

Valuing Statistical Life Using Seniors' Medical Spending*

Jonathan D. Ketcham Nicolai V. Kuminoff Nirman Saha[†]

October 2, 2021

Abstract

This study provides the first revealed preference evidence on the value of statistical life (VSL) for US seniors aged 67–97 from the rates at which they choose to consume medical care relative to other private goods, and by the effects of their choices on their survival probabilities. These effects are estimated from individuals' survey responses linked with their Medicare records. Instrumental variables estimators provide robust evidence that the mean VSL is below \$1 million and that it decreases with age, and, given age, increases with income, education, and health and is higher for women and people who never smoked.

JEL classification: D90, J14, J17, Q51

Keywords: Value of statistical life, VSLY, benefit cost analysis, environmental valuation, returns to medical spending

*We are grateful for insights and suggestions from Joe Aldy, Kelly Bishop, Glenn Blomquist, Mary Evans, Chad Jones, Alan Krupnick, David McCarthy, Alvin Murphy, Julian Reif, Dan Silverman, Kerry Smith, Jonathan Skinner, David Slusky, Greg Veramendi, Jeffrey Wooldridge, and seminar audiences at Camp Resources, the World Congress of Environmental and Resource Economists, the AERE Summer meeting, the ASHEcon conference, the WEAI summer conference, Arizona State University, Hong Kong Polytechnic University, Michigan State University, Stanford University, University of Birmingham, University of British Columbia, University of Illinois Urbana-Champaign, University of Kansas, and University of Southern California.

[†] Jonathan D. Ketcham : Arizona State University, Dept. of Economics, Tempe, AZ 85287 (e-mail: ketcham@asu.edu). Kuminoff: Arizona State University, Dept. of Economics, Tempe, AZ 85287 and National Bureau of Economic Research (email: kuminoff@asu.edu). Saha: University of Surrey, Dept. of Economics, Guildford, Surrey, GU2 7XH (email: n.saha@surrey.ac.uk).

1 Introduction

Mortality rates are affected by many government activities such as regulating air pollution, setting speed limits, and funding health programs. Evaluating the equity and efficiency of these activities requires policymakers to weigh their benefits, including mortality reductions, against their costs. US federal agencies are required to evaluate many of these activities by reporting the expected monetary costs and benefits of every major regulation they propose. The standard approach to monetizing changes in mortality rates due to a regulation is to multiply the expected change in the number of premature deaths by a constant value per statistical life (VSL). The monetized mortality effects often dominate cost–benefit analysis. In fact, [Lee and Taylor \(2019\)](#) report that survival gains represent up to 70% of all monetary benefits calculated for all federal regulations.

VSL measures are typically derived from econometric estimates of the compensating differentials paid to workers to induce them to perform jobs with higher risks of accidental death ([Cropper, Hammitt and Robinson \(2011\)](#)). Workers whose choices generate this evidence are almost entirely under age 65, but people over 65 often account for a large share of policies’ survival benefits. For example, senior citizens represent 75% of annual premature deaths avoided by regulating air pollution ([U.S. Environmental Protection Agency \(2011\)](#)). This discrepancy between the age group used to calculate VSL and the age group to whom VSL is applied may yield substantial mis-measurement of the benefits of mortality reductions. Economic theory predicts that the VSL will evolve over the life cycle with changes in health, wealth, risk aversion and remaining life expectancy (e.g. [Arthur \(1981\)](#), [Rosen \(1988\)](#), [Evans and Smith \(2006\)](#), [Murphy and Topel \(2006\)](#), [Bauser, Lakdawalla and Reif \(2018\)](#)). However, the net effect of how these changes evolve with age is largely an empirical question, and researchers have not provided any revealed preference evidence on this evolution beyond age 65.¹ As a result, the sign and magnitude of mis-measurement under the status quo are unknown.

This paper is the first study to develop revealed preference evidence on how much Americans over age 65 are willing to pay to reduce their own mortality risks. Our evidence comes from reconciled Medicare records linked to survey data on the rates at which people choose to consume medical care relative to other private goods. We view their choices through the lens of a life–cycle model in which people on fixed incomes make repeated decisions about how much to spend on medical care while facing uncertainty about their future health and survival.

¹Prior research on age—related variation in the VSL has either employed stated-preference designs (e.g. [Krupnick \(2007\)](#), [Blomquist, Dickie and O’Conor \(2011\)](#)) or stratified hedonic wage regressions by age bins (e.g. [Smith et al. \(2004\)](#), [Evans and Smith \(2006\)](#), [Viscusi and Aldy \(2007\)](#), [Aldy and Viscusi \(2008\)](#), [Evans and Schaur \(2010\)](#)). Both approaches have yielded mixed results with no consensus for predicting how the VSL evolves beyond age 65.

This setup is similar to the representative agent model in [Hall and Jones \(2007\)](#) but focuses on individual decisions made by people with heterogeneous health, wealth, and preferences. Their optimal choices will equate their marginal cost of medical care (conditional on insurance coverage) with its marginal benefits as determined by the discounted expected utility of future life, where this utility depends on medical care's expected effects on the quantity and quality of remaining life. This equality yields a key insight: the marginal effect of medical expenditures on the probability of survival can reveal how much people are willing to pay for marginal increases in their survival probability. Measures of individual willingness to pay can then be aggregated to calculate the VSL for groups of people who differ by age, health, income, and other characteristics.

We estimate VSL measures for a nationally representative random sample of about 22,000 people aged 67 to 97 who participated in the Medicare Current Beneficiary Survey (MCBS) from 2005–2011. The MCBS provides up to three years of reconciled Medicare records on each person's total and out-of-pocket (OOP) medical spending. These are the most comprehensive and accurate data on OOP spending for US seniors. They track all medical expenditures processed by Medicare, Medicaid, Medigap, employer-sponsored plans and other insurance plans, as well as any expenditures paid entirely OOP. The MCBS also reports each participant's insurance coverage, smoking history, income, education, employment status, knowledge of Medicare programs, utilization of assistance in making medical decisions, and their self-reported health and limitations in activities of daily living. We further link the MCBS to Medicare administrative records for the surveyed individuals, allowing us to additionally observe each person's demographics, residential location, medically diagnosed illnesses, and death date. Then we use the linked data to estimate survival functions.

Our survival functions measure how an individual's medical spending in the current year affects her probability of surviving through the next year. A key identification challenge is that medical spending is likely to be positively correlated with latent morbidity. This may bias the estimated return to spending toward zero if people who are sicker in unobserved ways tend to spend more on health care and die sooner. We overcome this challenge by adapting the approach of [Finkelstein, Gentzkow and Williams \(2016\)](#) to derive an instrument for medical spending from geographic variation in the supply of medical care. Intuitively, some of the variation in individuals' medical expenditures arises from similar individuals facing sets of treatment options that differ in costs due to differences in health care supply across markets. We construct an instrument for individuals' medical spending using this supply-side variation by using a separate, larger dataset describing within-person changes in annual Medicare expenditures for just under half a million people who moved between Dartmouth Atlas hospital referral

regions. When we use this index to instrument for medical expenditures, our main specification of the survival function implies that an additional \$1,000 in spending reduces mortality in the following year by about 0.4 percentage points on average. This average marginal effect varies from about 0.2 to 2 percentage points across groups of people who differ in their health, demographic, and socioeconomic characteristics. This range is consistent with the range of local average treatment effects found in prior studies, as is our finding that the returns to spending increase with illness and age.

Combining our main estimates for the return to medical spending with each person's observed coinsurance rate yields a mean VSL of about \$402,000 (year 2010 dollars) at age 67. Adjusting this value to account for estimated life expectancy and assuming a 7% discount rate implies a value per statistical life year (VSLY) of about \$39,000. Our estimates for age-specific VSL and VSLY measures decline as near-monotonic functions of age. We take a systematic and comprehensive approach to evaluating the sensitivity of these main findings to our research design, following [Leamer \(1983\)](#), [Banzhaf and Smith \(2007\)](#), and [Greenstone, Kopits and Wolverton \(2013\)](#). First, we define alternative analytic decisions along five dimensions: (i) sample criteria, (ii) source of data on medical expenditures, (iii) choice of instrument for medical expenditures, (iv) parametric form of the survival function, and (v) choice of covariates and spatial fixed effects. Then we estimate VSL measures for every possible combination of these decisions, yielding 200 sets of estimates. All of them produce mean VSLs below \$1 million, and for all of them the VSL declines with age. Further, they all produce mean VSLYs below \$100,000 at each age from 67 to 97 whether we assume a 3% or 7% discount rate.

Next, we investigate heterogeneity. At age 67 the VSL is higher for women compared to men, people who never smoked compared to those who have, and for people with more income and education. In each of these comparisons, people with higher VSL tend to be healthier. These patterns are also evident when we stratify the VSL by subjective measures of health, objective measures of health, and limitations in activities of daily living. For example, 67 year old people who describe their health as "excellent" for their age on a Likert scale have an average VSL of \$843,000, which is more than double the average VSL among all 67 year olds (\$402,000), and more than twenty times the average for 67 year olds who describe their health as "poor" (\$36,000). These group-wise differences appear to be due to differences in both quantity and quality of remaining life. As age increases, the VSL ranking across groups persists but the differences between their levels decline, consistent with declining differences in remaining life expectancy conditional on survival.

While the patterns of conditional heterogeneity in our VSL estimates can be rationalized by a life-cycle model,

the levels of our estimates fall an order of magnitude below the range commonly used to monetize mortality reductions for seniors. Federal agencies typically assume a constant VSL between \$6 and \$10 million (year 2010 dollars) for every avoided death, regardless of age and health.² The \$6 to \$10 million range is consistent with evidence on average VSL from hedonic wage regressions of workers aged 18 to 65 (e.g. [Costa and Kahn \(2004\)](#), [Cropper, Hammitt and Robinson \(2011\)](#), [Deleire, Kahn and Timmins \(2013\)](#), [Kneisner et al. \(2012\)](#), and [Lee and Taylor \(2019\)](#)).


Figure 1 summarizes how our estimates compare to the most closely related prior studies. The diamonds show our age-specific estimates for the mean VSL among relatively healthy people who report no major illnesses or functional limitations. Our findings are closest to a VSL-age function that [Deschenes, Greenstone and Shapiro \(2017\)](#) derived by calibrating the life-cycle model from [Murphy and Topel \(2006\)](#) to match a VSL measure constructed from evidence on speed limit changes ([Ashenfelter and Greenstone \(2004\)](#)). For example, Figure 1 shows that their calibrated value of \$640,000 at age 67 is similar to our revealed preference estimate of \$700,000; their calibrated value of \$280,000 at age 75 is near our estimate of \$240,000.

The solid lines in Figure 1 show projections made by regressing our age-specific estimates on third- and fourth-order polynomial functions of age and then projecting the fitted values back to age 40. This backward extrapolation of our estimates, while speculative, yields a predicted VSL range from \$6 to \$10 million for people in their early 40s. This range is consistent with evidence from well-identified hedonic wage studies ([Kneisner et al. \(2012\)](#), and [Lee and Taylor \(2019\)](#)). Further, the decline from age 40 to age 62 tracks with the findings of [Aldy and Viscusi \(2008\)](#). While additional research on medical consumption decisions among younger workers would be particularly insightful, these results suggest that the differences between our VSL estimates for US seniors and labor market evidence on younger healthier workers may be due to declines in health, functionality, and life expectancy that occur as people age rather than by methodological differences.

Our approach shares methodological features with the wage-hedonic method while differing in some important ways. Both approaches assume that people make informed choices for how to continuously trade consumption for mortality risk. The estimates may diverge if the conventional revealed preference assumptions common to both settings are less applicable to either labor market decisions or medical care choices. On one hand, our

²The one-size-fits-all approach to monetizing life extension in federal cost-benefit analyses is based on the US Office of Management and Budget's ([U.S. Office of Management and Budget \(2003\)](#)) judgment that there is insufficient evidence to guide age-based adjustments to VSL. They state that: "The age of the affected population has also been identified as an important factor in the theoretical literature. However, the empirical evidence on age and VSL is mixed. In light of the continuing questions over the effect of age on VSL estimates, you should not use an age adjustment factor in an analysis using VSL estimates." Political sensitivity to age adjustments came to light after the EPA proposed to reduce the VSL for seniors by 37% when calculating benefits of The Clear Skies Act, which became known as the "senior death discount" and was ultimately abandoned following controversy and opposition from interest groups ([Viscusi and Aldy \(2007\)](#)).

Figure 1: Comparing our Estimates to Prior Literature



Note: The figure reports measures for the VSL in year 2010 dollars by age. The dashed line at \$8 million is a benchmark one-size-fits-all VSL used by the US EPA and other agencies based on a review of academic studies comprised mainly of wage-hedonic regressions of workers age 18 to 65. The circle and square mark point estimates from wage-hedonic studies by [Kneisner et al. \(2012\)](#) and [Lee and Taylor \(2019\)](#) that take steps to mitigate threats to identification. The dotted line shows calibrated values from [Deschenes, Greenstone and Shapiro \(2017\)](#). The triangles show estimates from [Aldy and Viscusi \(2008\)](#) for workers aged 35-44, 45-54, and 55-62. They note that fitting a third-order polynomial to their estimates implies a VSL close to \$2 million at age 62. Finally, the diamonds show how our study fills the gap in knowledge about the VSL's evolution beyond age 65. Each diamond is an age-specific mean VSL calculated from our main IV survival function for people who report no functional limitations or chronic illnesses. We focus on this relatively healthy subsample to enable comparability with younger workers.

medical care setting provides stronger incentives for people to make careful choices and greater means to do so. Mortality risks are higher than in labor markets, information about risk is more accessible, and medical professionals are tasked with helping patients make informed decisions. Further, in contrast with a worker's job opportunities, seniors face essentially a continuum of options for the intensity and cost of medical care. In particular, patients can choose their medical care providers, including which physicians, specialties and hospitals as well as their treatment plans such as intensity and frequency of testing, use of newer versus older technologies, and medical versus surgical interventions.

On the other hand, health insurance is complex and everyone may not fully understand their treatment options and billing procedures. To explore the sensitivity of our estimates to maintaining revealed preference assumptions

for everyone in our sample, we investigate how VSL estimates vary with seniors' medical decision-making processes and knowledge. We find that conditional on age there is virtually no difference in VSL measures between people who usually make their own health insurance decisions, people who get help making decisions from family, and people who rely on others to make decisions for them. In addition, we implement the strategy suggested in [Bernheim and Rangel \(2009\)](#) by using ancillary data on each person's cognitive functioning, decision-making, and knowledge of Medicare institutions to divide people into groups for whom revealed preference assumptions are more or less likely to hold. VSL measures are lower on average among the group for whom ancillary data provide reasons to suspect that choices may not reveal preferences, but the differences are too small to reconcile our estimates with those from the wage-hedonic literature. Further analysis indicates that these differences cannot be reconciled by physicians overtreating patients, or by patients ignoring the fact that insurance covers substantial shares of their medical care costs.

Our results have two broad implications for evaluating the efficiency and equity of a wide range of activities that reduce mortality including environmental regulations, safety regulations, health insurance programs such as Medicare and Medicaid, and medical technology. First, because our VSL and VSLY estimates for seniors are far below the wage-hedonic estimates for younger healthier workers that have traditionally been used to monetize mortality reductions among seniors, using our estimates would reduce the monetized benefits of policies to varying degrees depending on the age and health of the beneficiaries. Second, our estimates imply that activities that improve health will increase the VSL and VSLY due to dynamic complementarity between the quantity and quality of life. The net effect of these two countervailing implications for estimated cost-benefit ratios likely varies across applications.

