to do answer these questions, we develop a more realistic model of Medicaid eligibility. To better capture key aspects of the Medicaid program, we match Medicaid eligibility rates which adds an important new source of identification. Because approximately $\frac{2}{3}$ of Medicaid payments are to those in a nursing home, we model the nursing home state explicitly. Furthermore, we compare Medicaid payments predicted by the model to those observed in the Medicare Current Beneficiary Survey (MCBS). We show that our model matches Medicaid payment flows well, although they are not matched by construction. This provides additional validation that the model is useful for Medicaid policy evaluation.

Hubbard et al. [29] and Scholz et al. [56] 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. Consistent with this evidence, Gardner and Gilleskie [25] exploit cross-state variation in Medicaid rules and find Medicaid has significant effects on savings. Brown and Finkelstein [7] develop a dynamic model of optimal savings and long-term care purchase decisions. They conclude that Medicaid crowds out private long-term care insurance for about two-thirds of the wealth distribution. Consistent with this evidence, Brown et al. [8] exploit cross-state variation in Medicaid rules and also find significant crowding out.

Several new papers (Hansen et al. [27], Paschenko and Porapakkarm [50], İmrohoroğlu and Kitao [32]) study the importance of medical expense risk in the aggregate. Kopecky and Koreshkova [36] find that old-age medical expenses, and the coverage of these expenses provided by Medicaid, have large effects on aggregate capital accumulation. Braun et al. [4] use a model with medical expense risk to assess the incentive and welfare effects of Social Security and other social programs. We focus on redistribution, and behavior and valuation at the individual level. Hence, consistent with the data, we use a partial equilibrium model that allows for much more heterogeneity. In addition, in our model people can adjust medical spending (as well as consumption and savings), and we estimate our model, rather than calibrating it.

We model endogenous medical expenditures so that we can consider individuals' valuation of quality of care. Some recent papers also contain life-cycle models where the choice of medical expenditures is endogenous. In addition to having different emphases, these papers model Medicaid in a more stylized way. Fonseca et al. [22] and Scholz and Seshadri [55] assume that the consumption floor is invariant to medical needs, whereas our specification allows for a more realistic link between medical needs and Medicaid transfers. Ozkan [47] studies health investments over the life cycle, but does not focus on the role of Medicaid.

This paper also contributes to the literature on the redistribution generated by government programs. Although there is a lot of research about the amount of redistribution provided by Social Security and a smaller amount of research about Medicare, to the best of our knowledge this is the first paper to comprehensively examine how Medicaid transfers to the elderly are distributed across income groups, and to document how even people with higher lifetime income can end up on Medicaid. Furthermore, we assess the valuation individuals place on their Medicaid benefits.[1] In this paper, we focus on the redistribution generated by Medicaid benefits and their valuation. Unlike Social Security, unemployment benefits, and disability insurance, Medicaid is not financed using a specific tax, but by general government revenue, making it difficult to determine how redistributive "Medicaid taxes" are.

3 Key features of the Medicaid program

In the United States, there are two major public insurance programs helping the elderly with their medical expenses. The first one is Medicare, a federal program that provides health insurance to almost every person over the age of 65. The second one is Medicaid, a means-tested program that is run jointly by the federal and state governments.[2]

An important characteristic of Medicaid is that it is the payer of "last resort": Medicaid contributes only after Medicare and private insurance pay their share, and

[1] Using a simpler, calibrated model, Brown and Finkelstein [7] analyze how Medicaid affects the valuation of long-term care insurance.

[2] De Nardi et al. [16] and Gardner and Gilleskie [25] document many important aspects of Medicaid insurance in old age.

the individual spends down his assets to a "disregard" amount. Whereas non-means-tested insurance reduces savings only by reducing risks, Medicaid's asset test provides an additional savings disincentive.

One area where Medicaid is particularly important is long-term care. Medicare reimburses only a limited amount of long-term care costs, and most elderly people do not have private long-term care insurance. As a result, Medicaid covers almost all nursing home costs of poor old recipients. More generally, Medicaid ends up financing 70% of nursing home residents (Kaiser Foundation [46]), and these costs are of the order of $60,000 to $75,000 a year (in 2005). Furthermore, 62% of Medicaid's $81 billion per year transfers for the elderly in 2009 are for nursing home payments (Kaiser Foundation [23]).

Medicaid-eligible individuals can be divided into two main groups. The first group comprises the *categorically needy*, whose income and assets fall below certain thresholds. People who receive SSI typically qualify under the categorically needy provision. The second group comprises the *medically needy*, who are individuals whose income is not particularly low, but who face such high medical expenditures that their financial resources are small in comparison.

The categorically needy provision thus affects the saving of people who have been poor throughout most of their lives, but has no impact on the saving of middle- and upper-income people. The medically needy provision, instead, provides insurance to people with higher income and assets who are still at risk of being impoverished by expensive medical conditions.

4 Some Data

We use two main data sets, the AHEAD and the MCBS. We now turn to discussing the main features of each.

4.1 The AHEAD dataset

We use data from the Assets and Health Dynamics of the Oldest Old (AHEAD) data set. The AHEAD is a survey of individuals who were non-institutionalized and aged 70 or older in 1994. It is part of the Health and Retirement Survey (HRS) conducted by the University of Michigan. We consider only single (i.e., never married, divorced, or widowed) retired individuals. A total of 3,872 singles 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, 2006, 2008, and 2010. This leaves us with 3,243 individuals, of whom 588 are men and 2,655 are women. Of these 3,243 individuals, 370 are still alive in 2010. We do not use 1994 assets or medical expenses. Assets in 1994 were underreported (Rohwedder et al. [54]) and medical expenses appear to be underreported as well.

A key advantage of the AHEAD relative to other datasets is that it provides panel data on health status, including nursing home stays. We assign individuals a health status of "good" if self-reported health is excellent, very good or good, and are assigned a health status of "bad" if self-reported health is fair or poor. We assign individuals to the nursing home state if they were in a nursing home at least 120 days since the last interview or if they spent at least 60 days in a nursing home before the next scheduled interview and died before that scheduled interview.

We break the data into 5 cohorts. The first cohort consists of individuals that were ages 72-76 in 1996; the second cohort contains ages 77-81; the third ages 82-86; the fourth ages 87-91; and the final cohort, for sample size reasons, contains ages 92-102. We calculate summary statistics (e.g., medians), cohort-by-cohort, for surviving individuals in each calendar year—we use an unbalanced panel. We then construct life-cycle profiles by ordering the summary statistics by cohort and age at each year of observation. Moving from the left-hand-side to the right-hand-side of our graphs, we thus show data for four cohorts, with each cohort's data starting out at the cohort's average age in 1996. Our graphs omit profiles for the oldest cohort because the sample sizes for this cohort are tiny.

Since we want to understand the role of income, we further stratify the data by post-retirement permanent income (PI). Hence, for each cohort our graphs usually display several horizontal lines showing, for example, average Medicaid status in each cohort and PI group in each calendar year. These lines also identify the moment conditions we use when estimating the model. To indicate PI rank, we vary the thickness of the lines on our graphs: thicker lines represent observations for higher-ranked PI groupings.

We measure post-retirement PI as the individual's average non-asset income over all periods during which he or she is observed. Non-asset income includes the value of Social Security benefits, defined benefit pension benefits, veterans benefits and annuities. Since we model social insurance explicitly, we do not include SSI transfers. Because there is a roughly monotonic relationship between lifetime earnings and the income variables that we use, our measure of post-retirement PI is also a good measure of lifetime permanent income.

4.2 Medicaid Recipiency

AHEAD respondents are asked whether they are currently covered by Medicaid. Figure 1 plots the fraction of the sample receiving Medicaid by age, birth cohort and income quintile for all the individuals alive at each moment in time. There are four lines representing PI groupings within each cohort. We split the data into PI quintiles, but then merge the richest two quintiles together because at younger ages no one in the top PI quintile is on Medicaid.

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

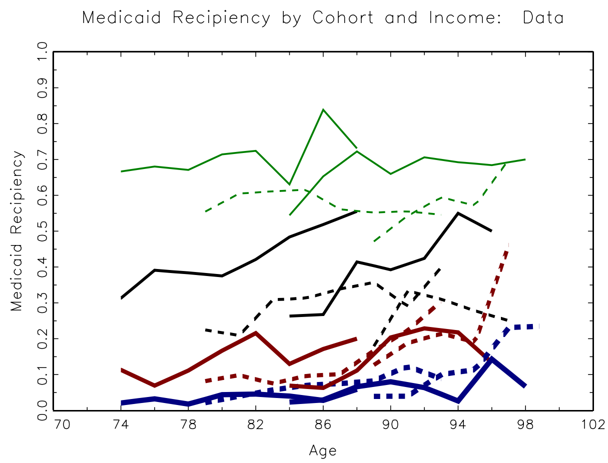

Figure 1: Medicaid recipiency rates by age, cohort, and permanent income. Thicker lines refer to higher PI groups.

then again every two years until 2010. The other cohorts start from older initial ages and are also followed for fourteen years. The graph reports the Medicaid recipiency rate for each cohort and PI grouping at eight dates over time.

Unsurprisingly, Medicaid recipiency is inversely related to permanent income: the thin top line shows the fraction of Medicaid recipients in the bottom 20% of the permanent income distribution, while the thick bottom line shows median assets in the top 40%. For example, the top left line shows that for the bottom PI quintile of the cohort aged 74 in 1996, about 70% of the sample receives Medicaid in 1996; this fraction stays rather stable over time. This is because the poorest people qualify for Medicaid under the categorically needy provision, where eligibility depends on income and assets, but not the amount of medical expenses.

The Medicaid recipiency rate tends to rise with age most quickly for people in the middle and highest PI groups. For example, Medicaid recipiency in the oldest cohort and top two permanent income quintiles rises from about 4% at age 89 to over 20% at age 96. Even people with relatively large resources can be hit by medical shocks severe enough to exhaust their assets and qualify them for Medicaid under the medically needy provision.

4.3 Medical expense profiles

In all waves, AHEAD respondents are asked about the medical expenses they paid out-of-pocket. Out-of-pocket medical expenses are the sum of what the individual spends out-of-pocket on insurance premia, drug costs, and costs for hospital, nursing home care, doctor visits, dental visits, and outpatient care. It includes medical expenses during the last year of life. It does not include expenses covered by insurance, either public or private.

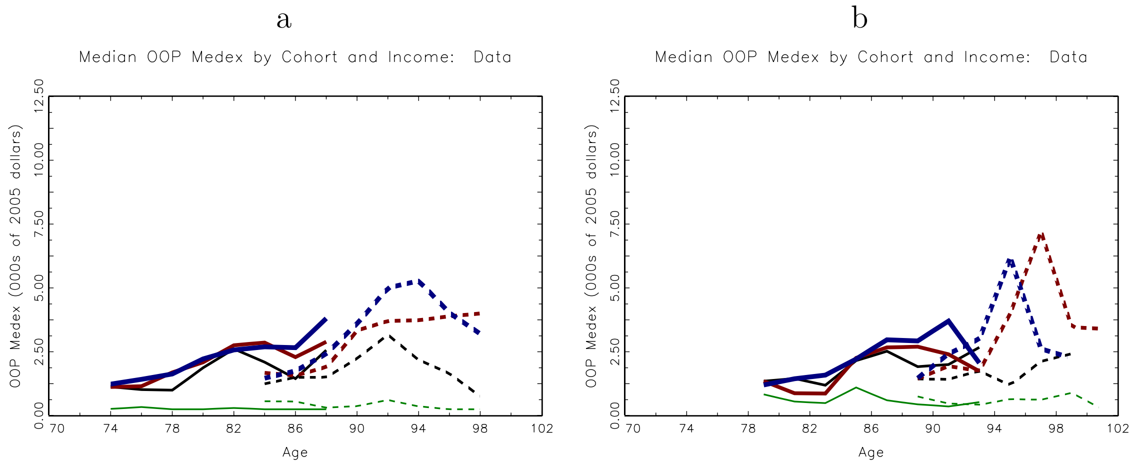

Figure 2: Median out-of-pocket medical expenditures by age, cohort, and permanent income. Thicker lines refer to higher PI groups.

French and Jones [24] show that the medical expense data in the AHEAD line up with the aggregate statistics. For our sample, mean medical expenses are $4,605 with a standard deviation of $14,450 in 2005 dollars. Although this figure is large, it is not surprising, because Medicare did not cover prescription drugs for most of the sample period, requires co-pays for services, and caps the number of reimbursed nursing home and hospital nights.

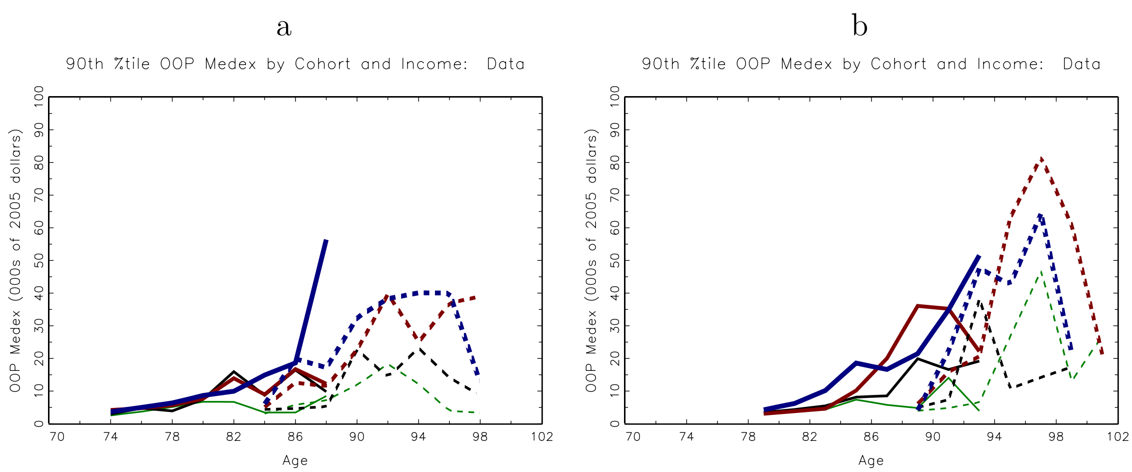

Figure 3: 90th percentile out-of-pocket medical expenditures by age, cohort, and permanent income. Thicker lines refer to higher PI groups.

