

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 Joe Altonji, Norma Coe, Courtney Coile, Amy Finkelstein, Victor Rios-Rull, Karl Scholz, Sam Schulhofer-Wohl, Jon Skinner, and Gianluca Violante. For excellent research assistance we thank David Benson, Taylor Kelley, Jeremy McCauley, An Qi, and Shani Schechter. De Nardi gratefully acknowledges support from the ERC, grant 614328 “Savings and Risks” and from the ESRC through the Centre for Macroeconomics. Jones gratefully acknowledges financial support from the Social Security Administration through the Michigan Retirement Research Center (MRRC grants UM10-16 and UM12-12). Jones also gratefully acknowledges the hospitality of the Federal Reserve Bank of Chicago and UCL. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research, any agency of the federal government, the Federal Reserve Bank of Chicago, the Social Security Administration, or the MRRC.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Mariacristina De Nardi, Eric French, and John Bailey Jones. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Medicaid Insurance in Old Age
Mariacristina De Nardi, Eric French, and John Bailey Jones
NBER Working Paper No. 19151
June 2013, Revised December 2015
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 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 currently single retirees.

We find that retirees with high lifetime incomes can end up on Medicaid and often value Medicaid insurance the most, as they face a larger risk of catastrophic medical needs at old ages and face the greatest consumption risk. Compensating variation calculations indicate that current retirees value Medicaid insurance at more than its actuarial cost, but that most would value an expansion of the current Medicaid program at less than its cost. These findings suggest that for current single retirees, the Medicaid program may be of the approximately right size.

Mariacristina De Nardi
Federal Reserve Bank of Chicago
230 South LaSalle St.
Chicago, IL 60604
and University College London
and Institute For Fiscal Studies - IFS
and also NBER
denardim@nber.org

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

Eric French
University College London and IFS
30 Gordon Street
London, WC1H 0AX
United Kingdom
Eric.French.econ@gmail.com

1 Introduction

Many elderly households benefit from Medicaid, a means-tested, public health insurance program. Despite the increasing importance of Medicaid in the presence of an aging population and rising medical costs, very little is known about how Medicaid payments are distributed among the elderly and how the elderly value these payments. Which elderly households receive Medicaid transfers? How redistributive are these transfers and the taxes needed to finance them? 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.

We focus on single retirees, who comprise about 50% of age 70+ people and 70% of age 70+ households. We 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 reciprocity rate in the bottom income quintile stays around 60%-70% throughout retirement. In contrast, the reciprocity rate of higher-income retirees is initially very low, but increases by age, reaching 20% by age 95. Moreover, data from the Medicare Current Beneficiary Survey (MCBS) shows that high income individuals, conditional on receiving Medicaid transfers, receive larger transfer than low income individuals.

The data show who ends up on Medicaid, how much they receive from Medicaid, how much wealth they hold and how much they spend on medical goods and services. However, to assess how much retirees value the insurance provided by Medicaid and to perform counterfactuals, we need a model. We thus develop and estimate a life-cycle model of consumption and endogenous medical expenditure that accounts for Medicare, Supplemental Social Insurance (SSI), and Medicaid. Agents in the model 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 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, one for the lifelong poor, and one for people impoverished by large medical expenses.

To appropriately evaluate Medicaid redistribution, we allow for heterogeneity in wealth, permanent income, 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. We estimate the model by matching life-cycle profiles of assets, out-of-pocket medical spending, and Medicaid reciprocity rates for elderly singles across different cohorts and permanent income groups. Matching Medicaid reciprocity introduces an unexpected angle in the identification of bequest motives. To match Medicaid reciprocity rates, Medicaid must be fairly generous, which in turn reduces medical expense risk. To reconcile high observed asset holdings with reduced medical expense risk, a bequest motive is necessary.

Our model matches its targets 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. [47] 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 the corresponding data from the MCBS survey.

Our model shows 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.

We also find that, with moderate risk aversion and realistic lifetime and medical needs risk, the value most retirees place on Medicaid insurance exceeds the actuarial value of their expected payments. For example, if we decrease the discounted present value of Medicaid payments by a dollar, to maintain the same level of utility a retired person at the bottom of the income distribution would have to be compensated by more than a dollar, and a person at the top of the distribution would have to be compensated by more than three dollars. On the other hand, we find that a Medicaid expansion would be valued by most retirees at less than its cost. These comparisons of the transfers' actuarial values to the recipient's valuations suggests that the current Medicaid program for most currently single retirees is about the right size.

Our calculations also show that it is the richer retirees who value Medicaid most highly and thus might be most in favor of a Medicaid expansion. However, this comparison does not take into account the subsidization implied by Medicaid taxes.

Using data from the Panel Survey of Income Dynamics (PSID), we estimate the distribution of Medicaid-related taxes. Our PSID computations indicate if we decrease the discounted present value of Medicaid payments by a dollar, a retired person at the bottom of the income distribution would save 0.2 dollars in taxes, and a person at the top of the distribution would save nearly five dollars. Therefore under the current tax system the rich would not support an expansion of Medicaid insurance, because the increase in their Medicaid tax burden would exceed the increase in their Medicaid valuation.

The paper thus contributes to the literature in multiple ways. First, it evaluates how Medicaid redistributes across people in a model with rich heterogeneity. Second, it uses the model to compute retirees' valuation of Medicaid insurance in a framework that matches the data well and explicitly models the response of savings and medical expenditures to the Medicaid rules. Finally, it provides additional identification of the bequest motive by carefully modeling risks and insurance and by matching Medicaid reciprocity and payment rates.

2 Literature review

This paper is related to previous work on savings, health risks, and social insurance. Kotlikoff [44] stresses the importance of modeling health expenditures when studying precautionary savings, but Hubbard et al. [35] and Palumbo [59] solve dynamic programming models of saving under medical expense risk and find that medical expenses have relatively small effects. However, Hubbard et al. [35] and Palumbo [59] likely underestimated medical spending risk, 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.

Using newer and more comprehensive data, De Nardi et al. [19] and Marshall, McGarry, and Skinner [49] find that late-in-life medical expenses are large and generate powerful savings incentives. Furthermore, Poterba, Venti, and Wise [62] show that those in poor health have considerably lower assets than similar individuals in good health. Lockwood [46], Nakajima and Telyukova [52], and Yogo [69] add to the literature by estimating life cycle models that include additional insurance choices, housing, and portfolio choices respectively. Laitner et al. [45] derive analytic expressions providing intuition for how uncertain longevity and medical expense risk affect savings decisions.

In this paper, we extend the endogenous medical spending model of De Nardi et al. [19] to measure the distribution of Medicaid transfers, the taxes used to fund the transfers, and the valuations retirees place on them. Consistent with the institutions, we explicitly model two separate ways to qualify for Medicaid: having low income and assets (the “categorically needy” pathway, which incorporates SSI) or becoming impoverished by high medical needs (the “medically needy” pathway). People at different points of the income distribution qualify for Medicaid benefits in different ways and thus receive different insurance. Because nearly two-thirds of Medicaid payments to the elderly are to those in nursing homes, we model the nursing home state explicitly. We expand our set of econometric targets to include Medicaid eligibility rates, adding an important new source of identification. We also compare the Medicaid payments predicted by the model to those observed in the MCBS. We show that our model matches Medicaid payment flows well, although they are not matched by construction.

Earlier studies of Medicaid include Hubbard et al. [36] and Scholz et al. [66], who 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 [31] exploit cross-state variation in Medicaid rules and find Medicaid has significant effects on savings. Brown and Finkelstein [10] develop a dynamic model of optimal savings and long-term care purchase decisions and 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. [12] exploit cross-state variation in Medicaid rules and also find significant crowding out. We also find that Medicaid encourages spending and reduces savings.

Several new papers study the importance of medical expense risk in general equilibrium, including Hansen et al. [33], Paschenko and Porapakarm [60], İmrohoroğlu and Kitao [39]. Kopecky and Koreshkova [43] 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. [7] use a model with medical expense risk to assess the incentive and welfare effects of Social Security and means-tested social insurance programs like Medicaid. They too find that Medicaid provides the elderly with valuable insurance. Compared to these papers, we focus more on valuations and redistribution at the individual level and include much more heterogeneity. We allow demographic transitions to depend on lifetime earnings, consistent with Hurd [37] and Hurd, McFadden, and Merrill. [38], who highlight the importance of accounting for the link between wealth and mortality in life-cycle models. We estimate our model

against life-cycle profiles, rather than calibrating it. Most important, in our model people can adjust medical spending – as well as consumption and savings – allowing the quality of care to vary.

Several 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. [28] and Scholz and Seshadri [65] assume that the consumption floor is invariant to medical needs, whereas our specification allows for more realistic links between medical needs and Medicaid transfers. Ozkan [57] 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 expected Medicaid transfers.¹ We also estimate the distribution of the taxes used to finance these transfers. 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. Adapting the approach of McClellan and Skinner [50], we assume that the Medicaid tax burden is proportional to the general tax burden.

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”:

¹Using a simpler, calibrated model, Brown and Finkelstein [10] analyze how Medicaid affects the valuation of long-term care insurance. Braun et al. [7] calculate the aggregate welfare effects of eliminating means-tested social insurance.

²De Nardi et al. [20] and Gardner and Gilleskie [31] document many important aspects of Medicaid insurance in old age.

Medicaid contributes only after Medicare and private insurance pay their share and 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 [56]) 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 were for nursing home payments (Kaiser Foundation [29]).

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 their medical conditions.

4 Some data

We use two main data sets, the AHEAD and the MCBS. We begin this section with an overview of each dataset.

4.1 The AHEAD dataset

The Assets and Health Dynamics of the Oldest Old (AHEAD) dataset 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 3,727 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. We drop 229 individuals

who were partnered with another individual at some point during the sample period or who did not remain single until death, and 252 individuals with labor income over \$3,000 at some point during the sample period. We are left with with 3,246 individuals, of whom 588 are men and 2,658 are women. Of these 3,246 individuals, 370 are still alive in 2010. We do not use 1994 assets or medical expenses. Assets in 1994 were underreported (Rohwedder et al. [64]) 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 on average 60 days per year) 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, each of which contains people born within a 5-year window. 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. Throughout, we will refer to each of these 5-year birth cohorts as a cohort.

Since we want to understand the role of income, we further stratify the data by post-retirement permanent income (PI). We measure PI as the individual’s average non-asset income over all periods during which he or she is observed. Non-asset income includes 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 non-asset income variables that we use, our measure of PI is also a good measure of lifetime permanent income.

4.2 The MCBS dataset

An important limitation of the AHEAD data is that it lacks information on other payers 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 1996-2010 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 paid 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, it gives extremely accurate data on Medicare payments and fairly accurate data on out-of-pocket and Medicaid payments. Both the AHEAD and the MCBS survey include information on those who enter a nursing home or die. This is an important advantage compared to the Medical Expenditure Panel Survey (MEPS), which does not capture late-life or nursing home expenses.

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 that we use for the AHEAD data. Specifically, we drop those who were observed to be married over the sample period, work, or be younger than 72 in 1996, 74 in 1998, etc. These sample selection procedures leave us 17,103 different individuals who contribute 40,157 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. In the MCBS, individuals are asked about total income, not annuitized income. Fortunately, we found that this variable lines up well with total income in the AHEAD. Furthermore, in the AHEAD, the correlation between total income and annuitized income is 0.8. Consistent with our computations in the AHEAD, we use average total income, over the time that we observe an individual, as our measure of permanent income (PI) in the MCBS.

4.3 Medicaid reciprocity and payments

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 PI quintile.

The approach we use to stratify the data behind Figure 1 is one we will use repeatedly throughout the paper. Recall that we stratify the data by PI quintile and

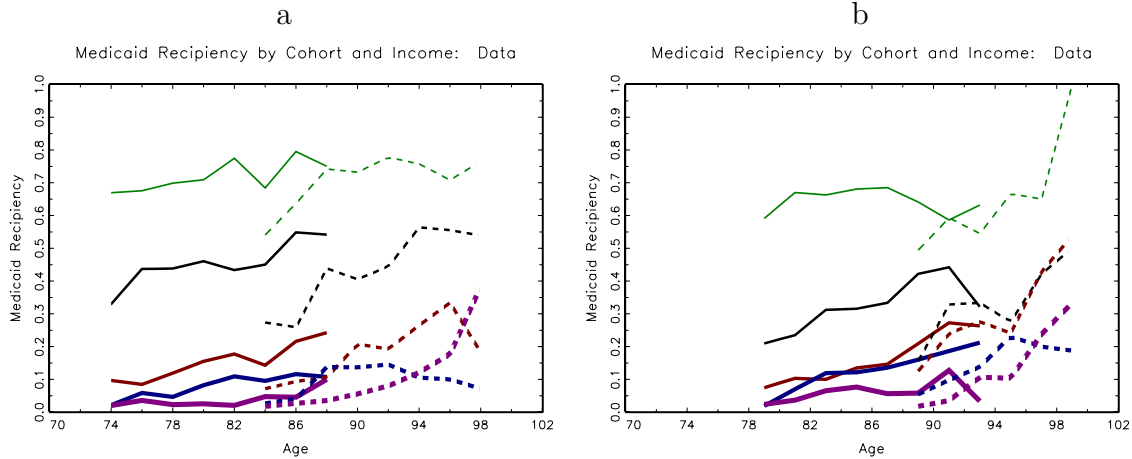


Figure 1: Each line represents Medicaid reciprocity rates for a cohort-income cell, traced over the time period 1996-2010. Thicker lines refer to higher permanent income groups. Panel a: cohorts aged 74 and 84 in 1996. Panel b: cohorts aged 79 and 89 in 1996.

cohort. For each cohort-quintile cell, we calculate the Medicaid reciprocity rate in each calendar year. We then construct life-cycle profiles by ordering the reciprocity rates 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. (We omit the profiles for the oldest cohort because the sample sizes are tiny.) For each cohort in the figure there are five horizontal lines, one for each PI quintile. To indicate PI rank, we vary the thickness of the lines on our graphs: thicker lines represent observations for higher-ranked PI groupings.

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 other year until 2010. The other cohorts start from older initial ages and are also followed for fourteen years. The graphs report the Medicaid reciprocity rate for each cohort and PI grouping at eight dates over time. At each sample date, we calculate the Medicaid reciprocity rate for individuals alive at that date — we use an unbalanced panel. Cohort-income-year cells with fewer than 10 observations are dropped.

Unsurprisingly, Medicaid reciprocity is inversely related to PI: the thin top line shows the fraction of Medicaid recipients in the bottom 20% of the PI distribution, while the thick bottom line shows median assets in the top 20%. 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 reciprocity rate tends to rise with age most quickly for people in the middle and highest PI groups. For example, in the oldest cohort and top two PI quintiles the fraction of people receiving Medicaid 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.

Permanent Income Quintile	Average Benefit	Reciprocity Rate	Average Benefit per Recipient
Bottom	9,080	.70	12,990
Fourth	5,720	.42	13,690
Third	2,850	.16	18,350
Second	1,950	.08	24,360
Top	1,280	.05	23,790

Table 1: Average Medicaid benefits, reciprocity, and benefits per recipient, MCBS.

Table 1 shows average Medicaid benefits, the reciprocity rate, and benefits per recipient in the MCBS data, conditional on PI quintile. Average payments decline with PI. However, this is because reciprocity rates also decline by PI. In fact, the payments received by each Medicaid recipient increases with PI, from \$12,990 at the bottom quintile to \$23,790 at the top.

4.4 Medical expense profiles

In all survey 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 private and Medicare part B insurance premia,

drug costs, and costs for hospital, nursing home care,³ doctor visits, dental visits, and outpatient care. It does not include expenses covered by insurance, either public or private. The AHEAD’s expenditure measure is retrospective, as it measures spending over the previous two years. We annualize the data by dividing spending over the last two years by two. It includes medical expenses during the last year of life, collected through interviews with the deceased’s children or other survivors.

