Medicaid Insurance in Old Age

Preliminary and Incomplete

Mariacristina De Nardi, Eric French, and John Bailey Jones*

December 19, 2011

Abstract

There are two main pathways to Medicaid eligibility for people over age 65: either having low assets and income, or being impoverished by large medical expenses. The first group of recipients mostly consists of the life-long poor, while the second group includes people who became poor only after incurring large medical expenses late in life. We document Medicaid take-up rates by age, permanent income, and gender in the data. We then construct a model that explicitly allows for these two pathways to Medicaid, and in which retired single people optimally choose consumption, Medicaid application if eligible, medical spending and saving. People in our model face uncertainty about their health, lifespan and needs for medical goods and services and nursing home stays. We show how well the model matches important features of the data and we analyze the degree of insurance provided

*Mariacristina De Nardi: Federal Reserve Bank of Chicago and NBER. Eric French: Federal Reserve Bank of Chicago. John Bailey Jones: University at Albany, SUNY. For useful comments and suggestions we thank Karl Scholz. De Nardi gratefully acknowledges financial support from NSF grant SES-0317872. 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 Federal Reserve System, the National Science Foundation, the Social Security Administration or the MRRC. Jones gratefully acknowledges the hospitality of the Federal Reserve Bank of Chicago.

by current programs for people with different lifetime incomes, assets, health status, and gender. We compute the costs and benefits of the Medicaid program for people of different ages and lifetime resources, and hence the degree of redistribution provided by Medicaid. Finally, we study the effects of different health care costs reforms.

1 Introduction

In the United States, there are two main 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. 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; in fact, Medicaid now assists 70 percent of nursing home residents.[1]

Medicaid-eligible individuals can be divided into two main groups. The first group comprises the *categorically needy*, whose assets and income fall below certain thresholds. People who receive Supplemental Social Insurance (SSI), typically qualify as categorically needy for the Medicaid program.

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 resources become small in comparison. Rather than being lifetime poor, these people consist of lower and middle-class individuals who become impoverished by medical shocks. This medically needy provision thus provides insurance only against catastrophic medical expenses, such as expenses resulting from long nursing home stays.

Another important characteristic of Medicaid is that it is the payer of "last resort," which means that Medicare and private insurance pay their shares of medical expenses in advance of Medicaid. The individual then spends down his assets to a "disregard" amount, and finally Medicaid contributes. The disregard amount typically includes the value of the individual's main residence, her car, some personal items, and a very small amount of financial wealth (typically $2,000). In addition, in the case of the medically needy, Medicaid contributes only if medical expenses are high relative to the person's income.

Because Medicaid provides some insurance against health shocks in old age, but restricts the benefits to those with assets below the disregard, it discourages saving. Hubbard et al. [19] and Scholz et al. [33] argue that means-tested social insurance programs (in the form of a minimum consumption floor) provide strong incentives for low-income individuals not to

[1] Statistics from the Kaiser Foundation [27].

save. De Nardi et al. [11] find that reducing the generosity of social insurance significantly increases the saving of elderly singles. Kopecky and Koreshkova [22] find that old-age medical expenses, and the coverage of these expenses provided by Medicaid, have large effects on aggregate capital accumulation. Brown and Finkelstein [5] conclude that Medicaid could explain the lack of private long-term care insurance for about two-thirds of the wealth distribution.

The existence of two pathways to Medicaid eligibility, with different requirements and generosities, implies that Medicaid affects different segments of the population in very different ways. The categorically needy provision affects the saving of people who have been poor throughout most of their lives, while it has no impact on the saving of middle- and upper-income people, who at a minimum fail the income test. In contrast, the medically needy provision provides insurance to people with higher income and assets who are still at risk of being impoverished by expensive medical conditions. However, eligibility under the Medically Needy provision is more stringent, and generosity of the benefits is sometimes more limited.

In this paper we extend the existing literature by constructing a richer and more realistic model of Medicaid and old age risks. We focus on the costs and benefits of Medicaid, and by analyzing its insurance role for people with different wealth levels, health, and gender. We estimate the parameters of our model rather than calibrating them to previous studies, which might have features which are inconsistent with the model at hand. We require our model to fit well across the entire income distribution, rather than simply explain mean or median behavior. We model medical expenditure as an endogenous choice. This allows us to consider how Medicaid reforms affect total medical spending. We model social insurance as providing a utility floor, rather than a fixed expenditure floor. This allows means-tested transfers to vary with medical needs in a way consistent with actual practice.[2] Due to the richness and complexity of our framework, we focus on the post-retirement part of the life-cycle and adopt a partial equilibrium approach.

[2]Three recent papers contain life-cycle models where the choice of medical expenditures also affects health outcomes. In addition to having different emphases, both papers model Medicaid in ways different from ours. Feng [13] models Medicaid as an insurance policy with no premiums and extremely low—possibly zero—co-payment rates. Fonseca et al. [16] assume that the consumption floor is invariant to medical needs (private conversation with Pierre-Carl Michaud). Ozkan [28] assumes that indigent individuals receive curative, but not preventative, care.

We make several contributions. First, we document the relevant patterns in the data using the Asset and Health Dynamics of the Oldest Old (AHEAD) data set. We show how Medicaid take-up rate varies with age, birth cohort cohort, net worth, and annuity income. We find that the average Medicaid recipiency rate for people in the bottom quintile of the permanent income distribution is just under 70% and stays more or less constant during retirement. Medicaid recipiency by higher-income retirees is significantly lower, but increases with age. Interestingly, this increase tends to happen at more advanced ages for people with higher permanent income quintiles, reflecting the fact that survivors with higher lifetime resources run out of savings (and thus qualify for Medicaid) later on in life.

Second, we construct and estimate a rich model of medical expense risk in old age. From the institutional standpoint, we explicitly model the two pathways to Medicaid that we have described and the Medicare program. From the agents' standpoint, we allow for heterogeneity in wealth, lifetime income, health, gender, and life expectancy. We allow people to optimally 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. The agents in the model face uncertainty about their, health, lifespan and medical needs and nursing home stays. This uncertainty is partially offset by insurance provided by the government and private institutions.

Third, we estimate the amount of insurance that Medicaid provides for retired single people and its degree of redistribution. Rich people tend to have higher lifetime income, but also to live longer. We use our estimated model to compute expected Medicaid payments by gender, permanent income, and health status.

Fourth, we compute the insurance value of Medicaid by comparing the individual's valuation of expected Medicaid transfers with the cost of the actual transfers that they receive.

Fifth, we consider how changes in Medicaid affect saving, consumption, and medical expenditures.

We find that the model closely matches the life-cycle profiles of assets, out-of-pocket medical spending, and Medicaid recipience rates for elderly singles in different cohort and permanent income groups. It also generates an elasticity of total medical expenditures to co-payment changes that is close to the one estimated in the data.

We also find that the current Medicaid system provides different kinds of insurance to households with different resources. Households in the lower per-

manent 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 pay-outs, but when they do, these pay-outs are very big and correspond to severe and expensive medical conditions. Therefore, Medicaid is an effective insurance device for the poorest, but also offers very valuable insurance to the rich by insuring them against catastrophic medical conditions.

More findings to come...

2 Some key features of the data

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 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, and 2006. This leaves us with 3,259 individuals, of whom 592 are men and 2,667 are women. Of these 3,259 individuals, 884 are still alive in 2006.

