MICHIGAN RETIREMENT AND DISABILITY RESEARCH CENTER
UNIVERSITY OF MICHIGAN

Promoting research on retirement, disability, and Social Security policy

How Redistributive Are Public Health Care Schemes? Evidence from Medicare and Medicaid in Old Age

Karolos Arapakis, Eric French, John Bailey Jones, Jeremy McCauley

MRDRC WP 2010-441

UM20-10

Preliminary: Please do not quote

How Redistributive Are Public Health Care Schemes? Evidence from Medicare and Medicaid in Old Age

Karolos Arapakis
University College London

John Bailey Jones
Federal Reserve Bank of Richmond

Eric French
University College London; Institute for Fiscal Studies

Jeremy McCauley
University of Bristol

June 2022

Michigan Retirement and Disability Research Center, University of Michigan, P.O. Box 1248. Ann Arbor, MI 48104, mrdrc.isr.umich.edu, (734) 615-0422

Acknowledgements

The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement and Disability Research Consortium through the University of Michigan Retirement and Disability Research Center Award RDR18000002-02. The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA or any agency of the federal government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States government or any agency thereof.

Preliminary: Please do not quote

How Redistributive Are Public Health Care Schemes? Evidence from Medicare and Medicaid in Old Age

Abstract

Most health care for the U.S. population 65 and older is publicly provided through Medicare and Medicaid. Despite the massive expenditures of these systems, little is known about how redistributive they are. Using data from the Health and Retirement Study matched to administrative Medicare, Medicaid, and Social Security earnings records, we estimate the distribution of lifetime Medicare and Medicaid benefits received and the distribution of lifetime taxes paid to finance these benefits. For the cohort who turned 65 between 1999 and 2004, we find that benefits are greater among those with high income, in large part because they live longer. Nonetheless, high-income people pay more in the way of taxes. Middle-income households gain the most from these programs as these people live long yet pay modest taxes. All income groups gain from these programs: This cohort's lifetime tax contribution did not cover the medical benefits it received. This deficit is paid by younger cohorts.

Citation

Arapakis, Karolos, Eric French, John Bailey Jones, Jeremy McCauley. 2022. "How Redistributive Are Public Health Care Schemes? Evidence from Medicare and Medicaid in Old Age." Ann Arbor, MI. University of Michigan Retirement and Disability Research Center (MRDRC) Working Paper; MRDRC WP 2022-441
https://mrdrc.isr.umich.edu/publications/papers/pdf/wp441.pdf

Authors' acknowledgements

Preliminary, please do not quote. We thank Gary Engelhardt for help with measuring tax payments, Bryan Tysinger, and Courtney Coile. Additional funding from the Economic and Social Research Council (Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (ES/T014334/1) and Grant Inequality and the insurance value of transfers across the life cycle (ES/P001831/1)) for this work is also gratefully acknowledged. Correspondence to eric.french.econ@gmail.com, jeremymccauley@gmail.com.

1 Introduction

The great majority of American households aged 65 and older receive health care benefits from Medicare, a mostly free single payer program that covers most health care services. Furthermore, many older households also receive benefits from Medicaid, which covers the health care costs of those with limited financial resources. Both programs are funded out of insurance premia and tax revenues. This paper estimates health care benefits received from these programs and taxes and insurance premia paid to finance these benefits, which contributes to the debate on reforming healthcare. Which income groups receive the most health care benefits? How redistributive are the benefits received and the taxes paid to finance those benefits? These are important questions to answer before reforming the programs currently in place. These questions are also of importance as some US policy makers advocate a single payer health care scheme. In this paper we seek to fill this gap.

We focus on households aged 65 and older, a group who is responsible for 33%[1] of all medical care received, whose expenditures are on average 2.6 times higher than the national average, and who rely much more on public funding than the average household.

Using data from the Health and Retirement Study matched to administrative Medicare, Medicaid, and Social Security earnings records, we estimate the distribution of lifetime medical benefits received and the distribution of lifetime taxes paid for that care among the cohort who turned 65 in 1999-2004. We forecast future health, longevity, and health care benefits received, allowing us to infer lifetime Medicare and Medicaid benefits received by this cohort. We also calculate total income, state and federal taxes paid at each age for this cohort. To do so, we calculate federal and state tax payment for each household in each year, then multiply these by the shares of aggregate federal and state taxes in that year that go towards paying Medicare and Medicaid for those over 65.

We estimate models of total medical spending, in addition to the share of medical spending paid by Medicare and Medicaid. We impute private insurance and other payments using data from the Medical Expenditure Panel Survey (MEPS). Thus our framework allows us to not only consider the amount of redistribution in the current system, but also how the amount of redistribution would change as a function of changes in reimbursement rate policy, for example.

For this cohort, we find that Medicare and Medicaid benefits received are 1.7 times greater

[1] See De Nardi et al. (2016) and Exhibit 1 of Lassman et al. (2014).

among those in the top lifetime income quintile than among those in the bottom quintile, in large part because they live longer. Nonetheless, high income people pay more in the way of taxes and as a result there are net transfers to those at the bottom of the income distribution. For example, we calculate that households in the top income quintile contribute 7.5 times more than those in the bottom quintile through taxes and Medicare insurance contributions. Those at the top of the income distribution contribute $248,000 on average to Medicare and Medicaid, and receive $401,000 in benefits over their lives. Those at the bottom of the income distribution contribute $33,000 on average to Medicare and Medicaid, and receive $229,000 in benefits over their lives. The largest beneficiaries of Medicare and Medicaid are those in the middle of the income distribution, as these people live long yet pay modest taxes. For example, those in the middle 20% of the income distribution contribute $82,000 on average to Medicare and Medicaid, and receive $337,000 in benefits over their lives.

All income groups within the cohort we study are net beneficiaries from Medicare and Medicaid. On average this cohort's lifetime tax contribution did not cover the medical benefits it received. Put differently, the key source of redistribution is from younger cohorts to the cohort we study. Thus, for the cohort we study, cross-cohort redistribution (from young to old) is greater than within cohort redistribution (from rich to poor).

2 Literature Review

This paper 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 government pension and other cash transfer schemes, there is less known about how health care systems redistribute resources.

A few papers (Bhattacharya and Lakdawalla 2006; McClellan and Skinner 2006; and Rettenmaier 2012) study Medicare progressivity. Most of these papers find that Medicare does redistribute resources from the lifetime income rich to the lifetime income poor, since the those with high income pay more in taxes. The amount of redistribution is lower than what one might think, however, because high income people consume more resources at each age and because they tend to live longer than their lower income counterparts. Furthermore, these estimates are sensitive to how lifetime income is measured. For example, when using local area income as a measure of lifetime income, McClellan and Skinner (2006) find the rich use more health care resources and are more likely to live until very old. In contrast, it does not appear that the high income use more resources per year when using either education (Bhattacharya and Lakdawalla 2006), or the individual's income (Rettenmaier 2012) measured over a long time interval. None of these papers consider the redistribution due to Medicaid in old age. One might expect Medicaid to be very progressive, since to be eligible an individual must have low assets and low income net of medical spending. While De Nardi, French, and Jones (2016) do find that the program does redistribute from rich to poor, they also find that even the highest income often receive benefits if they live long enough to run down their assets paying for expensive medical care. To the best of our knowledge this is the first paper to comprehensively examine how Medicaid and Medicare health care transfers are jointly distributed across income groups. While Auerbach et al. (2017) considers the distribution of benefits received from these programs, as well as from other programs such as Supplementary Security Income (SSI) and Social Security, they do not consider the distribution of taxes paid.

We also estimate the distribution of the taxes used to finance Medicare and Medicaid. Although Medicare Part A is funded by a dedicated tax, Medicaid and some Medicare health care benefits are not financed using a specific tax, but by general government revenue, making it difficult to determine how redistributive "Medicare and Medicaid taxes" are. For services that do not have a dedicated tax, we adapt the approach of McClellan and Skinner (2006), and De Nardi, French, and Jones (2016) and assume that the Medicare and Medicaid tax burden is proportional to the general tax burden.

This paper also contributes to a literature on cross-country comparisons of health care funding schemes. Two important papers are Wagstaff and Van Doorslaer (1992) and Wagstaff et al. (1999), who show that tax-financed schemes provide more redistribution than other schemes, such as private insurance.

3 An Accounting Framework for Measuring Redistribution

3.1 The Financing of Medicare and Medicaid

Like Social Security, Medicare and Medicaid are financed largely on a pay as you go basis. Since its inception in 1966, Medicare Part A, which provides hospital care, has been financed by a payroll tax on wages. The Part A payroll tax rate has increased over time and is currently 2.9%, half levied on the employee and half on the employer. Medicare Part B, which provides physician and outpatient services, is financed partly by beneficiary premia that cover one-quarter of expenditures

(Cubanski et al. 2015, Figure 29), with the rest primarily financed by general federal tax revenues. Because payroll and income taxes are increasing functions of earnings and income, respectively, an individual's overall Medicare contributions increase with their lifetime income. Medicaid, which was established at the same time as Medicare, provides income- and asset-tested health care benefits to families with young children, those who are disabled, or those who are over 65. We focus on the provisions for those over 65. Unlike Medicare, Medicaid is primarily financed by general federal tax revenues, with the remainder being paid by state tax revenues. As such, an individual's lifetime contributions to Medicaid should increase with their lifetime income.

While Medicare eligibility is nearly universal for those 65 and older, Medicaid eligibility is targeted towards the poor, which means that Medicaid benefits are falling in income. Eligiblity is not the only factor affecting the distribution of benefits, however. Lifetime Medicare and Medicaid benefits also vary because medical care usage differs across the income distribution at any given age and because lifespans differ across the income distribution. It is well known that those at the top of the income distribution live longer than their poorer counterparts, as shown both in the HRS (De Nardi, French, and Jones 2016; Pijoan-Mas and Ríos-Rull 2014) and in administrative tax records (Waldron 2007; Chetty et al. (2016)). But there is considerable debate about whether health care expenditures are positively or negatively correlated with lifetime income. These discrepancies seem to depend heavily on the time period, the measure of lifetime income, as well as for the measure of medical spending (see the aforementioned debate between McClellan and Skinner 2006 and Bhattacharya and Lakdawalla 2006). For this reason it is particularly important to have full lifetime income measures and detailed medical spending data on all components of Medicare, including parts A, B, C, and D.

For older cohorts born before the inception of Medicare and Medicaid, lifetime tax payments did not come close to financing their lifetime benefits. Individuals who were 65 or over in 1966 were clearly better off with the programs, since they contributed little to the programs but received significant benefits. More generally, previous research on Medicare has found that the rapid growth in expenditures over time, coupled with the largely pay as you go nature of its financing, has likely led to substantial intergenerational transfers to older beneficiaries (Vogel 1988).

3.2 Redistribution Within and Across Cohorts

For any individual household, the redistribution attributable to Medicare and Medicaid is the difference between present value of the Medicare- and Medicaid- related benefits that the household received over its lifetime and the present value of the Medicare- and Medicaid- related taxes that it paid. We use the accounting framework employed by McClellan and Skinner (2006) to decompose this sum into within-cohort and across-cohort components,[2] which allows us to measure the extent of cross-cohort redistribution among the multiple cohorts appearing in our sample.

Consider household i of birth cohort c. Our fundamental accounting identity is

$$\begin{aligned} \text{PV(Benefits)}_{ci} - \text{PV(Taxes)}_{ci} \quad &= \quad \text{PV(Within-Cohort Redistribution)}_{ci} \\ &\quad + \text{PV(Cross-Cohort Redistribution)}_{c}. \end{aligned} \tag{1}$$

We make two definitional assumptions. The first is that the amount of cross-cohort redistribution is the same for all members of a given cohort. The second is that within each cohort, redistribution sums to zero:

$$\sum_{i=1}^{N_c} \text{PV(Within-Cohort Redistribution)}_{ci} = 0,$$

where we use N_c to denote the size of cohort c. Inserting this result into equation (1) and summing, we have:

$$\text{PV(Cross-Cohort Redistribution)}_{c} = \frac{1}{N_c} \sum_{i=1}^{N_c} \left(\text{PV(Benefits)}_{ci} - \text{PV(Taxes)}_{ci} \right). \tag{2}$$

Let C denote the total number of cohorts. For each cohort c, $c = 1, 2, ..., C$, we proceed as follows. First, we measure PV(Benefits)_{ci} and PV(Taxes)_{ci} for every household ($i = 1, 2, ..., N_c$) in cohort c, using the procedures described below. Next, we use equation (2) to find $\text{PV(Cross-Cohort Redistribution)}_{c}$ for each cohort c. Finally, we insert $\text{PV(Cross-Cohort Redistribution)}_{c}$ back into equation (1), allowing us to calculate $\text{PV(Within-Cohort Redistribution)}_{ci}$ for every household i in cohort c.