Figures 2 and 3 display the median and 90th percentile of the out-of-pocket medical expense distribution, respectively. The bottom two quintiles of permanent income are merged as there is very little variation in out-of-pocket medical expenses in the lowest quintile until very late in life: at younger ages, most of the expenses in the bottom quintile are bottom-coded at $250. The graphs highlight the large increase in out-of-pocket medical expenses that occurs as people reach very advanced ages, and show that this increase is especially pronounced for people in the highest PI quintiles.

4.4 Net worth profiles

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, T-bills, etc.), IRAs, Keoghs, stocks, the value of a farm or business, mutual funds, bonds, and "other" assets.

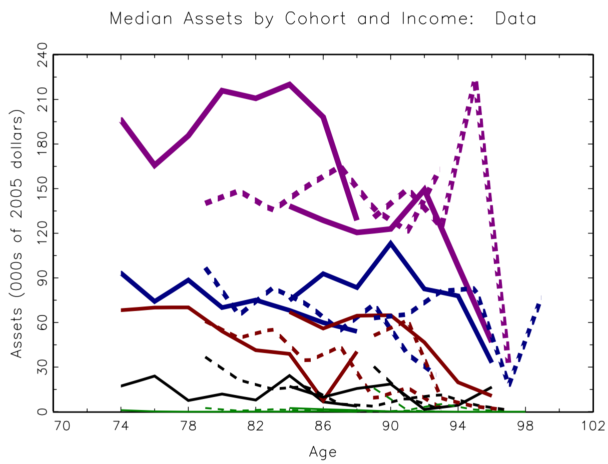

Figure 4: Median assets by age, cohort, and permanent income. Thicker lines refer to higher PI groups.

Figure 4 reports median assets by cohort, age, and PI quintile. However, the fifth, bottom line is hard to distinguish from the horizontal axis because households in this PI quintile hold few assets. Unsurprisingly, assets turn out to be monotonically increasing in income, so that the thin bottom line shows median assets in the lowest PI quintile, while the thick top line shows median assets in the top quintile. For example, the top left line shows that for the top PI quintile of the cohort age 74 in 1996, median assets started at $200,000 and then stayed rather stable until the final time period: $170,000 at age 76, $190,000 at age 78, $220,000 at age 80, $210,00 at age 82, $220,000 at age 84, $200,00 at age 86, and $130,000 at age 88.[3]

[3]The jumps in the profiles are due to the fact that there is dispersion in assets within a cell, and very rapid attrition due to death, especially at very advanced ages. For example, for the highest permanent income grouping in the oldest cohort, the cell count goes from 29 observations, to 20, and finally to 12 toward the end of the sample. Our GMM criterion weights each moment condition in proportion to the number of observations, so these cells have little effect on the GMM criterion function and thus the estimates.

For all PI quintiles in these cohorts, the assets of surviving individuals do not decline rapidly with age. Those with high income do not run down their assets until their late 80s, although those with low income tend to have their assets decrease throughout the sample period. The slow rate at which the elderly deplete their wealth has been a long-standing puzzle (see for example, Mirer [42]). However, as De Nardi, French, and Jones [15] show, the risk of medical spending rising with age and income goes a long way toward explaining this puzzle.

4.5 The MCBS dataset

An important limitation of the AHEAD data is that it lacks information on other payors of medical care, such as Medicaid and Medicare. Although there there are some self-reported survey data on total billable medical expenditures in the AHEAD, these data are mostly imputed, and are considered to be of low quality. To circumvent this issue, we use data from the 1996-2006 waves of the Medicare Current Beneficiary Survey (MCBS).

The MCBS is a nationally representative survey of disabled and elderly Medicare beneficiaries. Respondents are asked about health status, health insurance, and health care expenditures made out-of-pocket, by Medicaid, by Medicare and by other sources. The MCBS data are matched to Medicare records, and medical expenditure data are created through a reconciliation process that combines survey information with Medicare administrative files. As a result, the survey gives extremely accurate data on Medicare payments and fairly accurate data on out-of-pocket and Medicaid payments. As in the AHEAD survey, the MCBS survey includes information on those who enter a nursing home or die. This is an important advantage of the MCBS relative to the Medical Expenditure Panel Survey (MEPS), which does not capture late-life or nursing home expenses. The reason why MEPS is so widely used is that obtaining the MCBS is difficult and costly to obtain.

MCBS Respondents are interviewed up to 12 times over a 4 year period, forming short panels. We aggregate the data to an annual level. We use the same sample selection rules in the MCBS as we use for the AHEAD data. Specifically, we drop those who were ever observed to be marry, work, or be younger than 72 in 1996, 74 in 1998, etc. These sample selection procedures leave us 15,041 different individuals who contribute 34,343 person-year observations. Details of sample construction, as well as validation of the MCBS relative to the aggregate national statistics, are in Appendix A.

As with the AHEAD data, we assign individuals a health status of "good" if self-reported health is excellent, very good or good, and are assigned a health status of "bad" if self-reported health is fair or poor. We define an individual as being in a nursing home if that individual was in a nursing home at least 60 days over the year. However, the income data in the MCBS is limited. Individuals are asked about total income, not annuitized income. Nevertheless, we found that this variable lines up well with total income in the AHEAD. Furthermore, the correlation between total income and annuitized income in the AHEAD is 0.8. We use average total income over the time we observe the individual as our measure of permanent income in the MCBS.

We use MBCS data set to measure co-pay rates and to compare model predicted payments to the data.

5 The model

We focus on single people, male or female, who have already retired. This allows us to abstract from labor supply decisions and from complications arising from changes in family size.

5.1 Preferences

Individuals in this model receive utility from the consumption of both non-medical and medical goods. Each period, their flow utility is given by

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

where t is age, c_t is consumption of non-medical goods, m_t is total consumption of medical goods, and $\mu(\cdot)$ is the medical needs shifter, which affects the marginal utility of consuming medical goods and services. The consumption of both goods is expressed in dollar values. The intertemporal elasticities for the two goods, $1/\nu$ and $1/\omega$, can differ.

We assume that $\mu(\cdot)$ shifts with medical needs, such as dementia, arthritis, or a broken bone. These shocks affect the utility of consuming medical goods and services, including nursing home care. Formally, we model $\mu(\cdot)$ as a function of age, the discrete-valued health status indicator h_t, and the medical needs shocks ζ_t and ξ_t. Individuals optimally choose how much to spend in response to these shocks.

A complementary approach is that of Grossman [26], in which medical expenses represent investments in health capital, which in turn decreases mortality (e.g., Yogo [57]) or improves health. While a few studies find that medical expenditures have significant effects on health and/or survival (Card et al. [10]; Doyle [13], Finkelstein et al. [20], Chay et al. [12]), most others find small effects (Brook et al. [5]; Fisher et al. [21]; Finkelstein and McKnight [19]; Khwaja [33]); see De Nardi et al. [15] for a discussion. These findings suggest that the effects of medical expenditures on the health outcomes are, at a minimum, extremely difficult to identify. Identification problems include reverse causality (sick people have higher health expenditures) and lack of insurance variation (most elderly individuals receive baseline coverage through Medicare). Given that older people have already shaped their health and lifestyle, we view our assumption that their health and mortality depend on their lifetime earnings, but is exogenous to their current decisions, to be a reasonable simplification.

5.2 Insurance Mechanisms

We model two important types of health insurance. The first one pays a proportional share of total medical expenses and can be thought of as a combination of Medicare and private insurance. Let $q(h_t)$ denote the individual's co-insurance (co-pay) rate, i.e., the share of medical expenses not paid by Medicare or private insurance. We allow the co-pay rate to depend on whether a person is in a nursing home ($h_t = 1$) or not. Because nursing home stays are virtually uninsured by Medicare and private insurance, people residing in nursing homes face much higher co-pay rates. However, co-pay rates do not vary much across other medical conditions.

The second type of health insurance that we model is Medicaid, which is means-tested. To link Medicaid transfers to medical needs, $\mu(h_t, \zeta_t, \xi_t, t)$, we assume that each period Medicaid guarantees a minimum level of flow utility $\underline{u}_i$, which potentially differs between categorically needy ($i = c$) and medically needy ($i = m$) recipients. In practice, the floors for categorically and medically needy recipients are very similar, and we will set them equal in the estimation. We will allow the floors to differ, however, in some policy experiments.

More precisely, once the Medicaid transfer is made, an individual with the state vector (h_t, ζ_t, ξ_t, t) can afford a consumption-medical goods pair (c_t, m_t) such that

$$\underline{u}_i = \frac{1}{1-\nu} c_t^{1-\nu} + \mu(h_t, \zeta_t, \xi_t, t) \, \frac{1}{1-\omega} m_t^{1-\omega}. \tag{2}$$

To implement our utility floor, for every value of the state vector, we find the expenditure level $\underline{x}_i = c_t + m_t q(h_t)$ needed to achieve the utility level $\underline{u}_i$ (equation (2)), assuming that individuals make intratemporally optimal decisions. This yields the minimum expenditure $\underline{x}_c(\cdot)$ or $\underline{x}_m(\cdot)$, which correspond to the categorically and medically needy floors. The actual amount that Medicaid transfers, $b_c(a_t, y_t, h_t, \zeta_t, \xi_t, t)$ or $b_m(a_t, y_t, h_t, \zeta_t, \xi_t, t)$, is then given by $\underline{x}_c(\cdot)$ or $\underline{x}_m(\cdot)$ less the individual's total financial resources (assets, a_t, and non-asset income, y_t).

5.3 Uncertainty and Non-Asset Income

The individual faces several sources of risk, which we treat as exogenous: health status risk, survival risk, and medical needs risk. At the beginning of each period, the individual's health status and medical needs shocks are realized, and need-based transfers are determined. The individual then chooses consumption, medical expenditure, and savings. Finally, the survival shock hits.

We parameterize the preference shifter for medical goods and services (the needs shock) as

$$\begin{aligned}
\log(\mu(\cdot)) &= \alpha_0 + \alpha_1 t + \alpha_2 t^2 + \alpha_3 t^3 + \alpha_4 h_t + \alpha_5 h_t \times t && (3)\\
& \quad + \sigma(h,t) \times \psi_t, && (4)\\
\sigma(h,t)^2 &= \beta_0 + \beta_1 t + \beta_2 t^2 + \beta_4 h_t + \beta_5 h_t \times t, && (5)\\
\psi_t &= \zeta_t + \xi_t, \quad \xi_t \sim N(0, \sigma_\xi^2), && (6)\\
\zeta_t &= \rho_m \zeta_{t-1} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma_\epsilon^2), && (7)\\
\sigma_\xi^2 + \frac{\sigma_\epsilon^2}{1-\rho_m^2} &\equiv 1, && (8)
\end{aligned}$$

where ξ_t and ϵ_t are serially and mutually independent. We thus allow the need for medical services to have temporary (ξ_t) and persistent (ζ_t) shocks. It is worth stressing that we not allow any component of $\mu(\cdot)$ to depend on permanent income; income affects medical expenditures solely through the budget constraint.

Health status can take on three values: good (3), bad (2), and in a nursing home (1). We allow the transition probabilities for health to depend on previous health, sex (g), permanent income (I), and age. The elements of the health status transition matrix are

$$\pi_{j,k,g,I,t} = \Pr(h_{t+1} = k | h_t = j, g, I, t), \quad j, k \in \{1, 2, 3\}. \tag{9}$$

Mortality also depends on health, sex, permanent income and age. 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.

Non-asset income y_t, is a deterministic function of sex, permanent income, and age:

$$y_t = y(g, I, t). \tag{10}$$

5.4 The Individual's Problem

Consider a single person seeking to maximize his or her expected lifetime utility at age t, $t = t_{r+1}, ..., T$, where t_r is the retirement age.

To be categorically needy, a person must be eligible for SSI, by satisfying the SSI income and asset tests:

$$y_t + ra_t - y_d \leq \underline{Y} \text{ and } a_t \leq A_d, \tag{11}$$

where: a_t denotes assets; r is the real interest rate; $\underline{Y}$ is the SSI income limit; y_d is the SSI income disregard; and A_d is the SSI asset limit *and* asset disregard. Note that SSI eligibility is based on income *gross* of taxes. Low-income individuals with assets in excess of A_d can spend down their wealth and qualify for SSI in the future.

If a person is categorically needy and applies for SSI and Medicaid, he receives the SSI transfer, $\underline{Y} - \max\{y_t + ra_t - y_d, 0\}$, regardless of his health; in addition to determining income eligibility, $\underline{Y}$ is the largest possible SSI benefit. A sick person, defined here as one who can not achieve the utility floor with expenditures of $\underline{Y}$, receives additional resources in accordance with equation (2). The combined SSI/Medicaid transfer for a categorically needy person is thus given by

$$b_c\big(a_t, y_t, \mu(\cdot)\big) = \underline{Y} - \max\{y_t + ra_t - y_d, 0\} \; + \; \max\big\{\underline{x}_c(\cdot) - \underline{Y},\; 0\big\}, \tag{12}$$

recalling the restrictions on y_t and a_t in equation (11).

If the person's total income is above $\underline{Y}$ and/or her assets are above A_d, she is not eligible for SSI. If the person applies for Medicaid, transfers are given by

$$b_m\big(a_t, y_t, \mu(\cdot)\big) = \max\big\{\underline{x}_m(\cdot) - \big(\max\{y_t + ra_t - y_d, 0\} + \max\{a_t - A_d, 0\}\big),\; 0\big\}, \tag{13}$$

where we assume that the income disregard y_d and the asset disregard A_d are the same as under the categorically needy pathway.

Each period eligible individuals choose whether to receive Medicaid or not. We will use the indicator function I_M to denote this choice, with $I_M = 1$ if the person applies for Medicaid and $I_M = 0$ if the person does not apply.