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 use data for 6 different years; 1996, 1998, 2000, 2002, 2004, and 2006. To construct the profiles, we calculate summary statistics (e.g., medians), cohort-by-cohort, for surviving individuals in each calendar year. We then order the summary statistics by cohort and age at year of observation. Moving from the left-hand-side to the right-hand-side of our graphs, we thus have data for up to five cohorts, with each cohort's data starting out at the cohort's average age in 1996.[3]

Since we want to understand the role of lifetime resources, we also stratify most of our variables by a measure of post-retirement, non-asset permanent income (PI). We measure post-retirement permanent income 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,

[3]Due to a lack of data, our graphs typically omit profiles for the oldest cohort.

defined benefit pension benefits, annuities, veterans benefits, welfare, and food stamps. Because there is a roughly monotonic relationship between lifetime earnings and the income variables that we use, our measure of post-retirement permanent income is also a good measure of lifetime permanent income.

Individuals are stratified according to permanent income and cohort. Hence, for a given cohort, we see several horizontal lines showing, for example, average Medicaid status in each permanent income group in each calendar year. These lines also identify the moment conditions we use when estimating the model.

A key advantage of the AHEAD data 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.

2.1 Medicaid

AHEAD respondents are asked whether they are currently covered by Medicaid. Figure 1 plots the fraction of Medicaid recipients by age, birth cohort and income quintile for those individuals that are still alive at each moment in time, that is, for an unbalanced panel. There are four lines for each cohort because we have split the data into permanent income quintiles. However, we have merged the top two quintiles together because in many cases no one in the top permanent income 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 then again every two years until 2006. The other cohorts start from older initial ages and are also followed for ten years. The graph reports the fraction of Medicaid recipients for each cohort and permanent-income grouping for six data points over time.

Unsurprisingly, Medicaid recipiency is inversely related to permanent income: the top line shows the fraction of Medicaid recipients in the bottom 20% of the permanent income distribution, while the bottom line shows median assets in the top 40%. For example, the top left line shows that for the

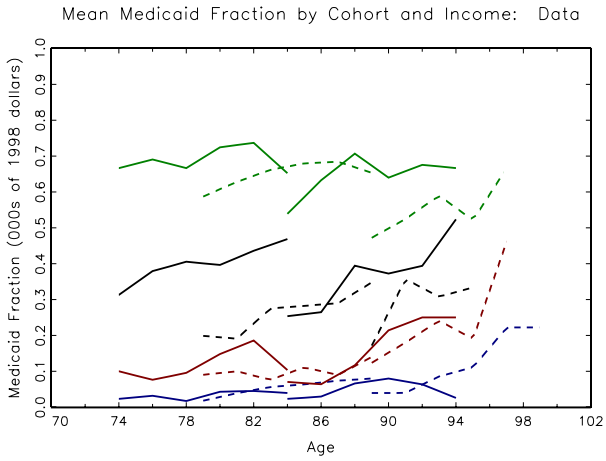

Figure 1: Medicaid take-up rates by age, cohort, and permanent income.

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.

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% over age 96. Even people with relatively large resources can be hit by medical shocks severe enough to to exhaust their assets and require the use of Medicaid.

2.2 Medical expense profiles

In all waves, AHEAD respondents are asked about what 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.

French and Jones [17] show that the medical expense data in the AHEAD line up with the aggregate statistics. For our sample, mean medical expenses are $3,712 with a standard deviation of $13,429 in 1998 dollars. Although this figure is large, it is not surprising, because Medicare did not cover prescription drugs for most of the sample period, requires co-pays for services, and caps

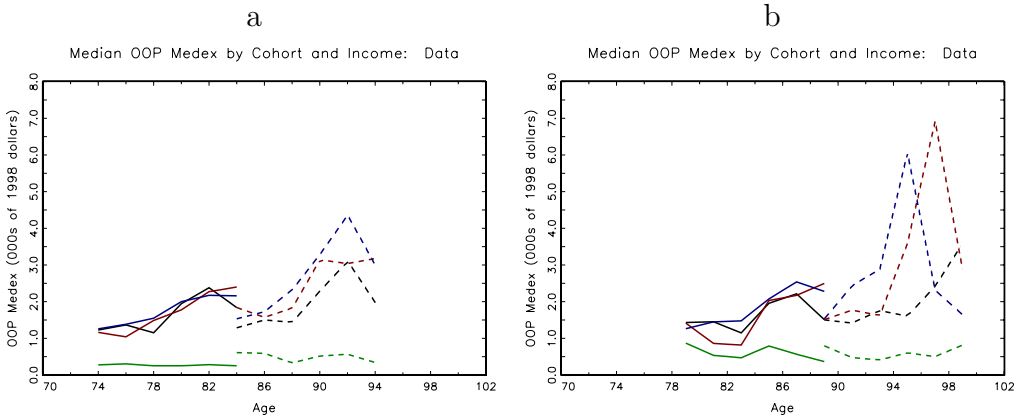

Figure 2: Median out-of-pocket medical expenditures by age, cohort, and permanent income.

the number of reimbursed nursing home and hospital nights.

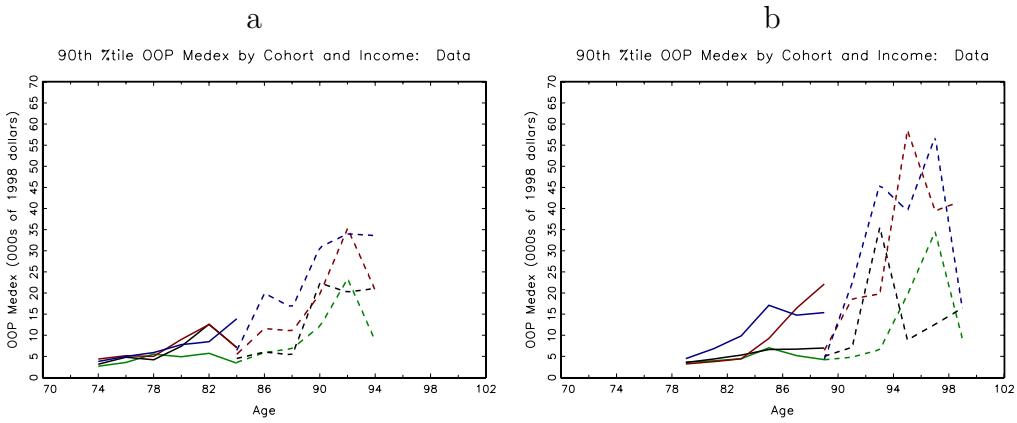

Figure 3: 90th percentile out-of-pocket medical expenditures by age, cohort, and permanent income.

Figures 2 and 3 display median and 90th percentile of out-of-pocket medical expenses by age, cohort, and permanent income, 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 as people reach very advanced ages and that

this increase is especially strong for people in the highest permanent income quintiles.

2.3 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. We do not use 1994 assets because they were underreported (Rohwedder et al. [32]).

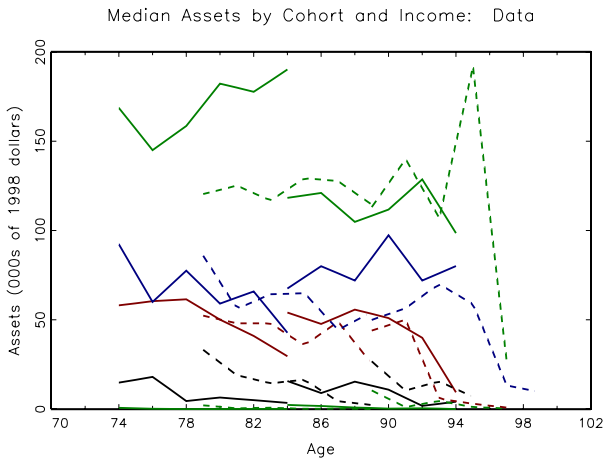

Figure 4: Median assets by age, cohort, and permanent income.

Figure 4 reports median assets by cohort, age, and permanent income. There are five lines for each cohort because we have split the data into permanent income quintiles. However, the fifth, bottom line is hard to distinguish from the horizontal axis because households in the lowest permanent income quintile hold few assets.