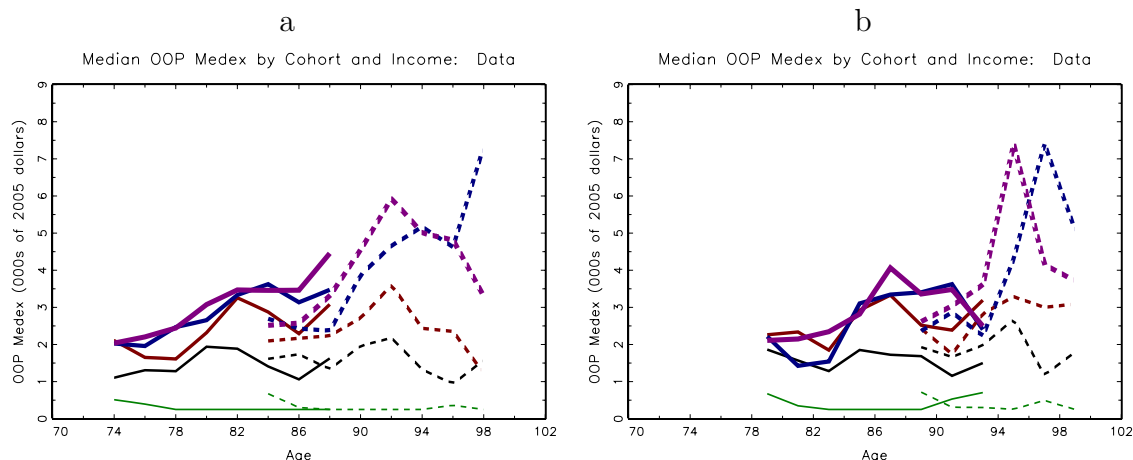


Figure 2: Each line represents median out of pocket medical expenditures for a cohort-income cell, traced over the time period 1996-2010. Thicker lines refer to higher permanent income groups. Panel a: cohorts aged 74 and 84 in 1996. Panel b: cohorts aged 79 and 89 in 1996.

French and Jones [30] show that the medical expense data in the AHEAD line up with the aggregate statistics. For our sample, mean out-of-pocket 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

³Nursing home costs include a food and shelter component, besides medical costs, thus raising the question of whether the food and shelter components should be eliminated from the nursing home costs to avoid double counting these items. There are two reasons why this is not as important as one might expect. First, the food and shelter component of nursing home costs make up for a small share of total nursing home costs. In fact, when we eliminate the food and shelter component of nursing home costs, our medical expense profiles do not change much. Second, many retirees in nursing homes keep their houses (whether owned or rented), expecting to go back to them. Hence, they are paying for two dwellings and it would be wrong to remove the shelter component of nursing homes from for these people. Finally, it should be noted that the shelter component is larger than the food component for most single retirees. For these reasons we believe that our approach most closely approximates reality.

reimbursed nursing home and hospital nights.

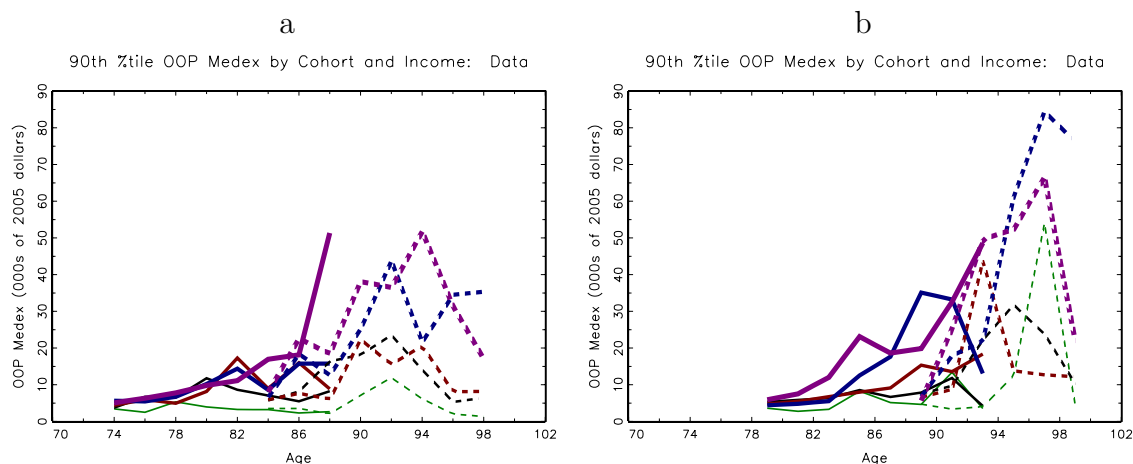


Figure 3: Each line represents the 90th percentile of out-of-pocket medical expenditures for a cohort-income cell, traced over the time period 1996-2010. Thicker lines refer to higher permanent income groups. Panel a: cohorts aged 74 and 84 in 1996. Panel b: cohorts aged 79 and 89 in 1996.

Figures 2 and 3 display the median and 90th percentile of the out-of-pocket medical expense distribution, respectively. 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. Protected by Medicaid, individuals in the bottom income quintiles pay less out-of-pocket.

4.5 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 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 PI, 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

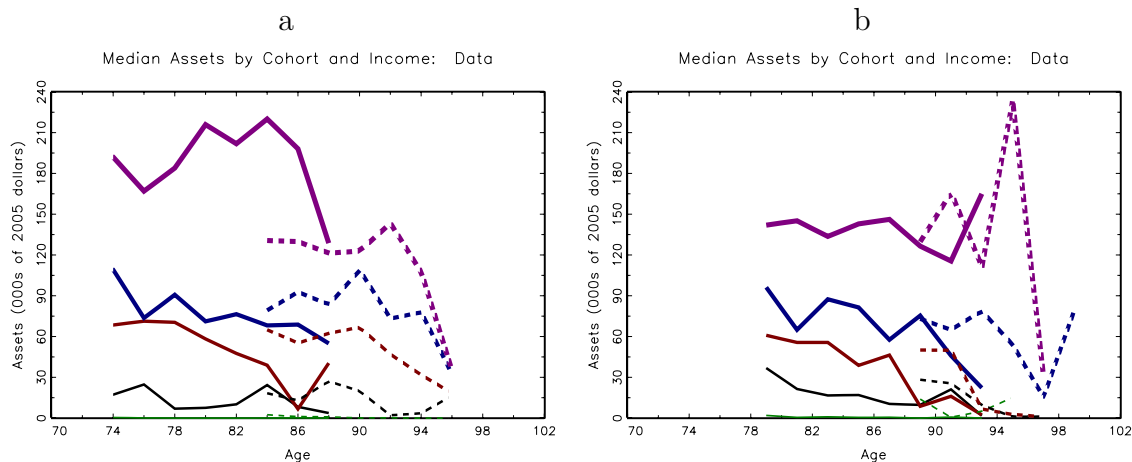


Figure 4: Each line represents median assets for a cohort-income cell, traced over the time period 1996-2010. Thicker lines refer to higher permanent income groups. Panel a: cohorts aged 74 and 84 in 1996. Panel b: cohorts aged 79 and 89 in 1996.

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, \$210,000 at age 82, \$220,000 at age 84, \$200,000 at age 86, and \$130,000 at age 88.⁴

For all PI quintiles in these cohorts, the assets of surviving individuals do not decline rapidly with age. Those with high PI do not run down their assets until their late 80s, although those with low PI 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 [51]). However, as De Nardi, French, and Jones [19] show, the risk of medical spending rising with age and income goes a long way toward explaining this puzzle.

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

⁴The 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 PI 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.

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

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

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.⁵ One way to interpret the medical spending in the utility function formulation is that medical spending improves within-period health. This is a simple way to capture endogenous medical spending, and is similar to other specifications used in the literature (Einav et al. [23], McClellan and Skinner [50], Bajari et al. [3]).

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 [32], in which medical expenses represent investments in health capital, which in turn decreases mortality (e.g., Yogo [69]) or improves health. While a few studies find that medical expenditures have significant effects on health and/or survival (Card et al. [14]; Doyle [17], Chay et al. [16]), most studies find small effects (Brook et al. [8]; Fisher et al. [27]; and Finkelstein and McKnight [25]). Interestingly, Finkelstein et al. [26] find that access to Medicaid increases medical total medical spending, but do not find that Medicaid reduces mortality for the under 65 population. Instead, they find that access to Medicaid reduces depression, which is consistent with our model that allows added health care to improve utility, but not longevity. These findings confirm that the effects of medical

⁵We assume that preferences are separable between medical and non-medical goods, which restricts the set of possible and income elasticities. The parameters of our current specification are identified largely by income elasticities, by matching the way in which out-of-pocket medical spending rises with income at multiple ages. However, our specification also generates reasonable price elasticities. Given that a simpler specification matches the facts well, we decided to not estimate a more complex non-separable specification, where identification would be less transparent.

expenditures on the health outcomes are 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). To get around these problems, Khwaja [40] estimates a structural model in which medical expenditures both improve health and provide utility. He finds (page 143) that medical utilization would only decline by less than 20% over the life cycle if medical care was purely mitigative and had no curative or preventive components. Blau and Gilleskie [6] also estimate a structural model and reach similar conclusions.

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 are 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

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

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).

In the standard consumption-savings model with exogenous medical spending (e.g., Hubbard et al. [36]), means-tested social insurance is typically modeled as a government-provided consumption floor. In that framework a consumption floor is equivalent to a utility floor, as a lower bound on consumption provides a lower bound on the utility that an individual can achieve. Our utility floor formulation is thus a straightforward generalization of means-tested insurance from the workhorse model, generalized to the case in which people choose their medical expenditures.

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

$$\log(\mu(\cdot)) = \alpha_0 + \alpha_1 t + \alpha_2 t^2 + \alpha_3 t^3 + \alpha_4 h_t + \alpha_5 h_t \times t \quad (3)$$

$$+ \sigma(h, t) \times \psi_t, \quad (4)$$

$$\sigma(h, t)^2 = \beta_0 + \beta_1 t + \beta_2 t^2 + \beta_4 h_t + \beta_5 h_t \times t, \quad (5)$$

$$\psi_t = \zeta_t + \xi_t, \quad \xi_t \sim N(0, \sigma_\xi^2), \quad (6)$$

$$\zeta_t = \rho_m \zeta_{t-1} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma_\epsilon^2), \quad (7)$$

$$\sigma_\xi^2 + \frac{\sigma_\epsilon^2}{1 - \rho_m^2} \equiv 1, \quad (8)$$

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 do not allow any component of $\mu(\cdot)$ to depend on PI, which 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

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

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 .

Since non-asset post-retirement income y_t , is mainly composed of social security and defined benefit pension income, it is not subject to shocks. For example, we found that negative health shocks have little effect on income changes in our AHEAD data. Thus, we model it a deterministic function of sex, permanent income, and age:

$$y_t = y(g, I, t). \quad (10)$$

5.4 The individual's problem

Consider a single person seeking to maximize his or her expected lifetime utility at age t , $t = t_{r+1}, \dots, T$, 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:

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

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

$$b_c(a_t, y_t, \mu(\cdot)) = \underline{Y} - \max\{y_t + ra_t - y_d, 0\} + \max\{-\underline{Y}, 0\}, \quad (12)$$

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

$$b_m(a_t, y_t, \mu(\cdot)) = \max\{x_m(\cdot) - (\max\{y_t + ra_t - y_d, 0\} + \max\{a_t - A_d, 0\}), 0\}, \quad (13)$$

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_{Mt} to denote this choice, with $I_{Mt} = 1$ if the person applies for Medicaid and $I_{Mt} = 0$ 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

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

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

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

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

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

$$a_{t+1} \leq \min\{A_d, a_t\}, \quad (17)$$

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 $x_i(\cdot)$, 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:

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

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

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

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 ($\nu, \omega, \beta, \underline{u}_c, \underline{u}_m$, 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 PI 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 4.⁶

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. Additional detail on our computational approach can be found in Appendix B.

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 C 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. [19]. The

⁶As 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.

procedures for estimating demographic transition probabilities and co-pay rates are new.

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 PI 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 PI. 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 D gives details on the procedure.

Permanent Income Percentile	<u>Males</u>			<u>Females</u>			All [†]
	Nursing Home	Bad Health	Good Health	Nursing Home	Bad Health	Good Health	
10	1.65	6.02	7.51	2.48	10.01	12.01	10.44
30	1.67	6.63	8.47	2.60	10.98	13.15	11.49
50	1.69	7.32	9.47	2.73	11.99	14.26	12.53
70	1.72	8.04	10.42	2.86	13.02	15.26	13.52
90	1.75	8.81	11.31	3.00	13.94	16.15	14.39
By gender: [‡]							
Men							9.71
Women							13.55
By health status: [◊]							
Bad Health							10.69
Good Health							13.99

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

Using the estimated transition probabilities, we simulate demographic histories, beginning at age 70, for different gender-PI-health combinations. Table 2 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 PI percentile in a nursing home expects to live only 1.65 more years, while a female at the 90th percentile in good health expects to live 16.15 more years.⁷

⁷Our predicted life expectancy at age 70 is about three years less than what the aggregate statistics imply. This discrepancy stems from using data on singles only: when we re-estimate the model for both couples and singles, predicted life expectancy is within a year of the aggregate statistics for both men and women. In addition, our estimated income gradient is similar to that in Waldron [68], who finds that those in the top of the income distribution live 3 years longer than

Another important driver of saving is the risk of needing nursing home care. Table 3 shows the probability at age 70 of ever entering a nursing home. The calculations show that 46.1% of women will ultimately enter a nursing home, as opposed to 30.6% for men. These numbers are similar to those from the Robinson model described in Brown and Finkelstein [9], which show 27% of 65-year-old men and 44% of 65-year-old women require nursing home care.

Permanent Income Percentile	<u>Males</u>		<u>Females</u>		All [†]
	Bad Health	Good Health	Bad Health	Good Health	
10	26.4	30.1	41.2	45.2	40.7
30	26.9	31.2	42.5	46.8	42.2
50	27.2	32.0	43.6	47.9	43.3
70	27.2	32.5	44.1	48.8	43.9
90	27.2	32.4	44.4	49.0	43.9
By gender: [‡]					
Men					30.6
Women					46.1
By health status: [◊]					
Bad Health					39.9
Good Health					45.0

Notes: Percentages 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 3: Percentage of people ever entering a nursing home, conditional on being alive at age 70.

those at the bottom, conditional on being 65.

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: $q_t = q(h_t)$.

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 27% and does not vary much with demographics. The co-pay rate for those in nursing homes is 68%. For every dollar spent on nursing homes, 34 cents come from Medicaid and 34 cents are from out-of-pocket, with 32 cents coming from Medicare or other sources. We cross-checked these co-pay rates with data from the 1997-2008 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

Table 4 presents our estimated parameters. Our estimate of β , the discount factor, is 0.994, which suggests a high level of patience. However, in our model individuals discount the future not only because of impatience, but also because they might not survive to the next period. The effective discount factor is the product $\beta s_{g,h,I,t}$. As Table 2 shows, the survival probability for our sample of older individuals is low, implying an effective discount factor much lower than β .

Our estimate of ν , the coefficient of relative risk aversion for “regular” consumption, is 2.8, while our estimate of ω , the coefficient of relative risk aversion for medical goods, is 3.0. Bajari et al. [3] estimate the same utility function in a static model of health insurance choice and medical care utilization. They estimate $\nu = 1.9$ and $\omega = 3.2$. Thus, they also find $\nu < \omega$. However, their estimated value for ν is lower than ours. Because we allow for self-insurance through savings, for any given set of parameters, demand for health insurance will be lower. Thus we need a higher coefficient of consumption risk aversion to explain health insurance and medical spending choices. Einav et al. [23] and McClellan and Skinner [50] also study two period

problems where utility depends on medical care.