When the person dies, any remaining assets are left to his or her heirs. We denote with e the estate net of taxes. Estates are linked to assets by

$$e_t = e(a_t) = a_t - \max\{0, \tau \cdot (a_t - \tilde{x})\}.$$

The parameter τ denotes the tax rate on estates in excess of $\tilde{x}$, the estate exemption level. The utility the household derives from leaving the estate e is

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

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

Using β to denote the discount factor, we can then write the individual's value function as

$$\begin{aligned} V_t(a_t, g, h_t, I, \zeta_t, \xi_t) = \max_{c_t, m_t, a_{t+1}, I_M} \Bigg\{ & u(c_t, m_t, \mu(\cdot)) \\ & + \; \beta s_{g,h,I,t} E_t\Big(V_{t+1}(a_{t+1}, g, h_{t+1}, I, \zeta_{t+1}, \xi_{t+1})\Big) \\ & + \; \beta(1 - s_{g,h,I,t}) \theta \frac{(e(a_{t+1}) + k)}{1-\nu}^{(1-\nu)} \Bigg\}, \end{aligned} \tag{14}$$

subject to the laws of motion for the shocks and the following constraints. If $I_M = 0$, i.e., the person does not apply for SSI and Medicaid,

$$a_{t+1} = a_t + y_n(ra_t + y_t) - c_t - q(h_t)m_t \geq 0, \tag{15}$$

where the function $y_n(\cdot)$ converts pre-tax to post-tax income. If $I_M = 1$, i.e., the person applies for SSI and Medicaid, we have

$$a_{t+1} = b_i(\cdot) + a_t + y_n(ra_t + y_t) - c_t - q(h_t)m_t \geq 0, \tag{16}$$

$$a_{t+1} \leq \min\{A_d, a_t\}, \tag{17}$$

where $b_i(\cdot) = b_c(\cdot)$ if equation (11) holds, and $b_i(\cdot) = b_m(\cdot)$ otherwise. Equations (15) and (16) both prevent the individual from borrowing against future income. Equation (17) forces the individual to spend at least $\underline{x}_i(\cdot)$, and to keep assets below the limit A_d up through the beginning of the next period.

To express the dynamic programming problem as a function of c_t only, we can derive m_t as a function of c_t by using the optimality condition implied by the intratemporal allocation decision. Suppose that at time t the individual decides to spend the total x_t on consumption and out-of-pocket payments for medical goods. The optimal intratemporal allocation then solves:

$$\mathcal{L} = \frac{1}{1-\nu}c_t^{1-\nu} + \mu(\cdot)\,\frac{1}{1-\omega}m_t^{1-\omega} + \lambda_t\left(x_t - m_t q(h_t) - c_t\right),$$

where λ_t is the multiplier on the intratemporal budget constraint. The first-order conditions for this problem reduce to

$$m_t = \left(\frac{\mu(\cdot)}{q(h_t)}\right)^{1/\omega} c_t^{\nu/\omega}. \tag{18}$$

This expression can be used to eliminate m_t from the dynamic programming problem in equation (14), and to simplify the computation of $b_i(\cdot)$.

6 Estimation procedure

We adopt a two-step strategy to estimate the model. In the first step, we estimate or calibrate those parameters that can be cleanly identified outside our model. For example, we estimate mortality rates from raw demographic data. In the second step, we estimate the rest of the model's parameters (ν,ω,β,$\underline{u}_c$,$\underline{u}_m$, and the parameters of $\ln\mu(\cdot)$) with the method of simulated moments (MSM), taking as given the parameters that were estimated in the first step. In particular, we find the parameter values that allow simulated life-cycle decision profiles to "best match" (as measured by a GMM criterion function) the profiles from the data. The moment conditions that comprise our estimator are:

1. To better evaluate the effects of Medicaid insurance, we match the fraction of people on Medicaid by PI quintile, 5 year birth cohort and year cell (with the top two permanent income quintiles merged together).

2. Because the effects of Medicaid depend directly on an individual's asset holdings, we match median asset holdings by PI-cohort-year cell.

3. We match the median and 90th percentile of the out-of-pocket medical expense distribution in each PI-cohort-year cell (the bottom two quintiles are merged). Because the AHEAD's out-of-pocket medical expense data are reported net of any Medicaid payments, we deduct government transfers from the model-generated expenses before making any comparisons.

4. To capture the dynamics of medical expenses, we match the first and second autocorrelations for medical expenses in each PI-cohort-year cell.

The first three sets of moment conditions are those described in section 4.[4]

The mechanics of our MSM approach are as follows. 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, a_t, g, h_t, I) drawn from the data distribution for 1996, and each is assigned the entire health and mortality history realized by the person in the AHEAD data with the same initial conditions. This way we generate attrition in our simulations that mimics precisely the attrition relationships in the data (including the relationship between initial wealth and mortality). The simulated medical needs 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 net worth, medical expenditures, health, and mortality. We then compute asset, medical expense and Medicaid profiles from the artificial histories in the same way as we compute them from the real data. We use these profiles to construct moment conditions, and evaluate the match using our GMM criterion. We search over the parameter space for the values that minimize the criterion. Appendix B contains a detailed description of our moment conditions, the weighting matrix in our GMM criterion function, and the asymptotic distribution of our parameter estimates.

7 First-step estimation results

In this section, we briefly discuss the life-cycle profiles of the stochastic variables used in our dynamic programming model. Using more waves of data, we update the procedure for estimating the income process described in De Nardi et al. [15]. The procedures for estimating demographic transition probabilities and co-pay rates are new.

[4] As was done when constructing the figures in section 4, we drop cells with less than 10 observations from the moment conditions. Simulated agents are endowed with asset levels drawn from the 1996 data distribution, and thus we only match asset data 1998-2010.

7.1 Income profiles

We model non-asset income as a function of age, sex, and the individual's PI ranking. Figure 5 presents average income profiles, conditional on permanent income quintile, computed by simulating our model. In this simulation we do not let people die, and we simulate each person's financial and medical history up through the oldest surviving age allowed in the model. Since we rule out attrition, this picture shows how income evolves over time for the same sample of elderly people. Figure 5 shows that average annual income ranges from about \$5,000 per year in the bottom PI quintile to about \$23,000 in the top quintile; median wealth holdings for the two groups are zero and just under \$200,000, respectively.

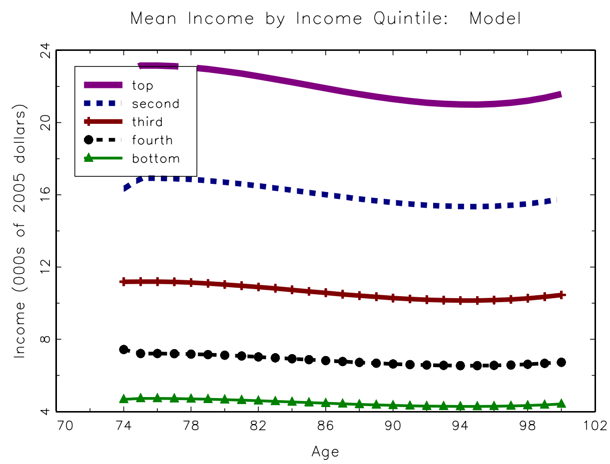

Figure 5: Average income, by permanent income quintile.

7.2 Mortality and health status

We estimate health transitions and mortality rates simultaneously by fitting the transitions observed in the HRS to a multinomial logit model. We allow the transition probabilities to depend on age, sex, current health status, and permanent income. We estimate annual transition rates: combining annual transition probabilities in consecutive years yields two-year transition rates we can fit to the AHEAD data. Appendix C gives details on the procedure.

Using the estimated transition probabilities, we simulate demographic histories, beginning at age 70, for different gender-PI-health combinations. Table 1 shows life expectancies. We find that rich people, women, and healthy people live much longer than their poor, male, and sick counterparts. For example, a male at the 10th permanent income percentile in a nursing home expects to live only 2.2 more years, while a female at the 90th percentile in good health expects to live 16.0 more years.

Income Percentile	Males Nursing Home	Males Bad Health	Males Good Health	Females Nursing Home	Females Bad Health	Females Good Health	All†
10	2.16	5.91	7.23	3.04	9.73	11.76	10.21
30	2.15	6.47	8.17	3.14	10.71	12.92	11.22
50	2.15	7.11	9.10	3.25	11.73	14.06	12.30
70	2.17	7.82	10.09	3.40	12.72	15.08	13.34
90	2.18	8.57	10.98	3.60	13.70	15.98	14.19
By gender:‡							
Men							9.42
Women							13.32
By health status:◇							
Bad Health							10.47
Good Health							13.74

Notes: Life expectancies calculated through simulations using estimated health transition and survivor functions. † Using gender and health distributions for entire population; ‡ Using health and permanent income distributions for each gender; ◇ Using gender and permanent income distributions for each health status group.

Table 1: Life expectancy in years, conditional on reaching age 70.

Another important saving determinant is the risk of requiring nursing home care. Table 2 shows the probability at age 70 of ever entering a nursing home. The calculations show that 40.9% of women will ultimately enter a nursing home, as opposed to 27.3% for men. These numbers are similar to those from the Robinson model described in Brown and Finkelstein [6], which show 27% of 65-year-old men and 44% of 65-year-old women require nursing home care. One possible reason we find a lower number for women is that the Robinson model is based on older data, and nursing home utilization has declined in recent years (Alecxih [1]).

7.3 Co-pay rates

The co-pay rate $q_t = q(h_t)$ is the share of total billable medical spending not paid by Medicare or private insurers. Thus, it is the share paid out-of-pocket or by Medicaid. We allow it to differ depending on whether the person is in a nursing home or not: $q_t = q(h_t)$.

Income Percentile	Males: Bad Health	Males: Good Health	Females: Bad Health	Females: Good Health	All†
10	26.7	28.1	37.3	39.9	36.7
30	26.3	28.1	38.3	41.0	37.8
50	26.1	27.9	38.9	42.5	38.5
70	25.4	28.1	39.5	43.1	39.3
90	24.8	28.2	39.9	43.8	39.5
By gender:‡					
Men					27.3
Women					40.9
By health status:◇					
Bad Health					36.0
Good Health					39.6

Notes: Entry probabilities calculated through simulations using estimated health transition and survivor functions; † Using gender and health distributions for entire population; ‡ Using health and permanent income distributions for each gender; ◇ Using gender and permanent income distributions for each health status group.

Table 2: Probability of ever entering a nursing home, people alive at age 70.

Using data from the MCBS, we estimate the co-pay rate by taking the ratio of mean out-of-pocket spending plus Medicaid payments to mean total medical expenses. The co-pay rate for people not in a nursing home averages 29% and does not vary much with demographics. The co-pay rate for those in nursing homes is 92%. For every dollar spent on nursing homes, 47 cents come from Medicaid and 45 cents are from out-of-pocket, with only 8 cents coming from Medicare or other sources. In our model, we round this number to 90%. We cross-checked these co-pay rates with data from the 1996-2008 waves of the Medical Expenditure Panel Survey (MEPS), again making the same sample selection decisions as in the AHEAD. For those not in a nursing home, the MCBS and MEPS estimated co-pay rates were very similar. However, MEPS does not contain information on individuals in nursing homes, so we rely on the estimated co-pay rates from MCBS.

8 Second step results, model fit, and identification

8.1 Parameter values

β: discount factor	.995
	(0.016)
ν: RRA, consumption	3.039
	(0.025)
ω: RRA, medical expenditures	3.367
	(0.034)
$\underline{Y}$: SSI income level,	\$6,420
	(235.9)
$\underline{u}_c = \underline{u}_m$: utility floor,†	\$5,261
	(194.4)
θ: bequest intensity	138.1
	(7.55)
k: bequest curvature (in 000s)	18.5
	(0.718)

† The estimated utility floor is indexed by the consumption level that provides the floor when $\mu = 0$.

Table 3: Estimated preference parameters. Standard errors are in parentheses below estimated parameters. NA refers to parameters fixed for a given estimation.

Table 3 presents our estimated parameters. Our estimate of β, the discount factor, is 0.995. This number has to be multiplied by the survival probability to obtain the effective discount factor. As Table 1 shows, the survival probability for our sample of older individuals is low, implying an effective discount factor much lower than β.

The estimate of ν, the coefficient of relative risk aversion for "regular" consumption, is 3.0, while the estimate of ω, the coefficient of relative risk aversion for medical goods, is 3.4; the demand for medical goods is less elastic than the demand for consumption. In a recent study, Fonseca et al. [22] calculate that the co-insurance elasticity for total medical expenditures ranges from -0.27 to -0.35, which they find to be consistent with existing micro evidence. Repeating their experiment (a 150% increase in co-pay rates) with our model reveals that elasticities range by age and income: richer and younger people have higher elasticities. To calculate a summary number, we use our model of mortality and an annual population growth rate of

1.5% to find a cross-sectional distribution of ages. Combining this number with our simulations, we find an aggregate cross-sectional elasticity of -0.27.

The SSI income generosity (which is also the income threshold to be medically needy) is estimated at \$6,420, a number very close to the \$6,950 statutory threshold used in many states.

In our baseline estimates, we constrain the two utility floors to be the same, as Medicaid generosity does not appear to be drastically different across the two categories of recipients. The utility floor corresponds to the utility from consuming \$5,261 a year when healthy. It should be noted that the medically needy are guaranteed a minimum income of \$6,420 (\$7,020 including the income disregard) so that their total consumption when healthy is at least \$7,020 a year. However, when there are large medical needs, transfers are determined by the Medicaid-induced utility floor.

The point estimates of θ and k imply that, in the period before certain death, the bequest motive becomes operative once consumption exceeds \$3,614 per year. (See De Nardi, French, and Jones [15] for a derivation.) For individuals in this group, the marginal propensity to bequeath, above the threshold level, is 83 cents out of every additional dollar. Several other authors have recently estimated bequest motives inside structural models of old age saving.[5] Imposing a linear bequest motive, Kopczuk and Lupton [35] find that agents with bequest motives (around three quarters of the population) would, when facing certain death, bequeath all wealth in excess of \$29,700. De Nardi, French and Jones [15] find that, depending on the specification, the bequest motive becomes active between \$31,500 and \$43,4000, and generates a marginal propensity to bequeath of 88-89%. Lockwood [38] finds a threshold of \$18,400 and a propensity to bequeath of 92%. While these studies suggest bequests are more of a luxury good than do our estimates, none of them seek to explain Medicaid usage. In contrast, Ameriks et al. [3] estimate their model using survey data questions, including hypothetical questions about bequests and long-term care insurance, in a model aimed at assessing Medicaid and medical expense risk. They find a terminal bequest threshold of \$7,100 and a propensity to bequeath of 98%. Compared to them, we find a lower threshold, but a much higher marginal propensity to consume.

We now turn to discussing how well the model fits the some key aspects of the data, the identification of the model's parameters, and to highlighting some interesting model implications.