Unsurprisingly, assets turn out to be monotonically increasing in income, so that the bottom line shows median assets in the lowest income quintile, while the 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 $170,000 and then stayed rather stable over time:

$150,000 at age 76, $160,000 at age 78, $180,000 at ages 80 and 82, and $190,000 at age 84.

For all permanent income quintiles in these cohorts, the assets of surviving individuals neither rise rapidly nor decline rapidly with age.[4] If anything, those with high income tend to have their assets increase as they age, whereas those with low income tend to have their assets decrease.

3 The model

We focus on single people who have already retired. This allows us to abstract from labor supply and retirement decisions and from complications arising from family dynamics such as the transition from two family members to one.

We assume that people are hit by medical needs shocks, such as cancer, diabetes, a heart attack, or a broken bone. These shocks affect their marginal utility of consuming medical goods and services. Individuals optimally choose how much to spend in response to these shocks.

A complementary approach is that of Grossman [18], in which medical expenses represent investments in health capital, which in turn decreases mortality (e.g., Yogo [34]) or improves health.

While some studies find that medical expenditures have significant effects on health and/or survival (Card et al. [7]; Doyle [9]), many others find small effects (Brook et al. [3]; Fisher et al. [15]; Finkelstein and McKnight [14]; Khwaja [20]); see De Nardi et al. [11] for a discussion. These findings suggest that the effects of medical expenditures on the health outcomes are at a minimum extremely difficult to identify.[5] 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.

[4]The low rate at which the elderly deplete their wealth is a long-standing puzzle (e.g., Mirer [24]).

[5]Identification problems include reverse causality—sick people have higher health expenditures—and a lack of insurance variation—most elderly individuals receive Medicare or Medicaid.

3.1 Modeling Medicaid and other public and private insurance mechanisms

We model two important types of health insurance. The first one pays a proportional share of the total medical expenses and can be thought of as a combination of Medicare and private insurance. Let $q(\cdot)$ 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 an individual's age, gender and health status. Health status takes on three values: good, bad, and in a nursing home. Allowing co-pay rates to depend on health status is particularly important because nursing home stays are virtually uninsured by Medicare and private insurance, leading people residing in nursing homes to face much higher co-pay rates.

The second type of health insurance that we model is means-tested, and includes Medicaid and Supplemental Social Insurance (SSI). As discussed above, we explicitly model both the categorically needy pathway to SSI and Medicaid, and the medically needy pathway to Medicaid.

Most analyses of Medicaid, and means-tested social insurance programs in general, assume that out-of-pocket medical expenses are exogenous resource shocks. In such a framework, social insurance covers the expenditure shocks of eligible individuals, and provides them with enough resources to purchase a baseline level of consumption. This effectively provides perfect insurance to people with low enough resources, by fixing their consumption (and thus utility) to a constant minimum level. To fix ideas, consider a person with no assets and no income. SSI and Medicaid will ensure that this person consumes the minimum consumption level every period, and hence has a constant utility flow every period, no matter how long this person lives and how high his medical expenses are.

In contrast, in our model medical expenses are chosen by the individual in response to medical shocks that affect their need for (and hence their utility from) medical treatment. We generalize the standard approach to social insurance to our current environment by modeling social insurance as providing a utility floor. This mechanism provides a mapping between shocks and transfers, and at the same time limits how much the insurance will pay for a given medical condition.

Our approach has two main advantages. First, and most important, it provides a realistic representation of the Medicaid program, because Medicaid provides larger transfers for more serious and expensive medical conditions.

Second, our formulation provides an insurance scheme similar to the one in in the standard model with exogenous shocks and a fixed consumption floor, at least for poorer people. In fact, with a fixed utility floor, poor people have the same amount of utility in every state of the world, just as in the standard framework with exogenous medical shocks.

To implement our utility floor formulation, for every state vector, we find the resource level $\underline{x}(\cdot)$ that is necessary to provide a flow utility of $\underline{U}_j$ for the current period. We allow the utility floor to be different for the categorically needy $j = c$, and the medically needy $j = m$.

If an individual's income and assets are below the categorically needy thresholds and he applies for Medicaid, he is given enough resources to achieve the categorically needy utility floor. If he is above the income threshold, he gets no SSI transfers and faces the utility floor for the medically needy. If his total resources net of the disregard are not enough to achieve the medically needy utility floor and he applies for Medicaid, he receives a transfer that pushes him up to the floor.

Because Medicaid is the payer of last resort, the amount it transfers is net of the individual's financial resources, less the resource disregard amount A_d. Moreover, the person's savings for the next period are constrained to be less than or equal to the resource disregard A_d, as the person needs to spend down his own resources in order to receive Medicaid. As a result, someone with a large amount of assets might not want to apply for Medicaid even if eligible, depending on how serious and persistent his medical needs are.

3.2 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 given. The individual then chooses consumption, medical expenditure, and saves. Finally, the survival shock hits.

Letting h_t denote the retiree's health status, we parameterize the prefer-

ence shifter for medical goods and services (the needs shock) as

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

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, permanent income, and age. The elements of the health status transition matrix are

$$\pi_{j,k,g,I,t} = \Pr(h_{t+1} = k | h_t = j, g, I, t), \quad j, k \in \{1, 2, 3\}. \tag{7}$$

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, g, permanent income, I, and age t:

$$y_t = y(g, I, t). \tag{8}$$

3.3 The individual's problem

Consider a single person, either male or female, seeking to maximize his or her expected lifetime utility at age t, $t = t_{r+1}, ..., T$, where t_r is the retirement age. His flow utility from consumption and medical expenditures is given by

$$v(c_t, m_t, h_t, \zeta_t, \xi_t, t) = \frac{1}{1-\nu} c_t^{1-\nu} + \mu(h_t, \zeta_t, \xi_t, t) \frac{1}{1-\omega} m_t^{1-\omega}, \tag{9}$$

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 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 thus solves:

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

where λ_t is the multiplier on the intratemporal budget constraint and $q(\cdot) = q(t, g, h_t)$ is the individual's co-insurance rate. The first-order conditions for this problem reduce to

$$m_t = \left(\frac{\mu(\cdot)}{q(\cdot)} \right)^{1/\omega} c_t^{\nu/\omega}. \tag{10}$$

Hence, combining the within period utility function (11) with equation (10), we get

$$u^*(c_t, \mu(\cdot), q(\cdot)) = \frac{1}{1-\nu} c_t^{1-\nu} + \mu(\cdot) \frac{1}{1-\omega} \left(\frac{\mu(\cdot)}{q(\cdot)} \right)^{1/\omega} c_t^{\nu(1-\omega)/\omega}. \tag{11}$$

Similarly, we can calculate the individual's total expenditures on consumption and medical co-payments as

$$x^*(c_t, \mu(\cdot), q(\cdot)) = c_t + q(\cdot) \left(\frac{\mu_t}{q(\cdot)} \right)^{1/\omega} {c_t}^{\nu/\omega}. \tag{12}$$

To express transfers as an explicit function of the utility floor $\underline{U}_j$, we use equations (11) and (12). First, we find the consumption level associated with the utility floor, $\underline{c}(\mu(\cdot), q(\cdot), \underline{U}_j)$, which can be done by solving the following equation numerically:

$$\underline{c}(\mu(\cdot), q(\cdot), \underline{U}_j) \equiv c : u^*(c, \mu(\cdot), q(\cdot) = \underline{U}_j.$$