3.3 Lifetime Benefits and Taxes

For each household we calculate the present value of Medicare and Medicaid benefits as:

$$\text{PV(Benefits)}_{i} = \lambda_{2014} \sum_{t=1965}^{2050} S_{it} M_{it} \left(\prod_{s=1965}^{t} (1 + r_s)^{-1} \right), \tag{3}$$

[2] An alternative approach is to calculate the internal rate of return, as used by Bhattacharya and Lakdawalla (2006). In their framework, however, the internal rate of return is selected such that the amount of cross-cohort redistribution is 0, ruling out any assessment of cross-cohort redistribution.

where: M_{it} is household i's Medicare and Medicaid benefits in year t (which may equal 0); r_s is the interest rate in year s; and $\lambda_{2014} \equiv \prod_{s=1965}^{2014}(1+r_s)$ rescales the total to 2014 dollars. For the years 1965-2016, the variable S_{it} is a 0-1 indicator equal to 1 if any member of household i in our data is alive at time t. From 2019 forward, we predict survival, which depends on each household member's age, gender and income, so that S_{it} is continuous. We describe our approach to estimating mortality probabilities below.

Turning to taxes, for each household i we calculate the present value of their Medicare Part A payroll taxes, their contributions to Medicare Parts B, C, and D and Medicaid from their federal and state income taxes, and their Medicare Part B insurance premia:

$$\text{PV(Taxes)}_i = \lambda_{2014} \sum_{t=1965}^{2050} \left[\left(\tau_t^m \min\{y_{it}, y_t^{\max}\} \right) + \psi_t^f T_{it}^f + \psi_t^s T_{it}^s + P_{it} \right] \left(\prod_{s=1965}^{t} (1+r_s)^{-1} \right), \tag{4}$$

where: τ_t^m is the Medicare payroll tax rate in year t; y_{it} is household i's earnings; $y_t^{\max}$ is the taxable maximum; ψ_t^f and ψ_t^s are the shares of federal and state income taxes that go towards Medicare and Medicaid; T_{it}^f and T_{it}^s are federal and state income tax payments; and P_{it} is the Medicare Part B insurance premium.

The first term in the square brackets in equation (4) captures the Medicare Part A component of the payroll tax. In earlier years, the payroll tax rate τ_t^m was only 0.7% for salaried and wage workers and 0.35% for the self employed, and the cap on taxable earnings $y_t^{\max}$ was modest: for example, $y_{1971}^{\max} = \$7{,}800$. Since then both the tax rate and the cap have risen, and in 1994 the cap was removed.[3] The second and third terms estimate the household's federal and state income tax payments that are devoted to Medicare Parts B, C and D and Medicaid, based on the fraction of total federal and state expenditures devoted to Medicare Parts B-D and Medicaid in that calendar year (ψ_t^f and ψ_t^s). The fourth term is the Part B premium P_{it}, which is zero prior to age 65 and is zero for Medicaid beneficiaries as well. Prior to 2007, the Part B premium was not income-tested (for example, it was \$381.60 per household member in 1992), but since then higher income households have paid higher premia. As with benefits, tax payments are scaled by the probability that the individual is alive at year t, S_{it}, and as before are normalized to 2014 dollars using the adjustment factor λ_{2014}. For our main results we do not discount.

[3]Because the payroll tax cap is applied to each family member separately, in households with multiple workers equation (4) is not completely accurate.

4 Data and Descriptive Statistics

Our principal data source is HRS survey data matched with administrative Medicare fee-for-service and Medicaid records. We use the MEPS data to impute payments for payors missing in the HRS, such as Medicare Part C, private insurers and other government programs. The result is a version of HRS medical spending data that is representative of all payors. We also match the HRS data to restricted earnings records provided by the Social Security Administration that date back to 1951.

4.1 The HRS

We use data from the HRS, which is a representative biennial survey of the population aged 51 and older and their spouses. In order to focus on the benefits from Medicaid and Medicare for the age-65+ population, we drop households whose head was younger than 65 at the end of the sample period in 2016. Although we focus on the cohort of individuals who turned 65 in 1999-2004, we estimate demographic transitions using the full sample of individuals aged 65 and older. Consistency with our demographic model leads us to drop a small number of households who, for example, are "partnered" or whose partner reports conflicting marital status. Furthermore, we drop households who do not consent to provide their Medicare and Medicaid information or are otherwise unable to be matched with administrative records. This leaves us with 8,777 households comprising 122,878 household-year observations for the full estimation sample. Appendix B documents our sample selection criteria.

The HRS conducts interviews every other year. Households are followed until both members die; attrition for other reasons is low. When a respondent dies, an "exit" interview with a knowledgeable party – usually another family member – is conducted in the next wave. This allows the HRS to collect data on end-of-life medical conditions and spending. Fahle, McGarry, and Skinner (2016) compare the medical spending data from the "core" and exit interviews and show that out-of-pocket spending rises significantly in the last year of life.

The HRS has a variety of health indicators. We assign individuals to the nursing home state if they were in a nursing home at least 120 days since the last interview or if they spent at least 60 days in a nursing home before the next scheduled interview and died before that scheduled interview. We exclude short-term visits: as Friedberg et al. (2015) and Hurd, Michaud, and Rohwedder (2017) document, many nursing home stays last only a few weeks and are associated with lower expenses. We focus instead on the longer and more expensive stays. We assign the remaining individuals a health status of "good" if their self-reported health is excellent, very good or good and a health status of "bad" if their self-reported health is fair or poor.

Although HRS respondents are drawn from the non-institutionalized population when first interviewed, they are tracked and reinterviewed as they enter nursing homes and other institutions. French and Jones (2004) show that by 1999, when our main sample begins, the HRS matches well the aggregate statistics on the share of the older population in a nursing home. (In 1999 roughly 2.8% of men and 6.1% of women in the data had entered nursing homes.)

The HRS collects data on all out-of-pocket medical expenses, including drug costs, hospital stays, nursing home care, doctor visits, dental visits, and outpatient care, including those incurred during the last year of life. To these data we append administrative Medicare and Medicaid data, which we describe next.

4.2 Administrative Medicare and Medicaid Records

The Centers for Medicare and Medicaid Services (CMS) have confidential administrative spending records for Medicare (Parts A, B and D) and Medicaid (including through Medicaid HMOs), that we link to the survey responses of consenting HRS respondents. We have Medicare data for each year between 1991 and 2016. These records include reimbursement amounts for inpatient, skilled nursing facility, home health, and hospice claims made under Medicare Part A, as well as outpatient, carrier (non-institutional medical care providers such as individual or group practitioners, non-hospital labs, and ambulances), and durable-medical-equipment payments made under Medicare Part B. To this total we add drug-related spending made under Medicare Part D, which began in 2006.

As with the Medicare records, we link restricted Medicaid data for those who give permission, allowing us to measure Medicaid spending for each year between 1999 and 2012. The Medicaid files contain information on enrollment, service use, and spending. Appendix C describes the Medicare and Medicaid data in more detail. Linking these data to the HRS results in a broad set of spending measures for the years 1999-2012. These will be the years used in our main analysis. For medical spending outside of this range, we predict medical spending using the methods described below.

4.3 Imputations Using MEPS Data

While the HRS contains accurate measures of out-of-pocket medical spending, and can be linked to Medicare and Medicaid records, it does not contain information on the payments made by Medicare Part C, private and other smaller public insurers (such as the Veterans Administration and state or local health departments). To circumvent this issue, we use data from the 1996-2017 waves of MEPS to impute these payments.

The MEPS is a nationally representative survey of non-institutionalized households. MEPS respondents are interviewed up to 5 times over a 2 year period, forming short panels. We aggregate the data to an annual frequency. MEPS respondents are asked about their (and their spouses') health status, health insurance, and the health care expenditures paid out-of-pocket, by Medicaid, by Medicare, private insurance, and by other public and private sources. The survey responses are matched to medical spending information provided by health care providers. Although the MEPS does not capture certain types of medical expenditures, such as nursing home expenditures, Pashchenko and Porapakkarm (2016) compare MEPS data to the aggregate statistics and show that MEPS captures most types of spending very well.

To impute medical spending not captured in the HRS, we proceed in two steps. First, we use the MEPS data to regress these payments on a set of observable variables found in both datasets. Variables include household income, a fourth order age polynomial, labor force participation status, education, marital status, doctor and hospital visits, race indicators, health measures, out-of-pocket spending and interactions. Second, we impute these expenses in the HRS data using a conditional mean-matching procedure. Applying the MEPS regression coefficients to the HRS data yields predicted values for each HRS household, to which we add residuals drawn from MEPS households with a similar levels of predicted spending. We describe our approach in more detail in Appendix D.

4.4 Total Medical Spending

Total medical spending is calculated as the sum of the measures described in Sections 4.1-4.3. This measure includes all types of medical care and all payors. To the best of our knowledge, this is the first paper that includes all types of coverage over a long panel.[4] All payors are measured on an annual basis, except out-of-pocket expenses which we annualize..

Table 1 summarizes the medical spending data contained in our HRS sample. The first two columns of the table show total spending by all payors for older households, sorted by the permanent income measure we describe in the next section. The columns show that medical spending is about \$3,000 higher at the bottom income quintile than at the top, \$19,400 as opposed to \$16,400. This reflects in large part the tendency of poorer households to be in worse health.

[4]Becacuse we are imputing insurance payments from MEPS, we are missing private payments for long-term care. We believe our measure captures all other payors and types of care.

The remaining 4 columns of Table 1 decompose this total by payor. Households in the top four income quintiles pay between 24% and 26% out of pocket, but those at the bottom pay only 13.4%. Most of this difference is accounted for by Medicaid, which is concentrated in the bottom income quintile.

In Appendix E we show that our medical spending measures against the MCBS. French et al. (2017) find that out-of-pocket medical spending in the HRS, MCBS, and MEPS match up well. We extend their exercise by comparing Medicare and Medicaid payments in the HRS restricted data to Medicare and Medicaid payments in the MCBS.

	Total Spending		Percent paid by				
PI Quintile	Average Amount	Pct. of Total	Medicare	Medicaid	Out-of-Pocket	Private	Other
All	17,077	100.0	62.9	7.5	22.6	3.5	3.4
Top	16,394	19.3	68.4	0.6	24.3	3.0	3.7
4^{th}	16,139	19.6	63.0	3.1	26.4	3.9	3.5
3^{rd}	16,115	19.5	62.1	2.8	25.8	5.3	4.0
2^{nd}	17,594	20.9	59.9	9.3	23.5	4.0	3.3
Bottom	19,416	20.7	61.5	20.8	13.4	1.4	2.9

Notes: This uses HRS data linked to Medicare and Medicaid records. Private and other imputed using MEPS. Data are from 1999-2012. Sample is cohort whose household head was 65 in 1999-2004.

Table 1: Household Medical Spending by PI Quintile and Payor

4.5 Social Security Earnings Records

In order to measure lifetime income, we link the HRS data to the Social Security earnings records of consenting respondents. In particular, we use the Respondent Cross-Year Summary Earnings dataset, which provides earnings records covering the years 1951-2016 for respondents and, when applicable, their former spouses. Accounting for former spouses is important because in our data many women with little or no lifetime earnings had husbands with high earnings. Given that a large share of our sample consists of older widows, omitting their deceased husbands' earnings would significantly understate their lifetime financial resources.

Because the SSA earnings records are used to calculate Social Security benefits, they are high quality data with little measurement error or attrition. However, earnings are top-coded for much of the sample period, as earnings above a taxable maximum are not subject to payroll taxes. We have non-top-coded earnings information for employed salaried and wage workers as far back as

1978, although self-employed earnings data were still top-coded until 1994. When present, these top-coding issues are not trivial: 30% of people had top-coded earnings records in 1966 and 22% had top-coded earnings in 1977.

For earnings that are top-coded or are "masked" (i.e., expressed as an interval rather than an exact value, which is done for high income individuals in some years), we use the imputed values provided by Fang (2018). For most top-coded or interval records, imputations are made using a lognormal or Pareto distribution, depending on the information we have about how much higher than the taxable maximum the earnings are likely to be, or a nearest neighbor-matching method. See Appendix F for more details.

4.6 Measuring Total Income

In order to measure lifetime taxes, we need measures of total income during each year of life. The HRS survey data include measures of capital income, pensions and Social Security benefits. Combining the HRS survey measures of capital and retirement income with the administrative earnings records gives us a comprehensive measure of total income. When the HRS income measures are missing or lie outside the sample period, we impute capital income, pensions and Social Security benefits.

4.7 Measuring Tax Contributions

We calculate payroll taxes using earnings data and historic tax rates, and we use NBER's TAXSIM program to compute federal and state taxes for each household. We allow these taxes to depend on the year, state of residence, marital status, number of dependents, number of taxpayers over age 65, and total household income. For calculating taxes after 2023, we use the tax schedules for 2023. The federal and state tax amounts for each household in each year are then multiplied by the shares of aggregate federal and state taxes in that year that go towards paying Medicare and Medicaid for those over 65. We describe these calculations in detail in Appendix G.

5 Measuring Permanent Income and Medical Spending

5.1 Permanent Income

Our goal is to measure Medicare and Medicaid redistribution by permanent income. We construct a measure of permanent income (PI) that captures substantial household ex-ante heterogeneity:

households with different PI ranks have different income levels when working, receive different flows of retirement income and face different processes for health, mortality and medical expenses. A contribution of this paper is distinguishing this ex-ante heteogeneity from risk.