8.2 Model fit

Figure 6 compares the Medicaid recipiency profiles generated by the model (dashed line) to those in the data (solid line) for the members of four birth-year cohorts. In panel a, the 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 second set of lines

[5] Assembling these figures requires a few derivations and inflation adjustments. Calculations are available on request.

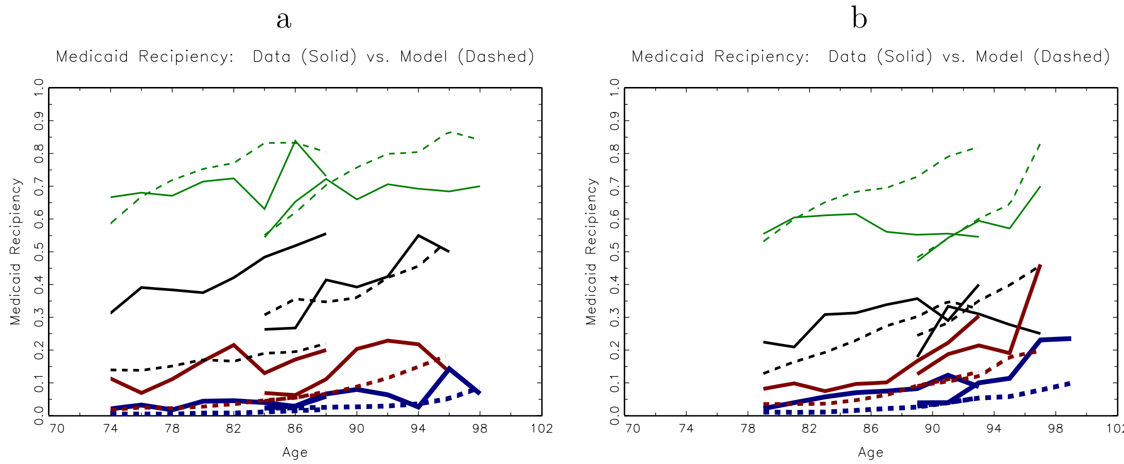

Figure 6: Medicaid recipiency by cohort and PI quintile: data (solid lines) and model (dashed lines).

are for the cohort aged 82-86 in 1996. Panel b displays the two other cohorts, starting respectively at age 79 and 89. The graphs show that the model matches the general patterns of Medicaid usage. The model tends to over-predict usage by the poor, and undepredict usage by the rich, especially at relatively younger ages.

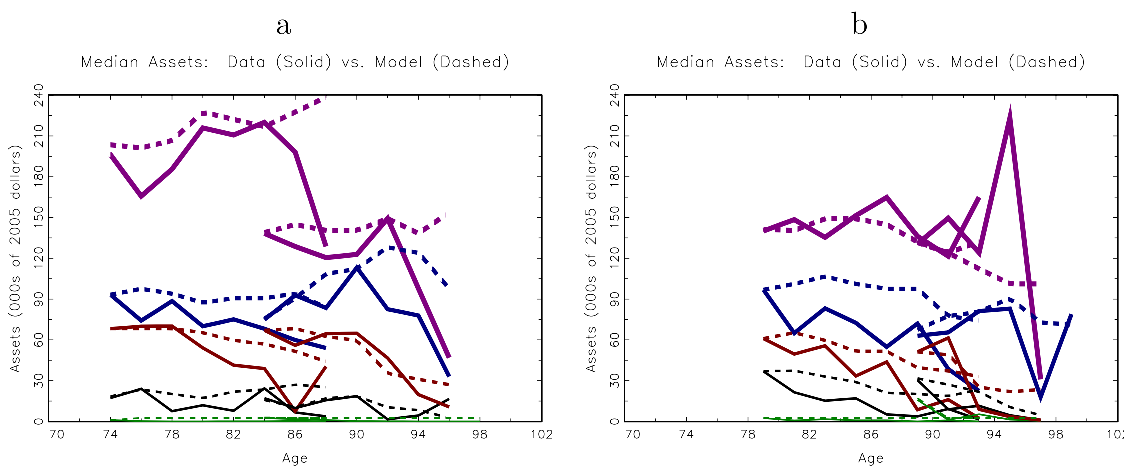

Figure 7: Median net worth by cohort and PI quintile: data (solid lines) and model (dashed lines).

Figure 7 plots median net worth by age, cohort, and permanent income. Here too the model does well, matching the observation that the savings patterns differ by permanent income and that higher PI people don't run down their assets until well past age 90. If anything, the model tends to over-predict saving in the top PI group.

Figure 8 displays the median and ninetieth percentile of out-of-pocket medical expenses (that is, net of Medicaid payments and private and public insurance co-pays) paid by people in the model and in the data. Permanent income has a large effect on out-of-pocket medical expenses, especially at older ages. Median medical expenses are less than $1,500 a year at age 75. By age 100, they stay flat for those in the bottom quintile of the income distribution but often exceed $5,000 for those at the top of the income distribution. Panels a and b show that the model does a reasonable job of matching the medians found in the data. The other two panels report the 90th percentile of out-of-pocket medical expenses in the model and in the data and thus provides a better idea of the tail risk by age and permanent income. Here the model reproduces medical expenses of $4,000 or less at age 74, staying flat over time for the lower PI people, but tends to understate the medical expenditures of high-PI people in their late nineties.

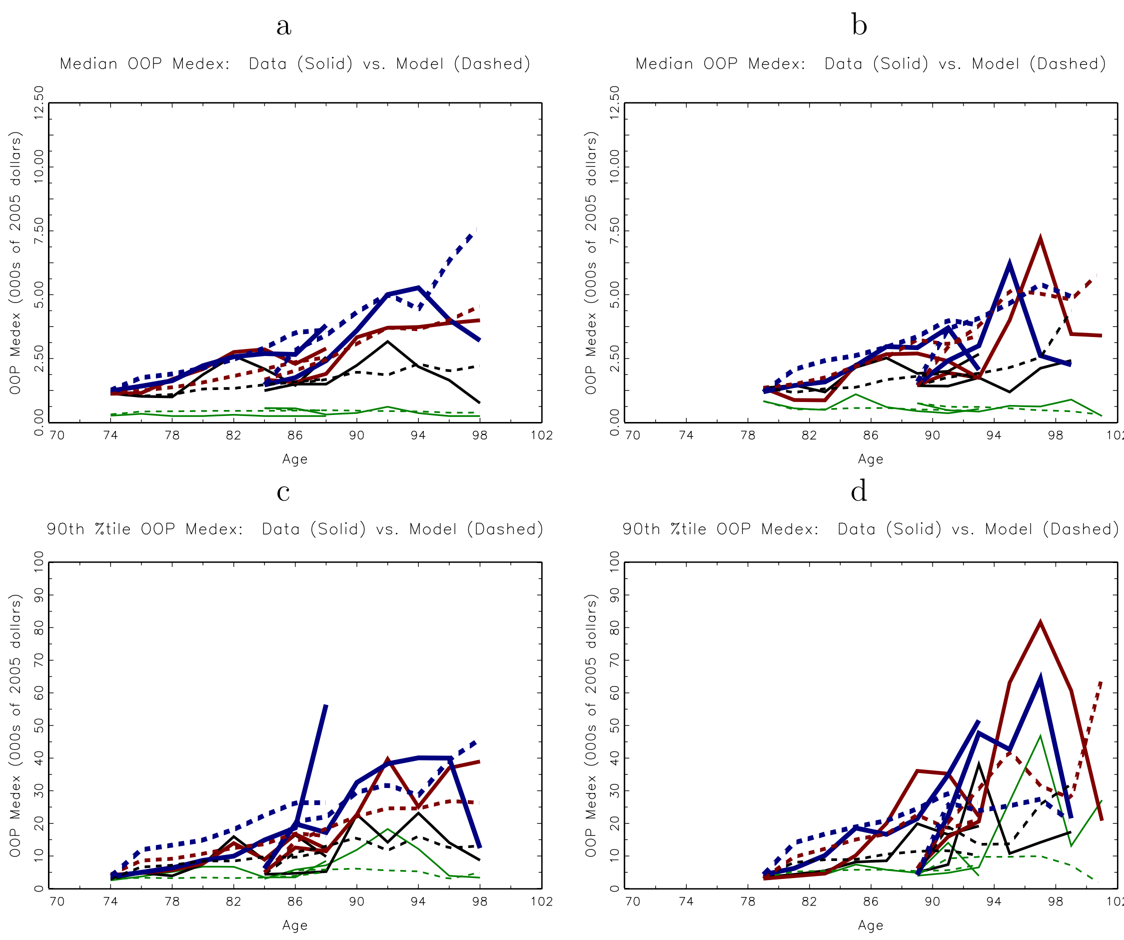

Figure 8: Median and ninetieth percentile of out-of-pocket medical expenses by cohort and PI quintile: data (solid lines) and model (dashed lines).

Turning to cross-sectional distributions of medical spending, Figure 9 presents three panels. Panel a, in the top left corner, presents the cumulative distribution function (CDF) of out-of-pocket medical expenditures found in the AHEAD and

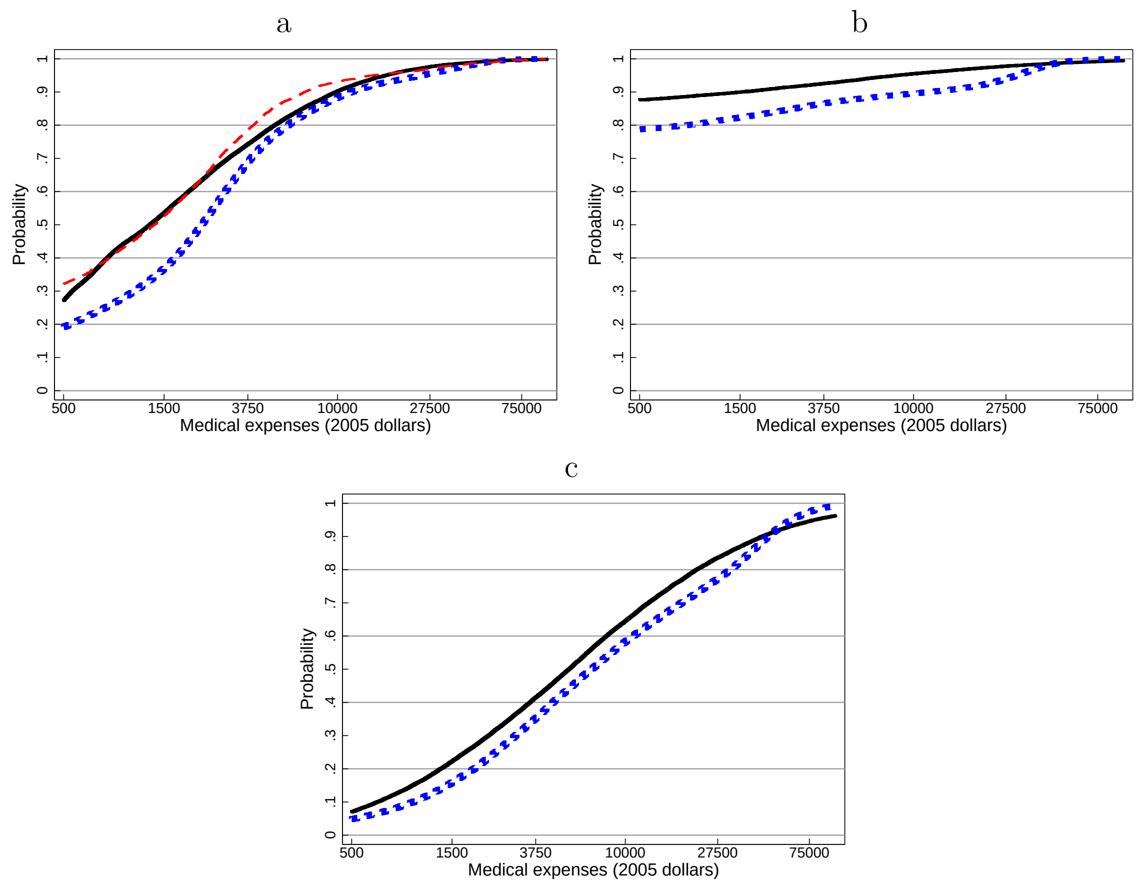

Figure 9: Cumulative distribution functions of medical spending: model (solid line), AHEAD data (dashed line) and and MCBS data (dotted line). Panel a: out-of-pocket expenditures. Panel b: Medicaid expenditures. Panel c: total expenditures.

MCBS data, as well as that produced by the model. The solid line is the model-predicted CDF, the dashed line is the AHEAD CDF, and the dotted line is the MCBS CDF. Because the model's parameters are estimated in part by fitting AHEAD out-of-pocket spending profiles—although not the CDF itself—it is not surprising that AHEAD and model-predicted CDFs are very similar. The model CDF also resembles the MCBS CDF, although out-of-pocket medical spending in the MCBS is higher up to the 98th percentile of the spending distribution.

Panel b shows the CDF of Medicaid payments, both as predicted by the model and in the MCBS data. Medicaid expenditures in the MCBS data are higher than those predicted by the model up to the 98th percentile, but are lower thereafter. Panel c, at the bottom, shows the CDF of total medical expenditures from all payors. Total expenditures in the MCBS are higher than the model predictions up to the 92nd percentile at \$48,000, and are lower thereafter. In summary, these differences are not large and the model fits well the distribution of-out-of pocket, Medicaid, and total medical spending. Because Medicaid and total medical expenditures are *not* part of the GMM criterion we use to estimate the model, the ability of the model to fit these data provides additional validation. This feature is important for policy analysis, as it means the model is able to match the possibility of catastrophic medical spending.

	Medicaid payments		Out-of-pocket expenses		
Income Quintile	MCBS Data	Model	MCBS Data	AHEAD Data	Model
Bottom	6,170	4,860	3,850	2,380	1,800
Fourth	4,220	3,610	4,740	3,810	3,510
Third	2,030	2,450	5,190	4,480	5,450
Second	1,080	1,460	5,880	5,520	6,470
Top	900	810	6,820	6,390	7,620
Men	2,040	1,820	4,540	4,090	4,780
Women	3,080	2,760	5,480	4,330	6,230
Good Health	220	1,200	2,360	2,550	4,260
Bad Health	620	2,480	2,660	3,260	4,400
Nursing Home	13,620	7,630	17,290	13,170	8,620

Table 4: Medicaid payments and out-of-pocket medical expenditures (2005 dollars), model, MCBS data and AHEAD data.