Our estimates imply that the demand for medical goods is less elastic than the demand for consumption. In a recent study, Fonseca et al. [28] 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 -0.25.

β : discount factor	.994 (0.013)
ν : RRA, consumption	2.825 (0.025)
ω : RRA, medical expenditures	2.986 (0.029)
\underline{Y} : SSI income level	\$6,670 (207)
$u_c = u_m$: utility floor [†]	\$4,600 (144)
θ : bequest intensity	39.71 (2.51)
k : bequest curvature (in 000s)	13.0 (0.650)

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

Table 4: Estimated preference parameters. Standard errors are in parentheses below estimated parameters.

The SSI income benefit (which is also the income threshold to be categorically needy) is estimated at \$6,670, a number 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 \$4,604 a year when healthy. It should be noted that the medically needy are guaranteed a minimum income of \$6,670 (\$7,270 including the income disregard) so that their total consumption when healthy is at least \$7,270 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 \$3,500 per year. (See De Nardi, French, and Jones [19] for a derivation.) For individuals in this group, the marginal propensity to bequeath, above the threshold level, is 78 cents out of every additional dollar. Several other authors have recently estimated bequest motives inside structural models of old age saving.⁸ Imposing a linear bequest motive, Kopczuk and Lupton [42] 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 et al. [19] 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 [46] 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. [2] 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 of the model's implications for medical and non-medical spending at older ages.

⁸Assembling these figures requires a few derivations and inflation adjustments. Calculations are available on request.

8.2 Model fit

Figure 6 compares the Medicaid reciprocity 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 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, especially at older ages, and to underpredict usage by the rich, especially at younger ages.

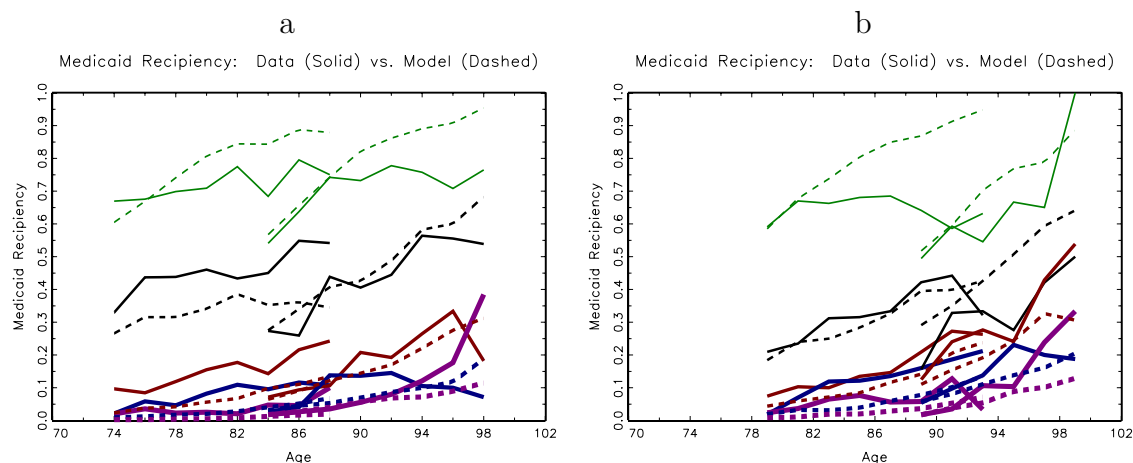


Figure 6: Each line represents Medicaid reciprocity for a cohort-income cell, traced over the time period 1996-2010: data (solid lines) and model (dashed lines). Thicker lines refer to higher permanent income groups. Panel a: cohorts aged 74 and 84 in 1996. Panel b: cohorts aged 79 and 89 in 1996.

Figure 7 plots median net worth by age, cohort, and PI. Here too the model does well, matching the observation that the savings patterns differ by PI and that higher PI people don't run down their assets until well past age 90.




Figure 7: Each line represents median net worth for a cohort-income cell, traced over the time period 1996-2010: data (solid lines) and model (dashed lines). Thicker lines refer to higher permanent income groups. Panel a: cohorts aged 74 and 84 in 1996. Panel b: cohorts aged 79 and 89 in 1996.

Figure 8 displays the median and ninetieth percentile of out-of-pocket medical expenses 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 PI distribution but often exceed \$5,000 for those at the top of the PI 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 PI. 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.

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 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

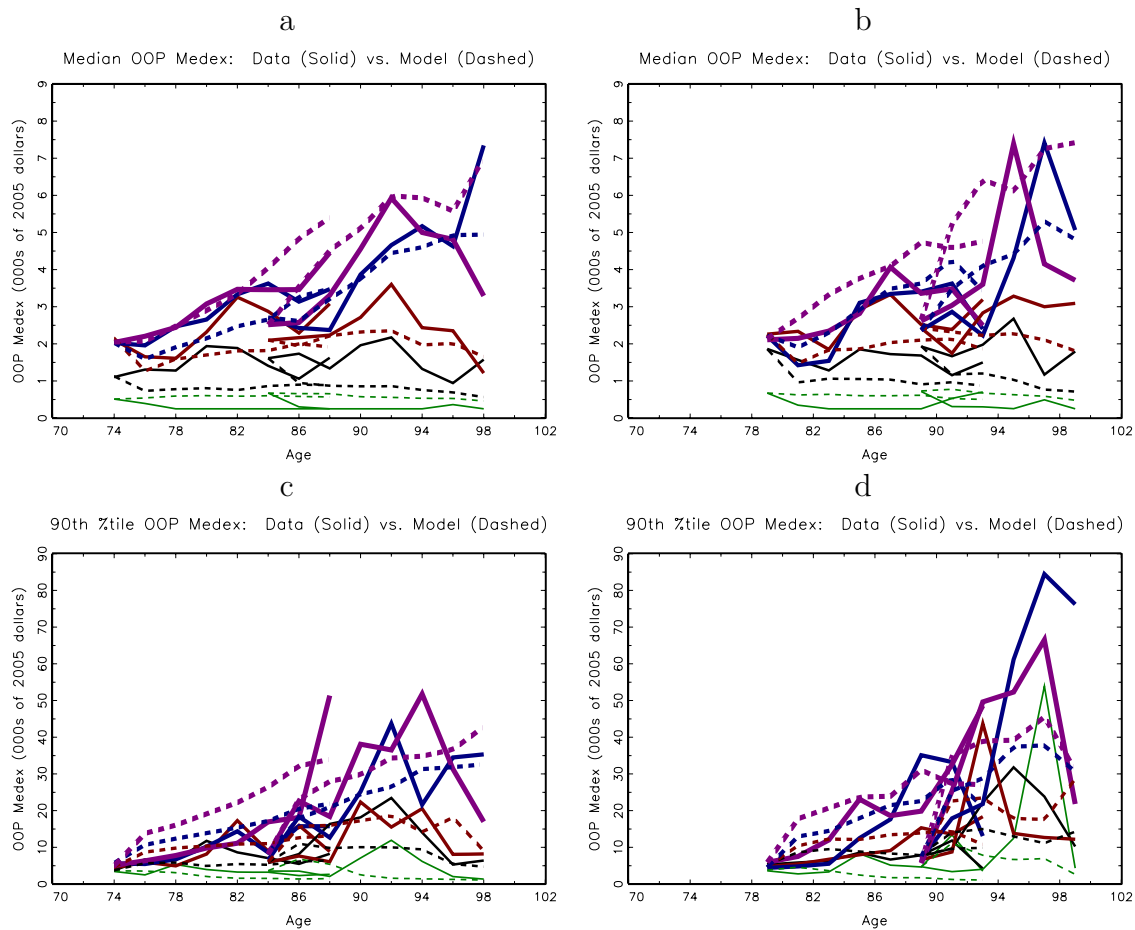


Figure 8: Each line represents median (top panels) and 90th percentile (bottom panels) of medical expenditures for a cohort-income cell, traced over 1996-2010: data (solid lines) and model (dashed lines). Thicker lines: higher permanent income groups. Panels a and b: different cohorts.

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-predicted 90th percentile of out-of-pocket spending is greater than what is observed in the AHEAD data, although it is very close to what is observed in the MCBS.

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 payers. Total expenditures in the MCBS are higher than the model predictions up to the 86th

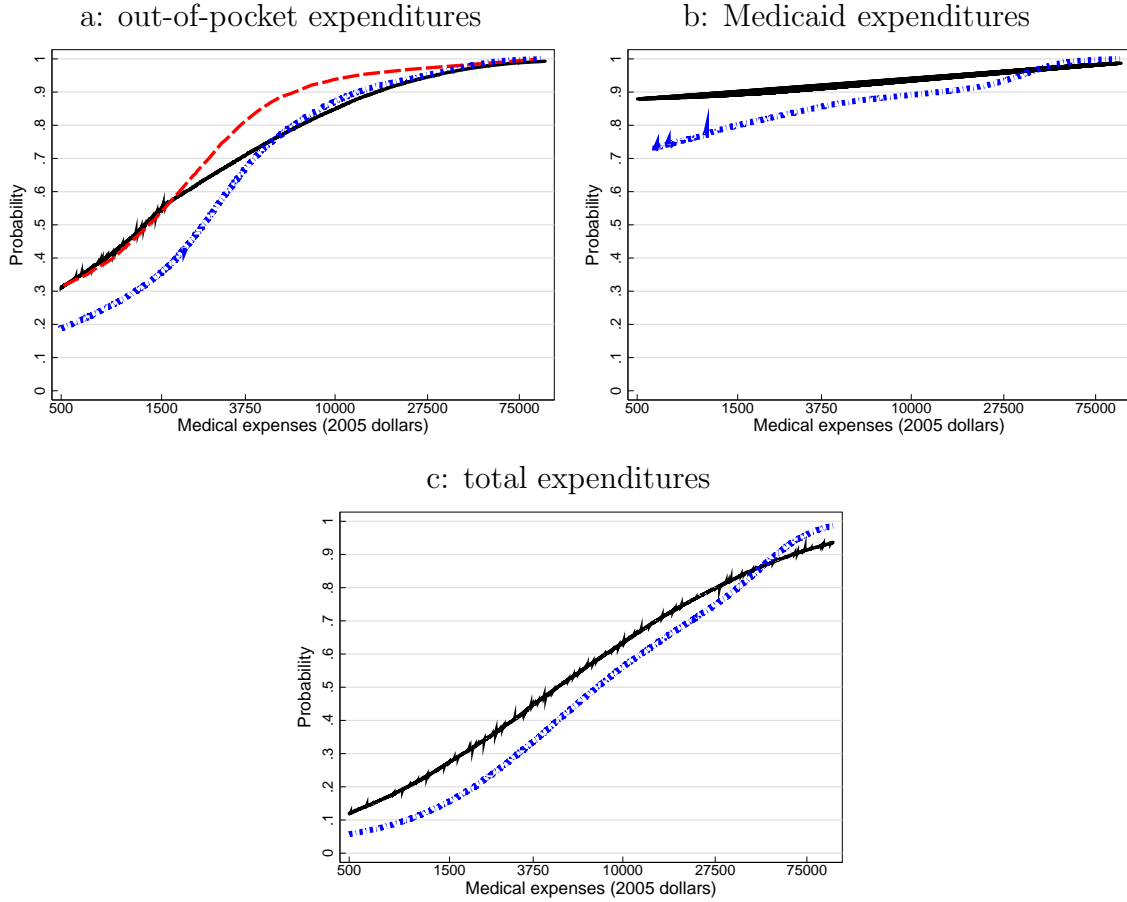


Figure 9: Cumulative distribution functions of medical spending: model (solid line), AHEAD data (dashed line) and MCBS data (dotted line). Panel a: out-of-pocket expenditures. Panel b: Medicaid expenditures. Panel c: total expenditures.

percentile at \$43,000, 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 risk of catastrophic medical spending.

Table 5 shows average Medicaid and out-of-pocket expenditures for each PI quintile, both as predicted by the model and as in the data. The first two columns of Table 5 compare Medicaid expenditures in the MCBS data to those predicted by the model. It shows that retirees at the bottom of the PI distribution have average

Permanent Income Quintile	Medicaid payments		Out-of-pocket expenses		
	MCBS Data	Model	MCBS Data	AHEAD Data	Model
Bottom	9,080	10,070	4,050	2,550	2,210
Fourth	5,720	7,960	5,340	4,270	3,800
Third	2,850	6,000	6,470	5,050	6,330
Second	1,950	3,910	7,300	6,360	8,500
Top	1,280	2,250	8,020	7,000	10,600
Men	2,850	3,780	5,440	4,760	8,280
Women	4,410	5,980	6,470	5,230	6,420

Table 5: Average Medicaid payments and out-of-pocket medical expenditures (2005 dollars), model, MCBS data, and AHEAD data, 1996-2010, for all individuals 72 and older in 1996.

Medicaid expenditures of \$9,080 and \$10,070 in the data and model, respectively. For those at the top of the PI distribution, Medicaid expenditures are \$1,280 and \$2,250 in data and model, respectively. It bears noting that the Medicaid payments reported in the MCBS are on average smaller than those reported in the administrative records: De Nardi et al. [18] find the administrative payments to be 24% higher. Keeping this in mind, Table 5 shows that the model matches Medicaid payments fairly well.

As shown in Table 1, although average Medicaid payments in the MCBS are smaller at the top of the PI distribution, conditional on receiving Medicaid those at the top of the PI distribution receive much larger payments. This is also true in the model. To the best of our knowledge, we are the first to document the progressivity of Medicaid payments among the elderly.⁹

The last three columns of Table 5 compare out-of-pocket expenditures from the MCBS, the AHEAD and the model. The MCBS data shows a less steep PI gradient than the AHEAD data or the model. Those at the bottom of the PI distribution spend \$4,050 in the MCBS data and \$2,380 in the AHEAD data, while expenditures at the top are \$8,020 in the MCBS versus \$6,390 in the AHEAD. Overall, however,

⁹Work by Bhattacharya and Lakdawalla [5] and McClellan and Skinner [50] studied Medicare progressivity.

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 Parameter identification

The preference parameters are identified jointly. There are multiple ways to generate high saving by the elderly: large values of the discount factor β , low values of the utility floors u_c and 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 [22] 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. All of these mechanics are documented in more detail in Appendix F, which shows how changing individual parameters, one at a time, affects the components of our GMM criterion and the life-cycle profiles of several key variables.

We acquire additional identification in several ways. We require our model to match Medicaid reciprocity 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 \$4,500 of consumption by the healthy. A lower floor would generate too few people on Medicaid, especially at higher PI quintiles: Table A5 in Appendix F shows that lowering the utility floor significantly worsens the model’s fit of its Medicaid reciprocity targets. By way of comparison, the model with endogenous medical expenses in De Nardi, French and Jones [19], the one most comparable with the model in this paper, was not estimated to match Medicaid reciprocity rates. That model was able to fit the asset data using a similar value of β , no bequest motives, and lower utility floors. This specification 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 reciprocity 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 reciprocity across age and PI. 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 factor and/or a stronger bequest motive.