To estimate a household's PI, we first sum all of its annuitized income sources (Social Security benefits, defined benefit pension benefits, veterans benefits and annuities). Because there is a roughly monotonic relationship between lifetime earnings and our annuity income measure, this measure reflects income during the working period. We then construct a PI measure comparable across households of different ages and sizes. To do so, we follow the approach of De Nardi et al. (2021) and regress annuity income on a household fixed effect, dummies for household structure, a polynomial in age, and interactions between these variables. The rank order of each household's estimated fixed effect provides our measure of its PI. This is a time-invariant measure that follows the household even after one of its members dies. See Appendix H for details.

When constructing lifetime income, we use the taxable earnings histories to calculate the age-65 present value of each household's earnings history:

$$\mathrm{PV(Earnings)}_i = \sum_{t=25}^{64} e_{it}\big(\Pi_{j=t}^{64}(1+r_j)\big), \tag{5}$$

where e_{it} is household i's earnings at age t and r_j is the real interest rate at time j, so that $e_{it}\big(\Pi_{j=t}^{64}(1+r_j)\big)$ is the value of earnings received at age t and saved until age 65. For our main analysis we do not discount and set $r_j = 0$. We construct this measure using the Social Security earnings records, which include the earnings of deceased spouses. We define household i's PI, I_i, as the percentile rank of $\mathrm{PV(Earnings)}_i$.

5.2 Medical Spending

Our approach is to use observed data whenever possible. However, our medical spending data cover the years 1999 to 2012, so we do not have lifetime medical spending histories. Thus, we must forecast lifetime medical spending using a model.

5.2.1 Modeling Medical Spending

Let $M^*_{i,t}$ denote total medical spending for household i at time t, and let $m_{i,t}$ denote its logarithm net of the observed variables contained in the vector $X_{i,t}$:

$$\ln M^*_{i,t} = X_{i,t}\gamma + m_{i,t}. \tag{6}$$

Below we describe how we map total medical spending $M^*_{i,t}$ into Medicare and Medicaid payments $M_{i,t}$. In practice the vector $X_{i,t}$ includes an age polynomial, health indicators, household structure and death-year indicators, PI percentile, and interactions among the aforementioned variables. We assume that $m_{i,t}$ follows the quantile-based panel data model proposed by Arellano, Blundell, and Bonhomme (2017) and extended by Arellano et al. (2021). In particular, we assume that $m_{i,t}$ can be expressed as the sum of the persistent first-order Markov component $\eta_{i,t}$ and the transitory component $\varepsilon_{i,t}$:

$$m_{i,t} = \eta_{i,t} + \varepsilon_{i,t}, \quad \forall i \in \{1, ..., N\},\ \forall t \in \{1, ..., T\}. \tag{7}$$

We assume the transitory component is i.i.d., but require only that it be zero-mean and satisfy the regularity conditions set forth in Arellano, Blundell, and Bonhomme (2017).

In the framework of Arellano, Blundell, and Bonhomme (2017), rewrite the conditional distribution for the persistent component $\eta_{i,t}$ as:

$$\eta_{i,t} = Q_\eta(\nu_{i,t} \,|\, \eta_{i,t-1}, a_{i,t}), \quad \nu_{i,t} \overset{iid}{\sim} U[0,1], \tag{8}$$

where $Q_\eta(\nu \,|\, \eta_{i,t-1}, a_{i,t})$ denotes the νth quantile of $\eta_{i,t}$ conditional on its lagged value and age $(a_{i,t})$. The quantile function Q_η maps $\eta_{i,t}$'s conditional rank, $\nu_{i,t}$, into a value of $\eta_{i,t}$ itself. To fix ideas, if we draw $\nu_{i,t} = 0.1$, the realized value of $\eta_{i,t}$ will equal the 10th percentile of the conditional distribution of $\eta_{i,t}$ at age $a_{i,t}$, given $\eta_{i,t-1}$. As a rank, $\nu_{i,t}$ is distributed uniformly over the $[0, 1]$ interval.

If $\eta_{i,t}$ followed a standard AR(1) process with normally distributed innovations, $\eta_{i,t} = \phi\eta_{i,t-1} + \zeta_{i,t}$, $\zeta_{i,t} \sim N(0, \sigma^2_\zeta)$, the quantile function would take the linearly separable form $\eta_{i,t} = Q_\eta(\nu_{i,t} \,|\, \eta_{i,t-1}, a_{i,t}) = \phi\eta_{i,t-1} + \sigma_\zeta\Phi^{-1}(\nu_{i,t})$, where $\Phi^{-1}(\cdot)$ is the inverse of the standard normal cumulative distribution function. (Conversely, with $\nu_{i,t} \sim U[0, 1]$, we have $\sigma_\zeta\Phi^{-1}(\nu_{i,t}) \sim N(0, \sigma^2_\zeta)$). Age-independence, normality, and linearity can thus be expressed as restrictions on the quantile function in equation (8).

In its most unrestricted form, this specification allows for a great degree of flexibility. One way to see this is to construct the persistence measure

$$\phi_\tau(\eta_{i,t-1}, a_{i,t}) = \frac{\partial Q_\eta(\tau \,|\, \eta_{i,t-1}, a_{i,t})}{\partial \eta_{i,t-1}}, \tag{9}$$

with τ denoting the conditional rank of interest. $\phi_\tau(\eta_{i,t-1}, a_{i,t})$ measures the effect of $\eta_{i,t-1}$ on the τth conditional quantile of $\eta_{i,t}$. Persistence can vary by rank (τ), age ($a_{i,t}$) and prior realization ($\eta_{i,t-1}$). In contrast, in the standard model persistence always equals the constant ϕ.

In estimation we parametrically approximate the conditional quantile function with low-order Hermite polynomials. Let $h_k^\eta(\cdot)$ denote the kth Hermite polynomial used in the approximation of $\eta_{i,t}$, with $\{h_k^\eta(\cdot)\}_{k=0}^{K_\eta}$ forming the polynomial basis for the approximation. $Q_\eta(\tau \,|\, \eta_{i,t-1}, a_{i,t})$ is thus a linear combination of the K_η Hermite polynomials, with the coefficients on the polynomials, $\{\beta_k^\eta(\tau)\}_{k=0}^{K_\eta}$ themselves functions of the quantile rank τ. We thus have

$$Q_\eta(\tau \,|\, \eta_{i,t-1}, a_{i,t}) = \sum_{k=0}^{K_\eta} \beta_k^\eta(\tau) \cdot h_k^\eta(\eta_{i,t-1}, a_{i,t}), \quad \tau \in (0,1]). \tag{10}$$

The distributions of the initial shock $\eta_{i,1}$ and the transitory shocks $\{\varepsilon_{i,t}\}_t$ are handled in analogous ways:

$$Q_1(\tau \,|\, a_{i,1}) = \sum_{k=0}^{K_1} \beta_k^1(\tau) \cdot h_k^1(a_{i,1}), \quad \tau \in (0,1), \tag{11}$$

$$Q_\varepsilon(\tau \,|\, , a_{i,t}) = \sum_{k=0}^{K_\varepsilon} \beta_k^\varepsilon(\tau) \cdot h_k^\varepsilon(a_{i,t}), \quad \tau \in (0,1). \tag{12}$$

For these distributions we do not condition on $\eta_{i,t-1}$ but only only on age.

Each of the coefficient functions $\left(\{\beta_k^\eta(\tau)\}_{k=0}^{K_\eta}, \{\beta_k^1(\tau)\}_{k=0}^{K_1}, \{\beta_k^\varepsilon(\tau)\}_{k=0}^{K_\varepsilon}\right)$ in equations (10)-(12) is modelled with a set of polynomial splines defined over the intervals $\{[\tau_{\ell-1}, \tau_\ell]\}_{\ell=1}^{L}$, along with two low-dimensional tail functions defined over $(0, \tau_1]$ and $[\tau_L, 1)$. It is the parameters for these weighting functions that we must estimate.

As both the persistent and transitory shocks are unobserved, we cannot estimate the parameters of the weighting functions directly using quantile regressions. Furthermore, our data are unbalanced because sample members die. We therefore follow the extension of the *E-M* algorithm described in and applied by Arellano et al. (2021).

- In the *E-step* we find the posterior distribution of the unobserved persistent shocks ($\{\eta_{i,t}\}_t$)

implied by the data and the current parameterization of the model. In particular, we use the coefficients of the Hermite polynomials $\left(\{\beta_k^{\eta}(\tau)\}_{k=0}^{K_\eta}, \{\beta_k^{1}(\tau)\}_{k=0}^{K_1}, \{\beta_k^{\varepsilon}(\tau)\}_{k=0}^{K_\varepsilon}\right)$, which fully determine the distributions of the shocks, and a Monte Carlo method to simulate draws from the distributions of the initial shock $\eta_{i,1}$ and the subsequent shocks $\{\eta_{i,t}\}_t$.[5] This part of the procedure is a special case of the Sequential Monte Carlo methods described in greater detail in Creal (2012).

- In the *M-step* we use quantile regressions to update the coefficient functions for the Hermite polynomials, using the distribution of $\{\eta_{i,t}\}_t$ found in the *E-step*. Once the coefficients have been updated, we return to the *E-step* and simulate new draws.

We iterate between the E and M steps until the parameters converge. See Appendix I for a more detailed description of the methodology.

5.3 Health and Mortality

Let $hs_{i,g,t}$ denote the health of member $g \in \{h, w\}$ in household i at age t. Health has four mutually exclusive possible values: dead; in a nursing home; in bad health; or in good health. We assume that the transition probabilities for an individual's health depend on his or her current health, age, permanent income I, and gender g.[6] It follows that the elements of the health transition matrix are given by

$$\pi_{q,r}(t, I_i, g) = \Pr\left(hs_{i,g,t+1} = r \quad hs_{i,g,t} = q; t, I_i, g\right), \tag{13}$$

with the transitions covering a one-year interval. Although the HRS interviews every other year, we adopt the approach in De Nardi, French, and Jones (2016), who fit annual models of health to the HRS data for singles. We extend their approach to account for the dynamics of two-person households. We estimate health/mortality transition probabilities by fitting the transitions observed in the HRS to a multinomial logit model.[7] See Appendix J for further details.

[5] This approach takes advantage of the Markovian structure of the model and has been shown to perform well in low-dimensional models.

[6] We do not allow health transitions to depend on medical spending. The empirical evidence on whether medical spending improves health, especially at older ages, is surprisingly mixed (De Nardi, French, and Jones 2016). Likely culprits include reverse causality – sick people have higher expenditures – and a lack of insurance variation – almost every retiree gets baseline insurance through Medicare. We also do not allow health transitions to depend on marital status. De Nardi et al. (2021) find that after controlling for income and past health, marital status has little added predictive power.

[7] We do not control for cohort effects. Instead, our estimates are a combination of period (cross-sectional) and cohort probabilities. This may lead us to underestimate the lifespans expected by younger cohorts as they age. Nevertheless, lifespans have increased only modestly over the sample period. Accounting for cohort effects would thus have at most a modest effect on our estimates.

5.4 Medical Spending Budget Sets

The variant of the Arellano, Blundell, and Bonhomme (2017) framework that we estimate provides us with a dynamic model of total medical medical spending $M^*_{i,t}$. To infer Medicare and Medicaid payments $M_{i,t}$ we disaggregate this total into its components, using budget sets that map total medical spending into Medicare and Medicaid. We estimate these budget sets in two steps. First, using a logit model, we estimate the probability of Medicaid receipt as a function of: age, marital status and gender, the health and nursing home status of the husband and wife, the log of total medical spending, and Medical spending interacted with PI. Second, using OLS we estimate the share of total medical spending paid by Medicare and Medicaid as a function of the same same variables we use to estimate Medicaid recipiency.

6 Estimation Results

6.1 Earnings Profiles

We calculate percentile ranks of lifetime earnings for the cohort whose household head turned 65 in 1999-2004. Figure 1 presents average earnings, conditional on PI quintile and age. For those at the bottom of the distribution, earnings are close to $20,000 between ages 40-55, and lower at other ages. For those at the top, earnings exceed $60,000 per year between ages 40 and 55.

Figure 2 presents the taxes and insurance premia used to fund Medicaid and Medicare, again conditional on PI quintile and age. To calculate taxes, we first calculate total income from all sources. From that we calculate total payroll, federal and state taxes. Then, using the methods described above, we attribute a portion of these taxes to the funding of Medicare and Medicaid. For those at the bottom of the distribution, total tax payments are very low, whereas the taxes paid to Medicare and Medicaid by those at the top exceed $3,000 per year. Prior to age 30, payments are close to $0, both because income is low at these ages and also because the Medicare and Medicaid programs were just starting at these ages for the cohort we study, and the tax rates for Medicare and Medicaid were low when the programs were new. At age 65, Medicare Part B premia begin to be paid, leading to a jump in payments. Because households at the bottom of the PI distribution are likely to have their Medicare premia covered by Medicaid, their total contributions jump less at age 65 than those of other groups.