2 Data: linking survey responses to administrative records

We link panel data on Medicare Current Beneficiary Survey (MCBS) participants to administrative records on the same individuals from the US Centers for Medicare and Medicaid Services (CMS). The MCBS is a nationally representative rotating panel survey that is administered to approximately 16,000 randomly chosen Medicare beneficiaries each year. Each respondent is interviewed for up to four consecutive years even if they change addresses or move to long-term care facilities, and if they become cognitively impaired then someone else responds as their proxy. The linked data provide a nationally representative sample of the 65+ population because all Americans

become eligible for Medicare benefits at age 65.³

Importantly, the MCBS provides comprehensive measures of each respondent's total and out-of-pocket medical spending. CMS develops these measures by combining federal administrative records on the respondent's Medicare claims with the respondent's financial records on expenditures that were not processed through Medicare. However, due to the time needed to collect and reconcile these measures, they are only available for the second year of the survey onward. The MCBS also provides detailed information on each person's socioeconomic characteristics, household composition, labor market participation and self-assessed health. This complements the information available in CMS administrative records on each person's demographics, diagnoses of medical conditions, residential address and timing of moves, and death dates.

2.1 Sample construction

We link MCBS interview data from 2005 to 2011 for respondents over age 65 to data extracted from each person's CMS administrative files from 2005 through 2012. The linked data contain 51,191 person-years with annual spending data for people who survived to the end of the calendar year. The minimum age is 67—the youngest age at which we observe MCBS respondents in their second full calendar year of survey participation. Then we make two sample cuts. First, we drop 730 person-years in which respondents declined to answer questions about their socioeconomic status or health, or their reported medical spending was zero, or their reported medical spending exceeded \$100,000.⁴ Second, we drop 5,764 person-years where the respondent was employed at the time of their MCBS interview. Dropping workers simplifies our analysis by allowing us to avoid modeling how current medical spending may affect future income through intermediate health shocks that could, in principle, affect labor productivity and the timing of retirement (Grossman (1972)). However, section 6.1.1 shows that adding these workers to the estimation sample does not meaningfully change the magnitude of our VSL estimates relative to the status quo estimates.

Our main sample is comprised of 22,206 people whom we observe for 44,697 person-years. Individuals are observed for one, two or three years. We cannot observe three years of spending for everyone because some people die while enrolled in the MCBS and others' MCBS enrollment cycles extend beyond the endpoints of our study

³ The linked data do not allow us to obtain a nationally representative sample of people under 65 because their Medicare eligibility stems from illness or poverty rather than age.

⁴ These data cuts retain 99% of our study population. Dropping the extreme tails of the expenditure distribution reduces the scope for outliers to affect our estimates. It also parallels labor market studies of the VSL such as Kneisner et al. (2012) which drops workers with real hourly wages below \$2 per hour or above \$100 per hour.

period. Finally, we use administrative data on death dates to observe one-year mortality for everyone in the sample, including those who exit the MCBS during our study period. Table 1 reports summary statistics. The average person is 78 years old and 5% die during the year after we observe their medical spending. The distribution of people by sex, race and educational attainment matches 2010 Census data on the US population age 65+.⁵ We also see that about half are married and 93% have living children.

Table 1: Summary Statistics

measure	summary statistic	data source
1-year mortality (%)	5	admin
mean age	78	admin
female (%)	58	admin
white, not-Hispanic (%)	85	admin
African American (%)	8	admin
Hispanic (%)	5	admin
education: less than high school (%)	26	MCBS
education: high school degree (%)	30	MCBS
education: some college (%)	22	MCBS
education: college degree (%)	21	MCBS
married (%)	52	MCBS
has living children (%)	93	MCBS
Gross annual medical spending (\$2010)	11,489	MCBS-admin
Out-of-pocket annual medical spending (\$2010)	1,817	MCBS-admin
ever smoked (%)	58	MCBS
underweight BMI (%)	4	MCBS
number of chronic conditions (out of 61)	7	admin
mean HCC score	-0.27	admin
self reported health = "poor" (%)	5	MCBS
self reported health = "fair" (%)	16	MCBS
self reported health = "good" (%)	33	MCBS
self reported health = "very good" (%)	31	MCBS
self reported health = "excellent" (%)	15	MCBS
one or more limitations on instrumental activities of daily living (%)	28	MCBS
one or more limitations on activities of daily living (%)	30	MCBS
number of people	22,206	
number of person years	44,697	

Note: Spending measures are adjusted to year 2010 US dollars using the CPI. Variables with the "MCBS" label are based on survey responses. Variables with the "admin" label are drawn from CMS administrative files. The spending variables are labeled "MCBS-admin" because they combine information from administrative files and MCBS-based tracking of respondents' medical and financial records.

⁵ American Community Survey data for 2010 identify 85% of the US population age 65+ as white, 57% as female, and 21% as having a bachelor's degree or higher.

2.2 Medical expenditures

The US Medicare program provides universal health insurance for Americans over age 65. Enrollees can choose between traditional “fee-for-service” Medicare that pays medical care providers a fixed fee for each service they perform and Medicare Advantage plans that charge a monthly premium in exchange for lower out-of-pocket (OOP) costs for certain services. Some people have additional health insurance provided by their former employers or spouses’ employers, and some people purchase private Medigap insurance plans to supplement their public Medicare coverage. The MCBS spending measure includes all of these public and private forms of coverage as well as expenditures paid entirely OOP.

The MCBS reports comprehensive measures of each respondent’s total and OOP medical spending during their second, third and fourth years of survey participation. These data are considered the best available for measures of OOP spending among the US Medicare population and they include costs for services not covered by Medicare. They account for all payments by third-party payers, including Medicaid, Medigap, or employer-sponsored insurance, which may cover some or all of the typical patient cost-sharing under Medicare. The data are collected from respondents who record medical events in calendars and keep documentation and receipts, e.g. from insurers, pharmacies, and Medicare explanations of benefits. CMS then reconciles these records with its administrative data on insurance claims. The resulting spending measures are more comprehensive than Medicare claims because they also include expenditures that were not processed through the Medicare system or not retained in CMS’s administrative files during our study period. Examples include prescription drug expenditures made before Medicare started subsidizing drugs in 2006, spending in Medicare Advantage and Medigap plans, and expenditures paid entirely OOP with no claim submitted, e.g. some generic drugs. Equally important is the fact that the reconciled spending measures provide a detailed accounting of how expenditures were divided across payees including the federal government, employer-sponsored plans, private insurers, and the beneficiary. This accounting allows us to observe the fraction of each MCBS respondent’s total annual medical expenditures that were paid OOP, i.e. their effective annual coinsurance rate.⁶

⁶ CMS’s official description of these files states: “The MCBS Cost and Use files link Medicare claims to survey-reported events and provides complete expenditure and source of payment data on all medical care services, including those not covered by Medicare. Expenditure data were developed through a reconciliation process that combines information from survey respondents and Medicare administrative files. The process produces a comprehensive picture of health services received, amounts paid, and sources of payment. The file can support a broader range of research and policy analyses on the Medicare population than would be possible using either survey data or administrative claims data alone. Survey-reported data include information on the use and cost of all types of medical services, as well as information on supplementary health insurance, living arrangements, income, health status, and physical functioning. Medicare claims data includes use and cost information on inpatient hospitalizations, outpatient hospital care, physician services, home medical care, durable medical equipment, skilled nursing home services, hospice care, and other medical services.”

Table 1 shows that the average person spent \$11,489 on medical care annually.⁷ OOP expenditures on medical services were on average \$1,817, which is equivalent to 7% of per capita income for the over-65 population in 2010 (US Current Population Survey, 2011).

2.3 Health

The lower part of Table 1 reports means for several measures of health. First, we track whether people face a statistically higher mortality risk because they have a history of smoking (58%) or were underweight based on their body mass index at the time of the survey (4%). Second, we use CMS Chronic Conditions Warehouse files to identify whether and when each person was first diagnosed with chronic illnesses based on insurance claims.⁸ The average person is diagnosed with 7 illnesses (out of 61). Third, we use data on CMS's hierarchical conditions categories (HCC) risk-adjustment score. HCC scores synthesize data on diagnosed illnesses, age, gender, and initial reason for Medicare eligibility into a normalized index of health risk that CMS uses to make capitation payments to Medicare Advantage plans.⁹

We augment the objective measures of health with subjective measures recorded in the MCBS. Respondents are asked, "In general, *compared to other people your age*, would you say that your health is ... excellent, very good, good, fair, or poor?" (emphasis added). Table 1 shows that the distribution of self-reported health is slightly left-skewed with 79% of people reporting that their health is good, very good, or excellent. We also track whether morbidity interferes with respondents' daily lives. The MCBS reports whether people say they are capable of performing various activities of daily living (ADL). Approximately 28% of respondents have difficulty performing at least one "instrumental" ADL, which includes activities that affect the ability to live independently such as managing money, doing household work, using the telephone and preparing meals. Approximately 30% of respondents report difficulty in performing one or more "basic" ADLs such as bathing, dressing, eating, walking, and using

⁷ This statistic is for 12 months of spending. To measure per capita expenditures consistently we exclude the calendar years in which people die. The median death occurs in early July.

⁸ The set of chronic conditions includes: acute myocardial infarction, ADHD and other conduct disorders, anemia, anxiety, asthma, atrial fibrillation, bipolar disorder, brain injury, cancer (breast, colorectal, prostate, lung, endometrial), cataract, cerebral palsy, chronic kidney disease, chronic obstructive pulmonary disease, congestive heart failure, dementia, depression, diabetes, epilepsy, fibromyalgia, glaucoma, hearing impairment, hip fracture, HIV, hyperlipidemia, hypertension, hypothyroidism, heart disease, intellectual disabilities, learning disabilities, leukemia, liver disease, mild cognitive impairment, migraine, mobility impairment, multiple sclerosis, muscular dystrophy, other development delays, personality disorders, post-traumatic stress disorder, obesity, osteoporosis, peripheral vascular disease, rheumatoid arthritis, schizophrenia, spina bifida and other congenital anomalies of the nervous system, spinal cord injury, stroke, tobacco disorder, ulcers, visual impairment, viral hepatitis.

⁹ Background information on CMS's HCC model can be found at <http://www.nber.org/data/cms-risk-adjustment.html>. We follow Finkelstein, Gentzkow and Williams (2016) in adjusting raw HCC scores for spatial and temporal trends. This adjustment is described in Appendix A.1.

the bathroom. These subjective variables may help to capture latent heterogeneity in health not captured by the objective measures. For instance, people who have difficulty performing tasks of daily living because of mobility limitations may also be more likely to suffer from more severe and debilitating symptoms of heart disease than other people with heart disease.

2.4 The evolution of health and medical spending

Figure 2 illustrates how health declines and medical spending increases with age. The figure documents the evolution of health and spending over MCBS years 2 through 4 for the subset of people in Table 1 whom we observe for all three years. As the average respondent ages from 77 to 79, they are more likely to be diagnosed with chronic conditions. For example, panel A shows that the share of people diagnosed with hypertension increases from 70% to 74%, the share diagnosed with ischemic heart disease increases from 42% to 47%, and the share diagnosed with Alzheimer's disease and related dementias increases from 6% to 10%. Panel B shows that the average person is diagnosed with a total of 6.3 chronic illnesses in year 2 and that this increases to 7.3 by year 4. Panel C shows that the average HCC morbidity score increases with the average number of chronic illnesses.


As people get older and sicker, Figure 2 shows that they are more likely to experience restrictions on instrumental and basic activities of daily living (panel D). Yet self-reported health status is relatively stable (panel E). This is consistent with the fact that the question is asked relative to others of the same age. Finally, panel F shows that per capita medical spending increases by 5% to 6% per year. While the reconciled MCBS measures of total medical spending that we rely on are larger than spending measures constructed from Medicare claims alone, their trends are nearly parallel.

3 A Dynamic Model of Medical Expenditures, Health, and Survival

We use a dynamic model to explain how retirees choose to adjust their medical spending as they experience health shocks that affect their expected future quantity and quality of life. People enter the model at age 65 with endowments of health and wealth.¹⁰ Each year they determine how much to spend on non-medical consumption and medical services which, in turn, affect their future health and wealth.

¹⁰ The median retirement age in the United States is 62. Individuals born before 1955 received full retirement benefits from the Social Security Administration if they retired at age 66. Among all individuals age 66 and over in the Medicare Current Beneficiary Survey, approximately 13 percent were working in 2010.

Figure 2: Evolution of Health and Medical Spending Over MCBS Years 2 to 4



Note: The figure summarizes the evolution of health and medical spending during years two through four of the Medicare Current Beneficiary Survey, the period for which we observe comprehensive spending measures. The figure is constructed from data on the subset of respondents whom we observe in all three survey years.

People face two sources of uncertainty when they make decisions. First, their health evolves through a partially stochastic process. The stock of health declines with age, on average, but the decline can be slowed or temporarily reversed by investing in medical care. Second, people can die at any time. Survival to the next period is modeled as a probabilistic function of age, health, and medical expenditures. Hence, people can increase the expected quantity and the quality of their lives by purchasing medical services to slow the degradation of their health stock and reduce their short-term probability of death. When people decide how much to invest in health they face an intertemporal tradeoff. Increasing medical expenditures decreases their current quality of life by reducing non-medical consumption, but it increases their expected future quantity and quality of life through the health stock and survival probability. Under standard assumptions, the way that people respond to this tradeoff reveals their willingness to pay for marginal changes in probabilistic life extension.

Many decisions about individual medical care may be made at the household level, or even without the individual's input in the case of people suffering from dementia and other cognitive impairments. Here we abstract from the complications of within-household bargaining and make no distinction between the decision-maker and the individual receiving care. Our model also embeds versions of the “continuous choice” and “full information” assumptions that are ubiquitous in the revealed preference literature on VSL estimation. Specifically, we assume that people are free to purchase medical services in continuous quantities and that they do so knowing how those purchases will affect their probability of surviving through the end of the following period. Section 8 takes a step toward relaxing these standard but strong revealed preference assumptions by stratifying the VSL measures based on whether people make their own medical care decisions and consider themselves to be well informed.

3.1 Preferences and health

In each time period, t , retiree i 's utility depends on her amount of non-medical consumption, c_{it} , and health, h_{it} :

$$U_{it} = u(c_{it}, h_{it}). \quad (1)$$

Health evolves over time as a function of medical expenditures, m_{it} . The retiree's stock of health in period $t + 1$ depends on her period t health stock, medical expenditures, age, and a random shock denoted by ϵ_{it} . In a slight abuse of notation we use t to index both age and time period so that evolution of the health stock can be

represented as

$$h_{it+1} = f(h_{it}, m_{it}, t, \epsilon_{it}). \quad (2)$$

Equation (2) captures the essence of a [Grossman \(1972\)](#) style health production function. Each retiree inherits a stock of health upon entering the model at age 65. Their initial stock may reflect genetic endowments and the cumulative effects of past medical consumption, lifestyle choices, pollution exposures, and health shocks. Health depreciates following negative shocks, but such declines can be partially offset by medical expenditures. For instance, health shocks may be caused by the arrival or worsening of conditions such as diabetes, heart disease, and cancer. Their negative effects on future health may be moderated by seeing doctors to obtain medications, surgeries, preventative care, or guidance on lifestyle choices such as diet and exercise. However, the return on medical care investments may decline with age.