The income gradient of medical expenditures helps us pin down the the coefficients of relative risk aversion for non-medical and medical goods, ν and ω . Dividing equation (18) by consumption allows us to obtain an equation governing the optimal ratio of medical goods and services to non-medical consumption goods

$$\frac{m_t}{c_t} = \left(\frac{\mu(\cdot)}{q(h_t)} \right)^{1/\omega} c_t^{\frac{\nu-\omega}{\omega}}. \quad (19)$$

This ratio depends on the relative size of the two risk aversion coefficients, ω and ν . As resources (and thus consumption) grow, $\frac{m_t}{c_t}$ falls if people are more risk averse over medical goods than over non-medical goods ($\omega > \nu$). Put differently, people with higher wealth and permanent income spend a smaller share of their resources on medical goods than on consumption goods when $\omega > \nu$. Our estimates suggest this is the empirically relevant case. Requiring our model to match the observed variation in out-of-pocket medical expenses across permanent income groups helps identify the size of these two parameters. In the data, out-of-pocket medical spending rises with permanent income. However, prior to age 90, the increase in medical expenditure is smaller than the increase in income, suggesting that medical expenditures share of total expenditure is falling in total expenditure, and thus $\omega > \nu$. Hence, inspection of equation (19) and Figures 2, 3, and 5 help explain how ω and ν are separately identified. Appendix F shows that these two parameters are tightly identified. Reducing either ν or ω by 10% leads to large changes in both our GMM criteria and in the age-profiles of assets, Medicaid reciprocity, and the consumption of non-medical goods and medical goods and services.

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 last two lines in Table A5 show the effects of first reducing the average of the medical needs shocks by 10% and then reducing their variance by 10%. Both changes worsen the fit of medical spending, but the first change also significantly worsens the fit of the Medicaid reciprocity moments.

8.4 Medical and non-medical spending in old age: present discounted values




Figure 10: Simulated consumption for the cohort aged 74 in 1996.

To assess the effects of Medicaid from a lifetime perspective, we simulate extended life histories for the youngest cohort. Each simulated individual receives a value of the state vector (t, a_t, g, h_t, I) drawn from the empirical 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 (potentially) age 100. Figure 10 uses these simulations to show the model’s implications for non-medical consumption, showing the trajectory of average consumption for each PI quintile. In contrast to medical expenditures, which rise rapidly with age, average non-medical consumption expenditures decline, albeit slightly, over retirement. This pattern is quite similar to the spending profiles found in the Consumer Expenditure Survey (see, e.g, Banks et al. [4]).

After simulating life histories, we convert the expenditure streams into present discounted values, using the model’s assumed pre-tax interest rate of 4%. Table 6 shows the present discounted value of both non-medical and medical consumption as of age 74. Table 6 reveals that the consumption of medical goods and services is large relative to the consumption of non-medical goods at all PI levels. However, non-medical consumption rises more quickly in PI than total medical spending, as $\nu < \omega$. Non-medical spending for the poorest is 25% of non-medical spending for the richest. In contrast, the total medical spending of the bottom PI quintile is nearly 50% of the total medical spending of the top quintile. In fact, for low PI individuals,

the present discounted value of total medical spending exceeds the present discounted value of non-medical consumption; for high PI individuals, the opposite is true.

Permanent Income Quintile	Non-medical consumption	Medical	
		<u>goods and services</u> Total	Out-of-pocket
Bottom	59,200	108,300	11,200
Fourth	79,700	121,100	20,400
Third	106,800	139,500	35,100
Second	163,900	178,800	55,200
Top	234,900	229,700	80,600
Men	136,000	133,900	42,700
Women	143,800	172,200	46,300
Good Health	173,200	182,200	54,300
Bad Health	97,500	144,000	33,000

Table 6: Present discounted value of non-medical consumption and the consumption of medical goods and services at age 74.

The final column of Table 6 shows that out-of-pocket medical expenses rise in PI even more quickly. This is because Medicaid covers a higher share of medical expenses for the poor. Over their lifetime, the out-of-pocket costs of medical goods and services for the income-richest are over 7 times as large as those of the income-poorest. The table also shows that the present discounted value of all spending, medical and non-medical, is larger for women than men, as they tend to live almost 4 years longer. Furthermore, those in good health also tend to spend more, as they tend to have longer lives and higher PI.

9 Medicaid benefits, taxes, and valuations

9.1 Medicaid benefits received and taxes paid

The first column of Table 7 shows the present discounted value of Medicaid benefits, beginning at age 74. Although the payments decrease by PI quintile, they are

non-trivial for all PI groups. For instance, those in the the highest PI quintile expect to receive \$8,800, which is about 40% of their yearly income. Although the poor are more likely to be receive Medicaid, even the rich are sometimes impoverished by expensive medical conditions, making them eligible for Medicaid benefits too.

Permanent Income Quintile	Medicaid Payments	Medicaid Taxes	Taxes/ Payments
Bottom	33,600	6,700	0.20
Fourth	29,400	8,600	0.29
Third	20,400	20,600	1.01
Second	15,100	30,300	2.00
Top	8,800	40,200	4.59
Men	8,600	28,800	3.34
Women	22,400	20,000	0.89
Good Health	17,800	26,400	1.48
Bad Health	23,700	12,500	0.53

Table 7: Present discounted value of Medicaid payments received (simulated from the model), Medicaid taxes paid (computed from the PSID), and the ratio of Medicaid taxes to Medicaid payments, all from the standpoint of age 74.

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

The middle column of Table 7 calculates the present discounted value at age 74 of the taxes paid to finance Medicaid transfers over all of one’s life, including the working period. Since we do not explicitly model the working period, to calculate Medicaid tax payments, we modify the approach found in McClellan and Skinner [50], who calculate tax payments for Medicare. We first use data from the Panel Study of Income Dynamics (PSID) to calculate lifetime taxes paid by different groups. Because Medicaid has no dedicated funding source, we assume that it is financed by a tax

schedule that is proportional to total tax payments, and that the average Medicaid tax rate in this progressive tax schedule balances the Medicaid budget for this cohort. See Appendix E for more details.

We use the PSID because it includes income from spouses who have died before the AHEAD sample begins. A large share of our sample consists of elderly widows. To capture the progressivity of the taxes they paid when young, we need a data source that includes income from their deceased husbands. Because high income women tend to marry high income men, ignoring the income and taxes paid by husbands would understate the taxes paid by higher-income widows relative to lower-income people, who might have not been married, or married with lower-earning spouses. Although the AHEAD has tax records from working years, information on taxes paid by deceased spouses is incomplete.

Those in the top PI quintile pay on average \$42,400 in taxes towards Medicaid, nearly 6 times as much as those in the bottom of the PI distribution. This reflects both higher income and higher marginal tax rates. As a result, those at the top of the PI distribution pay in much more than they receive in Medicaid payments. The rightmost column of Table 7 shows the ratio of taxes paid to transfers received. Those at the top of the distribution pay on average \$5 in taxes for every \$1 in payments received, whereas those at the bottom of the distribution pay \$0.21 for every \$1 in transfers received.

9.2 Household valuations of Medicaid

In this section, we simulate changes in Medicaid generosity and compare the resulting increases (or decreases) in government costs to the resulting gains (or losses) in consumer welfare.

To measure the costs of a Medicaid reform we compute by how much the present discounted value of Medicaid payments changes when the program changes. This represents the increase (or decrease) in the lifetime actuarial cost of providing Medicaid insurance and is an ex-post measure.

To measure the welfare gains, we compute the compensating variation; that is, the immediate payment after the Medicaid reform that would leave the retiree as well off as before the reform. This is an ex-ante measure. More specifically, the compensating variation at age 74, $\lambda_{74} = \lambda(a_{74}, g, h_{74}, I, \zeta_{74}, \xi_{74})$, is computed as:

$$V_t(a_t, g, h_t, I, \zeta_t, \xi_t; \text{current Medicaid}) = V_t(a_t + \lambda_{74}, g, h_t, I, \zeta_t, \xi_t; \text{Medicaid reform}),$$

where $V_t(a_t, g, h_t, I, \zeta_t, \xi_t; \cdot)$ is the value function evaluated at a given set of state

variables, either in the world with current Medicaid (left hand side of the equation above) or in a world with a reformed Medicaid program. Our measure is similar to the ones computed for Medicare by Finkelstein and McKnight [25] and McClellan and Skinner [50] but uses a forward-looking value function, rather than a static utility function. When considering a group, we simply take averages across all its members.

If Medicaid provides retirees with valuable insurance, the compensating variation may exceed the change in the actuarial value of Medicaid payments. On the other hand, people may value the transfer flows at less than their actuarial value. For example, if they are very impatient, they might prefer having the cash today, to dispose of as they wish, over receiving Medicaid transfers in the future. Furthermore, assets are taxed at 100% for those receiving Medicaid transfers, which in turn distorts savings decisions.

To distinguish the insurance provided by the categorically and the medically needy programs, we first analyze a 10% decrease in the categorically needy utility floor. This corresponds to the consumption of the categorically needy when healthy dropping from \$4,610 to \$4,140. Columns (1) and (2) of Table 8 show that this change only affects people in the bottom two PI quintiles, as people with higher incomes never qualify as categorically needy. The discounted present value of Medicaid payments drops by \$4,100 and \$2,100, respectively, for people in the two bottom PI quintiles. Column (2) reports the compensating variation.

Column (3) presents the ratio of column (2) to column (1), and reveals that the categorically needy people value their lost Medicaid insurance at more than the cost of providing it. However, the ratio is not very large, suggesting that the insurance value of these transfers, at the margin, is not very large. Nonetheless, because this group pays only a small fraction of the transfers' cost (see Table 7), the value they place on their Medicaid benefits almost surely exceeds their associated tax burden.

We next cut the consumption value of both utility floors (that is, both the categorically and medically needy floors) by 10%, and simulate our model again. The right-hand-side panel of Table 8 shows the resulting reductions in Medicaid payments and their compensating variations. A striking feature of this table is that while people in the lowest three PI quintiles value Medicaid fairly close to its cost, people in the top two PI quintiles value Medicaid at two to three times its cost. In fact, the compensating variation for retirees in the top PI quintile, \$4,400, is as big as that of the middle quintile, and is two-thirds as big as the compensating variation at the bottom. The insurance value of Medicaid is very high for these people because of two reasons. First, because these people are high-income they have a high lifetime level of consumption, and thus have more consumption to lose should it fall. Second, they

	Categorical floor down 10%			Both floors down 10%		
	(1)	(2)	(3)	(4)	(5)	(6)
Permanent Income Quintile	Reduction in PDV of Payments	Compensating Variation	Ratio of (2)/(1)	Reduction in PDV of Payments	Compensating Variation	Ratio of (5)/(4)
Bottom	4,100	5,600	1.37	4,500	6,300	1.40
Fourth	2,100	2,200	1.05	4,000	5,000	1.25
Third	0.0	0.0	NA	2,900	4,400	1.52
Second	0.0	0.0	NA	2,200	4,100	1.86
Top	0.0	0.0	NA	1,400	4,400	3.14
Men	300	200	0.67	1,300	1,100	0.85
Women	1,200	1,600	1.33	3,100	5,600	1.81
Good Health	700	900	1.29	2,600	4,800	1.85
Bad Health	1,700	2,200	1.29	3,300	5,000	1.52

Notes. Left panel: the categorically needy floor is cut by 10%. Right panel: both Medicaid floors are cut by 10%. Columns (1) and (4): decrease in the present discounted value of Medicaid payments as of age 74. Columns (2) and (5): dollar amount needed to compensate people for the Medicaid benefit cut. Columns (3) and (6): ratio of column (2) to column (1) and column (5) to column (4), respectively, which give the average compensating variation per dollar of reduced Medicaid benefits.

Table 8: The costs and benefits of cutting Medicaid by 10%.

face the double compounded risk of living well past their life expectancy, and facing extremely high medical needs. It is in those states of the world that insurance is most valuable. Offsetting these insurance gains, however, is a redistributive tax system. While individuals in the top income quintile place a value of \$3.14 on each dollar of transfers, they pay \$4.59 of taxes (Table 7).

In Table 9, we analyze the benefits of making the Medicaid program more generous, by increasing the Medicaid consumption floor by 10% (from \$4,610 to \$5,070). Table 9 shows that people at the bottom PI quintiles value these Medicaid increases at less than their cost, people in the next two quintiles value them at slightly above cost, and people in the top quintile value them by twice as much. In the aggregate, taking averages over all retirees reveals that the cost increase associated with a more

Permanent Income Quintile	Both floors up 10%		
	(1) Payment Increase	(2) Compensating Variation	(3) Ratio (2)/(1)
Bottom	4,700	2,600	0.55
Fourth	4,200	3,100	0.74
Third	3,100	3,600	1.16
Second	2,300	2,900	1.26
Top	1,300	2,600	2.00
Men	1,400	600	0.43
Women	3,300	3,500	1.06
Good Health	2,500	3,000	1.20
Bad Health	3,500	3,000	0.86

Notes. Column (1): increase in the present discounted value of Medicaid payments at age 74. Column (2): dollar amount people would be willing to pay to receive the higher Medicaid benefits. Column (3) is the ratio of column (2) to column (1), which show the average compensating variation per dollar of reduced Medicaid benefits.

Table 9: The costs and benefits of increasing Medicaid by 10%.

generous Medicaid program slightly exceeds the average valuation. Comparing the valuations to the associated tax burdens (see Table 7), however, produces different implications. Even though high-income retirees would receive the most “bang per buck” from a Medicaid expansion, under the current redistributive tax system they would not support it, as their tax burden would rise more than their valuation. In contrast, low-income retirees, who receive the least bang per buck from a Medicaid expansion, would support the expansion, as their tax burden would rise by even less. Only people in the middle quintile value a Medicaid expansion in excess of both its cost and their tax burden.

Put together, the results in Tables 8 and 9 indicate that under current programs rules people value Medicaid transfers at more than their actuarial cost, but that increasing Medicaid’s generosity would raise its insurance value by less than its cost. Our model therefore suggests that the current Medicaid system is of about the right size for most currently retired singles.

9.3 Long-term-care insurance

While our model includes endogenous medical spending and several dimensions of individual-level heterogeneity, it abstracts from the decision to purchase long term care insurance (LTCI). Only about 9% of elderly singles have LTCI (Lockwood [46]) and only 4% of LTC expenditures are paid for by LTCI (Congressional Budget Office [55]). Given that our results suggest that the elderly, and especially the high income elderly, value Medicaid insurance heavily, it is surprising that the market for LCTI is so small.

Brown and Finkelstein argue that that one major reason that the LTCI market is so small is that Medicaid crowds out LTCI and thus that major reductions in Medicaid would increase LTCI use. This is due to the fact that Medicaid is a payer of last resort and is subject to asset and income tests, which implies that LTCI payments for nursing home care would often crowd out Medicaid payments for the same services.

If there are fixed costs to acquiring/providing or discarding LTCI, larger changes in Medicaid generosity are more likely to induce changes in LTCI holdings. Our experiments thus involve relatively small changes to the Medicaid program, which imply smaller incentives to change LTCI positions. But even in the absence of transaction costs, there are other important reasons why LTCI use is limited and why it would likely stay limited even if Medicaid generosity was reduced by a reasonable amount. These factors include:

1. Lack of efficiency in the private market for long-term care insurance. Prices are high: Brown and Finkelstein [10] report that imperfect competition and transaction costs result in prices that are marked up substantially above expected claims, with loads on typical policies from 18 to 51 cents on the dollar, depending on whether one takes into account lapsed policies. These loads are much higher than loads that have been estimated in other private insurance markets and point to the existence of one or more supply side imperfections.
2. Limited insurance against nursing home risk. Brown and Finkelstein [11] report that comprehensive LTCI contracts exist but are not purchased. The typical LTCI contracts held by households cap both the maximum number of days covered over the life of the policy and the maximum daily payment for a nursing home stay, a daily payment that is often fixed in nominal terms (Fang [24]). Even the policies that provide some kind of indexation of the daily maximum payment are typically linked to aggregate price indexes rather than actual nursing home costs, thus generating substantial purchasing power risk between the

time a person purchases the policy and the time she enters a nursing home. As a result, most available policies do not provide insurance against tail risk, which is exactly the risk that the richest in our model fear the most, due to longer longevity and higher risk of large medical needs when very old.