This expression tells us how big c must be, given $q(\cdot)$ and $\mu(\cdot)$, to achieve utility $\underline{U}_j$. Using this result, we can then find medical expenditures and total

personal expenditures

$$\begin{aligned}
\underline{m}(\mu(\cdot), q(\cdot), \underline{U}_j) &= \left(\frac{\mu(\cdot)}{q(\cdot)}\right)^{1/\omega} \underline{c}(\mu(\cdot), q(\cdot), \underline{U}_j)^{\nu/\omega}, \\
\underline{x}(\mu(\cdot), q(\cdot), \underline{U}_j) &= \underline{c}(\mu(\cdot), q(\cdot), \underline{U}_j) + q(\cdot)\left(\frac{\mu_t}{q(\cdot)}\right)^{1/\omega} \underline{c}(\mu(\cdot), q(\cdot), \underline{U}_j)^{\nu/\omega}.
\end{aligned}$$

To be categorically needy a person's income and assets need to be below $\underline{Y}$ and $\underline{A}$, respectively. Besides being the maximum amount of income (excluding disregards) that one can have and still qualify for SSI/Medicaid $\underline{Y}$ is also the maximum SSI benefit that one can receive. Let $y_n(ra_t + y_t)$ denote the individual's after-tax income, where a_t denotes assets and r is the real interest rate. The SSI benefit equals $\underline{Y} - \max\{y_n(ra_t + y_t) - p_a, 0\}$, where p_a is a personal income allowance.

If a person is categorically needy and applies for SSI and Medicaid, he receives the $\underline{Y}$ transfer every period, and Medicaid goods and services as dictated by his medical needs (and the utility floor $\underline{U}_j = \underline{U}_c$). The transfer scheme for the Medicaid categorically needy and SSI recipients is thus given by:

$$\begin{aligned}
b_c\big(a_t, y_t, \mu(\cdot), q(\cdot)\big) &= \underline{Y} - \max\{y_n(ra_t + y_t) - p_a, 0\} \\
&\quad + \max\big\{0, \underline{x}(\cdot, \underline{U}_c) - \max\{a_t + y_n(ra_t + y_t) + \underline{Y} - A_d, 0\}\big\}, \qquad (13)
\end{aligned}$$

where A_d is the resources disregard level.

If the person's total income is above $\underline{Y}$, there will be no SSI transfer, and $\underline{U}_j = \underline{U}_m$. Thus, if the person applies for Medicaid, transfers are given by

$$b_m(t, a_t, g, h_t, I, \zeta_t, \xi_t) = \max\big\{\underline{x}(\cdot, \underline{U_m}) - \max\{a_t + y_n(ra_t + y_t) - A_d, 0\}\big\}. \quad (14)$$

Each period the person will decide whether to be on Medicaid or not. Let us use the indicator function I_m 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. 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:

$$V_t(a_t, g, h_t, I, \zeta_t, \xi_t) = \max_{c_t, a_{t+1}, I_m} \Bigg\{ u^*(c_t, \mu(\cdot), q(\cdot)) + \beta s_{g,h,I,t} E_t \Big(V_{t+1}(a_{t+1}, g, h_{t+1}, I, \zeta_{t+1}, \xi_{t+1}) \Big) + \beta(1 - s_{g,h,I,t})\theta \frac{(e+k)}{1-\nu}^{(1-\nu)} \Bigg\}, \tag{15}$$

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

$$a_{t+1} = a_t + y_n(ra_t + y_t) - x^*(c_t, \mu(\cdot), q(\cdot)) \geq 0, \tag{16}$$

If $I_m = 1$, i.e., the person does apply for Medicaid, we have

$$a_{t+1} = b_j(\cdot) + a_t + y_n(ra_t + y_t) - x^*(c_t, \mu(\cdot), q(\cdot)) \geq 0, \tag{17}$$

$$a_{t+1} \leq A_d, \tag{18}$$

where $b_j(\cdot) = b_c(\cdot)$ if $Y_t \leq \underline{Y}$ and $b_j(\cdot) = b_m$ otherwise.

4 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,$\underline{c}$,θ, k, 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. Because the effects of Medicaid depend directly on an individual's asset holdings, we match median asset holdings by birth-year cohort, permanent income, and calendar year. We sort individuals into PI quintiles, and the 5 birth-year cohorts described in section 2. We then compare data and model- generated cell medians in 5 different years (1998, 2000, 2002, 2004, and 2006).[6]

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

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

4. To improve our model's policy predictions, we match Medicaid usages rate in each year-cohort-PI "quintile" cell (the top two quintiles are merged).

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 a series of health, medical expense, and mortality shocks consistent with the stochastic processes described in the model section. We give each simulated person the entire health and mortality history realized by a person in the AHEAD data with the same initial conditions. 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

[6]Simulated agents are endowed with asset levels drawn from the 1996 data distribution. Cells with less than 10 observations are excluded from the moment conditions.

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 A contains a detailed description of our moment conditions, the weighting matrix in our GMM criterion function, and the asymptotic distribution of our parameter estimates.

5 First-step estimation results

In this section, we briefly discuss the life cycle profiles of the stochastic variables used in our dynamic programming model. The processes for income and co-pay rates were estimated for our analysis in De Nardi et al. [11], and are described in more detail there. The demographic transition probabilities are new.

5.1 Income profiles

We model non-asset income as a function of age, sex, health status, 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 $20,000 in the top quintile; median wealth holdings for the two groups are zero and just under $200,000, respectively.

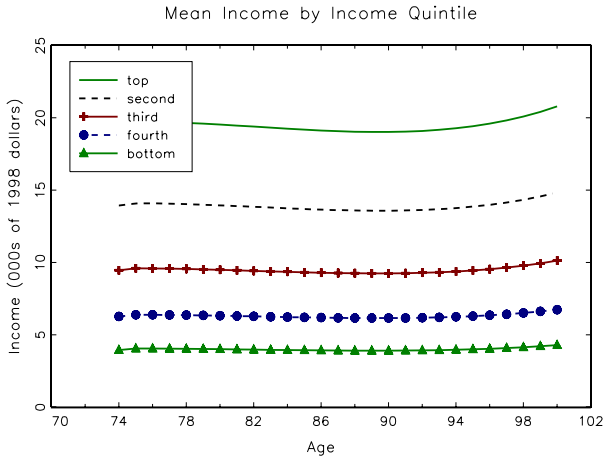

Figure 5: Average income, by permanent income quintile.

5.2 Mortality and health status

We estimate health transitions and mortality rates simultaneously, using a variant of the Robinson model described in Brown and Finkelstein [4]. Treating death as a fourth health state ($h_t = 0$), we fit 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.

Using the estimated transition probabilities, we simulate demographic histories, beginning at age 70. Table 1 shows life expectancies.[7] 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 3.5 more years, while a female at the 90th percentile in good health expects to live 16.1 more years. Such dramatic differences in life expectancy should have large effects on saving.

[7]Consistent with our discrete time model, in our calculations we treat people that live n+1 periods as living n+1 years, rather than n+0.5 years.

Income Percentile	Nursing Home	Males Bad Health	Good Health	Nursing Home	Females Bad Health	Good Health	All†
10	3.53	5.86	7.22	6.02	9.99	11.97	10.38
30	3.61	6.40	8.14	6.41	10.92	13.04	11.36
50	3.77	7.05	9.11	6.85	11.94	14.18	12.36
70	3.98	7.81	10.10	7.38	12.96	15.20	13.40
90	4.26	8.61	11.01	8.05	13.97	16.14	14.31
By gender:‡							
Men							9.41
Women							13.54
By health status:◇							
Nursing Home						NA	
Bad Health						10.56	
Good Health						13.93	

Notes: Life expectancies calculated through simulations using estimated health transition and survivor functions; †Calculations for aggregate ("all") results use the gender and health distributions observed for entire population; ‡Calculations use the health and permanent income distributions observed for each gender; ◇Calculations use the gender and permanent income distributions observed for each health status group. The initial distribution contains no nursing home residents.