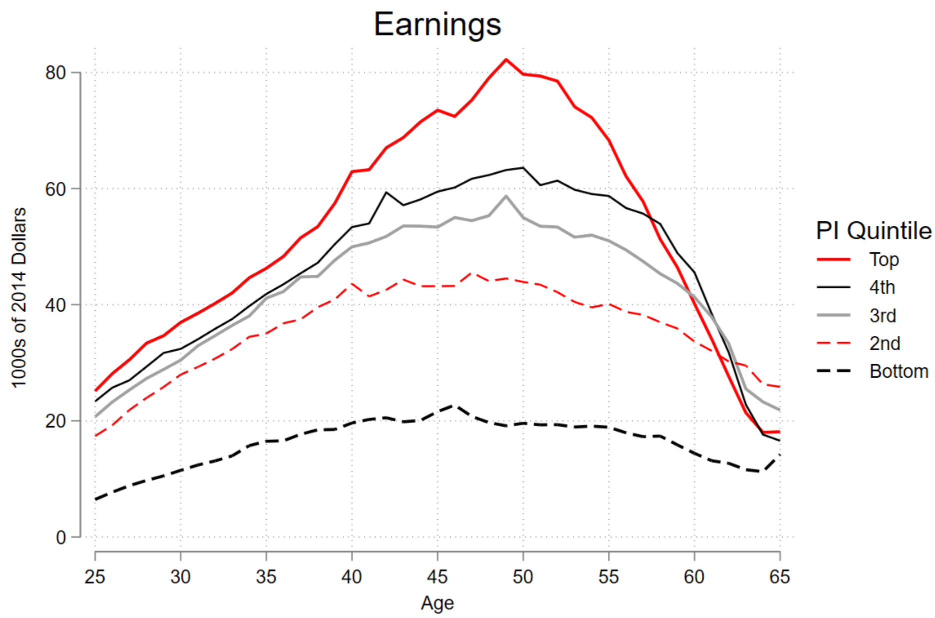

Figure 1: Average Household Earnings from Wage and Self-Employment Income by PI Quintile. Cohort whose household head turned 65 in 1999-2004

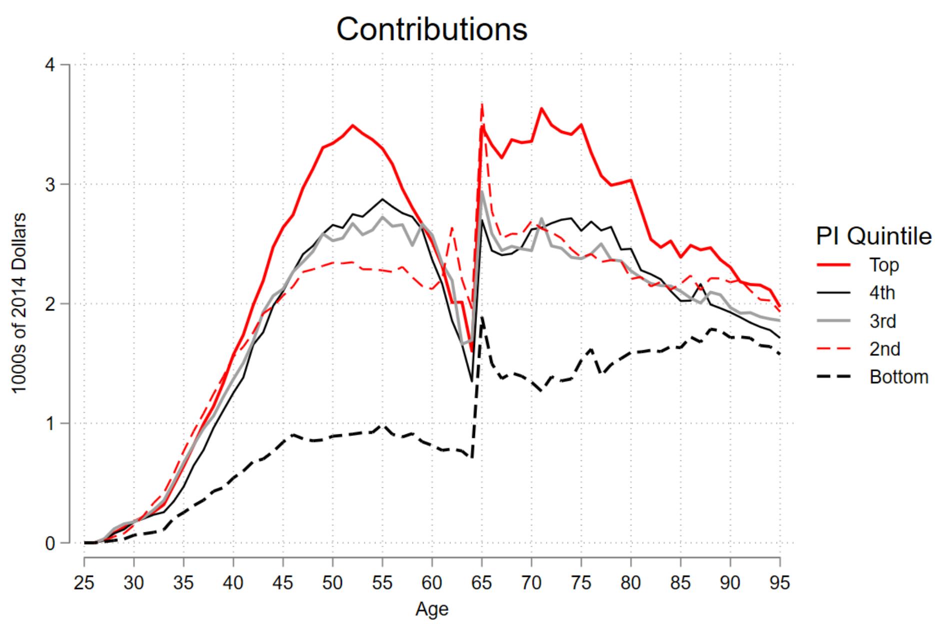

Figure 2: Contributions are the sum of: state & federal taxes going to Medicare/Medicaid+Medicare payroll contributions + premia. Cohort whose household head turned 65 in 1999-2004

6.2 Mortality and Health Status

We estimate health transitions and mortality rates simultaneously by fitting the transitions observed in the HRS to the multinomial logit model described in Section 5.3.

Table 2 shows the life expectancies implied by our demographic model for those still alive at age 65. The first panel of the table shows the life expectancies for individuals under different configurations of gender, PI percentile, and age-65 health. The healthy live longer than the sick, the rich (higher PI) live longer than the poor, and women live longer than men. For example, a man at the 10th PI percentile in a nursing home expects to live only 5.4 more years, while a single woman at the 90th percentile in good health expects to live 20.9 more years. The second panel of the table shows life expectancies for married *households*, that is, the average length of time that at least one member of the household is still alive or, equivalently, the life expectancies for the oldest survivors. While wives generally outlive husbands, a non-trivial fraction of the oldest survivors are men, and the life expectancy for a married household is roughly three years longer than that of a married woman with same initial health and PI quantile.

		Men			Women		
Permanent Income Percentile	Nursing Home	Bad Health	Good Health	Nursing Home	Bad Health	Good Health	All*
Individuals							
10	5.4	13.7	15.8	8.7	17.1	18.6	16.6
50	5.5	15.8	18.1	9.2	19.3	20.8	19.1
90	4.7	15.6	18.2	8.2	19.2	20.9	19.2
Couples (oldest survivor)†							
10	10.8	20.2	21.4				21.0
50	11.7	22.8	24.0				23.7
90	10.5	22.7	24.0				23.8
All Men							17.1
All Women							19.9
All Couples (oldest survivor)							23.3

* Averages taken over initial health found in the data. Results indexed by PI percentile are taken over the associated PI quintile.

† Life expectancy of oldest survivor in a household. Health-specific results for couples assume that both spouses have the same health at age 65.

Table 2: Life expectancy in years, conditional on reaching age 65

Most of the results shown in Table 2 are disaggregated by PI and initial health. When we average over all these factors, we find that a man alive at age 65 will on average live an additional 17.1 years, while a woman alive at age 65 will on average live an additional 19.9 years.[8] Our predicted life expectancy at age 65 is close to what the aggregate statistics imply. In addition, our estimated income gradient is similar to that in Waldron (2007) and De Nardi et al. (2021), who find that those in the top of the income distribution live 3 years longer than those at the bottom, conditional on being 65.

Another key statistic for our analysis is the probability that a 65-year-old will spend significant time (a stay of more than 120 days) in a nursing home before he or she dies. We estimate this probability to be 22.6% for men and 36.1% for women. Nursing home incidence differs relatively modestly across the PI distribution. Although high-income people are less likely to be in a nursing home at any given age, they live longer, and older individuals are much more likely to be in a nursing home. The nursing home risk also varies relatively little with age-65 health status (good or bad), for similar reasons.

6.3 Budget Sets

We estimate the probability of Medicaid receipt as a logistic function of: age, marital status and gender, the health and nursing home status of the husband and wife, the log of total medical spending and its square, PI percentile and its square, and medical spending interacted with PI. The estimated coefficients can be seen in Appendix Table 8. Medicaid receipt is increasing in medical spending and, for older households, decreasing in PI. Households with members in bad health are more likely to receive Medicaid, as are older households with a member in a nursing home.[9] Table 3 shows how Medicaid recipiency varies by PI quintile, both in the data and as predicted by the logistic model. Households at the bottom of the PI distribution are more likely to receive Medicaid, especially at older ages. Among those 65 and older, 39% of households in the bottom PI quintile receive Medicaid. The table also shows that the model fits the observed income gradient of receipt well.

Using OLS, we estimate the share of total medical spending paid by Medicare and Medicaid as a function of the same variables we use to predict Medicaid recipiency. Figure 3 shows the

[8]We construct these distributions of PI and initial health with bootstrap draws of people aged 63-67 in HRS.

[9]Although the specifications differ in many respects, our results share a number of similarities with those reported in Borella, De Nardi, and French (2018), who also find that Medicaid use decreases with PI and increases as health worsens.

estimated budget sets for older households, holding all health and household composition indicators fixed at their sample averages. The share of total medical spending paid by Medicare is rising in total spending, reflecting the fact that Medicare pays for close to 100% of high-cost hospital stays. Among Medicaid recipients, having a member in a nursing home leads the Medicare share to decline by 20 percentage points and the Medicaid share to rise by a similar amount. This reflects the fact that Medicare does not pay for extended nursing home stays but Medicaid does. For those not on Medicaid, having a member in a nursing home leads to a somewhat smaller decrease in the Medicare share, of around 10-12 percentage points. However, residing in a nursing home significantly increases the likelihood of Medicaid receipt among older households.

PI Quintile	Data	Model
Top	1.2%	1.5%
4th	3.2%	2.7%
3rd	6.1%	5.8%
2nd	14.3%	15.0%
Bottom	39.3%	39.0%

Notes: Model fractions are mean logit probabilities.

Table 3: Probability of Receiving Medicaid for over 65s by PI Quintile

7 Lifetime Taxes Paid and Healthcare Benefits Received

To assess taxes paid and benefits received from a lifetime perspective, we calculate the present values of taxes paid and medical spending from all payors. We use observed values of income, health and survival when they are available, which is up through 2016. For the years after 2016, we simulate life histories of health and mortality for each household member using our estimated transition probabilities and calculate income and thus tax payments for these households. Next, we simulate medical spending consistent with the stochastic processes described in the model section. We then use our estimated budget sets to calculate Medicare and Medicaid expenditures at each age, and can thus calculate the present value of these expenditures.

Figure 4 presents total medical spending (panel (a)), Medicare spending (panel (b)), Medicaid spending (panel (c)), and the sum of out-of-pocket spending, private and other insurance payments (panel (d)). Total medical spending rises with age. It also rises with income, reflecting the fact that high-income households have lower mortality, and are thus more likely to be two-person households,

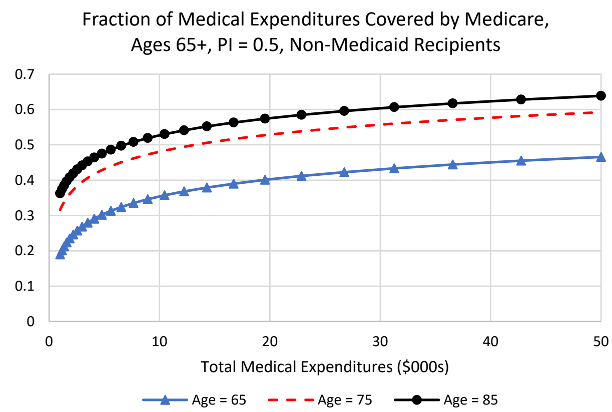

(a)

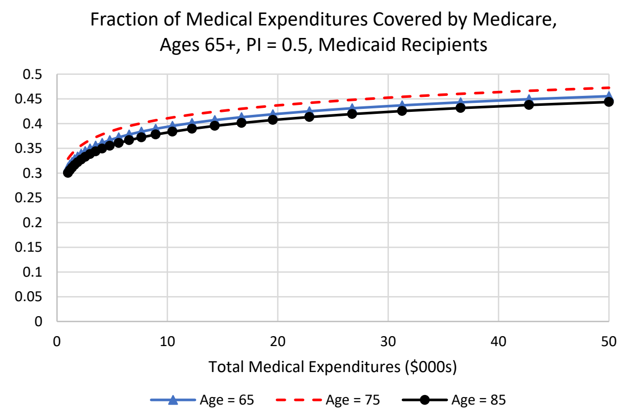

(b)

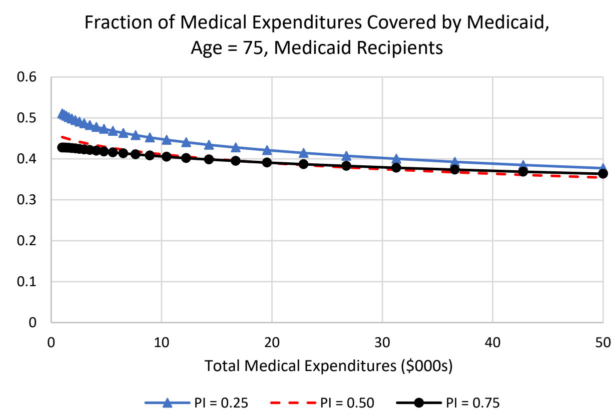

(c)

Figure 3: Fraction of Medical Expenditures Paid by Medicare (panels (a) and (b)) or Medicaid (panel (c))

Notes: These fractions average over all health states, marital status and gender. PI is permanent income percentile, so PI=.5 is median income.

even at older ages. Because high-income households have high medical spending, and they are less likely to receive Medicaid, they benefit heavily from the Medicare program. Panel (c) shows that it is mostly those at the bottom of the distribution who benefit from the Medicaid program, although at very old ages some high-income households also benefit from Medicaid, often because they face catastrophic nursing home expenses that deplete their assets (De Nardi, French, and Jones 2016). Panel (d) shows that that the low income pay less out of pocket.

