Elderly Health and Longevity in The US: Evidence and Implications

Liran Einav and Amy Finkelstein*

April 14, 2026

Abstract

Rising elderly life expectancy is a well-known source of fiscal pressure on Social Security and Medicare – but how have reductions in elderly mortality and morbidity affected the *relative* finances of these two programs? Using nearly three decades of Medicare Current Beneficiary Survey data (1992-2019), we estimate that these demographic changes raised expected lifetime Social Security spending for a 66-year-old by 14%, but Medicare spending by only 6%. This differential growth reflects two features of how longevity has increased: the additional 2.4 years (14% increase) of remaining life expectancy at age 66 were entirely healthy years – free of physical or cognitive limitations – while time spent with severe health limitations actually declined by about 30%, pulling down lifetime nursing-home and home-health spending. We then write down a stylized life-cycle model of a risk-averse retiree facing stochastic mortality and health to analyze the optimal allocation of public funds across these two programs, and to illuminate the key tradeoffs involved in allocating a fixed budget across these two distinct forms of consumption-smoothing.

*Einav: Stanford University and NBER, leinav@stanford.edu; Finkelstein: MIT and NBER, afink@mit.edu. We are grateful to Ro Huang, Ken Lin, and Sam Wolf for exceptional research assistance, to Kate Bundorf, David Cutler, Lee Lockwood, Petra Persson, Jim Poterba, Mark Shepard, and seminar participants at CEPR Ageing and Longevity, Columbia-NYU, Duke, Institute for Fiscal Studies, NBER Aging, NBER Public Economics, Singapore Management University, Stanford, Toulouse, and Whistler for helpful comments, and to David Cutler for data access.

1 Introduction

Over the last several decades, adult life expectancy has increased considerably in most high-income countries. While naturally a major source of improved well-being, this has also put substantial fiscal pressure on the public programs that provide pensions and health insurance to the elderly. In the United States, between 1990 and 2019 real per-capita public pension spending on annuities (Social Security) roughly doubled, while real per-capita public health insurance spending on the elderly (Medicare) quadrupled.[1] By 2024, these two large entitlement programs together accounted for over one third of federal expenditures (CBO 2025). There has been considerable attention to the fiscal pressure on both programs created by demographic trends, as well as discussion of possible entitlement reforms, such as raising the age of benefit eligibility, cutting benefit levels, and/or raising payroll taxes on the working population (e.g., National Research Council et al. 2013; Poterba 2014; Sheiner 2021). Yet, there has been much less attention paid to how demographic changes have affected the *relative* finances of Social Security and Medicare – let alone how to think about the optimal allocation of public funds across the two programs. This paper begins to fill these gaps.

Social Security and Medicare were created at different times, are administered by different agencies, and have separate budgets – and so are typically analyzed independently. Yet both programs provide social insurance to primarily the same population of the elderly and for the same basic purpose of consumption smoothing over stochastic mortality and health risks. Moreover, the design of both programs has been essentially fixed for at least the last several decades. We first show how demographic changes in elderly mortality and morbidity have shifted the relative allocation of spending between the two programs, and then write down and calibrate a stylized dynamic consumption-saving model to illustrate the economic forces affecting the optimal split.

The backbone of our analysis is almost three decades (1992-2019) of data from the Medicare Current Beneficiary Surveys (MCBS). This annual survey covers approximately 10,000 panelists each year in a rotating panel, with each participant followed for 3 to 4 years. It contains detailed individual-level information on 31 health conditions and functional limitations, cross-validated data on overall and category-specific annual health-care spending, as well as linkages to administrative data to individuals' dates of death through 2019. We focus primarily on changes between the early part of our data (1993) and the later part (2017).

We start, in Section 2, by replicating and extending prior findings of increased elderly life expectancy and improved elderly health conditional on age. To do so, we first classify the more granular, 31-dimensional data on health into five coarser (ranked) health states (that is, morbidity categories); this allows us to use the panel element of the MCBS, together with the linked information on subsequent death, in order to track how health evolves over time. Specifically, we estimate year-by-year, age-specific 5-by-6 transition matrices between the five coarse health states (and an absorbing state of death). We use these data to analyze, for a given calendar year, the life of a hypothetical 66 year old who faces the age-specific annual rates of morbidity and mortality experienced by the cross-section of

[1] Authors' calculations based on data from the Medicare Current Beneficiary Survey and the Social Security Administration (SSA 2025). Medicare spending includes both Traditional Medicare and Medicare Advantage spending among those 65 years old and older.

individuals at different ages.

Between 1993 and 2017, remaining life expectancy for the average 66 year-old increased by 2.4 years, from 17.4 to almost 20 years, increasing expected lifetime Social Security payments by about 14%. Heuristically, this increase in life expectancy can be attributed to some combination of "delayed" aging (individuals remaining healthier for more years) and "prolonged" dying (individuals surviving for more years in an unhealthy state through, say, technological improvements in extending the life of the sick). Strikingly, we find that the health trajectory in the years *immediately preceding* death has remained largely unchanged, and that the 2.4-year increase in life expectancy consists *entirely* of increases in healthy life-years; that is, expected years spent in the healthiest of the five morbidity categories – free of any of the physical or cognitive functional limitations that commonly accrue with advanced age. Moreover, expected time spent in the worst of the five morbidity states, in which individuals have severe physical or cognitive limitations, actually *declines* by 0.6 years (about 30%).

In Section 3 we develop and implement an approach for translating the observed improvements in morbidity and mortality into implications for expected lifetime health-care spending for a 66 year old. In particular, given the improvement in health conditional on age over our study period, we want to distinguish the impact of health (which naturally deteriorates with age) on health-care spending from any "direct," residual impact of age, conditional on health. We therefore use the 31-dimensional granular health measures (explicitly excluding age) as inputs to train a random forest algorithm to predict annual spending (both overall and by spending category). We then fit the residuals from this prediction model to a (proportional) age-specific shifter to capture any residual impact of age on health-care spending. Interestingly, we find that conditional on the granular health measures, there is little incremental role for age in predicting health-care spending; in other words, the role of age in driving health-care spending can be largely accounted for by observed health changes with age.

We use these estimates to analyze how the estimated changes in age-specific morbidity and mortality transitions have affected (predicted) health-care spending, holding the mapping from health and age to health-care spending fixed. This allows us to isolate the role of the changes in morbidity and mortality patterns, and abstract from the many changes in the practice and pricing of medicine – due to changes in technology or health-care policy – that occurred over the same period. The combination of increased life expectancy and improved health conditional on age between 1993 and 2017 increased lifetime health-care spending per elderly individual by about 6%. Increases in expected lifetime health-care spending are an order of magnitude greater for men than women, reflecting both larger increases in life expectancy for men and larger declines in years spent in really poor health for women. The increase in spending for higher-income individuals is more than twice as large as the increase for lower-income individuals, which is primarily driven by higher-income individuals experiencing greater increases in life expectancy.

Contrary to conventional wisdom that rising elderly life expectancy and the increasing prevalence of the "oldest old" will put substantial upward pressure on the use of nursing homes and other long-term care (e.g., Gruber et al. 2023), we estimate that the net impact of changes in mortality and morbidity is, in fact, to (slightly) *decrease* spending (per elderly individual) on nursing-home care and home-health care over our sample period. This stems from the fact that morbidity, rather than age, is the key predictor of nursing-home and home-care use; in particular, expected nursing-home and home-care use is heavily concentrated in the poorest morbidity state (the worst of the five morbidity states), regardless of age, and the expected number of years in this morbidity state actually declines, despite the overall increase in remaining life expectancy.

Taken together, the demographic changes in mortality and morbidity increased lifetime public spending on the average elderly individual between 1993 and 2017 by about 11%, with expected Social Security spending growing more than twice as fast as expected Medicare spending (14% compared to 6%). These increases reflect the "mechanical" impact of demographic changes under passive, status-quo program rules. In the final part of the paper, we sketch a simple dynamic model that allows us to analyze forces that affect the optimal allocation of a given public budget across these two social insurance programs. Specifically, we consider a standard life-cycle consumption model for a risk-averse retiree, who benefits from a fixed Social Security annuity income and government-funded health insurance while facing stochastic mortality and stochastic morbidity-induced medical spending.[2] We use this model to analyze consumer welfare under alternative budget-neutral policy parameters that re-allocate expected lifetime public spending across Social Security and Medicare.

The model clarifies two advantages of an incremental expected dollar of Medicare coverage relative to Social Security. First, while Social Security annuity payments only help individuals to reduce consumption volatility across states of the world in which they are alive, Medicare coverage also helps reduce consumption volatility across states of differential health realizations, and therefore provides a more complete insurance product. Second, as we show in Einav and Finkelstein (2026), the optimal public annuity is back-loaded while current Social Security payments are constant with respect to age. This creates a second advantage for Medicare in our context: because health deteriorates with age, Medicare coverage naturally delivers more resources to older individuals, providing an indirect form of back-loading payments. As a result, the optimum is to allocate the government budget entirely to Medicare until Medicare covers all medical expenses, and only then allocate any remaining funds to Social Security. To arrive at an interior solution, we also allow for moral hazard effects of health insurance. This creates a disadvantage for Medicare relative to Social Security, since it induces inefficient medical spending. We then illustrate via stylized calibrations how the demographic changes we estimate can affect this interior solution.

Our paper contributes to a number of related literatures. The first part of our paper is closely related to a large existing literature that goes back almost five decades, and has investigated whether increased elderly life expectancy has meant that the elderly are experiencing additional years in poor or good health (e.g., Gruenberg 1977; Fries 1980; Manton 1982; Fuchs 1984). These studies typically construct a "health deficit index" (also called a "frailty index") which measures the (unweighted) fraction of a list of aging-related health conditions that an individual has, and has been shown to be predictive of mortality. In general, the findings have indicated improved health conditional on age, both in the US (Freedman et al. 2002, 2013; Cutler et al. 2014; Abeliansky et al. 2020; Cutler et al. 2022) and in other high-income countries (e.g., Abeliansky and Strulik 2019; Utkus and Mitchell

[2] As is common in these models, we also assume a government-funded guaranteed minimum consumption floor, which approximates Medicaid.

2025), prompting aphorisms such as "70 is the new 60" (National Research Council et al. 2013; Lee 2016).[3] The spirit of our analysis in the first part of the paper is similar, and our findings are generally consistent with this existing literature. However, since our primary focus is on translating health and age into implications for health-care spending, we pay much closer attention to how different health conditions contribute to spending, and therefore our spending index is richer and places more weight on "expensive" health conditions that are associated with greater resource utilization.

There has been considerably less attention to the implications of the improvement in morbidity for medical spending. Qualitatively, it has long been recognized that rising life expectancy of the elderly increases their expected lifetime health-care spending, while improvements in health will move their health-care spending in the opposite direction (Poterba and Summers 1987). Quantitative analyses, however, have been rare. An important exception is Lubitz et al. (2003) who, in an approach closely related in spirit to ours, analyze MCBS data from 1992-1998 on health conditions, health transitions, and health-care spending to conclude that 70-year-old individuals who are in poorer health have similar expected lifetime spending to 70-year-old individuals who are in better health, due to the offsetting effects of greater longevity and better health on health-care spending. In a similar spirit, we replicate prior findings that relative to lower-socioeconomic-status individuals, higher-socioeconomic-status individuals have experienced greater increases in life expectancy and greater decreases in morbidity (e.g., Case and Deaton 2015; National Academies of Sciences 2015; Chetty et al. 2016; Case and Deaton 2020; Hudomiet et al. 2021, 2022), but we also estimate that they have experienced similar changes in lifetime health-care spending since these two forces offset each other.

Our model-based exercise, which studies the optimal allocation of public funds between Social Security and Medicare, is closely connected to an extensive literature in public finance that uses this type of quantitative modeling to study the impact of different public-policy interventions aimed at the elderly. Methodologically, our approach follows the well-established tradition of using calibrated, life-cycle models of individuals facing stochastic mortality and morbidity to model elderly behaviors, such as savings, bequests, and insurance demand (e.g., Mitchell et al. 1999; Scholz et al. 2006; De Nardi et al. 2010; Lockwood 2018). More closely related to our substantive analysis are papers that use calibrated life-cycle models to analyze various social insurance programs, such as Medicaid (e.g., Brown and Finkelstein 2008; De Nardi et al. 2016), Medicare (e.g., McClellan and Skinner 2006; Khwaja 2010), and Social Security (e.g., French 2005; Van der Klaauw and Wolpin 2008). While many of these papers model multiple social insurance programs, they tend to focus their analysis on the behavioral impact, the value, or the optimal design of one particular program, rather than on the relative value of different programs, which is our focus.

[3] The underlying causes of these improvements in mortality and morbidity are undoubtedly multi-faceted and complex. Nonetheless, estimates suggest that over the 1990-2015 time period, about two-fifths of the increase in life expectancy reflects changes in behavior (such as smoking and diet), and about one-third reflects the development and adoption of new pharmaceutical therapies, and only about 13% is attributed to intensive medical treatment for acute diseases (Buxbaum et al. 2020).

2 Data and descriptive patterns in health

2.1 Data and key health measures

Below we briefly describe the data and key variables. More details are provided in Appendix A.

Medicare Current Beneficiary Survey. We use almost three decades of data, from 1992 through 2019, from the Medicare Current Beneficiary Survey (MCBS).[4] Every year, the MCBS surveys about 10,000 individuals enrolled in Medicare (covering both Traditional Medicare and Medicare Advantage enrollees), including residents of nursing homes and other long-term care facilities. It operates as a rotating panel with a three-year or four-year period of inclusion. The panel dimension will be crucial for estimating transition probabilities across health states.

The MCBS essentially consists of two distinct surveys. One is a sample of people who are representative of the population that was enrolled in Medicare from the beginning of the calendar year through their interview date, which is typically in the fall. The second is a sample of people who are representative of Medicare beneficiaries who were enrolled for any portion of the survey year.

For our analysis purpose, there are four key components of the data. First, for both surveys, the MCBS collects basic demographic characteristics including age as of January 1 of the survey year, gender, education, and income. Second, both surveys contain a range of health-related survey questions, which we describe in more detail below.[5] Third, the second survey contains detailed information on annual health-care expenditure – overall and by type of care – across all payers. Finally, for the first survey, we have information on the beneficiary's date of death (if any) through 2019. The first three elements are available through typical MCBS data access while the last one requires additional permissions.

Sample definition. Since our exercise requires both data on mortality and data on health-care spending, we use the intersection of the people who appear in both MCBS samples described above. We restrict our analysis to Medicare recipients aged 66 and over. These two restrictions produce a sample of 226,647 individual-years, or an average of 8,384 individuals per year. In all of our analyses, we use the survey's individual-level survey weights, adjusted so that the sum of the weights is consistent in each year.[6] For most of our analyses, we compare data from the (pooled) first three years of our sample (1992-1994) to the pooled last three years (2016-2018), and occasionally to the middle three years (2011-2013); throughout we (respectively) refer to these samples as 1993, 2017, and 2012.

[4]The MCBS has been conducted every year since 1991, with the exception of 2014, and our sample covers 1992 through 2019, without 2014.

[5]The MCBS typically administers its questions directly to the individual of interest, but will survey family members if an individual is incapacitated, or facility staff if the individual is a resident of a long-term care facility.

[6]Specifically, as recommended by CMS (2021), we use the survey weights from the sample of individuals who are representative of people who were enrolled in Medicare for any portion of the survey year as weights when analyzing the intersection of the two samples.

Health-care spending. The MCBS collects self-reported health-care spending by category and payer.[7] The MCBS validates the self-reports against administrative data on health-care claims and utilization and uses this comparison to adjust its estimates of spending. We collapse all health spending data to the person-year level and standardize all expenses to 2024 dollars.[8] For some of our analyses we dis-aggregate spending into specific sub-categories. Specifically, we analyze inpatient, nursing home, home-health care, prescription drug, and the residual category which we call physician/outpatient.[9]

Granular morbidity measures. The MCBS provides a number of different morbidity measures, including limitations to six Activities of Daily Living (ADLs), limitations to six Instrumental Activities of Daily Living (IADLs), five functional limitations, and whether the individual has ever been diagnosed with a range of different diseases. Appendix Table OA.1 lists the full set of 31 ADLs, IADLs, functional limitations, and diseases employed in our analysis, as well as the 2017 sample mean for each of these measures. Throughout the paper, we use h'_{it} to denote this 31-dimensional vector of binary indicators associated with individual i in year t.

Limitations to ADLs measure individuals' ability to independently carry out physical tasks that are typically part of everyday life, such as eating, dressing, and bathing. Limitations to IADLs measure individuals' ability to carry out somewhat more complex tasks that may require greater executive functioning skills, such as managing money, shopping, and preparing meals. Each one is a binary measure based on whether the individual reports experiencing any difficulties performing the given task. ADLs and IADLs are common measures of disability across the literature on aging and morbidity (e.g., Freedman et al. 2002; Cutler et al. 2014).[10]

Functional limitations consist of comparatively more difficult physical activities than ADLs that may or may not come up in the course of everyday life, such as lifting ten pounds, raising one's hands above their head, and walking a quarter of a mile. Unlike the binary nature of ADLs and IADLs, the functional limitation questions are asked on a sliding scale ranging from "no difficulty" to complete inability to perform the task; we convert these to binary measures by recoding individuals as having the functional limitation if they report any level of difficulty above the minimum.[11] Finally, the disease measures are binary indicators for whether the individual has ever been told by a medical professional that they have a particular condition. We record responses on 14 such diseases, such as cancer, diabetes, and heart disease.

Coarser morbidity measures. For many of our analyses, we employ a coarser set of morbidity measures, categorizing individuals into five ascending (i.e. worsening) "morbidity groups" based on their total number of ADLs and IADLs. This coarser measure of health is denoted throughout by

[7] Spending categories include inpatient, outpatient, physician, long-term care, prescriptions, dental, home health, and hospice care. Payers include Traditional Medicare, Medicare Advantage, Medicaid, private insurers, and self-pay.

[8] Throughout the paper, all dollar values are in 2024 US dollars.

[9] This residual category also includes dental and hospice care but these are minuscule compared to physician and outpatient.

[10] More specifically, the literature measures *limitations* to ADLs and IADLs, which are frequently described by the shorthand "ADLs" and "IADLs" with the "limitations to ..." left implicit.

[11] This practice follows the same recoding procedure as previously adopted by Cutler et al. (2014).

h_{it}. Morbidity group 1 (the healthiest group) consists of individuals with zero limitations across all 12 ADL and IADL measures.[12] About half of the sample is in this lowest morbidity group ($h_{it} = 1$). We divide the remaining population across four (roughly equally-sized) morbidity groups: morbidity group 2 consists of individuals with 1 limitation to ADLs or IADLs; morbidity groups 3 and 4 have, respectively, 2-3 and 4-7 limitations, and the highest morbidity group (morbidity group 5) consists of individuals with 8-12 limitations. The share of individuals in higher morbidity groups increases with age (Appendix Figure OA.2); for example, roughly 70% of 66-year-old respondents are in morbidity group 1, while less than 20% of respondents aged 90 or older are in morbidity group 1.

Transitions across health states. Our analyses focus on changes over time in the life of a hypothetical 66 year-old individual who faces the age-specific annual rates of mortality and morbidity experienced by the current cross-section of individuals at different ages. This is analogous to a standard period life expectancy table in which we imagine a hypothetical individual living their life within the year, experiencing that year's age-specific annual transition rates across morbidity states (including death) as they go.

To construct these age-specific transition matrices, we leverage the rotating panel structure of the MCBS; this enables us to observe how individuals transition across morbidity groups at each age. Rather than consider transitions between all 2^{31} possible combinations of the more granular health measure h'_{it}, which is not practical, we consider health transitions at the level of the coarser morbidity groups measure h_{it}. For a given survey year t and a given age a, we estimate the one-year transition rates from year t's morbidity group $h \in \{1, 2, 3, 4, 5\}$ to year $t + 1$'s morbidity group or mortality. That is, we construct a series of transition matrices, in which individuals start at age $a - 1$ in one of five morbidity states, and end in age a in either one of the five morbidity states or in death. We use $Pr^t(h_a|h_{a-1})$ and $d_a^t(h_{a-1})$ to denote the year-specific and age-specific health transition matrix and mortality risk, respectively. We estimate a separate 5-by-6 transition matrix for each age, resulting in 34 total transition matrices per year (covering ages of 66 to 99); we assume that all individuals die after age 100 if not already deceased. We estimate these 34 transition matrices separately for each survey year; changes in these transition matrices over time capture the key demographic changes in mortality and morbidity that are the focus of our analysis.[13] Appendix Tables OA.3 and OA.4 provide some examples of transition matrices from 1993 and 2017 for 66 year-old and 80 year-old individuals, respectively. We will use these transition matrices to simulate life trajectories for a counterfactual 66-year-old individual experiencing the age-specific morbidity and mortality rates of a given year.

Socioeconomic status. To examine how results vary by socioeconomic status, we classify individuals based on whether they are above or below median annual income (from all sources) for people of their age, gender, and time period (1993 or 2017). This follows prior work that uses measures of

[12] As mentioned, there are six measures of each, so an individual may have limitations with up to 12 ADLs and IADLs.