3. Severe adverse selection. Hendren [34] shows that when private information problems are sufficiently large within certain subgroups, insurance markets fail to emerge. His main empirical findings are that a large fraction of those applying for insurance are rejected by underwriters, and that those who are rejected hold significant private information. He also finds that 23% of 65 year olds have health conditions that preclude them from purchasing LTCI.
4. Bequest motives (Lockwood [46]). In a framework with exogenous medical spending, Lockwood argues that reasonably estimated bequest motives, together with medical expense risk, help fit the patterns of both asset decumulation and LTCI purchases seen in the data. We also estimate a significant bequest motive, which reduces the value of LTCI.

9.4 Unpacking the Results: Moral hazard and exogenous expenditures

An important and open question is the extent to which people impoverish themselves in order to qualify for Medicaid. Because our model includes savings and medical spending choices, it is well suited to address the quantitative importance of this form of moral hazard. Moral hazard arises within our model both contemporaneously, in that retirees may purchase too much subsidized health care, and dynamically, in that people might be over-spending over a number of periods to qualify for Medicaid in the future. To better understand the quantitative importance of moral hazard, we analyze further the 10% cut in Medicaid generosity considered in columns (4)-(6) of Table 8.

A change in Medicaid generosity has two effects. First, it mechanically changes eligibility and transfers at any given level of individual resources. Second, it changes the degree of moral hazard by changing the incentives to consume and save. To help disentangle these effects, Table 10 shows total and out-of-pocket medical spending, non-medical spending, and Medicaid reciprocity rates at age 85 for the simulated life histories used to construct Table 8. The top panel of Table 10 shows quantities for the estimated baseline model, while the second panel shows the quantities after a

10% reduction in the Medicaid utility floor. The bottom two panels show the differences between the two cases, in absolute and then relative (percentage) terms. Table 10 shows that a 10% Medicaid cut would lead to lower non-medical spending. Non-medical spending at age 85 falls on average \$100- \$290, depending on the income group. However, this decline in non-medical spending is modest relative to the decline in total medical spending, which falls by an average of \$430-\$1,790. The decline in total medical spending is driven by lower government transfers: while cutting Medicaid reduces total medical expenditures, out-of-pocket expenditures *increase* slightly for all groups but the richest. Although people save more after a Medicaid cut, the savings response is modest relative to the decline in total medical spending from reduced transfers. These findings indicate that the mechanical effects of changing Medicaid are larger than the moral hazard effects.¹⁰

Another way to assess the importance of moral hazard to treat medical expenditures as an exogenous, rather than endogenous, variable. To do this we find the stochastic process for exogenous medical expenses that allows the model to best fit its estimation targets. We hold preference parameters fixed, so that our experiment focusses solely on changes to medical spending. Although our process for exogenous medical spending does not depend on wealth or permanent income, the medical spending that it generates is fairly similar to that of the endogenous medical spending model. We show model predictions for the exogenous medical spending specification in Figure 16.

¹⁰Using data from a large self-insured employer, Bajari et al. [3] find significant moral hazard. They focus on how changes in the co-insurance rate q changes the allocation of medical versus non-medical spending, but do not focus on savings. In contrast, in our Medicaid reforms, the coinsurance rates remain at their baseline values – the intratemporal allocation still obeys equation (18) – and we focus on the effect of Medicaid insurance on savings.

Permanent Income	Consumption	Medical Expenditures		Medicaid Recipiency (%)
		Total	OOP	
Baseline				
Bottom	7,890	22,110	1,430	89.1
Fourth	9,810	21,280	2,860	42.9
Third	12,380	22,710	5,310	9.9
Second	18,520	25,630	8,210	3.7
Top	26,330	32,790	11,860	1.3
Medicaid Floors Decreased 10%				
Bottom	7,600	20,320	1,580	88.6
Fourth	9,600	19,870	3,010	41.0
Third	12,240	21,580	5,420	8.9
Second	18,420	24,980	8,260	3.3
Top	26,210	32,360	11,850	1.2
Difference				
Bottom	-290	-1,790	150	-0.5
Fourth	-210	-1,410	150	-2.0
Third	-140	-1,130	110	-1.0
Second	-100	-650	50	-0.4
Top	-120	-430	-10	-0.1
Percentage Differences				
Bottom	-3.7	-8.1	10.5	-0.6
Fourth	-2.1	-6.6	5.2	-4.5
Third	-1.1	-5.0	2.1	-9.6
Second	-0.5	-2.5	0.6	-10.6
Top	-0.5	-1.3	-0.1	-7.1

Table 10: The effects of decreasing Medicaid payments by 10%, Age 85.

We then use the exogenous medical spending model to re-evaluate the 10% Medicaid cut considered in Tables 8 and 10. The compensating variations associated with this experiment are, from the bottom PI quintile to the top: \$5,200, \$5,300, \$5,000, \$5,700, and \$8,400. The similarity between these valuations and the valuations in the fifth column of Table 8 is not surprising. The utility floor in the endogenous spending model is indexed with consumption, as $u = \frac{1}{1-\nu}c^{1-\nu}$, which is identical to utility in the exogenous spending model when the consumption floor is \underline{c} . A 10% cut in \underline{c} thus represents the same reduction in guaranteed utility for both medical spending specifications, and should be valued similarly under both specifications.¹¹ Our finding that high-income retirees often value Medicaid as much as poorer retirees is robust to making medical expenses exogenous.

When medical spending is endogenous, cuts to the utility floor reduce both consumption and medical expenditures; recall that the transfers are allocated optimally between consumption and out-of-pocket medical spending.¹² When medical spending is exogenous, cuts to the utility floor can only reduce consumption. The transfer reductions associated with a cut to the utility floor are thus smaller, and the valuations per unit of spending higher, when medical spending is exogenous. We still find, however, that high-PI people have the highest valuation per dollar of spending.

9.5 Adverse selection

Our model generates considerable heterogeneity in health, mortality, and medical needs. To quantify the extent to which the Medicaid population are adversely selected on the basis of their medical needs, Table 11 compares Medicaid recipients to other retirees along several dimensions. We construct Table 11 by simulating the baseline model over the sample period 1996-2010 and taking cross-sectional averages. The first row of the table shows, unsurprisingly, that Medicaid recipients are considerably poorer. The second row shows that Medicaid recipients indeed have much higher total medical spending.

Because medical spending in our model represents the convolution of medical needs (μ) and financial incentives, we also consider the distribution of the medical preference shifter μ . Bajari et al. [3] measure adverse selection in a similar, if more

¹¹Differences in model dynamics, along with differences in the estimated spending processes, mean the valuations will not be identical.

¹²When ν and ω are close in value, as they are in our estimates, equation (18) can be approximated as $m = \left(\frac{\mu}{q}\right)^{1/\nu} c$, so that medical spending is proportional to consumption. Cuts in the utility floor thus reduce medical and non-medical spending by similar proportions.

detailed, way. Compared to non-recipients, the values of μ that confront Medicaid recipients are 4.7 times as likely to lie in the top decile, and 28 times as likely to lie in the top percentile. In short, Medicaid recipients are more likely to be sick, and far more likely to be very sick, consistent with Medicaid’s role as the payer of last resort, and consistent with our argument that Medicaid provides valuable insurance against catastrophic medical events.

	All Retirees	Medicaid Recipients	Non- Recipients
Mean Net Income	16,270	6,240	19,020
Mean Medical spending			
Total	28,930	61,820	19,920
Out-of-pocket	6,680	2,940	7,700
Medicaid	5,680	26,440	0
Medical needs shifter (μ)			
Mean	1,270	5,150	210
Fraction in top 50% of distribution (%)	50.0	67.9	45.1
Fraction in top 10% of distribution (%)	10.0	26.0	5.6
Fraction in top 1% of distribution (%)	1.0	4.0	0.1
Fraction of population (%)	100.0	19.7	80.3

Table 11: Comparison of Medicaid and non-Medicaid recipients: Simulations

10 Conclusions and future research

In this paper we assess the effects of Medicaid insurance on single retirees. Although Medicaid payments decrease with permanent income, even higher income people can receive sizeable Medicaid payments because they tend to live longer and face higher medical needs in very old age. Furthermore, our compensating variation calculations show that many higher income retirees value Medicaid insurance as much or more than lower-income ones. Our compensating variation calculations also 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. This suggests that the Medicaid program may currently be of the approximate right size for currently single retirees.

In the interest of tractability, our framework does not allow households to adjust their holdings of LTCL. Although only 9% of the households in our AHEAD sample hold such insurance, cuts to Medicaid may compel households to increase their coverage. Introducing this additional margin of portfolio choice to our model could lower our estimates of the value households place on Medicaid. While in Section 9.3 we argue that there are many reasons to think that introducing LTCL decisions would not significantly affect our results, it is worth studying this question more formally.

By focussing on the retirement period, we are able to explicitly model many dimensions of uncertainty and heterogeneity and to treat medical expenditures as a choice variable. However, it would be valuable to model the entire life cycle, the distortions generated by the income taxes needed to finance Medicaid, and the anticipated effects of Medicaid changes at younger ages.

By concentrating on single retirees, we study the population that is most likely to receive Medicaid transfers. The data shows that couples tend to be richer and less likely to end up in nursing homes and thus receive much smaller Medicaid payments. For example, singles in our MCBS sample on average receive \$3,760 in Medicaid transfers a year, while couples in the same age range on average receive \$2,140, or \$1,070 per person. It nonetheless would be interesting to extend our analysis to include the valuation of Medicaid insurance by couples.

References

- [1] 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.
- [2] 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. *Journal of Finance*, 66(2):519–561, 2011.
- [3] Patrick Bajari, Han Hong, Ahmed Khwaja, and Christina Marsh. Moral hazard, adverse selection and health expenditures: A semiparametric analysis. *RAND Journal of Economics*, 45(4):747–763, 2014.
- [4] James Banks, Richard Blundell, Peter Levell, and James P. Smith. Life-cycle consumption patterns at older ages in the US and the UK. RAND Working Paper WR-1100, 2015.
- [5] Jay Bhattacharya and Darius Lakdawalla. Does Medicare benefit the poor? *Journal of Public Economics*, 90(1):277–292, 2006.
- [6] David M. Blau and Donna Gilleskie. The role of retiree health insurance in the employment behavior of older men. *International Economic Review*, 49(2):475–514, 2008.
- [7] R. Anton Braun, Karen A. Kopeccky, and Tatyana Koreshkova. Old, sick, alone and poor: A welfare analysis of old-age social insurance programs. Mimeo, Federal Reserve Bank of Atlanta, 2015.
- [8] 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.
- [9] 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.
- [10] 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.

- [11] Jeff Brown and Amy Finkelstein. The private market for long-term care insurance in the u.s.: A review of the evidence. *Journal of Risk and Insurance*, 76(1):5–29, 2009.
- [12] 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.
- [13] Moshe Buchinsky. Recent advances in quantile regression models: A practical guideline for empirical research. *Journal of Human Resources*, 33:88–126, 1998.
- [14] David Card, Carlos Dobkin, and Nicole Maestas. Does Medicare save lives? *The Quarterly Journal of Economics*, 124(2):597–636, 2009.
- [15] Gary Chamberlain. Comment: Sequential moment restrictions in panel data. *Journal of Business & Economic Statistics*, 10(1):20–26, 1992.
- [16] Kenneth Y. Chay, Daeho Kim, and Shailender Swaminathan. Medicare, hospital utilization and mortality: Evidence from the programs origins. Mimeo, 2010.
- [17] 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.
- [18] Mariacristina De Nardi, Eric French, John Jones, and Jeremy McCauley. Medical spending on the U.S. elderly. Mimeo, University College London, 2015.
- [19] 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.
- [20] Mariacristina De Nardi, Eric French, John B. Jones, and Angshuman Gooptu. Medicaid and the elderly. NBER Working Paper 17689, 2011.
- [21] Darrell Duffie and Kenneth J. Singleton. Simulated moments estimation of markov models of asset prices. *Econometrica*, 61(4):929–952, 1993.
- [22] 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.

- [23] Liran Einav, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen. Selection on moral hazard in health insurance. *American Economic Review*, 103(1):178–219, 2013.
- [24] Hanming Fang. Insurance markets for the elderly. In John Piggott and Alan Woodland, editors, *Handbook of the Economics of Population Ageing*, page forthcoming. Elsevier, 2014.
- [25] 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.
- [26] Amy Finkelstein, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, 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.
- [27] 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.
- [28] Raquel Fonseca, Pierre-Carl Michaud, Titus Galama, and Arie Kapteyn. On the rise of health spending and longevity. Rand Working Paper WR-722, 2009.
- [29] The Kaiser Foundation. *Medicaid: A Primer*. <http://www.statehealthfacts.org/>, 2013.
- [30] Eric French and John Bailey Jones. On the distribution and dynamics of health care costs. *Journal of Applied Econometrics*, 19(4):705–721, 2004.
- [31] 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.
- [32] Michael Grossman. On the concept of health capital and the demand for health. *Journal of Political Economy*, 80(2):223–255, 1972.
- [33] 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.

- [34] Nathaniel Hendren. Private information and insurance rejections. *Econometrica*, 81(5):1713–1762, 2013.
- [35] 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.
- [36] R. Glenn Hubbard, Jonathan Skinner, and Stephen P. Zeldes. Precautionary saving and social insurance. *Journal of Political Economy*, 103(2):360–399, 1995.
- [37] Michael D. Hurd. Mortality risk and bequests. *Econometrica*, 57(4):779–813, 1989.
- [38] Michael D. Hurd, Daniel McFadden, and Angela Merrill. Predictors of mortality among the elderly. Working Paper 7440, National Bureau of Economic Research, 1999.
- [39] Selahattin İmrohoroğlu and Sagiri Kitao. Social security, benefit claiming and labor force participation: A quantitative general equilibrium approach. Mimeo, 2009.
- [40] 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.
- [41] Ruud H. Koning. Kernel: A gauss library for kernel estimation. <http://www.xs4all.nl/>
- [42] Wojciech Kopczuk and Joseph Lupton. To leave or not to leave: The distribution of bequest motives. *Review of Economic Studies*, 74:207–235, 2007.
- [43] Karen Kopecky and Tatyana Koreshkova. The impact of medical and nursing home expenses and social insurance policies on savings and inequality. Mimeo, 2009.
- [44] Laurence J. Kotlikoff. Health expenditures and precautionary savings. In Laurence J. Kotlikoff, editor, *What Determines Saving?* Cambridge, MIT Press, 1988, 1988.
- [45] John Laitner, Daniel Silverman, and Dmitriy Stolyarov. Annuitized wealth and post-retirement saving. Working Paper 20547, National Bureau of Economic Research, October 2014.

- [46] Lee Lockwood. Incidental bequests: Bequest motives and the choice to self-insure late-life risks. Mimeo, 2014.
- [47] 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.
- [48] Charles F. Manski. *Analog Estimation Methods in Econometrics*. Chapman and Hall, 1988.
- [49] Samuel Marshall, Kathleen McGarry, and Jonathan S. Skinner. The risk of out-of-pocket health care expenditure at the end of life. In *Explorations in the Economics of Aging*, pages 101–128. University of Chicago Press, 2011.
- [50] Mark McClellan and Jonathan Skinner. The incidence of Medicare. *Journal of Public Economics*, 90(1):257–276, 2006.
- [51] Thad Mirer. The wealth-age relation among the aged. *American Economic Review*, 69:435–443, 1979.
- [52] Makoto Nakajima and Irina Telyukova. Home equity in retirement. Mimeo, 2012.
- [53] Whitney K. Newey. Generalized method of moments specification testing. *Journal of Econometrics*, 29(3):229–256, 1985.
- [54] 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.
- [55] Congressional Budget Office. Financing long-term care for the elderly. <http://www.cbo.gov/ftpdocs/54xx/doc5400/04-26-longtermcare.pdf>, 2004.
- [56] The Kaiser Commission on Medicaid and the Uninsured. *Medicaid: A Primer*. Menlo Park, CA, The Henry J. Kaiser Family Foundation, 2010.
- [57] Serdar Ozkan. Income inequality and health care expenditures over the life cycle. Mimeo, Federal Reserve Board, 2011.
- [58] Ariel Pakes and David Pollard. Simulation and the asymptotics of optimization estimators. *Econometrica*, 57(5):1027–1057, 1989.