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 30.1% of women will ultimately enter a nursing home, as opposed to 17.9% for men. These numbers are lower than those reported in Brown and Finkelstein [4] which show 27% of 65-year-old men and 44% of 65-year-old women require nursing home care. One reason we find lower numbers is that the Robinson model is based on older data, and nursing home utilization has declined in recent years (Alecxih [1]).

	Males		Females		
Income Percentile	Bad Health	Good Health	Bad Health	Good Health	All†
10	15.9	17.6	26.6	28.8	26.0
30	15.8	17.8	27.3	29.6	26.4
50	15.7	18.1	27.8	30.6	27.1
70	16.1	19.0	29.0	32.0	27.9
90	16.4	18.8	30.2	33.2	29.4
By gender:‡					
Men					17.9
Women					30.1
By health status:◇					
Bad Health					25.4
Good Health					29.0

Notes: Entry probabilities calculated through simulations using estimated health transition and survivor functions; †Calculations for aggregate ("all") results use the gender and health distributions observed for entire population; ‡Calculations use the health and permanent income distributions observed for each gender; ◇Calculations use the gender and permanent income distributions observed for each health status group. The initial distribution contains no nursing home residents.

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

5.3 Co-pay rates

After being asked about out-of-pocket medical expenses, HRS respondents are asked to estimate their total billable medical expenses. Total medical expenses average $22,000 with a standard deviation of $75,000. We calculate the co-insurance rate, $q(\cdot)$, as the amount spent out-of-pocket (less insurance premia) divided by total billable medical expenses. We allow the co-pay rate to vary with gender, health status and age (Yogo [34] also allows it to depend on health and age).

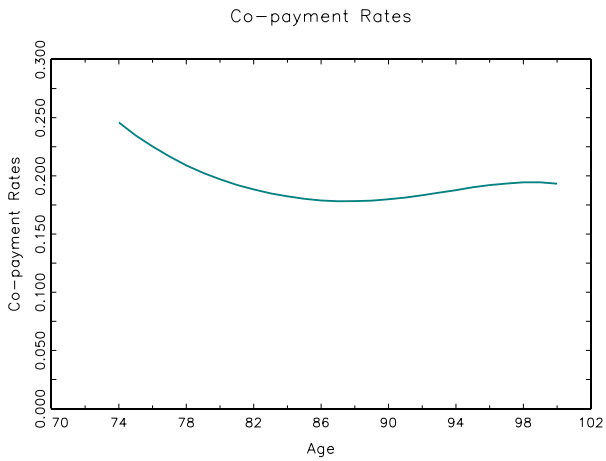

Figure 6: Co-pay rates by age.

Figure 6 presents average co-pays by age for people in our youngest cohort. Our estimated co-pays profiles are lower for people in bad health and display a non-linear pattern in age. On average, the co-pays display a large drop (from 25% to less than 18%) between ages 74 and 87, but then rise again to about 19% after age 87. This is likely capturing a shift in the composition of medical goods and services that people consume. The raise, in particular, could be due to a larger and larger fraction of people who, as they age and become more fragile, enter expensive nursing homes that they pay for out of pocket.

There are two key problems with inferring co-pay rates using out of pocket medical expenses and total billable medical expenses. First, the total billable medical expense information are largely imputed in AHEAD. Second, we wish to measure the share of total medical expenses not paid by Medicare

or private insurance: this includes both out of pocket expenses as well as Medicaid payments. The AHEAD data provide no information on Medicaid payments. In order to understand both of these problems we use data from the Medical Expense Panel Survey (MEPS). MEPS provides high quality information on total billable medical expenses as well as the payor of those expenses, including Medicaid. Estimates using MEPS are somewhat similar to results from the AHEAD: for example, the copay rate averages 29% in MEPS and does not vary much with demographics (although the co-pay rate is somewhat lower for men). Estimated total medical expenses are lower in MEPS than in the AHEAD, however. In the MEPS total billable medical expenses are $7,900 per year. Part of the reason for the discrepancy comes from nursing home expenses. An important limitation of the MEPS data is that the MEPS does not interview those in nursing homes. In the future we plan to better understand these discrepancies.

6 Second step results and model fit

At this point, we have not yet estimated the model. To illustrate some of the model's mechanics and point out some of the estimation issues, we perform some simulations using the endogenous medical expenditure model we developed in De Nardi et al. [10]. This model uses a simpler, 2-state health process, and assumes that everyone qualifies for Medicaid through the Medically needy pathway and applies whenever eligible.

6.1 Parameter values

Our estimate of β, the discount factor is 0.99.[8] The estimate of ν, the coefficient of relative risk aversion for "regular" consumption, is 2.15, while the estimate of ω, the coefficient of relative risk aversion for medical goods, is 3.19; the demand for medical goods is less elastic than the demand for consumption. The utility floor is the utility level that one gets when the medical needs shifter μ equals 1 and an individual consumes 202 dollars apiece of consumption and medical goods. The bequest motive is set to zero.

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

[8] Standard errors for these estimates can be found in De Nardi et al. [10].

ζ_t and ξ_t. The estimates for these parameters (available from the authors on request) imply that the demand for medical services rises rapidly with age.

We now turn to discussing how well the model fits the net worth and medical expense data and the model's implications for means-tested transfers and total medical expenditures.

6.2 Net worth profiles

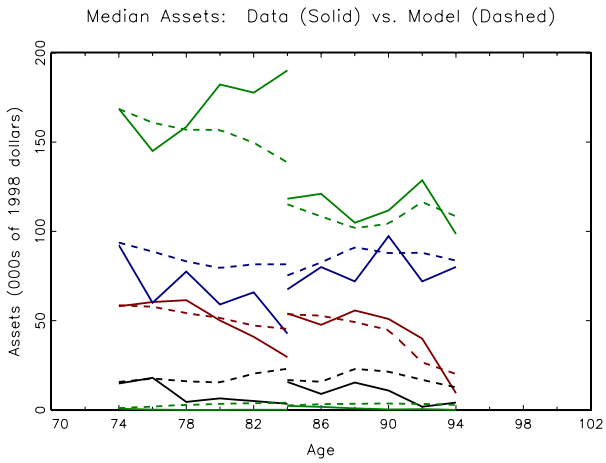

Figure 7: Median net worth by cohort and PI quintile: data (solid lines) and model (dashed lines).

Figure 7 compares the net worth profiles generated by the model (dashed line) and those in the data (solid line) for the members of two birth-year cohorts. 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 are for the cohort aged 82-86 in 1996. For the most part, the model replicates the main patterns found in the asset data: the most notable exception is that the model overstates asset holdings in the second-lowest permanent income quintile.

6.3 Medical expenses

Figure 8 displays average out-of-pocket medical expenses (that is, net of Medicaid payments and private and public insurance co-pays) paid by people

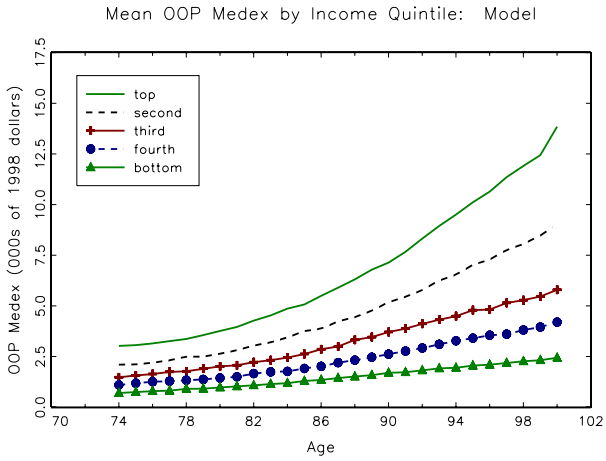

Figure 8: Average out-of-pocket medical expenses by age and permanent income.

in the model. Permanent income has a large effect on out-of-pocket medical expenses, especially at older ages. Average medical expenses are less than $3,000 a year at age 75 and vary little with income. By age 100, they rise to $2,400 for those in the bottom quintile of the income distribution and over $14,000 for those at the top of the income distribution.