[13] In practice, for sample size reasons, each year's transition matrices are estimated using data from the previous and following year as well (in other words, the transition matrix for, say, 73-year-olds in 1993 is estimated by examining the morbidity transitions for 73-year-olds interviewed by the MCBS in 1992, 1993, and 1994, and who are also interviewed the following year). Each transition matrix therefore represents a 3-year moving average. See more details in Appendix A.4.

income rank for analyses of heath by socioeconomic status (e.g., Chetty et al. 2016; Hudomiet et al. 2021).[14]

2.2 Changes in health and longevity over time

Figure 1 displays the average number of ADLs and IADLs (out of a maximum total of 12), separately for 1993, 2012, and 2017. Panel (a) shows the average number of ADLs and IADLs by age while Panel (b) shows them by years prior to death. Consistent with a compression of morbidity, the results indicate that, over our observation period, individuals continued to face heightened morbidity in the years leading up to their death, but over time remained healthier for longer prior to the lead-up to their death.[15]

Panel (a) of Figure 1 shows that over the 25-year period from 1993 to 2017, morbidity conditional on age declined substantially. For example, in 1993, the average 80-year-old individual had difficulty with 2.6 ADLs and IADLs, whereas by 2017 this had declined by 38% to 1.6. Put differently, a 75 year-old individual in 1993 had, on average, approximately the same number of ADLs (1.6) as an 80 year-old individual in 2017, suggesting that "80 is the new 75." Morbidity declines are even larger at older ages. These improvements in morbidity, conditional on age, are also visible in the health transition matrices. For example, in 1993, 70.0% of 66 year-old individuals started in morbidity group 1 (no ADLs or IADLs) and, conditional on starting there, 86.2% of them remained in group 1 the following year. In 2017, the corresponding numbers were 72.0% and 90.4%, respectively (Appendix Tables OA.3). Once again, the changes are even starker at older ages; in 1993, only 43.6% of 80 year-old individuals were in morbidity group 1 compared to 56.4% in 2017, and only 74.7% remained in group 1 the following year compared to 79.1% in 2017 (Appendix Table OA.4). These patterns hold for a range of other measures of morbidity. Appendix Figure OA.5 shows changes between 1993 and 2017 in the (age-adjusted) rates of each of the 31 health conditions in our data; the rates of most health problems declines, including physical health problems such as difficulty reaching or walking and cognitive health problems such as difficulty in managing money.

By contrast, Panel (b) of Figure 1 shows that declines in morbidity are much more modest when we condition on years relative to subsequent death. In particular, declines in morbidity are absent entirely in the last two years of life. Five years before death, the average number of ADLs and IADLs declined from 2.57 in 1993 to 2.30 in 2012, a decline of only 10%.

Figure 2 shows the remaining life expectancy for a hypothetical 66 year-old individual who experiences the age-specific mortality and morbidity transitions of a given period (that is, 1993 or 2017). To estimate remaining life expectancy, we draw the morbidity group of an individual of age 66 in a given period based on the corresponding empirical distribution of morbidity groups for that age and

[14] A challenge in our setting is that differential survival by socioeconomic status could skew percentiles at older ages. We considered classifying individuals based on education (as in, e.g., Case and Deaton 2015, 2020; Hudomiet et al. 2022) but found this challenging because of the substantial increase in educational attainment across cohorts and time over our study period.

[15] Appendix Figure OA.3 and OA.4 show that the patterns are similar if we look at the presence of any ADL/IADLs at each age or the number of ADL/IADLs conditional on having at least one. These patterns of how morbidity has changed both relative to years since birth and relative to years prior to death are consistent with prior findings by Cutler et al. (2014) who analyzed the same MCBS data through 2009.

period, and then simulate the individuals forward using that period's age-specific transition probabilities across morbidity states until death. We then average across 100,000 simulated life trajectories to compute expected remaining life-years; we also compute expected remaining life-years in different health states.

For this hypothetical 66 year-old individual, remaining life expectancy increased by 2.36 years between 1993 and 2017. This represents a 14% increase in remaining life expectancy as of age 66, from 17.37 to 19.73 years. Strikingly, Figure 2 also show that the increase in remaining life expectancy comes entirely from an increase in life-years in morbidity group 1 – i.e. years spent without any ADLs or IADLs – which increased by 2.49 years, from 9.04 to 11.53 years. Moreover, the number of years spent in the highest morbidity group (8-12 ADL or IADLs) declined by 0.58 years (30%), with a corresponding increase of 0.49 years spent in the second lowest morbidity group (a single ADL or IADL). This pattern is consistent with the earlier findings (Figure 1) that health conditional on age has improved substantially, but that individuals continued to face similar morbidity in the years prior to death.[16]

Heterogeneity by gender and socioeconomic status. We briefly examine how the changes in health and longevity in Figure 3 vary by gender and by whether the individual is above or below the median income. Figure 3 shows a slightly larger increase in life expectancy for men than women (with men gaining 2.5 years and women 2.4 years). The patterns by socioeconomic status are more striking. As shown in Figure 3, remaining life expectancy increased by 2.8 years for above-median-income individuals, compared to 2.0 years for below-median-income individuals.

3 Implication for healthcare spending

In this section we translate the improvements in morbidity and mortality we documented in Section 2 into implications for expected lifetime public health-care spending for a 66 year-old individual. This is more subtle than the increase in expected lifetime Social Security payments, which is simply a scaler multiple of the increase in life expectancy. While increases in life expectancy also increase expected lifetime health-care spending (by adding life-years of positive expected spending), reductions in morbidity conditional on age serve to reduce lifetime per-capita spending, as health-care spending rises with morbidity. In other words, the existence of more people who live to the age of 85 – when spending is higher than at age 75 – drives spending up, but the typical 75 year-old individual is now healthier, which affects spending in the other direction.

3.1 Rising age and improving health: opposing demographic forces

To get a sense of the quantitative nature of each of these forces, we begin by examining how spending varies with age and with health. Figure 4 shows the relationship between health-care spending and age (Panel (a)) or years till death (Panel (b)) in 2017; it shows results for overall spending as well as for each type of spending category. Panel (a) shows that health-care spending is increasing (and

[16] Appendix Figure OA.6 show the time trends in health and life expectancy over the 1993 to 2017 period.

appears convex) in age, with average annual health-care spending rising from about \$14,000 at age 66 to over \$21,000 at age 80, and to about \$33,000 for those at age 90. Nursing-home spending in particular rises sharply with age, from close to zero at age 66 to about \$15,000 at age 90. By contrast, inpatient spending displays much less of an age pattern. Panel (b) shows spending relative to years to death; it shows rapidly increasing spending as death approaches, with nursing home spending in particular rising sharply in the years prior to death.

Figure 5 shows spending – both overall (Panel (a)) and by type of spending (Panels (b) through (f)) – by age, separately for each morbidity group. Several striking findings emerge. First, age is hardly associated with overall spending once morbidity is accounted for; this is consistent with Lee (2016), who finds that health-care spending depends much more on time relative to death than on chronological age. Second, while spending increases as morbidity increases, there is a particularly large increase in spending between morbidity group 4 (4-7 limitations to ADLs or IADLs) and morbidity group 5 (8-12 limitations). This is true both for overall spending (Panel (a)) and particularly pronounced for nursing home and home health spending (Panels (e) and (f)) compared to outpatient or prescription drug spending (Panels (c) and (d)). As we will see, this means that improvements in morbidity combined with decreases in mortality will put upward pressure on lifetime outpatient and prescription drug spending, while the decline in the expected number of years spent in morbidity group 5 (see Figure 2) will drive *down* lifetime nursing home and home health use.

3.2 Empirical approach

In order to map morbidity and mortality to lifetime health-care spending we combine two sets of estimates: estimates of dynamic age-specific transitions across health states and estimates of a static mapping from health states to annual spending. With these in hand, we can construct expected lifetime elderly spending using simulations in the same manner as we constructed life expectancy in Section 2. We can then evaluate how spending changes in response to observed or counterfactual demographic trends, summarized by the age-specific health transition matrices.

A practical complication is that we wish to use the granular health information on all 31 possible binary health variables (h'_{it}) to develop a flexible mapping from health to spending, but, as noted in Section 2, our health transition matrices rely, out of necessity, on a coarser measure of health based on five morbidity groups (h_{it}). The two-step conceptual procedure described above thus becomes, in practice, a three-step procedure in which we estimate (i) age-specific transitions across the coarser morbidity groups (and mortality), $Pr_a^t(h_a|h_{a-1})$ and $d_a^t(h_{a-1})$; (ii) the age-specific distribution of the granular, 31-dimensional health measures for each morbidity group, which we denote by $k_a^t(h'|h)$; and (iii) a mapping from granular health and age to predicted (annual) spending, which we denote by $\tilde{G}_a^t(m|h')$.

In principle, we can allow any of these three components to be time varying. In practice, for our baseline approach, we allow two of the components to be time varying: the age-specific transitions across health states ($Pr_a^t(h_a|h_{a-1})$ and $d_a^t(h_{a-1})$) and the distribution of granular health for each combination of age and health state ($k_a^t(h'|h)$). We hold the third component – the mapping from health to spending – constant; i.e. rather than having $G_a^t(m|h') = \tilde{G}_a^t(m|k_a^t(h'|h))$ we do not allow the distribution $\tilde{G}(\cdot)$ to be time varying and instead have $G_a^t(m|h') = \tilde{G}_a(m|k_a^t(h'|h))$. In other words, we focus on how changes in mortality and morbidity affect health-care spending for a fixed (stochastic) mapping $G(\cdot)$ from morbidity and age to spending. This has the desirable property of focusing on the role of demographics by abstracting from changes in the practice of medicine that may have been driven by technological, economic, or health policy changes.

To construct this mapping from morbidity and age to spending, we use a standard random forest algorithm, with one specific tweak. Appendix B describes the procedure in detail; we summarize it here. We use all 31 of the measures of individual health indicators (h_i') – all of the ADLs, IADLs, functional limitations, and disease measures shown in Appendix Table OA.1 – as inputs to a random forest algorithm, in order to predict annual spending. We denote the observed annual spending by y_i, and the predicted spending by $\widehat{y_i}$. The resultant predicted spending $\widehat{y_i}$ only reflects variation in spending due to morbidity, not age. We do this because we want to conceptually separate the impacts of health and aging on spending in the context of this exercise.

To accommodate the fact that, conditional on health, spending may vary with age, we supplement the predictive model to allow for separable, proportional age-specific spending effects by averaging the prediction residuals for each age. Specifically, we define the average, age-specific residual m_a as

$$m_a = \left(\sum_{a_i = a} w_i y_i \right) / \left(\sum_{a_i = a} w_i \widehat{y_i} \right), \tag{1}$$

where w_i are survey weights. Our final prediction is then given by $m_{a_i}\widehat{y_i}$; this is the prediction from the random forest, adjusted for multiplicative age effects. We apply this prediction algorithm to different years of the data (in particular, 1993 and 2017) and different components of annual spending (e.g., nursing-home spending or inpatient spending).

Our baseline analysis focuses on total spending predicted using data from 2016 through 2018 pooled together. Patterns of predicted spending make sense given the observed data. The age residuals are relatively small in practice (Appendix Figure OA.8), consistent with the pattern in Figure 5 showing that age is hardly associated with overall spending once accounting for morbidity. Like actual spending, predicted spending, and particularly nursing home spending, rises in the run-up prior to death (Appendix Figure OA.9). And as expected, predicted spending is increasing in morbidity group (Appendix Figure OA.10(a)); the relationship between morbidity group and average predicted spending will drive our estimates of how expected lifetime spending has changed over time. However, conditional on morbidity group, there is still considerable variation in predicted spending; this will feature in our model of the value of Medicare insurance coverage in Section 4.[17]

[17] Appendix Figure OA.10(b) shows the relationship between predicted spending and the number of different conditions the individual has, which ranges from 0 to 31 in our data. This is a version of the "frailty index" that is typically used in the literature (e.g., Abeliansky and Strulik 2019; Abeliansky et al. 2020; Strulik 2023). It shows that, conditional on the frailty index, there is still considerable variation in predicted spending.

3.3 Net impacts of demographic changes for health-care spending

Figure 6 shows expected lifetime health-care spending for a 66 year-old individual in 1993 and in 2017. The differences reflect the net effect of differences in mortality and morbidity, as we allow these to vary over time, holding constant the mapping from morbidity to spending. We find that, overall, the net impact of these demographic trends is to increase expected lifetime health-care spending for the elderly by $21,700 (5.6%), from $388,900 to $410,600. Moreover, Appendix Figure OA.13 shows that lifetime predicted health care spending was relatively flat over between 1993 and about 2006, with most of the increase in lifetime predicted health care spending occurring in the last 10 years of our study period. Recall that we are holding the mapping from health to spending fixed, so these changes in lifetime healthcare spending are driven by changes in the underlying distribution of mortality and morbidity.

Thus, despite the 14% increase in life expectancy between 1993 and 2017, the substantial improvements in health (in particular, the large decrease in life-years spent in morbidity group 5 seen in Figure 2) result in a much less than proportional (6%) increase in lifetime health care spending. Put differently, if the added life years had been typical years in terms of elderly health, lifetime health care spending would have grown by 14%, while if the added years were disproportionately unhealthy years, it could have grown by even more; however, because of the declines in morbidity conditional on age, lifetime health care spending grew by substantially less.[18]

Figure 6 also shows that the change in expected lifetime health care spending is not uniform across health care categories. In particular, pharmaceutical drugs and outpatient spending have been the key drivers behind the rise in expected lifetime spending, rising by $12,500 (18%) and $16,800 (11%), respectively.[19] By contrast, inpatient spending remains roughly stable, rising by only $1,000 (2%) while nursing home and home health care spending decline by about $4,600 and $2,700 (5% and 18%, respectively). These differential changes across types of spending are the joint result of the underlying morbidity changes (Figure 2) and the relationship between morbidity and different spending categories (Figure 5). In particular, we have added *healthy* (morbidity groups 1 and 2) years of life, while reducing years spent in morbidity group 5, and nursing home, and home care rise much more rapidly with morbidity than prescription drug or outpatient/physician spending.[20]

Finally, note that all of the preceding analyses have used the 2017 mapping from granular health to spending. While this choice strikes us as a natural, it is not inconsequential. For example, if

[18]As another way of seeing this, Appendix Figure OA.11 displays predicted spending under a counterfactual in which mortality declines are reduced by a constant, age-specific proportional factor across morbidity groups, so that morbidity is held constant, as well as under an analogous counterfactual in which morbidity rates change but mortality does not. Increases in life expectancy, holding morbidity constant, would have increased lifetime healthcare spending by $74,200 (about 19%), while reduced morbidity holding life expectancy constant would have *decreased* lifetime spending by $40,300 (about 12%). The net increase in spending of $21,700 rather than $33,900 reflects the fact that these two forces are not additively separable.

[19]Note that changes in predicted spending *only* reflect changes in the morbidity transition matrices from 1993 to 2017. They do not reflect the impacts of technological or policy changes, such as the introduction of Medicare prescription drug coverage (part D) in 2006. This is because, as noted, we use the same, fixed mapping from health to spending (based on 2016-2018 data) to predict lifetime spending given 1993 morbidity transition matrices and given 2017 morbidity transition matrices.

[20]Recall that there is also (potentially time-varying) heterogeneity in predicted spending within each morbidity group (see Appendix Figure OA.10), which also contributes to the results.

we instead were to use the 1993 mapping from health to spending instead, the observed changed in mortality and morbidity between 1993 and 2017 would generate a net, slight *decline* in predicted health-care spending of about $6,400 rather than our baseline result of an increase of$21,700 (see Appendix Figure OA.14). This primarily reflects larger predicted declines in nursing home spending and smaller predicted increases in prescription drug and outpatient spending when using the 1993 mapping from health to spending (see Appendix Figure OA.15).

Heterogeneity by gender and socioeconomic status. We briefly examine variation across demographic groups in the changes in expected lifetime health-care spending shown in Figure 6.

Figure 7 shows that the increase in lifetime spending is more than five times as high for men (an increase of $48,100, or 15%) than for women (an increase of $9,500, or 2%). This reflects two forces (discussed in the last section) that push toward increasing lifetime spending for men relative to women: a slightly larger increase in life expectancy for men than women and a larger decline in the number of years spent in the worst health group for women compared to men; as seen in Figure 5(a), spending is so much higher for people in this highest morbidity group. Figure 7 also shows that the increase in lifetime health-care spending is more than three times as high for high-income individuals relative to low-income individuals, with an increase of $35,700 (9%) for above-median-income individuals compared to $10,900 (3%) for below-median-income individuals. This is driven by the previously-documented pattern that life expectancy increased substantially more for above-median-income individuals.

Implications for public spending. Figure 8 uses the increase in life expectancy estimated in Figure 2 and the increase in lifetime health-care spending estimated in Figure 6 to assess the implications of these demographic changes on overall public spending, specifically health-care spending through Medicare, and public pension spending through the Social Security annuity. Based on 1993 data, we assume a $18,000 fixed social security payment per year, and that Medicare covers 50 cents of each dollar of elderly medical expenditure (see Appendix C for details). As a result, the estimated 2.4 year increase in life expectancy translates into an increase in expected lifetime Social Security payments to a 66 year-old individual of about $42,500 (or about 14%), and the estimated $21,700 increase in expected lifetime health-care spending implies an increase of about $10,800 in lifetime *Medicare* spending (or about 6%). Thus, overall, the demographic changes imply that the average 66 year-old individual is expected to see an increase of about $53,300 in the public outlays they receive, with about 80% of this increase taking the form of Social Security payments, and the remainder coming through Medicare.

Thus under the status quo policies, demographic changes led to expected lifetime Social Security spending growing more than twice as fast as expected lifetime Medicare spending. We now turn to considering the forces that affect the optimal allocation of spending across these two programs.

4 The welfare impact of alternative public spending policies

To assess the factors that affect the relative welfare impacts of (budget neutral) shifts in public funding between Social Security and health-insurance ("Medicare") benefits, we specify and calibrate a standard life-cycle consumption model in which risk averse elderly individuals choose an optimal intertemporal consumption path in the presence of stochastic and exogenous health and mortality risks (e.g., Scholz et al. 2006; Brown and Finkelstein 2008; De Nardi et al. 2010). We first consider a case with no moral hazard, so that health insurance coverage does not affect total health-care spending. We provide intuition for why the optimal allocation is a corner solution; conditional on a given government budget, consumer surplus is maximized by spending the entire budget on Medicare until either the budget runs out or Medicare covers all medical expenses, and then allocating any remaining budget to Social Security. We then extend the model to allow for moral hazard, which leads to a more reasonable interior solution on which one can perform comparative statics.

4.1 A stylized model

Setting and notation. We consider a population of 66-year-old individuals, each associated with an initial wealth level w_{66} and an initial morbidity level h_{66}. They receive a constant, annual, government-provided Social Security benefits which provide an annual income payment ss conditional on survival. We omit individual subscripts to simplify the exposition, but all of these objects can be as individual-specific. Their morbidity and mortality follows an exogenous Markov process, outside of the control of the individual. As earlier, we denote by $d_a(h_{a-1})$ the probability that the individual dies at age a if their morbidity level the year before is h_{a-1}, and we denote by $Pr_a(h_a|h_{a-1})$ the age-specific probability distribution of morbidity at age a if their morbidity level the year before is h_{a-1} and they survive to age a. Health in turn translates into requisite medical spending m drawn from an age-specific distribution that is a function of morbidity, denoted by $G_a(m|h)$. This requisite health-care spending m is non-discretionary and "must" be paid.

Individuals have per-period CRRA utility from consumption, $u(c) = \frac{c^{1-\gamma}-1}{1-\gamma}$, which is assumed to be invariant to health, and a utility of zero after death (i.e., no bequest motive). The individual faces an annual consumption vs savings decision, which they make subject to a budget constraint and a no-borrowing constraint (to eliminate the possibility that the individual can die in debt). We model government-provided health insurance ("Medicare") as a linear coverage policy that covers a fixed proportion, $\lambda \in [0, 1]$, of m. The individual must pay the remaining portion, $(1 - \lambda)m$, but the government guarantees a consumption floor $\underline{c}$. That is, if the individual's income (ss) and assets (w_a) net of their required out-of-pocket medical spending $(1-\lambda)m_a$ produce a maximum feasible consumption level (given the no-borrowing constraint) that is below $\underline{c}$, the government will supplement the individual's income so that their feasible consumption at age a (after having paid their required out-of-pocket medical expenditures) reaches the threshold $\underline{c}$. This additional government spending can heuristically be thought of as Medicaid spending (Brown and Finkelstein 2008; De Nardi et al. 2016).

In other words, once the individual realizes their requisite total healthcare spending m_a, they choose their optimal consumption $c \in [\underline{c}, \max\{\underline{c}, w_a + ss - (1-\lambda)m_a\}]$ in order to maximize

$$V_a(m, w_a, h_a) =$$
$$max_c \left[u(c) + \beta(1 - d_{a+1}(h_a)) \sum_{h_{a+1}} \left(Pr_{a+1}(h_{a+1}|h_a) \int_{m'} V_{a+1}(m', w_{a+1}(c), h_{a+1}) dG_{a+1}(m'|h) \right) \right], \tag{2}$$

where β is the per-period discount rate, $V_a(m, w_a, h_a)$ is the value function, and wealth evolves according to

$$w_{a+1}(c) = \max\{0, (1+r)(w_a + ss - c_a - (1-\lambda)m_a)\}, \tag{3}$$

where r is the per-period interest rate.

Measuring welfare. To assess utility under alternative public policy arrangements, we monetize utility by mapping each policy to the amount individuals would be required to pay or receive at age 66 in order to be indifferent between a given policy and the baseline policy levels. We consider two public policy instruments: the level of the annual Social Security payment (ss), and a linear (that is, fixed proportion of spending) Medicare coverage (λ). We denote the average age-66 value function of individuals with initial wealth w_{66} under the baseline policy levels λ and ss by

$$V_{66}(w_{66}|\lambda, ss) = \sum_{h_{66}} \left(Pr(h_{66}) \int_{m'} V_{66}(m', w_{66}, h_{66}|\lambda, ss) dG_{66}(m'|h) \right). \tag{4}$$

The monetized utility from an alternative policy combination (λ', ss') relative to the baseline policy (λ, ss) is then given by the willingness to pay z, such that the individual is indifferent between paying that amount and remaining with the baseline policy or receiving the alternative policy:

$$V_{66}(w_{66} - z|\lambda, ss) = V_{66}(w_{66}|\lambda', ss'). \tag{5}$$

This approach follows an existing literature that calculates similar measures of willingness to pay for annuities (e.g., Mitchell et al. 1999; Davidoff et al. 2005) and for long-term care insurance (e.g., Brown and Finkelstein 2008). We refer to z as the "wealth equivalent" of policy (λ', ss') relative to the baseline policy. A positive (negative) value of z suggests that policy (λ', ss') is welfare enhancing (reducing) relative to the baseline policy.