Table 4 shows average Medicaid and out-of-pocket expenditures conditional on permanent income quintile, both as predicted by the model and as in the data. The first two columns of Table 4 compare Medicaid expenditures in the MCBS data to those predicted by the model. It shows that retirees at the bottom of the income distribution have average Medicaid expenditures of \$6,170 and \$5,050 in the data and model, respectively. For those at the top of the income distribution, Medicaid expenditures are \$900 and \$810 in data and model, respectively. Overall the model matches Medicaid payments well. It bears noting that although average Medicaid payments are smaller at the top of the income distribution, conditional on receiving Medicaid those at the top of the income distribution receive much larger payments than lower income groups. This is true both in the model and in the data.

The last three columns of Table 4 compare out-of-pocket expenditures from the MCBS, the AHEAD and the model. The MCBS data shows a less steep income gradient than the AHEAD data. Those at the bottom of the income distribution spend \$3,850 in the MCBS data and \$2,360 in the AHEAD data, while expenditures at the top are \$6,820 in the MCBS versus \$6,390 in the AHEAD. Overall, however, the gradients are similar. This similarity in average out-of-pocket expenditures gives us confidence that our facts are robust across datasets. The final column shows the average out-of-pocket expenditures predicted by the model. Overall the model fits the data well for both out-of-pocket and Medicaid expenditures. Details on the construction of these cross-sectional comparisons, and additional comparisons, can be found in Appendix A.

8.3 Identification

The preference parameters are identified jointly. There are multiple ways to generate high saving by the elderly: large values of the discount rate β, low values of the utility floors $\underline{u}_c$ and $\underline{u}_m$, large values of the curvature parameters ν and ω, or strong and pervasive bequest motives (high values of θ and small values of k). Dynan, Skinner and Zeldes [18] point out that the same assets can simultaneously address both precautionary and bequest motives. There are also multiple ways to ensure that the income-poorest elderly do not save, including high utility floors and bequest motives that become operative only at high levels of consumption.

We acquire additional identification in several ways. First, and importantly, we require our model to match Medicaid recipiency rates, which helps pin down the utility floors and the SSI threshold $\underline{Y}$. To be able to match the fraction of people on Medicaid by PI, cohort, and age, the Medicaid insurance floors have to be substantial, in excess of \$5,000 of consumption by the healthy. A lower floor would generate too few people on Medicaid, especially at higher permanent income quintiles. For example, the model with endogenous medical expenses in De Nardi, French and Jones [15], the one most comparable with the model in this paper, was not estimated to match Medicaid fractions. That model was able to fit the asset data using a similar value of β, no bequest motives, and lower utility floors. A similar combination of parameters matches the asset data very well even with our current, richer specification of the Medicaid program; the combination in fact matches the asset data better than our baseline estimates. However, the Medicaid program implied by those estimates is too stingy to generate the Medicaid fractions observed in the data. Requiring the model to match Medicaid recipiency thus introduces a tension in the estimation process: Medicaid needs to be fairly generous to generate both a high fraction of people on Medicaid and the pattern of Medicaid recipiency across age and income. However, a more generous Medicaid program reduces the need to accumulate assets. To match the same asset profiles under a more generous insurance system we need a higher discount rate and/or a stronger bequest motive.

Requiring the model to match observed out-of-pocket medical expenses, in addition to the other moments discussed above, helps identify the discount factor and the bequest motive. While the two have similar implications for asset holdings, they have different implications for the pattern of non-medical consumption and medical expenditures over time. A person saving because of high patience will tend to consume more at relatively later ages, and hence will on average, in an environment in which medical needs increase with age, incur more out-of-pocket medical costs in old age. If, instead, people save more due to bequest motives, consumption of medical goods and services does not need to rise as much in old age. Our combination of the Medicaid floor, discount factor, and bequest motives, yields the best fit of these three sets of moment conditions.

We also estimate the coefficients for the mean of the logged medical needs shifter $\mu(h_t, \psi_t, t)$, the volatility scaler $\sigma(h_t, t)$ and the process for the shocks ζ_t and ξ_t. As we show in the graphs that follow, the estimates for these parameters (available from the authors on request) imply that the demand for medical services rises rapidly with age. Matching the median and 90th percentile of out-of-pocket medical expenditures, along with their first and second autocorrelations, is the principal way in which we identify these parameters. The fact that the medical needs shocks do not depend directly on income—the only link is through the health transition probabilities—also helps us identify other parameters, as the expenditure profiles we match are disaggregated by income. Most notably, the income gradient of medical expenditures helps us pin down the curvature parameters ν and ω.

8.4 Other interesting implications of the model

Figure 10 presents profiles that arise when the youngest cohort is simulated from ages 74 to (potentially) 100. Each simulated individual receives a value of the state vector (t, a_t, g, h_t, I) drawn from the data distribution of 72- to 76-year-olds in 1996. He or she then receives a series of health, medical expense, and mortality shocks consistent with the stochastic processes described in the model section, and is tracked to age 100.

Panel a of Figure 10 shows average out-of-pocket medical expenses. Out-of-pocket expenditures rise rapidly with age for people in the top income groups, but remain low throughout retirement for those at the bottom. Panel b shows the sum of medical expenses paid out-of-pocket and the expenses paid by Medicaid. Because the costs picked up by Medicaid are co-pays (part of $q_t m_t$), the sums in panel b are in many ways a better measure of an individual's co-payment expenses than the "pure" out-of-pocket expenditures in panel a. These sums also increase rapidly with age, going from around \$4,000 at age 74 to \$20,000 at age 100. Medicaid allows poorer people to consume far more medical goods and services than they pay for. As a result, the expense sums shown in panel b rise much more slowly with income than do the out-of-pocket expenses shown in panel a.

Panel c displays non-medical consumption, including that funded by government transfers. The consumption profiles differ from the medical expense profiles shown in panel b in two important ways. First, while medical expenditures rise over retirement by a factor of five, non-medical consumption expenditures decline, albeit slightly, over

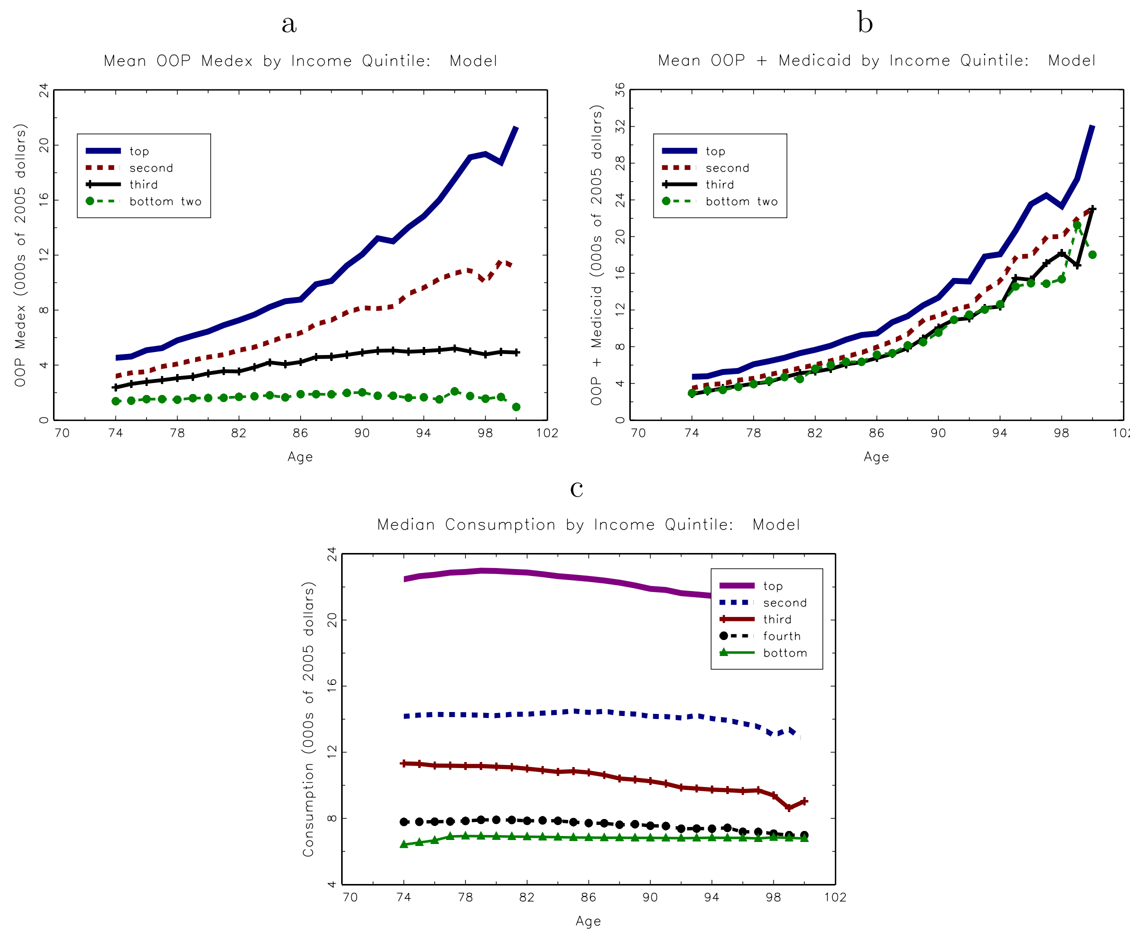

Figure 10: Average expenditures by age and permanent income. Panel a: medical expenses paid out-of-pocket. Panel b: medical expenses paid out-of-pocket or by Medicaid. Panel c: consumption.

the same horizon. Second, non-medical consumption rises much more rapidly with income than do medical expenses. This is consistent with our parameter estimates, which imply that the demand for medical goods is less elastic than the demand for consumption.

Figure 11 describes the Medicaid transfers generated by the model, and illuminate the interaction of the utility floor and medical needs shocks. Panel a of this figure shows the fraction of individuals receiving transfers, while panel b shows average transfers, taken across both recipients and non-recipients. Panel a shows that people in the bottom two permanent income quintiles receive Medicaid at fairly high rates throughout their retirement. Most of these people qualify through the categorically needy pathway. People in the top income quintiles, in contrast, use Medicaid much more heavily at older ages, when large medical expenditures make them eligible through the medically needy pathway.

Panel b of Figure 11 shows average Medicaid transfers. While low-income people are much more likely to qualify for Medicaid, the categorically needy provision allows them to qualify with small medical needs. The medically needy provision allows high-income people to qualify only when their medical expenses are high relative to their resources. Although the poor on average receive relatively more Medicaid benefits than the rich at younger ages, at very old ages both groups receive significant benefits.

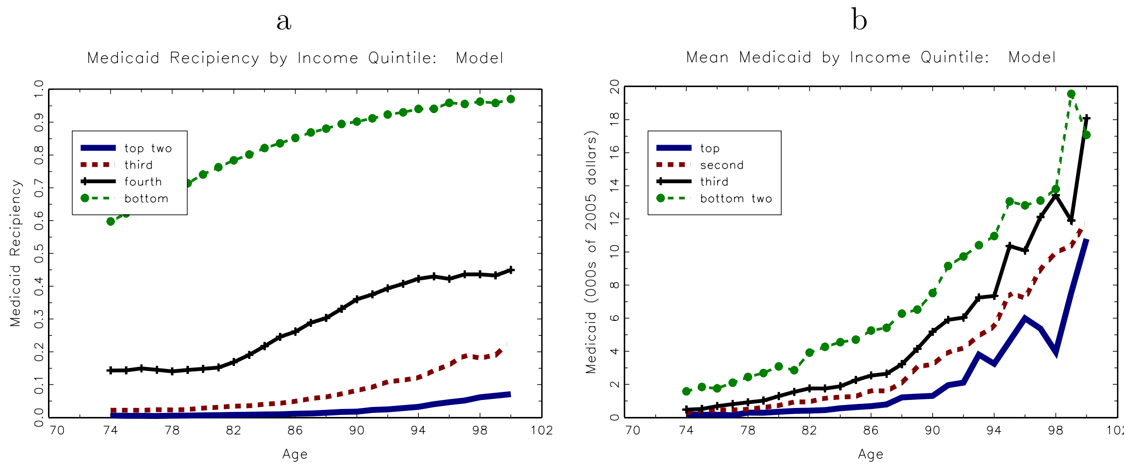

Figure 11: Medicaid. Panel a: fraction receiving Medicaid. Panel b: average Medicaid transfers.

9 Medicaid Insurance, other Insurance, and Lifetime Consumption

Once we have simulated life histories, it is straightforward to convert the simulated expenditure streams into present discounted values, using the model's assumed pre-tax discount rate of 4%. We also calculate the annuity values associated with the discounted sums. Each group's annuity value equals the group's average present discounted sum divided by its average lifespan (adjusted for discounting). Table 5 contains two pairs of columns, which contain results, respectively, for Medicaid payments and out-of-pocket medical expenditures. Within each pair, the left-hand-side column reports average present discounted values of payments for that income, gender or health group at age 74, while the right-had-side column shows annuity values.

The leftmost column of Table 5 shows present discounted values of Medicaid transfers. Although these transfers decrease by income quintile, they are non-trivial for all income groups. The present discounted value of Medicaid payments received by people in the highest income quintile is $4,300, which is about one fifth of their yearly income. Although the poor are more likely to be receiving Medicaid, they tend to

Income Quintile	Medicaid payments: Present Discounted Value	Medicaid payments: Annuity Value	Out-of-pocket medical costs: Present Discounted Value	Out-of-pocket medical costs: Annuity Value
Bottom	25,200	3,540	7,500	1,050
Fourth	19,200	2,470	16,800	2,160
Third	12,600	1,500	27,700	3,300
Second	8,800	1,010	41,300	4,720
Top	4,300	490	59,600	6,770
Men	6,400	970	31,400	4,750
Women	14,200	1,650	35,000	4,040
Good Health	11,000	1,210	41,300	4,580
Bad Health	15,900	2,170	23,800	3,250

Table 5: Medicaid payments and out-of-pocket medical costs at age 74.

die before they develop the most costly health conditions. On the other hand, the richest, while having the most medical expenses, have the most resources to pay for medical care themselves. The interaction of these two mechanisms leaves people in the middle income quintiles, who have more expensive medical conditions, but still modest financial resources, also receiving significant benefits.