Figure 9 compares the out-of-pocket medical expenses generated by the model to those found in the data. The version of the model used here underestimates out-of-pocket medical risk at very old ages. As an example, average out-of-pocket medical expenses for the oldest and richest people peak at over $30,000, while the model generates just $16,000. This discrepancy in part reflects the absence of a third, high co-pay, nursing home state in this earlier model.

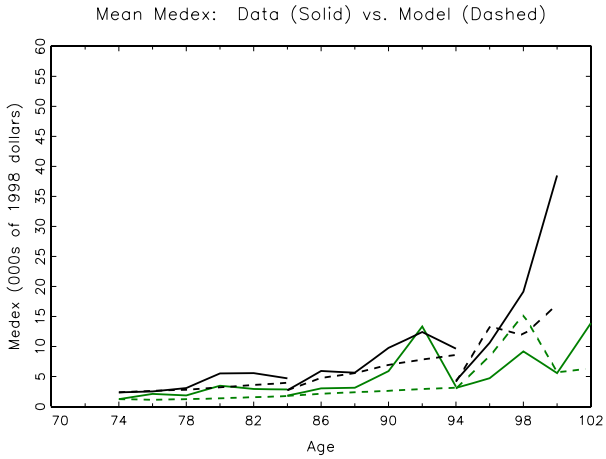

Figure 9: Mean out-of-pocket medical expenses: data (solid lines) and model (dashed lines).

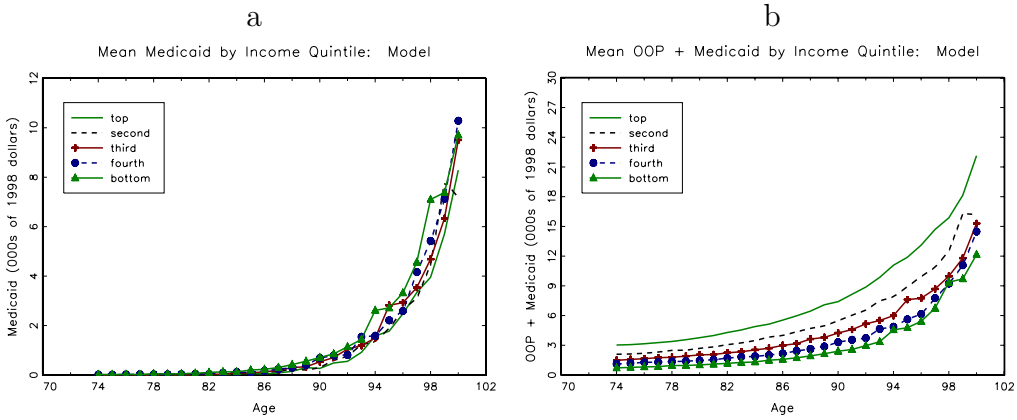

Figure 10: Average medical expenses by age and permanent income. Panel a: paid by Medicaid. Panel b: paid by Medicaid or out-of-pocket.

Panel a of Figure 10 displays the average medical expenses covered by our means-tested social insurance program, measured as the increase in $q_t m_t$ generated by government transfers. "Medicaid" payments rapidly increase with age, going from roughly zero at age 74 to nearly \$9,000 at age 100. Consistent with the redistributive nature of the program, these payments are quite close across people of different permanent incomes, but are higher for the poor. Panel b of Figure 10 shows the sum of medical expenses paid out-of-pocket and the expenses paid by Medicaid. Medicaid allows poorer people to consume proportionally much more medical goods and services than they pay for. As a result, the expense sum shown in panel b rises more slowly with income than the out-of-pocket expenditures shown in Figure 8. At age 100, people in the top permanent income quintile spend 470% more out-of-pocket than people in the bottom quintile. Once Medicaid is included, the difference narrows to 80%.

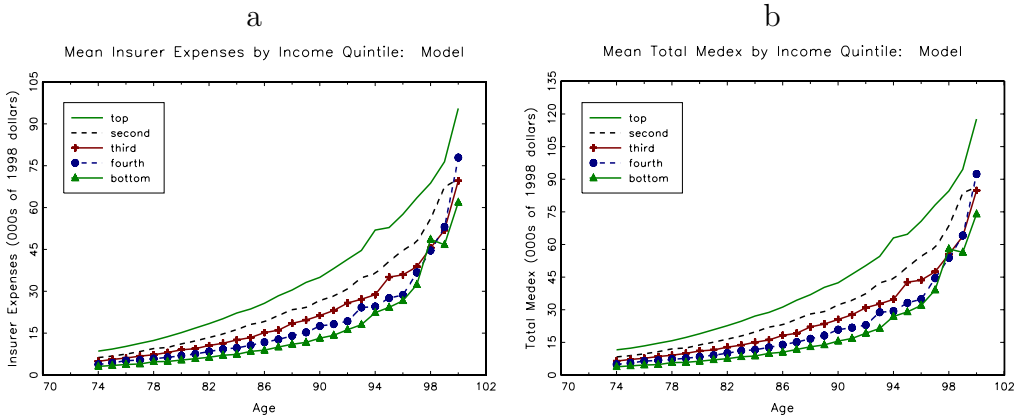

Figure 11: Average medical expenses by age and permanent income. Panel a: paid by insurers. Panel b: total.

Panel a of Figure 11 displays average medical expenses covered by private and public insurers. These payments are very large and also increase by age and permanent income, reaching over $90,000 for the oldest members of the top permanent income quintile. The oldest in the poorest permanent income quintile, however, also benefit from these payments, which reach over $60,000 at age 100. Panel b of Figure 11 displays total medical expenses, which in this case also coincides with total consumption of medical goods and services. Comparing the two panels makes it clear that most elderly individuals consume far more medical care than they for pay out-of-pocket. The increase in total medical expenses after retirement is very large, going from around $10,000 at age 74 to $60 to $100 thousand at age 99.

6.4 Utility floor, preference shocks, and implied insurance system

Through the interaction of the utility floor and medical needs shocks, the model has interesting implications on the insurance provided by means-tested programs. Our utility floor is based on the consumption of $202 in medical goods and $202 of non-medical goods with the medical preference shifter equalling 1. The interpretation of this number is not obvious, however, because people with higher medical needs receive larger transfers.