Table 4 shows the lifetime value of income, tax contributions, and medical benefits. Column (1) shows the present value of income, the sum of earnings, capital income, Social Security, and pensions. Those in the highest PI quintile have an average lifetime income of $8.0 million, whereas the

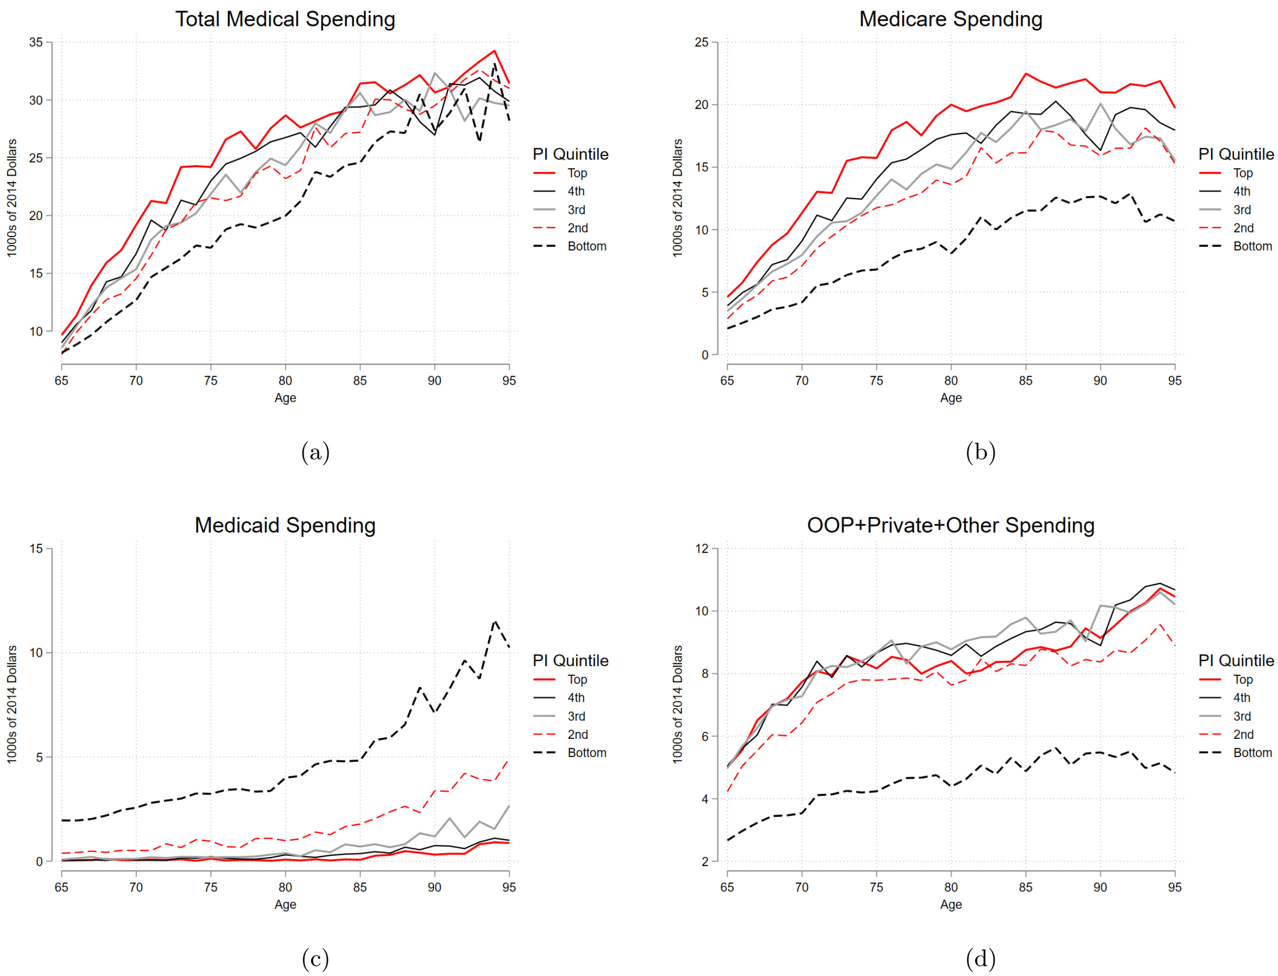

Figure 4: Mean Annual Medical Spending for Surviving Households, by permanent income (PI) quintile. Panel (a): Total Medical Spending, Panel (b): Medicare Payments, Panel (c): Medicaid Payments, Panel (d): Out-of-Pocket+Private+Other

average for those in the bottom is \$0.9 million. Column (2) shows the sum of Medicare payroll tax contributions, state and federal income taxes that go to Medicare and Medicaid, and Medicare Part B premia. Those at the top pay much more in taxes than those at the bottom of the PI distribution. For example, those in the highest quintile on average contribute \$248,000, whereas those at the bottom contribute \$33,000. This reflects both higher income and higher marginal tax rates for those in the top PI groups. Despite much higher contributions, those at the top of the PI distribution do not pay as much as they receive in benefits from Medicare and Medicaid. Columns (3)-(5) show lifetime medical spending payments from Medicare and Medicaid. Those in the top PI quintile receive on average \$396,000 in Medicare benefits, whereas those in the bottom receive

$137,000, reflecting the longer lifespans of those at the top. Those at the bottom of the distribution receive more in the way of Medicaid benefits, however. Whereas those at the top of the distribution on average receive $5,000 from Medicaid, those at the bottom receive $92,000.

Column (6) reports net benefits, which is the difference between Medicare and Medicaid benefits received and total contributions paid. All income groups in our cohort are net beneficiaries, on average receiving a net benefit of $220,000. This is the amount of cross-cohort redistribution that will be paid by younger cohorts. Different income groups within our sample cohort receive different net benefits, however. Those at the top and at the bottom receive fewer net benefits than those in the middle. Those at the top of the income distribution pay the most in the way of contributions, whereas those in the bottom receive the least in the way of benefits. As a result, the amount of within-cohort redistribution is -$67,000 for those at the top and -$25,000 for those at the bottom.

	Lifetime	Lifetime	Medical Benefits After Age 65			Net	Within-Cohort
PI	Income	Contributions	Medicare	Medicaid	Total	Benefits	Redistribution
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
All	*3,306*	*107*	*296*	*31*	*327*	*220*	–
Top	8,003	248	396	5	401	153	-67
4th	4,082	131	354	9	362	231	11
3rd	2,445	82	321	15	337	255	34
2nd	1,310	47	272	36	308	261	41
Bottom	906	33	137	92	229	196	-25
Cross-Cohort Redistribution: 220							

Notes: Amounts in 1000s of 2014 dollars.
Lifetime Income = earnings + capital income + Social Security + pension income.
Lifetime Contributions = Part B premia + Medicare payroll contributions + state and federal income taxes devoted to Medicare/Medicaid. Net Benefits = Medicare + Medicaid benefits – Lifetime Contributions.

Table 4: Income, Tax Contributions and Medical Spending by Permanent Income Quintile

8 Comparison to England and the NHS

An interesting comparison to the US Medicare and Medicaid systems is the National Health Service (NHS) system in the United Kingdom. In the UK all households have access to health care benefits from the NHS, which is (mostly) free at the point of use and (mostly) financed out of general tax revenues. Yet not all households receive equal benefits from the NHS over their lifetime, and not all households contribute equally in tax payments. Figure 5 presents evidence from survey data from

ELSA matched to hospital records (Hospital Episode Statistics, HES). The figure uses two measures of socioeconomic status: education level and local income in the area the individual lived. It shows that those at the bottom of the distribution received more health care benefits than those at the top. Comparing this figure to Figure 4 suggests there is more redistribution to the low income in the NHS than there is in the US Medicare and Medicaid systems.

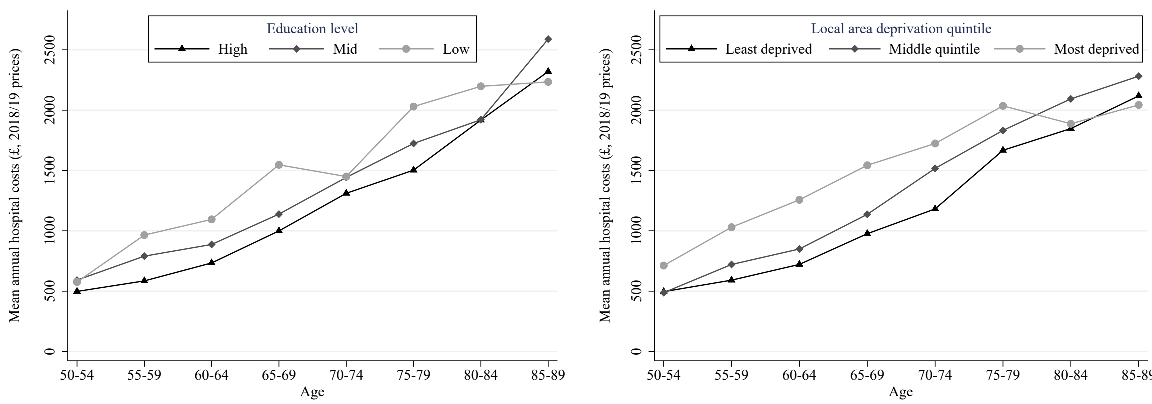

(a) (b)

Figure 5: Mean hospital spending by education group (left panel) and local deprivation level (right panel)

Note: Hospital spending refers here to the sum of inpatient, outpatient and Accident & Emergency spending. Data are pooled for waves 4–8 of ELSA (covering the period from 2008 to 2017). Source: Authors' calculations using linked ELSA-HES data.

9 Conclusion

Using data from the Health and Retirement Study, matched to administrative Medicare, Medicaid, and Social Security earnings records, we estimate the distribution of lifetime medical benefits received and the distribution of lifetime taxes paid for that care among the cohort who turned 65 in 1999-2004. The largest beneficiaries of Medicare and Medicaid are those in the middle of the income distribution, as these people live long yet pay modest taxes. In contrast, while those at the top of the income distribution live the longest, they contribute more in taxes. Those at the bottom contribute the least in taxes, but also have the shortest lifespans.

However, all income groups within the cohort we study are net beneficiaries from Medicare and Medicaid. On average this cohort's lifetime tax contributions did not cover the medical benefits it received. This deficit must be paid by younger cohorts. Thus, for the cohort we study, cross-cohort redistribution is greater than within cohort redistribution.

References

Arellano, M., R. Blundell, and S. Bonhomme (2017). Earnings and consumption dynamics: A nonlinear panel data framework. *Econometrica 85*(3), 693–734.

Arellano, M., R. Blundell, S. Bonhomme, and J. Light (2021). Heterogeneity of consumption responses to income shocks in the presence of nonlinear persistence.

Auerbach, A. J., K. K. Charles, C. C. Coile, W. Gale, D. Goldman, R. Lee, C. M. Lucas, P. R. Orszag, L. M. Sheiner, B. Tysinger, et al. (2017). How the growing gap in life expectancy may affect retirement benefits and reforms. *The Geneva Papers on Risk and Insurance-Issues and Practice 42*(3), 475–499.

Bhattacharya, J. and D. Lakdawalla (2006). Does Medicare benefit the poor? *Journal of Public Economics 90*(1), 277–292.

Borella, M., M. De Nardi, and E. French (2018). Who receives medicaid in old age? rules and reality. *Fiscal Studies 39*(1), 65–93.

Chetty, R., M. Stepner, S. Abraham, S. Lin, B. Scuderi, N. Turner, A. Bergeron, and D. Cutler (2016). The association between income and life expectancy in the united states, 2001-2014. *JAMA 315*(16), 1750–1766.

Creal, D. (2012). A survey of sequential monte carlo methods for economics and finance. *Econometric reviews 31*(3), 245–296.

Cubanski, J., C. Swoope, C. Boccuti, G. Jacobson, G. Casillas, S. Griffin, and T. Neuman (2015). A primer on Medicare: Key facts about the Medicare program and the people it covers. Technical report, The Henry J. Kaiser Family Foundation.

David, M., R. J. A. Little, M. E. Samuhel, and R. K. Triest (1986). Alternative methods for cps income imputation. *Journal of the American Statistical Association 81*(393), 29–41.

De Nardi, M., E. French, and J. B. Jones (2016). Medicaid insurance in old age. *American Economic Review 106*(11), 3480–3520.

De Nardi, M., E. French, J. B. Jones, and J. McCauley (2016). Medical spending on the US elderly. *Fiscal Studies 37*(3-4), 717–747.

De Nardi, M., E. French, J. B. Jones, and R. McGee (2021). Why do couples and singles save during retirement? Technical report, National Bureau of Economic Research.

Doucet, A., A. M. Johansen, et al. (2009). A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of Nonlinear Filtering 12*(656-704), 3.