These flows reinforce the view that middle- and higher-income people also benefit from Medicaid transfers in old age. Women benefit more than men from Medicaid, both because they live longer and because they tend to be poorer. Finally, those in good health at age 74 receive almost as many benefits as those in bad health at 74, because they tend to live long enough to require costly procedures and long nursing home stays.

The second column in Table 5 reports the annuity value of the same Medicaid payments. Even according to this measure, Medicaid transfers are non-negligible at all permanent income levels, especially when compared to the out-of-pocket expenditures for the medical goods and services that are being consumed. The final two columns of this table shows that out-of-pocket medical expenses rise quickly with income. Over their lifetime, the out-of-pocket costs of medical goods and services for income richest are almost 8 times as much as the income-poorest. Moving to the annuity values in the final column lowers the ratio to 6, as it removes the effects of the rich's longer life expectancy.

Table 6 reports values for consumption (left-hand-side columns) and consumption of medical goods and services (right-hand-side columns). Comparing the consumption

Income Quintile	Consumption: Present Discounted Value	Consumption: Annuity Value	Medical goods and services: Present Discounted Value	Medical goods and services: Annuity Value
Bottom	57,300	8,060	90,800	12,760
Fourth	78,100	10,040	103,200	13,260
Third	104,000	12,410	119,000	14,200
Second	159,600	18,270	148,000	16,940
Top	235,600	26,740	189,800	21,540
Men	134,100	20,280	113,800	17,200
Women	142,100	16,430	143,300	16,570
Good Health	171,200	18,970	154,000	17,070
Bad Health	95,300	13,020	116,500	15,930

Table 6: Consumption and consumption of medical goods and services at age 74.

of medical goods and services with their out-of-pocket costs shows that gap in total medical consumption between the income-rich and income-poor is much smaller the gap in out-of-pocket expenditures. The table also shows that while men and women consume more or less the same amount of medical goods and services per period of life, the discounted present value is much larger for women, as they tend to live almost 4 years longer. It also shows that non-medical consumption rises much more quickly in income than total medical spending, with the ratio of this variable for the poorest to the richest being less than 30%. Finally, the table highlights how large the consumption of medical goods and services is, relative to non-medical consumption, for people in all permanent income levels.

9.1 Changing Medicaid

To help inform policy reform, we conduct a few experiments in which we change the Medicaid program rules and compute the wealth transfers that would make each retiree as well off as she had been in the the base program behind our estimates. If Medicaid provides retirees with valuable insurance, the value that retirees place on Medicaid may greatly exceed the actuarial value of expected benefits. On the other hand, people may value the benefit flows at less than their face value, if they would prefer having the cash today, to dispose of as they wish, over receiving Medicaid benefits in the future. In the tables below, we present compensating differentials for different groups, taking averages across each group's distribution of the relevant state variables, and comparing them to the changes in Medicaid transfers associated with the same experiments.

To explore the differences in insurance provided by the categorically and the medically needy program, we start by decreasing the categorically needy utility floor by 25%, that is the consumption of the categorically needy when healthy drops from \$5,300 to \$3,900. The two leftmost columns of Table 7 show that this change only affects people in the bottom two permanent income quintiles, as those with higher incomes never qualify as categorically needy. The discounted present value of Medicaid payments drops by \$9,600 and \$2,300, respectively, for people in the two bottom PI quintiles. The second column of the pair reports how much the people in each cell need to be compensated in terms of a lump-sum wealth transfer at age 74. Comparing the two columns reveals that the categorically needy people value their lost Medicaid insurance at more than the cost required to provide it, but that the difference is not very big.

	Categorical floor down 25%		Both Medicaid floors down 25%	
Income Quintile	Increased Present Discounted Value	Individual Valuation	Increased Present Discounted	Individual Valuation
Bottom	7,200	9,600	8,300	10,900
Fourth	1,200	2,300	6,500	6,600
Third	0.0	0.0	4,400	6,100
Second	0.0	0.0	3,100	7,200
Top	0.0	0.0	1,600	9,000
Men	400	500	2,200	4,500
Women	1,700	2,500	4,800	8,900
Good Health	1,000	1,600	3,800	8,500
Bad Health	2,300	3,100	5,400	7,700

Table 7: Decrease in Medicaid Payments at age 74, categorically needy floors cut by 25%, and both Medicaid floors cut by 25%, compared to individuals' valuations of the cuts.

We next cut the consumption value of both utility floors (that is, both the categorically and medically needy floors) by 25%, and simulate our model again. The results in the right-hand-side columns of Table 7 show the resulting reductions in Medicaid payments and the cash compensation that people would require to be indifferent to such cuts. A striking feature of this table is that while people in the lowest three income quintiles value value Medicaid close to its cost, people in the top two permanent income quintile value Medicaid insurance two to over five times its cost. The insurance value is very high for these people because of two reasons. First, these people are high income, so they have a high lifetime value of consumption, and thus have more consumption to lose when they fall. Second, they face the double compounded risk of living a long time, well past one's life expectancy, and facing extremely high medical needs. It is in those states of the world that insurance is most valuable.

Figure 12 compares the saving profiles and the Medicaid recipiency rates for this experiment (solid line) to those from our estimated baseline model (dashed line). Figure 12 shows that while median assets increase only modestly for all permanent income quintiles, the Medicaid recipiency rate declines more significantly, especially at lower permanent income levels. Making Medicaid less generous causes the Medicaid recipiency rate to drop, and leads people to increase their precautionary savings.

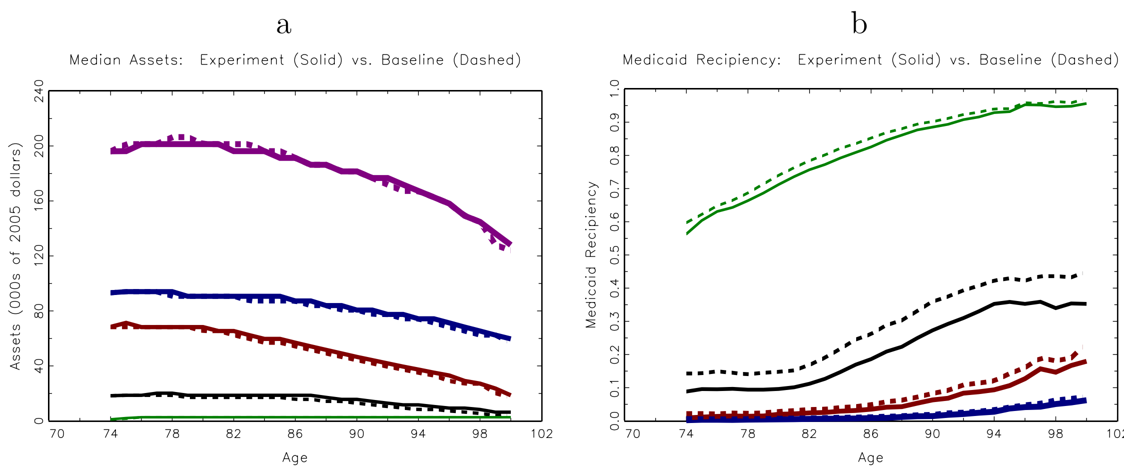

Figure 12: Assets (panel a) and Medicaid recipiency rates (panel b) by age and permanent income. Dashed line: benchmark, solid line: less generous floor

In Table 8, we analyze the benefits of making the Medicaid program more generous, by increasing the Medicaid consumption floor for the healthy by 10% (from \$5,300 to \$5,800). Table 8 shows that people at most permanent income levels value Medicaid flows less than their cost, with higher permanent income people valuing them by a bit more. Put together, the results in Tables 7 and 8 indicate that at current programs rules people value Medicaid transfers at more than their actuarial cost, but that further increasing the size of the program would not raise its insurance value as much as its cost.

Income Quintile	Both Medicaid floors up 10% Increased Present Discounted Value	Valuation
Bottom	3,900	1,700
Fourth	3,000	1,400
Third	1,800	800
Second	1,300	800
Top	700	1,300
Men	1,000	200
Women	2,200	1,400
Good Health	1,600	1,200
Bad Health	2,400	1,200

Table 8: Increased Medicaid Payments at age 74, both Medicaid floors raised by 10%, increased Medicaid payments compared to individuals' valuations of the increases.

10 Conclusion

We find that even higher income retirees end up on Medicaid if they live a long life and face large medical expenses. Although the lifetime discounted present value of Medicaid payments does decrease with permanent income, even higher income people can receive sizeable Medicaid payments, as they tend to live longer and face higher medical needs in very old age. Furthermore, our compensating differential calculations show that many higher income retirees value Medicaid insurance as much or more than lower-income ones.

Our findings of the importance of Medicaid spreading across people of all income levels are consistent with the views of the American families. The Henry Kaiser Family Foundation, an independent research institution, conducts surveys about health and insurance for the U.S. families. The July 2012 Kaiser Health tracking poll includes a representative cross-section of the U.S. adult population. Kaiser asked "How important for you and your family is the Medicaid program" Over 52% of respondents stated that Medicaid is either very important (35%) or somewhat important (17%). Also of interest is the breakdown of the answers to the question by family income. Of those making $40,000 or less a year, 69% answered that Medicaid was either very important (51%) or somewhat important (18%). Of those in the $40,000 to $90,000 income bracket, 43% replied that medicaid is either very important (22%) or somewhat important (21%). Of those making more than $90,000 a year, 36% replied that Medicaid is either very important (24%) or somewhat important (12%). These results confirm our finding that even families with higher income can end up on Medicaid, and thus value the insurance benefits that it provides.

Finally, our compensating differential calculations indicate that retirees value Medicaid insurance at more than its actuarial cost, but that most would value expansions of the current Medicaid program at less than cost.

References

[1] Lisa Alecxih. Nursing home use by "oldest old" sharply declines. Mimeo, 2006.

[2] Joseph G. Altonji and Lewis M. Segal. Small sample bias in gmm estimation of covariance structures. *Journal of Business and Economic Statistics*, 14(3):353–366, 1996.

[3] John Ameriks, Andrew Caplin, Steven Laufer, and Stijn Van Nieuwerburgh. The joy of giving or assisted living? using strategic surveys to separate bequest and precautionary motives. 2011.

[4] R. Anton Braun, Karen Kopecky, and Tatyana Koreshkova. Old, sick, alone, and poor: A welfare analysis of old-age insurance programs. Mimeo, 2013.

[5] Robert Brook, John E. Ware, William H. Rogers, Emmett B. Keeler, Allyson R. Davies, Cathy A. Donald, George A. Goldberg, Kathleen N. Lohr, Patricia C. Masthay, and Joseph P. Newhouse. Does free care improve adults' health? results from a randomized trial. *New England Journal of Medicine*, 309(23):1426–1434, 1983.

[6] Jeff Brown and Amy Finkelstein. The interaction of public and private insurance: Medicaid and the long term care insurance market. NBER Working Paper 10989, 2004.

[7] Jeff Brown and Amy Finkelstein. The interaction of public and private insurance: Medicaid and the long term care insurance market. *American Economic Review*, 98(5):837–880, 2008.

[8] Jeffrey R. Brown, Norma B. Coe, and Amy Finkelstein. Medicaid crowd-out of private long-term care insurance demand: Evidence from the health and retirement survey. In James M. Poterba, editor, *Tax Policy and the Economy*. MIT Press, 2007.

[9] Moshe Buchinsky. Recent advances in quantile regression models: A practical guideline for empirical research. *Journal of Human Resources*, 33:88–126, 1998.

[10] David Card, Carlos Dobkin, and Nicole Maestas. Does medicare save lives? *The Quarterly Journal of Economics*, 124(2):597–636, 2009.

[11] Gary Chamberlain. Comment: Sequential moment restrictions in panel data. *Journal of Business & Economic Statistics,*, 10(1):20–26, 1992.

[12] Kenneth Y. Chay, Daeho Kim, and Shailender Swaminathan. Medicare, hospital utilization and mortality: Evidence from the programs origins. Mimeo, 2010.

[13] Joseph Doyle. Returns to local-area healthcare spending: Using health shocks to patients far from home. *American Economic Journal: Applied Economics*, 3(3):221–243, 2011.

[14] Mariacristina De Nardi, Eric French, and John B. Jones. Life expectancy and old age savings. *American Economic Review: Papers and Proceedings*, 99(2):110–115, 2009.

[15] Mariacristina De Nardi, Eric French, and John B. Jones. Why do the elderly save? the role of medical expenses. *Journal of Political Economy*, 118(1):39–75, 2010.

[16] Mariacristina De Nardi, Eric French, and John B. Jones. Medicaid and the elderly. NBER Working Paper 17689, 2011.

[17] Darrell Duffie and Kenneth J. Singleton. Simulated moments estimation of markov models of asset prices. *Econometrica*, 61(4):929–952, 1993.

[18] Karen E. Dynan, Jonathan Skinner, and Stephen P. Zeldes. The importance of bequests and life-cycle saving in capital accumulation: A new answer. *American Economic Review*, 92(2):274–278, 2002.

[19] Amy Finkelstein and Robin McKnight. What did medicare do? the initial impact of medicare on mortality and out of pocket medical spending. *Journal of Public Economics*, 92(7):1644–1669, 2008.

[20] Amy Finkelstein, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Heidi Allen Joseph P. Newhouse, Katherine Baicker, and Oregon Health Study Group. The oregon health insurance experiment: Evidence from the first year. *Quarterly Journal of Economics*, 127(3):1057–1106, 2012.

[21] Elliott S. Fisher, David E. Wennberg, Therese A. Stukel, Daniel J. Gottlieb, F.L. Lucas, and Etoile L. Pinder. The implications of regional variations in medicare spending. part 2: Health outcomes and satisfaction with care. *Annals of Internal Medicine*, 138(4):288–322, 2003.

[22] Raquel Fonseca, Pierre-Carl Michaud, Titus Galama, and Arie Kapteyn. On the rise of health spending and longevity. Rand Working Paper WR-722, 2009.

[23] The Kaiser Foundation. *Medicaid: A Primer*. http://www.statehealthfacts.org/, 2013.

[24] Eric French and John Bailey Jones. On the distribution and dynamics of health care costs. *Journal of Applied Econometrics*, 19(4):705–721, 2004.