4.2 Analysis without moral hazard

Absent moral hazard, an additional expected dollar of Medicare coverage has two advantages over an additional expected Social Security dollar. These both push the optimal allocation to be spending the government budget on Medicare until either the budget runs out or Medicare covers everything (that is, $\lambda = 1$), and then allocating any remaining funds to Social Security.

The first goes back to Arrow (1963). The individual faces two sources of risk: mortality and health. Social Security payments address only the first source of risk, helping individuals reduce consumption volatility across states of the world in which they are alive (Yaari 1965). However, Medicare – which of course also only pays out in states of the world in which the individual is alive – also helps reduce consumption volatility *within* the alive states across different realizations of health (and hence medical expenditure) shocks. As a result, it is always optimal for a risk averse individual to receive the marginal incremental (expected) dollar in the form of Medicare coverage rather than Social Security.

The second advantage of Medicare coverage is more nuanced. It stems from the fact that for a retiree facing stochastic mortality in the (assumed) absence of private annuities markets, the optimal Social Security benefit schedule is back-loaded (i.e. increasing with age) for individuals with sufficient wealth; we derive this result and discuss it in more detail in Einav and Finkelstein (2026). Therefore, given that Social Security payments are assumed constant (i.e. do not vary with age), which is consistent with current policy, the fact that health deteriorates (on average) with age implies that tying government payments to health – through Medicare coverage – provides an indirect way for the government to back-load its payments, and get it closer to the optimal schedule. This ability of Medicare relative to Social Security to better target payments to higher marginal utility states of the world is, in some sense, a dynamic version of Lieber and Lockwood (2019)'s argument that, relative to cash, in-kind transfers through health insurance can improve the targeting of transfers to higher marginal utility states of the world.

To illustrate this result quantitatively, we calibrate the model and solve for wealth-equivalents for different budget-neutral policy combinations.

Model calibration and computation. We briefly describe the model's calibration and implementation; Appendix C provides more details. For the baseline policy parameters, we use $ss = \$18,000$ and $\lambda = 0.5$. We set the consumption floor ($\underline{c}$) to \$11,000 (the mean annual payout from Supplemental Security Income (SSI) in 1993). We assume no discounting ($\beta = 1$) and no interest ($r = 0$) and set the initial wealth to $w_{66} = \$150,000$, roughly the median of total non-annuitized wealth, excluding housing and other real estate, held by 65-69 year olds in 2008. We assume $\gamma = 3$.

For the distribution of risks facing the individual, we use the estimates of age-specific transitions across the five morbidity states (h) (as well as mortality) calculated using the MCBS panel data, as described in Section 2.1. Specifically we use the estimated transitions from 1993 and 2017 and denote them by $Pr_a^t(h_a|h_{a-1})$ and $d_a^t(h_{a-1})$, where $t \in \{1993, 2017\}$.[21] The distribution of health-care spending $G_a(m|h)$ is also based on the MCBS data, specifically the distribution of predicted spending for each age-morbidity combination (see Section 3.2). Recall that we use a constant mapping from granular health (h') to spending m across time periods.[22] We assume that death is certain after age

[21] Recall from Section 2.1 that 1993 morbidity and mortality transitions are estimated using data from 1992-1994, and 2017 estimates use data from 2016-2018.

[22] $G_a(m|h)$ is thus approximately time-invariant. Following the discussion in Section 3.2, the mapping of (granular) health to spending ($\tilde{G}_a(m|h')$) relies on 2017 spending data and therefore does not change over time in all our analyses below. However, $k_a^t(h'|h)$ – the mapping from the five health states (h) to the more granular, 31-dimensional health vector (h') (see Section 3.2) – varies between 1993 and 2017, so the distribution $G_a^{2017}(m|h) = \tilde{G}_a(m|h')k_a^{2017}(h'|h)$ can vary between 1993 and 2017 due to changes over time in the distribution of granular health at a given age conditional

100, allowing us to solve the model backwards as it means that $V_{100}(m, w_{100}, h_{100}) = 0$ for any value of (m, w_{100}, h_{100}).

We solve the model to compute two different levels for the government budget: expected lifetime government spending (across Medicare, Social Security, and maintaining the consumption floor $\underline{c}$ ("Medicaid")) for a 66 year old in 1993 (B_{1993}) and in 2017 (B_{2017}). In both cases, we use the baseline policy parameters and other calibrated parameters described above. To calculate B_{1993}, we solve the model using the health transitions in 1993 (i.e. $Pr^{1993}(h_a|h_{a-1})$ and $d_a^{1993}(h_{a-1})$); likewise to calculate B_{2017} we use the 2017 versions of these objects. We find B_{1993} = \$562 thousand and B_{2017} = \$609 thousand. The two main components of the budget – expected lifetime spending on Medicare and Social Security – correspond, by design, to the estimates already shown in Figure 8. The third component of the budget (which, using the calibrated model, is \$56.4 thousand in 1993 and \$49.4 thousand in 2017) is driven by the individual's endogenous consumption choices and hence the government "Medicaid" spending on maintaining the consumption floor.

We then solve for the wealth-equivalent z (equation (5)) that would make the individual indifferent between the baseline Medicare and Social Security parameters and alternative policy parameters (λ', ss'). Our main counterfactuals use the 2017 health transitions ($Pr^{2017}(h_a|h_{a-1})$, $d_a^{2017}(h_{a-1})$), and estimate z under alternative policy parameters that hold the total government spending fixed at either B_{1993} or B_{2017}.

Results. We compute the willingness to pay (see equation (5)) for different policy combinations (λ', ss') relative to the baseline policy of $(\lambda = 0.5, ss = 18,000)$, focusing on the policy combinations that give rise to the two budget levels, B_{1993} and B_{2017}. Note that determining a set of budget-neutral combinations of λ' and ss' requires solving the consumer's dynamic optimization problem in equation (2) because the government spending on "Medicaid" (to maintain the guaranteed consumption floor $\underline{c}$) depends on the individual's endogenous consumption choices; all else equal, higher consumption (lower savings) increases expected government Medicaid spending.

Our key object of interest is how the willingness to pay for alternative policy combinations change as we change morbidity and mortality patterns from the 1993 estimates to the 2017 estimates. Figure 9 illustrates the key results. Panel (a) presents iso-budget lines (solid) and iso-utility curves (dashed) in the space of the two policy instruments, λ (x-axis) and ss (y-axis). The red line collects all the (λ', ss') pairs that hold total government spending at the 1993 level B_{1993}. The blue line collects all the (λ', ss') pairs that hold total government spending at the 2017 level B_{2017}. Naturally, B_{2017} is always higher than B_{1993}; this reflects the implications for Social Security spending and Medicare spending of the changes in morbidity and mortality documented in Section 2 (recall Figure 8.) The dashed lines present iso-utility curves under the 2017 mortality and morbidity transitions in the same space, with the utility increasing as move up and to the right (as either λ and/or ss increase).

The key point to notice is that – without moral hazard – both the iso-budget and iso-utility curves are approximately linear, with the iso-utility curves being steeper. This implies that the optimal allocation of public funds maximizes the health insurance (that is, Medicare) benefits, which

on the morbidity group at that age (i.e. $k_a^t(h'|h)$).

is consistent with the economic intuition highlighted earlier. This is shown explicitly in Panel (b) of Figure 9, which plots the wealth equivalent (y-axis) associated with different (budget neutral) choices of λ (x-axis). It shows that the highest wealth equivalent is obtained at maximum Medicare coverage of $\lambda = 1$. Recall that this figure shows budget-neutral policies, so an increase in λ must come with a decrease in Social Security (ss) benefits to make the choice budget neutral. Panel (c) shows the mirror image of this results by plotting the wealth equivalent of (budget neutral) changes to the Social Security (ss) benefits.

4.3 Analysis with moral hazard

The preceding analysis illustrated the economic forces affecting the value of Medicare compared to Social Security. However, they also produced a "corner solution" in which the consumer always prefers to maximize the Medicare benefits. Incorporating moral hazard, a realistic feature of health insurance (Einav and Finkelstein 2018), introduces a downside of Medicare relative to Social Security since Medicare now induces excessive (and inefficient) health-care spending; indeed, moral hazard is a classic disadvantage of in-kind transfers such as health insurance when evaluating the welfare properties of in-kind relative to cash transfers (Lieber and Lockwood 2019). As a result, we can now find an interior solution to the optimal allocation of public funds between Medicare and Social Security and assess how this interior solution is affected by the changing mortality and morbidity patterns described in Section 3.

Incorporating moral hazard into the model. Moral hazard is typically modeled by considering an individual's utility-maximizing choice between medical and non-medical consumption, with insurance coverage affecting the price of that medical consumption; individual utility is a function of non-medical consumption and health, and medical consumption is an input into health (e.g., Einav et al. 2013; De Nardi et al. 2016). In these models, any counterfactual Medicare coverage that affects the individual's choice of medical spending would therefore also affect the health transition matrices in some way we would need to parameterize and calibrate. Since our primary goal is to examine how changes in these matrices affect the value of health insurance compared to Social Security, we want to abstract from any direct impacts of health insurance on these health transitions.

Therefore, instead of modeling optimal medical spending as part of the individual's optimal choice, we introduce moral hazard into the model in a "reduced form" way by assuming that increased Medicare coverage leads to greater medical spending without any effect on the individual health transitions. Specifically, we assume that if Medicare covers a share λ of total health-care expenditure, then a share $q(\lambda) = \theta\lambda^{\kappa} \in [0, 1)$ of total medical expenditure is wasted, where the parameter $\theta \geq 0$ controls the level of moral hazard and the parameter $\kappa > 1$ controls how fast moral hazard increases with coverage levels. We denote by m_e the efficient amount of health-care spending for a given individual in a given year, i.e. how much the individual would spend with no coverage ($\lambda = 0$), or equivalently with no moral hazard ($q = 0$); in the context of our model, m_e is a function of the individual's morbidity level and spending realization. Under these assumptions, the observed actual spending $m(\lambda)$ can be written as a function of the efficient medical spending and moral hazard:

$m(\lambda) = \frac{m_e}{1-q(\lambda)}$. That is, the observed medical spending is the efficient amount of medical spending m_e, inflated by a factor that is increasing in the level of moral hazard q; the difference between the observed level of spending and the efficient level of spending ($m(\lambda) - m_e = q(\lambda)m$) is "wasted."

Calibrating moral hazard. We use the same calibration choices as in the previous section but now must also calibrate the moral hazard parameters θ and κ. The value of θ is not very consequential so we calibrate it to 0.2,[23] but κ affects the curvature of moral hazard so it has a non-trivial effect on counterfactuals. We calibrate the value of κ by requiring it to make the baseline parameters ($ss = \$18,000, \lambda = 0.5$) optimal under the 1993 demographics. In other words, we calibrate κ by finding the value that would make it maximize the wealth equivalent z in equation (5), given the 1993 mortality and morbidity transition patterns and total government spending of B_{1993}. We make this (obviously heroic) assumption to make it more conceptually straightforward to analyze the impact of demographic changes. Specifically, we can ask: if policy choices were optimal in 1993, how is the optimal policy affected by the subsequent demographic changes?

Appendix Figure OA.12 illustrates this calibration by showing that at the calibrated value of $\kappa = 6$ the wealth equivalent is maximized at the observed policy ($\lambda = 0.5$). We note that this value of $\kappa = 6$, combined with $\theta = 0.2$, implies an elasticity of health spending (m) with respect to insurance coverage (λ) of approximately -0.02., which is an order of magnitude lower than the oft-quoted RAND elasticity of -0.2 (Manning et al. 1987; Keeler and Rolph 1988). Put differently, in order to obtain the (obviously heroic) benchmark that the 1993 Medicare and Social Security allocations were optimal requires a moral hazard elasticity much lower than the "conventional wisdom." One interpretation of this of course is that the 1993 allocation was not in fact optimal. Another is that there are other, unmodeled forces that favor Medicare.

Results. We repeat the exercise in Section 4.2, but now use the calibrated values of θ and κ to account for moral hazard. Naturally, once moral hazard is present, full Medicare coverage suboptimal; by design, the calibrated level of moral hazard means that under the 1993 budget and the 1993 transition matrices the wealth equivalent maximizing policy is the baseline policy of $\lambda = 0.5$ and $ss = \$18,000$.

Figure 10 shows the implications for the wealth equivalent maximizing policy *under the 2017 transition matrices*. Under the 2017 transition matrices, Panel (a) shows how wealth equivalents (y-axis) vary with the Medicare share (x-axis) for different budget levels; it is the moral hazard analog of Panel (b) of Figure 9. Panel (c) presents the analogous result for Social Security. Black dots indicate the wealth equivalent maximizing choice of the Medicare coverage and Social Security amount for each budget level *under the 2017 transition matrices.*

If – as entitlement programs do – we retain the status quo public benefits of $\lambda = 0.5$ and $ss = 18,000$, and allow total government spending to expand to B_{2017} as demographics change to the 2017 level, retaining a Medicare coverage of $\lambda = 0.5$ remain approximately optimal. This is shown by the

[23] Results are very similar when we use a broad range of alternative values for θ, as long as κ is re-calibrated (as described below) conditional on θ.

black dot on the blue line; the optimal Medicare share of 49% and the optimal Social Security budget of $\$18,175$ is very similar to the baseline policy.

By contrast, if we take the other extreme and require that the public benefits λ and ss shrink in response to the 2017 demographics in order to maintain government spending at the 1993 level (B_{1993}), then the optimal policy now involves a non-trivial decline in both Medicare and Social Security benefits. This is seen by the black dots on the red line. Specifically, if we require the budget to stay at the 1993 level, the optimal policy changes from 50% Medicare coverage and $18,000 Social Security to 35% Medicare coverage and $14,910 Social Security. This lower level of benefits makes individual "poorer," and requires the government to more than double their "Medicaid" benefits (from $57.2 thousands per individual to $125.3 thousands), reducing overall Medicare expenditure by 27.2% (from $196.4 thousands to $143.0) and Social Security expenditure by 6.5% (from $314.4 thousands to $294.0).

5 Conclusions

In this paper we have calculated the implications for public spending of increased elderly life expectancy and improved elderly health over the 1992-2018 period. We estimate that under current program rules for public health insurance and public pensions, these demographic trends increased expected lifetime Social Security spending for a 66 year old by 14%, more than twice as fast as the 6% increase in expected lifetime Medicare spending for a 66 year old. Interestingly, despite the increase in life expectancy, expected spending on home health and nursing homes actually declines slightly, due to the decrease in morbidity.

We then write down and calibrate a model that allows us to consider the optimal allocation of public funding between Social Security and Medicare. While these two programs are administered with separate revenue streams and through separate agencies, the model underscores that they are functionally highly related, providing two distinct sources of consumption smoothing for elderly retirees facing stochastic mortality and morbidity. The model clarifies several key economic forces that affect the relative consumer surplus from allocating an incremental expected dollar to Medicare (health insurance) coverage relative to Social Security (longevity insurance via annuities), and allows us to examine via stylized quantitative calibrations how demographic changes can affect this tradeoff. We hope this is a useful first step for further analysis of how public policy regarding not just the *level* of expenditures on social programs but also their relative allocation across programs should respond to major demographic trends.

References

Abeliansky, Ana Lucia and Holger Strulik, "Long-Run Improvements in Human Health: Steady but Unequal," *The Journal of the Economics of Ageing*, 2019, *14*, 100189.

_ **, Devin Erel, and Holger Strulik**, "Aging in the USA: Similarities and Disparities Across Time and Space," *Scientific Reports*, 2020, *10*, 14309.

Arrow, Kenneth J, "Uncertainty and the Welfare Economics of Medical Care," *The American Economic Review*, 1963, *53* (5), 941–973.

Brown, Jeffrey R and Amy Finkelstein, "The interaction of public and private insurance: Medicaid and the long-term care insurance market," *American Economic Review*, 2008, *98* (3), 1083–1102.

Buxbaum, Jason D, Michael E Chernew, A Mark Fendrick, and David M Cutler, "Contributions Of Public Health, Pharmaceuticals, And Other Medical Care To US Life Expectancy Changes, 1990-2015: Study examines the conditions most responsible for changing US life expectancy and how public health, pharmaceuticals, other medical care, and other factors may have contributed to the changes.," *Health Affairs*, 2020, *39* (9), 1546–1556.

Case, Anne and Angus Deaton, "Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century," *Proceedings of the National Academy of Sciences*, 2015, *112* (49), 15078–15083.

_ **and** _ , "Deaths of Despair and the Future of Capitalism," 2020.

CBO, *The Long-Term Budget Outlook: 2025-2055*, Congressional Budget Office Washington (DC), 2025.

Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler, "The association between income and life expectancy in the United States, 2001-2014," *Jama*, 2016, *315* (16), 1750–1766.

CMS, "MCBS Advanced Tutorial on Weighting and Variance Estimation," Technical Report 2021.

Cutler, David, Kaushik Ghosh, and Mary Beth Landrum, "Evidence for Significant Compression of Morbidity in the Elderly U.S. Population," *Discoveries in the Economics of Aging*, 2014, pp. 21–51.

Cutler, David M, Kaushik Ghosh, Kassandra L Messer, Trivellore Raghunathan, Allison B Rosen, and Susan T Stewart, "A satellite account for health in the United States," *American Economic Review*, 2022, *112* (2), 494–533.

Davidoff, Thomas, Jeffrey R Brown, and Peter A Diamond, "Annuities and individual welfare," *American Economic Review*, 2005, *95* (5), 1573–1590.

der Klaauw, Wilbert Van and Kenneth I Wolpin, "Social security and the retirement and savings behavior of low-income households," *Journal of econometrics*, 2008, *145* (1-2), 21–42.

Einav, Liran, Amy Finkelstein, Stephen P Ryan, Paul Schrimpf, and Mark R Cullen, "Selection on moral hazard in health insurance," *American Economic Review*, 2013, *103* (1), 178–219.

_ **and** _ , "Moral hazard in health insurance: what we know and how we know it," *Journal of the European Economic Association*, 2018, *16* (4), 957–982.

_ **and** _ , "On the Optimality of Deferred Public Annuities," *Journal of Public Economics*, 2026, *256*, 105601.

Freedman, Vicki A, Brenda C Spillman, Patti M Andreski, Jennifer C Cornman, Eileen M Crimmins, Ellen Kramarow, James Lubitz, Linda G Martin, Sharon S Merkin, Robert F Schoeni et al., "Trends in late-life activity limitations in the United States: an update from five national surveys," *Demography*, 2013, *50* (2), 661–671.

Freedman, Vicki, Linda Martin, and Robert Schoeni, "Recent Trends in Disability and Functioning among Older Adults in the United States: A Systematic Review," *JAMA*, 2002, *288* (24), 3137–3146.

French, Eric, "The effects of health, wealth, and wages on labour supply and retirement behaviour," *The Review of Economic Studies*, 2005, *72* (2), 395–427.

Fries, James F., "Aging, Natural Death, and the Compression of Morbidity," *New England Journal of Medicine*, 1980, *303*, 130–135.

Fuchs, Victor R, "" Though Much is Taken"–Reflections on Aging, Health, and Medical Care," 1984.

Gruber, Jonathan, Kathleen M. McGarry, and Charles Hanzel, "Long-term Care Around the World," *NBER Working Paper 31882*, 2023.

Gruenberg, Ernest M, "The failures of success," *The Milbank Memorial Fund Quarterly. Health and Society*, 1977, pp. 3–24.

Hudomiet, Péter, Michael D Hurd, and Susann Rohwedder, "Forecasting mortality inequalities in the US based on trends in midlife health," *Journal of health economics*, 2021, *80*, 102540.

_ , _ , and _ , "Trends in Health in Midlife and Late Life," *Journal of human capital*, 2022, *16* (1), 133–156.

Keeler, Emmett B and John E Rolph, "The demand for episodes of treatment in the health insurance experiment," *Journal of health economics*, 1988, *7* (4), 337–367.

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

Lee, Ronald, "Macroeconomics, aging, and growth," in "Handbook of the economics of population aging," Vol. 1, Elsevier, 2016, pp. 59–118.

Lieber, Ethan MJ and Lee M Lockwood, "Targeting with in-kind transfers: Evidence from Medicaid home care," *American Economic Review*, 2019, *109* (4), 1461–1485.

Lockwood, Lee M, "Incidental bequests and the choice to self-insure late-life risks," *American Economic Review*, 2018, *108* (9), 2513–2550.

Lubitz, James, Liming Cai, Ellen Kramarow, and Harold Lentzner, "Health, life expectancy, and health care spending among the elderly," *New England Journal of Medicine*, 2003, *349* (11), 1048–1055.

Manning, Willard G, Joseph P Newhouse, Naihua Duan, Emmett B Keeler, and Arleen Leibowitz, "Health insurance and the demand for medical care: evidence from a randomized experiment," *The American economic review*, 1987, pp. 251–277.

Manton, Kenneth G, "Changing concepts of morbidity and mortality in the elderly population," *The Milbank Memorial Fund Quarterly. Health and Society*, 1982, pp. 183–244.

McClellan, Mark and Jonathan Skinner, "The incidence of Medicare," *Journal of Public Economics*, 2006, *90* (1-2), 257–276.

Mitchell, Olivia S, James M Poterba, Mark J Warshawsky, and Jeffrey R Brown, "New evidence on the money's worth of individual annuities," *American economic review*, 1999, *89* (5), 1299–1318.

Nardi, Mariacristina De, Eric French, and John B Jones, "Why do the elderly save? The role of medical expenses," *Journal of political economy*, 2010, *118* (1), 39–75.