Figure 12 describes the transfers generated by the model. Panel a of this figure shows the fraction of individuals receiving transfers, while panel b

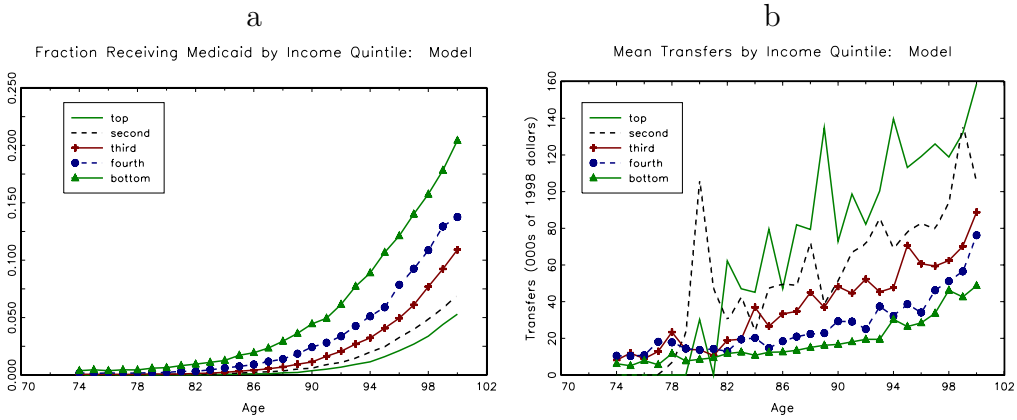

Figure 12: Means-tested transfers. Panel a: fraction receiving transfers. Panel b: average transfers per recipient.

shows transfers per recipient. The model generates lower Medicaid participation rates than found in the data, in part because it lacks a categorically needy pathway, which allows poor people with low medical medical expenses to receive Medicaid. On the other hand, the model captures the way in which Medicaid usage increases with age. Initially, very few people receive transfers, but as people age, and medical needs increase, more people become eligible. By age 100, over 10% of people receive transfers. The vast majority of the transfers are received by people with large medical needs and are thus spent on medical goods and services, rather than on non-medical consumption. Because people in the top permanent income quintile receive transfers only when their medical needs are extremely severe, very few of them receive transfers, but the average transfer is high. Even after age 95, only about 4% of this group receive transfers, with an average transfer in excess of $100,000. In contrast, after age 95, the average transfer in the bottom quintile is less than $50,000, but over 16% of this group receive transfers. Because the distribution of the medical needs shifter μ_t does not depend directly on income, the increased rate of recipiency found in the bottom income quintile means that the poor on average receive more transfers than the rich; see the discussion of Figure 10.

At any income level, however, the baseline parameterization of the consumption floor, coupled with our estimated medical needs shocks, can lead to very large transfers, which provide significant insurance against devastating medical illnesses.

7 Policy experiments

In the spirit of the recently debated health reforms, we study the effects of making public health insurance more generous. In the first experiment, we analyze an increase in the generosity of Medicaid by raising the utility floor by 50%. In the second experiment, we analyze an increase in the number of insured individuals by reducing co-payment rates by 25%.

7.1 A more generous means-tested program

In this policy experiment, we increase the generosity of means-tested insurance by increasing the level of the utility floor by 50%.[9] Figure 13 shows how transfers vary with the preference shifter μ_t for both the benchmark and the experiment with the higher utility floor. Figure 13 shows that in order to maintain a higher utility floor, transfers become much larger at all levels of medical need.

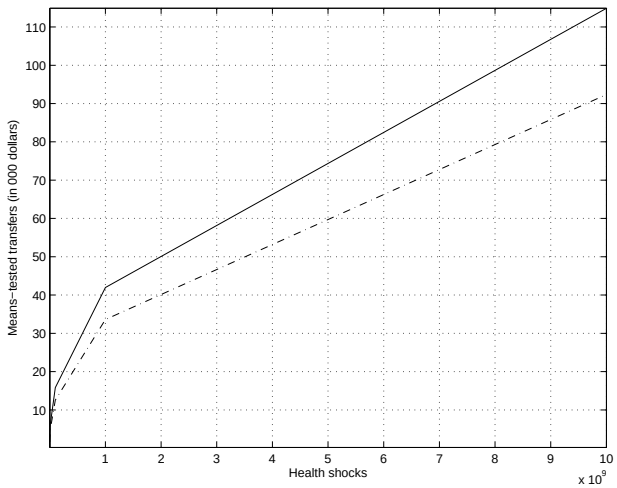

Figure 13: Means-tested transfers as a function of $\mu(\cdot)$. Dashed line: benchmark, solid line: experiment with more generous utility floor.

This increase in the insurance provided by the government leads people to save less for medical needs, and generates large reductions in net worth. Panel a of Figure 14 plots the net worth of the youngest cohort for the benchmark calibration (dashed line) and for the experiment with more generous

[9] More precisely, we increase the consumption equivalent $\underline{c}$ by 50%.

means-tested programs (solid line). Households deplete their assets more quickly in the specification with more generous insurance. The median asset holding of 95-year-old people in the highest permanent income quintile drops 22%, from $57,000 to $44,000; the median asset holding of 95-year-old people in the lowest permanent income quintile drops 51%, from $1,900 to $900. The declines are thus proportionally much larger for poorer people, who are the ones most likely to benefit to benefit from a means-tested transfer program. However, richer people also risk being wiped out by large medical expenses and thus benefit from the increased insurance provided by the higher consumption floor.

Raising the utility floor affects medical expenditures in several ways. The obvious direct effect of a higher floor is to raise the medical expenditures of individuals eligible for assistance. Moreover, raising the floor reduces the need to save, which will, holding assets fixed, lead individuals to increase their consumption of both medical and non-medical goods. Reduced saving, on the other hand, will lower medical expenditures in the future. Panel b of Figure 14 shows that a more generous utility floor increases total consumption of medical goods and services, especially after age 90 and for those in the bottom two permanent income quintiles.

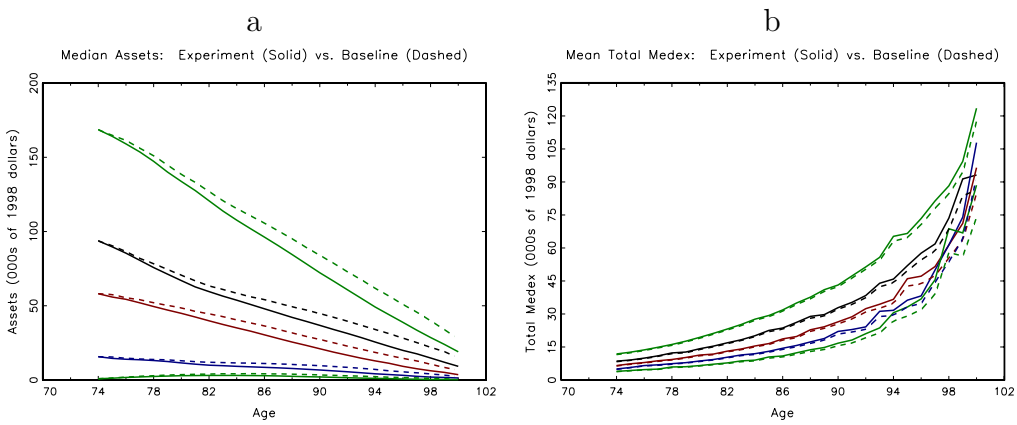

Figure 14: Net worth (panel a) and total medical expenses (panel b) by age and permanent income. Dashed line: benchmark, solid line: experiment with more generous utility floor.

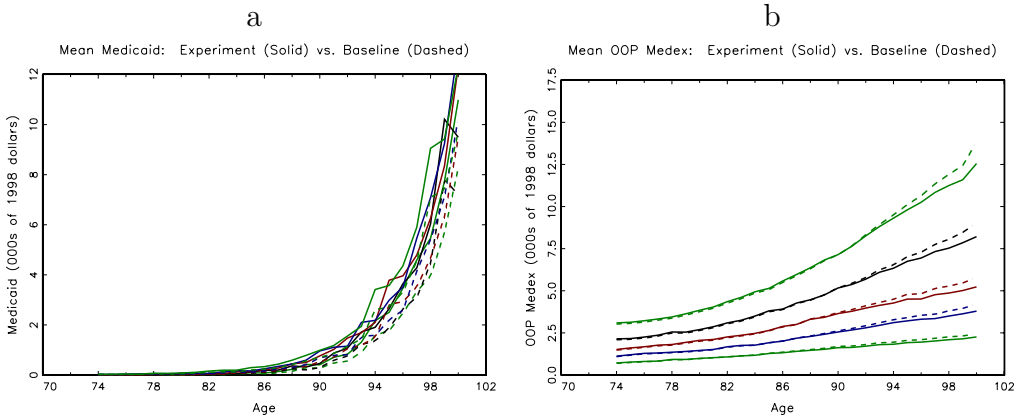

Figure 15: Medicaid payments (panel a) and out-of-pocket medical expenses (panel b) by age and permanent income. Dashed line: benchmark, solid line: experiment with more generous utility floor.