We model death using a distinct probabilistic function. Let s_{it} represent person i 's probability of survival to period $t + 1$. Survival is assumed to be a deterministic function of period t medical expenditures, health, age, and a random shock, μ_{it} . Integrating over the shock yields the survival probability.

$$s_{it} = g(m_{it}, h_{it}, t). \quad (3)$$

Together, equations (2) and (3) illustrate how medical expenditures may increase both the quantity and quality of life in future periods. Investing in medical care may increase the future quality of life through the health stock, and increasing the health stock may lower the short-term probability of death. Medical care investments may also increase the survival probability directly without affecting the health stock. Examples include medications and surgeries that reduce the chances of a fatal heart attack or stroke. Overall, the dynamic, stochastic nature of the return to investment in medical care presents budget-constrained retirees with a tradeoff between increasing current utility through non-medical consumption and investing in future utility through medical consumption.

3.2 Intertemporal budget constraint

For simplicity, we abstract from credit markets and require people to maintain non-negative assets each period. Equation (4) shows the intertemporal budget constraint.

$$a_{it+1} = (1 + r)a_{it} + y_i - c_{it} - \gamma_{it}m_{it} \geq 0 \quad \forall t. \quad (4)$$

The retiree's total assets in period $t+1$ are equal to the assets retained from the prior period, a_{it} , which are assumed to grow at a risk-free interest rate of r , plus non-asset income from all other sources, y_i , less expenditures on medical and non-medical consumption. Medical expenditures are subsidized by the government so that person i 's out-of-pocket costs are equal to $\gamma_{it}m_{it}$. The subsidy rate, γ_{it} , varies across people depending on their mix of medical services. Finally, retirees' incomes are assumed to be fixed.

3.3 The dynamic optimization problem

The retiree's dynamic optimization problem can be expressed as the following Bellman equation:

$$V_t(a_{it}, y_i, h_{it}) = \max_{\{c_{it}, m_{it}\}} u(c_{it}, h_{it}) + \alpha_i s_{it}(m_{it}, h_{it}, t) E[V_{t+1}(a_{it+1}, y_i, h_{it+1})]. \quad (5)$$

Each period, the agent allocates assets toward medical and non-medical consumption to maximize expected utility over the remaining lifetime, with a discount factor of α_i . The expectation operator is taken with respect to the following period's health stock. The maximization problem is subject to the budget constraint in (4), the survival function in (3) and the health production function in (2).¹¹

The agent maximizes utility by choosing levels of medical and non-medical consumption such that their marginal utilities are equalized at each age. Solving the optimization problem in period t and combining the first-order conditions yields the following expression:

$$\frac{1}{\gamma_{it}} \left[\alpha_i EV_{t+1}(a_{it+1}, y_i, h_{it+1}) \frac{\partial s_{it}}{\partial m_{it}} + \alpha_i s_{it} \frac{\partial E[V_{t+1}(a_{it+1}, y_i, h_{it+1})]}{\partial m_{it}} \right] = u_c(c_{it}, h_{it}). \quad (6)$$

The first term inside brackets reflects the discounted stream of benefits from increasing the survival probability by investing an additional dollar in medical expenditures. The second term captures the associated return in terms of improved future health. Dividing by the marginal utility of income and rearranging terms yields an equilibrium condition equating the marginal benefits and costs of investing in medical care for probabilistic life extension:

$$\frac{\alpha_i EV_{t+1}(a_{i,t+1}, y_i, h_{it+1})}{u_c(c_{i,t}, h_{i,t})} + \alpha_i \frac{s_{i,t}}{u_c(c_{i,t}, h_{i,t})} E \left[\frac{\partial V_{t+1}(a_{i,t+1}, y_i, h_{it+1})(f_m/g_m)}{\partial h_{it+1}} \right] = \gamma_{it} \frac{\partial m_{it}}{\partial s_{it}}. \quad (7)$$

The expression to the left of the equality in (7) defines the expected marginal private benefit of increasing

¹¹ Because we do not explicitly model bequests, the utility value of transferring wealth to others is implicitly included as a form of non-medical consumption.

the survival probability via medical expenditures. The first term is the benefit of surviving to the next period conditional on the health stock. The second term captures the co-benefit of increasing the future health stock via the same investment in medical consumption that increases the survival probability. The ratio $f_m/g_m = \frac{\partial h_{it+1}}{\partial m_{it}} \frac{\partial m_{it}}{\partial s_{it}}$ tracks how the increase in medical expenditures that is used to marginally increase the survival probability affects the future health stock which, in turn, influences both the quality of life and the survival probability in future periods.

The expression to the right of the equality in (7) defines the marginal private cost of increasing the survival probability. This statistic is proportional to the cost of saving a statistical life. For instance, if increasing medical expenditures by \$10,000 increases the survival probability by 0.001, then the total cost of avoiding the death of one type i individual at age t is \$10 million. If $\gamma_{it} = 0.5$ then the private out-of-pocket cost of avoiding that death for type i individuals is \$5 million.

The equilibrium condition in (7) relates two important statistics for evaluating public policies that may affect people's health and survival: the marginal return to medical spending and the VSL. The relationship between them suggests a simple "sufficient statistics" approach to estimating VSL. First estimate the survival function, then differentiate with respect to medical spending to calculate $\frac{\partial s_{it}}{\partial m_{it}}$, which can be rescaled by the coinsurance rates to calculate private VSL for a type i individual at age t , as in (8):

$$VSL_{it} \propto \frac{\gamma_{it}}{\partial s_{it} / \partial m_{it}}. \quad (8)$$

This VSL measure recognizes that an agent's willingness to pay for a marginal change in statistical life extension may depend on their expected future health. Assuming that flow utility is strictly increasing in health and that expected future health is weakly increasing in medical spending, equation (7) implies that the VSL revealed by medical spending will exceed a hypothetical health-neutral VSL.¹² This feature is also common to the market environments used to estimate VSL for younger people. For example, improving workplace safety is likely to reduce the risk of a variety of non-fatal injuries as a co-benefit to reducing the risk of death on the job.

Under additional assumptions, the model can be used to formalize hypotheses about sources of heterogeneity in the VSL. For instance, under mild restrictions on similar life-cycle models, (Murphy and Topel (2006) and Hall and Jones (2007)) predict that the VSL will increase in wealth and decline in age among retirees on fixed incomes

¹²Alternatively, a health-neutral VSL could exceed the measure in (8) if medical spending reduces future health, for example through undesirable side effects of prescription drugs.

because the health stock and survival probability both tend to decline beyond age 65. Similarly, [Dow, Philipson and Sala-I-Martin \(1999\)](#) and [Murphy and Topel \(2006\)](#) predict that complementarity between different types of health investments will cause the VSL to decline as people experience negative health shocks. For example, a diagnosis of heart failure may reduce the VSL by accelerating the expected decline of health. We test these and other hypotheses about heterogeneity in the VSL in Section 7.

4 Econometric Model of Survival

As our theoretical model demonstrates, a key empirical object for our approach to measuring VSL is the individual-specific returns to medical spending in terms of reduced mortality. Prior work on estimating the return to medical spending used long-term aggregate measures ([Hall and Jones \(2007\)](#)) or estimated local average treatment effects for specialized cohorts of patients or specific types of medical spending ([Huh and Reif \(2017\)](#), [Clayton \(2018\)](#)), [Doyle et al. \(2015\)](#), [Romley and Sood \(2013\)](#)) such as hospital spending for patients who were hospitalized through the emergency department for heart attacks while visiting Florida ([Doyle \(2011\)](#)). Because prior literature does not provide the information needed for our approach, one contribution of this article is to provide the full set of estimates of the marginal returns to medical spending across the full range of age, health and other characteristics.

We model survival as a discrete-time process over the annual intervals at which we observe individuals' medical expenditures. Each year, death is predicted by lagged values for medical expenditures and health. Formally, let s_{it}^* be a latent variable that determines survival, scaled so that person i lives through period $t + 1$ if and only if $s_{it}^* > 0$. Survival depends on medical expenditures, health, age, and a random shock:

$$s_{it}^* = \beta + \beta_m m_{it} + \beta_h h_{it} + \beta_t t - \mu_{it}. \quad (9)$$

The probability of survival, s_{it} , can be represented as

$$s_{it} = Pr(s_{it}^* > 0) = Pr(\beta + \beta_m m_{it} + \beta_h h_{it} + \beta_t t > \mu_{it}). \quad (10)$$

Under the assumption that survival shocks are *iid* draws from a Type I extreme value distribution, the survival

probability takes the complementary log–log form,

$$s_{it} = \exp(-\exp(\beta + \beta_m m_{it} + \beta_h h_{it} + \beta_t t)). \quad (11)$$

This parametric form is an intuitive choice for modeling death among older adults because the model’s asymmetry allows the probability to approach 1 (survival) slowly relative to the rate at which it approaches 0 (death).¹³ We measure the explanatory variables at the end of calendar year t so that we are modeling how survival during a particular year depends on health and age at the start of that same year, along with total medical expenditures for the prior year.

Latent health presents a key challenge to identifying the survival model parameters. While our data contain a rich set of measures of each person’s health (Table 1), any function of those variables is still likely to contain some error in measuring the true stock of health that determines survival. This problem is magnified by the potential for the latent component of health to be correlated with both medical expenditures and survival. That is, people who are sicker in unobserved ways are likely to have higher medical spending and lower survival rates.¹⁴ All else constant, this simultaneity will lead to a downward bias in our estimator for the marginal effect of medical spending on survival and upward bias in our estimator for the VSL. We use several different instrumental variables approaches to mitigate this threat.

4.1 Constructing an instrument for medical expenditures

Economists have often used geographic variation in medical treatment style to construct instruments for measuring how medical care affects survival (e.g. McClellan, McNeil and Newhouse (1994), Stukel et al. (2007), Currie and Slusky (2020)). Our featured instrument adapts the method developed in Finkelstein, Gentzkow and Williams (2016) to decompose geographic variation in medical expenditures across the 306 US Hospital Referral Regions (HRRs) into demand–side factors and place–specific supply factors.¹⁵ Making this decomposition enables us to

¹³ If we rescale the dependent variable to be 1 in the case of death the resulting mortality function is commonly known as the Gompit model because of its similarity to the Gompertz model of human mortality. Section 6.1.4 shows that our main findings are robust to estimating a Gompertz mortality function.

¹⁴ For example, consider the severity of disease. CMS records allow us to observe if and when each individual is first diagnosed with chronic kidney disease, but we are unable to directly observe whether the individual’s kidneys are mildly damaged (stage 1) or have already failed so that the individual requires costly dialysis treatments or a kidney transplant to survive (stage 5).

¹⁵ “Hospital Referral Regions” represent regional medical care markets for tertiary medical care as determined by the Dartmouth Atlas. Each HRR contains at least one hospital that performs major cardiovascular procedures and neurosurgery. HRRs were defined by assigning Hospital Service Areas to the region where the greatest proportion of major cardiovascular procedures were performed, with minor modifications to achieve geographic contiguity, a minimum population size of 120,000, and a high localization index. The Dartmouth Atlas defines a Hospital Service Area as a collection of ZIP Codes whose residents receive most of their hospitalizations from hospitals in the area. For further details

identify the survival model parameters from variation in medical expenditures that is unrelated to patient health. Intuitively, people with identical morbidities who live in different geographic areas face menus of treatment options with different costs, which leads to variation in their medical expenditures and survival probabilities. We develop an index of this supply–side variation to use as an instrument for individuals’ medical expenditures.

The logic for our instrument starts from the observation that per capita annual medical expenditures vary greatly across the United States among the Medicare population.¹⁶ Some of this variation may reflect patient preferences and health, whereas some of the variation reflects differences in the supply of medical care. For example, [Cutler et al. \(2019\)](#) points to the importance of differences in physician practice style, highlighting that aggressive treatment practices increase spending. [Chandra and Staiger \(2007\)](#) highlights the importance of productivity spillovers and physician migration. Other supply factors that may contribute to spatial variation in expenditures include peer effects among physicians, differences in physical capital, and institutional features of local medical care markets. Against this background, [Finkelstein, Gentzkow and Williams \(2016\)](#) use Medicare administrative records on patients who move between HRRs to implement a regression–based procedure to decompose the spatial variation in patients’ expenditures into supply and demand factors, finding that each source accounts for about half of the total variation. We follow their approach to estimation and use the resulting measure of regional supply–side variation in expenditures to instrument for individuals’ total expenditures. In Section 6.1.3 we describe other ways of constructing this instrument and alternative instruments altogether and show that the VSL estimates are robust across these analytic decisions.

We construct the instrument from 3.2 million person–years of claims-based data on expenditures for 484,000 people over age 65 who were enrolled in traditional Medicare and changed their residential address from one HRR to another exactly once between 1999 and 2013. These data were extracted from a 10% random sample of Medicare beneficiaries. We use CMS administrative records for this random sample of movers to estimate the supply–side component of their Medicare expenditures that varies across HRRs,

$$m_{ijt} = \sigma_i + \phi_j + \chi_t + \psi X_{it} + o_{ijt}. \tag{12}$$

The dependent variable in the regression is total medical expenditures in year t for a person living in HRR j . Covariates include an individual fixed effect, σ_i , an HRR fixed effect, ϕ_j , a year fixed effect, χ_t , a vector of time–

see: <http://www.dartmouthatlas.org/downloads/methods/geogappdx.pdf>.

¹⁶ This fact is well documented. For evidence, see [Finkelstein, Gentzkow and Williams \(2016\)](#), [Cutler et al. \(2019\)](#) and references therein.

varying person-specific covariates, X_{it} , and an orthogonal error term, o_{ijt} . We follow [Finkelstein, Gentzkow and Williams \(2016\)](#) in defining X_{it} to include dummies for 5-year age bins and dummies for the current year relative to the year in which an individual is observed moving between HRRs. These relative-year dummies allow migration decisions to coincide with unobserved health shocks that simultaneously affect the demand for medical care and the desire to live close to caregivers such as children. The resulting vector of HRR fixed effects, $\hat{\phi}_1, \dots, \hat{\phi}_{306}$, provides an index for spatial variation in medical expenditures driven by supply factors.

The validity of using $\hat{\phi}$ to instrument for medical expenditures rests on the assumption that it is uncorrelated with latent health. While this assumption cannot be tested directly, the specification in (12) is designed to reduce the scope for such correlation. [Finkelstein, Gentzkow and Williams \(2016\)](#) provide a detailed discussion of how the HRR fixed effects in equation (12) are identified by supply-side variation, holding health fixed. In summary, the index $\hat{\phi}$ is identified by the ways in which changes in movers' medical expenditures differ between their origin and destination locations. To see this, first notice that if people never moved then $\hat{\phi}$ could not be identified because the individual fixed effects would absorb all of the spatial variation in average per/capita expenditures. Second, the year-relative-to-move fixed effects included in X_{it} absorb average trends in medical expenditures around the move. This forces the identification to come from differential changes in expenditures across people undertaking different migration patterns. While equation (12) allows expenditures to differ arbitrarily between movers (via σ_i) and to differ systematically around their moves (via X_{it}), [Finkelstein, Gentzkow and Williams \(2016\)](#) point out that it maintains the assumption that health shocks leading to expenditure changes do not precisely coincide with the timing of moves. We relax this assumption by the way we construct the estimation sample. Specifically, our estimation sample excludes movers who were newly diagnosed with any chronic conditions during their move year.¹⁷ Thus, the instrument is identified by differential changes in medical expenditures among people who move between HRRs and do not experience observed health shocks during their move years.