Fahle, S., K. McGarry, and J. Skinner (2016). Out-of-pocket medical expenditures in the United States: Evidence from the Health and Retirement Study. *Fiscal Studies 37*(3-4), 785–819.

Fang, C. (2018). Health and retirement study imputations of lifetime earnings records.

Feenberg, D. and E. Coutts (1993). An introduction to the TAXSIM model. *Journal of Policy Analysis and Management 12*(1), 189–194.

French, E. and J. B. Jones (2004). On the distribution and dynamics of health care costs. *Journal of Applied Econometrics 19*(4), 705–721.

French, E. and J. B. Jones (2011). The effects of health insurance and self-insurance on retirement behavior. *Econometrica 79*(3), 693–732.

French, E., J. B. Jones, and J. McCauley (2017). The accuracy of economic measurement in the Health and Retirement Study. *Forum for Health Economics and Policy 20*(2).

Friedberg, L., W. Hou, W. Sun, A. Webb, and Z. Li (2015). New evidence on the risk of requiring long-term care. WP 2014-12, Boston College Center for Retirement Research.

Hurd, M. D., P.-C. Michaud, and S. Rohwedder (2017). Distribution of lifetime nursing home use and of out-of-pocket spending. *Proceedings of the National Academy of Sciences 114*(37), 277–297.

Koenker, R. and G. Bassett Jr (1978). Regression quantiles. *Econometrica: Journal of the Econometric Society*, 33–50.

Lassman, D., M. Hartman, B. Washington, K. Andrews, and A. Catlin (2014). Us health spending trends by age and gender: selected years 2002–10. *Health Affairs 33*(5), 815–822.

McClellan, M. and J. Skinner (2006). The incidence of Medicare. *Journal of Public Economics 90*(1), 257–276.

Nicholas, L. H., K. M. Langa, T. J. Iwashyna, and D. R. Weir (2011). Regional variation in the association between advance directives and end-of-life Medicare expenditures. *JAMA 306*(13), 1447–1453.

Pashchenko, S. and P. Porapakkarm (2016). Medical spending in the US: facts from the Medical Expenditure Panel Survey data set. *Fiscal Studies 37*(3-4), 689–716.

Pijoan-Mas, J. and J.-V. Ríos-Rull (2014). Heterogeneity in expected longevities. *Demography 51*(6), 2075–2102.

Rettenmaier, A. (2012). The distribution of lifetime medicare benefits, taxes and premiums: Evidence from individual level data. *Journal of Public Economics 96*(9-10), 760–772.

Vogel, R. J. (1988). An analysis of the welfare components and intergenerational transfers under the medicare program. *Lessons from the First Twenty Years of Medicare*, 73–114.

Wagstaff, A. and E. Van Doorslaer (1992). Equity in the finance of health care: some international comparisons. *Journal of Health Economics 11*(4), 361–387.

Wagstaff, A., E. Van Doorslaer, H. Van Der Burg, S. Calonge, T. Christiansen, G. Citoni, U.-G. Gerdtham, M. Gerfin, L. Gross, U. Häkinnen, et al. (1999). Equity in the finance of health care: some further international comparisons. *Journal of Health Economics 18*(3), 263–290.

Waldron, H. (2007). Trends in mortality differentials and life expectancy for male social security-covered workers, by socioeconomic status. *Social Security Bulletin 67*(3), 1–28.

Appendix A: How Medicare is Financed

Medicare is divided into four parts: each part provides different coverage and is funded differently. Medicare originally was made up of Hospital Insurance (HI, Part A) and Supplemental Medical Insurance (SMI, Part B). Since then it has expanded to include Medicare Advantage (Part C) in which beneficiaries can select an HMO or PPO plan offered by private companies approved by Medicare, and drug coverage (Part D), which was first offered in 2006. Part A is primarily financed using a payroll tax, whereas Parts B and D are financed primarily out of general tax revenues. Part C is financed out of transfers from both Parts A and B (and is thus funded through a combination of general revenues and the payroll tax). We take these differences in financing into account when calculating how much each individual pays into the Medicare and Medicaid systems.

Part A covers inpatient hospital services and is funded primarily (88% in 2019) through a payroll tax, with most of the rest coming from taxation of OASDI benefits. Part A is available for free if an individual or their spouse contributed 40 quarters of payroll tax during the time they spent working. Otherwise, they must pay a premium or, if they are low income, Medicaid pays the premium. As a result, most people (99%) [https://www.cms.gov/newsroom/fact-sheets/2020-medicare-parts-b-premiums-and-deductibles] do not pay the Part A premium. As a result, we assume nobody pays the Part A premium.

Part B mostly covers doctor office visits and other outpatient services. It is financed roughly 75% by general revenue and 25% by premia. Most Medicare beneficiaries enroll in Medicare Part B (in 2007, 93% of Medicare beneficiaries were enrolled in both Parts A and B, 6% were enrolled in Part A only, and 1% were enrolled in Part B only) and thus we assume everyone is enrolled in both Part A and Part B.

Part B beneficiaries pay a monthly premium ($144.60 is the standard monthly premium, though some get their Part B premia paid through Medicaid). However, this premium can vary based on a person's circumstances.

Higher earning beneficiaries have had to pay larger premia since 2007; this affects about 7% of people with Medicare Part B. $491.60 is the maximum monthly premium for the highest income individuals.

Medicaid also plays a role in what beneficiaries have to pay for Medicare coverage. In 2016, 5.3 million traditional Medicare beneficiaries (traditional Medicare refers to Parts A and B) were "dual-eligibles" who received both full Medicaid benefits and payments of their Medicare premia and cost sharing. An additional 1.7 million traditional beneficiaries did not qualify for full Medicaid benefits but qualified for the Medicare Savings Programs that cover Part A and/or Part B premia and/or cost sharing.

We calculate the Part B premia in our HRS data using information on individuals' earnings and Social Security benefits. We assume that if a household receives Medicaid, Medicaid pays the premium.

Part C (Medicare Advantage) can be chosen instead of Parts A and B and covers the same services, in addition to possibly including prescription drugs and other additional benefits. Dual-eligible Medicare and Medicaid beneficiaries can also choose to enroll in Medicare Advantage plans. In 2019, 34% of Medicare beneficiaries were enrolled in a Medicare Advantage plan, which is a significant increase from 18% in 1999 (https://www.kff.org/medicare/fact-sheet/medicare-advantage/). Part C monthly premia vary by plan and over time, averaging $36 in 2020. In addition to this, the Part B premium must be paid as well because a person has to have coverage from Parts A and B to join a Part C plan. The Medicare Advantage program is not separately financed; Medicare payments to private health plans come from both the HI (Part A) trust fund and Part B account of the SMI trust fund. The split in 2019 was 43% from Part A and the rest from Part B.

Part D covers prescription drugs and is financed by a combination of beneficiary premia (16% in 2019), general revenues (71% in 2019) and transfers from states (12% in 2019) in part to pay for those enrolled in both Medicare and Medicaid. The Part D monthly premia average $30 for 2020, but also vary by plan and higher income beneficiaries can pay up to $76.40 more than their standard plan premia. In 2019, 70% of Medicare beneficiaries were enrolled in Part D.[10]

Appendix B: Sample Selection

Our initial sample comprises 26,598 households. We drop households with multiple members of the same sex, who share the same household but are not married, or who disagree about their marital status, or who marry or divorce over the sample period. We also drop households where members disagree about their marital status, problematic mortality transitions (e.g., who are reported to be dead in one period and are reported to be alive subsequently), whose household head is under 65, or whose Medicare and Medicaid records were not provided or were missing. Table 5 below denotes the HRS sample size after every drop. Our final sample is 8,777 households comprising 122,878 household-year observations for the full estimation sample.

Appendix C: Our Medicare and Medicaid Data

[10] Of this 70%, 39% had coverage through Medicare Advantage plans

Selection criteria (reason for dropping)	Remaining sample size
Initial sample	26,598
Partnered households or marital status contradictions[1]	23,936
Inconsistent marital status transitions[2]	23,270
Wrong number of people in the household[3]	22,270
Attrition or missing values[4]	18,709
Household head under 65 in 2016	12,530
Incomplete Medicare and Medicaid records	8,777

Notes: Sample size refers to the number of unique households after a sample selection criteria.
[1] We drop households who are partnered but not married, disagree about their marital status, or get divorced over the sample period.
[2] We drop households with the following transitions: never married to widowed, married to never married, never married to divorced, or who get married.
[3] We drop same sex households, married but spouse is not observed, unmarried people sharing a household.
[4] Household drops out of sample for reasons other than death, or who have missing marital status, or who have all sample members die before a second interview.

Table 5: Sample Selection

Medicare

We link restricted Medicare fee-for-service (Parts A and B), and Part D data for the years 1999-2012 (2006 was the first year of Medicare Part D and thus our Part D data begins then) to our HRS survey data for respondents who consent to allow their Medicare data to be linked to their survey responses (approximately 64.7% percent of persons in our study population). These records have enrollment information and data on reimbursement amounts for inpatient, skilled nursing facility, home health, and hospice claims (Medicare Part A), as well as outpatient, carrier (non-institutional medical care providers such as individual or group practitioners, non-hospital labs, and ambulances), and durable medical equipment claims for Medicare Part B.

We use the Beneficiary Annual Summary File (BASF), which summarizes information from the micro-level claims records. The BASF contains annual information for each individual on the number of months of enrollment in Medicare Part A, Part B, and non-fee-for-service plans. The BASF has information on Medicare fee-for-service (FFS) claims. Almost all claims for services used by non-FFS Medicare patients are not observed in these data, so all analyses exclude an individual in a given year if they were enrolled in a non-FFS Medicare plan for more than half the year.

Medicare Part D is the prescription drug benefit. We calculate the Medicare Part D payment using the Part D event files. For the Part D contribution we subtract from the gross drug cost the payments paid by the beneficiary, family, or friends the drug costs at point of sale over the whole year.

Medicaid

As with the Medicare data, we are able to link restricted Medicaid data (CMS Medicaid Analytic eXtract, or "MAX" files) for those in the HRS who gave permission, allowing us to measure Medicaid expenditures for the Medicaid beneficiaries in our dataset for the years 1999-2012. The MAX files contain personal summaries (which contain eligibility, enrollment, and demographic information) and claims data across four service categories (inpatient, long-term care, prescription drugs, and other services). Other services include a variety of services (e.g., physician services and lab work) that do not fit under the other three service categories. The inpatient, long-term care, prescription drugs, and other services files contain the primary variable of interest, "Medicaid Payment Amount," which is the total amount of money paid by Medicaid for a particular service. We sum over all the claims for all the different service categories for a particular individual in each year.

Appendix D: Imputing Missing Medical Expenditures

Our goal is to measure all medical spending: the variable M_{it} in equation (6) of the main text is defined to include out-of-pocket spending, Medicare and Medicaid payments, and private and other public (such as Veterans Administration benefits, and care provided by local and state health departments) insurance payments. While the HRS includes information on out-of-pocket spending and can be linked to Medicare and Medicaid payments, it does not include Medicare Part C, private, or other public insurance payments. In this appendix we describe how we use data from the Medical Expenditure Panel Survey (MEPS) to impute these payments in the HRS. Although the MEPS has extremely high quality information on all payors for all household members, it lacks the long-panel dimension of the HRS. Our imputation procedures allows us to exploit the best of both datasets.

Our imputation procedure has two steps. First, we use the MEPS to infer private and other public insurance payments, conditional on variables that are observed in both datasets. Second, we impute private and other public insurance payments in the HRS data using a conditional mean matching procedure (which is a procedure very similar to hot-decking).

First Step of Imputation Procedure

We use the MEPS to infer payments of other payors, conditional on the observable variables that exist in both the MEPS and the HRS datasets.

Let i index individuals in the HRS and j index individuals in the MEPS. Define M_{it}^{obs} as out of pocket, Medicaid, and Medicare (Part A, B, and D, but not Part C) payments which are observed in both the HRS and MEPS datasets, M_{it}^{miss} as the components of medical spending that are missing in the HRS but observed in the MEPS, and $M_{it} = M_{it}^{miss} + M_{it}^{obs}$ as total medical spending. To impute M_{it}^{miss}, which is missing in the HRS, we follow David et al. (1986), French and Jones (2011), and De Nardi et al. (2021) and use a predictive mean-matching regression approach. There are two steps to our procedure. First, we use the MEPS data to regress M_{it}^{miss} on observable variables that exist in both datasets. This regression has an R^2 statistic of 0.30 for Medicare Part C, 0.15 for private insurance payments and 0.18 for other payors. Second, we impute M_{it}^{miss} in the HRS data using a conditional mean-matching procedure, a procedure very similar to hot-decking.

First, for every member of the the MEPS sample, we regress the variable of interest M^{miss} on the vector of observable variables z_{jt}, yielding $M_{jt}^{miss} = z_{jt}\beta + \varepsilon_{jt}$. Second, for each individual j in the MEPS we calculate the predicted value $\widehat{M^{miss}}_{jt} = z_{jt}\hat{\beta}$, and for each member of the sample we calculate the residual $\widehat{\varepsilon}_{jt} = M_{jt}^{miss} - \widehat{M^{miss}}_{jt}$. Third, we sort the predicted value $\widehat{M^{miss}}_{jt}$ into deciles and keep track of all values of $\hat{\varepsilon}_{jt}$ within each decile. We use this procedure separately to impute Medicare Part C benefits, private payments, and other payments.