_ , _ , and John Bailey Jones, "Medicaid insurance in old age," *American Economic Review*, 2016, *106* (11), 3480–3520.

National Academies of Sciences, "The growing gap in life expectancy by income: Implications for federal programs and policy responses," 2015.

National Research Council, Division of Behavioral and Social Sciences and Committee on Population and Division on Engineering and Physical Sciences, Board on Mathematical Sciences and Their Applications, and Committee on the Long-Run Macroeconomic Effects of the Aging US Population, *Aging and the Macroeconomy: Long-Term Implications of an Older Population*, National Academies Press, 2013.

Poterba, James M, "Retirement security in an aging population," *American Economic Review*, 2014, *104* (5), 1–30.

Poterba, James M. and Lawrence H. Summers, *Public Policy Implications of Declining Old-Age Mortality*, Brookings Institution Press, 1987.

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun, "Are Americans saving "optimally" for retirement?," *Journal of political economy*, 2006, *114* (4), 607–643.

Sheiner, Louise, "The long-term impact of aging on the federal budget," in "Fiscal Accountability and Population Aging," Edward Elgar Publishing, 2021, pp. 93–117.

SSA, "Benefits Paid by Type of Beneficiary," Technical Report 2025.

Strulik, Holger, "Measuring ageing," in "The Routledge Handbook of the Economics of Ageing," Routledge, 2023, pp. 455–473.

Utkus, Stephen P and Olivia S Mitchell, "Extending Healthspans in an Aging World," *NBER Working Paper 33992*, 2025.

Yaari, Menahem E, "Uncertain lifetime, life insurance, and the theory of the consumer," *The Review of Economic Studies*, 1965, *32* (2), 137–150.

Figure 1: Trends in ADL/IADLs by Age and Years to Death

(a) By Age

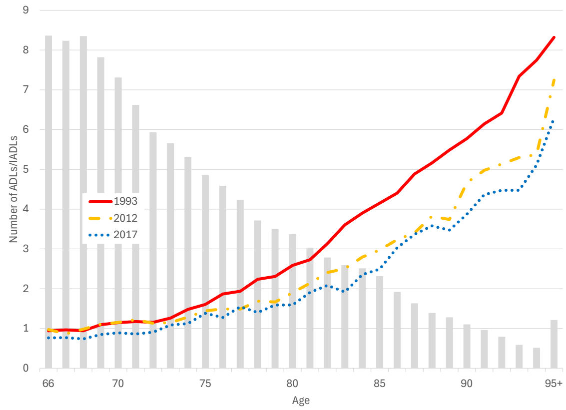

(b) By Years to Death

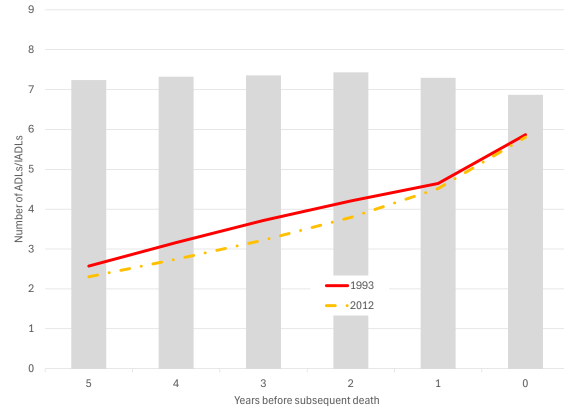

Notes: This figure displays the mean number of ADL/IADLs by age (Panel (a)) and by number of years to death (Panel (b)), where year 0 consists of anywhere from 365 to 0 days before death. We show results separately for the 1993 survey cohort (in red), the 2012 survey cohort (in yellow), and the 2017 survey cohort (in blue). The 2017 survey cohort is not displayed in Panel (b) because our mortality data end in 2019. All ages above 95 are collapsed into a 95+ category. A gray histogram displays the proportion of observations in each age category in 2017 (Panel (a)) or time-to-death category in 2012 (Panel (b)). Standard survey weights are used. The sample includes individuals aged 66+ surveyed by the MCBS in 1993 (N = 26,413), 2012 (N = 23,669), or 2017 (N = 19,917).

Figure 2: Life Expectancy at 66 by Morbidity Group

(a) Life Expectancy

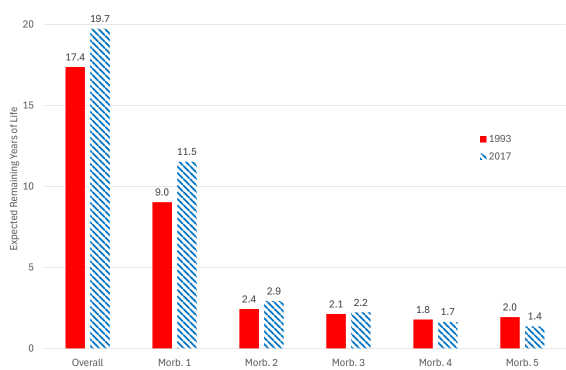

(b) Change in Life Expectancy, 1993-2017

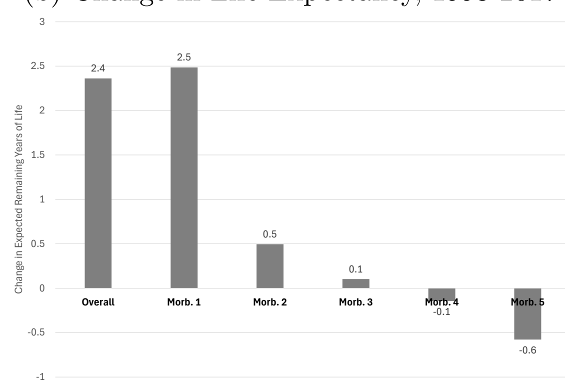

Notes: This figure displays the life expectancy of the typical 66-year-old, both overall and calculated as the number of years expected to be spent within each morbidity group. Panel (a) displays these statistics in years separately for 1993 (in red) and 2017 (in blue). Panel (b) displays the difference between these two years. Life expectancies for each survey year are calculated using that year's transition matrices. N = 100,000 simulated observations for each year.

Figure 3: Lifetime Expectancy at 66 by Gender and Socioeconomic Status

(a) Men Life Expectancy

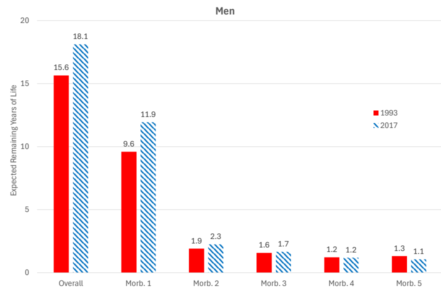

(b) Change in Men Life Expectancy

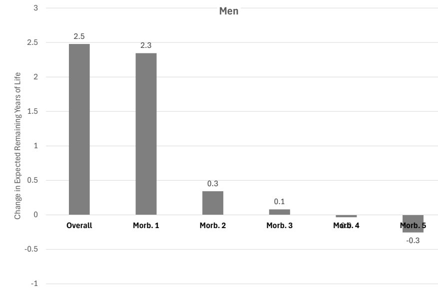

(c) Women Life Expectancy

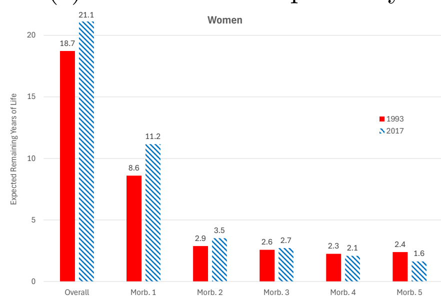

(d) Change in Women Life Expectancy

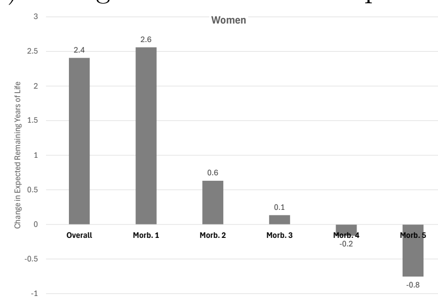

(e) Low-Income Life Expectancy

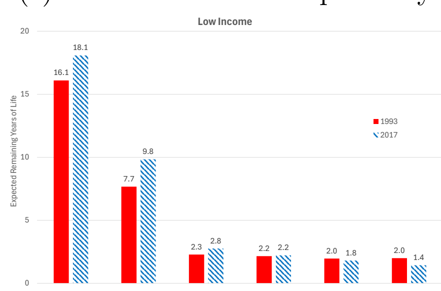

(f) Change in Low-Income Life Expectancy

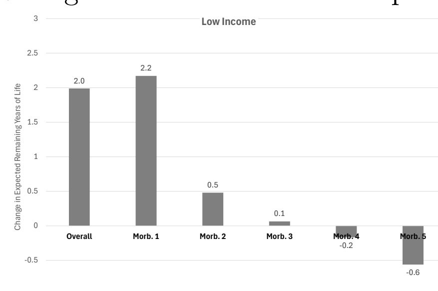

(g) High-Income Life Expectancy

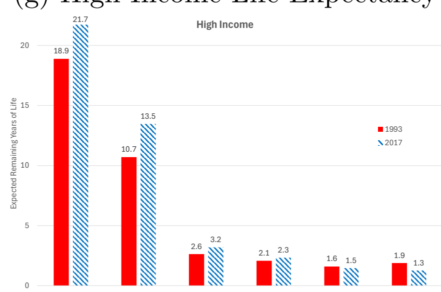

(h) Change in High-Income Life Expectancy

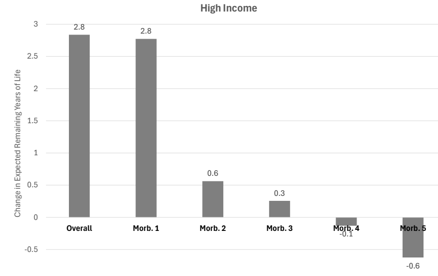

Notes: These figures display the life expectancy (first column) and difference in life expectancy from 1993 to 2017 (second column) for a 66-year old by gender and socioeconomic status. We define an individual as high- or low-income depending on whether they are above or below median income for their initial age in the data, gender, and pooled years. The starting morbidity breakdown and morbidity group transitions are allowed to vary by survey year and subgroup. N = 100,000 simulated observations for each year.

Figure 4: Trends in Healthcare Spending by Age and Years to Death

(a) By Age (b) By Years to Death

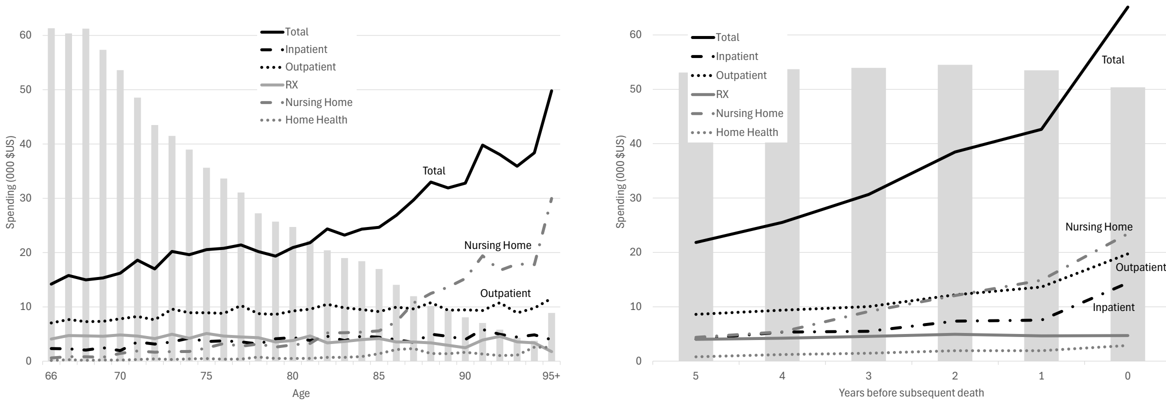

Notes: Figure displays mean spending in 2024 dollars by age (Panel (a)) and by years before death (Panel (b)), where year 0 consists of anywhere from 365 to 0 days before death. We show results by age in panel (a) for the 2017 cohort and results by years to death in panel b for the 2012 cohort. All ages above 95 are collapsed into a 95+ category. Spending is displayed for all care types (solid black line), and for individual subcategories of care. The care category labeled "Outpatient" includes hospice and dental spending as well as outpatient and physician spending. A gray histogram displays the proportion of observations in each age category in 2017 (Panel (a)) or time-to-death category in 2012 (Panel (b)). Standard survey weights are used. The sample includes individuals aged 66+ surveyed by the MCBS in 2011-2013 (N = 23,669) and in 2016-2018 (N = 19,917).

Figure 5: Health Care Spending by Morbidity Group Within Age Bins, 2017

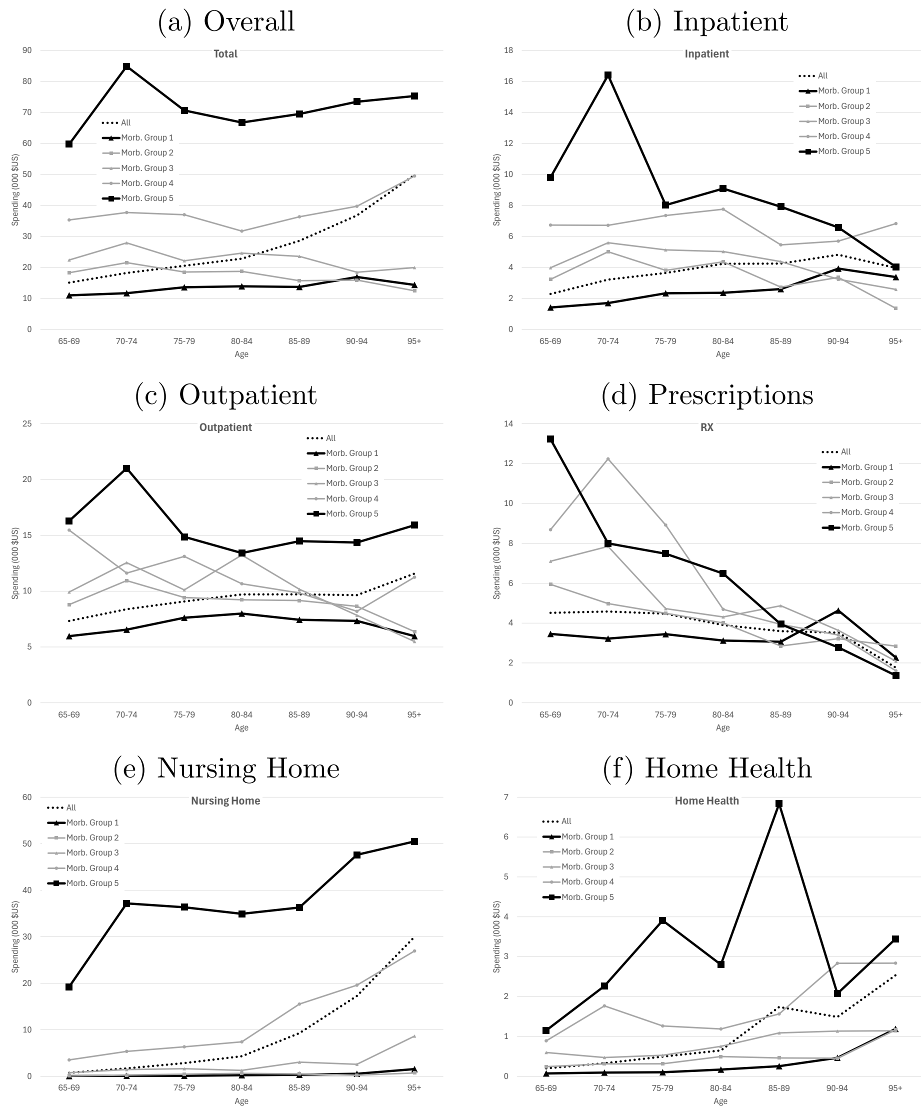

Notes: Figure displays mean spending in 2017 by morbidity group within five-year age bins, in 2024 dollars. The dotted line indicates mean spending within an age bin across all morbidity groups. All ages above 95 are collapsed into a 95+ category. Spending by morbidity group is displayed separately for all care types together (Panel (a)) and for individual subcategories of care (Panels (b) through (f)); the care category labeled "Physicians / Outpatient" includes hospice and dental spending as well as outpatient and physician spending. Standard survey weights are used. The sample includes 19,917 individuals aged 66+ surveyed by the MCBS between 2016 and 2018.

Figure 6: Lifetime Health-Care Spending at 66 by Spending Category

(a) Lifetime Spending　　(b) Changes in Lifetime Spending, 1993-2017

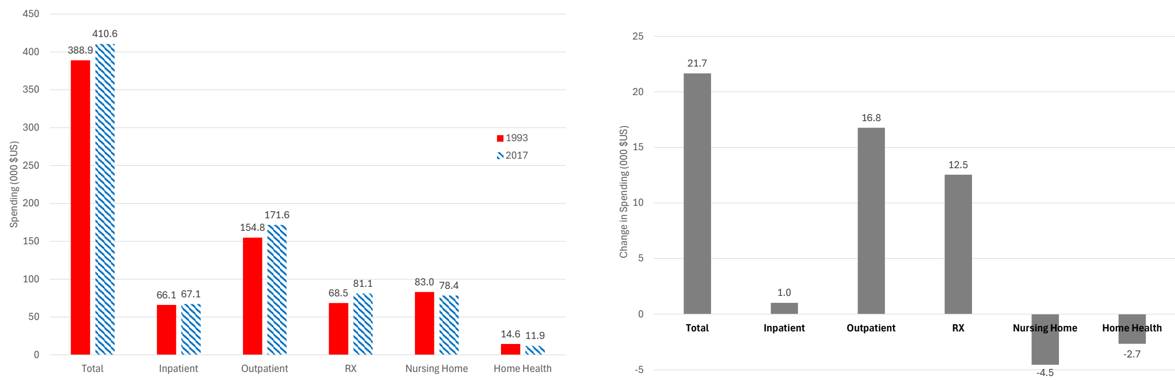

Notes: This figure displays the lifetime health-care spending of the typical 66-year-old, both overall and within each subcategory of spending. Panel (a) displays these statistics in years separately for 1993 (in red) and 2017 (in blue). Panel (b) displays the difference between these two years. The category labeled "Outpatient" includes hospice and dental spending as well as outpatient and physician spending. All spending is in 2024 dollars. The starting morbidity breakdown, morbidity group transitions, and mapping from coarse morbidity groups to granular morbidity predictors are all allowed to vary by survey year. The mapping from granular morbidity to predicted spending is held fixed by imposing the mapping generated by random forests trained on 2017 data. N = 100,000 simulated observations for each year.

Figure 7: Lifetime Health-Care Spending at 66 by Gender and Socioeconomic Status

(a) Men

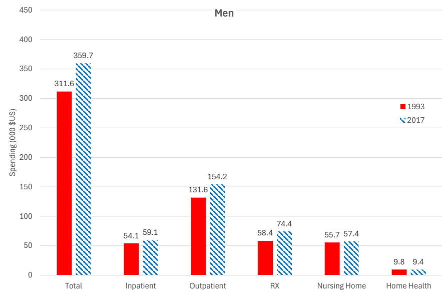

(b) Change for Men

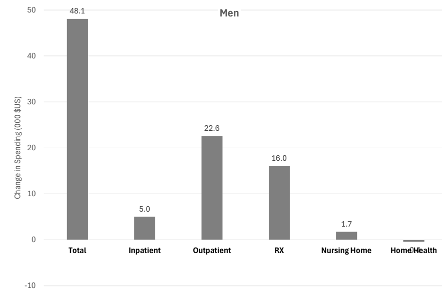

(c) Women

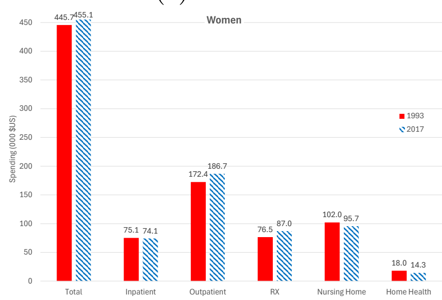

(d) Change for Women

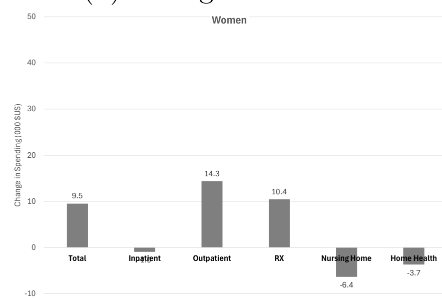

(e) Low-Income Individuals

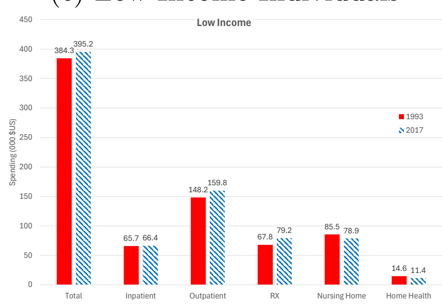

(f) Change for Low-Income Individuals

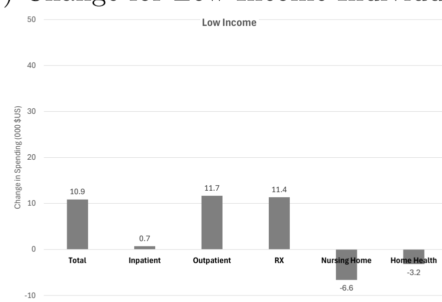

(g) High-Income Individuals

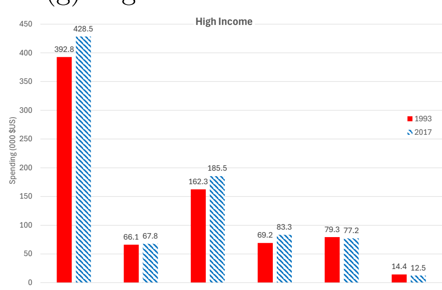

(h) Change for High-Income Individuals

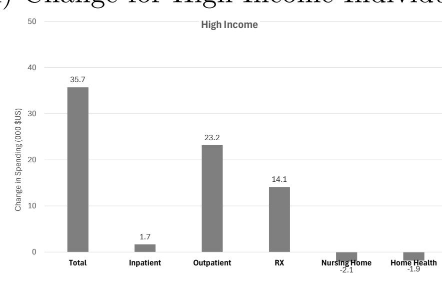

Notes: These figures display the lifetime health-care spending (first column), and difference in lifetime health-care spending from 1993 to 2017 (second column) for a 66-year old by gender and socioeconomic status. We define an individual as high or low income depending on whether they are above or below median income for their initial age in the data, gender, and pooled years. The category labeled "Outpatient" includes hospice and dental spending as well as outpatient and physician spending. All spending is in 2024 dollars. The starting morbidity breakdown, morbidity group transitions, and mapping from coarse morbidity groups to granular morbidity predictors are all allowed to vary by survey year and subgroup. The mapping from granular morbidity to predicted spending is held fixed by imposing the mapping generated by random forests trained on 2017 data. N = 100,000 simulated observations for each year.