4.2 Main econometric model

We instrument for total medical expenditures in a linear first-stage regression,

$$m_{it} = \pi + \pi_h h_{it} + \pi_t t + \pi_z z_{it} + \omega_{it}, \quad (13)$$

¹⁷ We also follow [Finkelstein, Gentzkow and Williams \(2016\)](#) in dropping the very small fraction of people who moved multiple times because such individuals complicate the definition of the year-relative-to-move fixed effects.

where the supply side expenditure index for person i living in HRR j in year t is defined as $z_{it} = \hat{\phi}_j - \hat{\phi}_k$, and k is used to index an arbitrary reference location. We then estimate the survival function as the second stage control function that includes the first-stage residuals:

$$s_{it} = \exp(-\exp(\beta + \beta_m m_{it} + \beta_h h_{it} + \beta_t t + \hat{\omega}_{it})), \quad (14)$$

where h_{it} includes measures of demographics, socioeconomic status, and health from Table 1 as proxy measures for the health stock. Finally, we use the parameter estimates to predict the marginal effect of medical expenditures on survival and rescale it by the coinsurance rate to calculate the VSL measure from equation (8).

In summary, our estimation approach proceeds as follows. First we use Medicare claims data to construct the expenditure instrument from equation (12). Then we estimate equations (13) and (14) using the linked MCBS-administrative data. The two-stage control function approach yields a consistent estimator for model parameters under the assumption that the instrument is valid and the survival function is correctly specified (Terza, Basu and Rathouz (2008), Wooldridge (2015)). We calculate standard errors and confidence intervals using a nonparametric bootstrap over sequential estimation of (13) and (14) with the errors clustered by HRR to coincide with the identifying source of variation in the instrument (Cameron, Gelbach and Miller (2008), Abadie et al. (2017)).

5 Main Results

5.1 Evidence on Supply-Side Variation in Medical Expenditures

We use equation (12) to estimate HRR fixed effects from claims-based data on movers. Then we normalize the estimates relative to Birmingham, AL. The normalized estimates range from +\$2,500 for Miami, FL to -\$1,150 for Greensboro, NC. Moving from the 10th percentile to the 90th percentile in the distribution of HRR effects is equivalent to increasing annual expenditures by \$1,870, or approximately 22% of the mean expenditure that we observe in claims-based data for traditional Medicare enrollees in MCBS data.¹⁸ Likewise, the between-HRR standard deviation is \$661 (7% of the mean expenditures). These results are consistent with Finkelstein, Gentzkow and Williams (2016) in suggesting that supply-side factors explain a substantial share of the between-HRR variation in

¹⁸ Figure A.1 shows the entire distribution of HRR estimates.

Medicare expenditures per person.¹⁹

Next we collapse the normalized fixed effects into an index of supply-side variation in average expenditures, assign index values to MCBS respondents based on their residential locations, and use this variable to instrument for their annual medical expenditures. To test the hypothesized supply-side mechanisms underlying the instrument we regress it on measures of physician treatment style constructed by [Cutler et al. \(2019\)](#). Specifically, we use their measures for each HRR's fraction of "cowboy" physicians who consistently recommend intensive care beyond clinical guidelines and the fraction of "comforter" physicians who consistently recommend palliative care for the seriously ill.²⁰ Consistent with the hypothesized mechanisms, a standard deviation increase in the cowboy share is conditionally associated with a 0.16 standard deviation increase in the IV, whereas a standard deviation increase in the comforter share is associated with a 0.17 standard deviation decrease in the IV.

Finally, we test whether the claims-based instrument for medical spending has power to explain variation in medical expenditures that are not reflected in Medicare claims, for example among the MCBS sample on Medicare Advantage. The answer is yes. For each MCBS respondent enrolled in traditional Medicare we calculate the difference between the comprehensive MCBS expenditure measure and the corresponding claims-based measure from CMS administrative files. Univariate regression reveals that a standard deviation increase in the instrument is associated with a 0.04 standard deviation increase in expenditures not processed by Medicare compared to a 0.07 standard deviation increase in expenditures processed by Medicare. These coefficients only decline by about 25% when we add the comprehensive set of covariates described below. Thus, the identifying variation in medical spending comes partly from services covered by Medicare and partly from services that are covered entirely by a combination of Medicare Advantage plans, Medigap plans, employer plans and OOP spending.

5.2 The Effect of Medical Expenditures on Survival

Table 2 reports average marginal effects for survival functions, using 1,000 bootstrap replications to calculate standard errors.²¹ In addition to the health covariates featured in the table, all specifications include the demographic and socioeconomic variables summarized in Table 1. First-stage coefficients on the instrument and associated F-statistics are reported at the bottom of the table and unabridged estimates from the first- and second-stage

¹⁹ Because these estimates exclude all movers who are diagnosed with new chronic conditions during their move years our estimates can also be interpreted as providing additional support for [Finkelstein, Gentzkow and Williams \(2016\)](#) by showing that their qualitative findings are robust to relaxing their maintained assumption that expensive health shocks do not coincide with moves.

²⁰ We thank Jon Skinner for sharing these data.

²¹ Each replication repeats both stages of estimation.

models are reported in Appendix Tables [A.1](#) and [A.2](#).

The model in column (1) ignores the potential endogeneity of medical spending. The positive marginal effect indicates that, all else constant, a \$1,000 increase in spending is associated with a 0.06 percentage point increase in the one-year mortality rate. This counterintuitive result is consistent with the idea that people who are sicker in unobserved ways will tend to spend more on medical care and die sooner. This counterintuitive result disappears when we instrument for medical spending.

Column (2) reports results from the control-function analog to column (1). The first-stage residual measure of latent morbidity has a positive, statistically significant coefficient. This supports the view that unobserved latent health is positively correlated with both medical expenditure and mortality and that the assumed exogeneity of medical expenditure in column (1) is unlikely to hold true.

The IV survival function implies that a \$1,000 increase in medical spending reduces the one-year mortality rate by just under half a percentage point. The first-stage F-statistic and IV coefficient reported toward the bottom of the table indicate that the IV is adequately powered and that a marginal dollar increase in the supply-side index of medical expenditures is associated with less than a dollar increase in total expenditures.

Columns (3)–(6) show results from repeating estimation of the IV survival function after incrementally adding additional covariates. These specifications address our concern that the marginal effect of medical spending in column (2) may be biased toward zero if our adaptation of the [Finkelstein, Gentzkow and Williams \(2016\)](#) decomposition does not fully purge latent health. For example, some of our estimated between-HRR “supply-side” variation in expenditures could be caused by people opting in or out of private insurance coverage at the time of their moves. These adjustments could introduce bias due to adverse or advantageous selection on latent health. Column (3) addresses this concern by adding separate indicators for whether each person was enrolled in a Medigap plan, a Medicare Advantage plan, or received Medicaid benefits. Adding these covariates reduces the estimated return to spending by about 10%.

A second concern is that we may understate the return to spending if higher-expenditure HRRs have a higher marginal impact on health per dollar spent, for example because they have higher-quality medical care providers. We test this hypothesis by adding controls for hospital quality in column (4). The additional covariates include HRR-specific measures of the number of hospital beds per capita, primary care physicians per capita, specialists per capita, discharges for ambulatory care-sensitive conditions among Medicare beneficiaries, and CMS’s “Hospi-

Table 2: Average Marginal Effects on Mortality

Outcome: Mortality in year t+1	(1)	(2)	(3)	(4)	(5)	(6)
\$1,000 in medical spending	0.06*** (0.01)	-0.47*** (0.20)	-0.42*** (0.18)	-0.64*** (0.26)	-0.88*** (0.63)	-0.82* (5.80)
1st stage residual morbidity		0.54*** (0.20)	0.49*** (0.18)	0.71*** (0.26)	0.95*** (0.63)	0.88* (5.80)
HCC index	3.10*** (0.31)	11.11*** (3.01)	10.72*** (2.67)	13.91*** (3.86)	17.77*** (9.42)	16.73** (86.50)
one or more ADL restrictions	1.70*** (0.23)	3.05*** (0.56)	2.96*** (0.52)	3.51*** (0.70)	4.13*** (1.57)	4.02** (14.31)
one or more IADL restrictions	0.50** (0.24)	1.28*** (0.40)	1.22*** (0.37)	1.54*** (0.49)	1.92*** (0.99)	1.86** (8.74)
ever smoked	1.19*** (0.22)	1.30*** (0.25)	1.27*** (0.24)	1.31*** (0.26)	1.34*** (0.30)	1.33*** (1.32)
underweight BMI	2.67*** (0.36)	2.26*** (0.43)	2.31*** (0.42)	2.15*** (0.49)	1.99** (0.71)	2.01* (3.74)
health = poor	3.09*** (0.32)	7.19*** (1.67)	6.85*** (1.52)	8.52*** (2.23)	10.46*** (5.02)	9.85** (46.17)
health = fair	1.48*** (0.26)	2.76*** (0.54)	2.67*** (0.52)	3.19*** (0.73)	3.81*** (1.61)	3.61** (14.96)
health = very good	-1.52*** (0.26)	-2.39*** (0.45)	-2.34*** (0.43)	-2.69*** (0.54)	-3.08*** (1.04)	-2.95** (9.44)
health = excellent	-2.40*** (0.39)	-3.81*** (0.67)	-3.70*** (0.62)	-4.27*** (0.79)	-4.91*** (1.64)	-4.70** (14.83)
insurance type covariates			x	x	x	x
health care quality covariates				x	x	x
environmental covariates					x	x
state dummies						x
1st-stage coefficient on IV		0.87*** (0.17)	0.96*** (0.17)	0.80*** (0.16)	0.65*** (0.19)	0.67** (0.29)
F-statistic on the IV		27	32	25	11	5
number of person-years	44,697	44,697	44,697	44,697	44,697	44,697
number of people	22,206	22,206	22,206	22,206	22,206	22,206

Note: The table reports average marginal effects expressed as percentage point changes in the one-year probability of death. All models include age, sex, age interacted with sex and with an indicator for over 90, and indicators for race, educational attainment, marital status, and living children. Columns (2)–(6) instrument for medical spending. Column (3) adds indicators for insurance coverages: Medigap, Medicaid, and Medicare Advantage. Column (4) adds HRR-level measures of CMS's hospital compare index, and per capita measures of the numbers of acute care hospital beds, primary care physicians, medical specialists, and hospital admissions for ambulatory care-sensitive conditions. Column (5) adds HRR-level measures of automobile mortality, homicide mortality, fine particulate matter, mean winter low temperature, mean summer high temperature, share urban, median income, high school graduation rate, and Census division dummies. Column (6) replaces the Census division dummies with state dummies. Standard errors are calculated using 1,000 bootstrap replications and clustered by hospital referral region. Asterisks indicate statistical significance at the 1% (***), 5% (**) and 10% (*) levels.

tal Compare” index of hospital quality.²² Indeed, we find that adding this set of proxy measures for hospital quality moderately increases our estimated return to spending to -0.64.

A related concern is that some regions may have higher medical expenditures due to environmental conditions that impair population health, attenuating our estimate for the return to spending. Column (5) addresses this concern by adding a set of proxy measures for local environmental conditions. These include automobile mortality, homicide mortality, average concentrations of fine particulate air pollution $PM_{2.5}$, average winter minimum temperature, average summer maximum temperature, the fraction of people living in urban areas, median income, the high school graduation rate, and dummy variables for the nine Census regions. Intuitively, we find that adding these controls further increases our estimated return to spending to -0.88. Finally, column (6) further tightens the controls for unobserved environmental conditions by replacing the Census region dummies with state dummies, forcing the identification to come from within-state variation in the HRR index (the average state has 6 HRRs). The resulting estimate for the return to spending is nearly the same at -0.82 despite the drop in statistical power.²³

Taken together, columns (2) through (6) define a range of estimates for the return to medical spending from -0.88 to -0.42. We use column (3) as our main specification for calculating VSL. This model utilizes all of the micro-data describing individual health and insurance coverages. In comparison, adding the additional HRR-level covariates and state dummies in columns (4), (5) and (6) presents a tradeoff. It mitigates potential bias from spending being correlated with other spatially varying determinants of health, but in the absence of bias it reduces identifying variation in the instrument and statistical precision. Indeed, the 95% confidence intervals on the estimated effects of medical spending in columns (4), (5), and (6) include the point estimate from column (3). Another reason why we feature the model in column (3) is that it generally yields the largest VSLs among the specifications in Table 2. Therefore, our choice to feature this specification works against the hypothesis that VSL measures based on seniors’ medical care choices fall below VSL measures based on workers’ occupation choices. Section 6 presents our full range of VSL results from a comprehensive sensitivity analysis to alternative specifications including those in Table 2 and many others. As we show there in Figure 5, this featured model falls near the middle of the distribution of VSL estimates for every given age.


²² This index is primarily derived from measures of the shares of patients who receive “timely and effective” care upon arrival at hospitals: <https://data.medicare.gov/data/archives/hospital-compare>. An example is the share of heart attack patients who are given aspirin. To measure average quality for each HRR, the shares are first averaged across all measures and hospitals in each HRR for each reporting period, and then averaged over all reporting periods in a year, and finally averaged over years.

²³The large standard errors in column (6) are driven by a few outliers.

5.3 Model Fit

The marginal effects of the health measures in Table 2 are intuitive and quantitatively important. For example, a standard deviation increase in the HCC morbidity index of observable chronic illnesses is associated with a 11 to 18 percentage point increase in the one-year probability of death across the IV models. Mortality is also conditionally higher among people with basic and instrumental limitations in activities of daily living, a history of smoking, a BMI that classifies them as being underweight, and a relatively poor subjective assessment of their own health. The reference category for self-reported health is “good”. Moving from “good” to “poor” is associated with a 7 to 10 percentage point increase in the probability of death, whereas moving from “good” to “excellent” is associated with a reduction of 4 to 5 percentage points.

Figure 3: Predicted and Actual One-Year Mortality Rates for Males and Females



Note: The dashed lines show one-year mortality rates by age and sex in the data. The solid lines show model predictions from col (3) of Table 2.


Figure 3 shows model fit by comparing its predictions for one-year mortality rates by integer age and sex to the data. Model predictions closely approximate mortality through age 87. Beyond age 87 the model continues to capture the upward trend in average mortality but it does not reproduce much of the idiosyncratic year-to-year

variation around the trend. This improves our assessment of model fit because idiosyncratic deviations from the trend after age 90 are likely to reflect statistical imprecision caused by declining sample size.

5.4 Value of a Statistical Life

Figure 4 plots the VSL profile from age 67 to age 97 based on the model in column (3) of Table 2. This figure highlights several important features of our results. First, our VSL estimates for seniors are an order of magnitude below the prevailing wage–hedonic estimates for workers who are, on average, in their early 40s and in much better health. The solid line in the figure shows the mean VSL by age. At age 67 the mean VSL is about \$402,000.

Figure 4: VSL Estimates: Age 67 to Age 97



Note: See the text for details.

Second, the mean VSL declines with age in a near-monotonic fashion. This curvature is driven by the data. It is worth reiterating that our econometric model does not embed assumptions for the parametric form of utility or

the rate of time preference. Rather, the shape of the curve reflects individuals' decisions about how much of their own money to spend on medical care, given their current health, wealth, preferences, and beliefs about the return to spending. Analyzing the underlying components of the VSL equation (8) reveals that the downward trend in age results from dividing the individual coinsurance rate, which is relatively flat in age, by the return to medical spending, which increases in age. These trends are shown in Figure A.2. Intuitively, the decision not to spend more on one's health when the return to doing so is relatively high reveals that the VSL must be relatively low.