In practice we include in z_{jt} a fourth-order age polynomial, marital status, gender, self-reported health (=1 if self reported health is good, very good, or excellent), race, visiting a medical practitioner (doctor, hospital or dentist), out-of-pocket medical spending, education of head (high school, some college, college), death of an individual, and total household income. We estimate this regression two times: once for the privately insured, and once for other payors.

Because the measure of medical spending in the HRS is medical spending over two years, we divide HRS out-of-pocket medical spending by 2 and assume that medical spending is equal across the two years.

Second Step of Imputation Procedure

For every observation in the HRS sample with a positive Medicaid indicator, we impute $\widehat{Med}_{it} = z_{it}\widehat{\beta}$, using the values of $\widehat{\beta}$ estimated from the MEPS. Then we impute ε_{it} for each observation of this subsample by finding a random observation in the MEPS with a value of $\widehat{Med}_{jt}$ in the same decile as $\widehat{Med}_{it}$, and setting $\hat{\varepsilon}_{it} = \hat{\varepsilon}_{jt}$. The imputed value of Med_{it} is $\widehat{Med}_{it} + \hat{\varepsilon}_{it}$.

As David et al. (1986) point out, our imputation approach is equivalent to hot-decking when the "z" variables are discretized and include a full set of interactions. The advantages of our approach over hot-decking are two-fold. First, many of the "z" variables are continuous. Second, to improve fit we use a large number of "z" variables. We find that adding extra variables is very important for improving fit when imputing payments. Because hot-decking uses a full set of interactions, this would result in a large number of hot-decking cells relative to our sample size. Thus, in this context, hot-decking is too data intensive.

Appendix E: Validating the Administrative Medical Spending Data

Here, we examine in greater detail the accuracy of the administrative medical spending data, as well as the out-of-pocket spending found in the AHEAD cohort of the HRS, comparing them to data from the Medicare Current Beneficiary Survey (MCBS) and Medical Expenditure Panel Survey (MEPS). See De Nardi, French, and Jones (2016) and De Nardi et al. (2016) for more details of the MCBS data and for example Nicholas et al. (2011) for details of the HRS linked data.

The MCBS is a nationally representative survey of Medicare beneficiaries, consisting of Disability Insurance recipients and Medicare recipients aged 65 and older. 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 spending (from all sources). The MCBS data are matched to Medicare records, and medical spending 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 HRS 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 both samples, we applied only modest sample selection restrictions. The key sample selection issue shown in Table 5 is that in the HRS we drop households with missing or erroneous Medicare or Medicaid records.

Table 6 compares distributions of total, out-of-pocket, Medicare, and Medicaid payments between the MCBS and the HRS data. In the table medical spending in the HRS is at an individual level (rather than household) to be comparable with the MCBS, which only has individual level data. Medical spending is higher in our HRS sample than in the MCBS sample. Furthermore, this higher level of spending is driven by higher out-of-pocket spending, Medicare, and Medicaid spending. These differences potentially are an advantage of the HRS data since, as noted in De Nardi et al. (2016), the MCBS clearly understates aggregate Medicare and especially Medicaid spending, potentially due to the issue that the MCBS does not have administrative data on Medicaid spending, and thus relies heavily on imputation.

The next set of benchmarking exercises that we perform is for out-of-pocket medical spending, Medicaid recipiency and income between the AHEAD cohort of the HRS and MCBS. For both the HRS and MCBS, we restrict the sample to singles (over the sample period) who meet the HRS/AHEAD age criteria (at least 70 in 1994, 72 in 1996, ...) and who are not working over the sample period. Because the MCBS sample lacks spousal information, for this analysis we focus

	Total Spending				OOP			
	HRS		MCBS		HRS		MCBS	
Total Spending Percentiles	Average Expenditure	Percentage of Total	Average Expenditure	Percentage of Total	Average Expenditure	Percentage of Total	Average Expenditure	Percentage of Total
All	17,091	100	14,120	100	3,825	100	2,740	100
95-100%	114,238	33.4	97,880	34.6	45,643	59.6	26,930	49.1
90-95%	59,000	17.3	48,890	17.3	8,619	11.3	6,700	12.2
70-90%	26,870	31.4	20,540	29.1	3,480	18.2	2,920	21.3
50-70%	9,025	10.6	7,750	11	1,394	7.3	1,360	9.9
0-50%	2,502	7.3	2,250	8	178	3.6	420	7.6

	Medicare				Medicaid			
	HRS		MCBS		HRS		MCBS	
Total Spending Percentiles	Average Expenditure	Percentage of Total	Average Expenditure	Percentage of Total	Average Expenditure	Percentage of Total	Average Expenditure	Percentage of Total
All	11,343	100	7,720	100	1,896	100	1,320	100
95-100%	85,268	37.6	67,560	43.7	33,773	89.1	24,980	94.7
90-95%	41,731	18.4	28,370	18.4	4,092	10.8	1,360	5.2
70-90%	17,251	30.4	10,280	26.6	230	0.2	10	0.1
50-70%	5,031	8.9	2,980	7.7	0	0	0	0
0-50%	1,076	4.7	550	3.5	0	0	0	0

Table 6: Individual Medical Spending Percentiles: HRS versus MCBS

	HRS/AHEAD Data				MCBS Data		
Income Quintile	Total Income	Annuity Income	Out-of-pocket Expenses	Medicaid Recipiency	Total Income	Out-of-pocket Expenses	Medicaid Recipiency
Top	33,580	26,300	7,000	3.0	44,150	8,020	5.4
4th	19,290	14,390	6,360	5.6	19,710	7,300	8.0
3rd	15,500	10,900	5,050	11.0	13,740	6,470	15.5
2nd	10,290	8,270	4,270	28.1	10,020	5,340	41.8
Bottom	7,740	4,820	2,550	60.9	6,750	4,050	69.9

Source: Table A.2 of De Nardi, French, and Jones (2016).
Notes: Calendar years 1996-2010, for those age 72 and older in 1996.

Table 7: Income, Out-of-pocket Spending, and Medicaid Recipiency Rates: HRS versus MCBS

only on singles. We use the De Nardi, French, and Jones (2016) measure of permanent income and construct a measure of permanent income that is the percentile rank of total income over the period we observe these individuals (the MCBS asks only about total income). The first four columns of Table 7 show sample statistics from the full HRS/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 HRS/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 HRS/AHEAD. Next, we compare out-of-pocket medical spending in the MCBS and HRS/AHEAD. Out-of-pocket medical spending (including insurance payments) averages $2,360 in the bottom PI quintile and $6,340 in the top quintile in the HRS/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 HRS/AHEAD are similar, which may be surprising given that the two surveys each have their own advantages in terms of survey methodology.[11] The share of the population receiving Medicaid transfers is also very similar in the HRS/AHEAD and MCBS. Sixty-one percent and 70 percent of those in the bottom PI quintile are on Medicaid in the the HRS/AHEAD and MCBS, respectively. In the top quintile, 3% of people are on Medicaid in the HRS/AHEAD whereas 5% are in the MCBS.

[11]There are more detailed questions underlying the out-of-pocket medical expense questions in the HRS, 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.

Appendix F: Our Lifetime Earnings Data

We link restricted Social Security Administration (SSA) earnings data for the years 1951-2013 to our survey data for HRS respondents who consented. In this appendix we describe our earnings data and how we handle it. We can combine husbands' and wives' earnings records to create household earnings. In addition, we also have earnings records of deceased spouses through the Deceased Spouse Earnings file.

The SSA earnings data are used to calculate Social Security benefits, and are thus of high quality, with little measurement error or attrition. However, for much of the sample period, earnings were taxable only up to a cap. We describe these issues, as well as how these issues are addressed, below. A longer description of the data, and how the data are handled, is described in Fang (2018).

Earnings can be broken down into Social Security "covered earnings" or "noncovered earnings", where covered earnings refers to earnings in a sector of the economy where workers make mandatory Social Security contributions. Our data only includes covered earnings. We have covered earnings data for wage earners in commerce and industry, non-farm and non-professional self-employed workers from 1951. Between 1955 and 1966 more groups were covered, such as farm and professional self-employed workers. Currently the great majority of all workers are covered apart from some workers in state government.

The raw covered SSA earnings data are top-coded at the taxable maximum in earlier years. This is because only summary earnings records, which come from Social Security contributions, were available between 1951 and 1977; detailed records, which use information from W-2 and 1040 Schedule SE records, have been available since 1978. Thus after 1978 we have non-top-coded earnings information for employed indivduals, although for self-employed individuals the data are top-coded until 1994.[12] Between 1978 and 1993 the nearest neighbor matching method was used to impute self-employment earnings, and since 1994, self-employment earnings can be treated as not being top-coded.

SSA earnings records were "masked" if above $250,000, and placed in intervals of between $250,000 and $300,000, $300,000 and $500,000, and above $500,000. Masked values were dealt with using the methods below.

In 1951 the taxable maximum was $3,600 and earnings above this level are top-coded. This maximum rose more rapidly than earnings prior to 1978, making it less of an issue over time.

[12] For some individuals who were initially interviewed in 1992 or 1993, the HRS only has their summary records until 1991 or 1992 because they consented when initially interviewed but then never re-consented.

During this period, information on quarters of coverage (QC) is also available and helps us to impute top-coded earnings. Individuals could earn up to four QC a year. The "pattern of coverage" reports which quarter the taxable maximum was hit. Assuming that earnings are equal in all quarters, we can infer earnings up to four times the taxable maximum–above that level we will only know that their earnings exceeded the taxable maximum in the first quarter of the year. Earnings between the taxable maximum and 200% of the taxable maximum are imputed using a log-normal distribution and a Pareto distribution is used beyond that.

For those who reached the taxable maximum through self-employment earnings alone and also those individuals who never re-consented after 1992/3, the "pattern of coverage" was not available and a different imputation strategy was used.

As described below, wage earners and the self-employed face different payroll tax rates. We cannot distinguish earnings from covered employment and covered self-employment when an individual has earnings from both prior to 1977, meaning that we may miss self-employment earnings for those whose employment earnings exceeded the taxable maximum. However, according to the Census Bureau's Consumer Income Publication Series, in 1966, only 5.7% of all families earned only self-employment income and 10.6% of all families' earnings included both self-employment and wages or salary income. For most of those families, self employment income is likely not top-coded, and if it is, the amount is likely modest. Therefore, we assume that the self employed pay the employed tax rate before 1977.

Appendix G: Calculating Tax Payments in the HRS Data

As noted in Appendix A, the Medicare Part A is financed using a payroll tax, whereas Part B Part D and Medicaid are financed through general revenues from federal and state taxes. Below we describe how we calculate tax payments.

Capital income, Pension and Social Security Benefit Predictions

We observe capital income, pension and Social Security benefits during the HRS survey years 1992-2016. For the year intervals 1951-1991 and 2017-2066, we predict them using a fixed-effect regression. In the regression we control for household age using a third order polynomial and for marital status. These regressions have R^2 statistics of 0.41 for Social Security benefits, 0.010 for pension benefits and 0.002 for capital income.

Payroll Taxes

Federal Insurance Contributions Act (FICA) taxes for Medicare Part A (Hospital Insurance) were first collected in 1966 and amounted to 0.7% of wages up to the taxable maximum, with employers and employees each paying 0.35% of wages. We assume that the employer contribution is ultimately borne by the worker. The payroll tax rate has increased over time, and is currently 2.9% with employers and employees each paying 1.45% of wages. Prior to 1984 the self employed only had to the employee part of the tax: afterwards they had to pay both the employee and employer shares. Prior to 1990 the Medicare and Social Security taxable maximum was the same, and was $51,300 in 1990. In 1991 the taxable maximum for Medicare jumped to $125,000 whereas for Social Security it rose with wage growth to $53,400. In 1994 the taxable maximum was repealed; Medicare Hospital Insurance payroll taxes are now paid on all wage and self-employment earnings. We apply these tax rates, τ_t^m, to earnings histories in each year t to calculate payroll tax contributions.

Federal and State Taxes

We calculate federal taxes, T_{it}^f, for individual i in each year t using TAXSIM (Feenberg and Coutts 1993) based on household income, marital status, number of taxpayers over age 65, and number of dependents. We assume that all households take the standard deduction, and for years after 2023 we apply the 2023 income tax schedule. For household income we include all earnings, self-employment, capital, pension and Social Security income. As we lack data on other components of income that incur federal tax, such as rental or dividend income, we understate tax payments. We assume that taxes to pay for Medicare and Medicaid are paid on a Pay as You Go basis, and thus today's taxpayers pay for today's beneficiaries. As such, we assume that a share ψ_t^f of federal taxes in year t go to paying Medicare Part B and D, and and also Medicaid. To do this for each year we calculate ratio of expenditures on these programs for the over 65s to federal tax collections. For state taxes, T_{it}^s, we also apply TAXSIM using the same method as above. Although we do not know the state in which each household resides, we know the household's Census Bureau division after 1990. If we do not observe the Census Bureau division for a household in a given year (e.g., before 1990), we assign a household to the census division in which they reside the first time they are observed in the data. After calculating the possible state tax for each household if they resided in every state in their census division, we calculate their state tax as a weighted average of the state taxes for states in their census division, based each state's population size in 1960. TAXSIM only produces state tax information from 1977 onwards. As with the federal tax calculations, we assume that a share ψ_t^s taxes in year t go towards paying Medicaid for those over 65.