Figure 8: Lifetime Per-Capita Government Spending at 66

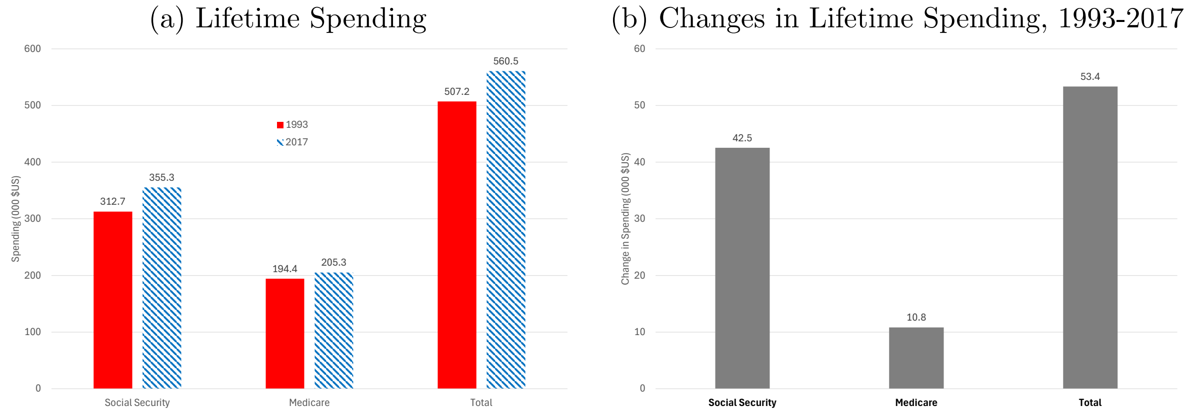

Notes: This figure displays the lifetime government spending on public pensions and public health insurance incurred by the typical 66-year-old. Panel (a) displays these statistics in years separately for 1993 (in red) and 2017 (in blue). Panel (b) displays the difference between these two years. All spending is in 2024 dollars. Public pensions are set at $17,634 per year in 2024 dollars. Public health insurance is set at 49.3% coverage. The starting morbidity breakdown, morbidity group transitions, and mapping from coarse morbidity groups to granular morbidity predictors are all allowed to vary by survey year. The mapping from granular morbidity to predicted spending is held fixed by imposing the mapping generated by random forests trained on 2017 data. N = 100,000 simulated observations for each year.

Figure 9: Iso-Budget Lines and Wealth Equivalence for 2017 Mortality and Morbidity: No Moral Hazard

Notes: This figure shows the model results under the assumption of no moral hazard. In Panel (a), the colored lines are plotted for the Medicare and Social Security combinations that, using the 2017 morbidity and mortality transitions, produce a budget that equal to the one found at 50% Medicare coverage and $18,000 Social Security in 1993 (i.e. B_{1993}, in red) and in 2017 (B_{2017}, in blue); the level of the 1993 (2017) budget is calculated using that year's morbidity-mortality transitions and the status quo policy of 50% Medicare and $18,000 Social Security payments; budget calculations include consumption floor spending. Panels (b) and (c) show the additional wealth that would make an individual indifferent between having 50% Medicare coverage and $18,000 Social Security in 2017 or having the policy combinations that generate the 1993 budget (in red) and the 2017 budget (in blue). In Panel (b), the x-axis represents Medicare coverage, with Social Security adjusted to satisfy the corresponding budget constraint; in Panel (c) the x-axis represents the Social Security policy, with Medicare coverage adjusted to satisfy the corresponding budget constraint. Budget calculations include consumption floor spending. All calculations are made for a 66 year old with $150,000 in wealth and of typical health.

Figure 10: Iso-Budget Lines and Wealth Equivalence for 2017 Morbidity and Mortality: with Moral Hazard

(a) Across Medicare Policy

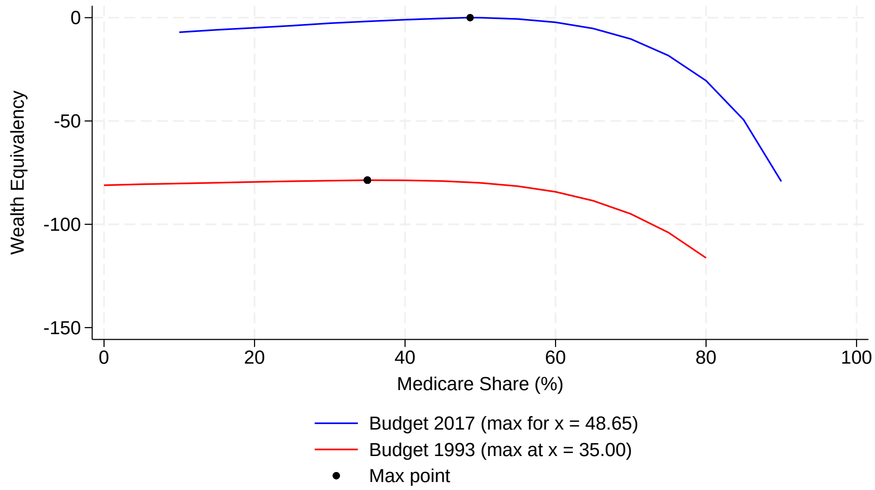

(b) Across Social Security Policy

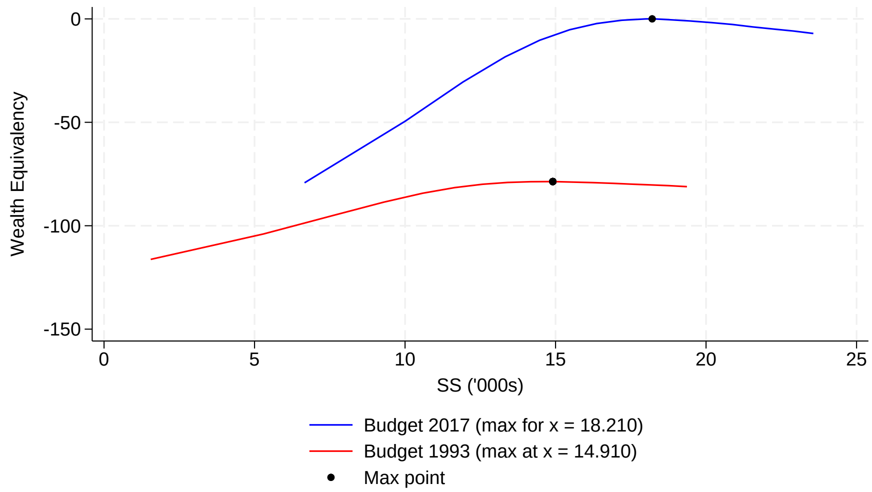

Notes: This figure shows the model results under an assumed level of moral hazard such that the status quo allocation of 50% Medicare coverage and $18,000 Social Security was optimal given 1993 mortality and morbidity. Panels (a) and (b) show the additional wealth that would make an individual indifferent between having 50% Medicare coverage and $18,000 Social Security or having the policy combinations that generate the 1993 budget (i.e, B_{1993}, in red) and the 2017 budget (i.e, B_{2017} in blue). In panel (a), the x-axis represents Medicare coverage, with Social Security adjusted to satisfy the corresponding budget constraint; in panel (b) the x-axis represents the Social Security policy, with Medicare coverage adjusted to satisfy the corresponding budget constraint. The level of the 1993 (2017) budget is calculated using that year's morbidity-mortality transitions and the status quo policy of 50% Medicare and $18,000 Social Security payments; budget calculations include consumption floor spending. All calculations are made for a 66 year old with $150,000 in wealth and of typical health.

Online Appendix

A Data and key measures

A.1 Data

Every year, the MCBS, which is sponsored by the Centers for Medicare and Medicaid Services (CMS), surveys on the order of 10,000 randomly selected current Medicare beneficiaries. These include residents of nursing homes and other long-term care facilities, who are often omitted from other comparable surveys; previous literature has often chosen to employ the MCBS for this reason (Cutler et al. 2014). The MCBS has been conducted annually since 1991, with the exception of 2014. The MCBS has a rotating panel with an inclusion period of three to four years. Some of our analyses will treat the data as a repeated cross-section, while in other we will exploit the panel dimension to estimate transition probabilities across health states.

The MCBS has undergone several evolutions since its inception. Until 2013, it was divided into two distinct surveys: an "Access to Care" survey file that started in 1991 and focused on questions around access to and satisfaction with health care, and a "Cost and Use" file that started in 1992 and focused on utilization and spending. The population of individuals included in the two surveys is similar but not identical: the Access to Care file was designed to be representative of the population that was enrolled in Medicare for the entire survey year in question, while the Cost and Use file is representative of the population enrolled in Medicare for any portion of the survey year.[24] After 2013, the survey was re-designed. It did not take place in 2014, and resumed in 2015 with a "Survey File" (covering similar information to the Access to Care file) and a "Cost Supplement File" (covering similar information to the Cost and Use file). On average from 1992 through 2019, 91% of those interviewed by the Access to Care survey (and its successor) were also interviewed by the Cost and Use survey in that year, and 65% of those interviewed by the Cost and Use survey were also interviewed by the Access to Care survey.

In both surveys, we observe basic demographic characteristics including age, gender, race/ethnicity, marital status, and education. In addition, we observe nursing home residence status and answers to a range of health-related survey questions, including activities of daily living (ADLs), instrumental activities of daily living (IADLs), and whether an individual has ever been diagnosed with a number of diseases. The MCBS typically asks questions directly to the individual of interest, but will instead survey family members if an individual is incapacitated; in addition, the MCBS surveys facility staff, rather than the individual in question, when the individual is a current resident of a long-term care facility.[25]

MCBS data can be linked to the Medicare denominator files, which include date of death. We have access to this Medicare linkage for death data (through 2019) for individuals in the Access to Care surveys (and its post 2013 successor, the Survey File).[26] In addition, the Cost and Use survey (and post-2013 successor Cost Supplement survey) ask a range of questions on annual health-related expenditures.

[24] Specifically, the "Access to Care" sample limits to those who were continuously enrolled in the MCBS from the first day of the calendar year through the date of their interview, which typically takes place in the fall (CMS 2021b). To the best of our knowledge, individuals are not removed from this sample if their Medicare coverage is terminated after their interview date but prior to the last date of the calendar year. Accordingly, we observe individuals with post-interview deaths within the same year for the "Access to Care" and "Cost and Use" sub-samples alike.

[25] For an overview on how the MCBS defines facilities and conducts facility interviews, see CMS (2019).

[26] This linkage consists of MCBS respondents' Medicare beneficiary ID: *bene_id* for later observations, and *ehic* for earlier observations.

A.2 Sample selection.

We use 27 years of data from 1992 through 2019. As noted, the MCBS was not conducted in 2014, and we start in 1992 (rather than in 1991) because we want to include individuals who appear in both the Access to Care and Cost and Use Surveys. Since our analysis requires data both on date of death (if any) as well as information on health care spending and its components, we limit our sample each year to the intersection of those interviewed by the Access to Care and Cost and Use surveys.[27] Although the MCBS surveys Medicare recipients of all ages, we restrict our attention to the elderly, specifically, individuals aged 66 and over.[28]

Appendix Figure OA.1 displays the distribution of age, morbidity, and spending for the Access to Care sample, the Cost and Use sample, and our intersection; comfortingly, the intersection sample exhibits highly similar trends to both the Access to Care and Cost and Use samples. We follow the data documentation's recommendation to use the survey weights from the Cost and Use file when analyzing the intersection sample (CMS 2021a).[29] We additionally adjust the weights such that every year's weight is the same; we do this so that the 3-year moving average time trend figures and the transition matrices perform a simple average across time.

After all of these restrictions, our analysis sample consists of 226,647 individual-years. On average, our sample size is 8,394 individuals per year. We will often show statistics for specific ages and specific morbidity groups using data pooled across three consecutive survey years – for specific ages and specific morbidity groups. To give a sense of sample size, Appendix Table OA.2 shows sample sizes by age and morbidity group for 2016-2018 (which we refer to throughout as '2017').

A.3 Key variables

Morbidity Using MCBS survey questions on ADLs, IADLs, functional limitations, and diseases, we develop annual measures of morbidity for the 66+ population between 1992 and 2019. Broadly, ADLs ask about individuals' ability to independently carry out physical tasks that are typically part of everyday life, including eating, dressing, and bathing. IADLs ask about somewhat more complex tasks that may require greater executive functioning skills, including managing money, shopping, and preparing meals. The MCBS asks about six ADLs and six IADLs, each of which is a binary question on whether individuals experience any difficulties performing the given task. The MCBS also asks about five "functional limitations." These consist of comparatively more difficult physical activities that may or may not come up in the course of everyday life, including lifting 10 pounds, raising one's hands above their head, and walking a quarter of a mile. Unlike ADLs and IADLs, these questions are asked on a sliding scale ranging from "no difficulty" to complete inability to perform the task; we recode individuals as having the functional limitation if they report any level of difficulty above the minimum.[30] Finally, the MCBS asks whether individuals have ever been diagnosed with 16 different

[27] The typical sequence is that individuals are interviewed in Year 1 by the Access survey, and in Years 2-4 by the Access and Cost surveys. Since we limit our analysis to person-years that are in both, we typically observe a given individual for three years.

[28] We limit the sample to ages 66 and older instead of 65 and older as the 65-year-olds interviewed by the MCBS are disproportionately those that qualify for Medicare through a disability, end-stage renal disease (ESRD), or amyotrophic lateral sclerosis (ALS). This occurs because Access to Care was designed to be representative of the population that was enrolled in Medicare for the entire survey year, and therefore it is not always administered to those who enter Medicare at age 65 through age-based eligibility. This policy changed in 2014, where the sampling frame was expanded to include those missing 65 year old individuals. To ensure continuity across the redesign, we also define age as of January 1 on the year the survey was administered. This guarantees that everyone who is considered 66 who qualified through the age pathway will have had at least one year being on Medicare.

[29] We also restrict to individuals with non-missing health information, and those with a unique crosswalk to Medicare data; these restrictions remove only a handful of observations.

[30] This follows the same recoding procedure as previously adopted by Cutler et al. (2014).

chronic diseases. These are binary questions, and ask individuals whether they have ever been told by a medical professional that they have a particular condition.

The health questions we leverage in this paper are those that maintain relatively consistent wording year-over-year and do not exhibit economically meaningful discontinuities between survey years that could be suggestive of changes in survey design. In practice, we retain all ADLs, IADLs, and functional limitation measures, and 14 out of the 16 disease measures. [31]

Morbidity groupings. For some of our analyses, we will find it useful to categorize individuals into five ascending "morbidity groups" based on the number of ADLs and IADLs currently experienced. ADLs and IADLs as the most typical measure of disability across the literature on aging and morbidity (e.g. Cutler et al. (2014); Freedman et al. (2002)). Accordingly, we construct morbidity groups based on total number of ADLs and IADLs present as follows: Group 1 has 0 ADLs/IADLs, Group 2 has 1 ADL/IADL, Group 3 has 2-3 ADLs/IADLs, Group 4 has 4-7 ADLs/IADLs, and Group 5 has 8+ ADLs/IADLs. Individuals can have up to 12 ADLs/IADLs (6 ADLs and 6 IADLs). These groups are defined to create a roughly flat overall distribution of groups between Groups 2-5, though approximately half of all observations are in Group 1.

Health-care spending. Our primary measurement of cost is total annual expenses across all payer types (which consist of Medicare (both Traditional Medicare and Medicare Advantage), Medicaid, a range of private insurers, and self-pay), and all MCBS-reported categories of care (inpatient, outpatient, long-term care, medical provider encounters, prescriptions, dental, home health, and hospice care). Unless stated otherwise, these expenses are reported in 1992 dollars.

Respondents are asked about their health expenditures by category of care (including inpatient, nursing home, and home health care), and payer (including Medicare, Medicaid, and self-pay). The MCBS then validates these self-reported costs against administrative data on claims from Traditional Medicare, as well as on home health care use and nursing home care regardless of payer from the Outcome and Assessment Information Set (OASIS) and Minimum Data Set (MDS) respectively, and uses this comparison to adjust its estimates of expenditures.[32] We collapse all health spending data to the person-year level and standardize all expenses to 1992 dollars.

Mortality. We obtain death dates by linking the MCBS to the Medicare Master Beneficiary Summary File, which has death dates (if any) for all Medicare enrollees, including those enrolled in Medicare Advantage. We say an individual transitions from a given morbidity state into death if the death date is within a calendar year of their last interview date with morbidity state information. In the rare circumstance that the death date is before the last interview date, we recode the death date as after the last interview date if it is less that 90 days before the last interview date. If the death date is more than 90 days before the last interview date, we assume that the death date is erroneous and assume the individual has not died.

[31] The only diseases omitted from our sample are those on the presence of arthritis and depression, both due to differences in wording across years. We also omit self-reported measures of sight and hearing due to similar inconsistencies. We note that, starting in 1997, MCBS respondents residing in SNFs receive a slightly different health questionnaire, where both question wording and potential responses are somewhat different. In the nursing-home-specific surveys, ADLs, IADLs, and functional limitations alike are measured on a sliding scale, ranging from total independence to total dependence rather than asking whether individuals experience any difficulty performing the activity. In this case, we recode response above the lowest disability level as "1", following Cutler et al. (2014). We perform similar transformations for other variables with changing formats between the non-nursing home and nursing home health questionnaires.

[32] For the validation, they observe claim-level data for all Traditional Medicare expenses, and comparatively more aggregated data for all non-Medicare expenses (typically at the encounter, stay, or person-year level, depending on the expense type).

We define interviews to be "0" years to death if the interview happened anywhere from 360 to 0 days before death. Note that it is possible that multiple interviews from one individual can occur within a 365 day window.

A.4 Transition matrix creation

As discussed in section 2, we construct age- and year-specific transition matrices by leveraging the rotating panel structure of the MCBS. However, given that we must construct these transition matrices with morbidity groups, there are a few additional steps that must be taken beyond the standard period life expectancy table.

First, we must account for the difference in year $t+1$ information about morbidity versus mortality. Each age-specific transition matrix in year t will have rows according to the starting, year t, morbidity group. However, in year $t+1$ we observe morbidity status only for those who are alive and remain on the MCBS panel. In contrast, we observe mortality outcomes for everyone who die in year $t+1$. To account for this difference, we scale down the number of people who died in $t+1$ by the probability of being interviewed if one were alive in $t+1$. We do so conditional on year t morbidity group before calculating the mortality rate within each starting morbidity group to account for the possibility that individuals of different morbidity are associated with different probability of remaining on the MCBS panel.

To give a concrete example, suppose that at age 80 there are 100 people who start in morbidity group 1. Say that at age 81, 9 of these people end up in morbidity group 1, 9 end up in morbidity group 2, 9 end up in morbidity group 3, 9 end up in morbidity group 4, and 9 end up in morbidity group 5. Say that there are 45 people of these people who are alive but are not interviewed at age 81. Finally, say that there are 10 people who are dead at age 81. The probability of being interviewed if the individual is alive at 81 is $\frac{45}{90} = 0.5$. That means that the number of people who would have been interviewed if they were alive is $10 \cdot 0.5 = 5$. That row of the transition matrix would use the $45 + 5 = 50$ as the denominator, with 9 people going to each morbidity group and 5 dying. That row's values would therefore be $[0.18, 0.18, 0.18, 0.18, 0.18, 0.1]$.

Second, because we estimate transition matrices for each age and year separately, certain combinations of age and morbidity group are associated with very few individuals who are available to inform the transition matrix proportions. For example, only a small number of 95-year-olds are in morbidity group 1. Moreover, unless these individuals either have an interview at age 96 or die, they provide no information for the transition matrix. We do two things to amend this. First, we pool years. A standard period life expectancy table for year t only uses transitions for each age from year t to year $t+1$. However, in our analysis, for year t, we use transitions for each age from year $t-1$ to t, from year t to $t+1$, and from year $t+1$ to $t+2$, and then take the average of the three transitions as the final transition matrix. Second, we pool ages. For a given age a and morbidity group m, if the number of transitions[33] we have is less than 15, we will also add the transitions from ages $a-1$ and $a+1$ and morbidity group m to inform the transition probabilities in the row for morbidity group m in age a's transition matrix.[34] If after pooling we still have no transitions to another morbidity group, we will continue pooling the surrounding ages until there is at least one transition to another morbidity group.

[33]The number of transitions includes both transitions to another morbidity group and transitions to death, with transitions to death scaled down by the probability of being interviewed if alive as described earlier in the section.