Third, our estimates for the mean VSL are reasonably precise. The shaded region in Figure 4 defines the bootstrapped 95% confidence interval on our estimate for the age-specific mean. It is asymmetric around the mean because the VSL is inversely proportional to the estimated return to medical spending. Even at age 67 the upper bound of the confidence interval is below \$1 million.

Finally, Figure 4 shows that there is substantial heterogeneity in the VSL conditional on age. The dotted and dashed lines denote the 5th and 95th percentiles in the distribution of our age-specific estimates. At age 67 the 95th percentile is approximately \$1.5 million.²⁴ However, the variation across people at a given age declines sharply with age.

Viewing these results through the lens of the life-cycle model can explain why the mean and variance of VSL both decline in age. As people get older their health tends to decline, as does the variance in remaining life expectancy. Negative health shocks reduce the expected future quantity and quality of life, creating a disincentive to invest in probabilistic life extension. While medical spending increases in age (Figure 2 panel F), it does not increase by enough to reduce the marginal return to further spending, yielding a decline in both the mean and variance of the VSL. In Section 7.1 we document similar patterns arising from heterogeneity in health conditional on age.

6 A Systematic Sensitivity Analysis

We take a systematic and comprehensive approach to testing the robustness of our main VSL estimates to modifying features of our research design. Our approach is inspired by Leamer (1983), Banzhaf and Smith (2007), and Greenstone, Kopits and Wolverton (2013). First we define a set of potential modeling decisions along each dimension of our research design. Then we report VSL estimates derived from every possible combination of modeling

²⁴ Figure A.3 further illustrates the within-age heterogeneity in VSL by showing the distribution among 70-year-old people. Within that group, 93.74% of people have VSL values below \$1 million, 5.52% have VSL values between \$1 and \$2 million, and 0.74% of people have values over \$2 million.

decisions.

6.1 Modifiable Features of the Research Design

6.1.1 Including or Excluding Workers

Our main estimation sample excludes data for 5,764 person–years where the beneficiary was employed at the time of their MCBS interview. This exclusion improves internal validity by sharpening our focus on medical care as the relevant market for trading consumption against mortality risk, but it threatens external validity. We can investigate this threat by adding workers to the estimation sample. Doing so tends to increase the VSL modestly. For example, repeating estimation of the model from column (3) of Table 2 yields higher VSL measures among workers at each age from 67 to 87 with an average age–specific differential of 38%.²⁵ One explanation for this differential is that at any given age healthier and wealthier workers may be willing to pay more for statistical life extension and be less likely to retire.

6.1.2 Using MCBS or Claims–Based Data on Medical Expenditures

MCBS provides the most comprehensive data on total and OOP medical spending. Unfortunately MCBS does not collect these data during the first year of survey participation, which is why our main specification uses data from survey years 2 through 4. Alternatively, we can use all four years of survey data if we are willing to swap the MCBS spending measures for CMS’s less comprehensive measures derived from the universe of claims processed under Medicare Parts A and B. Swapping the spending measures alters our sample size in countervailing ways. It expands the sample to include some observations from MCBS year 1 while simultaneously excluding people in each year who enrolled in Medicare Advantage plans (for whom claims–based spending data are unavailable). The net effect of these adjustments is to increase our sample size by 6%. More importantly, repeating the estimation for each sample reveals whether our findings are driven by how people respond to price adjustments by insurers that cannot be observed in claims data from fee-for-service Medicare.

²⁵ This age range accounts for 99% of workers in the MCBS sample. There is a near monotonic decline in the number of workers per integer age. We observe 28 people working at age 87, but never more than 15 people at older ages.

6.1.3 Alternative Instrumental Variables for Medical Expenditures

Our main specification for the instrument in equation (12) followed [Finkelstein, Gentzkow and Williams \(2016\)](#) in using dummies for 5-year age bins to absorb unobserved changes in health that could have occurred around each migrant's move year. As a robustness check, we incrementally relax the exclusion restriction on the IV to allow for additional forms of sorting on unobserved health. First we reconstruct the IV after replacing the dummies for 5-year age bins in (12) with dummies for integer age. Then we reconstruct the IV a second time using sex-by-integer-age dummies. As a third alternative, we reconstruct the IV after extending the sample to include people who never moved.²⁶ This increases statistical power and yields a more nationally representative sample of seniors. Finally, we construct a fourth alternative instrument from data on end-of-life spending based on evidence from [Skinner, Fisher and Wennberg \(2005\)](#) and [Cutler et al. \(2019\)](#) that a significant fraction of spatial variation in end-of-life spending is explained by variation in physician practice style. Specifically, we use average per-patient spending during the last 6 months of life reported by the Dartmouth Atlas at the HRR level.

6.1.4 Alternative Parametric Forms of the Survival Function

As an alternative to our featured Gompit specification for the survival function, we repeat the estimation using a Gompertz specification similar to the ones used by [Chetty et al. \(2016\)](#) and [Finkelstein, Gentzkow and Williams \(2019\)](#). The Gompertz form assumes that the log of the mortality rate is linear in covariates. This restriction reduces model fit based on the log-likelihood function value in our main specification from Table 2. Nevertheless, the Gompertz model has the benefit of simplicity and familiarity, having served as a common approach to modeling mortality for almost 200 years.

6.1.5 Alternative Covariates in the Survival Function


All of our IV specifications include the covariates shown in Table 2 and summarized in the footnote to the table. We repeat estimation as we incrementally add each of the five sets of augmented covariates. Corresponding to the columns in Table 2, these include: (2) no additional covariates, (3) insurance plan enrollment covariates, (4) medical care quality covariates, (5) environmental covariates, and (6) state dummies.

²⁶ This matches the featured specification in [Finkelstein, Gentzkow and Williams \(2016\)](#). They exclude movers as a robustness check.

6.2 Results

Altogether we consider five different sets of covariates, two parametric forms for the survival function, five different instruments for medical spending, two ways of measuring medical spending, and models including and excluding workers. Considering all permutations of these modeling decisions yields 200 different models. We estimate each one and calculate the mean VSL by age.

Figure 5: Sensitivity of VSL Estimates to Model Features



Note: The figure shows the estimated mean VSL by age for 200 different specifications of the survival function. Each line corresponds to a different combination of modeling decisions as described in the main text. The large dashed line is our main specification from col (3) of Table 2.

Figure 5 shows results from the 200 models. The dashed line highlights our main specification from Figure 4. It sits near the middle of the range of estimates. Readers who disagree with our preferred modeling decisions can see how much those decisions matter relative to the alternatives outlined above.

At age 67 our preferred estimate is \$402,000. The 5th and 95th percentiles in the distribution of models are \$189,000 and \$555,000, and the maximum is \$899,000. These moments provide a partial measure of the model uncertainty in our VSL estimates. They have practical relevance because federal agencies use such moments to

define benchmarks for sensitivity analysis, e.g. when using the social cost of carbon in policy evaluations (Greenstone, Kopits and Wolverton (2013)). Notably, every one of the 200 specifications yields mean VSL estimates that lie below \$1 million at ages 67 and above, and substantially below the mean VSL estimates derived from occupation choices made by younger, healthier workers.

We use an internal meta-analysis (Banzhaf and Smith (2007)) to determine which factors cause the variation seen in Figure 5. Specifically, to summarize how modeling decisions affect the estimated VSL, we regress the log of mean VSL from each of the 200 models on indicators for model features. Relative to our main specification, VSL estimates tend to increase if we add workers to the estimation sample (+24%), switch to using claims-based spending measures (+30%), or instrument using end-of-life spending (+39%). VSL estimates tend to decrease slightly if we switch to the Gompertz specification for mortality (-8%), fail to control for selection into insurance coverages (-11%), or include never-movers in the sample used to construct the instrument (-12%). As we previewed earlier, the VSL estimates decrease more substantially if we expand the covariate set to include HRR-level measures of medical care quality and environmental quality (-28% to -42%). The complete meta-regression results are reported in Table A.3.

Overall, we find that the level and curvature of the VSL-age profile is somewhat sensitive to modeling decisions, but two of its most important features are thoroughly robust. First, the VSL declines with age. Second, \$1 million provides an upper bound on the VSL implied by seniors' medical expenditures.


7 Heterogeneity

7.1 Heterogeneity in the Return to Medical Spending by Health

Figure 6 summarizes how our estimates for the return to medical spending vary with subjective and objective measures of health. Each of the four panels reports the estimated average percentage point increase in one-year survival from a \$1,000 increase in medical spending. Panels A and B stratify by self-reported measures of health. Panel (A) shows that conditional on age, the return to medical spending increases as self-assessed health declines. For example, at age 72 a \$1,000 increase in spending reduces mortality by 0.8 percentage points for the average person who reports their health as "poor" compared to 0.08 for the average person who reports their health as "excellent". Panel (B) shows the same qualitative pattern such that, conditional on age, the return to further spending is lowest among those with no restrictions on activities of daily living, followed by those with restrictions on instru-

mental activities (e.g. managing money) but not basic activities (e.g. eating), followed by those with restrictions on basic but not instrumental activities, followed by those with restrictions on both basic and instrumental activities. Markers along the trend lines in each panel denote statistical significance of differences in returns between adjacent health categories. The presence of a marker indicates that the mean return at that integer age exceeds the mean return on the next lower trend line in at least 99% of bootstrap samples. Statistical precision declines with age due partly to the decline in age-specific sample sizes.

Figure 6: Survival Gains from Marginal Increase of \$1,000 in Medical Spending



Note: Each panel shows the average marginal effect (AME) of a \$1,000 increase in medical spending on the probability of surviving to the end of the following year measured in percentage points on the vertical axis and calculated from the model shown in col (3) of Table 2. Markers along each trend line denote ages at which the AME exceeds the AME for the next lower trend line in at least 99% of 1,000 bootstrap samples with errors clustered by hospital referral region.

Panels (C) and (D) show that the basic pattern persists if we instead stratify by objective measures of health.

In Panel (C) the age-specific return is always lower among people who have been diagnosed with fewer than the median number of chronic conditions for people of their age. In Panel (D) the age-specific return is always lower among people with HCC scores below the median for their age.

Our estimates for the levels of returns and their variation with respect to health and age in Figure 6 essentially span the range of local average treatment effects that prior studies estimated from quasi-experimental sources of variation in expenditures within the Medicare population. For example, [Huh and Reif \(2017\)](#), [Clayton \(2018\)](#), [Doyle et al. \(2015\)](#), [Romley and Sood \(2013\)](#), and [Doyle \(2011\)](#) collectively suggest a range of marginal returns to \$1,000 of medical spending from about 0.2 to 2, with relatively higher returns among sicker cohorts. Thus, even if the reader is unconvinced by our instrument for medical spending, substituting the reader's preferred estimate for the return to spending from prior literature into our VSL equation (8) would yield VSL measures of the same order of magnitude as the ones we report.²⁷

The stratification patterns in Figure 6 are not likely to be entirely causal. They may also reflect other socioeconomic factors that are correlated with health. For example, at any given age, healthier people are more likely to have a college degree and higher incomes. Attempting to disentangle these mechanisms is worthwhile but tangential to our main objective of estimating the VSL, so we leave it to future research. Regardless of the mixture of casual mechanisms underlying Figure 6, the observed negative relationship between health and the return to spending implies that our VSL estimates will tend to be lower for people in worse health. This occurs because the VSL is defined by the ratio of the coinsurance rate, which tends to increase in health, to the returns to survival from medical spending, which tends to decrease in health.

²⁷ We summarize these prior results here, converting all dollars to 2010 by adjusting by the CPI. At the lower end, [Huh and Reif \(2017\)](#) find that each additional \$1,000 spending on prescription drugs due to the implementation of Medicare Part D reduced mortality by 0.15 percentage points. Among the younger, poorer Medicaid population, however, an additional \$1,000 spending on prescription drugs led to a 2.1 percentage point reduction in mortality ([Clayton \(2018\)](#)). [Doyle \(2011\)](#) uses a similar identification strategy as ours that leverages geographic variation in treatment intensity. Using Medicare beneficiaries who experience heart-related emergencies that lead to hospital admission through the emergency department while visiting Florida, his estimates imply that an additional \$1,000 in spending (in 2010 dollars) reduced annual mortality of 0.2 percentage points. [Doyle et al. \(2015\)](#) relies on quasi-random variation in treatment intensity due to ambulance referral patterns to evaluate the returns to spending among Medicare patients who are experiencing their first hospital admission while on Medicare and arrive at the hospital via ambulance with a subset of illnesses that have high admission rates. They estimate that an additional \$1,000 in spending (in 2010 dollars) reduced annual mortality by about 1.9 percentage points. [Romley and Sood \(2013\)](#) relies in instruments to additionally account for unobserved heterogeneity in hospital productivity and estimate that an additional \$1,000 spending lowered 30-day mortality by 4.7 percentage points for Medicare patients admitted to the hospital due to pneumonia, 2.2 percentage points for those admitted due to congestive heart failure, and 1.8 percentage points for those admitted due to heart attacks.

7.2 Heterogeneity in the VSL by Health

Figure 7 summarizes how VSL varies by subjective and objective health. We focus first on the extremes in Panel (A). At age 67, people who state they are in excellent health have an average VSL close to \$850,000, more than 10 times the average VSL for people who state that they in poor health. People in poor health are more likely to have ever smoked and are diagnosed with more chronic illnesses such as kidney disease (29% compared to 6% of those in excellent health) and congestive heart failure (50% versus 11%). Conditional on age, across all categories the mean VSL increases monotonically with self-reported health.²⁸ The differences between health categories declines in age as relatively healthier groups experience sharper declines: VSL for people in excellent health in their early 90s is similar to VSL for people in good health in their early 80s and to VSL for people in poor health in their late 60s. Panels (B), (C) and (D) show the same patterns emerge when the age-specific mean VSL is stratified by restrictions on activities of daily living, the number of diagnosed chronic medical conditions, or the CMS HCC score. The markers show that the differences between adjacent categories are almost always statistically significant at the 1% level from ages in the late sixties through the late eighties.

7.3 Heterogeneity in the VSL by Behavioral and Socioeconomic Factors


7.3.1 The VSL is Lower for Smokers

Figure 8 summarizes how the VSL-age profile varies with behavioral and socioeconomic factors. Panel (A) highlights a large VSL gap between ever-smokers and never-smokers. At age 67 the VSL among never-smokers is approximately twice as large as among ever-smokers. This gap narrows with age as the differences in remaining life expectancy decline and is statistically indistinguishable from zero beyond age 92. These trends are consistent with the fact that smoking habits are associated with a 10-year reduction in life expectancy (Jha et al. (2013)) and lower quality of life. For example, COPD is twice as common among ever-smokers and lung cancer is six times as common among ever-smokers. These and other chronic illnesses may significantly reduce their expected remaining quantity and quality of life, providing an incentive to shift consumption from medical care to other forms of private consumption.²⁹

²⁸ In addition to the differences in the prevalence of chronic conditions, this within-age pattern is consistent with the fact that the survey question asks people to compare themselves against peers of the same age.

²⁹ Our evidence of the VSL smoking gap late in life diverges from findings reported in wage-hedonic studies. For example, Viscusi and Hersch (2008) augmented a hedonic wage model with data on smoking status and found virtually no difference in the VSL estimated for workers who smoked compared to those who did not. The divergence in results could be explained by the fact that we study people at older ages at which smoking-related morbidities are more likely to have manifested.

Figure 7: Heterogeneity in VSL by Age and Health




Note: Each panel shows the mean age-specific VSL in \$1,000 (2010) dollars stratified by measures of health. The VSL is calculated from the model shown in col (3) of Table 2. Markers along each trend line denote ages at which the VSL exceeds the VSL for the next lower trend line in at least 99% of 1,000 bootstrap samples with errors clustered by hospital referral region.