In practice, as we are working with restricted data, we do not run TAXSIM directly on our data but instead on a large simulated dataset with millions of observations. The resulting federal and state taxes from the simulated data are then merged to the restricted dataset based on the listed characteristcs.

Appendix H: Computing Permanent Income

We assume that log income after age 65 evolves as following

$$\ln y_{it} = \kappa(t, f_{it}) + h(I_i) + \omega_{it}, \tag{14}$$

where $\kappa(t, f_{it})$ is a flexible function of age t and family structure f_{it} (i.e., couple, single man or single woman) and ω_{it} represents measurement error. The variable I_i is the household's percentile rank in the permanent income (PI) distribution. Since it is a summary measure of lifetime income at retirement, it should not change during retirement and is thus a fixed effect over our sample period. However, income could change as households age and potentially lose a family member.

To estimate equation (14) we first estimate the fixed effects model

$$\ln y_{it} = \kappa(t, f_{it}) + \alpha_i + \omega_{it}, \tag{15}$$

which allows us to obtain a consistent estimate of the function $\kappa(t, f_{it})$. Next, note that as the number of time periods over which individual i is observed (denoted T_i) becomes large,

$$\operatorname*{plim}_{T_i \to \infty} \frac{1}{T_i} \sum_{t=1}^{T_i} \left[\ln y_{it} - \kappa(t, f_{it}) - \omega_{it}\right] = \frac{1}{T_i} \sum_{t=1}^{T_i} \left[\ln y_{it} - \kappa(t, f_{it})\right] = \alpha_i = h(I_i). \tag{16}$$

We thus calculate the PI ranking I_i for every household in our sample by taking the percentile ranking of $\frac{1}{T_i}\sum_{t=1}^{T_i}[\ln y_{it} - \widehat{\kappa}(t, f_{it})]$, where $\widehat{\kappa}(t, f_{it})$ is the estimated value of $\kappa(t, f_{it})$ from equation (15). Put differently, we take the mean residual per person from the fixed effects regression (where the residual includes the estimated fixed effect), then take the percentile rank of the mean residual per person to construct I_i.

However, we also need to estimate the function $h(I_i)$, which converts the estimated index I_i back to a predicted level of income that can be used in the dynamic programming model. To do this we estimate the function

$$[\ln y_{it} - \kappa(t, f_{it})] = h(I_i) + \omega_{it} \tag{17}$$

where the function $h(I_i)$ is a flexible functional form. In practice we model $\kappa(t, f_{it})$ as a third order polynomial in age, dummies for family structure, and family structure interacted with an age trend. When estimating equation (17), we replace the function $\kappa(t, f_{it})$ with its estimated value. We model $h(I_i)$ as a fifth order polynomial in our measure of permanent income percentile.

Given that we have, for every member of our sample, t, f_{it}, and estimates of I_i and the functions $\kappa(.,.), h(.)$, we can calculate the predicted value $\ln \widehat{y}_{it} = \widehat{\kappa}(t, f_{it}) + \widehat{h}(\widehat{I}_i)$. It is $\ln \widehat{y}_{it}$ that we use when simulating the model for each household. A regression of $\ln y_{it}$ on $\ln \widehat{y}_{it}$ yields a R^2 statistic of 0.74, suggesting that our predictions are accurate.

Appendix I: Estimation Methodology

We estimate the model using the extension of the *E-M* algorithm employed by Arellano, Blundell, and Bonhomme (2017). Recall from equations (10)-(12) that the functions $Q_\eta(\tau \mid \eta_{i,t-1}, a_{i,t})$, $Q_1(\tau \mid a_{i,1})$ and $Q_\varepsilon(\tau \mid , a_{i,t})$ are constructed from Hermite polynomials $\left(\{h_k^\eta(\cdot)\}_{k=0}^{K_\eta}, \{h_k^1(\cdot)\}_{k=0}^{K_1}, \{h_k^\varepsilon(\tau)\}_{k=0}^{K_\varepsilon}\right)$, using the coefficient functions $\{\beta_k^\eta(\tau)\}_{k=0}^{K_\eta}$, $\{\beta_k^1(\tau)\}_{k=0}^{K_1}$, and $\{\beta_k^\varepsilon(\tau)\}_{k=0}^{K_\varepsilon}$. The coefficient functions are in turn modeled with a set of polynomial splines defined over the intervals $\left\{[\tau_{\ell-1}, \tau_\ell]\right\}_{\ell=1}^{L}$, along with two tail functions for $(0, \tau_1]$ and $[\tau_L, 1)$. It is the parameters for these weighting functions that we must estimate.

Define θ as the vector of all parameters (the β parameters) in equations (10)-(12). The procedure to estimate θ is as follows. Starting with the vector $\hat{\theta}^{(0)}$ we iterate between the following two steps until $\hat{\theta}^{(j)}$ converges:

1. *Stochastic E-Step*: For each observation i, draw S values of $\eta_i^{(s)} = (\eta_{i1}^{(s)}, ..., \eta_{iT}^{(s)})$ from $f_i(\,.\,; \hat{\theta}^{(j)})$ (derived from $Q_1^{(j)}(\cdot)$, $Q_\eta^{(j)}(\cdot)$ and $Q_\varepsilon^{(j)}(\cdot)$).

2. *M-step*: Find

$$\underset{\beta^\eta_{\ell 0}, \ldots, \beta^\eta_{\ell K_\eta}}{\text{argmin}} \quad \sum_{i=1}^{N} \sum_{s=1}^{S} \sum_{t=2}^{T} \rho_{\tau_\ell}\left(\eta_{it}^{(s)} - \sum_{k=1}^{K_\eta} \beta^\eta_{\ell k}\, h_k^\eta(\eta_{i,t-1}^{(s)}, a_{it}) \right), \quad \ell = 1, ..., L.$$

We use $\rho_\tau(\cdot)$ to denote Koenker and Bassett Jr's (1978) quantile "check" function. To identify the full set of splines, this function is minimized at each point ℓ on the grid over τ. The coefficients for $\varepsilon_{i,t}$ and $\eta_{i,1}$ likewise solve

$$\underset{\beta^{\varepsilon}_{\ell 0},\ldots,\beta^{\varepsilon}_{\ell K_\varepsilon}}{\operatorname{argmin}} \sum_{i=1}^{N}\sum_{s=1}^{S}\sum_{t=1}^{T} \rho_{\tau_\ell}\left(m_{i,t} - \eta_{it}^{(s)} - \sum_{k=1}^{K_\varepsilon} \beta^{\varepsilon}_{\ell k}\, h^{\varepsilon}_k(a_{it}) \right), \quad \ell = 1, ..., L,$$

$$\underset{\beta^{1}_{\ell 0},\ldots,\beta^{1}_{\ell K_1}}{\operatorname{argmin}} \sum_{i=1}^{N}\sum_{s=1}^{S} \rho_{\tau_\ell}\left(\eta_{i1}^{(s)} - \sum_{k=1}^{K_1} \beta^{1}_{\ell k}\, h^{1}_k(a_{i1}) \right), \quad \ell = 1, ..., L.$$

There are also moment conditions related to the tails of the distribution: See Arellano, Blundell, and Bonhomme (2017). These estimates give us $\hat{\theta}^{(j+1)}$.

For longer panels, settings with unbalanced data, or when estimating more complicated models the *E-step* can perform poorly when using standard samplers (e.g., Metropolis-Hastings). We therefore employ the sequential Monte-Carlo (SMC) approach implemented by Arellano et al. (2021). Comprehensive surveys of these methods can be found in Doucet et al. (2009) and Creal (2012).

Step 1: SMC Stochastic E-Step to sample from $f(\eta_{i,1}, ..., \eta_{i,T} | Y_i^T, a_i^T)$. For $i = 1, ..., N$:

At $t = 1$:

1. Sample S particles $\eta_1^{(s)} \sim g(\eta_1 | y_1)$, where $g(\cdot)$ is the closed form posterior from the Gaussian model.
2. Compute the weights $w_1(\eta_1^{(s)})$ and apply a self-normalization to obtain $W_1^{(s)} \propto w_1(\eta_1^{(s)})$.
3. If $Var(W^{(s)})$ exceeds some threshold, re-sample $\{W_1^{(s)}, \eta_1^{(s)}\}$ to obtain S equally weighted particles.

At $t > 1$:

1. Sample S particles $\eta_t^{(s)} \sim g(\eta_t | \eta_{t-1}, y_t)$, where $g(\cdot)$ is the closed-form posterior from the Gaussian model.
2. Compute the weights $w(\eta_{1:t}^{(s)})$ and apply a self-normalization to obtain $W_t^{(s)} \propto w_t(\eta_{1:t}^{(s)})$.
3. If $Var(W^{(s)})$ exceeds some threshold, re-sample $\{W_t^{(s)}, \eta_t^{(s)}\}$ to obtain S equally weighted particles.
4. If $t = T$, sample P particles to be used in the *M-Step*. (We set $P = 1$).

Step 2: M-Step

1. Update quantile regressions for equations (10), (11) and (12) .
2. Update Laplace parameters for the tail functions.

3. Update parameters for Gaussian proposal distributions.

Appendix J: Demographic Transition Probabilities in the HRS

Let $hs_{i,j,t} \in \{0,1,2,3\}$ denote death ($hs_{i,j,t} = 0$) and the 3 mutually exclusive health states of the living (nursing home = 1, bad = 2, good = 3, respectively) of household member j, household i, time t. Let $x_{i,j,t}$ 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 $q \in \{1,2,3\}$, $r \in \{0,1,2,3\}$, we rewrite equation (13) as

$$\begin{aligned}
\pi_{q,r,t} &= \Pr(\, hs_{i,g,t+1} = r \,|\, hs_{i,g,t} = q; x_{i,g,t}) \\
&= \gamma_{qr} \Big/ \sum_{s\in\{0,1,2,3\}} \Big(\gamma_{qs}, \\
\gamma_{qs} &\equiv 1, \quad s = 0 \\
\gamma_{qs} &= \exp\left(x_{i,g,t}\beta_s\right), \quad s \in \{1,2,3\},
\end{aligned}$$

where $\{\beta_s\}_{s=1}^{3}$ are coefficient vectors for each future state s and $x_{i,g,t}$ is the explanatory variable vector which depends on the current state q.

Appendix K: Additional Tables and Figures

Table 8 shows the coefficient estimates for the logit model of Medicaid receipt.

	(1)	(2)
log(total medical spending)	0.470	0.453
	(0.07)	(0.11)
$\log^2$(total medical spending)	-0.006	-0.007
	(0.004)	(0.006)
PI percentile	-13.012	-10.471
	(0.43)	(0.67)
PI percentile^2	2.982	1.502
	(0.23)	(0.38)
Age	-1.021	-2.136
	(0.25)	(0.46)
Age^2	0.0105	0.0243
	(0.003)	(0.006)
Age^3	-3.55×10^{-5}	-9.15×10^{-5}
	(1.25×10^{-5})	(2.26×10^{-5})
Woman in bad health	0.622	0.414
	(0.03)	(0.05)
Man in bad health	0.564	0.535
	(0.04)	(0.07)
Woman in nursing home	1.026	0.621
	(0.05)	(0.09)
Man in nursing home	1.570	1.148
	(0.08)	(0.14)
Couple	-0.255	-0.019
	(0.06)	(0.11)
Single male	0.612	0.624
	(0.03)	(0.06)
log(total medical spending)×PI	0.577	0.607
	(0.04)	(0.07)
Couple×PI	-1.806	-1.532
	(0.15)	(0.24)
Lagged Medicaid receipt		5.972
		(0.06)
Constant	28.261	55.749
	(6.90)	(12.39)
Observations	89,604	89,604
Pseudo-R^2	0.349	0.739

Table 8: Probability of Receiving Medicaid for over 65s, Logit Specification

	Lifetime Income and Contributions			Medical Spending (Age 65+)			
PI	Income	Earnings	Contributions	Total	Medicare	Medicaid	OOP+Private+Other
All	*3,306*	*1,632*	*107*	*536*	*296*	*31*	*177*
Top	8,003	3,824	248	611	396	5	204
4th	4,082	2,300	131	587	354	9	216
3rd	2,445	1,437	82	561	321	15	212
2nd	1,310	641	47	516	272	36	176
Bottom	906	56	33	406	137	92	79

Table 9: Earnings, Tax Contributions and Medical Spending (in 000s) by Permanent Income Quintile