[25] Lara Gardner and Donna Gilleskie. The effects of state medicaid policies on the dynamic savings patterns of the elderly. Working Paper 12208, National Bureau of Economic Research, May 2006.

[26] Michael Grossman. On the concept of health capital and the demand for health. *Journal of Political Economy*, 80(2):223–255, 1972.

[27] Gary D. Hansen, Minchung Hsu, and Junsang Lee. Health insurance reform: The impact of a medicare buy-in. Working Paper 18529, National Bureau of Economic Research, November 2012.

[28] R. Glenn Hubbard, Jonathan Skinner, and Stephen P. Zeldes. Expanding the life-cycle model: Precautionary saving and public policy. *American Economic Review*, 84:174–179, 1994.

[29] R. Glenn Hubbard, Jonathan Skinner, and Stephen P. Zeldes. Precautionary saving and social insurance. *Journal of Political Economy*, 103(2):360–399, 1995.

[30] Michael D. Hurd. Mortality risk and bequests. *Econometrica*, 57(4):779–813, 1989.

[31] Michael D. Hurd, Daniel McFadden, and Angela Merrill. Predictors of mortality among the elderly. Working Paper 7440, National Bureau of Economic Research, 1999.

[32] Selahattin İmrohoroğlu and Sagiri Kitao. Social security, benefit claiming and labor force participation: A quantitative general equilibrium approach. Mimeo, 2009.

[33] Ahmed Khwaja. Estimating willingness to pay for medicare using a dynamic life-cycle model of demand for health insurance. *Journal of Econometrics*, 156(1):130–147, 2010.

[34] Ruud H. Koning. Kernel: A gauss library for kernel estimation. http://www.xs4all.nl/

[35] Wojciech Kopczuk and Joseph Lupton. To leave or not to leave: The distribution of bequest motives. *Review of Economic Studies*, 74:207–235, 2007.

[36] Karen Kopecky and Tatyana Koreshkova. The impact of medical and nursing home expenses and social insurance policies on savings and inequality. Mimeo, 2009.

[37] Laurence J. Kotlikoff. Health expenditures and precautionary savings. In Laurence J. Kotlikoff, editor, *What Determines Saving?* Cambridge, MIT Press, 1988, 1988.

[38] Lee Lockwood. Incidental bequests: Bequest motives and the choice to self-insure late-life risks. Mimeo, 2012.

[39] Willard G. Manning, Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, and Arleen Leibowitz. Health insurance and the demand for medical care: Evidence from a randomized experiment. *American Economic Review*, 77:251–277, 1987.

[40] Charles F. Manski. *Analog Estimation Methods in Econometrics*. Chapman and Hall, 1988.

[41] Samuel Marshall, Kathleen M. McGarry, and Jonathan S. Skinner. The risk of out-of-pocket health care expenditure at the end of life. Working Paper 16170, National Bureau of Economic Research, 2010.

[42] Thad Mirer. The wealth-age relation among the aged. *American Economic Review*, 69:435–443, 1979.

[43] Makoto Nakajima and Irina Telyukova. Home equity in retirement. Mimeo, 2012.

[44] Whitney K. Newey. Generalized method of moments specification testing. *Journal of Econometrics*, 29(3):229–256, 1985.

[45] Whitney K. Newey and Daniel L. McFadden. Large sample estimation and hypothesis testing. In Robert Engle and Daniel L. McFadden, editors, *Handbook of Econometrics, Volume 4*. Elsevier, Amsterdam, 1994.

[46] The Kaiser Commission on Medicaid and the Uninsured. *Medicaid: A Primer*. Menlo Park, CA, The Henry J. Kaiser Family Foundation, 2010.

[47] Serdar Ozkan. Income inequality and health care expenditures over the life cycle. Mimeo, Federal Reserve Board, 2011.

[48] Ariel Pakes and David Pollard. Simulation and the aysmptotics of optimization estimators. *Econometrica*, 57(5):1027–1057, 1989.

[49] Michael G. Palumbo. Uncertain medical expenses and precautionary saving near the end of the life cycle. *Review of Economic Studies*, 66:395–421, 1999.

[50] Svetlana Paschenko and Ponpoje Porapakkarm. Quantitative analysis of health insurance reform: Separating regulation from redistribution. Mimeo, 2012.

[51] Jorn-Steffen Pischke. Measurement error and earnings dynamics: Some estimates from the PSID validation study. *Journal of Business & Economics Statistics*, 13(3):305–314, 1995.

[52] James M. Poterba, Steven F. Venti, and David A. Wise. The asset cost of poor health. Working Paper 16389, National Bureau of Economic Research, 2010.

[53] James Powell. Estimation of semiparametric models. In Robert Engle and Daniel L McFadden, editors, *Handbook of Econometrics, Volume 4*. Elsevier, Amsterdam, 1994.

[54] Susann Rohwedder, Steven J. Haider, and Michael Hurd. Increases in wealth among the elderly in the early 1990s: How much is due to survey design? *Review of Income and Wealth*, 52(4):509–524, 2006.

[55] John Karl Scholz and Ananth Seshadri. Health and wealth in a lifecycle model. Mimeo, 2012.

[56] John Karl Scholz, Ananth Seshadri, and Surachai Khitatrakun. Are americans saving optimally for retirement? *Journal of Political Economy*, 114:607–643, 2006.

[57] Motohiro Yogo. Portfolio choice in retirement: Health risk and the demand for annuities, housing, and risky assets. Mimeo, University of Pennsylvania, 2009.

[58] Katherine Young, Rachel Garfield, MaryBeth Musumeci, Lisa Clemans-Cope, and Emily Lawton. Medicaid's role for dual eligible beneficiaries. *Kaiser Commission on Medicaid and the Uninsured*, pages 1–19, 2012.

Appendix A: The MCBS data

In order to assess the accuracy of the model's predictions, we compare model-predicted distributions of out-of-pocket and Medicaid medical spending to the distributions observed in the AHEAD and MCBS data in the main text of the paper. Here, we describe in greater detail the construction and accuracy of the MCBS data.

The MCBS is a nationally representative survey of disabled and age-65+ Medicare beneficiaries. The survey contains an over-sample of beneficiaries older than 80 and disabled individuals younger than 65. Respondents are asked about health status, health insurance, and health care expenditures (from all sources). The MCBS data are matched to Medicare records, and medical expenditure data are created through a reconciliation process that combines information from survey respondents with Medicare administrative files. As a result, the survey is thought to give extremely accurate data on Medicare payments and fairly accurate data on out-of-pocket and Medicaid payments. As in the AHEAD survey, the MCBS survey includes information on those who enter a nursing home or die. Respondents are interviewed up to 12 times over a 4 year period. We aggregate the data to an annual level.

In order to assess the quality of the medical expenditure data in the MCBS, we benchmark it against administrative data from the Medicaid Statistical Information System (MSIS) and survey data from the AHEAD. For Medicare payments, the match is close. For example, when using population weights, the number of Medicare beneficiaries lines up almost exactly with the aggregate statistics. More important, Medicare expenditures per beneficiary are very close. Over the 1996-2006 period, MCBS Medicare expenditures per capita for the age 65+ population are \$6,070, only 11% smaller than the value of \$6,820 in the official statistics.[6]

The MCBS also accurately measures the share of the population receiving Medicaid benefits.[7] However, MCBS Medicaid payments for the age 65+ population are on average 32% smaller than what administrative data from the MSIS suggest. Table 9 compares the distribution of the MSIS administrative payment data (taken from Young et al. [58]) to data from the MCBS. We show the MCBS distribution for all Medicare/Medicaid beneficiaries, the set closest to the the sample in the MSIS data. Table 9 shows both means and means conditional on the distribution of payments. The MSIS data show that the least costly 50% of all Medicaid enrollees account for only 0.9% of total Medicaid payments, whereas the most costly 5% of all beneficiaries are responsible for 41% of payments. Although the MCBS data match the MSIS data

[6]Medicare statistics are located at http://www.census.gov/compendia/statab/cats/health_nutrition/medicare_medicaid.html.

[7]According to MCBS data, there were on average 5.1 million Medicaid beneficiaries over the 1996-2006 period, versus 4.7 million Medicaid beneficiaries in the MSIS data. This difference potentially reflects a small number of Medicaid age 65+ individuals who are classified as "disabled" instead of "aged" in the MSIS data. Medicaid MSIS statistics are located at https://www.cms.gov/Research-Statistics-Data-and-Systems/Computer-Data-and-Systems/MedicaidDataSourcesGenInfo/MSIS-Tables.html.

Expenditure Percentile	Percentage of Medicaid enrollees	Percentage of Medicaid expenditures (MSIS)	Average expenditure per enrollee (MSIS)	Average expenditure per enrollee (MCBS)
everyone	100%	100%	13,410	9,090
95-100%	5%	40.1%	100,060	69,810
90-95%	5%	20.4%	50,180	36,890
70-90%	20%	32.4%	21,940	14,460
50-70%	20%	5.5%	3,690	3,090
0-50%	50%	0.9%	240	490
Note: 2008 MSIS data, adjusted to 2005 dollars				

Table 9: Medicaid Enrollment and Expenditures by Enrollee Spending Percentile, MSIS versus MCBS.

well across the bottom 70% of the distribution, the top 5% of all payments in the MSIS average \$100,060, whereas in the MCBS they are \$69,810. Limiting the MCBS sample to our estimation sample (retired singles who meet our age selection criteria: greater than 70 in 1994, 72 in 1996, 74 in 1998, etc.) leads to higher payments: average Medicaid payments of Medicaid beneficiaries in this MCBS subsample are \$13,620.

	AHEAD data				MCBS data		
Income Quintile	Total income	Annuity income	Out-of-pocket expenses	Medicaid recipiency	Total income	Out-of-pocket expenses	Medicaid recipiency
1	7,740	4,820	2,360	60.9%	6,680	3,540	69.6%
2	10,290	8,270	3,830	28.1%	9,880	4,800	41.5%
3	15,500	10,900	4,470	11.0%	13,570	5,710	15.3%
4	19,290	14,390	5,550	5.6%	19,520	6,290	7.7%
5	33,580	26,300	6,340	3.0%	44,430	7,020	5.4%

Table 10: Income, out-of-pocket spending, and Medicaid recipiency rates, AHEAD versus MCBS.

One might be concerned with large medical expenses outliers. We thus truncate the top and bottom 1 percent of the expenditure distributions. The cross-sectional data behind Figure 9 and Table 4 are constructed accordingly.

The next set of benchmarking exercises that we perform is for out-of-pocket medical spending, Medicaid recipiency and income between the AHEAD and MCBS.We restrict the sample to singles (over the sample period) who meet the AHEAD age criteria (at least 70 in 1994, 72 in 1996, ...) and who are not working over the sample period, just as we do in the AHEAD data. We construct a measure of permanent income, which is the percentile rank of total income over the period we observe these individuals (the MCBS asks only about total income). The first four columns of Table 10 show sample statistics from the full AHEAD sample while the final three columns of the table shows sample statistics from the MCBS sample. The first statistics we compare are income. Total income in the AHEAD data (including asset and other non-annuitized income) lines up well with total income in the MCBS data, although income in the top quintile of the MCBS is higher than in the AHEAD. Next we compare out-of-pocket medical spending in the MCBS and AHEAD. Out-of-pocket medical expenditure (including insurance payments) averages \$2,360 in the bottom PI quintile and \$6,340 in the top quintile in the AHEAD. In comparison, the same numbers in the MCBS data are \$3,540 and \$7,020. Overall, out-of-pocket medical spending in the MCBS and AHEAD are similar, which may be surprising given that the two surveys each have their own advantages in terms of survey methodology.[8] The share of the population receiving Medicaid benefits is also very similar in the AHEAD and MCBS. 61% and 70% of those in the bottom PI quintile receive Medicaid in the the AHEAD and MCBS, respectively. 3% of those in the top quintile receive Medicaid benefits in the AHEAD whereas it is 5% in the MCBS. The higher Medicaid recipiency rate in the MCBS might reflect that the MCBS data has administrative information on whether individuals are receiving Medicaid benefits, which eliminates underreporting problems.

We also assessed the usefulness of the Medicaid-related data in MEPS. A key problem with the MEPS data, however, is that that it does not include information on nursing home stays or expenses in the last few months of life. Using data from MSIS, Young et al. [58] report that among those aged 65 and older, 79% of all Medicaid expenses are for long term care (although only 14% of these beneficiaries are receiving long term care). The MEPS data are only useful for understanding the remaining 21% of Medicaid payments. Consistent with this fact, mean Medicaid payments in the MEPS for beneficiaries are \$3,499, whereas they are \$13,414 according to the administrative data from the Medicaid Statistical Information System (MSIS).For non-nursing home expenses, however, we can use the MEPS data to understand the distribution of Medicaid payments and to verify the accuracy of the models predictions.

Appendix B: Moment conditions and asymptotic distribution of parameter estimates

[8]There are more detailed questions underlying the out-of-pocket medical expense questions in the AHEAD, including the use of "unfolding brackets". Respondents can give ranges for medical expense amounts, instead of a point estimate or "don't know" as in the MCBS. The MCBS has the advantage that forgotten medical out-of-pocket medical expenses will be imputed if Medicare had to pay a share of the health event.

Recall that we estimate the parameters of our model in the two steps. In the first step, we estimate the vector χ, the set of parameters than can be estimated with explicitly using our model. In the second step, we use the method of simulated moments (MSM) to estimate the remaining parameters, which are contained in the $M \times 1$ vector Δ. The elements of Δ are ν, ω, β, $\underline{Y}$, $\underline{u}$, θ, k, and the parameters of $\ln \mu(\cdot)$. Our estimate, $\hat{\Delta}$, of the "true" parameter vector Δ_0 is the value of Δ that minimizes the (weighted) distance between the life-cycle profiles found in the data and the simulated profiles generated by the model.

For each calendar year $t \in \{t_0, ..., t_T\} = \{1996, 1998, 2000, 2002, 2004, 2006, 2008, 2010\}$, we match median assets for $Q_A = 5$ permanent income quintiles in $P = 5$ birth year cohorts.[9] The 1996 (period-t_0) distribution of simulated assets, however, is bootstrapped from the 1996 data distribution, and thus we match assets to the data for 1998, ..., 2006. In addition, we require each cohort-income-age cell have at least 10 observations to be included in the GMM criterion.