7.3.2 The VSL is Higher for Females

Panel (B) in Figure 8 shows a VSL gender gap. At age 67, the VSL is approximately twice as high for females, consistent with the higher female life expectancy. The differential declines as the difference in remaining life expectancy falls with age. This evidence validates the out-of-sample predictions made by [Aldy and Smyth \(2014\)](#) and [Murphy and Topel \(2006\)](#) based on life cycle models that incorporate expected longevity.

Our evidence on the VSL-gender gap is novel. Hedonic wage studies rarely stratify VSL estimates by gender due to data limitations. [Leeth and Ruser \(2003\)](#) show that females are less likely to work in high-risk occupations

Figure 8: Heterogeneity in VSL by Smoking, Gender, Income, and Education



Note: Each panel shows the mean age-specific VSL in \$1,000 (2010) dollars stratified by measures of health. The VSL is calculated from the model shown in col (3) of Table 2. Markers along each trend line denote ages at which the VSL exceeds the VSL for the next lower trend line in at least 99% of 1,000 bootstrap samples with errors clustered by hospital referral region.

and, conditional on occupation, females have substantially lower fatality rates. According to the Census of Fatal Occupational Injuries, males account for more than 90% of all accidental deaths on the job. This makes it difficult to calculate precise occupation-by-gender fatality rates, motivating researchers to focus exclusively on males (e.g. [Costa and Kahn \(2004\)](#), [Kneisner et al. \(2012\)](#)). An exception is [Deleire, Kahn and Timmins \(2013\)](#) who report mixed evidence on the VSL gender gap for workers age 18–60 based on combining large worker samples from the Current Population Survey with non-gender-specific data on fatality risk. In contrast, our data capture the important differences in gender-specific health and fatality risk.

7.3.3 The VSL is Increasing in Income and Education

The hedonic wage literature suggests that the VSL is increasing in worker income with a cross-sectional elasticity over one (e.g. Cropper, Hammitt and Robinson (2011)), Evans and Schaur (2010), Viscusi (2010), Aldy and Smyth (2014)). We excluded income from the survival model because we expect income to affect mortality risk indirectly through the covariates describing health and/or medical expenditures. Nevertheless, we can stratify our VSL estimates based on MCBS respondents' income bins to bound the cross-sectional income elasticity.

Panel (C) in Figure 8 shows the expected relationship between VSL and income, with the stratification between bins declining in age. This validates the prediction from Aldy and Smyth (2014) that the VSL-income elasticity will decline late in life as the scope for differences in remaining life expectancy declines. Focusing on the minimum difference in income between people in the top and bottom bins defines upper bounds on the income elasticity of 1.28 at age 67, 0.93 at age 77, and 0.77 at age 87.³⁰

Because income is increasing with education, it is unsurprising to see the VSL increasing in education as well in Panel (D). Nonetheless, the magnitudes are striking. At age 67 the mean VSL among people with a college degree is more than three times as large as for people who did not finish high school and more than 50% larger than for those who did not finish college. Meanwhile, we see virtually no difference between people who finished high school and did not attend college and people who attended some college but did not complete a degree.

8 Assessing the Influence of Revealed Preference Assumptions

So far we have followed the convention in the VSL literature and assumed that people make informed tradeoffs between consumption and mortality risk. In our context, the assumption is that Medicare beneficiaries accurately assess their OOP costs of reducing their mortality risk, perhaps with the assistance of family members and physicians. In reality, this assumption is unlikely to always hold true because some people do not fully understand their treatment options and billing procedures, even with help from family and physicians. We capitalize on the fact that the MCBS includes ancillary questions that allow us to assess how the VSLs implied by medical spending vary across people based on their knowledge and decision autonomy.


³⁰ For example, if we assume that the difference in income between people in the "above \$40,000" and "below \$20,000" bins is approximately \$20,000, then doubling income at age 67 is associated with multiplying VSL by 2.56, yielding an upper bound on the elasticity of 1.28.

8.1 The VSL is Insensitive to Who Makes Health Care Decisions

The MCBS asks people whether they usually make health insurance decisions on their own, receive help making decisions and who helps them, or rely on others to make decisions for them. In cases of Alzheimer's disease or other impairments, the proxy who makes health insurance decisions also responds to the MCBS. For these patients our VSL measures are best interpreted as a reflection of family-level valuations because the proxy decision-makers are almost always family members.

Panel (A) of Figure 9 stratifies our VSL estimates based on who usually makes health insurance decisions. There is little difference between the 67.6% of beneficiaries who usually make their own decisions, the 27.6% who get help, and the 4.8% who rely on someone else to make decisions for them. As a result, narrowing the focus to the subset of people who make their own medical care decisions yields virtually the same VSL measures as our featured specification (Figure 4). Thus, consistent with the simplifying assumptions of our theoretical model, the distinction between individual- and family-level valuations of mortality risk reductions for beneficiaries does not appear to be quantitatively important for our estimates.

Figure 9: Heterogeneity in VSL by Decision Process and Knowledge



Note: Each panel shows the mean age-specific VSL in \$1,000 (2010) dollars. The VSL is calculated from the model shown in col (3) of Table 2. Panel (A) stratifies the results based on who makes medical care decisions for the beneficiary. Panel (B) stratifies the results based on whether we observe evidence causing us to suspect that the beneficiary may not be fully informed. See the text for definitions. Markers along each trend line denote ages at which the VSL exceeds the VSL for the next lower trend line in at least 99% of 1,000 bootstrap samples with errors clustered by hospital referral region.

8.2 The VSL Increases Slightly with Health Care Knowledge

The MCBS also allows us to evaluate potential effects of some people not being fully informed about their costs and benefits of medical care. For information frictions to attenuate our VSL measures, the frictions would have to increase the marginal return to medical spending. This could occur, for example, if “behavioral hazard” causes people to systematically underuse beneficial treatments (Baicker, Mullainathan and Schwartzstein (2015)). However, we note that our estimated returns to spending span the local average treatment effects from prior studies of contexts in which undertreatment due to behavioral hazard seems unlikely, e.g. inpatient spending on heart-attack patients admitted through the emergency room (Doyle (2011)).

MCBS data do not facilitate investigation into specific information frictions, but they do contain several signals about whether beneficiaries are likely to be more or less informed. We use these signals to classify decisions about annual medical care expenditures as “suspect” or “non-suspect” for the purpose of revealing preferences, borrowing terminology from Bernheim and Rangel (2009). We classify decisions as “non-suspect” if we have no reason to suspect that the decision-maker is less than fully informed. We classify decisions as “suspect” if we suspect that conventional revealed preference assumptions may not strictly hold in the data because one or more of the following statements about the beneficiary is true: (1) does not make their own health insurance decisions, (2) has assistance managing money, (3) does not realize that OOP costs vary across Medicare Part D prescription drug plans, (4) suffers from dementia and/or depression, or (5) does not think they know most of what they need to know about Medicare.³¹ These criteria lead us to classify 82% of all person-years of expenditure decisions as suspect. Importantly, this classification does not mean that revealed preference logic necessarily fails for these observations, only that we have reason to suspect that it might.

Panel (B) of Figure 9 shows that non-suspect choices are associated with slightly higher VSL measures even conditional on age. This is consistent with the hypothesis that less informed beneficiaries have higher marginal returns to medical spending, e.g. due to behavioral hazard. However, it can also be explained by the fact that the people making non-suspect choices are slightly healthier (e.g. 1.8 fewer chronic conditions, 36% of a standard deviation reduction in the HCC score). In any case, the differences between the suspect and non-suspect VSL measures are small, suggesting that heterogeneity in information frictions is unlikely to substantially attenuate

³¹ The Part D knowledge question asks respondents whether it is true or false that “Your out-of-pocket costs are the same in all Medicare prescription drug plans.” The correct answer is false. The Medicare general knowledge question asks people to report “How much do you think you know about the Medicare program? Do you know... [just about everything/most/some/a little/almost none] of what you need to know about the Medicare program?”

our main VSL estimates.

8.3 Imperfect Physician Agency Would Bias the VSL Estimates Upwards

The presence of physicians differentiates medical decisions from the occupation choices that have traditionally been used to infer VSL. Ideally, physicians would help patients understand their options, strengthening the credibility of the revealed preference assumptions. However, a potential concern is that Medicare’s fee-for-service payment methods may incentivize some physicians to recommend more treatment than under perfect agency. While we cannot directly evaluate the importance of this concern for our results, we expect it to work against our finding that the VSL is low compared to wage-hedonic studies. All else constant, higher medical spending due to breakdowns in physician agency driven by fee-for-service payment would lower our estimated returns to marginal medical spending and subsequently inflate our VSL estimates.

8.4 Deriving an Upper Bound on the VSL


As a final step in evaluating the scope for information frictions to attenuate our estimates, we derive upper bounds on the VSL by making an extreme assumption about beliefs. We assume that people ignore insurance when they make spending decisions and instead falsely believe that they will pay the entirety of their medical bills out of pocket. This increases our estimate for the mean VSL to approximately \$1.3 million for those in their late 60’s and \$100,000 for those in their late 90’s.³² These values approximately double if we make the additional ad hoc adjustment of dividing our estimated returns to spending by a constant that forces our model to match the lower bound point estimate on average returns to spending from recent quasi-experimental studies of the Medicare population (Huh and Reif (2017), Clayton (2018)), Doyle et al. (2015), Romley and Sood (2013), and Doyle (2011)). Even under these extreme assumptions, the resulting upper bound on VSL for ages in the high sixties is approximately one-quarter of the standard wage regression estimates derived from younger workers’ occupation choices.

³² These values are also analogous to what Hall and Jones (2007) call the “social value of life” because they can be reinterpreted as extending revealed preference logic to incorporate taxpayer expenditures on Medicare. Our \$1.3 million estimate for those age 67–69 is very similar to the measures they calibrate from macro data on medical spending. Our use of microdata on health, demographic and socioeconomic characteristics introduces more curvature so that our estimates decline more steeply with age. Mechanically, we calculate this simply by replacing each person’s observed coinsurance rate with a coinsurance rate of one.

9 The Value of a Statistical Life Year

Policy analyses often rely on annuitized VSL estimates to monetize the benefits of policies that modify life expectancy among older populations. A conventional but arbitrary value for one statistical life year (VSLY) is \$100,000. Revealed preference evidence from workers suggests that the VSLY is an inverse U-shaped function of age (Aldy and Viscusi (2008))³³. We add to this literature by using our VSL estimates to provide direct evidence on the VSLY for seniors.

Figure 10: Value of a Statistical Life Year by Age and Gender



Note: Each panel shows the mean age-by-gender-specific value per statistical life year in \$1,000 (2010) dollars. The VSL is calculated from the model shown in col (3) of Table 2. Panel (A) uses a discount rate of 3% and Panel (B) uses a discount rate of 7%. Markers along each trend line denote ages at which the underlying VSL estimate exceeds the VSL estimate for the next lower trend line in at least 99% of 1,000 bootstrap samples with errors clustered by hospital referral region.

Figure 10 shows the value of a statistical life year by age and gender, and Appendix Table A.4 reports the values underlying the figure. We calculate these measures by combining our age-by-gender-specific VSL measures with age-by-gender-specific information on expected life years remaining from the US life tables. Panels (A) and (B) of the figure report the VSLY for the US Office of Management and Budget's recommended range of discount rates for valuing mortality reductions: 3% to 7% (U.S. Office of Management and Budget (2003)). Even at the upper-bound discount rate, our VSLY estimates at age 67 are well below the commonly used benchmark values. The differences by gender and the decline with age are consistent with a strong relationship between the health stock

³³ Despite the evidence in Aldy and Viscusi (2008) the US Office of Management and Budget (U.S. Office of Management and Budget (2003)) gives the following instructions to economists tasked with cost-benefit analysis of federal programs who choose to use VSLY measures: "you should adopt a larger VSLY estimate for senior citizens because senior citizens face larger overall health risks from all causes and they may have accumulated savings to spend on their health and safety".

and the value of statistical life extension.

With a 7% discount rate our estimates imply a mean VSLY of \$51,000 for women at age 67. Figure A.4 shows the sensitivity of our VSLY estimates to the 200 alternative specifications for the survival function discussed in Section 6. The 95th percentile in the distribution of models for females age 67 is \$71,000 and the maximum is \$105,000. In fact, this is the only specification that ever yields an average VSLY over \$100,000 for men or women of any age beyond 67.

10 Conclusion

We linked US seniors' Medicare records to survey data on their health and medical spending, estimated their value of a statistical life, and analyzed heterogeneity in VSL measures by age, health, income, demographics, knowledge and agency. Our results imply that the conventional wage–hedonic estimates for the VSL overstate by an order of magnitude what seniors are willing to pay for medical care that marginally increases their own survival probabilities. Likewise, under standard assumptions for discount rates we find values per statistical life year (VSLY) that are less than half the size of values that are commonly used to assess the benefits of technologies, policies, and regulations that affect the health of US seniors. We also find that the VSL increases in health, income and remaining life expectancy. These findings have potential to improve the efficiency and equity of a wide range of government activities that affect seniors' health and longevity.

Our finding that the VSL for seniors increases sharply with their health and life expectancy implies that the VSL should not be treated as a statistic that is invariant to many of the policies that it is used to evaluate. Simply multiplying VSL by the number of premature deaths avoided by a policy will bias the benefit measure toward zero for policies that also reduce morbidity. Such life–saving policies may trigger a virtuous cycle in which deaths are averted directly, but health is also improved, the value of life increases, and people make greater subsequent investments in their health. Extending our analysis to directly model how this dynamic complementarity works through the VSL to modify the benefits of regulations that simultaneously affect morbidity and mortality is an important task for future research.

References

- Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey Wooldridge.** 2017. "When Should You Adjust Standard Errors for Clustering?" Working Paper.
- Aldy, Joseph E, and Seamus J Smyth.** 2014. "Heterogeneity in the Value of Life." NBER Working Paper 20206.
- Aldy, Joseph E, and W Kip Viscusi.** 2008. "Adjusting the Value of a Statistical Life for Age and Cohort Effects." *The Review of Economics and Statistics*, 90(3): 573–581.
- Arthur, W.B.** 1981. "The Economics of Risks to Life." *American Economic Review*, 71(1): 54–64.
- Ashenfelter, Orley, and Michael Greenstone.** 2004. "Using mandated speed limits to measure the value of a statistical life." *Journal of Political Economy*, 112(S1): S226–S267.
- Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein.** 2015. "Behavioral Hazard in Health Insurance." *Quarterly Journal of Economics*, 130(4): 1623–1667.
- Banzhaf, H. Spencer, and V. Kerry Smith.** 2007. "Meta-Analysis in Model Implementation: Choice Sets and the Valuation of Air Quality Improvements." *Journal of Applied Econometrics*, 22: 1013–1031.
- Bauser, Daniel, Darius Lakdawalla, and Julian Reif.** 2018. "Mortality Risk, Insurance, and the Value of Life." *NBER Working Paper*, 25055.
- Bernheim, B. Douglas, and Antonio Rangel.** 2009. "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics." *Quarterly Journal of Economics*, 124(1): 51–104.
- Blomquist, Glenn C., Mark Dickie, and Richard M. O'Connor.** 2011. "Willingness to pay for improving fatality risks and asthma symptoms: Values for children and adults of all ages." *Resource and Energy Economics*, 33(2): 410 – 425.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller.** 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics*, 90(3): 414–427.
- Chandra, Amitabh, and Douglas O Staiger.** 2007. "Productivity spillovers in health care: evidence from the treatment of heart attacks." *Journal of Political Economy*, 115(1): 103–140.

- Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler.** 2016. "The Association Between Income and Life Expectancy in the United States, 2001-2014." *Journal of the American Medical Association*, 315(16): 1750–1766.
- Clayton, Denise Hammock.** 2018. "The Effect of Prescription Drug Coverage on Mortality: Evidence from Medicaid Implementation." *Journal of Health Economics*.
- Costa, Dora L, and Matthew E Kahn.** 2004. "Changes in the Value of Life, 1940–1980." *Journal of Risk and Uncertainty*, 29(2): 159–180.
- Cropper, Maureen, James K Hammitt, and Lisa A Robinson.** 2011. "Valuing mortality risk reductions: progress and challenges." *Annual Review of Resource Economics*, 3(1): 313–336.
- Currie, Janet, and David Slusky.** 2020. "Does the Marginal Hospitalization Save Lives? The Case of Respiratory Admissions for the Elderly." *NBER Working Paper 26618*.
- Cutler, David, Jonathan S. Skinner, Ariel Dora Stern, and David Wennberg.** 2019. "Physician Beliefs and Patient Preferences: A New Look at Regional Variation in Health Care Spending." *American Economic Journal: Economic Policy*, 11(1): 192–221.
- Deleire, Thomas, Shakeeb Kahn, and Christopher Timmins.** 2013. "Roy Model Sorting and Nonrandom Selection in the Valuation of a Statistical Life." *International Economic Review*, 54(1): 279–306.
- Deschenes, Olivier, Michael Greenstone, and Joseph S. Shapiro.** 2017. "Defensive Investments and the Demand for Air Quality: Evidence from the NOx Budget Program." *American Economic Review*, 107(10): 2958–2989.
- Dow, William H, Tomas J. Philipson, and Xavier Sala-I-Martin.** 1999. "Longevity Complementarities Under Competing Risks." *American Economic Review*, 89(5): 1358–1371.
- Doyle, Joseph J.** 2011. "Returns to Local-Area Health Care Spending: Evidence from Health Shocks to Patients Far from Home." *American Economic Journal: Applied Economics*, 3(3): 221–43.
- Doyle, Joseph J, John A Graves, Jonathan Gruber, and Samuel A Kleiner.** 2015. "Measuring Returns to Hospital Care: Evidence from Ambulance Referral Patterns." *Journal of Political Economy*, 123(1): 170–214.
- Evans, Mary F, and Georg Schaur.** 2010. "A quantile estimation approach to identify income and age variation in the value of a statistical life." *Journal of Environmental Economics and Management*, 59(3): 260–270.

- Evans, Mary F, and V Kerry Smith.** 2006. "Do we really understand the age–VSL relationship?" *Resource and Energy Economics*, 28(3): 242–261.
- Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams.** 2016. "Sources of geographic variation in health care: Evidence from patient migration." *The Quarterly Journal of Economics*, 131(4): 1681–1726.
- Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams.** 2019. "Place-Based Drivers of Mortality." NBER Working Paper 25975.
- Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton.** 2013. "Developing a Social Cost of Carbon for US Regulatory Analysis: A Methodology and Interpretation." *Review of Environmental Economics and Policy*, 7(1): 23–46.
- Grossman, Michael.** 1972. "On the concept of health capital and the demand for health." *Journal of Political Economy*, 80(2): 223–255.
- Hall, Robert E., and Charles I. Jones.** 2007. "The Value of Life and the Rise in Health Spending." *The Quarterly Journal of Economics*, 122(1): 39–72.
- Huh, Jason, and Julian Reif.** 2017. "Did Medicare Part D reduce mortality?" *Journal of Health Economics*, 53: 17 – 37.
- Jha, Prabhat, Chinthanie Ramasundarahettige, Victoria Landsman, Brian Rostron, Michael Thun, Robert N Anderson, Tim McAfee, and Richard Peto.** 2013. "21st-century hazards of smoking and benefits of cessation in the United States." *New England Journal of Medicine*, 368(4): 341–350.
- Kneisner, Thomas J, W Kip Viscusi, Christopher Woock, and James P Ziliak.** 2012. "The value of a statistical life: Evidence from panel data." *Review of Economics and Statistics*, 94(1): 74–87.
- Krupnick, Alan.** 2007. "Mortality–risk Valuation and Age: Stated–Preference Evidence." *Review of Environmental Economics and Policy*, 1(1): 261–282.
- Leamer, Edward E.** 1983. "Let's Take the Con Out of Econometrics." *American Economic Review*, 73(1): 31–43.
- Lee, Jonathan M., and Laura O. Taylor.** 2019. "Randomized Safety Inspections and Risk Exposure on the Job: Quasi-Experimental Estimates of the Value of a Statistical Life." *American Economic Journal: Economic Policy*, 11(4): 350–374.

- Leeth, John D, and John Ruser.** 2003. "Compensating wage differentials for fatal and nonfatal injury risk by gender and race." *Journal of Risk and Uncertainty*, 27(3): 257–277.
- McClellan, Mark, Barbara J McNeil, and Joseph P Newhouse.** 1994. "Does more intensive treatment of acute myocardial infarction in the elderly reduce mortality?: analysis using instrumental variables." *Journal of the American Medical Association*, 272(11): 859–866.
- Murphy, Kevin M, and Robert H Topel.** 2006. "The value of health and longevity." *Journal of Political Economy*, 114(5): 871–904.
- Romley, John A., and Neeraj Sood.** 2013. "Identifying the Health Production Function: The Case of Hospitals." NBER Working Paper 19490.
- Rosen, Sherwin.** 1988. "The Value of Changes in Life Expectancy." *Journal of Risk and Uncertainty*, 1: 285–304.
- Skinner, Jonathan S., Elliott S. Fisher, and John Wennberg.** 2005. "The Efficiency of Medicare." *Analyses in the Economics of Aging*, 129–160. University of Chicago Press.
- Smith, V Kerry, Mary F Evans, Hyun Kim, and Donald H Taylor Jr.** 2004. "Do the near-elderly value mortality risks differently?" *Review of Economics and Statistics*, 86(1): 423–429.
- Song, Yunjie, Jonathan Skinner, Julie Bynum, Jason Sutherland, John E Wennberg, and Elliott S Fisher.** 2010. "Regional variations in diagnostic practices." *New England Journal of Medicine*, 363(1): 45–53.
- Stukel, Thérèse A, Elliott S Fisher, David E Wennberg, David A Alter, Daniel J Gottlieb, and Marian J Vermeulen.** 2007. "Analysis of observational studies in the presence of treatment selection bias: effects of invasive cardiac management on AMI survival using propensity score and instrumental variable methods." *Journal of the American Medical Association*, 297(3): 278–285.
- Terza, Joseph V., Anirban Basu, and Paul J. Rathouz.** 2008. "Two-Stage Residual Inclusion Estimation: Addressing Endogeneity in Health Econometric Modeling." *Journal of Health Economics*, 27(3): 531–43.
- U.S. Environmental Protection Agency.** 2011. "Benefits and Costs of the Clean Air Act from 1990 to 2020."
- U.S. Office of Management and Budget.** 2003. "Circular A-4 Memo on Regulatory Analysis."

- Viscusi, W. Kip.** 2010. "The heterogeneity of the value of statistical life: Introduction and overview." *Journal of Risk and Uncertainty*, 40(1): 1–13.
- Viscusi, W Kip, and Joni Hersch.** 2008. "The mortality cost to smokers." *Journal of Health Economics*, 27(4): 943–958.
- Viscusi, W Kip, and Joseph E Aldy.** 2007. "Labor market estimates of the senior discount for the value of statistical life." *Journal of Environmental Economics and Management*, 53(3): 377–392.
- Welch, H Gilbert, Sandra M Sharp, Dan J Gottlieb, Jonathan S Skinner, and John E Wennberg.** 2011. "Geographic variation in diagnosis frequency and risk of death among Medicare beneficiaries." *Journal of the American Medical Association*, 305(11): 1113–1118.
- Wooldridge, Jeffrey M.** 2015. "Control Function Methods in Applied Econometrics." *Journal of Human Resources*, 50(2): 420–445.

A Supplemental Appendix for Online Publication

A.1 Hierarchical condition categories risk adjustment scores

CMS uses the Hierarchical Condition Categories (HCC) risk adjustment score to adjust capitation payments to Medicare Advantage plans based on their enrollees' health expenditure risk. The HCC score is designed to synthesize information about individuals' chronic illnesses and demographics from CMS administrative records³⁴ The index is a function of age, gender, indicators for numerous chronic illnesses and the initial reason for Medicare eligibility.

Raw HCC scores may embed some measurement error. In particular, there is evidence that some of the spatial and temporal variation in diagnosis rates for the chronic illnesses used to compute HCC scores actually reflects differences in medical care providers' diagnostic and treatment decisions rather than differences in patients' health (Song et al. (2010); Welch et al. (2011)).³⁵ We reduce the scope for such errors by adjusting HCC scores using the procedure from Finkelstein, Gentzkow and Williams (2016). This involves regressing HCC score on dummies for year and geographic area, individual fixed effects, and a vector of covariates used to proxy for latent health.

We follow Finkelstein, Gentzkow and Williams (2016) in defining geographic areas as $j = 1, \dots, 306$ Dartmouth Atlas of medical care hospital referral regions (HRR) and in defining the vector of covariates, x_{it} , to include dummies for five-year age bins and relative-year fixed effects, ρ_{it} , for people who change their residential location, where $\rho_{it} = t - t^*$, t^* denotes the year of move, and $\rho_{it} = 0$ for people who do not move at any point during the study period.³⁶ Including these relative-year migration dummies in the covariate vector recognizes that migration late in life may coincide with negative health shocks that induce people to move closer to caregivers. We use the resulting predicted health index as a measure of objective health in the survival function.

³⁴ Additional background information on the risk adjustment model can be found at <http://www.nber.org/data/cms-risk-adjustment.html>.


³⁵ For example, Song et al. (2010) uses movers to examine how diagnosis rates change as people move across quintiles of the distribution of spending. Results showed a significantly larger increase in diagnosis rates for those who moved to higher intensity regions compared to those who moved to lower intensity regions and those who did not move at all.

³⁶ "Hospital Referral Regions" (HRRs) represent regional medical care markets for tertiary medical care as determined by the Dartmouth Atlas. Each HRR contains at least one hospital that performs major cardiovascular procedures and neurosurgery. HRRs were defined by assigning Hospital Service Areas to the region where the greatest proportion of major cardiovascular procedures were performed, with minor modifications to achieve geographic contiguity, a minimum population size of 120,000, and a high localization index. The Dartmouth Atlas defines a Hospital Service Area as a collection of ZIP Codes whose residents receive most of their hospitalizations from hospitals in the area. For further details see: <http://www.dartmouthatlas.org/downloads/methods/geogappdx.pdf>.

A.2 Additional tables and figures


Figure A.1 provides a graphical representation of our first stage results. The histogram shows the density of the instrument constructed from the HRR fixed effects in (12). The solid line shows the conditional variation in medical expenditures predicted by the HRR-based instrument. Specifically, it shows the expenditure levels predicted by regressing residual medical spending on residual variation in the instrument, after controlling for the HCC index, presence of ADL restrictions, presence of IADL restrictions, self-reported health categories, smoking history, sex, an age spline that allows the marginal effect of age to vary by sex and by whether people are under or over 90, marital status, living children, underweight BMI, indicators for race, and indicators for educational attainment. The dashed lines represent a 95% confidence interval on the prediction. Intuitively, the slope suggests that a one dollar increase in the instrument is associated with approximately a one dollar increase in medical expenditure.

Figure A.1: Identifying Variation in the Instrument



Note: The histogram shows the variation in medical spending due to place effects estimated for 306 hospital referral regions for 2005–2011. The right vertical axis plots conditional variation in medical spending against conditional variation in the instrument after removing the variation in each that is explained by individual measures of health and age. Dashed lines show 95% confidence intervals on predicted values.

Figure A.2: Coinsurance Rate and Return to Spending: Age 67 to Age 97



Note: The dashed line shows the average coinsurance rate from the data, i.e., the ratio of out-of-pocket to total medical expenditures. The solid line shows the average marginal effect of a \$1,000 increase in medical spending on the probability of surviving to the end of the following year measured in percentage points and calculated from the model shown in col (3) of Table 2.

Table A.1: First Stage Results from Survival Function Estimation

	(2)	(3)	(4)	(5)	(6)
instrument	0.868 (0.166)	0.961 (0.171)	0.796 (0.160)	0.646 (0.195)	0.670 (0.293)
HCC index	14.926 (0.414)	14.849 (0.396)	14.857 (0.397)	15.098 (0.398)	15.105 (0.400)
one or more ADL restrictions	2.556 (0.195)	2.644 (0.199)	2.613 (0.196)	2.552 (0.196)	2.562 (0.196)
one or more IADL restrictions	1.493 (0.214)	1.449 (0.211)	1.482 (0.212)	1.490 (0.209)	1.493 (0.207)
health = poor	7.633 (0.461)	7.812 (0.466)	7.825 (0.469)	7.868 (0.467)	7.850 (0.469)
health = fair	2.353 (0.241)	2.470 (0.242)	2.475 (0.241)	2.486 (0.240)	2.487 (0.240)
health = very good	-1.599 (0.149)	-1.635 (0.143)	-1.643 (0.143)	-1.625 (0.146)	-1.631 (0.147)
health = excellent	-2.641 (0.186)	-2.657 (0.185)	-2.670 (0.183)	-2.656 (0.184)	-2.670 (0.184)
ever smoked	0.195 (0.147)	0.220 (0.143)	0.193 (0.143)	0.145 (0.143)	0.139 (0.141)
male	-0.553 (1.934)	-1.128 (1.869)	-1.030 (1.850)	-0.812 (1.840)	-0.855 (1.823)
age x {male} x {under 90}	-0.369 (0.024)	-0.378 (0.022)	-0.382 (0.022)	-0.394 (0.022)	-0.393 (0.022)
age x {female} x {under 90}	-0.353 (0.019)	-0.369 (0.018)	-0.372 (0.018)	-0.381 (0.019)	-0.381 (0.018)
age x {male} x {over 90}	-0.361 (0.022)	-0.372 (0.021)	-0.376 (0.021)	-0.387 (0.021)	-0.387 (0.021)
age x {female} x {over 90}	-0.337 (0.018)	-0.353 (0.017)	-0.356 (0.017)	-0.364 (0.018)	-0.364 (0.018)
insurance type covariates		x	x	x	x
health care quality covariates			x	x	x
environmental covariates				x	x
state dummies					x
number of person-years	44,697	44,697	44,697	44,697	44,697
number of people	22,206	22,206	22,206	22,206	22,206

Note: The table reports coefficients and bootstrapped standard errors clustered by hospital referral region from the first stage regressions corresponding to Table 2 columns (2) through (6). The dependent variable is annual gross medical expenditures, measured in \$1,000. Continued on the next page.