- [59] Michael G. Palumbo. Uncertain medical expenses and precautionary saving near the end of the life cycle. *Review of Economic Studies*, 66:395–421, 1999.
- [60] Svetlana Pashchenko and Ponpoje Porapakarm. Quantitative analysis of health insurance reform: Separating regulation from redistribution. *Review of Economic Dynamics*, 16(3):383 – 404, 2013.
- [61] 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.
- [62] 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.
- [63] James Powell. Estimation of semiparametric models. In Robert Engle and Daniel L McFadden, editors, *Handbook of Econometrics, Volume 4*. Elsevier, Amsterdam, 1994.
- [64] 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.
- [65] John Karl Scholz and Ananth Seshadri. Health and wealth in a lifecycle model. Mimeo, 2013.
- [66] John Karl Scholz, Ananth Seshadri, and Surachai Khitatrakun. Are americans saving optimally for retirement? *Journal of Political Economy*, 114:607–643, 2006.
- [67] George Tauchen. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2):177–181, 1986.
- [68] Hilary Waldron. Trends in mortality differentials and life expectancy for male social security-covered workers, by socioeconomic status. *Social Security Bulletin*, 67(3):1–28, 2007.
- [69] Motohiro Yogo. Portfolio choice in retirement: Health risk and the demand for annuities, housing, and risky assets. Mimeo, University of Pennsylvania, 2009.
- [70] 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: for online publication: 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.¹³

The MCBS also accurately measures the share of the population receiving Medicaid payments.¹⁴ However, MCBS Medicaid payments for the age 65+ population are on average 32% smaller than what administrative data from the MSIS suggest.

¹³Medicare statistics are located at http://www.census.gov/compendia/statab/cats/health_nutrition/medicare_medicaid.html.

¹⁴According to MCBS data, there were on average 5.1 million age 65+ Medicaid beneficiaries over the 1996-2006 period, versus 4.7 million “aged” (which mostly refers to aged 65+) 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/MedicareDataSourcesGenInfo/MSIS-Tables.html>. See De Nardi et al. [18] for further comparisons of the MCBS data to administrative data on Medicare and Medicaid beneficiaries and payments.

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	8,630
95-100%	5%	40.5%	100,060	69,410
90-95%	5%	20.1%	50,180	37,510
70-90%	20%	32.5%	21,940	13,150
50-70%	20%	5.9%	3,690	2,460
0-50%	50%	1.0%	240	330

Note: 2010 MSIS data, adjusted to 2005 dollars

Table A1: Medicaid enrollment and expenditures by enrollee spending percentile, MSIS versus MCBS.

Table A1 compares the distribution of the MSIS administrative payment data (taken from Young et al. [70]) to data from the MCBS. We show the MCBS distribution for all dual Medicare/Medicaid beneficiaries, the set closest to the the sample in the MSIS data. 59% of all dual eligibles are age 65+, the other 41% being disabled individuals under age 65 who are potentially more costly than the age 65+ dual eligibles. Table A1 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 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 for Medicaid beneficiaries in this MCBS subsample are \$13,620.

The next set of benchmarking exercises that we perform is for out-of-pocket medical spending, Medicaid reciprocity 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

Income Quintile	AHEAD data				MCBS data		
	Total income	Annuity income	Out-of-pocket expenses	Medicaid reciprocity	Total income	Out-of-pocket expenses	Medicaid reciprocity
1	7,740	4,820	2,550	60.9%	6,750	4,050	69.9%
2	10,290	8,270	4,270	28.1%	10,020	5,340	41.8%
3	15,500	10,900	5,050	11.0%	13,740	6,470	15.5%
4	19,290	14,390	6,360	5.6%	19,710	7,300	8.0%
5	33,580	26,300	7,000	3.0%	44,150	8,020	5.4%

Table A2: Income, out-of-pocket spending, and Medicaid reciprocity rates, AHEAD versus MCBS, 1996-2010, for those age 72 and older in 1996.

individuals (the MCBS asks only about total income). The first four columns of Table A2 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.¹⁵ The share of the population receiving Medicaid transfers is also very similar in the AHEAD and MCBS. 61% and 70% of those in the bottom PI quintile are on Medicaid in the the AHEAD and MCBS, respectively. In the top quintile, 3% of people are on Medicaid in the AHEAD whereas 5% are in the MCBS. The higher Medicaid reciprocity rate in the MCBS might reflect that the MCBS data has administrative information on whether individuals are on Medicaid, which eliminates underreporting

¹⁵There 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.

problems.

We also assessed the usefulness of the Medicaid-related data in MEPS. A key problem with the MEPS data, however, is 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. [70] 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 useful for understanding the remaining 21% of Medicaid payments. Consistent with this fact, mean Medicaid payments in the MEPS for elderly beneficiaries are only \$3,499, whereas they are \$8,630 in the MCBS, and \$13,414 according to the administrative data from the MSIS.

Appendix B: for online publication: computational details

This Appendix details our simulation procedure.

1. To find optimal decision rules, we solve the model backwards using value function iteration. The state variables of the model are assets, gender, health status, permanent income, and the permanent and transitory components of medical spending (ζ and ξ). At each age, we solve the model for 200 grid points for assets, two points for gender (male and female), three points for health (good, bad, and nursing home), 13 grid points for permanent income, five points for the persistent component of medical needs shocks, and four points for the idiosyncratic component of the medical needs shocks. Our approach for discretizing the medical needs shocks follows Tauchen [67], with the grid spaced over the percentile range [0.175, 0.825], a specification we found to work well.
2. Our initial sample of simulated individuals is large, consisting of 150,000 random draws from the individuals in the first wave of our data. Given that we randomly simulate a sample of individuals that is larger than the number of individuals observed in the data, most observations will be used multiple times.
3. All initial conditions are taken from the data, except for the split between the permanent and transitory components of the medical spending shifters (ζ and ξ). Regarding the final two variables, we only observe out-of-pocket medical expenses, which in our model are a function of not only the spending shifters, but all the other state variables. Recall that forward-looking retirees will respond differently to persistent and transitory shocks of the same size. Inferring the two shocks would thus involve a costly filtering procedure utilizing the model's

decision rules. We instead draw the initial values of ζ and ξ from their invariant distributions.

4. For each draw, not only we take the joint distribution of the initial conditions for the state variables, but we also use the observed health and mortality history experienced by that particular individual. We assign entire health and mortality histories to insure that we properly match how our sample composition changes with age. One concern is that our sample is fairly small, so that the medians (or 90th percentiles) of wealth or medical spending in some cohort-income groups can change with the deaths of a few individuals. While we expect these effects to average out if we forward-simulated demographic transitions, it is simpler to match the data if we base our simulations on actual life histories. A more fundamental issue is that the processes for health and mortality that we feed into the model do not depend on wealth, because wealth is an endogenous variable in our model. However, we know that high wealth is a good predictor of longevity, conditional on the other state variables. Our simulation procedure captures the initial wealth/mortality gradient by construction, whereas our estimated health and mortality transition models do not.
5. Given the optimal decision rules, and the initial conditions of the state variables, we calculate life histories for savings, consumption, Medicaid reciprocity, and medical spending.
6. We aggregate the simulated data in the same way we aggregate the observed data, and construct moment conditions. We describe these moments in greater detail in appendix C. Our method of simulated moments procedure delivers the model parameters that minimize a GMM criterion function, which we also describe in Appendix C.

Appendix C: for online publication: moment conditions and asymptotic distribution of parameter estimates

Recall that we estimate the parameters of our model in two steps. In the first step, we estimate the vector χ , the set of parameters than can be estimated without 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 $t \in \{t_0, \dots, t_T\} = \{1996, 1998, 2000, 2002, 2004, 2006, 2008, 2010\}$, we match median assets for $Q_A = 5$ permanent income quintiles in $P = 5$ birth year cohorts.¹⁶ 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 q th 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

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

The preceding equation can be rewritten as a moment condition (Manski [48], Powell [63] and Buchinsky [13]). In particular, applying the indicator function produces

$$E(1\{a_{it} \leq a_{pqt}(\Delta_0, \chi_0)\} - 1/2 \mid p, q, t, \text{individual } i \text{ observed at } t) = 0. \quad (20)$$

Letting \mathcal{I}_q denote the values contained in the q th permanent income quintile, we can convert this conditional moment equation into an unconditional one (e.g., Chamberlain [15]):

$$E \left([1\{a_{it} \leq a_{pqt}(\Delta_0, \chi_0)\} - 1/2] \times 1\{p_i = p\} \times 1\{I_i \in \mathcal{I}_q\} \times 1\{\text{individual } i \text{ observed at } t\} \mid t \right) = 0 \quad (21)$$

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

We also include several moment conditions relating to medical expenses. Recall that within the model medical expenses are chosen annually and are forward-looking (i.e., for the calendar year in which they are chosen). In contrast, medical expenditures in the AHEAD are averages of spending over the preceding two years. To reconcile the two measures, we first simulate medical expenses at an annual frequency, take two-year averages, and move the resulting averages back one year, to produce a

¹⁶Because we do not allow for macro shocks, in any given cohort t is used only to identify the individual's age.

measure of medical expenditures comparable to the ones contained in the AHEAD. This means that the AHEAD measure for medical spending in 2000 will be compared to averages of model-simulated spending for 1998 and 1999. Using lagged values also allows us to account for people who died prior to the most current wave. This too ensures consistency with the AHEAD, which collects end-of-life medical spending data through survivor interviews.

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 $m_{pqt}^{50}(\Delta, \chi)$ 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:

$$E\left([1\{m_{it} \leq m_{pqt}^{50}(\Delta_0, \chi_0)\} - 0.5] \times 1\{p_i = p\} \times 1\{I_i \in \mathcal{I}_q\} \times 1\{\text{individual } i \text{ observed at } t\} \mid t\right) = 0 \quad (22)$$

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

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 $m_{pqt}^{90}(\Delta, \chi)$ denote the model-predicted 90th percentile, we have the following moment condition:

$$E\left([1\{m_{it} \leq m_{pqt}^{90}(\Delta_0, \chi_0)\} - 0.9] \times 1\{p_i = p\} \times 1\{I_i \in \mathcal{I}_q\} \times 1\{\text{individual } i \text{ observed at } t\} \mid t\right) = 0 \quad (23)$$

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

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

$$R_{it} = \ln(m_{it}) - \overline{\ln m_{pqt}},$$

$$\overline{\ln m_{pqt}} = E(\ln(m_{it}) \mid p_i = p, q_i = q, t)$$

and define the standard deviation σ_{pqt} as

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

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 j th 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

$$\begin{aligned} E \left(\left[\frac{R_{it} R_{i,t-1}}{\sigma_{pqt} \sigma_{pq,t-1}} - AC_{pqt1}(\Delta_0, \chi_0) \right] \times 1\{p_i = p\} \times 1\{I_i \in \mathcal{I}_q\} \right. \\ \left. \times 1\{\text{individual } i \text{ observed at } t \text{ \& } t-1\} \middle| t \right) = 0. \end{aligned} \quad (24)$$

The corresponding moment condition for the second autocorrelation is

$$\begin{aligned} E \left(\left[\frac{R_{it} R_{i,t-2}}{\sigma_{pqt} \sigma_{pq,t-2}} - AC_{pqt2}(\Delta_0, \chi_0) \right] \times 1\{p_i = p\} \times 1\{I_i \in \mathcal{I}_q\} \right. \\ \left. \times 1\{\text{individual } i \text{ observed at } t \text{ \& } t-2\} \middle| t \right) = 0. \end{aligned} \quad (25)$$

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 $\bar{u}_{pqt}(\Delta, \chi)$ 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

$$\begin{aligned} E \left([u_{it} - \bar{u}_{pqt}(\Delta_0, \chi_0)] \times 1\{p_i = p\} \times 1\{I_i \in \mathcal{I}_q\} \right. \\ \left. \times 1\{\text{individual } i \text{ observed at } t\} \middle| t \right) = 0 \end{aligned} \quad (26)$$

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

To summarize, the moment conditions used to estimate model with endogenous medical expenses consist of: the moments for asset medians described by equation (21); the moments for median medical expenses described by equation (22); the moments for the 90th percentile of medical expenses described by equation (23); the moments for the autocorrelations of logged medical expenses described by equations (24) and (25); and the moments for the Medicaid utilization rates described by equation (26). 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 $\hat{\varphi}_I(\cdot)$ denote its sample analog. Letting $\widehat{\mathbf{W}}_I$ denote a $J \times J$ weighting matrix, the MSM estimator $\hat{\Delta}$ is given by

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

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 [58] and Duffie and Singleton [21], 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;

$$\mathbf{D} = \left. \frac{\partial \varphi(\Delta; \chi_0)}{\partial \Delta'} \right|_{\Delta = \Delta_0} \quad (27)$$

is the $J \times M$ gradient matrix of the population moment vector; and $\mathbf{W} = \operatorname{plim}_{I \rightarrow \infty} \{\widehat{\mathbf{W}}_I\}$. Moreover, Newey [53] shows that if the model is properly specified,

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

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 [1].) We thus use a “diagonal” weighting matrix, as suggested by Pischke [61]. 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.

An additional problem is that in cells with small numbers of observations, a moment condition will occasionally have a variance of zero. In one particular cell of the current specification, every person receives Medicaid. Rather than exclude these cells from the moment criterion, we add a small amount of measurement error to the moment condition, so that the weight on the moment (the inverse of the variance) is large but finite.

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 (21). This means that we cannot consistently estimate \mathbf{D} as the numerical derivative of $\hat{\varphi}_I(\cdot)$. 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 [58], Newey and McFadden [54] (section 7), and Powell [63].

To find \mathbf{D} , it is helpful to rewrite equation (21) as

$$\Pr \left(p_i = p \ \& \ I_i \in \mathcal{I}_q \ \& \ \text{individual } i \text{ observed at } t \right) \times \left[\int_{-\infty}^{a_{pqt}(\Delta_0, \chi_0)} f(a_{it} \mid p, I_i \in \mathcal{I}_q, t) \, da_{it} - \frac{1}{2} \right] = 0. \quad (28)$$

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

$$\Pr \left(p_i = p \ \& \ I_i \in \mathcal{I}_q \ \& \ \text{individual } i \text{ observed at } t \right) \times f \left(a_{pqt} \mid p, I_i \in \mathcal{I}_q, t \right) \times \frac{\partial a_{pqt}(\Delta_0; \chi_0)}{\partial \Delta'}. \quad (29)$$

In practice, we find $f(a_{pqt} \mid p, q, t)$, the conditional p.d.f. of assets evaluated at the median a_{pqt} , with a kernel density estimator written by Koning [41]. The gradients for equations (22) and (23) are found in a similar fashion.

Appendix D: for online publication: demographic transition probabilities in the 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 $i \in \{1, 2, 3\}$, $j \in \{0, 1, 2, 3\}$,

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

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.

Appendix E: For online publication: the PSID data and our tax calculations

The lifetime contribution towards Medicaid is calculated using data on household federal tax payments from the PSID. Our calculations require two steps. The first one creates a PSID sample that is comparable to the AHEAD sample. The second

step computes the present discounted value of lifetime taxes for each individual and aggregates it by PI quintile, gender, and health status.

To generate a sample from the PSID that matches that from the AHEAD as closely as possible, we use only individuals that by 1996 are single, make no significant labor income, and are aged 70 to 79. In the AHEAD sample the cohort is aged 72 to 76 but for sample size reasons in the PSID we increase the window from 5 years to 10 years. This leaves a sample of 112 individuals, who are then sorted by permanent income into income quintiles as is done with the AHEAD data.