Suppose that individual i belongs to birth cohort p and his permanent income level falls in the qth permanent income quintile. Let $a_{pqt}(\Delta, \chi)$ denote the model-predicted median asset level for individuals in individual i's group at time t, where χ includes all parameters estimated in the first stage (including the permanent income boundaries). Assuming that observed assets have a continuous conditional density, a_{pqt} will satisfy

$$\Pr\Big(a_{it} \leq a_{pqt}(\Delta_0, \chi_0)\,|p, q, t, \text{individual } i \text{ observed at } t\Big) = 1/2.$$

The preceding equation can be rewritten as a moment condition (Manski [40], Powell [53] and Buchinsky [9]). In particular, applying the indicator function produces

$$E\big(1\{a_{it} \leq a_{pqt}(\Delta_0, \chi_0)\} - 1/2\,|p, q, t, \text{individual } i \text{ observed at } t\big) = 0. \tag{19}$$

Letting $\mathcal{I}_q$ denote the values contained in the qth permanent income quintile, we can convert this conditional moment equation into an unconditional one (e.g., Chamberlain [11]):

$$\begin{aligned} E\Big([1\{a_{it} \leq a_{pqt}(\Delta_0, \chi_0)\} - 1/2] \times 1\{p_i = p\} \times 1\{I_i \in \mathcal{I}_q\} \\ \times 1\{\text{individual } i \text{ observed at } t\} \bigm| t\Big) = 0 \end{aligned} \tag{20}$$

for $p \in \{1, 2, ..., P\}$, $q \in \{1, 2, ..., Q_A\}$, $t \in \{t_1, t_2..., t_T\}$.

We also include several moment conditions relating to medical expenses. To use these moment conditions, we first simulate medical expenses at an annual frequency, and then take two-year averages to produce a measure of medical expenses comparable to the ones contained in the AHEAD.

[9] Because we do not allow for macro shocks, in any given cohort t is used only to identify the individual's age.

As with assets, we divide individuals into 5 cohorts and match data from 7 waves covering the period 1998-2010. (Because the model starts in 1996, while the medical expense data are averages over 1995-96, we cannot match the first wave.) The moment conditions for medical expenses are split by permanent income as well. However, we combine the bottom two income quintiles, as there is very little variation in out-of-pocket medical expenses in the bottom quintile until very late in life; $Q_M = 4$.

We require the model to match median out-of-pocket medical expenditures in each cohort-income-age cell. Let $m^{50}_{pqt}(\Delta, \chi)$ denote the model-predicted 50th percentile for individuals in cohort p and permanent income group q at time (age) t. Proceeding as before, we have the following moment condition:

$$E\Big([1\{m_{it} \leq m^{50}_{pqt}(\Delta_0, \chi_0)\} - 0.5] \times 1\{p_i = p\} \times 1\{I_i \in \mathcal{I}_q\} \\ \times 1\{\text{individual } i \text{ observed at } t\} \;\Big|\; t\Big) = 0 \tag{21}$$

for $p \in \{1, 2, ..., P\}$, $q \in \{1, 2, ..., Q_M\}$, $t \in \{t_1, t_2..., t_T\}$.

To fit the upper tail of the medical expense distribution, we require the model to match the 90th percentile of out-of-pocket medical expenditures in each cohort-income-age cell. Letting $m^{90}_{pqt}(\Delta, \chi)$ denote the model-predicted 90th percentile, we have the following moment condition:

$$E\Big([1\{m_{it} \leq m^{90}_{pqt}(\Delta_0, \chi_0)\} - 0.9] \times 1\{p_i = p\} \times 1\{I_i \in \mathcal{I}_q\} \\ \times 1\{\text{individual } i \text{ observed at } t\} \;\Big|\; t\Big) = 0 \tag{22}$$

for $p \in \{1, 2, ..., P\}$, $q \in \{1, 2, ..., Q_M\}$, $t \in \{t_1, t_2..., t_T\}$.

To pin down the autocorrelation coefficient for ζ (ρ_m), and its contribution to the total variance $\zeta + \xi$, we require the model to match the first and second autocorrelations of logged medical expenses. Define the residual R_{it} as

$$R_{it} = \ln(m_{it}) - \overline{\ln m}_{pqt}, \\ \overline{\ln m}_{pqt} = E(\ln(m_{it})|pi = p, qi = q, t)$$

and define the standard deviation σ_{pqt} as

$$\sigma_{pqt} = \sqrt{E\big(R^2_{it}|p_i = p, q_i = q, t\big)}.$$

Both $\overline{\ln m}_{pqt}$ and σ_{pqt} can be estimated non-parametrically as elements of χ. Using these quantities, the autocorrelation coefficient AC_{pqtj} is:

$$AC_{pqtj} = E\left(\frac{R_{it}R_{i,t-j}}{\sigma_{pqt}\ \sigma_{pq,t-j}} \;\middle|\; p_i = p, q_i = q\right).$$

Let $AC_{pqtj}(\Delta, \chi)$ be the jth autocorrelation coefficient implied by the model, calculated using model values of $\overline{\ln m}_{pqt}$ and σ_{pqt}. The resulting moment condition for the first autocorrelation is

$$E\left(\left[\frac{R_{it}R_{i,t-1}}{\sigma_{pqt}\ \sigma_{pq,t-1}} - AC_{pqt1}(\Delta_0, \chi_0)\right] \times 1\{p_i = p\} \times 1\{I_i \in \mathcal{I}_q\} \right.$$
$$\left. \times\, 1\{\text{individual } i \text{ observed at } t \;\&\; t-1\} \;\middle|\; t \right) = 0. \tag{23}$$

The corresponding moment condition for the second autocorrelation is

$$E\left(\left[\frac{R_{it}R_{i,t-2}}{\sigma_{pqt}\ \sigma_{pq,t-2}} - AC_{pqt2}(\Delta_0, \chi_0)\right] \times 1\{p_i = p\} \times 1\{I_i \in \mathcal{I}_q\} \right.$$
$$\left. \times\, 1\{\text{individual } i \text{ observed at } t \;\&\; t-2\} \;\middle|\; t \right) = 0. \tag{24}$$

Finally, we match Medicaid utilization (take-up) rates. Once again, we divide individuals into 5 cohorts, match data from 5 waves, and stratify the data by permanent income. We combine the top two quintiles because in many cases no one in the top permanent income quintile is on Medicaid: $Q_U = 4$.

Let $\overline{u}_{pqt}(\Delta, \chi)$ denote the model-predicted utilization rate for individuals in cohort p and permanent income group q at age t. Let u_{it} be the $\{0, 1\}$ indicator that equals 1 when individual i receives Medicaid. The associated moment condition is

$$E\Big(\big[u_{it} - \overline{u}_{pqt}(\Delta_0, \chi_0)\big] \times 1\{p_i = p\} \times 1\{I_i \in \mathcal{I}_q\}$$
$$\times\, 1\{\text{individual } i \text{ observed at } t\} \;\big|\; t\Big) = 0 \tag{25}$$

for $p \in \{1, 2, ..., P\}$, $q \in \{1, 2, ..., Q_U\}$, $t \in \{t_1, t_2..., t_T\}$.

To summarize, the moment conditions used to estimate model with endogenous medical expenses consist of: the moments for asset medians described by equation (20); the moments for median medical expenses described by equation (21); the moments for the 90th percentile of medical expenses described by equation (22); the moments for the autocorrelations of logged medical expenses described by equations (23) and (24); and the moments for the Medicaid utilization rates described by equation (25). In the end, we have a total of $J = 631$ moment conditions.

Suppose we have a dataset of I independent individuals that are each observed at up to T separate calendar years. Let $\varphi(\Delta; \chi_0)$ denote the J-element vector of moment conditions described immediately above, and let $\hat{\varphi}_I(.)$ denote its sample analog. Letting $\widehat{\mathbf{W}}_I$ denote a $J \times J$ weighting matrix, the MSM estimator $\hat{\Delta}$ is given by

$$\underset{\Delta}{\operatorname{argmin}} \frac{I}{1+\tau}\ \hat{\varphi}_I(\Delta; \chi_0)' \widehat{\mathbf{W}}_I \hat{\varphi}_I(\Delta; \chi_0),$$

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

In practice, we estimate χ_0 as well, using the approach described in the main text. Computational concerns, however, compel us to treat χ_0 as known in the analysis that follows. Under regularity conditions stated in Pakes and Pollard [48] and Duffie and Singleton [17], the MSM estimator $\hat{\Delta}$ is both consistent and asymptotically normally distributed:

$$\sqrt{I}\left(\hat{\Delta} - \Delta_0\right) \rightsquigarrow N(0, \mathbf{V}),$$

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

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

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

$$\mathbf{D} = \frac{\partial \varphi(\Delta; \chi_0)}{\partial \Delta'}\bigg|_{\Delta=\Delta_0} \tag{26}$$

is the $J \times M$ gradient matrix of the population moment vector; and $\mathbf{W} = \text{plim}_{I\to\infty}\{\widehat{\mathbf{W}}_I\}$. Moreover, Newey [44] shows that if the model is properly specified,

$$\frac{I}{1+\tau}\hat{\varphi}_I(\hat{\Delta}; \chi_0)'\mathbf{R}^{-1}\hat{\varphi}_I(\hat{\Delta}; \chi_0) \rightsquigarrow \chi^2_{J-M},$$

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

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

The asymptotically efficient weighting matrix arises when $\widehat{\mathbf{W}}_I$ converges to $\mathbf{S}^{-1}$, the inverse of the variance-covariance matrix of the data. When $\mathbf{W} = \mathbf{S}^{-1}$, $\mathbf{V}$ simplifies to $(1+\tau)(\mathbf{D}'\mathbf{S}^{-1}\mathbf{D})^{-1}$, and $\mathbf{R}$ is replaced with $\mathbf{S}$.

But even though the optimal weighting matrix is asymptotically efficient, it can be biased in small samples. (See, for example, Altonji and Segal [2].) We thus use a "diagonal" weighting matrix, as suggested by Pischke [51]. This diagonal weighting scheme uses the inverse of the matrix that is the same as $\mathbf{S}$ along the diagonal and has zeros off the diagonal of the matrix. This matrix delivers parameter estimates very similar to our benchmark estimates.

We estimate $\mathbf{D}$, $\mathbf{S}$, and $\mathbf{W}$ with their sample analogs. For example, our estimate of $\mathbf{S}$ is the $J \times J$ estimated variance-covariance matrix of the sample data. When estimating this matrix, we use sample statistics, so that $a_{pqt}(\Delta, \chi)$ is replaced with the sample median for group pqt.

One complication in estimating the gradient matrix $\mathbf{D}$ is that the functions inside the moment condition $\varphi(\Delta; \chi)$ are non-differentiable at certain data points; see equation (20). This means that we cannot consistently estimate $\mathbf{D}$ as the numerical derivative of $\hat{\varphi}_I(.)$. Our asymptotic results therefore do not follow from the standard GMM approach, but rather the approach for non-smooth functions described in Pakes and Pollard [48], Newey and McFadden [45] (section 7), and Powell [53].

To find $\mathbf{D}$, it is helpful to rewrite equation (20) as

$$\Pr\Big(p_i = p \;\&\; I_i \in \mathcal{I}_q \;\&\; \text{individual } i \text{ observed at } t\Big)\times$$
$$\left[\int_{-\infty}^{a_{pqt}(\Delta_0,\chi_0)} f\Big(a_{it} \;\big|\; p, I_i \in \mathcal{I}_q, t\Big)\, da_{it} \;-\; \frac{1}{2}\right] = 0. \tag{27}$$

It follows that the rows of $\mathbf{D}$ are given by

$$\Pr\Big(p_i = p \;\&\; I_i \in \mathcal{I}_q \;\&\; \text{individual } i \text{ observed at } t\Big)\times$$
$$f\Big(a_{pqt} \;\big|\; p, I_i \in \mathcal{I}_q, t\Big) \times \frac{\partial a_{pqt}(\Delta_0;\chi_0)}{\partial \Delta'}. \tag{28}$$

In practice, we find $f\big(a_{pfqt}|p,q,t\big)$, the conditional p.d.f. of assets evaluated at the median a_{pqt}, with a kernel density estimator written by Koning [34]. The gradients for equations (21) and (22) are found in a similar fashion.

Appendix C: Demographic transition probabilities in the HRS/AHEAD

Let $h_t \in \{0,1,2,3\}$ denote death ($h_t = 0$) and the 3 mutually exclusive health states of the living (nursing home = 1, bad = 2, good = 3, respectively). Let x be a vector that includes a constant, age, permanent income, gender, and powers and interactions of these variables, and indicators for previous health and previous health interacted with age. Our goal is to construct the likelihood function for the transition probabilities.

Using a multivariate logit specification, we have, for $i \in \{1,2,3\}$, $j \in \{0,1,2,3\}$,

$$\begin{aligned}
\pi_{ij,t} &= \Pr(h_{t+1} = j |\, h_t = i) \\
&= \gamma_{ij} \Big/ \sum_{k \in \{0,1,2,3\}} \gamma_{ik}, \\
\gamma_{i0} &\equiv 1, \quad \forall i, \\
\gamma_{1k} &= \exp\left(x\beta_k\right), \quad k \in \{1,2,3\}, \\
\gamma_{2k} &= \exp\left(x\beta_k\right), \quad k \in \{1,2,3\}, \\
\gamma_{3k} &= \exp\left(x\beta_k\right), \quad k \in \{1,2,3\},
\end{aligned}$$

where $\{\beta_k\}_{k=0}^{3}$ are sets of coefficient vectors and of course $\Pr(h_{t+1} = 0 |\, h_t = 0) = 1$.

The formulae above give 1-period-ahead transition probabilities, $\Pr(h_{t+1} = j \mid h_t = i)$. What we observe in the AHEAD dataset, however, are 2-period ahead probabilities, $\Pr(h_{t+2} = j \mid h_t = i)$. The two sets of probabilities are linked, however, by

$$\begin{aligned}
\Pr(h_{t+2} = j \mid h_t = i) &= \sum_k \Pr(h_{t+2} = j \mid h_{t+1} = k) \Pr(h_{t+1} = k \mid h_t = i) \\
&= \sum_k \pi_{kj,t+1} \pi_{ik,t}.
\end{aligned}$$

This allows us to estimate $\{\beta_k\}$ directly from the data using maximum likelihood.