Table A.1: (continued) First Stage Results from Survival Function Estimation

	(2)	(3)	(4)	(5)	(6)
married	0.412 (0.153)	0.114 (0.150)	0.127 (0.150)	0.135 (0.151)	0.121 (0.151)
has living children	0.634 (0.264)	0.617 (0.255)	0.638 (0.255)	0.597 (0.254)	0.572 (0.254)
underweight BMI	-0.741 (0.361)	-0.778 (0.356)	-0.754 (0.355)	-0.692 (0.356)	-0.682 (0.356)
African-American	-2.136 (0.332)	-1.150 (0.307)	-1.102 (0.303)	-1.039 (0.309)	-1.079 (0.310)
Hispanic	-1.579 (0.649)	0.010 (0.679)	-0.031 (0.704)	-0.026 (0.679)	-0.067 (0.664)
race = other	-2.647 (0.527)	-1.602 (0.509)	-1.719 (0.492)	-1.548 (0.496)	-1.494 (0.501)
education = less than high school	-1.493 (0.205)	-1.057 (0.204)	-1.031 (0.202)	-0.933 (0.203)	-0.924 (0.205)
education = some college	0.631 (0.193)	0.576 (0.191)	0.553 (0.192)	0.575 (0.193)	0.557 (0.194)
education = college	1.810 (0.215)	1.485 (0.204)	1.402 (0.203)	1.396 (0.202)	1.377 (0.204)
Medicare advantage coverage		-3.234 (0.252)	-3.271 (0.246)	-3.326 (0.230)	-3.284 (0.234)
Medigap coverage		1.519 (0.184)	1.525 (0.183)	1.544 (0.184)	1.571 (0.187)
Medicaid coverage		-1.622 (0.261)	-1.564 (0.257)	-1.628 (0.261)	-1.652 (0.261)
insurance type covariates		x	x	x	x
health care quality covariates			x	x	x
environmental covariates				x	x
state dummies					x
number of person-years	44,697	44,697	44,697	44,697	44,697
number of people	22,206	22,206	22,206	22,206	22,206

Note: The table reports coefficients and bootstrapped standard errors clustered by hospital referral region from the first stage regressions corresponding to Table 2 columns (2) through (6). The dependent variable is annual gross medical expenditures, measured in \$1,000. Continued on the next page.

Table A.1: (continued) First Stage Results from Survival Function Estimation

	(2)	(3)	(4)	(5)	(6)
hospital compare index			-0.249 (2.523)	-6.643 (3.177)	-0.544 (5.175)
hospital beds / capita			-0.906 (0.274)	-0.528 (0.372)	-0.769 (0.512)
primary care physicians / capita			-0.006 (0.012)	-0.022 (0.017)	-0.027 (0.025)
medical care specialists / capita			0.017 (0.006)	0.007 (0.009)	0.014 (0.013)
ambulatory discharges / capita			0.023 (0.010)	0.026 (0.011)	0.032 (0.018)
automobile mortality				-0.180 (0.027)	-0.187 (0.032)
homicide mortality				-0.038 (0.036)	-0.078 (0.049)
fine particulate matter				-0.173 (0.050)	-0.142 (0.054)
mean summer high temperature				0.030 (0.020)	0.036 (0.023)
mean winter low temperature				-0.017 (0.021)	-0.013 (0.032)
share urban				0.093 (1.230)	-1.489 (1.791)
median household income				0.000 (0.000)	0.000 (0.000)
high school graduation rate				6.045 (3.766)	4.026 (5.351)
insurance type covariates		x	x	x	x
health care quality covariates			x	x	x
environmental covariates				x	x
state dummies					x
number of person-years	44,697	44,697	44,697	44,697	44,697
number of people	22,206	22,206	22,206	22,206	22,206

Note: The table reports coefficients and bootstrapped standard errors clustered by hospital referral region from the first stage regressions corresponding to Table 2 columns (2) through (6). The dependent variable is annual gross medical expenditures, measured in \$1,000.

Table A.2: IV Survival Functions, Full Results

	(2)	(3)	(4)	(5)	(6)
\$1,000 in medical spending	-0.471 (0.201)	-0.424 (0.178)	-0.638 (0.263)	-0.884 (0.628)	-0.815 (5.804)
1st stage residual morbidity	0.537 (0.203)	0.493 (0.180)	0.707 (0.264)	0.954 (0.629)	0.885 (5.804)
HCC index	11.114 (3.014)	10.718 (2.673)	13.910 (3.864)	17.771 (9.418)	16.726 (86.501)
one or more ADL restrictions	3.049 (0.558)	2.958 (0.518)	3.514 (0.698)	4.128 (1.572)	4.021 (14.305)
one or more IADL restrictions	1.281 (0.396)	1.216 (0.367)	1.544 (0.487)	1.919 (0.993)	1.860 (8.739)
health = poor	7.190 (1.670)	6.854 (1.516)	8.518 (2.231)	10.462 (5.020)	9.845 (46.173)
health = fair	2.756 (0.544)	2.666 (0.520)	3.187 (0.733)	3.810 (1.613)	3.613 (14.962)
health = very good	-2.390 (0.453)	-2.343 (0.426)	-2.694 (0.542)	-3.080 (1.043)	-2.953 (9.442)
health = excellent	-3.810 (0.673)	-3.699 (0.618)	-4.269 (0.786)	-4.910 (1.643)	-4.696 (14.832)
ever smoked	1.302 (0.250)	1.271 (0.244)	1.307 (0.257)	1.337 (0.299)	1.327 (1.322)
male	6.761 (3.366)	6.646 (3.322)	6.436 (3.558)	6.567 (4.279)	6.454 (16.835)
age x {male} x {under 90}	0.019 (0.077)	0.029 (0.069)	-0.053 (0.099)	-0.159 (0.239)	-0.130 (2.197)
age x {female} x {under 90}	0.090 (0.072)	0.098 (0.066)	0.018 (0.102)	-0.080 (0.247)	-0.054 (2.204)
age x {male} x {over 90}	0.044 (0.074)	0.053 (0.068)	-0.028 (0.098)	-0.132 (0.236)	-0.103 (2.158)
age x {female} x {over 90}	0.112 (0.068)	0.119 (0.063)	0.042 (0.098)	-0.052 (0.235)	-0.027 (2.122)
insurance type covariates		x	x	x	x
health care quality covariates			x	x	x
environmental covariates				x	x
state dummies					x
number of person-years	44,697	44,697	44,697	44,697	44,697
number of people	22,206	22,206	22,206	22,206	22,206

Note: The table reports coefficients and bootstrapped standard errors clustered by hospital referral region from the survival functions in Table 2 columns (2) through (6). The table reports average marginal effects expressed as percentage point changes in the one-year probability of death. Continued on the next page.

Table A.2: IV Survival Functions, Full Results

	(2)	(3)	(4)	(5)	(6)
married	-0.394 (0.232)	-0.520 (0.212)	-0.491 (0.230)	-0.430 (0.293)	-0.441 (0.520)
has living children	0.594 (0.411)	0.554 (0.406)	0.701 (0.456)	0.837 (0.591)	0.759 (3.851)
underweight BMI	2.258 (0.425)	2.307 (0.417)	2.155 (0.487)	1.994 (0.709)	2.006 (3.743)
African-American	-0.799 (0.530)	-0.427 (0.438)	-0.661 (0.528)	-0.817 (0.740)	-0.798 (7.508)
Hispanic	-1.966 (0.668)	-1.362 (0.601)	-1.334 (0.736)	-1.436 (0.972)	-1.294 (5.946)
race = other	-2.873 (1.178)	-2.372 (1.007)	-2.716 (1.111)	-2.958 (1.447)	-2.957 (7.807)
education = less than high school	-0.752 (0.432)	-0.508 (0.349)	-0.740 (0.408)	-0.902 (0.663)	-0.835 (4.487)
education = some college	0.238 (0.277)	0.175 (0.267)	0.293 (0.313)	0.455 (0.434)	0.409 (3.498)
education = college	-0.123 (0.473)	-0.280 (0.399)	-0.004 (0.478)	0.351 (0.910)	0.248 (8.898)
Medicare advantage coverage		-0.878 (0.578)	-1.564 (0.818)	-2.430 (1.970)	-2.190 (18.238)
Medigap coverage		-0.227 (0.432)	0.118 (0.559)	0.542 (1.112)	0.442 (9.634)
Medicaid coverage		-1.501 (0.479)	-1.811 (0.564)	-2.208 (1.105)	-2.128 (8.932)
insurance type covariates		x	x	x	x
health care quality covariates			x	x	x
environmental covariates				x	x
state dummies					x
number of person-years	44,697	44,697	44,697	44,697	44,697
number of people	22,206	22,206	22,206	22,206	22,206


Note: The table reports coefficients and bootstrapped standard errors clustered by hospital referral region from the survival functions in Table 2 columns (2) through (6). The table reports average marginal effects expressed as percentage point changes in the one-year probability of death. Continued on the next page.

Table A.2: IV Survival Functions, Full Results

	(2)	(3)	(4)	(5)	(6)
hospital compare index			-1.888 (3.224)	-5.581 (7.205)	-2.347 (20.157)
hospital beds / capita			-0.694 (0.497)	-0.255 (0.693)	-0.385 (3.274)
primary care physicians / capita			-0.017 (0.017)	-0.020 (0.031)	-0.011 (0.126)
medical care specialists / capita			0.019 (0.012)	0.006 (0.017)	0.008 (0.127)
ambulatory discharges / capita			0.026 (0.016)	0.037 (0.025)	0.026 (0.187)
automobile mortality				-0.101 (0.115)	-0.080 (0.960)
homicide mortality				-0.108 (0.061)	-0.117 (0.557)
fine particulate matter				-0.109 (0.150)	-0.107 (0.910)
mean summer high temperature				0.033 (0.042)	0.042 (0.221)
mean winter low temperature				-0.040 (0.030)	-0.037 (0.288)
share urban				3.527 (1.822)	2.495 (10.999)
median household income				0.000 (0.000)	0.000 (0.000)
high school graduation rate				6.359 (6.425)	7.006 (58.410)
insurance type covariates		x	x	x	x
health care quality covariates			x	x	x
environmental covariates				x	x
state dummies					x
number of person-years	44,697	44,697	44,697	44,697	44,697
number of people	22,206	22,206	22,206	22,206	22,206

Note: The table reports coefficients and bootstrapped standard errors clustered by hospital referral region from the survival functions in Table 2 columns (2) through (6). The table reports average marginal effects expressed as percentage point changes in the one-year probability of death.

Figure A.3: Heterogeneity in the VSL at Age 70



Note: The histogram shows the variation in VSL estimates based on 2,698 people who we observe at age 70. Conditional on age, the VSL differs across person-types due to differences in their health and demographics.

Table A.3: Internal Meta-Analysis of Sensitivity to Analytic Decisions

	ln(mean VSL) (all ages)	ln(mean VSL) (age 67)	ln(mean VSL) (age 77)	ln(mean VSL) (age 87)
Gompertz specification	-0.088*** (0.026)	-0.134*** (0.026)	-0.055** (0.026)	0.011 (0.026)
claims-based spending measure	0.261*** (0.026)	0.226*** (0.026)	0.299*** (0.026)	0.382*** (0.026)
include workers in estimation sample	0.214*** (0.026)	0.243*** (0.026)	0.124*** (0.026)	0.084*** (0.026)
covariates: Table 2, Column 2	-0.119*** (0.042)	-0.129*** (0.042)	-0.110*** (0.042)	-0.101** (0.042)
covariates: Table 2, Column 4	-0.324*** (0.042)	-0.322*** (0.042)	-0.327*** (0.042)	-0.324*** (0.042)
covariates: Table 2, Column 5	-0.427*** (0.042)	-0.420*** (0.042)	-0.429*** (0.042)	-0.426*** (0.042)
covariates: Table 2, Column 6	-0.552*** (0.042)	-0.542*** (0.042)	-0.554*** (0.042)	-0.553*** (0.042)
instrument: FGW + integer age	-0.002 (0.042)	-0.002 (0.042)	-0.002 (0.042)	-0.002 (0.042)
instrument: FGW + integer age x gender	-0.003 (0.042)	-0.003 (0.042)	-0.003 (0.042)	-0.003 (0.042)
instrument: FGW, including never-movers	-0.129*** (0.042)	-0.130*** (0.042)	-0.129*** (0.042)	-0.128*** (0.042)
instrument: end of life spending	0.328*** (0.042)	0.327*** (0.042)	0.327*** (0.042)	0.328*** (0.042)
R ²	0.74	0.74	0.74	0.76
number of models	200	200	200	200


Note: The table reports coefficients and standard errors from a regression of VSL estimates on indicators for features of 200 alternative specifications of the survival function. In column (1) the dependent variable is the log of mean VSL. In columns (2), (3) and (4) the dependent variables are the logs of mean VSL for the subsets of people aged 67, 77, and 87. The excluded indicators define the reference model as the one summarized in Table 2 column (3). It uses a Gompit specification for mortality with MCBS spending data, workers excluded, and the instrument based on Finkelstein, Gentzkow and Williams (2016). Coefficients define the conditional effects of deviating from those analytic decisions as explained in the main text. Because all covariates are binary the Halvorsen-Palmquist formula can be used to convert any coefficient shown in the table, β , into its percentage point effect on the VSL: $100 * (e^{\beta} - 1)$. Asterisks indicate the coefficients are statistically distinguishable from zero at the 10% (*), 5% (**) and 1% (***) levels.

Table A.4: VSL and VSLY Estimates by Age for Males and Females


Age	male			female		
	VSL	VSLY		VSL	VSLY	
		3% discount	7% discount		3% discount	7% discount
67	260,904	19,837	26,742	528,917	36,221	50,532
68	246,966	19,410	25,907	519,140	36,677	50,625
69	232,789	18,937	25,024	490,645	35,674	48,776
70	205,763	17,349	22,699	403,258	30,354	41,048
71	176,659	15,462	20,030	362,179	28,150	37,697
72	174,812	15,907	20,404	317,122	25,624	33,926
73	156,340	14,816	18,818	292,854	24,531	32,154
74	144,655	14,301	17,986	265,266	23,085	29,951
75	120,096	12,406	15,453	213,372	19,338	24,830
76	109,743	11,865	14,639	191,382	18,110	23,010
77	95,056	10,776	13,171	180,766	17,912	22,517
78	93,301	11,112	13,457	174,803	18,197	22,630
79	86,224	10,809	12,971	157,665	17,172	21,156
80	72,024	9,522	11,325	129,189	14,766	18,019
81	68,922	9,628	11,352	118,230	14,229	17,198
82	63,562	9,401	10,991	109,484	13,927	16,670
83	54,333	8,523	9,883	99,651	13,456	15,950
84	52,276	8,717	10,027	88,994	12,678	14,904
85	51,609	9,171	10,469	80,811	12,340	14,362
86	44,077	8,353	9,465	71,298	11,593	13,380
87	40,113	8,116	9,133	61,120	10,490	12,024
88	35,795	7,739	8,652	62,148	11,472	13,038
89	40,245	9,306	10,340	62,104	12,213	13,784
90	31,254	7,733	8,543	47,550	10,012	11,222
91	21,504	5,696	6,258	38,784	8,796	9,790
92	22,956	6,510	7,116	35,824	8,629	9,554
93	16,924	5,137	5,590	37,818	9,940	10,928
94	19,586	6,362	6,892	34,704	9,799	10,715
95	14,500	5,037	5,434	26,508	7,878	8,584
96	14,929	5,540	5,956	23,409	7,557	8,189
97	11,419	4,522	4,845	29,332	10,050	10,852

Note: All measures are reported in constant year 2010 dollars. See the main text for explanation of the underlying calculations.


Figure A.4: Sensitivity of VSLY Estimates to Model Features




(a) Female, 3% discount rate



(b) Female, 7% discount rate



(c) Male, 3% discount rate



(d) Male, 7% discount rate

Note: The figures show the estimated mean VSLY by age, gender, and discount rate for each of 200 different specifications of the survival function. Each line corresponds to a different combination of modeling decisions as described in the main text. The large dashed line is our main specification from col (3) of Table 2.