	AHEAD data		PSID data	
	Number	%	Number	%
Men	138	19.4	19	17.0
Women	573	80.6	93	83.0
Good Health	433	60.9	72	64.3
Bad Health	258	36.3	40	35.7
Nursing Home	20	2.8	0	0.0
Total Observations	711	100	112	100

Table A3: Sample size comparison, AHEAD versus PSID, 1996-2010.

PI Quintile	AHEAD data	PSID data
Bottom	4,830	4,530
Fourth	8,900	8,960
Third	12,550	11,920
Second	16,930	16,970
Top	32,250	31,160
Overall Average	15,710	15,880

Table A4: Annuity income comparison, AHEAD versus PSID, 1996-2010.

To compute taxes, we start by computing permanent income, which is the average

annuity income for each person, where annuity income is calculated as the sum of Social Security, VA Pensions, non-VA Pensions, and Annuities. To match the AHEAD data this is calculated for the years the individual remained alive in 1996, 1998, 2000, 2002, 2004, 2008, and 2010.

Table A4 compares mean annuity income for each income quintile in the PSID sample and the AHEAD sample and shows that they match closely. After being sorted by income quintiles, the PDV of total household federal taxes (value in 1995, measured in 2005 dollars) is calculated for each income quintile-gender group g , as follows:

$$PDV(taxes, g) = \frac{\sum_{t=1967}^{2015} w(g, t) \frac{1}{I(g)} (\sum_{i \in g} tax(i, t) \prod_{j=t+1}^{2015} (1 + r(j)))}{(\prod_{z=1995}^{2015} (1 + r(z))) \cdot (\prod_{q=2005}^{2015} (1 + i(q)))}$$

where $w(g, t)$ is the probability that a member of group g is alive at time t , conditional on being alive in 1967. The mortality rates behind $w(g, t)$ are taken from McClellan and Skinner (2006) until age 70 and are then updated using data from the US Life Tables for 2009. $I(g)$ is the number of people in group g , $tax(i, t)$ is the household federal taxes of individual i in year t , $r(j)$ is the nominal interest rate in year j , and $i(j)$ is the inflation rate. Since after 1990 the PSID no longer reports the value of taxes paid, we assume that tax payments after that year equal those paid in 1990, inflation-adjusted. We also assume a 3% real interest rate. We sum across all individuals to calculate the aggregate PDV of federal taxes. Given the total taxes paid for each group, we need to determine what fraction of these taxes was related to Medicaid.

To determine the average Medicaid tax rate necessary to balance the Medicaid budget for this cohort, we sum the present discounted value of Medicaid transfers reported in Table 7 across individuals. The ratio of the present discounted value of Medicaid transfers to the present value of total taxes paid is \aleph , the share of total taxes used to fund Medicaid for the elderly.

Finally, the PDV of contributions to Medicaid for each PI quintile (or gender and health group) is calculated for each group as \aleph multiplied by the PDV of federal taxes for that group.

Appendix F: For online publication: identification and sensitivity to parameter values

In this appendix we consider how changes in key parameters affect the model's implications for outcomes such as savings, out-of-pocket medical spending, and Medicaid reciprocity. We change one parameter at a time, holding all other parameters

at their baseline values. Table A5 shows how the parameter changes affect the asset, out-of-pocket medical spending, and Medicaid reciprocity moments, as well as the total GMM criterion (the sum of all the moments). Figures 11-15 show how the parameter changes affect the life-cycle profiles of assets, out-of-pocket medical spending, Medicaid reciprocity, and non-medical consumption. This appendix also includes Figure 16, which shows the same profiles for the version of the model where medical spending is exogenous.

The top row of Table A5 shows the moment contributions for our baseline model. The second row shows the moment contributions that result when we reduce the consumption curvature parameter ν by 10%. This specification fits the data much worse: the GMM criterion in the baseline model is 1,217, whereas it is 3,513 when we reduce ν by 10%. Figure 11 reveals that this specification produces much lower medical spending and Medicaid reciprocity, and Table A5 shows that this leads to a much worse model fit.

Decreasing the curvature parameter ω by 10% leads the model to over-predict medical spending and Medicaid reciprocity. Reducing the discount factor β by 10% leads to much more rapid asset decumulation, which is not consistent with the data. The next two rows of Table A5 show the effects of changing the bequest motive parameters, that is the marginal propensity to consume out of wealth in the final period before certain death (MPC) and the threshold where the bequest motive becomes operative. Both of these objects are functions of the bequest parameters θ and k . Changing the bequest parameters does not necessarily make the model fit the asset moments less well, but it does make the model fit the medical spending and Medicaid reciprocity moments less well. Next, we decrease the utility floor and the Medicaid income threshold by 10%. Reducing these parameters worsens the model's fit of the Medicaid moments. Finally, reducing either the mean or the variance of the medical needs shocks causes the model to fit the data less well.

Specification	Medical Spending				Total
	Asset Quantiles	Quantiles	Autocor- relations	Medicaid Reciprocity	
Baseline	166	543	174	335	1,217
ν decreased 10%	202	2,355	189	767	3,513
ω decreased 10%	424	1,853	252	1,207	3,736
β decreased 10%	213	696	169	316	1,394
MPC decreased 10%	179	541	174	351	1,246
Bequest threshold doubled	146	718	182	372	1,418
Utility floors decreased 10%	175	532	201	364	1,271
Medicaid income threshold decreased 10%	165	595	150	345	1,254
Medical shocks decreased 10%	175	580	174	378	1,308
Variance of shocks decreased 10%	163	581	173	321	1,238

Table A5: Effects of Parameter Changes on GMM Criteria

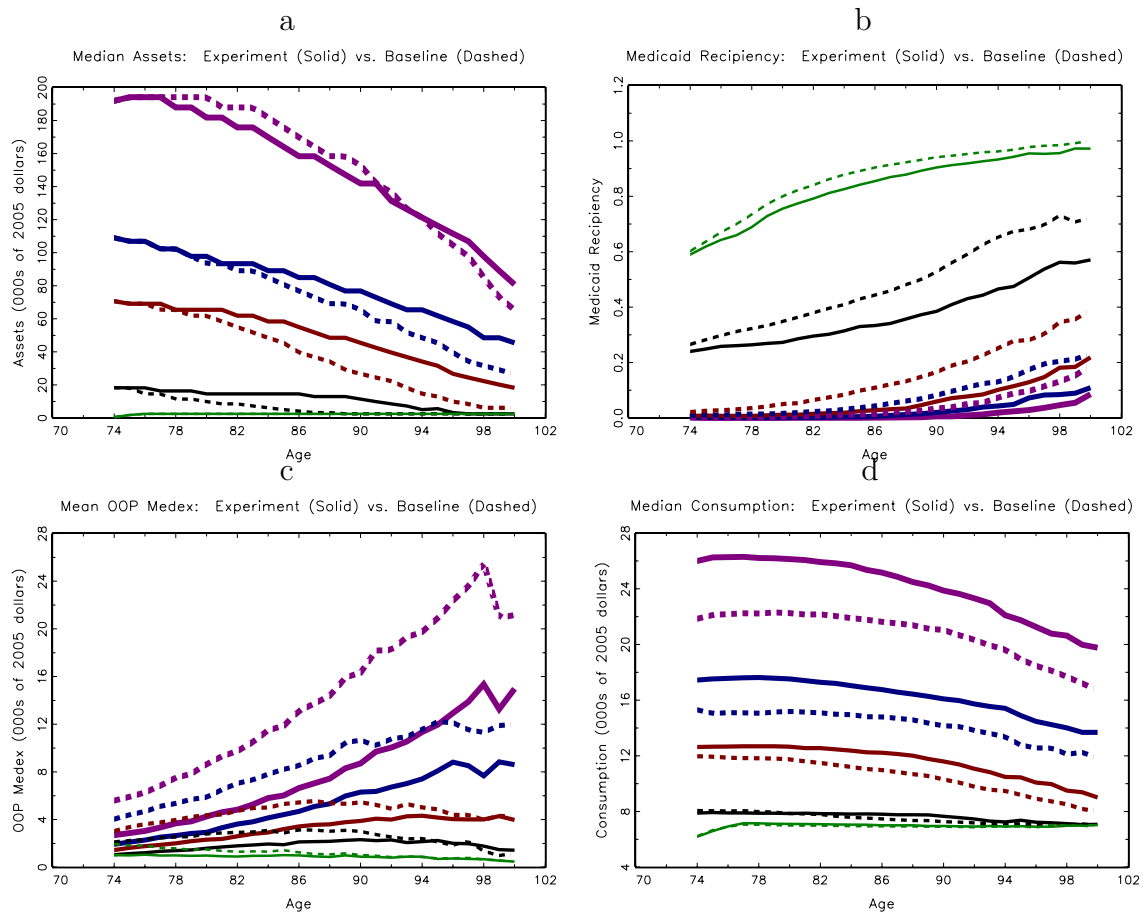


Figure 11: Assets (panel a), Medicaid reciprocity (panel b), out-of-pocket medical spending (panel c), and non-medical consumption (d) by age and permanent income. Dashed line: benchmark, solid line: ν decreased 10%.

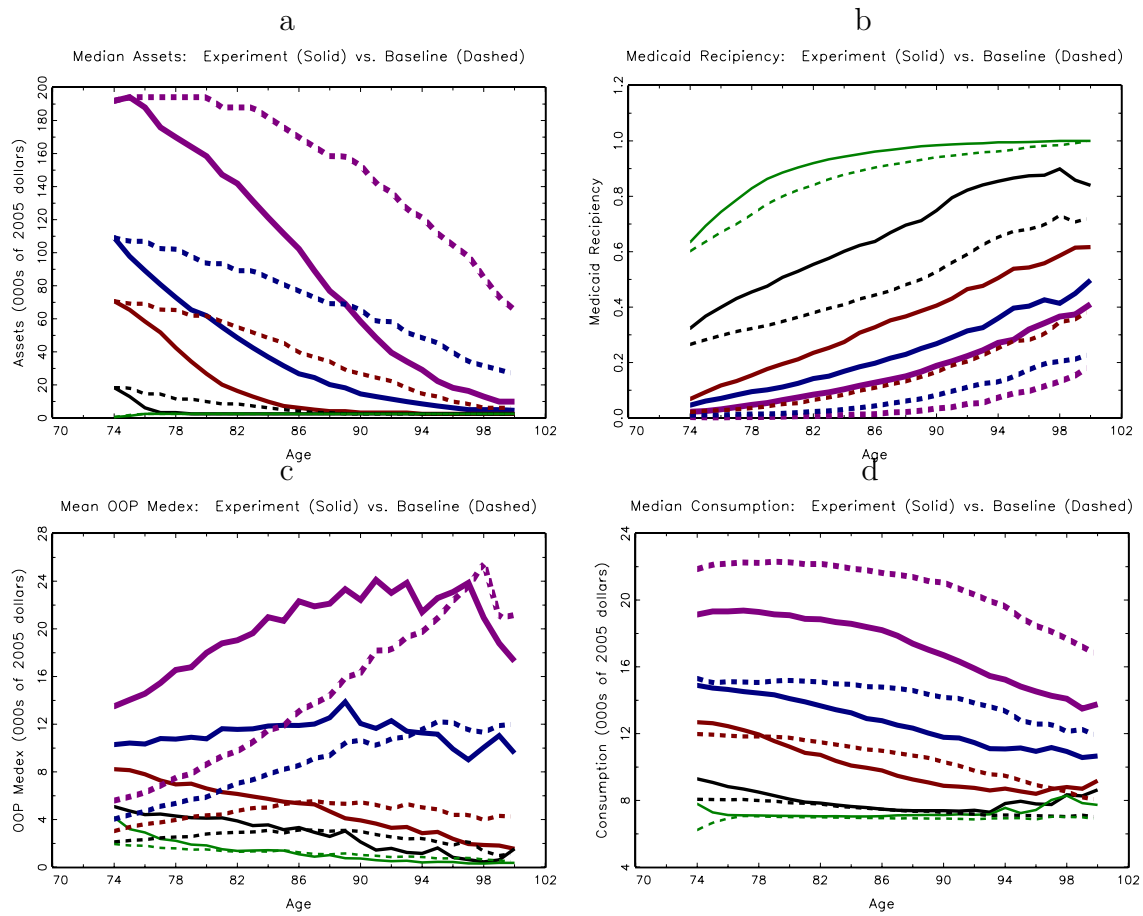


Figure 12: Assets (panel a), Medicaid reciprocity (panel b), out-of-pocket medical spending (panel c), and non-medical consumption (d) by age and permanent income. Dashed line: benchmark, solid line: ω decreased 10%.

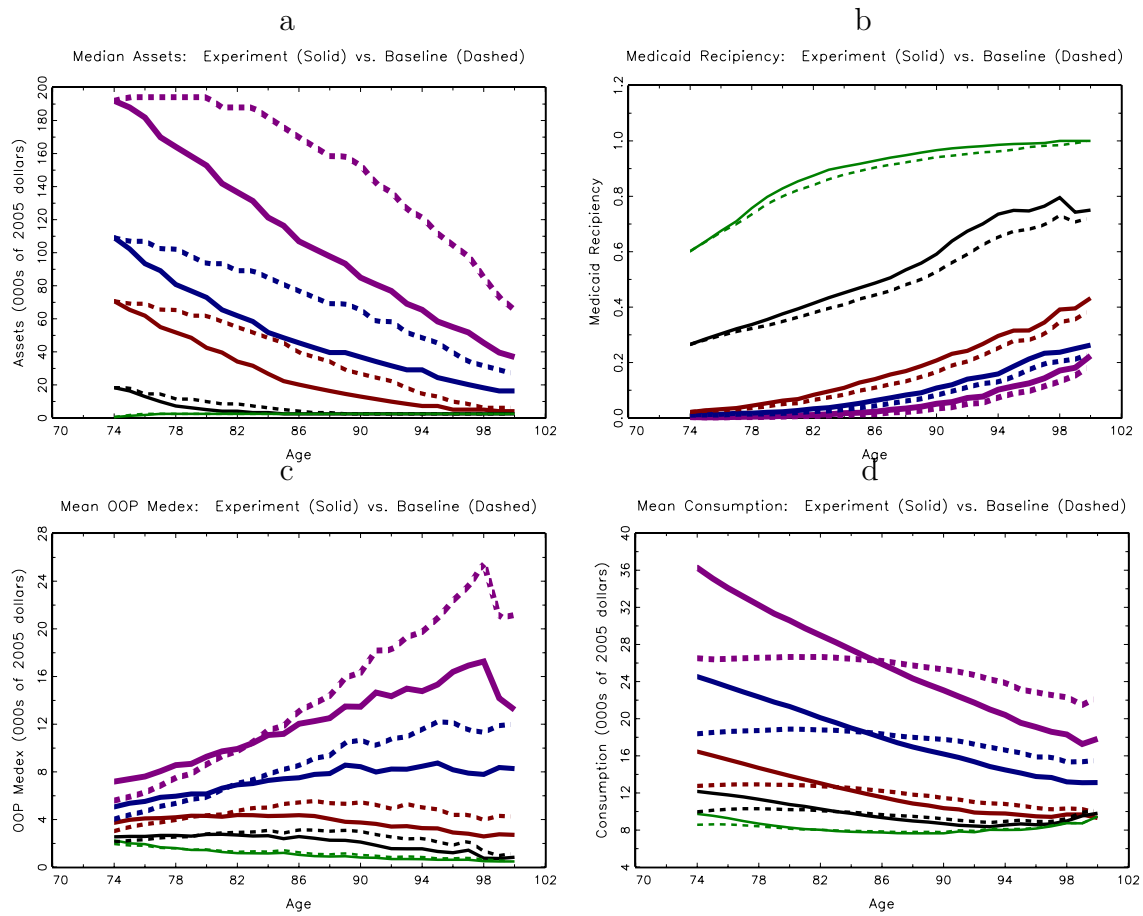


Figure 13: Assets (panel a), Medicaid reciprocity (panel b), out-of-pocket medical spending (panel c), and non-medical consumption (d) by age and permanent income. Dashed line: benchmark, solid line: β decreased 10%.