[34]For the edge ages 66 and 99, we will simply pool the three nearest available years. So if morbidity group m needing pooling for age 66, we would add transitions from individuals starting in morbidity group m at age 67 and age 68

B Mapping from health and age to health-care spending

To allow flexible interaction between predictors, we use random forests to predict spending. Specifically, we create 200 trees, and each tree subsamples from our full sample, with replacement, until arriving at a sample equal to our full sample size. To determine the minimum node size and the number of predictors considered at each branch, we create distinct random forests from a two-dimensional grid of these hyperparameters, employing 10-fold cross-validation for each random forest, and choose the set of hyperparameters with the lowest out-of-sample RMSE.

When first generating random forests, we include all ADLs, IADLs, functional limitations, and diseases as predictors, but no demographic variables. The outcome being predicted is total yearly health spending in 2024 dollars, aggregated across all providers and for all utilization types. Figure OA.7 plots the importance of these predictors in the random forest.

Once these initial predictions are generated, we incorporate a proportional age-specific shifter. We do this so that the effect of morbidity status and age on spending are multiplicatively separable from one another. Note that all ages have at least 15 observations, and so none are pooled as we do in the transition matrices.

To incorporate this proportional age-specific shifter, we first define a multiplier m_a as the weighted average actual spending of individuals of age a divided by the weighted average predicted spending of individuals of age a:

$$m_a = \frac{\sum\limits_{i,t:\alpha(i,t)=a} y_{it}}{\sum\limits_{i,t:\alpha(i,t)=a} \hat{y}_{it}} \tag{6}$$

where $\alpha(i,t)$ is defined to be the age of individual i in year t. Note that we calculate these multiplier fixed effects using only data in the years we trained our random forest on (i.e., 2016-2018). We can then take this multiplier and apply it to the original random forest spending predictions:

$$\hat{\hat{y_{it}}} = \hat{y_{it}} \cdot m_{\alpha(i,t)} \tag{7}$$

where $\hat{\hat{y_{it}}}$ is predicted spending from the random forest after applying the proportional age-specific shifter.

Figure OA.8 plots these multiplier fixed effects. Consistent with observed patterns, it turns out that age is hardly associated with overall spending once accounting for morbidity.

C Model calibrations and computations

C.1 Calibration.

Baseline policy parameters. We set the baseline Medicare policy level (λ) to 50%, which is based on the 49.3% estimate of the share of total health expenditures covered by Medicare in the MCBS data in 1993 (i.e. using the 1992–1994 data). We analyze individual-years in which individuals were between 66 and 100 years old and since we only see Traditional Medicare expenses, we further restrict to individual-years in which the individual is not enrolled in an HMO. We estimate an average Medicare coverage share by dividing the sum of all Traditional Medicare expenses by the sum of total expenses. Similar values have been found by De Nardi et al. (2016) and Chulis et al. (1993).[35]

[35]This average Medicare coverage share of 50% masks considerable variation in cost sharing across different types of care. For example, we estimate that in 1993 Medicare covered 0% of prescription drug spending and 7% of nursing home spending, compared to 85% of inpatient spending.

We set the baseline Social Security annuity payment (ss) to \$18,000, which is based on the \$17,634 value derived from the average annual Social Security benefit amount for retired workers from 1993 after cost of living adjustments, inflated to 2024 dollars (SSA 2025a).

For the consumption floor ($\underline{c}$), we follow Brown and Finkelstein (2008) who use the level of SSI benefits to approximate a consumption floor. We set the baseline consumption floor to \$11,000, which is based on the \$11,342 value derived from the maximum annual Federal SSI payment amounts for an eligible individual in 1993 after cost of living adjustments, inflated to 2024 dollars (SSA 2025b).

Other calibrated parameters. We set baseline wealth (w_{66}) to be \$150,000. This is the median total non-annuitized wealth, excluding housing and other real estate, held by 65-69 year olds in 2008 (see Table 9 in Poterba (2014)) inflated to 2024 dollars. For risk aversion, we follow Einav et al. (2010) and the references therein and set $\gamma = 3$.

Smoothed medical spending distribution. The distribution of health-care spending $G_a(m|h)$ comes from the distribution of predicted spending for each age-morbidity combination. For computational purposes, we smooth this predicted spending distribution by running the following regression for each morbidity group $h \in \{1, 2, 3, 4, 5\}$:

$$\log y_i = \beta_{0,h(i)} + \beta_{1,h(i)}\,(\text{age}_i - 66) + \varepsilon_{i,h(i)}, \qquad \varepsilon_{i,h(i)} \sim N(0, \sigma^2_{h(i)}) \tag{8}$$

where y_{ih} is predicted spending from the random forest, adjusted for multiplicative age effects, for individual i in morbidity group h. We subtract age by 66 so that the intercept is more interpretable.

Using these regression, we obtain an estimated distribution $N(\hat{\mu}_{ah}, \hat{\sigma}^2_h)$ for each age–morbidity combination (a, h). We create bins of $< \log(500), \log(500) - \log(1500), \log(1500) - \log(2500), ...$ and find the probabilities for each bin according to the normal distribution $N(\hat{\mu}_{ah}, \hat{\sigma}^2_h)$. We then map the bins back to levels (i.e., $< 500, 500 - 1500, 1500 - 2500, ...$) and preserve the probabilities we calculated for each bin. In the model, an individual will get hit by the "center point" of each of these bins (i.e., $0, 1000, 2000, ...$) with the respective probability. Finally, we truncate the distributions at \$160,000 by cutting off the probability mass beyond that point and re-normalize so that the whole distribution still sums to one.[36]

C.2 Calculating Budget-neutral policy combinations and wealth equivalents

Our main counterfactuals use the 2017 morbidity, mortality, and spending distributions, and estimate wealth equivalents z under alternative policy parameters (λ', ss') that hold total government spending fixed at either B_{1993} or B_{2017}. This requires that we first *find* these budget-neutral policy regimes. To do so, we solve the model to find the possible sets of policy parameters (λ', ss') that, given 2017 morbidity, mortality and spending distributions, *as well as* the individual's optimal consumption choices, yield the total government spending amount B_{1993} or B_{2017}. We do this using a grid of possible policy combinations – given by (λ', ss') – and linearly interpolate the overall government spending between the grid points. We use the same interpolation procedure to find the expected lifetime utility at each budget-neutral policy regimes, which in turn are used to calculate z.

[36] The actual maximum predicted spending is around \$144,200.

Appendix References

Brown, Jeffrey R and Amy Finkelstein, "The interaction of public and private insurance: Medicaid and the long-term care insurance market," *American Economic Review*, 2008, *98* (3), 1083–1102.

Chulis, George S., Franklin J. Eppig, Mary O. Hogan, Daniel R. Waldo, and Ross H. Arnett, "Health Insurance and the Elderly: Data from MCBS," *Health Care Financing Review*, 1993, *14* (3), 163–181.

CMS, "MCBS Advanced Tutorial on Using Community and Facility Data," 2019.

_ , "MCBS Advanced Tutorial on Weighting and Variance Estimation," Technical Report 2021.

_ , "Medicare Current Beneficiary Survey (MCBS): Frequently Asked Questions," Technical Report 2021.

Cutler, David, Kaushik Ghosh, and Mary Beth Landrum, "Evidence for Significant Compression of Morbidity in the Elderly U.S. Population," *Discoveries in the Economics of Aging*, 2014, pp. 21–51.

Einav, Liran, Amy Finkelstein, and Paul Schrimpf, "Optimal Mandates and The Welfare Cost of Asymmetric Information: Evidence from The U.K. Annuity Market," *Econometrica*, 2010, *78*, 1031–1092.

Freedman, Vicki, Linda Martin, and Robert Schoeni, "Recent Trends in Disability and Functioning among Older Adults in the United States: A Systematic Review," *JAMA*, 2002, *288* (24), 3137–3146.

Nardi, Mariacristina De, Eric French, John Bailey Jones, and Jeremy McCauley, "Medical Spending of the US Elderly," *Fiscal Studies*, 2016, *37* (3-4), 717–747.

Poterba, James M, "Retirement security in an aging population," *American Economic Review*, 2014, *104* (5), 1–30.

SSA, "Benefits Paid by Type of Beneficiary," Technical Report 2025.

_ , "SSI Federal Payment Amounts," Technical Report 2025.

Figure OA.1: Yearly Summary Statistics by MCBS Subsample

(a) Age

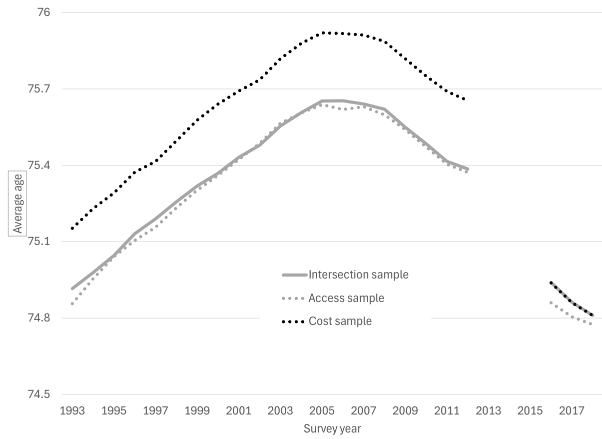

(b) #ADL/IADLs

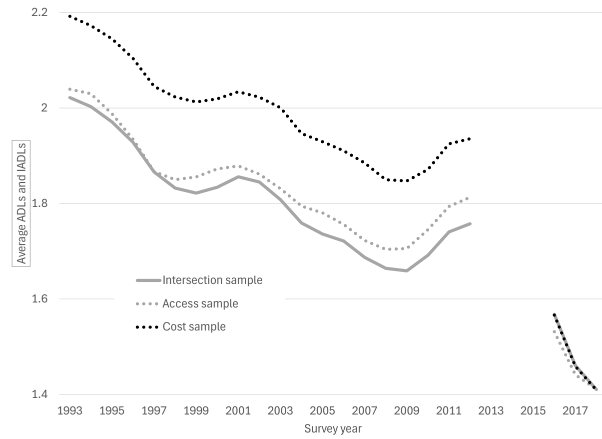

(c) Overall Spending

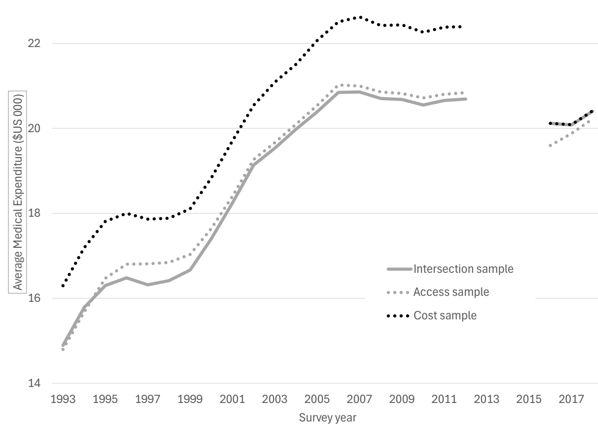

Notes: Figure displays, by 3-year moving average, summary statistics for three distinct subsamples of the MCBS: the Access to Care sample, weighted with Access to Care weights (grey dotted line); the Cost and Use sample, weighted with Cost and Use weights (black dotted line); and the intersection between the two that forms our primary subsample, weighted with Cost and Use weights (grey solid line). Yearly means are shown for each sample for age (Panel (a)), total number of ADL/IADLs (Panel (b)), and total per-capita health spending (Panel (c)). The sample includes 348,743 individuals aged 66+ surveyed by the MCBS for either the Access to Care or the Cost and Use sample between 1992 and 2019. The MCBS survey does not exist for 2014, so (3-year moving average) is missing for 2013-2015.

Figure OA.2: Morbidity Group Distribution by Age, 2017

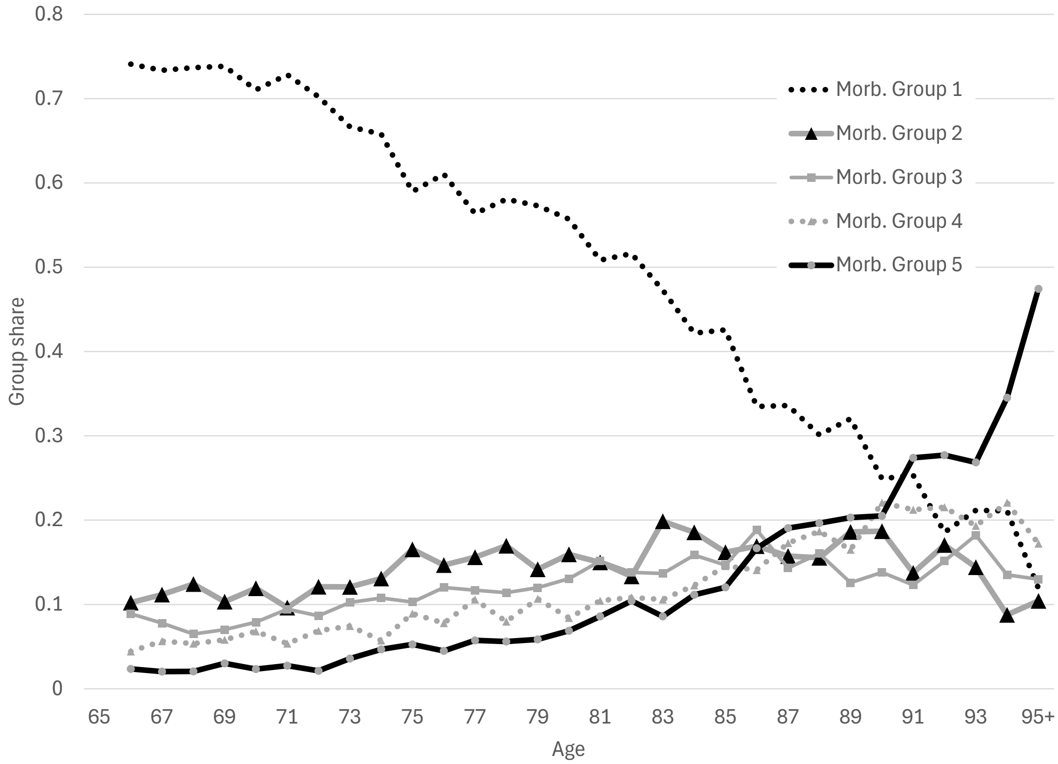

Notes: Figure displays the proportion of the surveyed population in each of five morbidity groups by age, restricting to years 2016-2018. All ages above 95 are collapsed into a 95+ category. Standard survey weights are used. N = 19,917 individuals aged 66+ surveyed by the MCBS in 2017.

Figure OA.3: Trends in ADL/IADLs by Age

(a) Any ADL/IADL (b) Mean ADL/IADLs (Conditional)

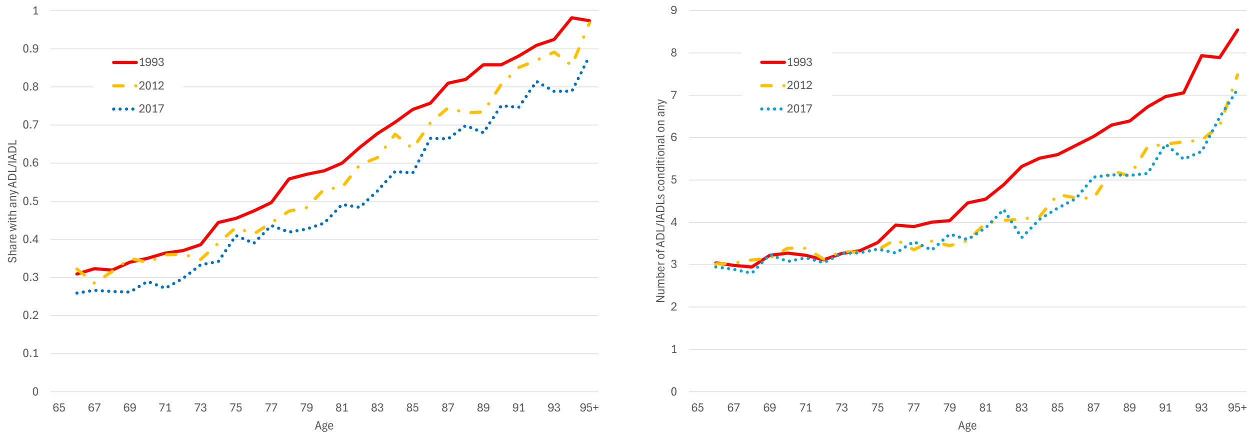

Notes: Figure displays, for each age, the presence of any ADL/IADL (Panel (a)), and mean ADLs/IADLs conditional on having at least one (Panel (b)), separately for the 1992-1994 survey cohort (in red), the 2011-2013 survey cohort (in yellow), and the 2016-2018 survey cohort (in blue). All ages above 95 are collapsed into a 95+ category. Standard survey weights are used. The sample includes individuals aged 66+ surveyed by the MCBS in 1993 (N = 26,413), 2012 (N = 23,669), or 2017 (N = 19,917).

Figure OA.4: Trends in ADL/IADLs by Years to Death

(a) Any ADL/IADL (b) Mean ADL/IADLs (Conditional)

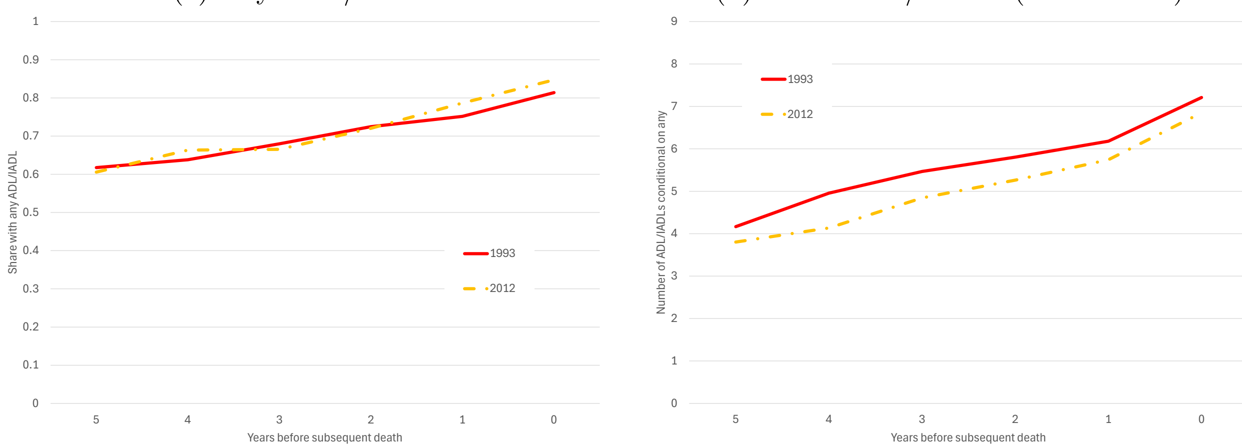

Notes: Figure displays, for years to death ranging from 0 to 5, the presence of any ADL/IADL (Panel (a)) and mean ADLs/IADLs conditional on having at least one (Panel (b)), separately for the 1993 survey cohort (in red) and the 2012 survey cohort (in yellow). The 2012, rather than 2017, survey cohort is used for the post-period so that deaths within five years can be observed without censoring for all respondents. Standard survey weights are used. The sample includes individuals aged 66+ surveyed by the MCBS in 1993 (N = 26,413) and in 2012 (N = 23,669).

Figure OA.5: Trends in Morbidity Measures, 1993 and 2017

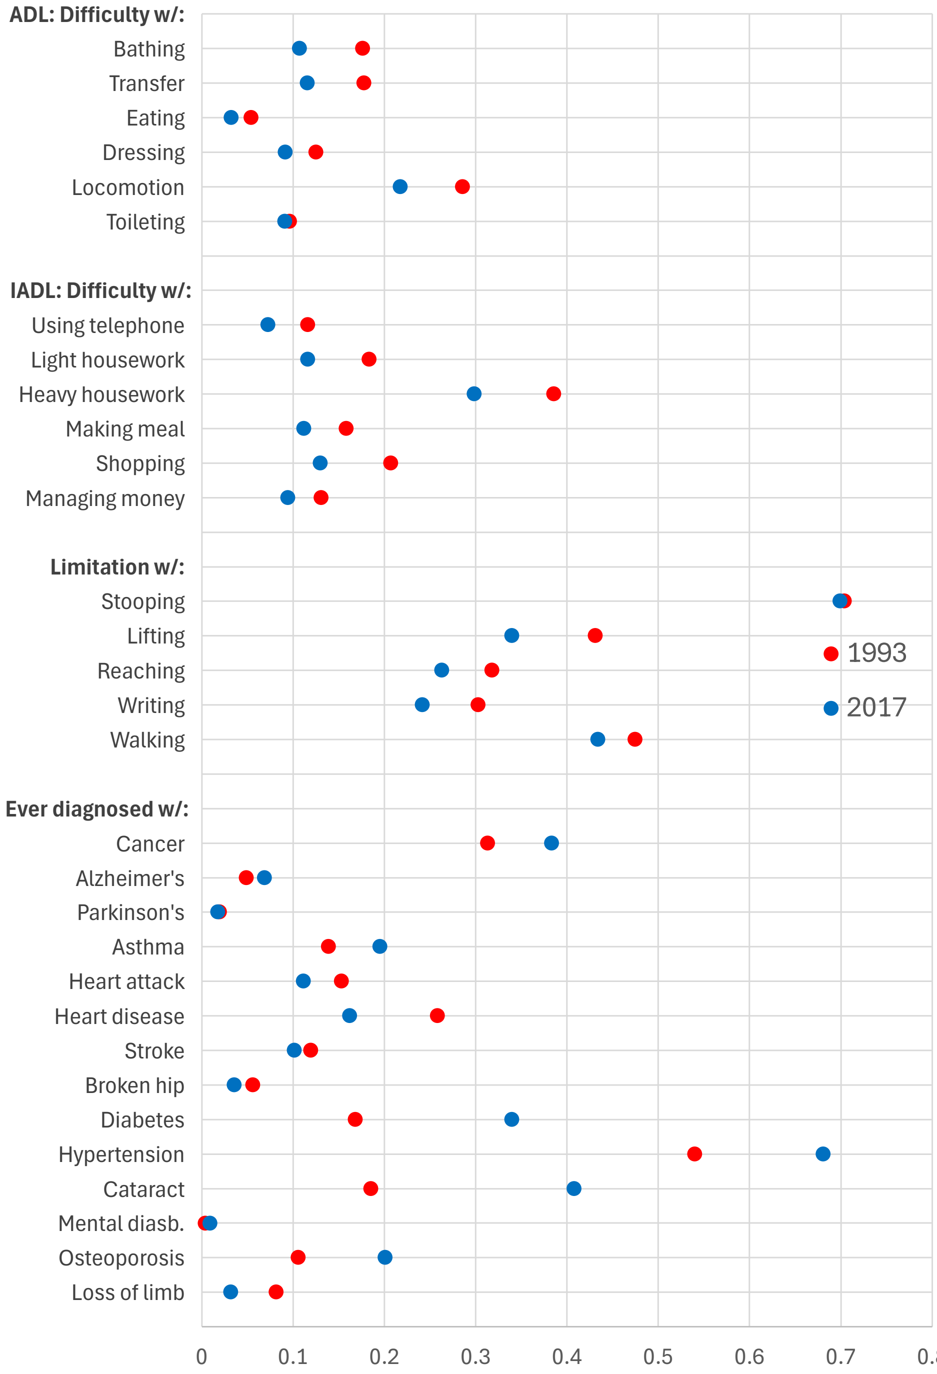

Notes: Figure displays, by 3-year moving average, the age-adjusted share of individuals with ADLs, IADLs, functional limitations, and diseases. Standard survey weights are used; in addition, the sample is re-weighted such that the age distribution in all years matches the age distribution in 1999-2001 and that all years have the same weight. The sample includes individuals aged 66+ surveyed by the MCBS in 1993 (N = 26,413) or 2017 (N = 19,917).