Panel a of Figure 15 shows that the higher utility floor, along with the resulting decrease in assets, increases Medicaid payments. Panel b shows that a more generous insurance system reduces out-of-pocket medical expenditures; the reduction in the consumers' cost share outweighs the increase in total medical expenditures.

7.2 A more generous co-insurance system

In this policy experiment we reduce the co-payment schedule by 25%. As in the previous experiment, the households react to the increased insurance by running down their assets more rapidly. This experiment, however, has smaller effects than the previous one (see Figure 16), and the largest effects occur at earlier ages. For example, the assets of 95-year-old people in the top permanent income group drop by 12%, while the assets in the bottom group change by 0.1%.

Panel b of Figure 16 shows that total medical expenses go up at all ages, especially for households in the highest permanent income quintiles. While the largest increases in absolute terms occur at the oldest ages, the increase at younger ages represent larger proportions.

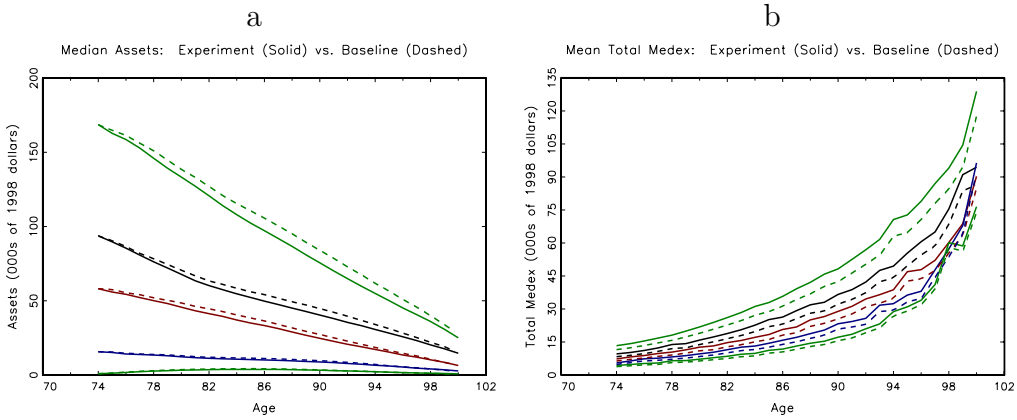

Figure 16: Net worth (panel a) and total medical expenses (panel b). Dashed: benchmark, solid: more generous co-insurance system.

Figure 17 shows that reducing the co-pay rates reduces out-of-pocket medical costs at all ages, especially for those with higher permanent income. Because individuals in the lower income quintiles rely more heavily on Medicaid, they are less likely to incur co-pays. As a result, their medical care decisions are less sensitive to changes in co-payment rates. While lower co-pay rates increase total medical expenditures, this increase in quantities is more than offset by the reduction in the consumers' out-of-pocket shares. Medicaid payments go down for similar reasons. The increase in total medical expenses is thus borne entirely by insurers, as shown in Figure 18. Given that Medicare is by far the principal insurer for retirees, we see an important interaction between the Medicaid and Medicare programs: increases in the generosity of Medicare reduce Medicaid payments.

Our finding that a decrease in the out-of-pocket price of medical expenditures leads to a reduction in out-of-pocket expenditures indicates that the elasticity for medical goods is fairly small. In a recent study, Fonseca et al. [16] 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 widely 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 -0.46.

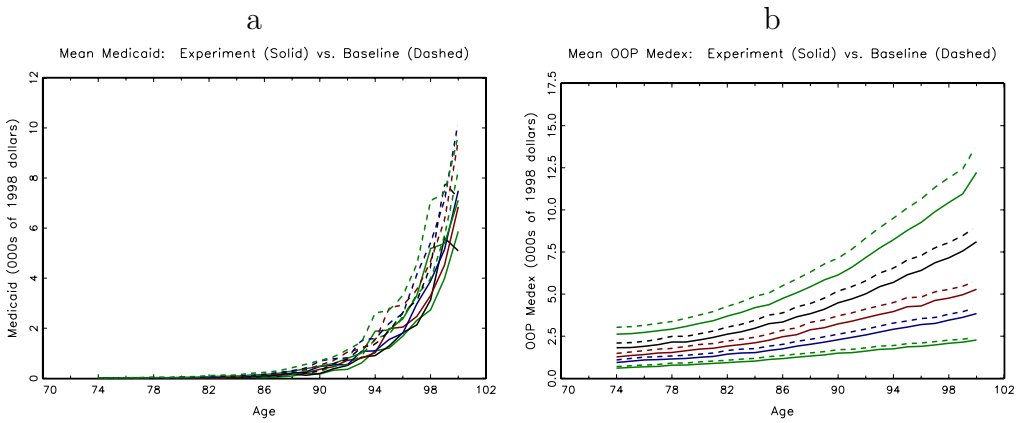

Figure 17: Medicaid payments (panel a) and out-of-pocket medical expenses (panel b) by age and permanent income. Dashed line: benchmark, solid line: experiment with more generous co-insurance system.

8 Conclusions and extensions

To come ...

Figure 18: Medical insurer payments by age and permanent income. Dashed line: benchmark, solid line: experiment with more generous co-insurance system.

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] 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.

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

[5] 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.

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

[7] David Card, Carlos Dobkin, and Nicole Maestas. Does medicare save lives? NBER Working Paper 13668, 2007.

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

[9] Joseph J. Jr. Doyle. Returns to local-area healthcare spending: Using health shocks to patients far from home. Mimeo, 2010.

[10] 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.

[11] 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.

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

[13] Zhigang Feng. Macroeconomic consequences of alternative reforms to the health insurance system in the u.s. Mimeo, 2009.

[14] Amy Finkelstein and Robin McKnight. What did medicare do (and was it worth it)? NBER Working Paper 11609, 2005.

[15] 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.

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

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

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

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

[20] Ahmed Khwaja. Estimating willingness to pay for medicare using a dynamic life-cycle model of demand for health insurance. *Journal of Econometrics, forthcoming*, 2009.

[21] Ruud H. Koning. Kernel: A gauss library for kernel estimation, 1996. http://www.xs4all.nl/

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

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

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

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

[26] 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.

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

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

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

[30] 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.

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

[32] 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? Working Paper 10862, National Bureau of Economic Research, 2004.

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

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

Appendix A: Moment conditions and asymptotic distribution of parameter estimates

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, $\underline{c}$, 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 $t \in \{t_0, ..., t_T\} = \{1996, 1998, 2000, 2002, 2004, 2006\}$, we match median assets for $Q_A = 5$ permanent income quintiles in $P = 5$ birth year cohorts.[10] 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 [23], Powell [31] and Buchinsky [6]). In particular, applying the indicator function produces

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

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

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

(e.g., Chamberlain [8]):

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

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.

As with assets, we divide individuals into 5 cohorts and match data from 5 waves covering the period 1998-2006. 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 the 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:

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

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:

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

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

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

and define the standard deviation σ_{pqt} as

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

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

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

The corresponding moment condition for the second autocorrelation is

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

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 $J = 478$ 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 [29] and Duffie and Singleton [12], 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;

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

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 [25] 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

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

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].) To check for robustness, we also use a "diagonal" weighting matrix, as suggested by Pischke [30]. 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 [29], Newey and McFadden [26] (section 7), and Powell [31].

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

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

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

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

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 [21]. The gradients for equations (21) and (22) are found in a similar fashion.

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

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

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