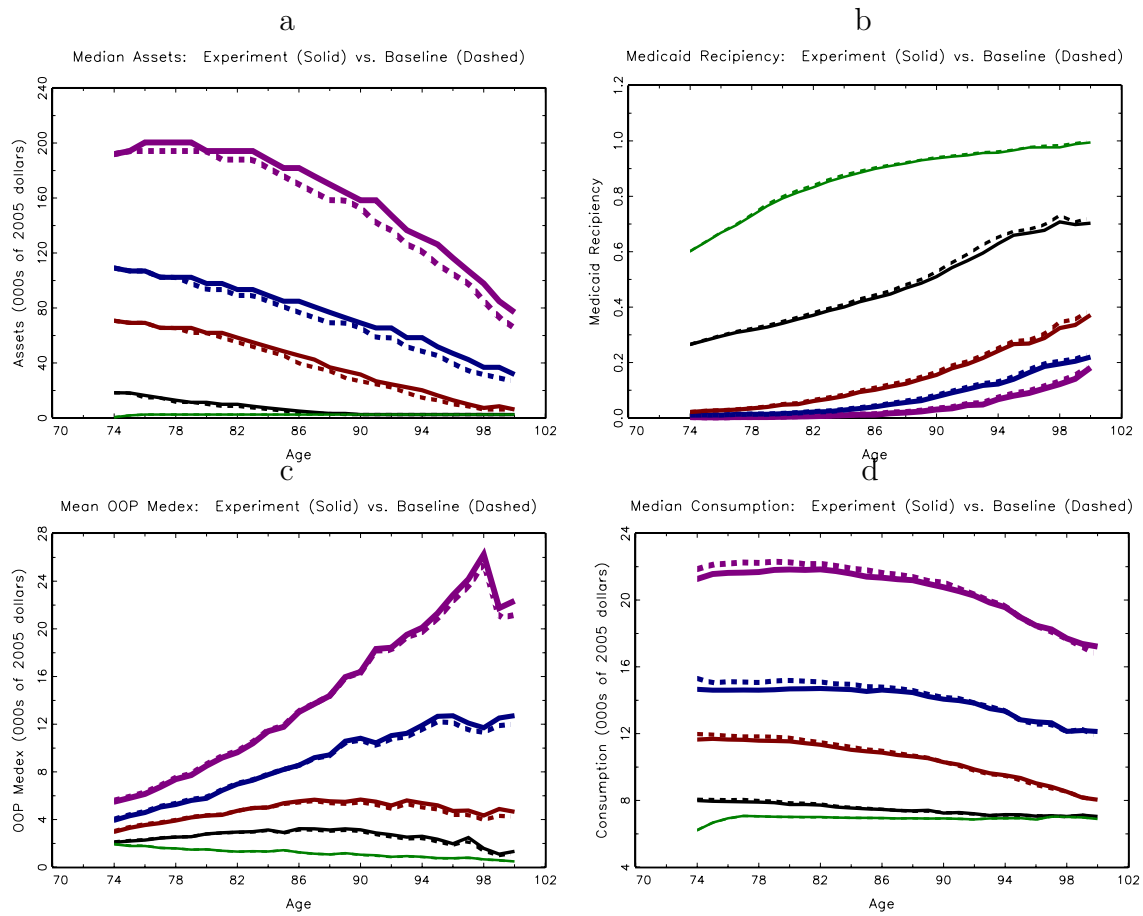


Figure 14: Assets (panel a), Medicaid reciprocity (panel b), out-of-pocket medical spending (panel c), and non-medical consumption (d) by age and permanent income. Dashed line: benchmark, solid line: MPC decreased 10%.

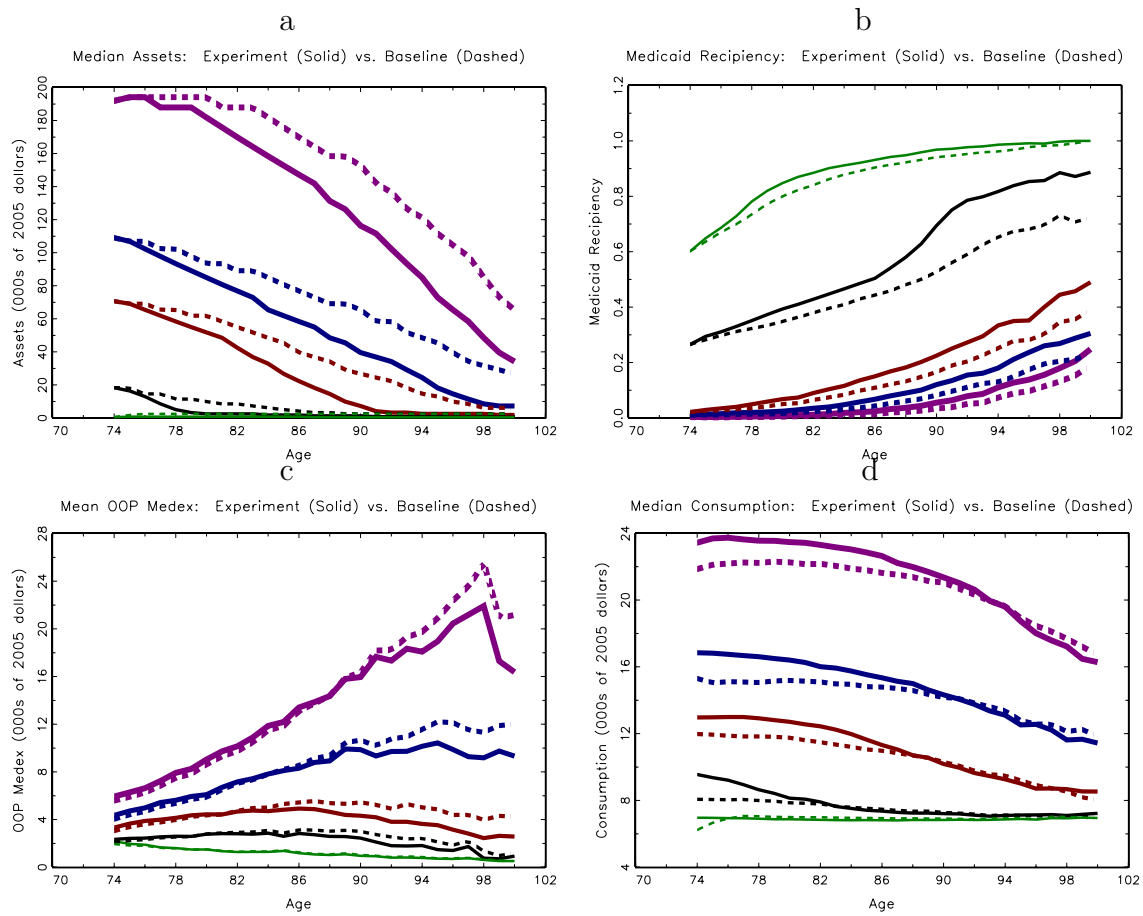


Figure 15: Assets (panel a), Medicaid reciprocity (panel b), out-of-pocket medical spending (panel c), and non-medical consumption (d) by age and permanent income. Dashed line: benchmark, solid line: Bequest threshold doubled.

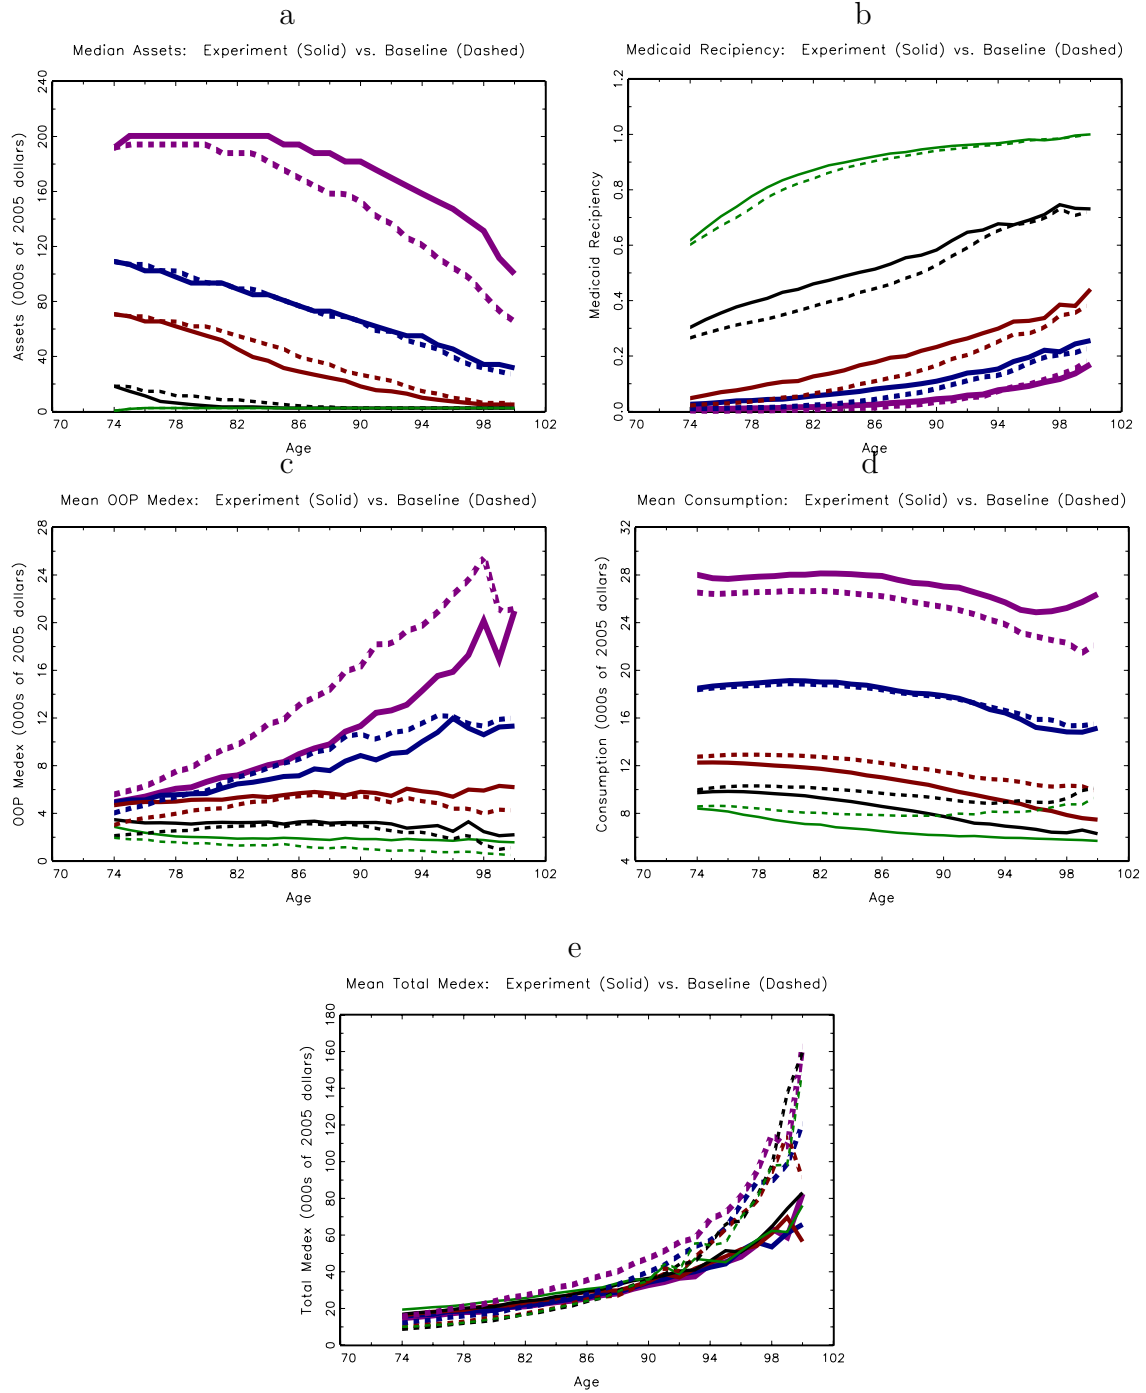


Figure 16: Assets (panel a), Medicaid reciprocity (panel b), out-of-pocket medical spending (panel c), and non-medical consumption (d) by age and permanent income. Dashed line: benchmark, solid line: exogenous medical spending.

Appendix G: The CDF of assets predicted by the model compared with the one in the data

To gauge model fit of the asset data, figure 17 presents both the CDF of assets in the AHEAD data and the CDF of assets predicted by the model. The model prediction is the solid line and the AHEAD data are the dotted line. The CDF for model predicted assets looks like a step function since we discretize the asset grid. Overall, the fit of the model is good. The model underpredicts the probability of having low assets slightly. For example, 47% of all AHEAD households have assets below \$30,000, whereas the model predicts that 41% of all households have assets below \$30,000. At higher asset levels the fit of the model is better. For example, 67% of all households have assets below \$100,000, whereas the model predicts that 66% of all households have assets below \$30,000.




Figure 17: Cumulative distribution functions of assets: model (solid line), AHEAD data (dotted line).

Appendix H: Robustness of compensating variations to parameter changes

For the purposes of better understanding what affects the compensating variations in our estimated model, we now turn to changing one parameter at a time and recomputing the compensating variation for a 10% decrease in Medicaid generosity. It should be noted that, as we have shown in Appendix F, some of these parameter changes (and notably those for risk aversion and the discount factor) lead to drastic deteriorations in the model's fit of the observed data.

The results in Table A6 of this appendix show that realistically small changes in Medicaid generosity and income eligibility generate small changes in the estimated compensating variations, relative to their benchmark values. In contrast, changes in parameter values that significantly worsen the fit of the model can generate significant changes in the compensating variations. It should also be noted that the direction of these changes makes economic sense. If, for instance, we decrease patience, the insurance value that households attribute to Medicaid decreases. Similarly, if we decrease risk aversion, our compensating variations tend to decrease.

	Baseline Model (1)	Initial floor reduced 10% (2)	\bar{Y} (for SSI) = \$6,950 (3)	ν reduced 10% (4)	ω reduced 10% (5)	β reduced 10% (6)
Change in Discounted Lifetime Spending						
Bottom	4,500	3,900	4,400	2,600	33,300	4,600
Fourth	4,000	3,500	4,000	2,100	34,900	4,200
Third	2,900	2,500	2,900	1,300	35,000	3,200
Second	2,200	1,900	2,200	800	33,800	2,500
Top	1,400	1,100	1,400	400	29,500	1,600
Compensating Variation						
Bottom	6,300	6,000	6,400	3,900	22,800	4,000
Fourth	5,000	4,800	5,000	2,400	20,600	3,600
Third	4,400	4,400	4,400	1,100	20,500	3,600
Second	4,100	4,600	4,100	600	22,500	2,400
Top	4,400	4,600	4,400	200	24,700	1,700
Compensating Variation / Change in Spending						
Bottom	1.40	1.54	1.45	1.50	0.68	0.87
Fourth	1.25	1.37	1.25	1.14	0.59	0.86
Third	1.52	1.76	1.52	0.85	0.59	1.13
Second	1.86	2.42	1.86	0.75	0.67	0.96
Top	3.14	4.18	3.14	0.50	0.84	1.06

Table A6: The costs and benefits of increasing Medicaid by 10%, alternative specifications.