Figure OA.6: Yearly Trends in ADL/IADLs and Morbidity Expectancies

(a) Age-Adjusted ADL/IADLs

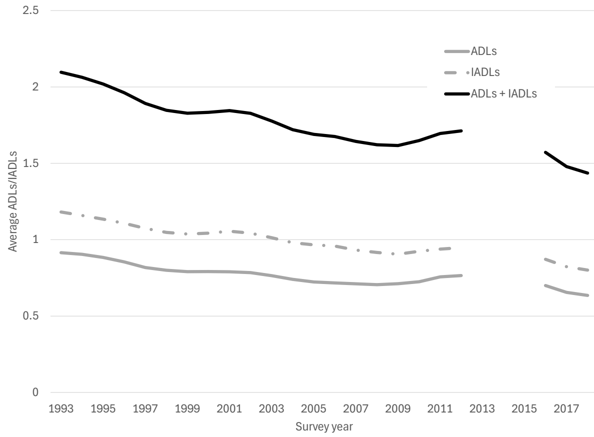

(b) Yearly Life Expectancy

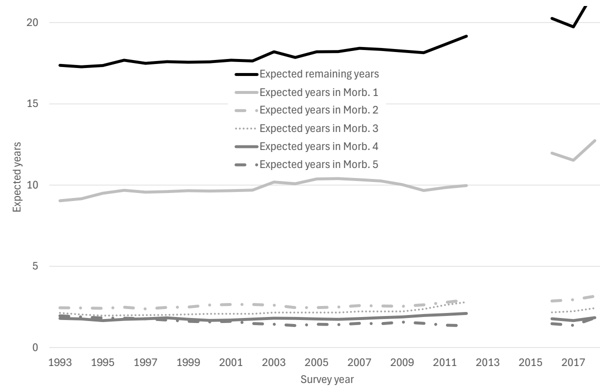

Notes: Figure displays, by 3-year moving average, the mean, age-adjusted number of ADLs and IADLs (Panel (a)) and life expectancy by morbidity group and overall (Panel (b)). Each year's life expectancy is derived using corresponding pooled transition matrices. Standard survey weights are used; in addition, the sample in Panel (a) is re-weighted such that the age distribution in all years matches the age distribution in 1999-2001 and that all years have the same weight. Data points 2013-2015 are not shown since there is no 2014 MCBS survey and three years of data must inform each data point. N = 226,647 individuals aged 66+ surveyed by the MCBS between 1992 and 2019 for Panel (a). N = 100,000 simulated observations for each year for Panel (b).

Figure OA.7: Random Forest Variable Importance of Predictors

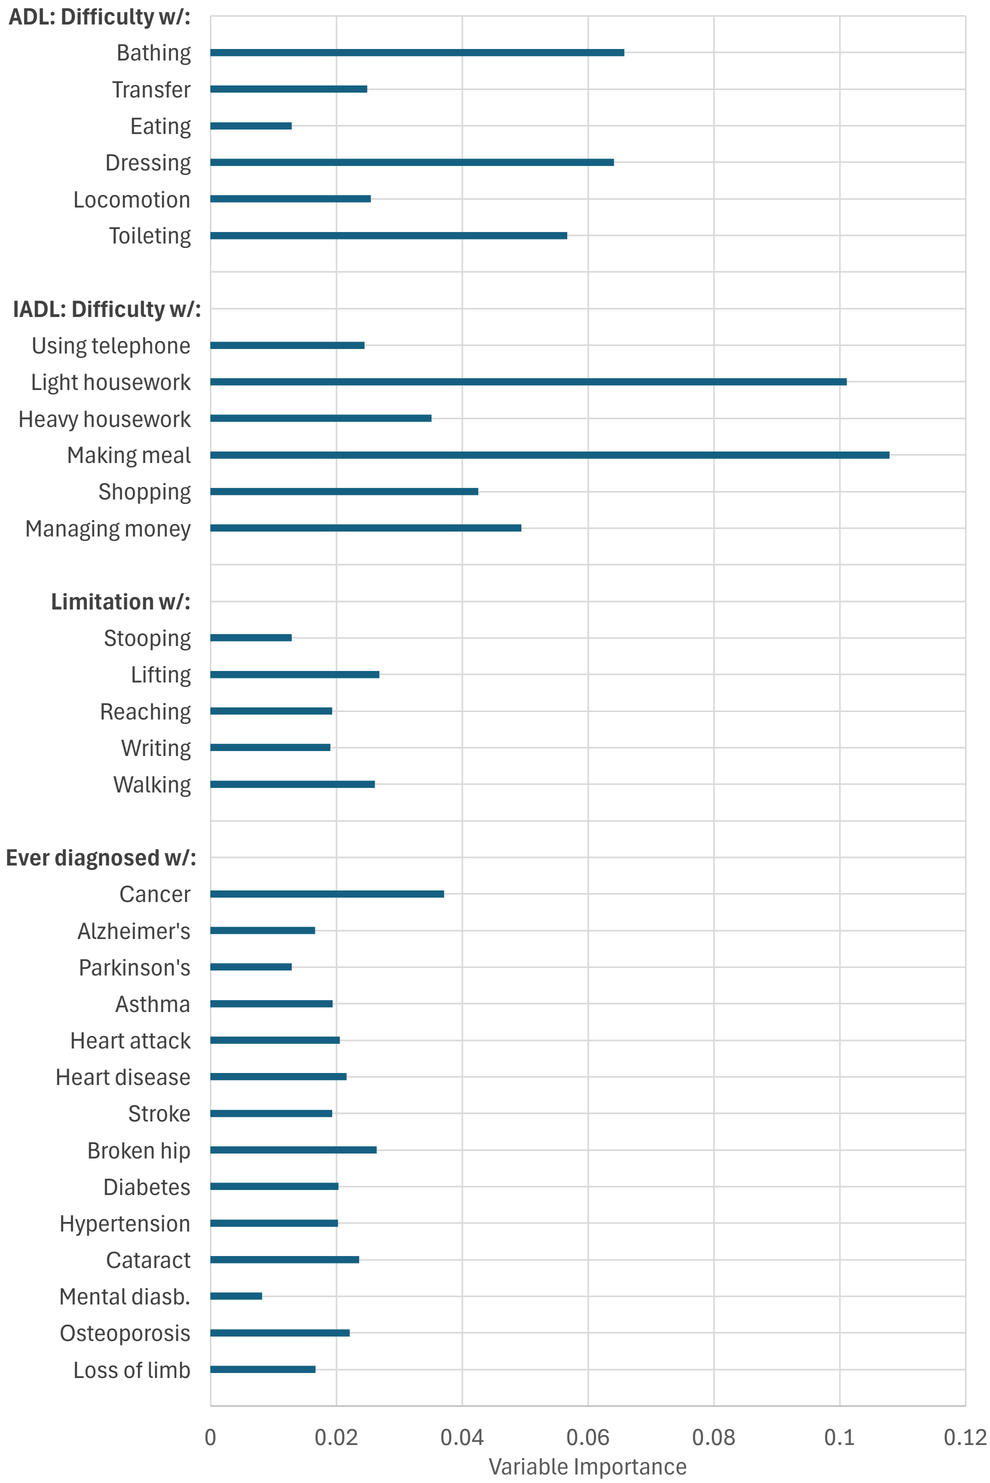

Notes: Figure displays, for a given variable used as a predictor for random forests, the variable's importance when used to predict total annual spending. The variable importance term corresponds to the reduction in MSE from the inclusion of that variable; variable importance measures in this figure are normalized to sum to 1 across all predictors. Predicted spending is calculated using random forests trained on 2017 data, wherein a multiplicative age effect is added after the fitting of random forests per the methodology described in Section 3.2. Standard survey weights are used. The sample includes 226,647 MCBS survey respondents between 1992 and 2019.

Figure OA.8: Random Forest Age Fixed Effects

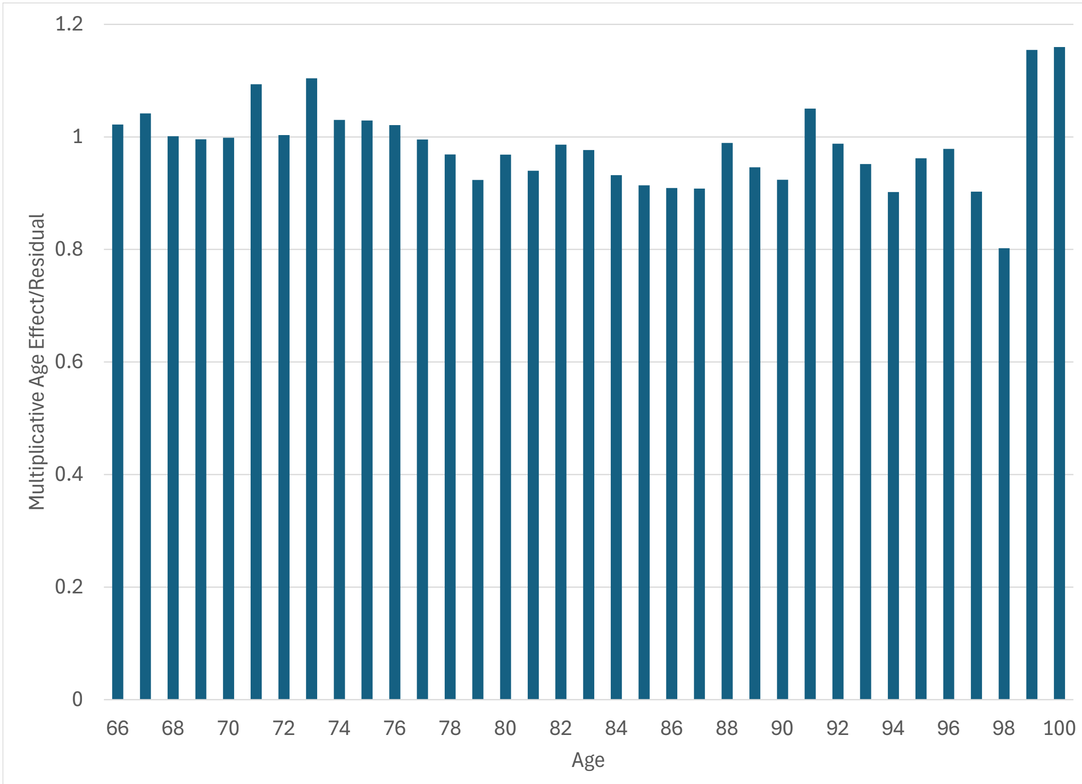

Notes: Figure displays, for a given age, the multiplicative fixed effect applied to predictions of total annual spending for that age. Predicted spending is calculated using random forests trained on 2017 data, wherein a multiplicative age effect is added after the fitting of random forests per the methodology described in Section 3.2. Standard survey weights are used. N = 226,647 MCBS survey respondents between 1992 and 2019.

Figure OA.9: Predicted Health-Care Spending by Years to Death, 2012

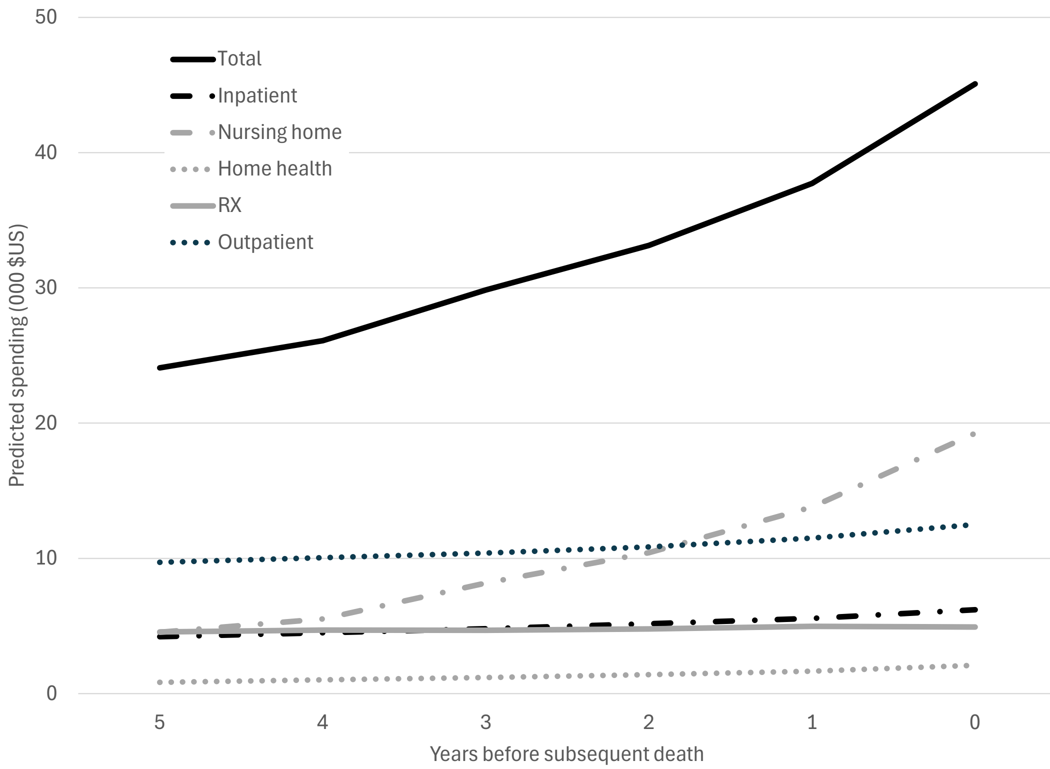

Notes: Figure displays mean predicted spending in 2012 (reported in 2024 dollars) by number of years to death, where year 0 consists of anywhere from 365 to 0 days before death. Spending is displayed for all care types (in solid black), and for individual subcategories of care. The care category labeled "Outpatient spending" includes hospice and dental spending as well as outpatient and physician spending. Predictions are made using random forests as discussed in Appendix Section B, and are calibrated using 2017 data. The sample includes 23,669 individuals aged 66+ surveyed by the MCBS in 2012.

Figure OA.10: Predicted Health-Care Spending by Morbidity Group and Number of Conditions Group

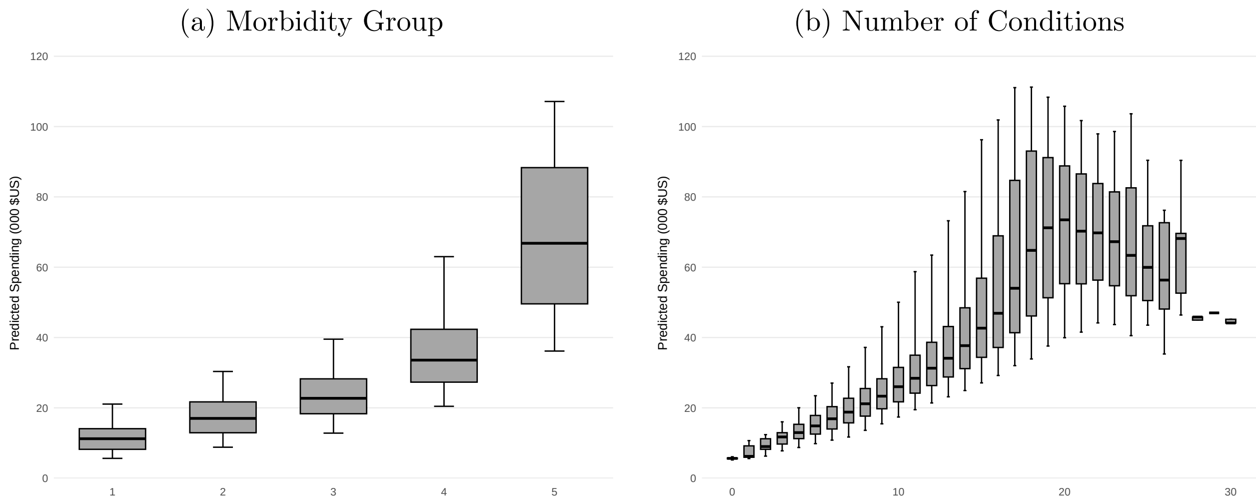

Notes: Figure displays a boxplot of predicted spending in 2024 dollars for individuals in a given morbidity group (Panel (a)) or with a given number of conditions (Panel (b)). Predicted spending is calculated using random forests trained on 2016-2018 data. For each boxplot, the line in the box denotes the median, the edges of the box denote the 25th and 75th percentiles, and the lower and upper ends of the line indicate the 5th and 95th percentiles. Standard survey weights are used. The sample includes 226,647 MCBS survey respondents between 1992 and 2019.

Figure OA.11: Lifetime Health-Care Spending at 66 for Alternative Mortality and Morbidity Combinations

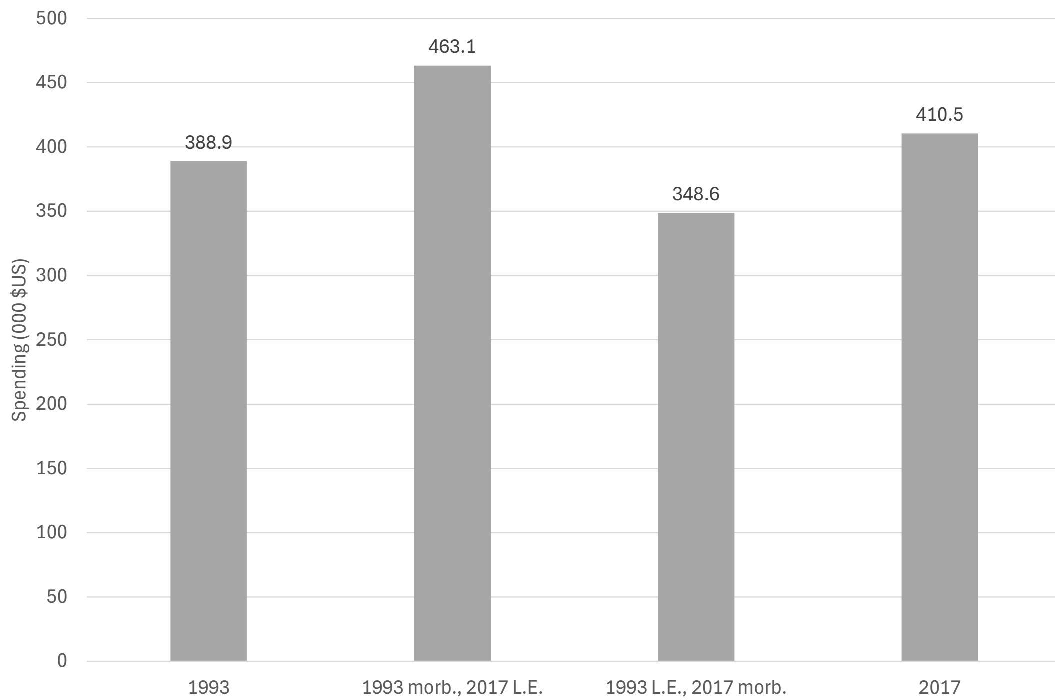

Notes: Figure displays lifetime predicted health care spending (in '000s) of the typical 66-year-old under various mortality and morbidity conditions. It shows the prior estimates of predicted health care spending (from Figure 6) using the 1993 morbidity and mortality transition matrices (first bar) and the estimates using the 2017 morbidity and mortality transitions (last bar), as well as predicted spending under two counterfactuals. One counterfactual (second bar) examines the impact on predicted spending of changes in mortality, holding morbidity constant. It uses the 1993 morbidity transition matrices but adjusts life expectancy to its 2017 level by scaling down (with an age-specific proportional factor) the exit rates to mortality from all five morbidity groups; this ensures that the morbidity-specific mortality rates at each age match the 2017 rates. We then scale down the mortality rates of each morbidity group at a given age by a common proportional factor so that the age-specific mortality rate (averaged across morbidity groups based on their 1993 age-specific shares) matches the 2017 age-specific mortality rate. The other counterfactual (third bar) examines the impact on predicted spending of changes in morbidity, holding mortality constant. It uses the 2017 morbidity transition matrices but hold life expectancy at its 1993 levels by an analogous procedure. Appendix Table OA.5 illustrates these different transition matrices for an 80 year old. All estimates are based on 100,000 simulated observations for each counterfactual.

Figure OA.12: WTA for Budget-Neutral Policies in 1993

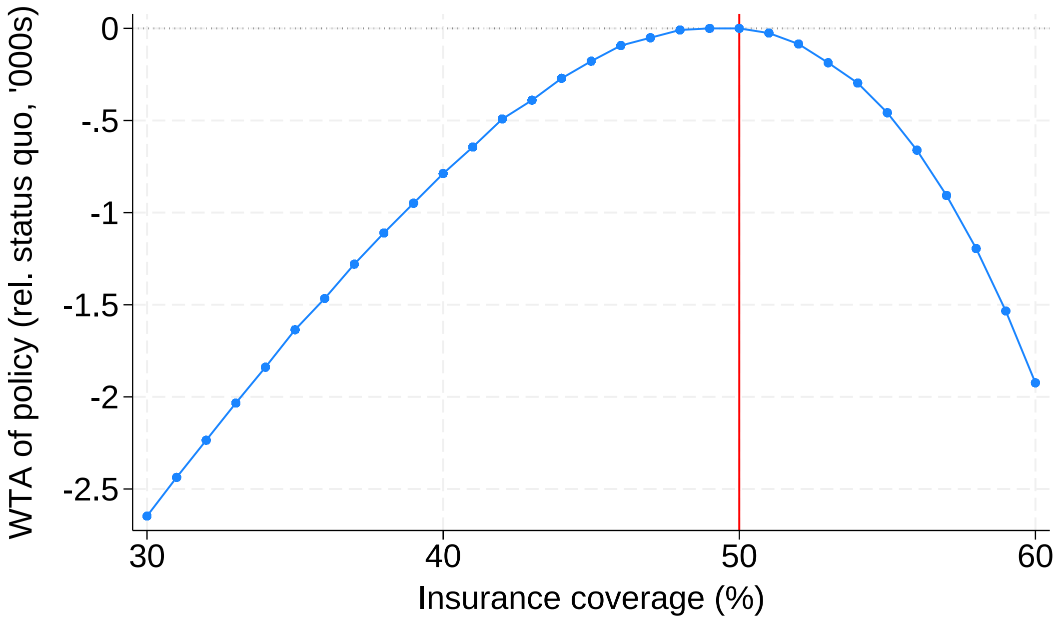

Notes: The y-axis displays the WTA for policy regimes relative to the baseline policy regime in 1993 under the assumed moral hazard parameters. The x-axis shows the proportion of medical expenses covered by the government. All policy regimes are budget-neutral, meaning that the total costs of each policy regime are equal to the total costs of the baseline policy regime in 1993. The vertical red line denotes the baseline policy regime. These values are calculated for a 66 year old in typical health with $150,000 in wealth.

Figure OA.13: Predicted Lifetime Health-Care Spending at Age 66, 1993-2017

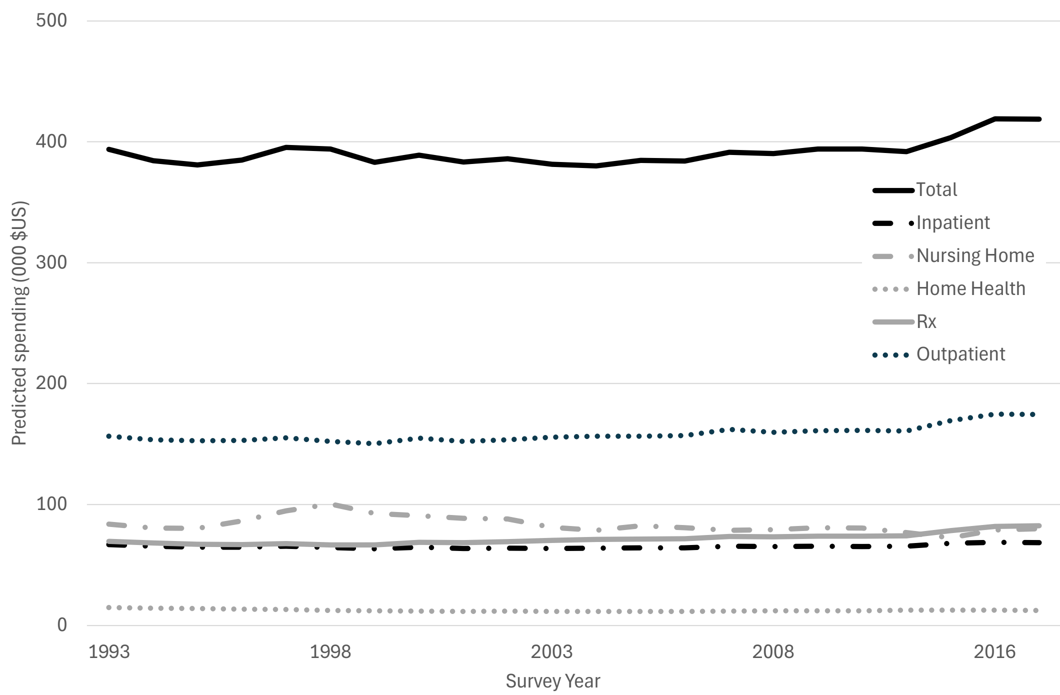

Notes: Figure displays predicted lifetime health-care spending for a 66 year old in typical health in 2024 dollars from 1993 to 2017. The starting morbidity breakdown, morbidity group transitions, and mapping from coarse morbidity groups to granular morbidity predictors are all allowed to vary by survey year. The mapping from granular morbidity to predicted spending is held fixed by imposing the mapping generated by random forests trained on 2016-2018 data. Spending is displayed for all care types (in solid black), and for individual subcategories of care. The care category labeled "Outpatient" includes hospice and dental spending as well as outpatient and physician spending. Predictions are made using random forests as discussed in Appendix Section B, and are calibrated using 2017 data. $N = 100,000$ simulated observations for each year.

Figure OA.14: Predicted Lifetime Health-Care Spending at 66: Random Forest Trained on 2017 vs. 1993 Data

(a) Random Forests Trained on 2017 Data

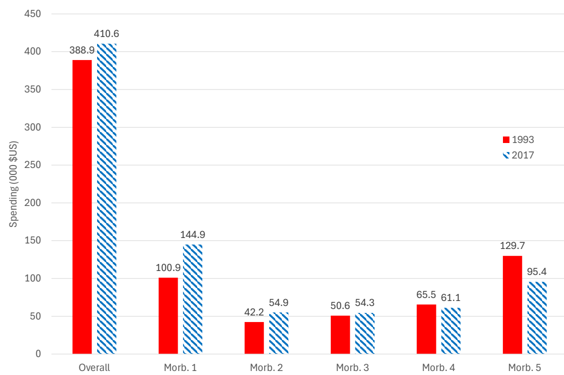

(b) Random Forests Trained on 1993 Data

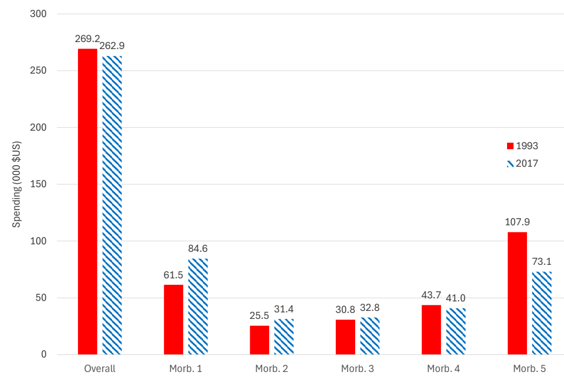

(c) Changes in Lifetime Spending, 1993-2017

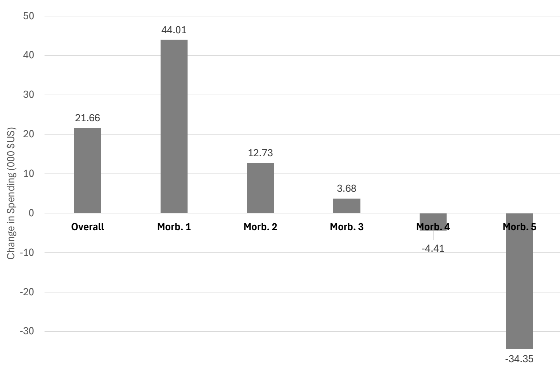

(d) Changes in Lifetime Spending, 1993-2017

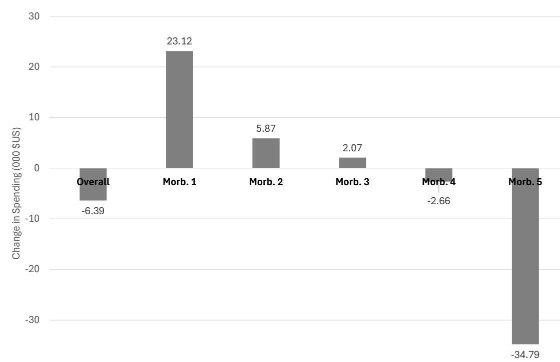

Notes: Figure displays predicted lifetime health-care spending of the typical 66-year-old, both overall and calculated as the the amount of predicted spending within each morbidity group. Panel (a) reports predictions generated using random forests trained on 2017 data, while Panel (b) reports predictions generated using random forests trained on 1993 data. In each panel, estimates are shown separately for 1993 (red) and 2017 (blue), using the morbidity and mortality transition matrices corresponding to each year. Panels (c) and (d) display the difference in lifetime spending between 1993 and 2017. N = 100,000 simulated observations for each year.

Figure OA.15: Predicted Lifetime Health-Care Spending at 66 by Spending Category: Random Forest Trained on 2017 vs. 1993 Data

(a) Random Forests Trained on 2017 Data

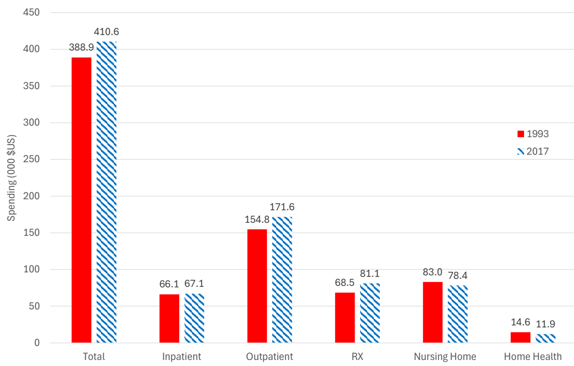

(b) Random Forests Trained on 1993 Data

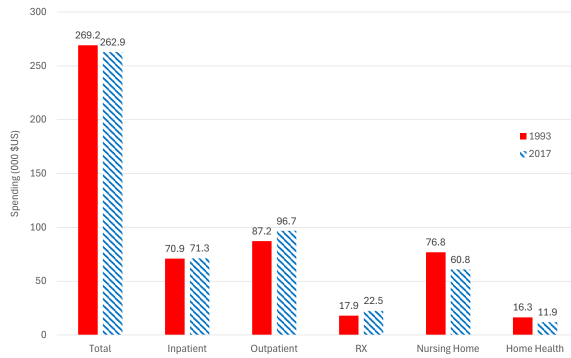

(c) Changes in Lifetime Spending, 1993-2017

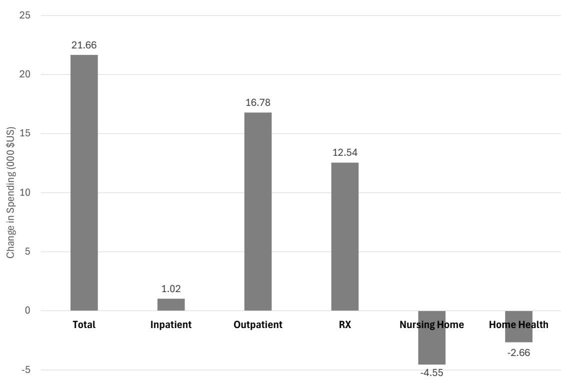

(d) Changes in Lifetime Spending, 1993-2017

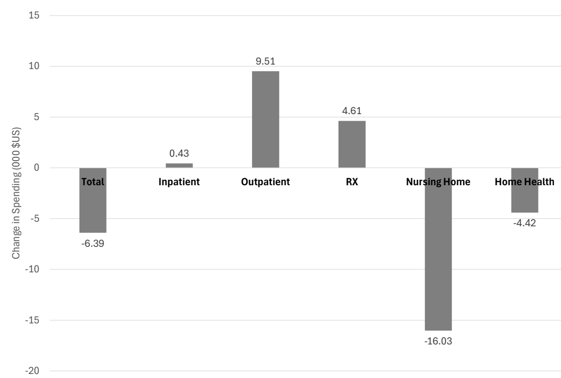

Notes: Figure displays predicted lifetime health-care spending of the typical 66-year-old, both overall and calculated as the the amount of predicted spending within each spending category. Panel (a) reports predictions generated using random forests trained on 2017 data, while Panel (b) reports predictions generated using random forests trained on 1993 data. In each panel, estimates are shown separately for 1993 (red) and 2017 (blue), using the morbidity and mortality transition matrices corresponding to each year. Panels (c) and (d) display the difference in lifetime spending between 1993 and 2017. N = 100,000 simulated observations for each year.

Table OA.1: Summary Statistics of Morbidity Measures, 2017

Condition	Prop.	Condition	Prop.
ADLs		Diseases	
(difficulty with...)		(ever diagnosed with...):	
Bathing	0.10	Cancer	0.38
Moving from bed/chair	0.11	Alzheimer's	0.06
Eating	0.03	Parkinson's	0.02
Dressing	0.08	Asthma/COPD	0.19
Walking around house	0.21	Heart attack	0.11
Using toilet	0.09	Heart disease	0.15
IADLs		Stroke	0.10
(difficulty with...)		Broken hip	0.03
Using telephone	0.06	Diabetes	0.34
Light housework	0.11	Hypertension	0.66
Heavy housework	0.28	Cataract	0.39
Making meal	0.10	Mental disability	0.01
Shopping	0.12	Osteoporosis	0.19
Managing money	0.09	Paralysis/loss of limb	0.03
Functional limitations			
(difficulty with...):			
Stooping/crouching	0.69		
Lifting 10 pounds	0.33		
Reaching arms above shoulder	0.25		
Writing	0.24		
Walking 2-3 blocks	0.41		

Notes: Table displays, for all ADLs, IADLs, functional limitations, and diseases which are used as predictors for spending in random forests, the incidence of that predictor across our full sample. Standard survey weights are used. N = 19,917 individuals aged 66+ surveyed by the MCBS between 2016 and 2018.

Table OA.2: Sample Size by Age and Morbidity Group, 2017

Age	Overall	Morbidity Group: 1	2	3	4	5
66	606	408	71	66	43	18
67	826	583	95	71	53	24
68	1,025	746	123	71	62	23
69	997	734	104	68	59	32
70	926	658	105	70	62	31
71	818	587	80	73	47	31
72	745	514	97	63	48	23
73	691	456	79	73	50	33
74	670	436	87	69	40	38
75	791	466	127	80	68	50
76	864	519	127	96	68	54
77	855	479	132	94	83	67
78	794	458	127	87	64	58
79	741	420	103	89	74	55
80	845	464	128	108	75	70
81	840	421	128	121	88	82
82	839	415	112	114	93	105
83	778	356	140	105	89	88
84	774	319	136	120	100	99
85	721	295	114	110	103	99
86	610	201	95	109	91	114
87	523	174	69	69	91	120
88	448	129	62	62	89	106
89	413	120	66	51	72	104
90	358	79	55	45	78	101
91	317	75	36	36	65	105
92	270	41	40	38	58	93
93	213	36	23	33	42	79
94	178	26	13	25	37	77
95	124	14	11	12	24	63
96	93	9	6	13	14	51
97	80	8	8	8	14	42
98	47	2	3	6	12	24
99	33	1	2	4	6	20
100+	64	7	3	4	11	39
Overall	19,917	10,656	2,707	2,263	2,073	2,218

Notes: Table displays, for each cell of age and morbidity level, the total number of MCBS observations in that cell between 2016 and 2018 (disregarding survey weights). All ages over 100 are collapsed into a single category. N = 19,917 individuals aged 66+ surveyed by the MCBS between 2016 and 2018 (referred to throughout as '2017').

Table OA.3: Sample Transition Matrices, Age 66

	Prop. Obs.	Ending group: 1	2	3	4	5	Death
		A. 1993					
Starting group:							
1	0.699	0.859	0.086	0.037	0.007	0.002	0.008
2	0.132	0.417	0.306	0.216	0.031	0.007	0.022
3	0.080	0.190	0.237	0.348	0.142	0.056	0.027
4	0.064	0.030	0.145	0.200	0.440	0.113	0.071
5	0.025	0.039	0.000	0.028	0.092	0.657	0.183
		B. 2017					
Starting group:							
1	0.724	0.899	0.067	0.026	0.005	0.000	0.003
2	0.110	0.369	0.291	0.245	0.040	0.014	0.040
3	0.078	0.148	0.270	0.445	0.095	0.000	0.042
4	0.074	0.155	0.029	0.247	0.462	0.070	0.037
5	0.014	0.000	0.026	0.000	0.209	0.688	0.078

Notes: The column of this table labeled "Prop. Obs." displays the breakdown of morbidity groups for those aged 66, separately in 1993 (Panel A) and 2017 (Panel B). Remaining columns display aggregate five-by-six transition matrices for those aged 66 in each of the three years, displaying the raw transition matrix for each age and year. Transition matrices are generated using the procedure outlined in Section 2.2. The sample includes individuals aged 66+ surveyed by the MCBS in 1992-1994 (used to estimate 1993 transition matrices; N = 35,088) or in 2016-2018 (used to estimate 2017 transition matrices; N = 27,384).

Table OA.4: Sample Transition Matrices, Age 80

	Prop.	Ending group:					
	Obs.	1	2	3	4	5	Death
		A. 1993					
Starting group:							
1	0.436	0.756	0.103	0.073	0.022	0.017	0.029
2	0.153	0.380	0.274	0.148	0.128	0.047	0.022
3	0.184	0.153	0.172	0.385	0.171	0.071	0.048
4	0.119	0.027	0.075	0.205	0.357	0.199	0.136
5	0.108	0.005	0.009	0.022	0.139	0.576	0.249
		B. 2017					
Starting group:							
1	0.564	0.807	0.103	0.046	0.012	0.005	0.027
2	0.125	0.320	0.353	0.185	0.094	0.014	0.034
3	0.170	0.104	0.248	0.294	0.238	0.048	0.068
4	0.074	0.033	0.064	0.157	0.366	0.281	0.098
5	0.068	0.000	0.000	0.051	0.162	0.588	0.200

Notes: The column of this table labeled "Prop. Obs." displays the breakdown of morbidity groups for those aged 66, separately in 1993 (Panel A) and 2017 (Panel B). Remaining columns display aggregate five-by-six transition matrices for those aged 80 in each of the three years, displaying the raw transition matrix for each age and year. Transition matrices are generated using the procedure outlined in Section 2.2. N = 46,855 individuals aged 66+ surveyed by the MCBS in 1992-1994 (used to estimate 1993 transition matrices; N = 26,413) or 2016-2018 (used to estimate 2017 transition matrices; N = 27,384).

Table OA.5: Alternative Transition Matrices, Age 80

	Ending group:						Ending group:					
	1	2	3	4	5	Death	1	2	3	4	5	Death
	A. 1993 morb. + 1993 mort.						B. 2017 morb. + 2017 mort.					
Starting group:												
1	0.756	0.103	0.073	0.022	0.017	0.029	0.807	0.103	0.046	0.012	0.005	0.027
2	0.380	0.274	0.148	0.128	0.047	0.022	0.320	0.353	0.185	0.094	0.014	0.034
3	0.153	0.172	0.385	0.171	0.071	0.048	0.104	0.248	0.294	0.238	0.048	0.068
4	0.027	0.075	0.205	0.357	0.199	0.136	0.033	0.064	0.157	0.366	0.281	0.098
5	0.005	0.009	0.022	0.139	0.576	0.249	0.000	0.000	0.051	0.162	0.588	0.200
	C. 1993 morb. + 2017 mort.						D. 1993 mort. + 2017 morb.					
Starting group:												
1	0.805	0.103	0.046	0.012	0.005	0.029	0.758	0.103	0.073	0.022	0.017	0.027
2	0.324	0.357	0.188	0.095	0.014	0.022	0.375	0.271	0.146	0.127	0.047	0.034
3	0.106	0.253	0.300	0.243	0.049	0.048	0.150	0.168	0.378	0.167	0.069	0.068
4	0.032	0.061	0.151	0.350	0.269	0.136	0.028	0.078	0.214	0.373	0.208	0.098
5	0.000	0.000	0.048	0.152	0.552	0.249	0.005	0.010	0.024	0.148	0.613	0.200

Notes: This table displays the "factual" transition matrices in 1993 (Panel A) and 2017 (Panel B), each for 80-year-olds. Panel C then displays a counterfactual transition matrix which combines the mortality rates from 2017 with the morbidity rates from 1993 (taking the 1993 matrix as a base, and reallocating observations to death proportionally from each ending morbidity level to reach the 2017 death rate at each starting morbidity level). Panel D then displays a counterfactual transition matrix which combines the mortality rates from 1993 with the morbidity rates from 2017 (taking the 2017 matrix as a base, and reallocating observations to death proportionally from each ending morbidity level to reach the 1993 death rate at each starting morbidity level). The sample includes individuals aged 66+ surveyed by the MCBS in 1992-1994 (used to estimate 1993 transition matrices; N = 26,413) or 2016-2018 (used to estimate 2017 transition matrices; N = 19,917).